THE INNER CIRCLE TRADERS

The ICT Trading Course

From Beginner to Advanced — Understand, Execute, and Manage Trades

A Complete, Structured Path Through the Inner Circle Trader Methodology

6 PARTS · 32 CHAPTERS · 2026 EDITION

theinnercircletraders.com

The ICT Trading Course

2026 Edition

Risk disclaimer. Trading Forex, futures, and CFDs carries a substantial risk of loss and is not suitable for every investor. This course is educational and does not constitute financial, investment, or trading advice. Past performance does not guarantee future results. You alone are responsible for your trading decisions.

How to use this book. This is a course designed to be read in order, with a chart open and the drills completed. The diagrams are schematic teaching illustrations, not real historical charts; you should study real charts alongside them as you practise.

Contents

Part VI · Execution, Risk & Becoming a Trader

PART I

Foundations & Mindset

CHAPTER 01 · FOUNDATIONS & MINDSET

How to Use This Course

This is a course, not a reference. It is built to be read in order, worked through with a chart open, and returned to. This short chapter explains how to get the most from it.

IN THIS CHAPTER YOU WILL

- → Understand how the course is structured and why the order matters
- → Know what each chapter's sections are for
- → Set up the tools and habits you'll need before Chapter 2

A course that builds on itself

Most trading education is a pile of disconnected tips. This course is deliberately the opposite. It is a single staircase: each chapter assumes you have understood the one before it. Part II teaches you to read structure; Part III shows you the liquidity that structure is built around; Part IV gives you the tools institutions leave behind; Part V adds time and assembles the models; Part VI turns all of it into an executable, risk-managed process. Skip a step and the later steps will feel arbitrary.

The single most important instruction in this entire course is this: read it in order, and do the drills. Reading alone will not make you a trader. The drills at the end of each chapter are where understanding actually forms.

How each chapter is built

Every chapter follows the same shape so you always know where you are. First, what it is — a plain definition. Then why it works — the institutional logic, because a concept you understand is one you can adapt, while a rule you merely memorise breaks the moment the market looks unfamiliar. Then how to spot it, with a diagram. Then how to trade it. Then the common mistakes that cost beginners money. Then drills to practise, and a recap to lock it in.

> KEY IDEA
>
> When you see a Key Points box, treat it as the minimum you must be able to recall before moving on. If you cannot, re-read the chapter. There is no benefit to rushing — the market will still be there.

What you need before Chapter 2

You need very little to start: a charting platform with candlestick charts (a free one is fine), the discipline to keep a simple written journal, and a demo or very small live account you are emotionally indifferent to. You do not need indicators, paid tools, or a signal service. This

methodology is built on price alone. If anything, remove the indicators you already have — they will distract you from learning to read price directly.

A word on honesty

This course will not promise you easy money, because there is none. What it offers is a coherent framework for understanding how price is delivered, and a realistic path to trading it. Whether you succeed depends on how seriously you take the practice. That begins now.

KEY POINTS

- The course is a staircase — each chapter depends on the previous one, so read in order
- Every chapter follows the same template: what, why, how to spot, how to trade, mistakes, drills, recap
- Doing the drills is not optional — understanding forms through practice, not reading
- You need only a candlestick chart, a journal, and a demo account; no indicators or paid tools

PRACTICE DRILLS

- ○ Set up a clean candlestick chart on one instrument you will follow throughout this course.
- ○ Remove all indicators from that chart. Get comfortable looking at raw price.
- ○ Start a simple journal document — you'll use it from Part VI onward, but begin the habit now.

RECAP

This is a build-in-order course where each part depends on the last, and the end-of-chapter drills are where real learning happens. Set up a clean chart and a journal, and commit to working through it sequentially.

CHAPTER 02 · FOUNDATIONS & MINDSET

What ICT Trading Really Is

Before any tool or setup, you need to understand the one idea the entire methodology rests on. Get this, and everything downstream becomes logical rather than arbitrary.

IN THIS CHAPTER YOU WILL

- → State the core premise of ICT in a single sentence
- → Understand why ICT focuses on liquidity rather than indicators
- → See the difference between reacting to price and anticipating delivery

The one idea everything rests on

ICT — the Inner Circle Trader methodology — is a framework for reading price the way large institutions actually move it. Its central premise can be stated in one sentence: price is engineered to reach liquidity. Markets are not a random walk of buyers and sellers meeting by chance. Price is delivered, deliberately, toward the places where orders are resting — and if you can identify those places, you can anticipate where price is likely to go rather than merely reacting after it has already moved.

> KEY IDEA
>
> Write this sentence somewhere you will see it: price is engineered to reach liquidity. Every concept in this course is a more detailed answer to two questions — where is the liquidity, and how will price be delivered to it?

Why not indicators?

Almost all retail education is built on indicators — moving averages, RSI, MACD, and the rest. Every one of them is calculated from past price. By definition, they lag: they tell you what price has already done, not what it is about to do. Worse, because they are so widely used, institutions know exactly where indicator-based traders will act, and can engineer price to trigger those actions. When an oscillator flashes 'oversold' at an obvious low, that low is often precisely the liquidity an institution intends to sweep.

ICT discards indicators entirely and reads price directly: the structure it forms, the liquidity it leaves, and the footprints — order blocks and fair value gaps — that institutional activity prints into the chart. These are not lagging calculations; they are visible at the moment they form.

Reacting versus anticipating

Consider the difference in practice. A reactive trader sees price break above a resistance level and buys, hoping momentum continues. An anticipatory trader saw, before the break, that a

pool of buy-side liquidity sat above that resistance — the stop-losses of short-sellers and the buy orders of breakout traders — and understood that price was likely being drawn up to take it. When price spikes above and then reverses, the reactive trader is trapped; the anticipatory trader expected exactly that and positioned for the reversal.

> EXAMPLE
>
> Two traders watch the same price reach an old high. The first buys the breakout and is stopped out minutes later. The second watched the same candle, recognised the high as buy-side liquidity being swept, waited for confirmation, and entered in the opposite direction. Same chart, same candle — opposite outcomes, decided entirely by which framework each was using.

What ICT is not

ICT is not a magic set of secret levels, and it is not a guarantee. It is a lens. It gives you a coherent, repeatable way to ask where price is going and why. It still requires judgement, practice, and risk management — all of which this course will build. But the lens itself is the foundation, and it starts with that single premise.

KEY POINTS

- The core premise: price is engineered to reach liquidity — markets are delivered, not random
- Every ICT concept answers two questions: where is the liquidity, and how will price reach it?
- Indicators lag because they're calculated from past price, and institutions exploit their predictability
- ICT reads price directly — structure, liquidity, and footprints visible as they form
- The goal is to anticipate delivery, not react to moves that have already happened

COMMON MISTAKES

Treating ICT as a collection of setups to memorise rather than one premise to understand

Keeping indicators 'just in case' — they pull your attention away from reading price

Expecting certainty; the framework improves your odds, it does not remove risk

PRACTICE DRILLS

- [ ] Write the core premise in your own words without looking at the chapter.
- [ ] On a chart, find one recent move and ask: what liquidity was price reaching for?
- [ ] Find one example where an 'obvious' breakout reversed — note what liquidity sat beyond it.

RECAP

ICT reads price the way institutions deliver it, built on one premise: price is engineered to reach liquidity. It replaces lagging indicators with a direct reading of structure, liquidity, and footprints, so you can anticipate delivery instead of reacting to it.

CHAPTER 03 · FOUNDATIONS & MINDSET

Why Price Is Engineered, Not Random

The claim that price is 'engineered' sounds like a conspiracy until you understand the mechanics. This chapter explains why large players must move price the way they do — it isn't a secret, it's a consequence of size.

IN THIS CHAPTER YOU WILL

- → Understand why large orders cannot be filled at a single price
- → See how institutions must create liquidity to enter and exit
- → Recognise engineering as a mechanical necessity, not a conspiracy

The problem of size

Imagine you need to buy a genuinely large position — far larger than the orders currently resting at the best price. If you simply market-buy, you will eat through every sell order available, pushing price sharply up against yourself and getting steadily worse fills. By the time you are filled, you have moved the market and paid a premium. For an institution moving size that dwarfs the visible order book, this problem is enormous. They cannot simply 'buy'. They must find, or manufacture, enough opposing orders to fill against without moving price to their disadvantage.

> KEY IDEA
>
> This is the mechanical root of everything in ICT. Large players need liquidity — a supply of opposing orders — to enter and exit. Where that liquidity is thin, they must engineer it. That engineering is what leaves the footprints you will learn to read.

Where the liquidity is

The most reliable place to find a pool of opposing orders is at obvious levels, because that is where retail traders cluster their orders. Below an obvious support sit the stop-losses of everyone who bought that support, plus the sell-stops of breakout traders — a pool of sell orders. Above an obvious resistance sit the stop-losses of short-sellers and the buy-stops of breakout buyers — a pool of buy orders. These clusters are predictable precisely because retail education teaches everyone to place orders in the same places.

So when an institution needs to buy, it is often efficient to first drive price down into the sell-side pool below support — triggering all those sell orders — and buy against them at a discount. To the retail trader it looks like support 'failed'. Mechanically, the institution just used that pool as fuel.

Engineering is a necessity, not a mystery

Notice that none of this requires a secret cabal. It is the unavoidable consequence of needing to transact in size in a market where liquidity clusters at predictable places. Any large participant faces the same problem and arrives at the same solution: go to where the orders are, and if necessary, push price there to trigger them. When enough large participants behave this way, price takes on the engineered, liquidity-seeking character that ICT describes.

> EXAMPLE
>
> A market rallies all morning into an obvious high where many shorts have placed stops. Price spikes through the high, triggers those buy-stops (which are sell-side fuel from the institution's perspective is actually buy fuel — the shorts buying to cover), the institution offloads its longs into that burst of buying, and price reverses. The 'breakout' was the exit liquidity.

Why this matters for you

Once you internalise that price moves toward liquidity because large players need it there, the whole chart re-organises in your mind. You stop asking 'is this support strong?' and start asking 'whose orders are resting here, and who needs them?' That shift — from levels to liquidity — is the heart of trading like an institution rather than against one.

KEY POINTS

- Large orders can't fill at one price — institutions need pools of opposing orders to transact
- Retail traders cluster stops and entries at obvious levels, creating those predictable pools
- Institutions often drive price into a pool to trigger it, then fill against it — 'support failing' is often this
- Engineering is a mechanical necessity of size, not a conspiracy
- The mental shift is from 'is this level strong?' to 'whose orders rest here, and who needs them?'

COMMON MISTAKES

Thinking 'engineered' means a secret manipulation rather than a consequence of size and liquidity

Still reading levels as simple support/resistance instead of as liquidity pools

Assuming a level 'failed' when it was actually used as fuel exactly as intended

PRACTICE DRILLS

- ○ Pick an obvious support on your chart. Ask: whose orders sit just below it, and what are they?
- ○ Find a case where price swept below support then rallied hard. Explain it in liquidity terms.
- ○ Re-describe three recent moves using the words 'pool', 'fuel', and 'draw' instead of 'support/resistance'.

RECAP

Institutions can't fill large orders at one price, so they must transact against pools of resting orders — which cluster at the obvious levels retail traders use. Driving price into those pools to trigger them is a mechanical necessity of size, and it's what gives price its engineered, liquidity-seeking behaviour.

Builds on: Ch 2

CHAPTER 04 · FOUNDATIONS & MINDSET

Who Moves the Market

To read a market you must know who is in it and what each participant needs. This chapter sketches the players and, crucially, which side you are learning to align with.

IN THIS CHAPTER YOU WILL

- → Identify the main categories of market participant and their goals
- → Understand why retail order flow is the liquidity institutions use
- → See where you fit — and how to stop being the fuel

The two sides that matter

You can divide market participants many ways, but for trading purposes two groups matter most: the large, informed players — banks, hedge funds, and the desks that move genuine size — and retail traders, the individuals trading their own small accounts. The first group must engineer liquidity to operate. The second group, taught by mainstream education to trade obvious patterns, unknowingly provides that liquidity. Almost everything in this course is about moving you from the second group's behaviour toward reading the first group's intent.

What the large players need

Institutions are not trying to be clever for its own sake. They have positions to build and unwind that are too large for the visible market. Their entire challenge is filling those positions at good prices without tipping their hand. That forces them to seek liquidity, to be patient, and to act at specific times when volume is high enough to hide their activity. Their needs — liquidity, timing, discretion — are the reason kill zones, sweeps, and the other concepts in this course exist.

Why retail is the fuel

Retail traders are not foolish; they are simply taught a framework that makes their behaviour predictable. Buy the breakout, sell the breakdown, put your stop below support — when millions do this, their orders form exactly the pools institutions need. The retail trader is not usually competing with the institution so much as feeding it. This is uncomfortable to hear, but it is the most useful thing a new trader can understand, because it points directly at the solution.

> KEY IDEA
>
> You do not beat the institution by out-muscling it — you cannot. You profit by recognising what it is doing and positioning alongside it. That means, quite literally, doing the opposite of the trapped retail crowd at the moments that matter.

Where you fit

You are a small, nimble participant — which is actually an advantage. You do not have the institution's problem of size; you can enter and exit in a heartbeat, wait as long as you like for the right setup, and take only the highest-quality opportunities. Your edge is not capital or speed of information. It is selectivity and understanding: the ability to recognise an engineered move and align with it, while the crowd around you does the opposite and provides the fuel.

EXAMPLE

At a session open, the crowd sees price break yesterday's high and piles in long. You recognise the high as buy-side liquidity, see the break as a sweep rather than a genuine breakout, wait for structure to shift down, and position short. The crowd's buy orders were the liquidity that let larger sellers fill — and you simply chose the right side of that exchange.

KEY POINTS

- The two groups that matter: large informed players who need liquidity, and retail who provide it
- Institutions must fill oversized positions discreetly — hence liquidity-seeking, timing, and patience
- Retail behaviour is predictable because mainstream education teaches everyone the same moves
- You don't beat institutions by force — you recognise their intent and position alongside them
- Your edge as a small trader is selectivity and understanding, not capital or speed

COMMON MISTAKES

Believing you're competing head-to-head with institutions rather than reading and joining them

Trading like the crowd (breakouts, obvious stops) and becoming the fuel

Underestimating the advantage of being small and able to wait for only the best setups

PRACTICE DRILLS

- [ ] For one session, note where 'the crowd' would obviously enter. Watch what happens next.
- [ ] Identify one moment where doing the opposite of the trapped crowd would have paid.
- [ ] Write one sentence on what your specific edge as a small trader is.

RECAP

Markets are driven by large players who must engineer liquidity and retail traders who unknowingly supply it. You profit not by fighting institutions but by recognising their intent and aligning with it — using your advantage of selectivity to avoid being the crowd that becomes the fuel.

Builds on: Ch 2, Ch 3

CHAPTER 05 · FOUNDATIONS & MINDSET

Realistic Expectations & The Learning Path

Most people who quit trading do so because their expectations were wrong, not because the method was. This chapter sets honest expectations and gives you a realistic path from here to competence.

IN THIS CHAPTER YOU WILL

- → Hold realistic expectations about timeline, win rate, and drawdown
- → Understand why most traders fail and how to avoid their path
- → Have a concrete study-and-practice plan for the rest of this course

The uncomfortable truth about timelines

Learning to trade this methodology to competence typically takes months of daily screen time, not weeks — and reaching consistent profitability takes longer still. This is not a flaw in the method; it is the nature of a skill. You would not expect to play an instrument competently after a fortnight. Trading is harder than most instruments, because the market actively punishes your mistakes with money and emotion. Accepting a realistic timeline is itself an edge, because it stops you quitting during the exact period when learning happens.

Win rate, risk-reward, and what 'edge' means

Beginners fixate on win rate, imagining professionals win almost every trade. They do not. A robust ICT approach might win only half its trades — or fewer — and still be highly profitable, because each win is worth several times each loss. This is the single most liberating idea in trading: you do not need to be right often; you need your winners to be much larger than your losers. A method that wins 45% of the time at an average of three-to-one reward-to-risk is enormously profitable. Chase a high win rate instead and you will cut winners short and let losers run — the exact opposite of an edge.

> KEY IDEA
>
> Internalise this now: profitability comes from the relationship between your win rate and your reward-to-risk, not from being right. A modest win rate with large winners beats a high win rate with small ones. Part VI builds your trading around this.

Why most traders fail

The common causes of failure are not mysterious. Traders over-leverage and blow up on a single bad run. They abandon a method after a normal losing streak, jumping to the next shiny

system just before the first one would have paid. They trade without a plan, so every decision is emotional and inconsistent. They refuse to keep records, so they never learn from their mistakes. And they expect certainty, so a losing trade feels like a betrayal rather than a normal cost of doing business. Every one of these is a choice, and this course is designed to steer you away from all of them.

Your path from here

Here is the realistic path. Work through this course in order, doing the drills. As you finish Part IV, begin backtesting — going through historical charts and marking setups — because this compresses months of screen time into weeks. As you reach Part VI, start forward-testing on a demo or tiny live account, journaling every trade. Expect to spend real time in each phase. Do not risk meaningful money until your journal shows a documented edge over a large sample of trades. That is not caution for its own sake; it is the only sane sequence.

> EXAMPLE
>
> A disciplined beginner spends two months studying and backtesting, then three months forward-testing on demo while journaling. Their journal eventually shows a small but consistent edge across two hundred trades. Only then do they risk modest real capital. A year in, they are profitable. An impatient beginner risks real money in week two, blows the account in a month, and concludes 'ICT doesn't work.' Same method — different path.

KEY POINTS

- Competence takes months of daily screen time; consistent profit takes longer — this is normal
- You don't need a high win rate; you need winners much larger than losers (reward-to-risk)
- Most failure comes from over-leverage, method-hopping, no plan, no journal, and expecting certainty
- Backtesting compresses learning; forward-testing on demo builds it safely
- Don't risk meaningful money until your journal documents an edge over a large sample

COMMON MISTAKES

Expecting profitability in weeks and quitting during the normal learning curve

Chasing a high win rate, which leads to cutting winners and holding losers

Risking real capital before a journal proves an edge

PRACTICE DRILLS

- ○ Write down your honest timeline and capital plan. Commit to not risking meaningful money early.
- ○ Calculate: at 45% win rate and 3R average winners, is the method profitable? Prove it to yourself.
- ○ List which of the common failure causes you're personally most at risk of, and how you'll guard against it.

RECAP

Competence takes months and profitability longer; that timeline is normal, not a defect. Your edge comes from large winners relative to losers, not a high win rate, and survival comes from avoiding the well-known failure paths — over-leverage, method-hopping, and trading without a plan or journal. Study, backtest, forward-test, then risk real money only once your records prove an edge.

Builds on: Ch 2, Ch 4

PART II

The Language of Price

CHAPTER 06 · THE LANGUAGE OF PRICE

Reading Candles & Charts Like an Institution

Before structure, liquidity, or any setup, you must read a single candle correctly and know what it tells you about who won the fight for that period. This is the alphabet of price.

IN THIS CHAPTER YOU WILL

- → Read a candlestick's open, high, low, and close and what they imply
- → Interpret wicks as rejection and bodies as conviction
- → Choose timeframes deliberately rather than by habit

What a single candle says

A candlestick summarises one period of trading with four prices: where it opened, the highest and lowest it traded, and where it closed. The body spans open to close; the thin wicks reach to the high and low. In this course a filled body means price closed above its open (a control-to-the-upside candle) and a hollow body means it closed below (control to the downside). Simple as this is, a candle is a story: it tells you where control ended up and how hard the other side fought.

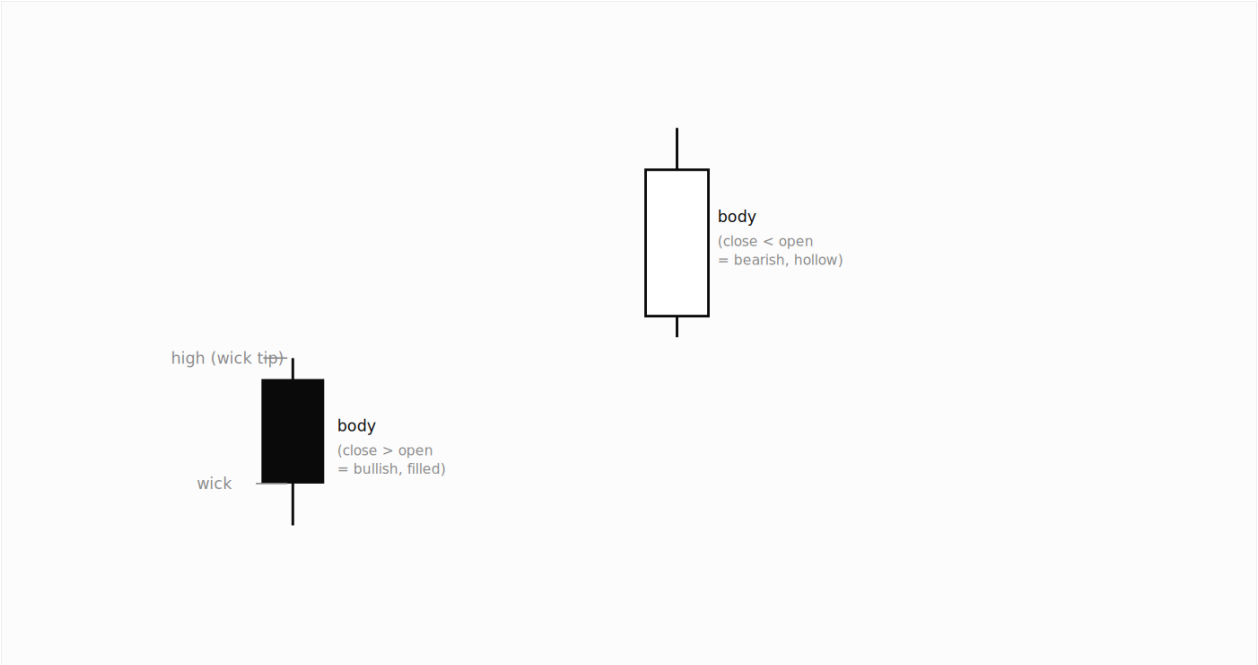

The anatomy of a candle: the body spans the open and close, the wicks reach the high and low. A filled body closed up (bullish); a hollow body closed down (bearish).

Wicks are rejection, bodies are conviction

The most useful habit you can build is reading wicks as rejection. A long wick means price tried to go somewhere and was forced back — the market rejected that level. A long lower wick says buyers overwhelmed a push down; a long upper wick says sellers slammed a push up. The body, by contrast, shows conviction: a large body closing near its extreme says one side was firmly in control for that period. When you later learn to spot sweeps and rejections, you are really just reading wicks at important levels.

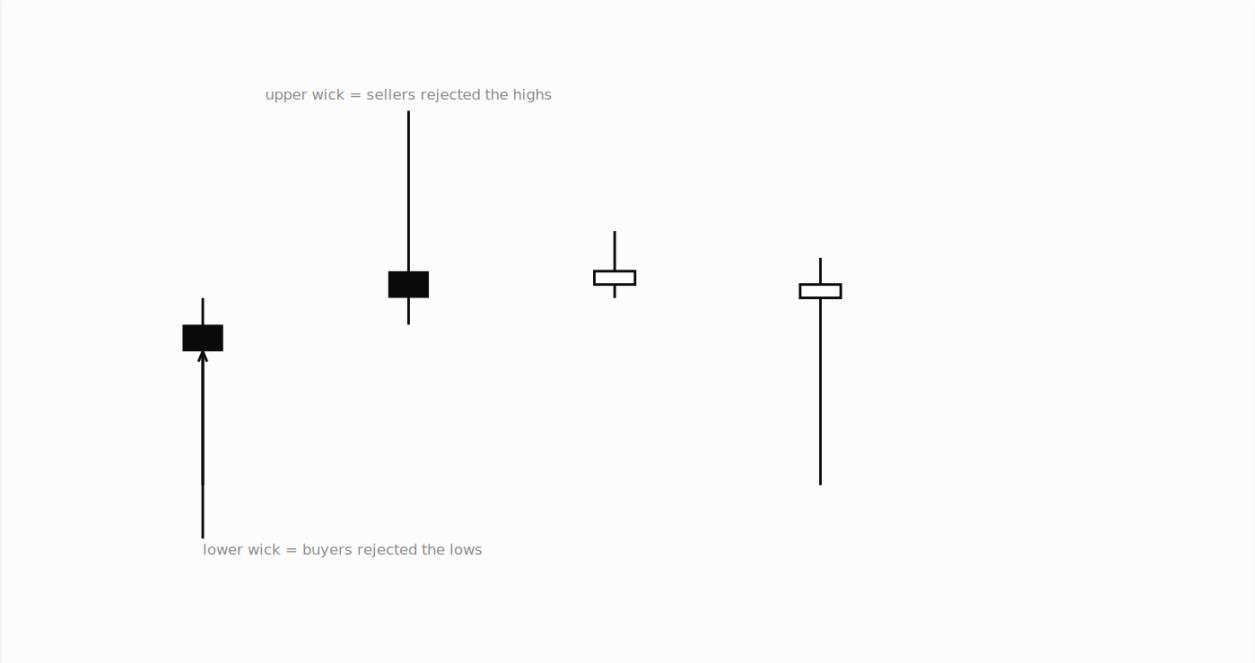

Wicks are rejection. A long lower wick shows buyers overwhelming a push down; a long upper wick shows sellers slamming a push up. Read wicks at key levels above all else.

Timeframes are a choice, not a default

The same market looks different on every timeframe, and each is telling a true story at its own scale. Higher timeframes — daily, four-hour — show the dominant direction and the levels that matter most. Lower timeframes — fifteen, five, one minute — show how price is delivered between those levels and give precise entries. Neither is 'right'. Professionals read the higher timeframe for context and direction, then drop to a lower one to time entries. This is called top-down analysis, and you will use it constantly.

> KEY IDEA
>
> A rule of thumb you'll refine later: let the higher timeframe tell you what to do (direction and key levels), and the lower timeframe tell you when to do it (the precise entry). Confusing the two — taking higher-timeframe-sized risk on a lower-timeframe whim — is a classic beginner error.

Reading a sequence, not a candle

No single candle is a signal on its own. Meaning comes from sequence and location. A long lower wick in the middle of nowhere is noise; the same candle rejecting a known liquidity level after a sweep is a potential entry. As you move through this course, you will always be reading candles in context — at structure, at liquidity, inside a time window. The candle is the word; structure and liquidity are the grammar that give it meaning.

KEY POINTS

- A candle shows open, high, low, close — the body is open-to-close, wicks reach the extremes
- Filled body = closed up (upside control); hollow = closed down (downside control)
- Wicks mean rejection; bodies mean conviction — reading wicks at key levels is core to ICT
- Use higher timeframes for direction and key levels, lower timeframes for precise entries
- Candles only mean something in context — at structure, liquidity, or a time window

COMMON MISTAKES

Reading a single candle as a signal without its location or the sequence around it

Trading one timeframe in isolation instead of top-down (context then timing)

Ignoring wicks — they carry the rejection information that matters most

PRACTICE DRILLS

- ○ On your chart, find five long-wicked candles and describe what each rejection means.
- ○ Take one instrument and read it on the daily, then the 15-minute. Note how the story differs.
- ○ Practice top-down: state the daily direction, then find where a lower timeframe would time an entry.

RECAP

A candle's body shows conviction and its wicks show rejection, but no candle means anything alone — only in sequence and at a location. Read higher timeframes for direction and key levels, then lower timeframes for entry timing. This top-down reading of candles-in-context is the alphabet everything else is written in.

Builds on: Ch 2

CHAPTER 07 · THE LANGUAGE OF PRICE

Market Structure: The Skeleton of Price

Market structure is the skeleton every other concept hangs on. If you can read structure, you can always answer the most important question: which direction is this market currently delivering?

IN THIS CHAPTER YOU WILL

- → Define an uptrend and downtrend in terms of swing points
- → Read structure as a sequence rather than a feeling
- → Identify the current structural direction on any chart

Structure is just a sequence of swings

Strip a chart of everything and market structure is simply the sequence of its turning points — its swing highs and swing lows. An uptrend is a market making higher highs and higher lows: each push up exceeds the last, and each pullback stops above the previous one. A downtrend is the mirror: lower highs and lower lows. A market doing neither cleanly is ranging — rotating between a top and a bottom. That is the entire vocabulary, and almost everything else in this course is layered on top of it.

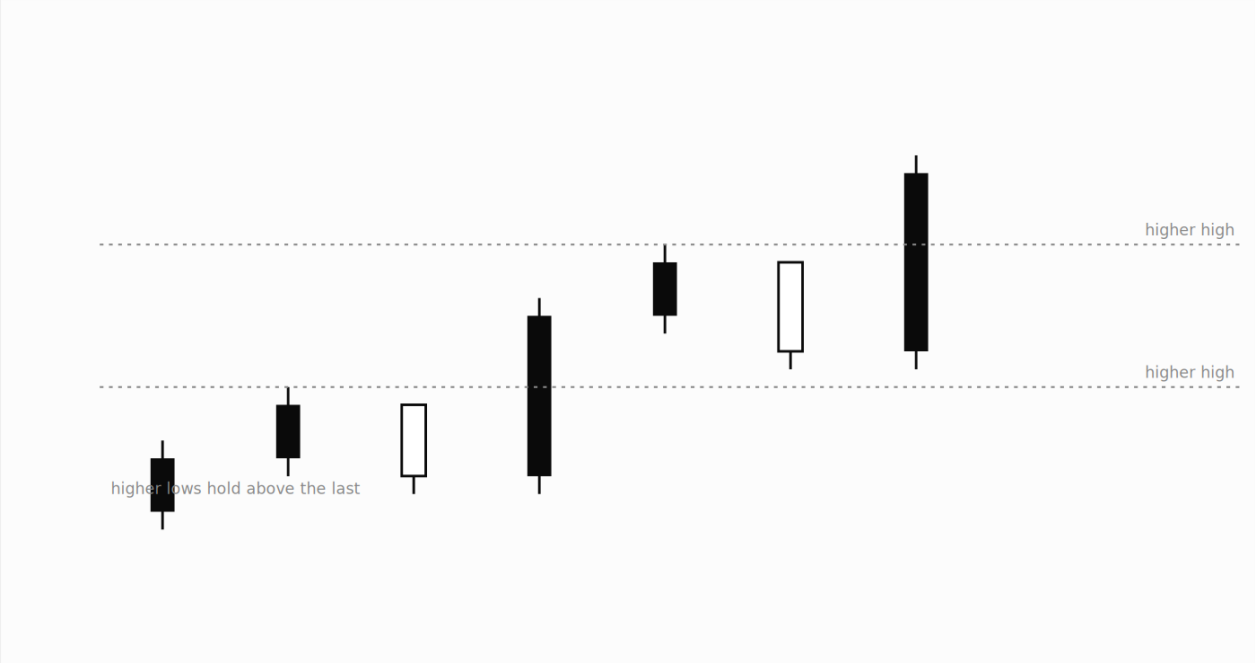

An uptrend is a sequence of higher highs and higher lows: each push exceeds the last, each pullback holds above the previous one. Structure is read from the swings, mechanically.

Why structure comes first

Structure comes before liquidity, before PD arrays, before models — because it sets the context that makes all of them tradeable. An order block in an uptrend is a place to buy; the identical order block in a downtrend is a place to expect failure. A liquidity sweep against the trend and a sweep with the trend mean different things. Without a read on structure, every other tool is ambiguous. With it, they resolve into clear decisions.

> KEY IDEA
>
> Before you look for any setup, answer one question: what is the structure doing here — up, down, or ranging? If you cannot answer confidently, you are not ready to trade this chart. The answer is your permission slip for everything that follows.

Structure is fractal

The same structural rules apply on every timeframe, and timeframes nest inside each other. A daily uptrend contains hour-by-hour pullbacks that are themselves small downtrends. This is not a complication — it is the key to top-down analysis. The higher timeframe's structure tells you the dominant direction; a counter-trend move on a lower timeframe is often just the pullback that sets up the next higher-timeframe entry. Reading structure fractally is how you reconcile a market that seems to be going up and down at once: it depends which scale you are asking about.

Reading structure honestly

Beginners tend to see the structure they wish were there. Discipline yourself to mark the actual swing highs and lows and read what they say, not what you hope. If the sequence is higher highs and higher lows, the market is in an uptrend even if you are bearish. Structure is the one part of analysis that should be almost mechanical — the swings are where they are. Your bias must bend to the structure, never the reverse.

KEY POINTS

- Structure is the sequence of swing highs and lows: HH+HL = uptrend, LH+LL = downtrend, neither = range
- Structure comes first because it gives every other tool its meaning (buy zone vs failure zone)
- The same rules apply on every timeframe, and timeframes nest — this enables top-down analysis
- Read the actual swings mechanically; your bias must bend to structure, not the reverse
- Always answer 'up, down, or ranging?' before looking for any setup

COMMON MISTAKES

Seeing the structure you want instead of the swings actually printed

Trading a setup without first establishing the structural direction

Confusing a lower-timeframe pullback with a change in the higher-timeframe trend

PRACTICE DRILLS

- On three charts, mark the swing highs and lows and label each as up, down, or ranging.
- Find a daily uptrend, then drop to a lower timeframe and find a counter-trend pullback inside it.
- Practice stating structure out loud before any other analysis, on ten different charts.

RECAP

Market structure is the sequence of swing points: higher highs and lows make an uptrend, lower ones a downtrend, and neither a range. It comes first because it gives every other tool meaning, it's fractal across timeframes, and it must be read mechanically — your bias bends to structure, never the reverse.

Builds on: Ch 6

CHAPTER 08 · THE LANGUAGE OF PRICE

Swing Points, BOS & CHoCH

Two events tell you everything about a trend's health: when it continues, and when it changes character. Break of Structure and Change of Character are how you read continuation versus reversal in real time.

IN THIS CHAPTER YOU WILL

- → Define a swing point precisely enough to mark it consistently
- → Distinguish a Break of Structure (continuation) from a Change of Character (reversal)
- → Use BOS and CHoCH to time when to trust a trend and when to doubt it

Defining a swing point

A swing high is a candle whose high is higher than the candles on either side of it; a swing low is the mirror. That is the simplest definition and it is enough to start. The important discipline is consistency: mark them the same way every time so your structure reads are objective. The most recent significant swing high and swing low are the levels that matter, because their breaks are what confirm continuation or reversal.

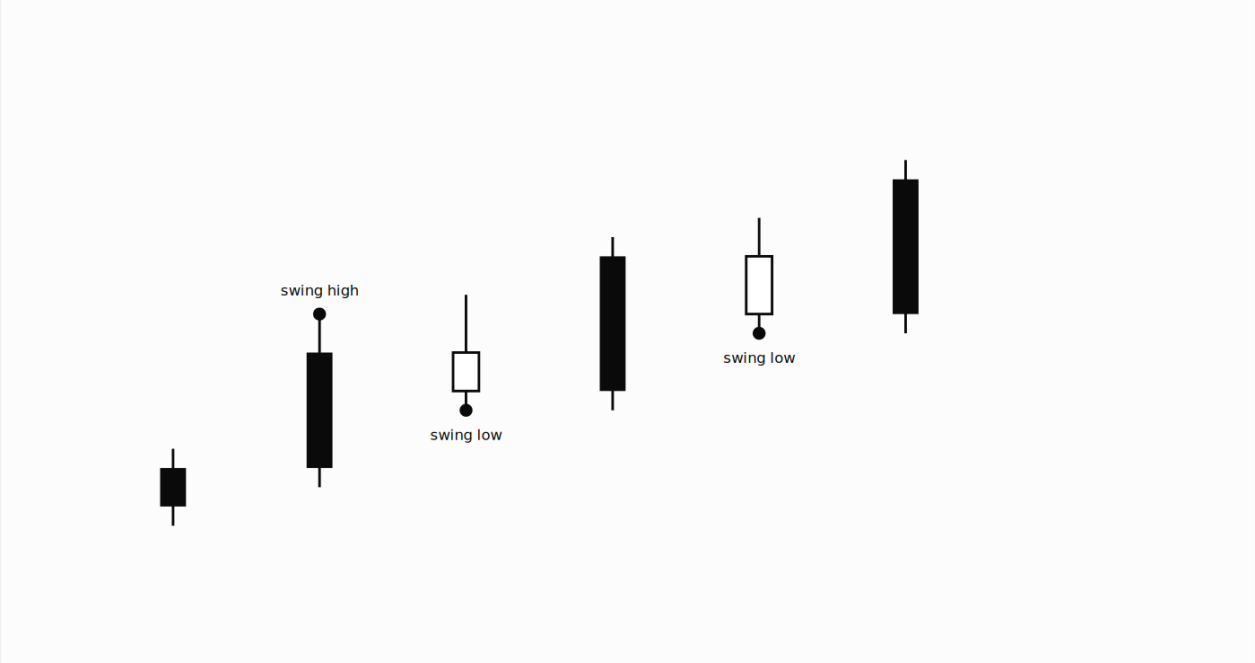

A swing high is a candle whose high exceeds its neighbours; a swing low is the mirror. The most recent significant swings are the levels whose breaks confirm continuation or reversal.

Break of Structure — the trend continues

A Break of Structure (BOS) happens when price breaks the most recent swing point in the direction of the trend. In an uptrend, price breaking above the last swing high is a bullish BOS

— it confirms the uptrend is intact and continuing. Each BOS is the market saying 'still going'. A series of clean BOS events in the same direction is the healthiest thing a trend can show, and it is the context in which you take continuation entries.

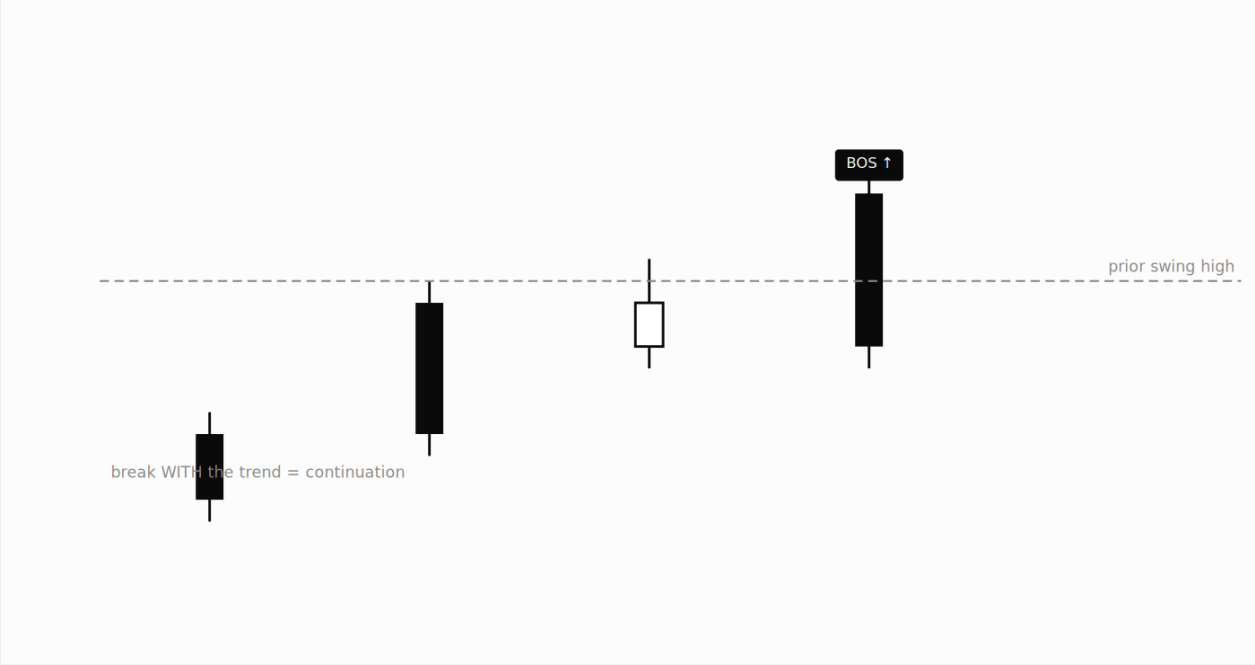

A Break of Structure breaks the prior swing high in the direction of the trend, confirming the uptrend is intact and continuing. A series of clean BOS events is a healthy trend.

Change of Character — the trend may reverse

A Change of Character (CHoCH) is the first break against the trend — the first crack. In an uptrend of higher highs and higher lows, the moment price breaks below the most recent higher low, the pattern of the trend is broken for the first time. That is the CHoCH: the earliest structural signal that control may be shifting. It does not guarantee a full reversal, but it is the first evidence, and it is what reversal setups are built around.

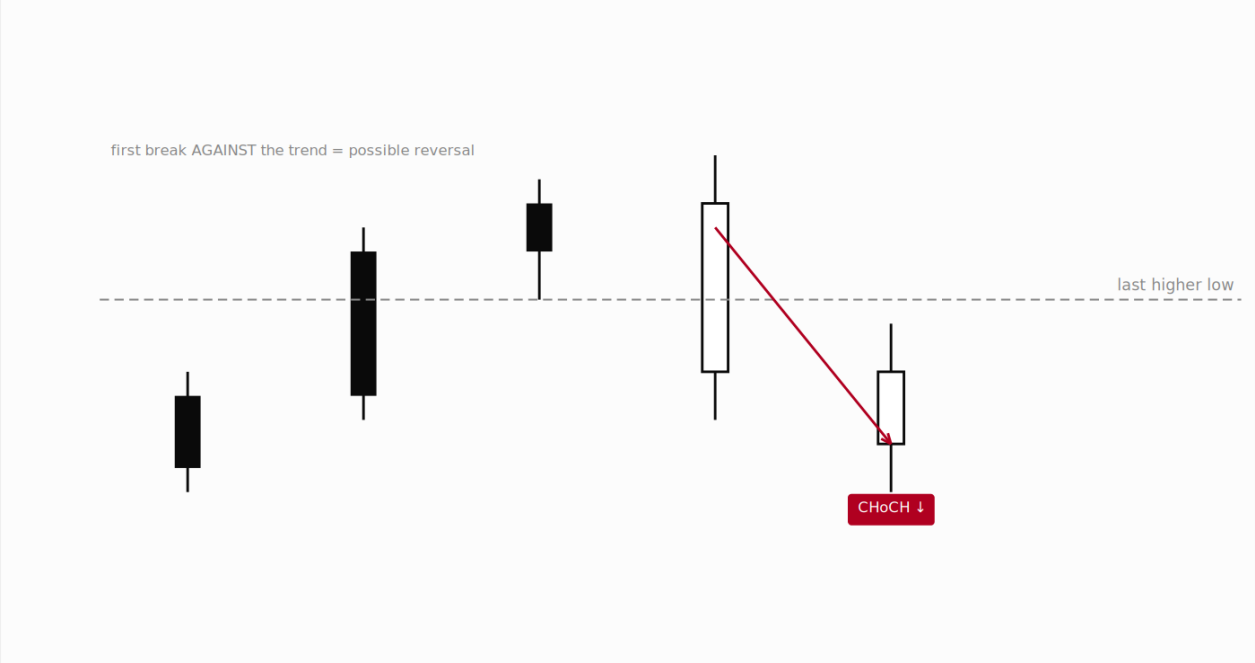

A Change of Character is the first break against the trend — here, price breaking below the last higher low. It is the earliest structural signal that control may be shifting.

KEY IDEA

The distinction in one line: a BOS breaks structure WITH the trend and signals continuation; a CHoCH breaks structure AGAINST the trend and signals a possible reversal. Confusing them is one of the most expensive beginner errors — you end up buying into a reversal thinking it's continuation.

Using them together

In practice you read BOS and CHoCH as a running commentary on the trend. While BOS events keep printing with the trend, you trust it and look for continuation entries. When a CHoCH appears, you become alert to a reversal and stop taking trend-continuation trades until the picture clarifies. Often a reversal reads as: liquidity swept, then CHoCH, then a new trend confirmed by BOS in the new direction. You now have the structural language to describe that whole sequence — and you will use it constantly in Parts IV and V.

KEY POINTS

- A swing high/low is a candle whose extreme exceeds its neighbours — mark them consistently
- BOS = break of structure WITH the trend = continuation confirmed
- CHoCH = first break AGAINST the trend = earliest signal of possible reversal
- While BOS prints with trend, trust it; when CHoCH appears, stop taking continuation trades
- A reversal often reads: sweep → CHoCH → new-direction BOS

COMMON MISTAKES

- Confusing a CHoCH (reversal warning) with a BOS (continuation) and trading the wrong way
- Marking swing points inconsistently, making structure reads subjective
- Treating a single CHoCH as a guaranteed reversal rather than a first warning

PRACTICE DRILLS

- ○ Mark the most recent swing high and low on five charts; identify which break would be BOS vs CHoCH.
- ○ Find one clean BOS and one clean CHoCH in history and label the full sequence around each.
- ○ Track a trend live and note each BOS as it continues, and the CHoCH when it eventually breaks.

RECAP

A Break of Structure breaks a swing point with the trend and confirms continuation; a Change of Character is the first break against the trend and warns of reversal. Read them as a running commentary — trust the trend while BOS prints, turn cautious at the first CHoCH — and you can describe both continuation and the sweep-CHoCH-BOS reversal sequence precisely.

Builds on: Ch 7

CHAPTER 09 · THE LANGUAGE OF PRICE

The Dealing Range, Premium & Discount

Institutions think in ranges, not lines. The dealing range — and its split into a premium half and a discount half — is how you decide whether a price is expensive or cheap, and therefore whether to buy or sell.

IN THIS CHAPTER YOU WILL

- → Define a dealing range from significant swing points
- → Split any range into premium and discount at equilibrium
- → Use premium/discount to filter which side of a setup to take

From lines to ranges

Retail traders think in horizontal lines — 'support here, resistance there'. Institutions think in ranges. A dealing range is simply the span between a significant swing low and a significant swing high — the corridor price is currently working within. Once you have a range, you can ask the question that actually matters for entries: within this range, is price expensive or cheap?

Equilibrium, premium and discount

Split the dealing range in half at its midpoint — the equilibrium, the 50% level. The upper half is premium: price here is relatively expensive, the region you want to be selling from. The lower half is discount: price is relatively cheap, the region you want to be buying from. This single filter — buy in discount, sell in premium — instantly improves the quality of your entries, because it stops you buying at the top of a range or shorting at the bottom, which is where retail habitually gets trapped.

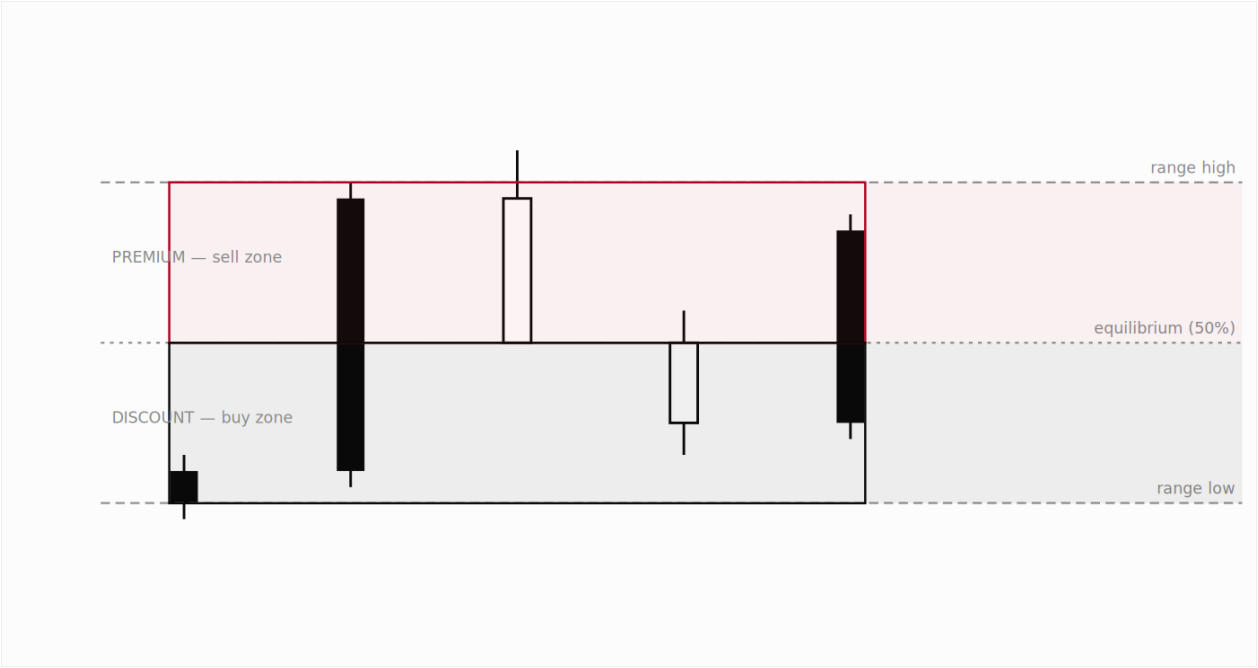

A dealing range splits at equilibrium into a premium half (expensive — sell) and a discount half (cheap — buy). The core filter: buy in discount, sell in premium.

KEY IDEA

Buy in discount, sell in premium. This is the institutional principle of getting a good price, reduced to a rule you can apply on any chart. A perfect PD array in the wrong half of the range is a trade to skip.

Combining with bias

Premium and discount become powerful when combined with your directional bias. If your higher-timeframe read is bullish and price is sitting in discount, everything aligns — you look for longs from the cheap half toward the liquidity above. If you are bullish but price is in premium, you wait; chasing into premium is how good directional reads turn into bad entries. The range gives your bias a location; the bias gives the range a direction. Together they tell you not just which way, but from where.

Ranges within ranges

Like structure, dealing ranges are fractal. A large higher-timeframe range contains smaller lower-timeframe ranges, each with its own equilibrium. A price that is in discount on the daily might be in premium on the five-minute during a pullback. This is not a contradiction — it is precisely the nuance that lets you time entries. You look for a lower-timeframe discount within a higher-timeframe discount to buy, stacking the odds of a good price.

KEY POINTS

- A dealing range is the span between a significant swing low and high — think ranges, not lines
- Split at equilibrium (50%): upper half is premium (sell), lower half is discount (buy)
- The core filter: buy in discount, sell in premium — skip good arrays in the wrong half
- Combine with bias: a bullish read plus a discount price is full alignment; bullish in premium means wait
- Ranges are fractal — seek a lower-timeframe discount within a higher-timeframe discount

COMMON MISTAKES

Buying in premium or selling in discount — chasing price to a bad location

Ignoring the range and thinking only in support/resistance lines

Forgetting to combine premium/discount with directional bias

PRACTICE DRILLS

- ○ Mark a dealing range on three charts, draw equilibrium, and label premium and discount.
- ○ For each, state where you'd look to buy or sell and why.
- ○ Find a case where price was in higher-timeframe discount but lower-timeframe premium during a pullback.

RECAP

A dealing range spans significant swing points and splits at equilibrium into a premium half to sell from and a discount half to buy from. Buying in discount and selling in premium is the institutional principle of a good price as a simple rule, and combining it with bias tells you not just which direction to trade but from where.

Builds on: Ch 7, Ch 8

CHAPTER 10 · THE LANGUAGE OF PRICE

Displacement & Delivery

When institutions commit, price doesn't drift — it displaces. Learning to see displacement is how you tell genuine intent from noise, and it's the engine that creates the PD arrays you'll trade in Part IV.

IN THIS CHAPTER YOU WILL

- → Define displacement and distinguish it from ordinary movement
- → Recognise displacement as the signature of institutional commitment
- → See how displacement creates fair value gaps and confirms structure breaks

What displacement is

Displacement is a strong, purposeful move — a run of large-bodied candles travelling decisively in one direction, usually leaving little overlap between them. It is the visual signature of institutions committing size. Ordinary price movement is choppy and two-sided; displacement is one-sided and urgent. When you see it, you are seeing intent, not noise. Learning to distinguish the two is one of the highest-value reads in this entire course.

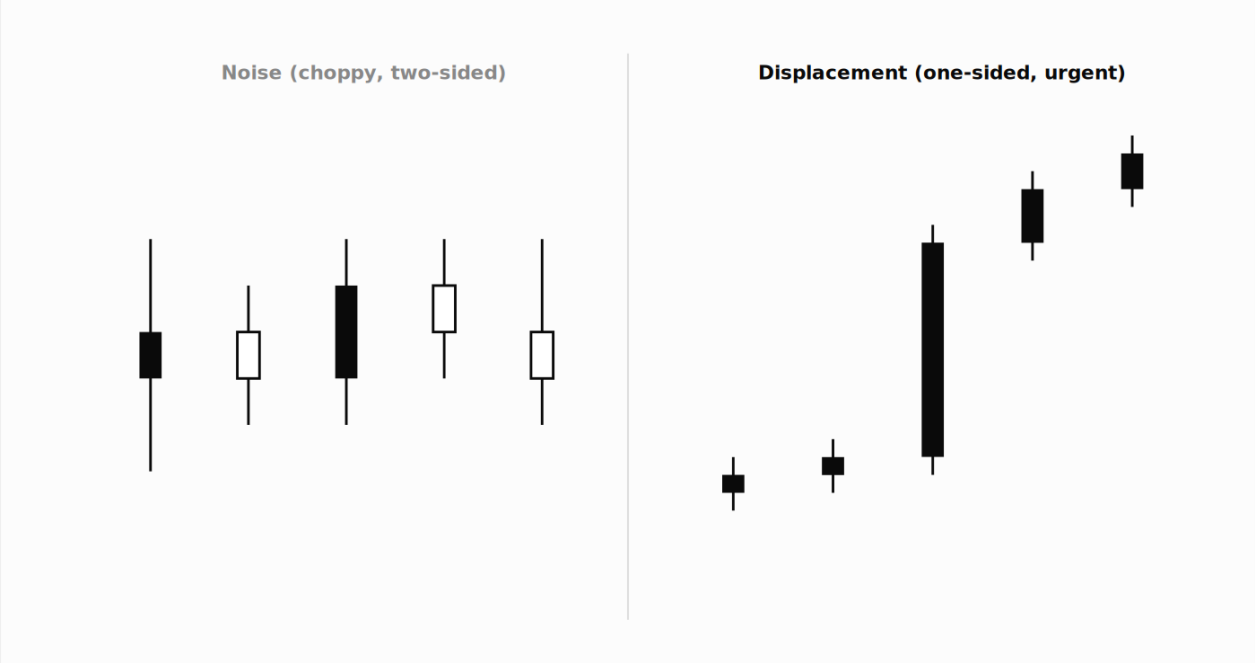

Displacement is a strong, one-sided run of large-bodied candles — the signature of institutional commitment — versus choppy, overlapping noise. A structure break on displacement is intent; without it, often a trap.

Why displacement matters so much

Displacement matters for three reasons, each of which you will use repeatedly. First, it confirms intent: a structure break on displacement is far more trustworthy than a limp,

overlapping break. Second, it creates fair value gaps — the speed of the move leaves imbalances in price that become your entry zones (Part IV). Third, it reveals direction: the side displacement favours is the side institutions are committing to, which informs your bias. A single displacement leg often tells you more than an hour of ranging.

> KEY IDEA
>
> Whenever you are about to trust a move — a break of structure, a reversal, a change of character — ask: did it happen on displacement? A break with displacement is intent. A break without it is often a trap that will be reclaimed.

Displacement and the structure break

Recall the CHoCH from the previous chapter — the first break against the trend. A CHoCH on displacement is a genuine warning of reversal; a CHoCH on weak, overlapping candles is far less reliable and frequently fails. The same is true of a BOS. This is why displacement is taught here, before the models: it is the quality filter you apply to every structural event. Strong break, trust it; weak break, doubt it.

The gap it leaves behind

Because displacement moves so fast, it often skips price levels that would normally trade — leaving a fair value gap, an imbalance between where one candle ended and the next began. You will study these in depth in Part IV, but plant the connection now: displacement is the cause, the fair value gap is the effect, and price frequently returns to that gap to rebalance before continuing. Every time you see a sharp displacement leg, you should already be looking for the gap it left — that gap is a future entry.

KEY POINTS

- Displacement is a strong, one-sided run of large candles — the signature of institutional commitment
- It matters because it confirms intent, creates fair value gaps, and reveals direction
- Apply it as a quality filter: a structure break ON displacement is trustworthy; without it, doubtful
- Displacement is the cause; the fair value gap it leaves is the effect and a future entry
- A single displacement leg often tells you more than an hour of choppy price

COMMON MISTAKES

Trusting a structure break that happened on weak, overlapping candles rather than displacement

Failing to look for the fair value gap a displacement leg leaves behind

Confusing ordinary choppy movement with genuine displacement

PRACTICE DRILLS

- Find five displacement legs and five choppy moves; articulate what makes each what it is.
- For each displacement leg, mark the fair value gap it left.
- Find a structure break on displacement and one without — compare how each resolved.

RECAP

Displacement is a strong, purposeful, one-sided move that signals institutions committing — it confirms intent, reveals direction, and leaves fair value gaps behind. Use it as the quality filter on every structural event: trust breaks that happen on displacement, doubt those that don't, and always look for the gap a displacement leg leaves as a future entry.

Builds on: Ch 8, Ch 9

PART III

Liquidity — The Engine

CHAPTER 11 · LIQUIDITY — THE ENGINE

What Liquidity Actually Is

Liquidity is the single most important concept in this course — the engine behind every move. If Part II taught you to read the skeleton, this chapter shows you the muscle that moves it.

IN THIS CHAPTER YOU WILL

- → Define liquidity in the specific way ICT uses it
- → See liquidity as resting orders, not abstract 'volume'
- → Understand why identifying liquidity lets you anticipate direction

Liquidity means resting orders

In everyday finance 'liquidity' means how easily you can trade. In ICT it means something more specific and more useful: pools of resting orders that price can be delivered to. These are real orders sitting in the market — stop-losses waiting to be triggered, and pending orders waiting to be filled. They cluster at predictable places, and because institutions need them to fill size, price is drawn toward them. When you learn to see liquidity, you stop seeing a random chart and start seeing a map of destinations.

Where the orders rest

Resting orders concentrate at obvious points. Above an old high sit buy orders: the stop-losses of traders who are short (a stop on a short is a buy), and the entry orders of breakout buyers. Below an old low sit sell orders: the stops of traders who are long, and the entries of breakout sellers. Equal highs or equal lows are especially rich pools, because every touch adds more orders at the same level. These clusters are the liquidity — and their locations are visible to you right now on any chart, in the form of prior highs, prior lows, and equal levels.

> KEY IDEA
>
> Reframe the whole chart: every prior high is a pool of buy-side liquidity, every prior low a pool of sell-side liquidity. Price is not wandering — it is travelling between these pools. Your job is to identify which pool is the current destination.

Why this lets you anticipate

Here is the payoff. If price is a delivery vehicle travelling toward liquidity, then identifying the liquidity tells you the likely destination before price gets there. Instead of reacting to moves, you form an expectation: 'there is untapped liquidity above; price is likely to be drawn up to it.' That expectation is testable, it gives your trades a target, and it is the foundation of the 'draw on liquidity' you will formalise at the end of this part. Everything — sweeps, PD arrays, models

— is in service of getting price to liquidity.

Liquidity is the 'why' behind the tools

As you move through the rest of this course, keep liquidity central. A PD array is only worth trading if it sits between price and a meaningful liquidity target. A time window matters because it is when price gets delivered toward liquidity. A model works because it sequences a liquidity grab with a structural confirmation. If you ever feel lost in the mechanics, return to the one question that organises all of it: where is the liquidity, and is this setup taking price toward it?

KEY POINTS

- In ICT, liquidity means pools of resting orders — real stops and pending orders price is drawn to
- Buy-side liquidity rests above old highs; sell-side rests below old lows; equal levels are rich pools
- Every prior high and low is a liquidity destination — the chart is a map of pools
- Identifying liquidity lets you anticipate the destination before price arrives, giving trades a target
- Liquidity is the 'why' behind every tool — always ask if a setup takes price toward it

COMMON MISTAKES

Thinking of liquidity as vague 'volume' rather than specific resting orders at specific levels

Trading a setup that isn't actually taking price toward any meaningful liquidity

Forgetting to identify the liquidity target before entering

PRACTICE DRILLS

- ○ On a chart, mark every significant prior high and low as a liquidity pool.
- ○ Identify one untapped pool that price seems to be drawn toward, and why.
- ○ For a recent move, name the liquidity it was reaching for.

RECAP

Liquidity in ICT means pools of resting orders — stops and pending orders that cluster above old highs and below old lows. Price is delivered toward these pools, so identifying them turns the chart into a map of destinations and lets you anticipate direction. It is the 'why' behind every other tool in the course.

Builds on: Ch 2, Ch 3

CHAPTER 12 · LIQUIDITY — THE ENGINE

Buy-Side & Sell-Side Liquidity

The two kinds of liquidity — buy-side and sell-side — are the poles price travels between. Getting crisp on which is which, and whose orders sit where, is what makes the rest of Part III click.

IN THIS CHAPTER YOU WILL

- → Distinguish buy-side from sell-side liquidity precisely
- → Identify whose orders form each pool
- → Read a chart as an alternation between the two

Buy-side liquidity

Buy-side liquidity rests above price — above old highs, above equal highs, above obvious resistance. It is made of buy orders: the stop-losses of short-sellers (closing a short means buying) and the buy-stop entries of breakout traders. When price runs up into this zone, all those buy orders fire at once. Counter-intuitively, that burst of buying is often the exit liquidity that lets larger players sell — which is why a run into buy-side liquidity so often precedes a move down.

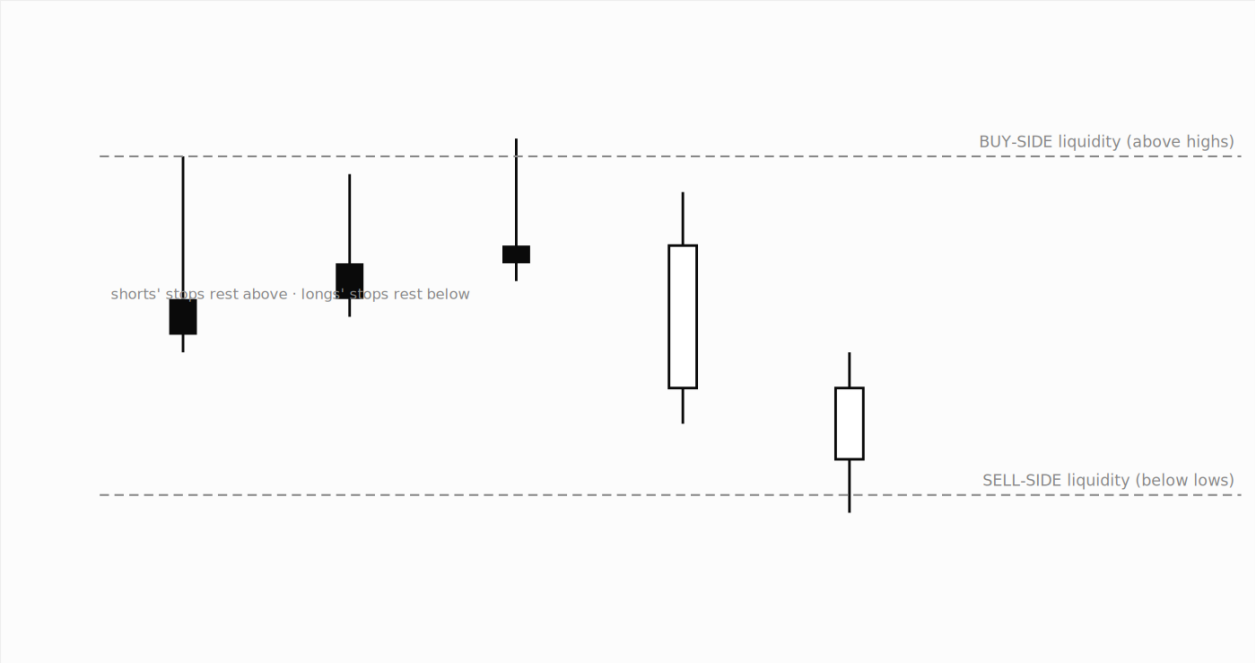

Buy-side liquidity rests above old highs (shorts' stops and breakout buys); sell-side rests below old lows (longs' stops and breakdown sells). Price hunts each in turn.

Sell-side liquidity

Sell-side liquidity rests below price — below old lows, equal lows, obvious support. It is made of sell orders: the stop-losses of longs (closing a long means selling) and the sell-stop entries of breakdown traders. When price runs down into this zone, those sell orders fire, providing the buying opportunity larger players need to accumulate longs. A run into sell-side liquidity often precedes a move up.

> KEY IDEA
>
> The counter-intuitive core: a run UP into buy-side liquidity often precedes a move DOWN, and a run DOWN into sell-side liquidity often precedes a move UP. Price reaches for the liquidity, takes it, and reverses. This is why chasing breakouts fails and why patience at liquidity pays.

The alternation

Once both are clear, a chart resolves into an alternation. Price reaches up to take buy-side liquidity, reverses, travels down to take sell-side liquidity, reverses again. Each raid provides the fuel for the next leg. Reading this rhythm — up to grab buy-side, down to grab sell-side — is reading the market's fundamental breathing. Your entries will typically come right after one of these raids, positioned for the reversal that follows, in the direction of your higher-timeframe bias.

Whose stops, whose fuel

Always keep track of whose orders you are watching. When price sweeps buy-side liquidity above a high, the short-sellers who placed stops there are being stopped out — their forced buying is the fuel. Simultaneously, breakout buyers are entering long at the worst possible price, about to be trapped. If you can name who is being hurt at a liquidity level, you can usually name the direction price will go next: away from the trapped crowd.

KEY POINTS

- Buy-side liquidity rests above old highs — shorts' stops and breakout buy-stops
- Sell-side liquidity rests below old lows — longs' stops and breakdown sell-stops
- A run UP into buy-side often precedes a move DOWN; a run DOWN into sell-side often precedes UP
- Charts alternate: up to grab buy-side, down to grab sell-side — each raid fuels the next leg
- Name who is trapped at a level and you can usually name the next direction

COMMON MISTAKES

Mixing up which liquidity sits where, or whose orders form it

Chasing a run into liquidity instead of waiting for the reversal it often precedes

Ignoring the trapped crowd whose forced orders provide the fuel

PRACTICE DRILLS

- ○ Label buy-side and sell-side pools on three charts and name whose orders form each.
- ○ Find a run into buy-side liquidity that preceded a move down, and the mirror case.
- ○ For one liquidity raid, write down who was trapped and which way price then went.

RECAP

Buy-side liquidity sits above old highs (shorts' stops, breakout buys) and sell-side below old lows (longs' stops, breakdown sells). Price alternates between raiding each pool, and a run into liquidity often precedes a reversal away from it — so naming who's trapped at a level tells you the likely next direction.

Builds on: Ch 11

CHAPTER 13 · LIQUIDITY — THE ENGINE

Liquidity Pools, Sweeps & Runs

This is where liquidity becomes tradeable. The difference between a sweep and a run — a wick past a level versus a close through it — is one of the most valuable distinctions you will ever learn to read.

IN THIS CHAPTER YOU WILL

- → Identify a liquidity pool from equal highs/lows and obvious levels
- → Distinguish a sweep (reversal) from a run (continuation)
- → Use sweeps as the trigger context for reversal entries

Spotting a pool

A liquidity pool is a concentration of orders at an obvious level. The richest pools form at equal highs or equal lows — where price has touched the same level repeatedly, stacking stops and pending orders each time. A double or triple top is not just a pattern; it is a visible advertisement of a buy-side pool sitting just above it. When you see equal levels, you are seeing exactly where a large player might send price to harvest orders.

Sweep versus run — the crucial distinction

When price reaches a pool, one of two things happens, and telling them apart is essential. A sweep (also called a raid or a stop hunt) pierces the level with a wick and then reverses — price pokes past the equal highs, triggers the orders, and closes back below. The liquidity was taken; there was no genuine intent to go higher. A run (a genuine breakout) closes decisively through the level and continues — the level is broken, not raided, and often flips to support or resistance.

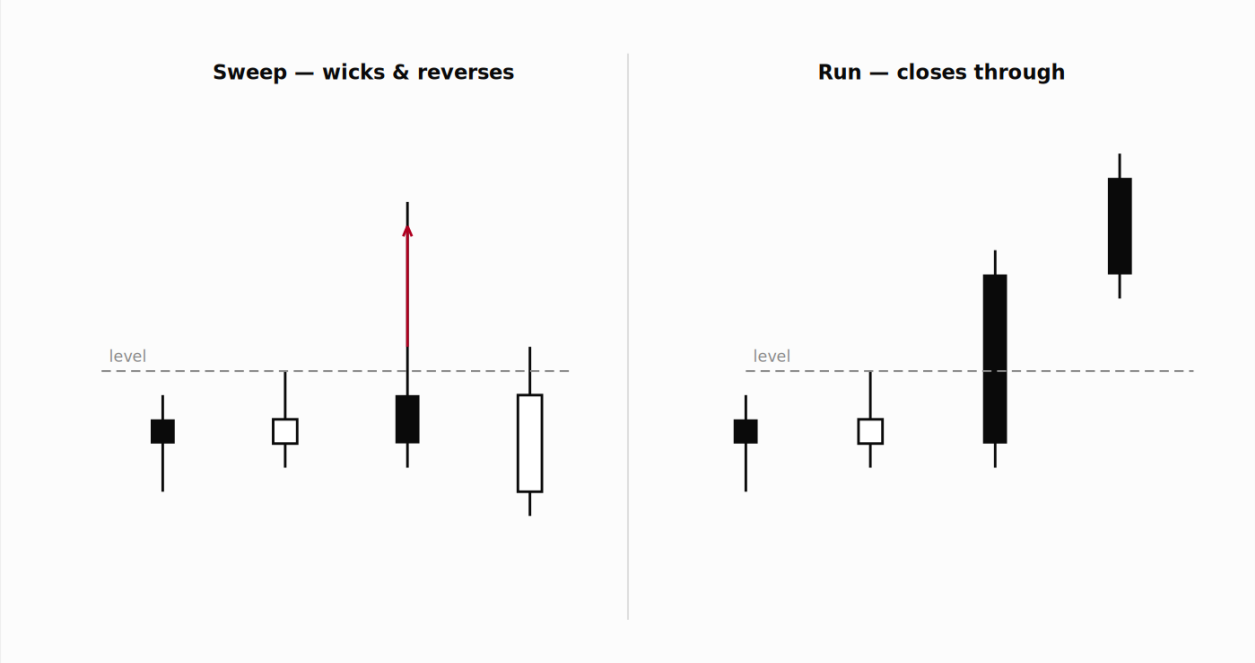

The test is the close, not the touch. A sweep wicks past the level and closes back (liquidity taken, expect reversal). A run closes decisively through and continues (level broken, expect continuation).

KEY IDEA

The test is the close, not the touch. Did price merely WICK past the level and close back (sweep → expect reversal), or did it CLOSE through and hold (run → expect continuation)? This one distinction will keep you out of countless bad trades and put you into countless good ones.

Why sweeps set up reversals

A sweep is the manipulation an institution needs: it grabs the resting orders, traps the breakout crowd, and provides the fuel to reverse. That is why a sweep is the starting gun for so many reversal setups. But a sweep alone is not an entry — price could still continue. In Part IV and V you will pair the sweep with a structural confirmation (a change of character on displacement) and a PD array entry. For now, plant the sequence firmly: liquidity swept, then reversal confirmed. The sweep provides the reason; the confirmation provides the trigger.

Reading the wick

Practically, sweeps announce themselves through wicks. A long wick poking through an obvious level and closing back is the visual fingerprint of a raid — orders triggered, price rejected. This is why Chapter 6 insisted you read wicks as rejection: at a liquidity pool, a rejection wick is the market telling you the pool has been harvested and the move that reached for it was not the real one. Train your eye to flag long wicks at equal highs and lows above all else.

KEY POINTS

- A liquidity pool is a cluster of orders at an obvious level; equal highs/lows are the richest
- A sweep wicks past a level and reverses (liquidity taken, no real intent) — expect reversal
- A run closes decisively through and continues (level broken, often flips) — expect continuation
- The test is the close, not the touch: wick-and-close-back vs close-through-and-hold
- A sweep is the starting gun for reversals, but needs confirmation — 'swept, then confirmed'

COMMON MISTAKES

Reading a sweep as a breakout (or vice versa) — judging by the touch instead of the close

Entering on the sweep alone without waiting for structural confirmation

Overlooking equal highs/lows, the most obvious liquidity pools on the chart

PRACTICE DRILLS

- ○ Find five equal-high or equal-low pools and mark them before price reaches them.
- ○ Collect five clear sweeps and five clear runs; for each, point to the close that decided it.
- ○ For one sweep, mark where confirmation (a CHoCH on displacement) appeared afterward.

RECAP

Liquidity pools cluster at obvious levels, especially equal highs and lows. A sweep wicks past the level and reverses — liquidity taken with no real intent — while a run closes through and continues; the close, not the touch, decides which. A sweep is the starting gun for reversals but must be paired with confirmation before you enter.

Builds on: Ch 12, Ch 10

CHAPTER 14 · LIQUIDITY — THE ENGINE

Internal vs External Range Liquidity

Liquidity doesn't only sit at the edges of a range — it also sits inside it. Understanding how price rotates between internal and external liquidity tells you what it's reaching for next.

IN THIS CHAPTER YOU WILL

- → Distinguish internal range liquidity from external range liquidity
- → Understand the rotation between them
- → Use the rotation to predict the next target

Two kinds of liquidity in every range

Within any dealing range there are two categories of liquidity. External range liquidity (ERL) sits at the edges — the highs and lows of the range, where stops rest beyond the extremes. Internal range liquidity (IRL) sits inside the range — principally the fair value gaps and imbalances left by moves within it. External liquidity is the obvious stuff at the boundaries; internal liquidity is the imbalance in the middle that price returns to rebalance.

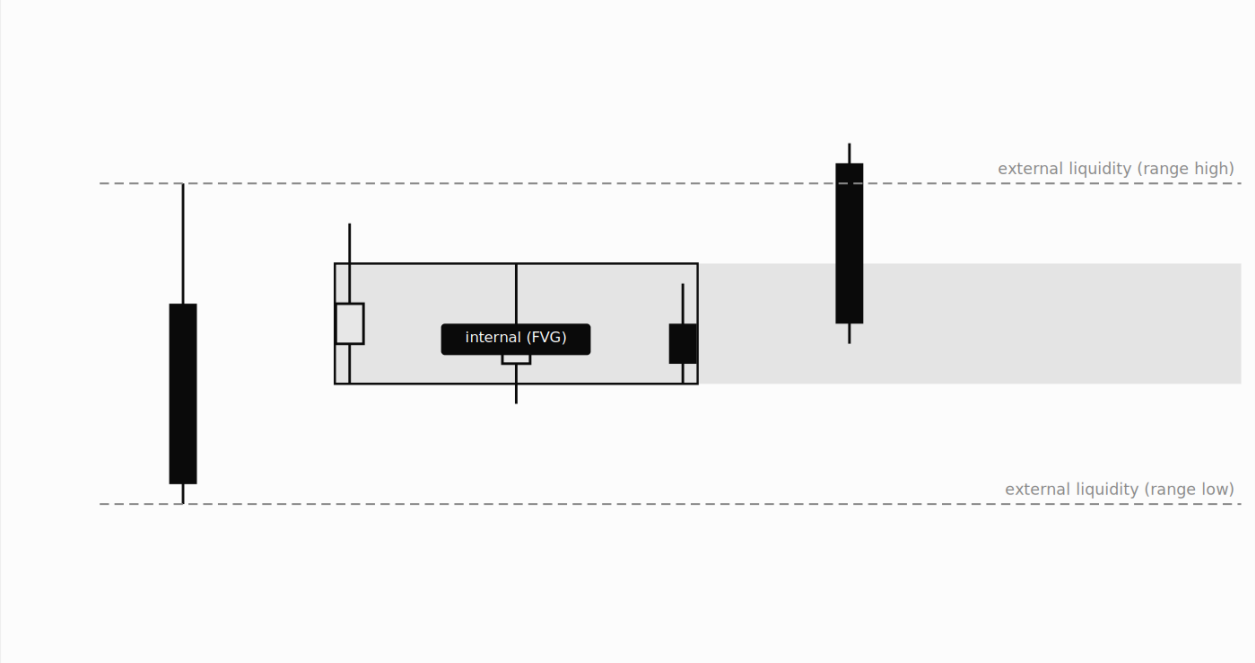

External range liquidity sits at the edges (the highs and lows); internal range liquidity sits inside (fair value gaps). Price rotates between them — taking an edge, returning to an internal gap, then pushing to the opposite edge.

The rotation

Price tends to rotate between the two. After taking external liquidity at one edge (say, sweeping the range high), price often turns and delivers back to internal liquidity — filling a fair

value gap inside the range. Having rebalanced that internal liquidity, it may then push to take external liquidity at the opposite edge. The delivery cycle reads: external, then internal, then external again. Recognising which one price is currently seeking tells you the immediate target.

> KEY IDEA
>
> Ask at any moment: has price just taken external liquidity (an edge)? Then the next draw is often internal (a gap inside). Has it just filled internal liquidity? Then the next draw is often external (the opposite edge). This rotation gives you a running expectation of the next target.

Why this matters for entries

The rotation is not academic — it shapes where you enter and where you target. If price has just swept the range high (external) and left a fair value gap on the way down, that gap (internal) becomes both a likely destination and, once reached, a potential entry toward the opposite edge (external). Your target for that entry is the external liquidity on the far side. In other words, IRL and ERL frequently serve as entry and target for one another. Reading the rotation lets you plan the whole trade — from which internal gap to enter, toward which external pool to target.

Fitting it together

This chapter quietly ties Part III to Part IV. The internal liquidity you will target and enter from is made of the fair value gaps you will study next. The external liquidity is the buy-side and sell-side pools you just learned. The draw on liquidity — the next chapter — is your judgement about which of these is the current destination. Hold the rotation in mind as you learn PD arrays: they are how you interact with internal liquidity, in service of reaching external liquidity.

KEY POINTS

- External range liquidity (ERL) sits at the range edges (highs/lows); internal (IRL) sits inside (gaps)
- Price rotates: take external at an edge, deliver back to internal, then push to external at the far edge
- Knowing which price just took tells you the next likely target
- IRL and ERL often serve as entry and target for each other — enter at a gap, target the opposite edge
- Internal liquidity is made of the fair value gaps you'll study in Part IV

COMMON MISTAKES

Watching only the range edges and missing the internal gaps price returns to

Failing to use the rotation to anticipate the next target

Targeting internal liquidity when price has just filled it and is due to seek external

PRACTICE DRILLS

- On a ranging chart, mark external liquidity at the edges and one internal gap inside.
- Identify which price most recently took, and predict the next target from the rotation.
- Trace one full rotation: external sweep, return to internal gap, push to opposite external.

RECAP

Every range holds external liquidity at its edges and internal liquidity (fair value gaps) inside. Price rotates between them — taking an edge, returning to an internal gap, then pushing to the opposite edge — and this rotation lets you anticipate the next target and often use internal and external liquidity as entry and target for one another.

Builds on: Ch 9, Ch 13

CHAPTER 15 · LIQUIDITY — THE ENGINE

The Draw on Liquidity: Where Price Is Going

This chapter turns everything in Part III into a single, decision-shaping judgement: the draw on liquidity — where price is most likely being delivered next. It is the question every trade begins with.

IN THIS CHAPTER YOU WILL

- → Define the draw on liquidity and why it's the master question
- → Weigh competing pools to judge the most likely destination
- → Use the draw to set direction and targets for every trade

The master question

The draw on liquidity is your judgement about which liquidity pool price is most likely being delivered toward right now. It is the master question of ICT trading, because it sets both your direction (you trade toward the draw) and your target (the draw is where you take profit). Every concept you have learned feeds into it: structure tells you the trend, premium/discount tells you if price is cheap or expensive, and the buy-side and sell-side pools are the candidate destinations. The draw is the synthesis.

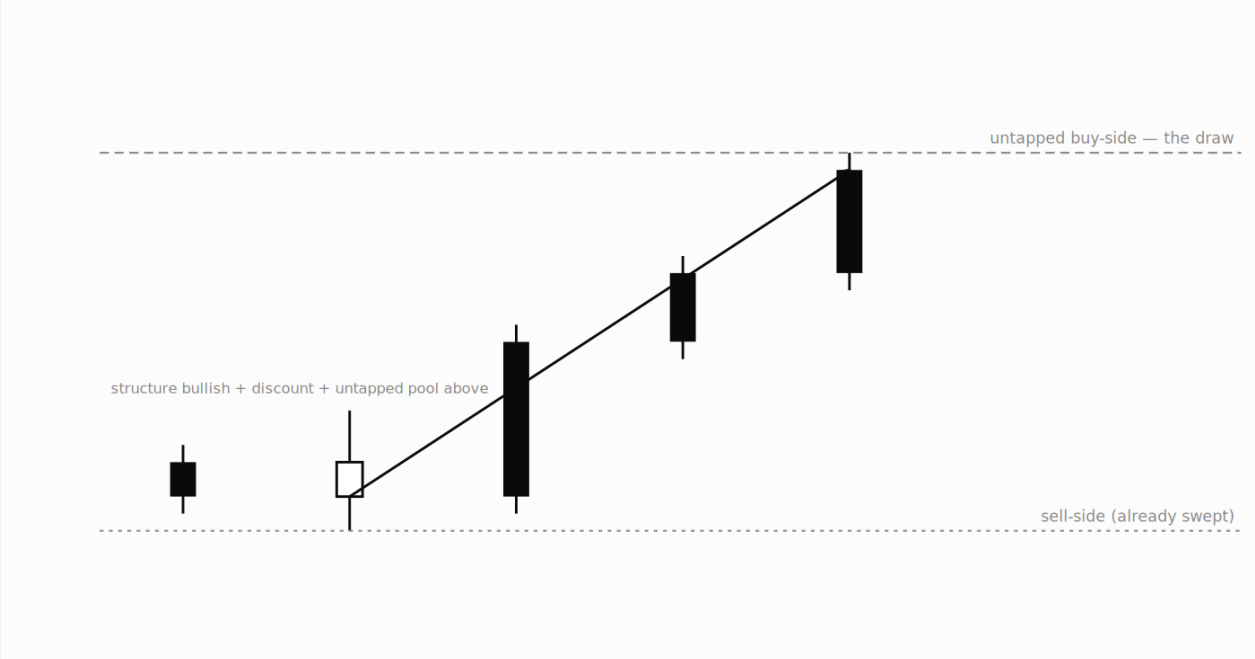

The draw on liquidity is your judgement of which pool price is being delivered toward. Here structure is bullish, price swept sell-side and reversed from discount, and buy-side above is untapped — the draw is up.

Weighing competing pools

There is usually liquidity in both directions — a pool above and a pool below. Judging the draw means deciding which is the more likely destination. Several things tip the balance. Structure: an uptrend biases the draw toward buy-side liquidity above. Premium/discount: from deep discount, the draw is more likely upward toward the range high. Untapped versus tapped: an obvious pool that has not yet been raided is a stronger magnet than one already swept. Recent sweeps: if price just took sell-side liquidity below and reversed, the draw has likely flipped to the buy-side above. You weigh these together into a single directional judgement.

> KEY IDEA
>
> Before every trade, state the draw out loud: 'price is in discount, structure is bullish, buy-side liquidity above is untapped — the draw is up, toward that high.' If you cannot state a clear draw, you do not yet have a trade. The draw is your thesis; the setup is merely how you enter it.

The draw sets direction and target

Once you have judged the draw, two of the hardest decisions are made for you. Direction: you look only for setups in the direction of the draw — if the draw is up, you take longs and ignore shorts. Target: you aim for the liquidity pool that is the draw. This is why traders who lead with the draw have such clean trades — they are not guessing at direction or targets in the heat of the moment; both were decided before they ever looked for an entry. The setup (Parts IV and V) is simply the precise way to join a move whose destination you already identified.

When the draw is unclear

Sometimes the draw is genuinely ambiguous — structure is ranging, price is near equilibrium, pools on both sides are untapped. The correct response is not to force a guess but to wait. Often the market resolves the ambiguity itself: a sweep of one side's liquidity followed by a reversal is the market telling you the draw is now toward the other side. Patience for a clear draw is one of the most profitable habits you can build. No clear draw, no trade — and there is no shame in an empty chart.

KEY POINTS

- The draw on liquidity is your judgement of which pool price is being delivered toward now
- It's the master question — it sets both your direction (trade toward it) and target (profit at it)
- Weigh structure, premium/discount, untapped vs tapped pools, and recent sweeps to judge it
- State the draw explicitly before every trade; no clear draw means no trade
- When the draw is ambiguous, wait — a sweep-and-reverse often resolves it for you

COMMON MISTAKES

- Entering a trade without first judging a clear draw on liquidity
- Ignoring evidence (recent sweeps, premium/discount) that points to one pool over the other
- Forcing a directional guess when the draw is genuinely ambiguous instead of waiting

PRACTICE DRILLS

- ○ On five charts, state the draw in one sentence citing structure, premium/discount, and untapped pools.
- ○ Find a case where a sweep flipped the draw to the opposite side, and describe it.
- ○ Practise identifying 'no clear draw' situations and consciously deciding not to trade them.

RECAP

The draw on liquidity is your synthesised judgement of where price is being delivered next, and it's the master question because it sets both direction and target before you look for an entry. Weigh structure, premium/discount, and which pools remain untapped; state the draw explicitly every time, and when it's ambiguous, wait rather than guess.

Builds on: Ch 12, Ch 14, Ch 9

PART IV

PD Arrays — The Toolkit

CHAPTER 16 · PD ARRAYS — THE TOOLKIT

PD Arrays: The Institutional Footprint

PD Arrays are the footprints institutions leave in price. This chapter frames the whole toolkit before you meet each tool individually, so you understand what they have in common and how to think about them.

IN THIS CHAPTER YOU WILL

- → Understand what a PD array is and why they exist
- → See the common logic behind every array
- → Know how arrays relate to liquidity and premium/discount

What a PD array is

PD Array stands for Price Delivery Array — a structure left in price by institutional activity, which price tends to react to when it returns. Order blocks, fair value gaps, breakers, and mitigation blocks are all PD arrays. They are not indicators or patterns imposed on the chart; they are the actual footprints of large orders being placed, visible after the fact and reliable enough to trade from when read in context.

The common logic

Every PD array shares one underlying idea: it marks a price level where institutions have unfinished business. An order block marks where they entered; a fair value gap marks an imbalance their speed created; a breaker marks a level that failed and flipped. In each case, when price returns to the array, there is a reason to expect a reaction — orders to be filled, imbalance to be rebalanced, a flipped level to hold. Understanding this shared logic means you are never memorising disconnected patterns; you are recognising variations on one theme.

> KEY IDEA
>
> A PD array is a location with a reason. The reason is always some form of unfinished institutional business — an entry to defend, an imbalance to fill, a level that flipped. When you trade an array, you are betting that business gets attended to when price returns.

Arrays live inside the framework you've built

A PD array is only worth trading in the context of Parts II and III. It must sit in the right half of the dealing range — you buy from arrays in discount, sell from arrays in premium. It must point price toward a meaningful draw on liquidity — an array that leads nowhere is not a trade. And it is strongest when it aligns with structure and follows a liquidity sweep. The array is the entry

mechanism; the draw, the structure, and the premium/discount location are what make that entry worth taking.

How to approach the next chapters

You will now meet the arrays one at a time: order blocks, fair value gaps, then breakers and mitigation blocks, then how to choose between them. Resist the urge to collect them like trophies. The goal is not to know ten arrays; it is to reliably trade two or three in the right context. As you learn each, keep asking the framework questions — which half of the range is this in, what liquidity does it point toward, does it follow a sweep? An array that cannot answer those is not yet a trade.

KEY POINTS

- A PD array (Price Delivery Array) is an institutional footprint price reacts to on return
- Order blocks, fair value gaps, breakers, and mitigation blocks are all PD arrays
- Common logic: every array marks a level with unfinished institutional business
- Arrays are only tradeable in context — right half of the range, pointing toward a draw, ideally after a sweep
- Aim to trade two or three arrays reliably, not to collect many

COMMON MISTAKES

Trading an array in isolation, ignoring premium/discount and the draw on liquidity

Collecting many arrays instead of mastering a few in context

Forgetting that the array is only the entry — the draw and structure justify the trade

PRACTICE DRILLS

- ○ Write the shared logic of PD arrays in one sentence.
- ○ On a chart, find any array and check it against the three framework questions.
- ○ List which two or three arrays you intend to focus on mastering.

RECAP

PD arrays are institutional footprints — order blocks, fair value gaps, breakers — that share one logic: they mark levels with unfinished institutional business. They're only tradeable in context: in the right half of the range, pointing toward a draw on liquidity, ideally after a sweep. Master a few in context rather than collecting many.

Builds on: Ch 10, Ch 15, Ch 9

CHAPTER 17 · PD ARRAYS — THE TOOLKIT

Order Blocks

The order block is the most important PD array and the one you'll use most. Master it properly — how it forms, how to validate it, how to trade it — and you have a complete entry method on its own.

IN THIS CHAPTER YOU WILL

- → Identify a valid order block precisely
- → Understand why the last opposing candle matters
- → Trade an order block with a defined entry, stop, and target

What an order block is

An order block is the last opposing candle before a strong displacement move — the last down-candle before a sharp rally (a bullish order block), or the last up-candle before a sharp drop (a bearish order block). It marks the price zone where institutions placed the orders that caused the displacement. Because those orders represent genuine institutional interest, price frequently returns to the block and reacts from it — which is what makes it a high-quality entry.

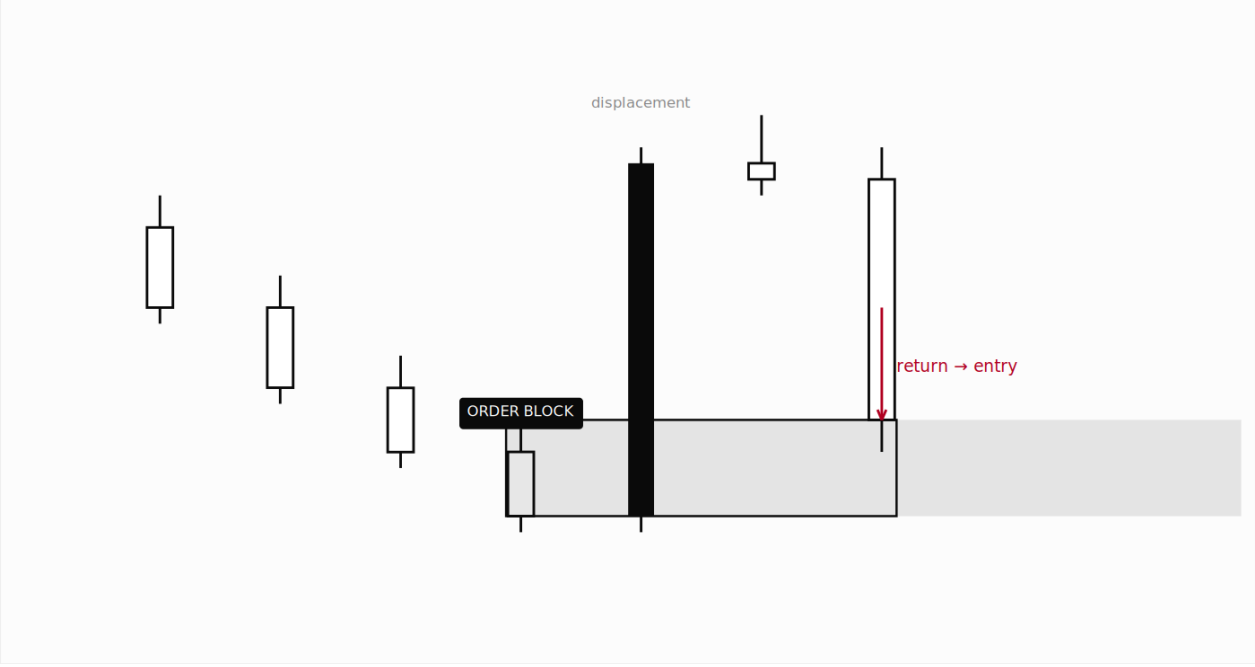

An order block is the last down-candle before strong displacement up. Price displaces away, then returns to the block — the return is the entry, stop just beyond its far side, target the draw.

Why the last opposing candle

The logic is precise. Before institutions drive price up, they must accumulate long positions — and they do so while price is still falling, in that last down-candle before the move. That candle

is the visible footprint of their accumulation. When price later returns to it, the institutions defend their entries, and price bounces. The same logic mirrors for bearish order blocks. This is why you look for the last opposing candle specifically: it is where smart money positioned before revealing its hand through displacement.

Validating an order block

Not every last-opposing candle is a valid order block. Three things raise the quality. First, the displacement that follows must be genuine — strong, large-bodied candles, ideally leaving a fair value gap. A weak move does not validate a block. Second, the block should ideally cause a break of structure, showing it produced real intent. Third, it is far stronger when it sits in the correct premium/discount location and follows a liquidity sweep. A bullish order block in discount, formed right after a sell-side sweep, that displaces up and breaks structure, is a textbook high-probability array.

> KEY IDEA
>
> Quality checklist for an order block: (1) strong displacement away from it, ideally with a fair value gap; (2) it caused a structure break; (3) correct premium/discount location; (4) follows a liquidity sweep. The more of these it satisfies, the more you can trust it.

How to trade it

The trade is clean. Your entry is the return of price into the order block — you can enter on the first touch of the block, or wait for a lower-timeframe confirmation inside it for more precision. Your stop goes just beyond the far side of the block: if price closes through the block, the institutional orders did not hold and your premise is void. Your target is the draw on liquidity — the opposing pool price is being delivered toward. Because the stop is tight (the block is small) and the target is the liquidity draw (often far), order block trades naturally produce good reward-to-risk.

KEY POINTS

- An order block is the last opposing candle before strong displacement (bullish = last down candle before a rally)
- It marks where institutions accumulated before revealing intent through the displacement
- Validate with: genuine displacement (ideally an FVG), a structure break, correct premium/discount, a prior sweep
- Entry = return into the block; stop = just beyond its far side; target = the draw on liquidity
- Tight stop plus a distant liquidity target gives order blocks naturally good reward-to-risk

COMMON MISTAKES

Marking any down/up candle as an order block without validating the displacement and context

Placing the stop inside the block instead of just beyond its far side

Trading a block that doesn't point toward a meaningful draw on liquidity

PRACTICE DRILLS

- Find five order blocks and run each through the four-point quality checklist.
- For one high-quality block, mark the exact entry, stop, and target.
- Backtest ten order-block returns and record how many held versus failed, and why.

RECAP

An order block is the last opposing candle before displacement, marking where institutions accumulated. Validate it with genuine displacement, a structure break, correct premium/discount, and a prior sweep. Enter on the return, stop just beyond its far side, and target the draw on liquidity — a structure that naturally yields strong reward-to-risk.

Builds on: Ch 16, Ch 10, Ch 15

CHAPTER 18 · PD ARRAYS — THE TOOLKIT

Fair Value Gaps & Imbalance

The fair value gap is the other essential array, and the most direct expression of imbalance. It's created by the very displacement you learned to read in Part II — and price is magnetically drawn back to fill it.

IN THIS CHAPTER YOU WILL

- → Identify a fair value gap from a three-candle sequence
- → Understand it as an imbalance price wants to rebalance
- → Trade an FVG entry and combine it with order blocks

What a fair value gap is

A fair value gap (FVG) is an imbalance left by displacement — specifically, the gap between the high of one candle and the low of the candle two positions later, created when a large middle candle moves so fast that price skips a range. In a bullish FVG, the first candle's high sits below the third candle's low, leaving an untraded gap. This gap represents inefficient delivery — price moved too fast to trade both sides fairly — and markets tend to return to rebalance it.

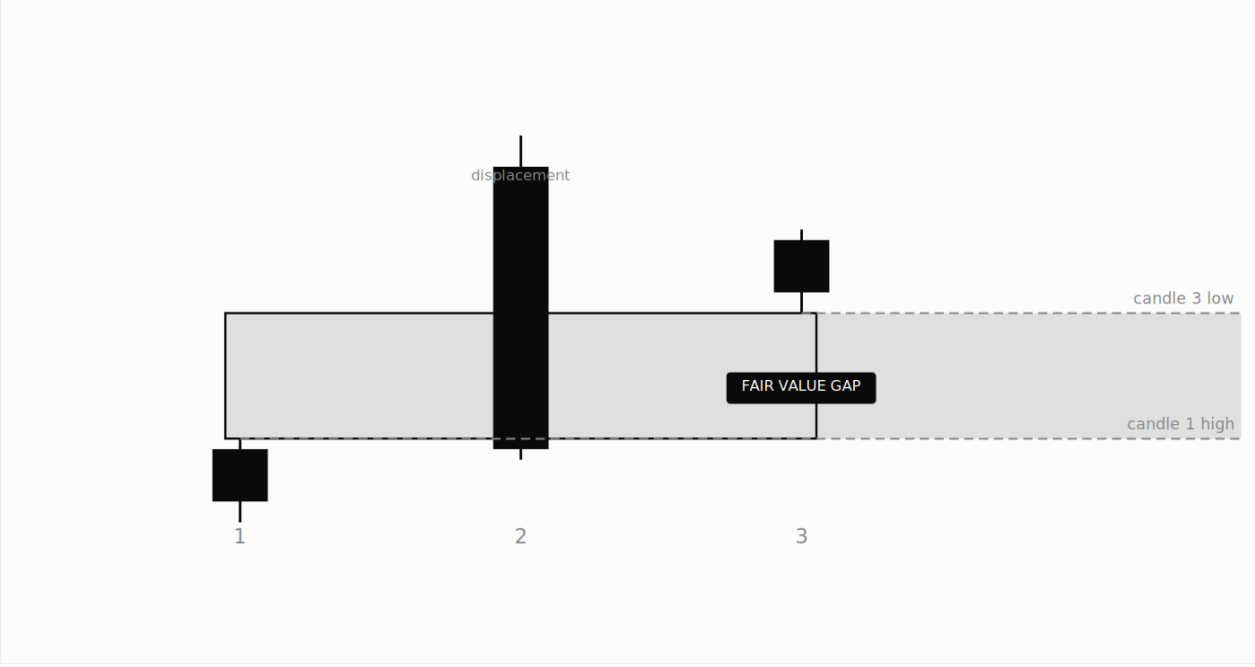

A fair value gap is the imbalance between candle 1's high and candle 3's low, left by the fast displacement of candle 2. Price returns to rebalance it before continuing — the gap is a precise entry.

Why price returns to fill it

The idea of rebalancing is central. When displacement leaves a gap, one side of the market did not get to transact there — the move was one-sided. Efficient markets dislike this imbalance,

and price frequently returns to the gap to let the missed orders fill before continuing in the original direction. This return is your opportunity: the gap is both a likely destination (internal range liquidity, from Part III) and, once reached, a precise entry in the direction of the original displacement.

> KEY IDEA
>
> An FVG is inefficiency made visible. Displacement created it by moving too fast; the market returns to fill it because the imbalance wants resolving. You enter as price fills the gap, betting it then continues in the displacement's direction.

How to trade it

Trading an FVG mirrors the order block. Your entry is price returning into the gap — often you enter at the edge of the gap or wait for price to partially fill it. Your stop goes beyond the gap: if price fills straight through and closes past it, the imbalance is fully resolved and your directional premise weakens. Your target is, again, the draw on liquidity. FVGs are especially valuable because they are objective — the three-candle sequence is either there or it isn't — which makes them one of the easiest arrays for a beginner to identify consistently.

Combining FVGs with order blocks

The two essential arrays are strongest together. Very often a genuine displacement leaves both an order block (the last opposing candle) and a fair value gap (the imbalance) right next to each other. When price returns to a zone that is simultaneously an order block and an FVG, you have two independent reasons for a reaction stacked at one price — a powerful confluence. Learning to spot where an order block and an FVG overlap is one of the highest-value skills in this part, and it will form the backbone of the entry model in Part V.

KEY POINTS

- A fair value gap is an imbalance from displacement — a gap between candle 1's high and candle 3's low
- It represents inefficient, one-sided delivery, and price returns to rebalance it
- The gap is both a likely destination (internal liquidity) and, once reached, a precise entry
- Entry = return into the gap; stop = beyond it; target = the draw on liquidity
- FVGs and order blocks often overlap — that confluence is a powerful, high-probability entry

COMMON MISTAKES

Miscounting the three-candle sequence and marking a gap that isn't there

Entering an FVG that doesn't point toward a meaningful draw

Missing the confluence when an FVG and order block overlap at the same price

PRACTICE DRILLS

- Find ten fair value gaps and confirm each with the three-candle rule.
- Mark a case where an order block and FVG overlap; note why that's stronger.
- For one FVG, define entry, stop, and target, then check how it resolved historically.

RECAP

A fair value gap is an imbalance left by displacement — the gap between candle 1's high and candle 3's low — that price returns to rebalance. It serves as both a destination and a precise entry toward the draw on liquidity, and it's especially powerful where it overlaps an order block, a confluence that anchors the entry model in Part V.

Builds on: Ch 10, Ch 17, Ch 14

CHAPTER 19 · PD ARRAYS — THE TOOLKIT

Breaker Blocks & Mitigation Blocks

Two more arrays complete your core toolkit: the breaker block and the mitigation block. Both involve price returning to a prior zone, but for different reasons — and knowing the difference sharpens your reads.

IN THIS CHAPTER YOU WILL

- → Define a breaker block and how a level flips polarity
- → Define a mitigation block and how it differs from a breaker
- → Trade both with appropriate context

The breaker block

A breaker block is a failed order block that flips polarity. Suppose a bullish order block forms, but instead of holding, price breaks down through it and shifts structure. That failure is information: the level that was support has now failed, and it often flips to resistance. The breaker is that flipped zone. When price retraces back up to it, you look for it to reject — now trading it as resistance rather than support. The defining feature of a breaker is the structure break that confirms the original level failed.

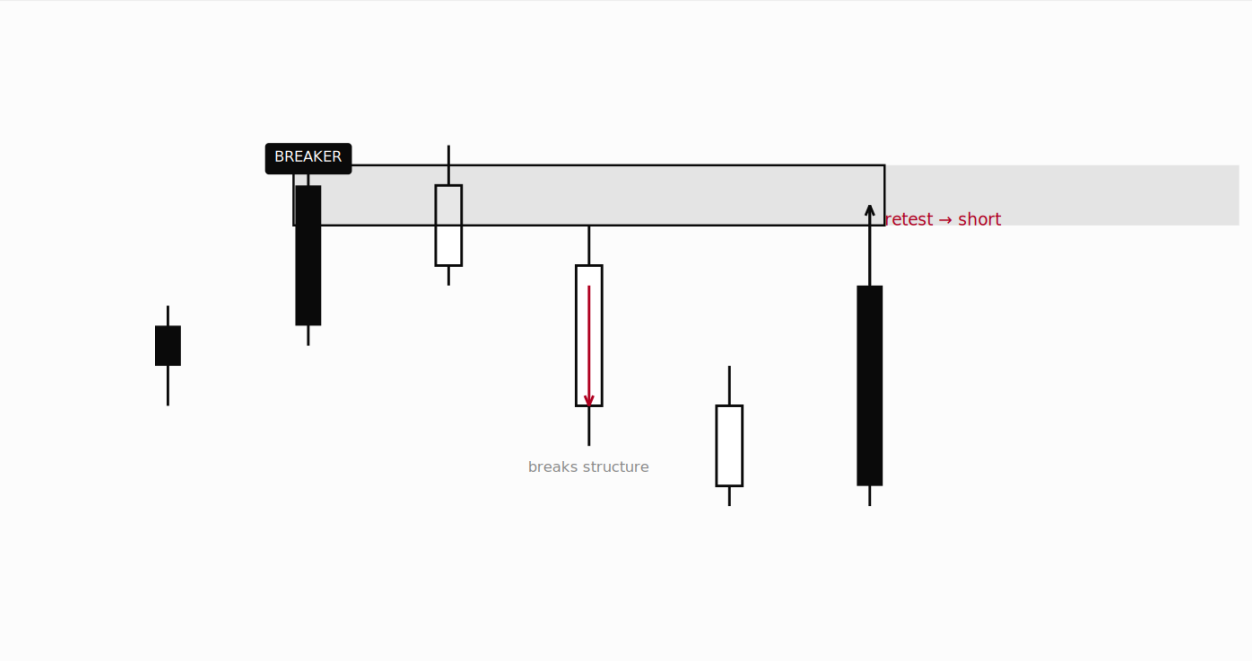

A breaker block is a failed order block that flips polarity. Price breaks structure through the old zone, which then flips from support to resistance — the retest into it is the entry.

The mitigation block

A mitigation block also involves price returning to a prior zone, but the mechanism is different. When institutions enter a position, part of their intended size may go unfilled as price runs away. A mitigation block is a zone price returns to so those institutions can mitigate — fill the remainder of their position or manage the trade — without a prior structural failure. Unlike a breaker, no structure break is required. The mitigation block is about completing unfinished orders; the breaker is about a level that flipped after failing.

> KEY IDEA
>
> The distinction in one line: a breaker is a level that FAILED and flipped polarity (needs a structure break); a mitigation block is a level institutions RETURN to in order to finish filling (no structure break needed). Both are entries on the return, but they arise from different stories.

How to trade them

Both are traded like the arrays you already know: entry on the return into the zone, stop beyond it, target the draw on liquidity. The breaker is particularly useful in reversals — because it forms after a structure break, it naturally appears when a trend is changing, giving you an entry in the new direction. The mitigation block is more of a continuation tool — price returns to complete institutional orders before carrying on. Knowing which story you are in tells you whether you are trading a reversal or a continuation, which in turn shapes your expectations for the trade.

Don't over-collect

With breakers and mitigation blocks you now have four arrays. Remember the warning from the overview: the goal is not to trade all four on every chart, but to recognise which array the current situation is offering and trade it well. In practice, order blocks and fair value gaps will be your bread and butter, with breakers appearing at reversals and mitigation blocks at continuations. Let the context tell you which array is in play rather than hunting for a specific one.

KEY POINTS

- A breaker block is a failed order block that flips polarity — it requires a structure break
- A mitigation block is a zone institutions return to to finish filling — no structure break needed
- Breakers suit reversals (they form after a structure break); mitigation blocks suit continuations
- Both are traded the same way: entry on return, stop beyond, target the draw
- Let context tell you which array is in play rather than hunting for one

COMMON MISTAKES

Confusing a breaker (failed and flipped) with a mitigation block (returning to fill)

Calling a zone a breaker without the confirming structure break

Hunting for one favourite array instead of trading whichever the context offers

PRACTICE DRILLS

- Find a breaker block and mark the structure break that confirmed the flip.
- Find a mitigation block and confirm there was no structure break — just a return to fill.
- For each, decide whether it implied a reversal or a continuation, and why.

RECAP

A breaker block is a failed order block that flipped polarity after a structure break — a reversal tool — while a mitigation block is a zone institutions return to in order to finish filling, with no break required, suiting continuations. Both are entries on the return toward the draw; let context tell you which is in play.

Builds on: Ch 17, Ch 8

CHAPTER 20 · PD ARRAYS — THE TOOLKIT

Choosing the Right Array & Confluence

You now know four arrays. The skill that separates competent traders from confused ones is choosing which to trade and stacking them into confluence rather than treating each as a lone signal.

IN THIS CHAPTER YOU WILL

- → Decide which array to trade in a given situation
- → Stack arrays and context into genuine confluence
- → Avoid the trap of over-analysis and conflicting signals

Confluence, not collection

Confluence means several independent reasons pointing to the same conclusion at the same price. It is the opposite of collecting signals. One order block is a reason to trade; an order block that overlaps a fair value gap, sits in discount, points toward untapped buy-side liquidity, and formed right after a sell-side sweep is confluence — five reasons stacked at one location. The more independent, aligned reasons at a single price, the higher the probability. Your job is not to find more arrays; it is to find where arrays and context agree.

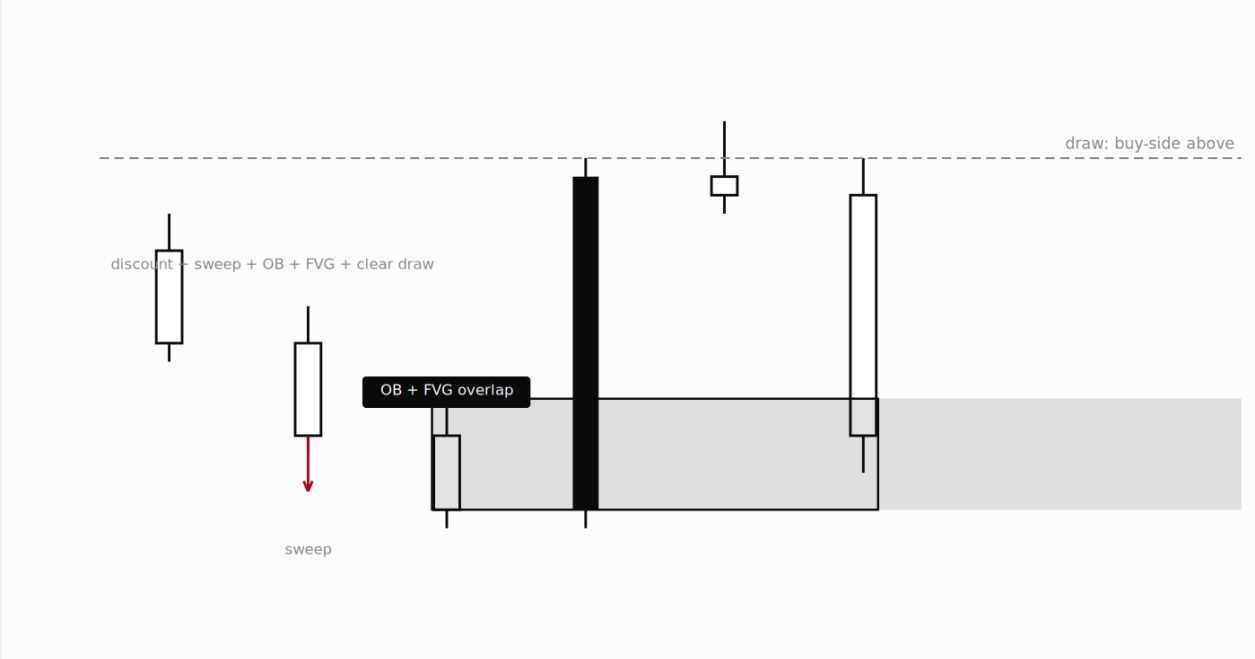

Confluence is several independent reasons at one price: an order block overlapping a fair value gap, in discount, after a sell-side sweep, pointing at untapped buy-side liquidity. The more aligned reasons, the higher the probability.

Which array to trade

When several arrays are present, prefer the one with the most context behind it. A hierarchy that serves beginners well: favour the array that (1) sits in the correct premium/discount half, (2) points toward the clearest draw on liquidity, (3) follows a clean liquidity sweep, and (4) overlaps another array. An order block that meets all four beats a lone fair value gap that meets one. When arrays conflict — one suggesting up, another down — that is the market telling you the picture is unclear, and the correct trade is often no trade.

> KEY IDEA
>
> Rank potential trades by how much context agrees, not by which array you prefer. The best trade on any chart is the one where premium/discount, the draw, a sweep, structure, and one or more arrays all point the same way. If they don't align, wait.

The danger of over-analysis

There is a failure mode at the opposite extreme from ignorance: over-analysis. Armed with four arrays and several context tools, a beginner can find a reason for any direction on any chart, paralysing themselves or, worse, justifying a trade they wanted to take anyway. Guard against this by leading with the framework in a fixed order — structure, draw, premium/discount, then array — and by demanding genuine alignment. If you find yourself hunting for an array to support a pre-decided direction, stop; you have reversed the process.

A simple decision routine

Here is a routine to apply on every chart until it becomes automatic. First, read structure — up, down, or ranging. Second, judge the draw on liquidity — where is price being delivered? Third, check premium/discount — is price positioned to enter cheaply toward that draw? Fourth, look for an array with a sweep in that location. Only when all four align do you have a trade. This routine, applied patiently, is most of what practical ICT trading actually is — and it sets up the models in Part V, which are simply this routine formalised around time.

KEY POINTS

- Confluence is several independent, aligned reasons at one price — not a collection of signals
- Prefer the array with the most context: correct premium/discount, clear draw, a sweep, and overlap
- Conflicting arrays mean the picture is unclear — often the correct trade is no trade
- Guard against over-analysis by leading with a fixed routine, not hunting to justify a bias
- Routine: structure → draw → premium/discount → array-with-sweep; trade only when all align

COMMON MISTAKES

- Treating arrays as a collection of signals rather than seeking genuine confluence
- Hunting for an array to justify a direction you'd already decided on
- Forcing a trade when arrays conflict instead of standing aside

PRACTICE DRILLS

- ○ On five charts, apply the four-step routine and note whether a trade qualifies.
- ○ Find one high-confluence setup and list every independent reason stacked at that price.
- ○ Catch yourself once trying to justify a pre-decided bias, and consciously reverse the process.

RECAP

Choosing arrays is about confluence — several independent, aligned reasons at one price — not collecting signals. Prefer the array with the most context behind it, treat conflicting arrays as a signal to wait, and guard against over-analysis by always leading with a fixed routine: structure, draw, premium/discount, then an array with a sweep. That routine is the foundation of the models in Part V.

Builds on: Ch 17, Ch 18, Ch 19, Ch 15

PART V

Time & The Models

CHAPTER 21 · TIME & THE MODELS

Time & Price: Why Timing Is Structure

Everything so far has been about price. Now we add the other half of the methodology: time. When a setup occurs matters as much as where — and understanding why turns good setups into high-probability ones.

IN THIS CHAPTER YOU WILL

- → Understand why time is a form of structure in ICT
- → See why the same setup varies in quality by time of day
- → Adopt time as a filter on every trade

Time is structure too

ICT is often summarised as 'time and price' — and until now you have only had price. The insight of the time half is that institutional activity is not spread evenly through the day. It concentrates at specific windows, driven by when major markets are open, when volume is high enough to hide large orders, and when the algorithm that delivers price is most active. Because of this, the same price setup is not equally reliable at all hours. Time is not a vague backdrop; it is a structural filter as real as premium and discount.

Why the same setup varies by time

Consider an identical order block, in discount, pointing at clear liquidity. Taken during a high-activity session when institutions are actually delivering price, it has a strong chance of working. Taken in a dead, low-volume period, the same block may simply be ignored — there is no one there to deliver price toward the draw. The setup did not change; the time did. This is why professional ICT traders are ruthless about when they trade: they know that price and time must both be right, and that a perfect setup at the wrong time is a low-probability trade.

> KEY IDEA
>
> Add one question to your routine: is this the right time? A high-quality price setup during a high-activity window is a trade; the same setup in a dead period is usually a pass. Time is a filter you apply to every potential entry.

The building blocks of time

The time half of ICT has a few key components, which the next chapters cover. Kill zones are the broad session windows when activity concentrates. Macros are short, precise windows within sessions when delivery is especially reliable. AMD (accumulation, manipulation,

distribution) is the rhythm a session tends to follow through time. Together these tell you not just where price is going, but when it is likely to go there — and that combination is the complete ICT read.

Time makes the models possible

The reason time is taught here, right before the models, is that the models are essentially price setups anchored to time windows. The Silver Bullet is an entry model that only operates in a specific hour. The complete entry model in the final chapter of this part combines a liquidity sweep, a structure shift, and an array entry — but insists they occur within a high-probability time window. Without the time dimension, you have setups; with it, you have models. That is the step you are about to take.

KEY POINTS

- ICT is 'time and price' — institutional activity concentrates at specific windows, not evenly
- The same price setup varies in reliability by time; a perfect setup at the wrong time is low-probability
- Add 'is this the right time?' to your routine as a filter on every entry
- Time's building blocks: kill zones (sessions), macros (precise windows), AMD (session rhythm)
- Models are price setups anchored to time windows — time is what turns setups into models

COMMON MISTAKES

Treating a setup as equally valid at any hour, ignoring the time filter

Trading in dead, low-activity periods where price isn't being delivered

Thinking of time as a vague backdrop rather than a real structural filter

PRACTICE DRILLS

- ○ Note the times of your last several setups and whether they were in active windows.
- ○ Compare one setup that worked in an active window with a similar one that failed in a dead period.
- ○ Add the time question to your written routine.

RECAP

Time is the second half of ICT: institutional activity concentrates at specific windows, so the same price setup varies in reliability by time. Adding 'is this the right time?' as a filter — via kill zones, macros, and the AMD rhythm — is what turns setups into models, the step Part V now takes.

Builds on: Ch 2, Ch 20

CHAPTER 22 · TIME & THE MODELS

Kill Zones & Trading Sessions

Kill zones are the windows when the market comes alive. Knowing them tells you when to be at your screen and when to stay away — the simplest, highest-value application of time.

IN THIS CHAPTER YOU WILL

- → Identify the major kill zones and what happens in each
- → Know which windows are highest-probability
- → Restrict your trading to the right hours

What a kill zone is

A kill zone is a session window during which institutional activity — and therefore the clean, tradeable delivery of price — concentrates. The major kill zones correspond to the opens of the world's main trading sessions and their overlaps, because that is when volume and participation peak. Outside these windows, price often drifts or ranges aimlessly; inside them, the liquidity raids, displacement, and structure shifts you have learned to read tend to occur with clarity.

Kill zones are the session windows where institutional activity concentrates. The London and New York AM windows carry the highest participation and produce the cleanest liquidity runs of the day.

The main windows

There are four kill zones to know. The Asian session is typically quiet and range-bound — it often sets the day's reference range rather than trending. The London session brings the first surge of real volume and frequently produces the day's initial manipulation move. The New York AM session, overlapping the tail of London, is the heavyweight — the highest participation and the cleanest liquidity runs of the day. The New York PM session offers a secondary window, often continuation or a late reversal. The exact clock times depend on your timezone and the instrument, which you must confirm for yourself.

> KEY IDEA
>
> The two windows that matter most for the majority of traders are the London and New York AM kill zones. If you did nothing but wait for high-quality setups in those two windows and ignore the rest of the day, you would be trading like a professional in one respect that matters enormously.

Why this is so valuable for beginners

Restricting yourself to kill zones solves several beginner problems at once. It stops you overtrading in dead periods where setups fail. It concentrates your attention on the hours that actually matter, so you can be fully present rather than watching all day and tiring. And it aligns you with when institutions are active, dramatically improving the odds that a good price setup will actually be delivered. Simply not trading outside kill zones is one of the fastest improvements a struggling trader can make.

Using kill zones in practice

In practice, kill zones frame your day. You do your preparation — structure, draw, key levels — before the session. When the kill zone opens, you watch for the sequence you know: a liquidity sweep, a shift in structure, an array entry in the correct premium/discount location. Outside the kill zone, you journal, review, and wait. This rhythm — prepare, engage during the window, review after — is the daily shape of a disciplined ICT trader, and it is built entirely on knowing when the kill zones are.

KEY POINTS

- A kill zone is a session window where institutional activity and clean delivery concentrate
- Four windows: Asian (quiet, sets range), London (first real volume), NY AM (heaviest), NY PM (secondary)
- London and New York AM are the highest-probability windows for most traders
- Trading only in kill zones stops overtrading, sharpens focus, and aligns you with institutional activity
- Daily rhythm: prepare before, engage during the window, review after

COMMON MISTAKES

- Trading all day, including dead periods, instead of concentrating on kill zones
- Not confirming the exact kill-zone times for your timezone and instrument
- Being fatigued from watching all day instead of being sharp during the key windows

PRACTICE DRILLS

- [ ] Mark the four kill zones on your chart in your own timezone.
- [ ] Track for a week where the day's high and low form relative to the kill zones.
- [ ] Trade (or paper-trade) only London and NY AM for a week and note the difference.

RECAP

Kill zones are the session windows — Asian, London, New York AM, New York PM — when institutional activity and clean delivery concentrate, with London and New York AM the most valuable. Restricting your trading to these windows stops overtrading, sharpens focus, and aligns you with when price is actually delivered, framing the disciplined daily rhythm of prepare, engage, review.

Builds on: Ch 21

CHAPTER 23 · TIME & THE MODELS

ICT Macros & The Silver Bullet

Inside the kill zones sit even sharper windows — macros — and the most famous timed setup of all, the Silver Bullet. This is time precision at its finest.

IN THIS CHAPTER YOU WILL

- → Understand what an ICT macro is and why it exists
- → Learn the Silver Bullet as a complete timed setup
- → See how precise timing sharpens an entry

Macros: precision windows

A macro is a short, fixed time window — often just a handful of minutes, frequently around the top of the hour — during which the algorithm that delivers price is especially active and reliable. If a kill zone is the session, a macro is a precise moment within it. Institutions and the delivery algorithm reliably do something in these windows: a sweep, a reversal, a push toward liquidity. Knowing the macros lets you narrow your attention from an hour-long session to the exact minutes when a move is most likely.

The Silver Bullet

The Silver Bullet is the most famous ICT timed setup, and it is essentially a complete entry model constrained to a one-hour window. Within that hour you look for a single, clean sequence: price sweeps a liquidity pool, then displaces to leave a fair value gap, and you enter on the retrace into that gap, targeting the opposing liquidity. The discipline is severe by design — one window, one high-quality trade. You are not looking for many opportunities; you are waiting for the specific sequence to appear in the specific hour, and passing if it does not.

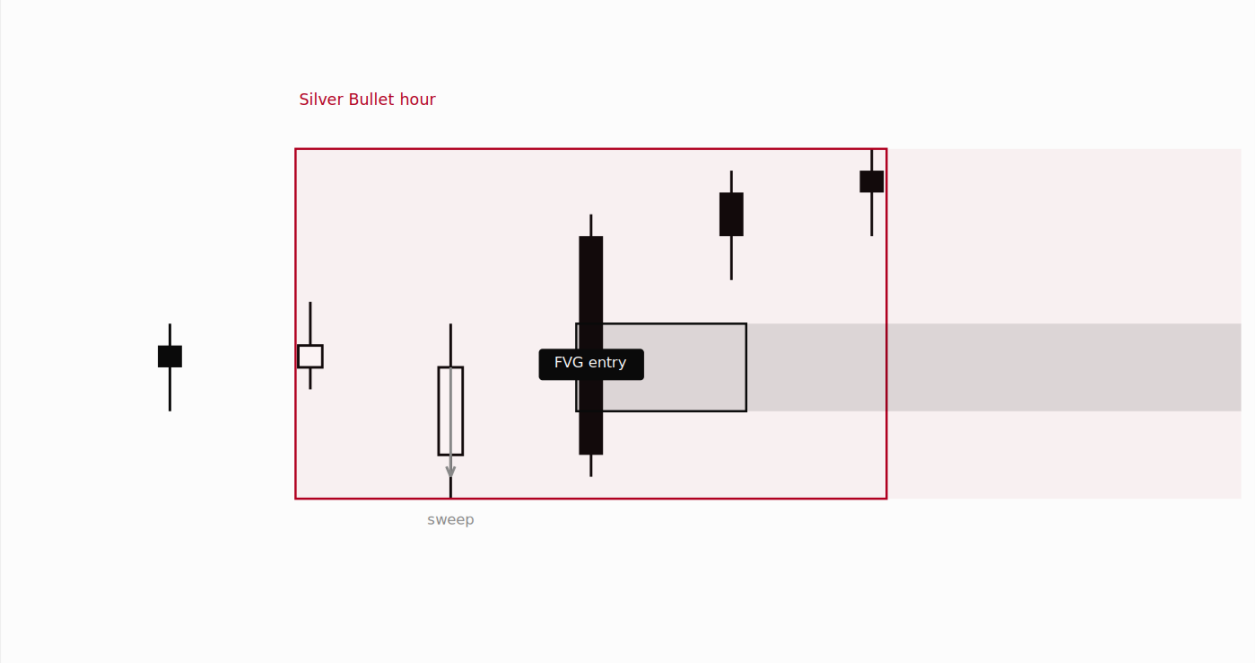

The Silver Bullet is a complete entry model inside a one-hour window: a liquidity sweep, then displacement leaving a fair value gap, entered on the retrace toward the opposing liquidity. One window, one trade.

KEY IDEA

The Silver Bullet embodies the whole philosophy of this course in miniature: a known time window, a liquidity sweep for the reason, displacement and a fair value gap for the entry, and the opposing liquidity for the target — all executed with the patience to take just one trade. Master it and you have a complete method.

Why precision timing helps

Narrowing to a macro or a Silver Bullet window helps in two concrete ways. First, it improves odds: you are acting exactly when the delivery algorithm is most likely to move price, so your setup is more likely to be honoured. Second, it reduces overtrading: by defining a tiny window in which you are permitted to act, you are structurally prevented from taking the dozens of mediocre trades that fill a beginner's day. Precision timing is as much a discipline tool as an edge — it forces the selectivity that is your true advantage as a small trader.

A note on confirming times

Exact macro and Silver Bullet times depend on your instrument and timezone, and you must verify them for your own charts rather than trusting a number from memory. What matters conceptually is the principle: within the broad kill zone, there are precise windows where delivery is most reliable, and constraining your trading to them stacks time-precision on top of your price setup. Confirm the windows once, mark them on your chart, and let them govern when you are allowed to act.

KEY POINTS

- A macro is a short, fixed window (often around the hour) when price delivery is especially reliable
- The Silver Bullet is a complete entry model constrained to a one-hour window
- Its sequence: sweep a pool, displace to leave an FVG, enter the retrace, target opposing liquidity
- Precision timing improves odds and structurally prevents overtrading
- Confirm exact macro/Silver Bullet times for your instrument and timezone yourself

COMMON MISTAKES

Taking many trades in a window meant for one high-quality setup

Trusting remembered macro times instead of confirming them for your instrument

Ignoring the sequence and entering just because the clock is in the window

PRACTICE DRILLS

- ○ Confirm and mark the Silver Bullet window for your instrument and timezone.
- ○ Watch that window daily for a week and note whether the sweep-then-FVG sequence appeared.
- ○ Paper-trade one Silver Bullet setup end to end: sweep, FVG entry, stop, target.

RECAP

Macros are precise windows within kill zones when delivery is most reliable, and the Silver Bullet is a complete entry model constrained to a single hour: sweep, displacement, FVG entry, opposing-liquidity target — one window, one trade. Precision timing both improves odds and enforces the selectivity that is a small trader's real edge.

Builds on: Ch 22, Ch 18, Ch 13

CHAPTER 24 · TIME & THE MODELS

AMD & The Power of Three

AMD — accumulation, manipulation, distribution — is the rhythm price moves through in time. Also called the Power of Three, it's the master template that every session and every model sits inside.

IN THIS CHAPTER YOU WILL

- → Define the three phases of AMD
- → Recognise AMD across a session and other timeframes
- → Use AMD to anticipate the shape of a trading day

The three phases

AMD stands for Accumulation, Manipulation, Distribution — the three phases price tends to move through, especially across a trading day. In accumulation, price consolidates in a range while institutions quietly build positions. In manipulation, price makes a sharp false move — sweeping liquidity on one side, trapping the crowd, and giving institutions their fill; this is the Judas swing. In distribution, price delivers in the true direction, toward the real draw on liquidity. Accumulate, trick, deliver — that is the rhythm.

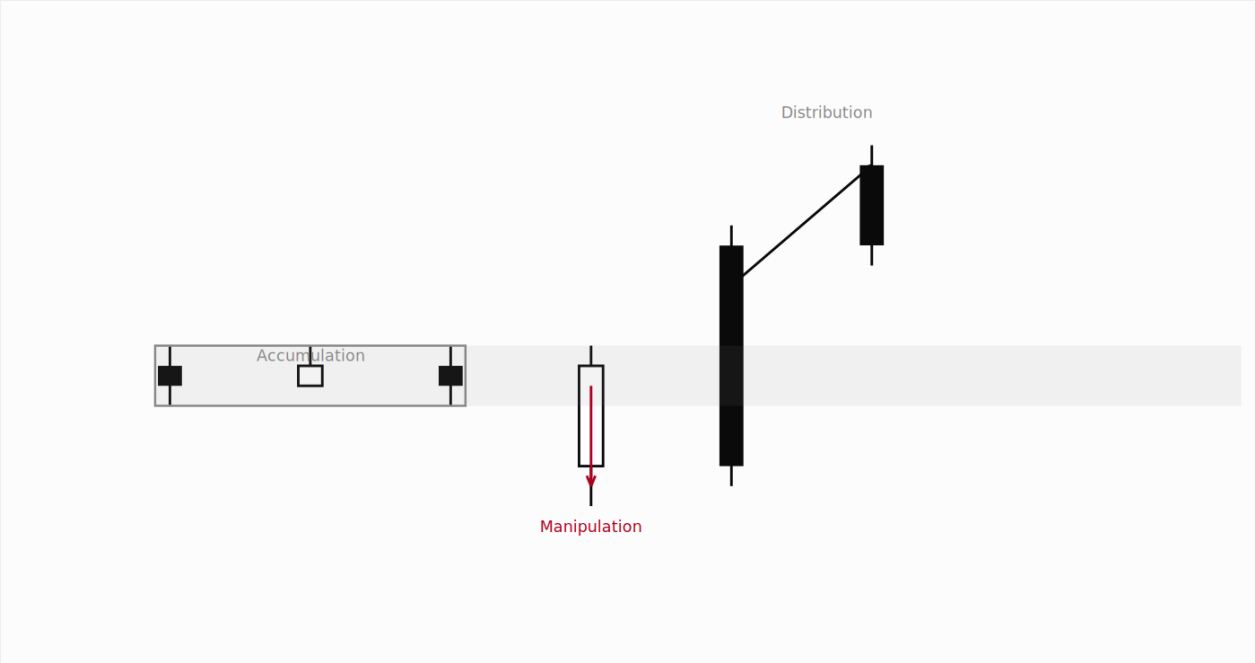

AMD — the Power of Three: price accumulates in a range, manipulates with a false sweep (the Judas move), then distributes in the true direction. Every session and model sits inside this rhythm.

Also called the Power of Three

The same idea is called the Power of Three because a candle on a higher timeframe contains all three phases: its open (accumulation), a wick to one side (manipulation), and its body toward the close (distribution). This is why a daily candle so often has a wick on the opposite side to its body — the wick is the manipulation, the body is the distribution. Seeing AMD in a single higher-timeframe candle is a powerful way to read intent: the wick tells you which side was manipulated, and the body tells you the true direction.

> KEY IDEA
>
> AMD reframes a trading day as a three-act play: consolidate, sweep the wrong way to trap and fill, then deliver the true move. Once you see it, you stop being surprised by the manipulation phase — you expect it, and you wait to trade the distribution.

Why AMD is the master template

AMD sits underneath everything else in this part. A kill zone often contains a full AMD cycle. The Silver Bullet's sweep-then-deliver is AMD's manipulation-then-distribution in miniature. The complete entry model in the next chapter is, at heart, a way to enter at the transition from manipulation to distribution. When you understand AMD, the disparate pieces — sweeps, Judas swings, structure shifts — resolve into a single story: the market accumulates, manipulates to grab liquidity and trap traders, then distributes in the true direction, and your job is to enter as distribution begins.

Trading with the rhythm

Practically, AMD tells you what to expect and when to act. Early in a session, expect accumulation — a range — and do not chase it. Expect a manipulation move that sweeps one side; do not be trapped by it, and instead treat it as the signal. Then position for distribution in the true direction, using an array entry after the manipulation sweep. Traders who know AMD are patient through accumulation, unfazed by manipulation, and decisive in distribution — while the crowd does the reverse, getting bored in accumulation, trapped in manipulation, and late to distribution.

KEY POINTS

- AMD = Accumulation (build in a range), Manipulation (false sweep, the Judas move), Distribution (true delivery)
- Also called the Power of Three: a higher-timeframe candle contains all three (open, wick, body)
- A daily candle's wick is often the manipulation; its body is the distribution
- AMD is the master template — kill zones, the Silver Bullet, and the entry model all sit inside it
- Trade the rhythm: patient in accumulation, unfazed by manipulation, decisive in distribution

COMMON MISTAKES

- Chasing price during accumulation instead of waiting for the range to resolve
- Being trapped by the manipulation sweep instead of treating it as the signal
- Entering late, after distribution is already well underway

PRACTICE DRILLS

- ○ Label accumulation, manipulation, and distribution on three past sessions.
- ○ Find a daily candle whose wick (manipulation) and body (distribution) show the Power of Three.
- ○ For one session, identify the manipulation sweep and where distribution began.

RECAP

AMD — accumulation, manipulation, distribution — is the rhythm price moves through: build in a range, sweep one side to trap and fill, then deliver the true move. Also called the Power of Three, it's the master template beneath kill zones, the Silver Bullet, and the entry model, and trading it means staying patient through accumulation, unfazed by manipulation, and decisive as distribution begins.

Builds on: Ch 13, Ch 21

CHAPTER 25 · TIME & THE MODELS

The Complete Entry Model

This is the chapter everything has been building toward: the complete entry model that combines liquidity, structure, arrays, and time into one repeatable sequence you can execute.

IN THIS CHAPTER YOU WILL

- → Assemble the full entry model step by step
- → See how every prior concept plays its role
- → Have a concrete, repeatable sequence to trade

The model in one sentence

Here is the complete model, stated once: in a high-probability time window, wait for price to sweep liquidity against your intended direction, confirm a shift in structure on displacement, and enter on the retrace into the resulting PD array — in the correct premium/discount location — targeting the draw on liquidity. Every clause of that sentence is a concept you have already learned. The model is not new material; it is the assembly of everything into a single, ordered sequence.

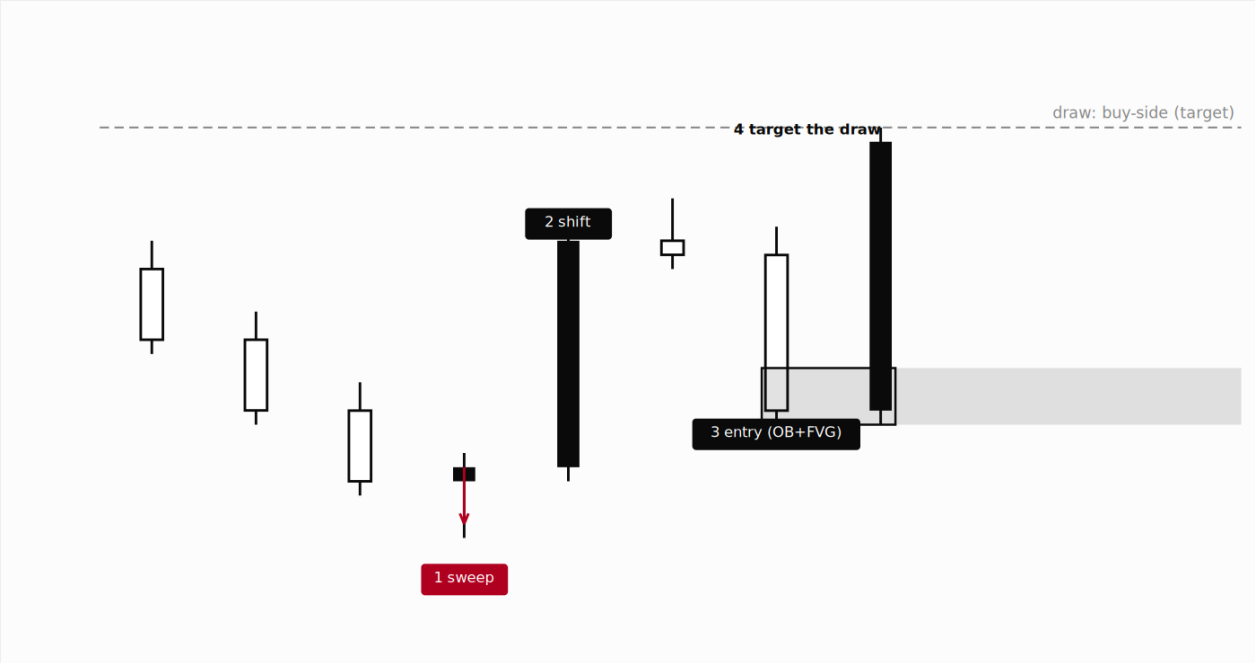

The complete model: in a time window, price sweeps liquidity against your direction (1), shifts structure on displacement (2), you enter the resulting array in discount (3), and target the draw on liquidity (4).

Step by step

Step 1 — Context. Before the window, establish structure, judge the draw on liquidity, and note the premium/discount split. You should be able to say: 'the draw is up, toward that buy-side pool; I want longs from discount.'

Step 2 — Time. Wait for a kill zone or macro. You do not act outside your window.

Step 3 — The sweep. Within the window, wait for price to sweep liquidity against your direction — for a long, a sweep of sell-side liquidity below. This is the manipulation phase; it grabs orders and traps the crowd.

Step 4 — The shift. Confirm a change of character on displacement in your intended direction. The sweep gave the reason; this gives the trigger. A weak, non-displacement shift does not count.

Step 5 — The entry. Enter on the retrace into the PD array left by the displacement — the order block or fair value gap, ideally where both overlap — in the correct premium/discount half.

Step 6 — Stop and target. Stop just beyond the sweep extreme; target the draw on liquidity you identified in Step 1. The trade is now fully defined before you are in it.

> KEY IDEA
>
> Notice that the model decides everything in advance: direction (the draw), location (premium/discount and the array), timing (the window), trigger (sweep then shift), and risk (stop beyond the sweep, target at the draw). Nothing is left to in-the-moment emotion. That is the entire point of having a model.

Why the order matters

The sequence is not arbitrary — each step gates the next. Context without a time window is a setup with no catalyst. A time window without a sweep is a clock with no reason. A sweep without a structure shift is a possible reversal with no confirmation. A shift without an array is a signal with no precise entry. Only when all are present, in order, do you act. This gating is what makes the model high-probability: it demands that price, structure, liquidity, and time all agree before you commit.

From model to mastery

This model will carry you a very long way, but a model is only as good as its execution — which is Part VI. Knowing the sequence is necessary but not sufficient; you must also size the trade correctly, manage it once live, journal it honestly, and keep your psychology in check. A trader who knows this model and executes it with discipline over hundreds of trades has a genuine edge. A trader who knows it but sizes recklessly, moves stops, and abandons it after two losses does not. The model is the what; Part VI is the how of actually turning it into results.

KEY POINTS

- The model: in a time window, sweep liquidity against your direction, confirm a structure shift on displacement, enter the resulting array in the right premium/discount half, target the draw
- It's not new material — it's the ordered assembly of everything you've learned
- Six steps: context, time, sweep, shift, entry, stop-and-target
- Each step gates the next; you act only when price, structure, liquidity, and time all agree
- The model decides direction, location, timing, trigger, and risk in advance — removing in-the-moment emotion

COMMON MISTAKES

Skipping a step — e.g. entering on a sweep without the confirming structure shift

Taking the model outside a valid time window

Knowing the model but failing on execution (sizing, management, discipline) — the subject of Part VI

PRACTICE DRILLS

- ○ Write the six-step model from memory and check it against the chapter.
- ○ Find three historical examples where all six steps aligned, and mark each step.
- ○ Paper-trade the full model five times, journaling each step as it occurs.

RECAP

The complete entry model assembles everything: in a time window, wait for a liquidity sweep against your direction, confirm a structure shift on displacement, enter the resulting array in the correct premium/discount half, and target the draw on liquidity. Each step gates the next so you act only when price, structure, liquidity, and time agree — but the model's edge is only realised through the execution discipline of Part VI.

Builds on: Ch 15, Ch 10, Ch 18, Ch 23, Ch 24

PART VI

Execution, Risk & Becoming a Trader

CHAPTER 26 · EXECUTION, RISK & BECOMING A TRADER

Entries, Stops & Targets

You know the model. Now we make its three most consequential numbers precise: exactly where you enter, where your stop goes, and where you target. Getting these right is the difference between a good idea and a good trade.

IN THIS CHAPTER YOU WILL

- → Choose between aggressive and confirmed entries
- → Place stops logically, not arbitrarily
- → Set targets at liquidity, and understand reward-to-risk

Two ways to enter

Once your model conditions are met, you have a choice of entry style. An aggressive entry takes the trade as price first touches the array — best price, best reward-to-risk, but higher chance of a deeper retrace against you first. A confirmed entry waits for a lower-timeframe signal inside the array — a small structure shift, a rejection — before committing; you sacrifice some reward-to-risk for a higher hit rate. Neither is universally correct. Beginners generally benefit from confirmed entries while learning, then add aggressive entries as their reads sharpen.

Placing the stop logically

A stop is not an arbitrary distance; it is the price at which your trade idea is proven wrong. In the model, that price is just beyond the sweep extreme: if price trades back through the level it swept, the manipulation-then-distribution premise has failed, and you want to be out. Place the stop there — a little beyond the wick, to allow for noise — and never inside the array, where normal fluctuation will clip you. A logical stop is tight because the invalidation is close, which is exactly what gives ICT setups their strong reward-to-risk.

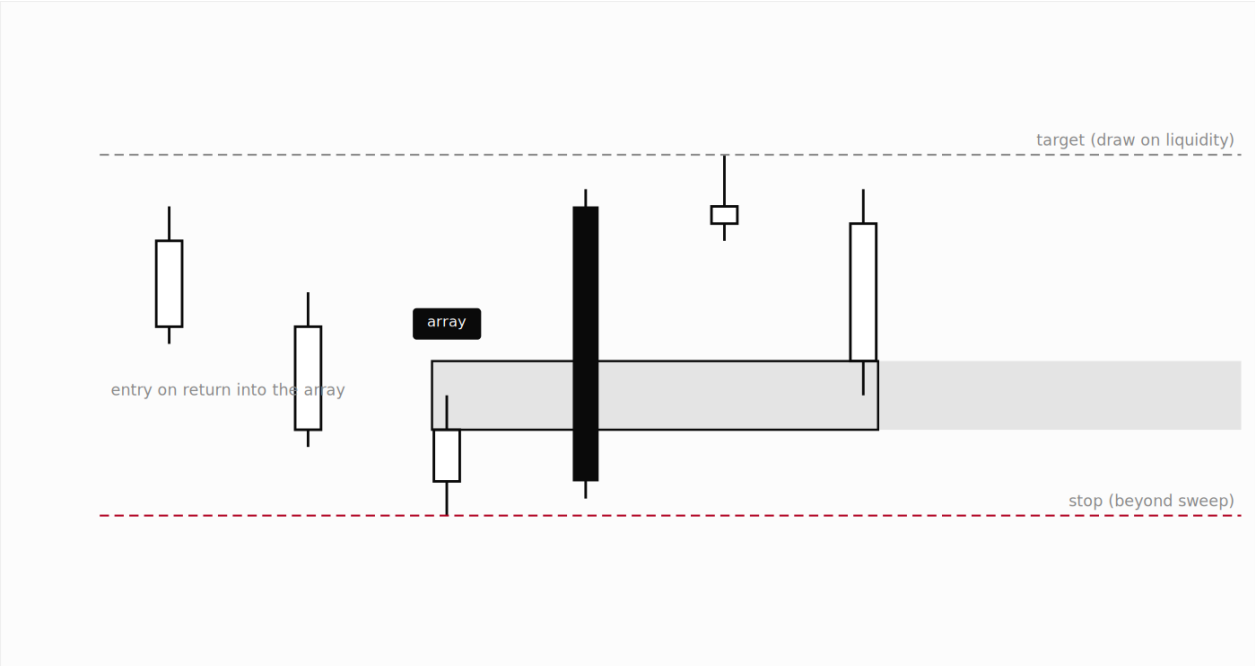

The three numbers that define a trade: entry on the return into the array, stop just beyond the sweep extreme (where the idea is wrong), and target at the draw on liquidity. A tight stop and distant target give strong reward-to-risk.

KEY IDEA

Your stop marks where you are wrong, not how much you are willing to lose. Find the price that invalidates the idea (beyond the sweep), place the stop there, and then size the position so that distance equals your chosen risk — never the other way around.

Targeting liquidity

Your target is the draw on liquidity you identified before entering — the opposing pool price is being delivered toward. This is why leading with the draw matters so much: your target was decided before the trade, at a logical destination, rather than guessed at emotionally once you were in profit. For a long entered after a sell-side sweep, the target is the buy-side liquidity above. You can take partial profit at intermediate liquidity and hold the rest for the full draw, but the primary target is always a real liquidity level, never a round number or a hopeful guess.

Reward-to-risk as a filter

With entry, stop, and target defined, you can measure the trade's reward-to-risk before taking it — the distance to target divided by the distance to stop. This becomes a filter: if a setup only offers, say, one-to-one, it is usually not worth taking, because your edge depends on winners being multiples of losers. Demanding a minimum reward-to-risk — many traders use two-to-one or better — automatically screens out marginal trades and enforces the mathematics that make you profitable. A setup that does not offer enough reward for its risk is a pass, however pretty it looks.

KEY POINTS

- Aggressive entry (first touch) maximises reward-to-risk; confirmed entry (lower-timeframe signal) raises hit rate
- The stop marks where the idea is wrong — just beyond the sweep extreme, never inside the array
- Find the invalidation first, place the stop there, then size the position to your chosen risk
- Target the draw on liquidity you identified before entering — a real level, never a round number
- Measure reward-to-risk before entering and pass setups that don't meet your minimum

COMMON MISTAKES

Placing stops by a fixed distance or dollar amount instead of at the logical invalidation

Putting the stop inside the array where noise clips it

Taking trades with poor reward-to-risk because the setup 'looks good'

PRACTICE DRILLS

- [ ] For five setups, mark the logical stop (beyond the sweep) and the liquidity target.
- [ ] Calculate reward-to-risk for each and decide which pass your minimum.
- [ ] Practise both an aggressive and a confirmed entry on the same setup and compare.

RECAP

Entries can be aggressive (first touch, best reward-to-risk) or confirmed (lower-timeframe signal, higher hit rate). The stop belongs just beyond the sweep, marking where the idea is wrong, and the target is the pre-identified draw on liquidity. Measuring reward-to-risk before entering and rejecting marginal trades enforces the maths that makes the edge real.

Builds on: Ch 25, Ch 15

CHAPTER 27 · EXECUTION, RISK & BECOMING A TRADER

Risk Management & Position Sizing

This is the most important chapter in the course. Not the most exciting — the most important. More traders are destroyed by poor risk management than by poor analysis. Master this and you survive long enough to let your edge work.

IN THIS CHAPTER YOU WILL

- → Understand risk as a fixed fraction of your account
- → Calculate position size from stop distance
- → Grasp why risk management, not analysis, determines survival

Why this chapter matters most

You can have the sharpest reads in the world and still blow your account if you size trades carelessly. Conversely, mediocre analysis with excellent risk management survives, learns, and improves. The market cannot take what you do not risk. Every professional treats risk management as the foundation, not an afterthought — because they know that survival is prerequisite to everything else. Read this chapter twice.

Risk a fixed small fraction

The core rule is simple: risk a small, fixed fraction of your account on every trade — commonly one percent, sometimes less while learning. 'Risk' means the amount you lose if the stop is hit. Fixing this fraction does two things. It caps the damage of any single loss, so no one trade can seriously hurt you. And it means a losing streak — which is inevitable — drains you slowly and survivably rather than catastrophically. A trader risking one percent can lose ten in a row and still have ninety percent of their account; a trader risking ten percent per trade is one bad run from ruin.

> KEY IDEA
>
> Fix your risk per trade as a small percentage of your account and never exceed it. This single rule prevents the account-ending blow-ups that end most trading careers. It is not negotiable, and it matters more than any setup in this book.

Position size follows the stop

Here is the mechanical link that ties everything together. You have a fixed risk amount (say one percent of your account) and a stop distance (from entry to the logical invalidation). Your position size is whatever makes that stop distance equal to that risk amount. A tight stop allows a larger position for the same risk; a wide stop requires a smaller position. This is why you place the stop logically first and size second — the stop distance determines the size, so that

whatever happens, a loss costs you exactly your fixed fraction and no more. Never widen a stop to avoid a loss; never oversize because a setup looks certain.

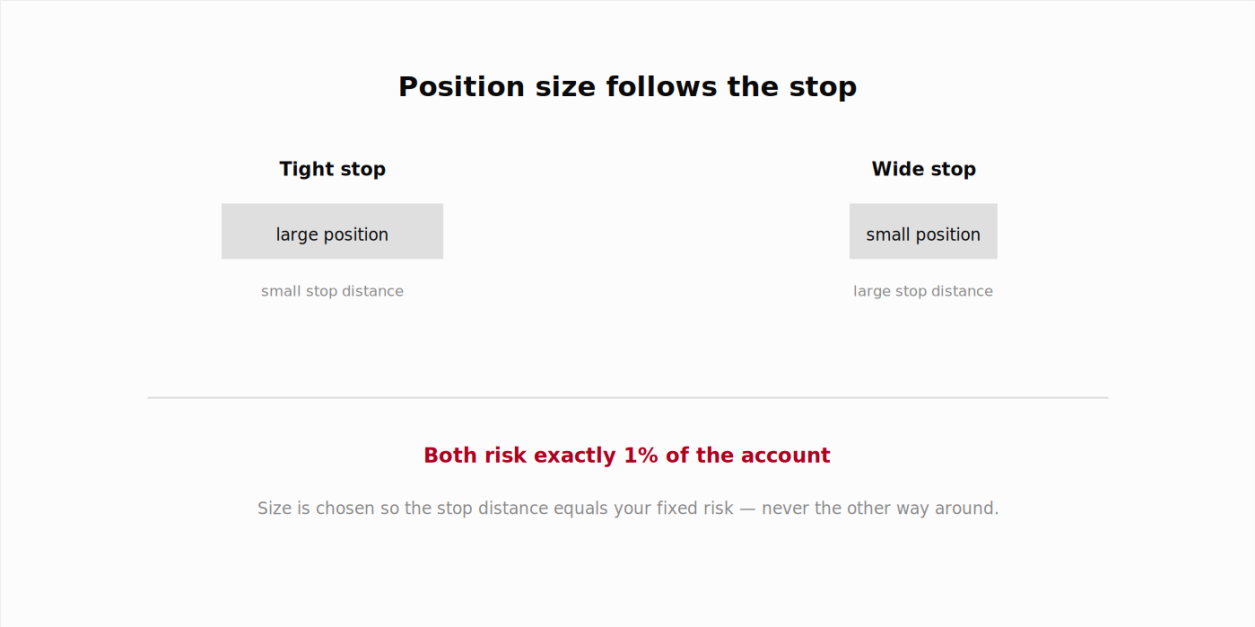

Position size follows the stop: a tight stop allows a larger position and a wide stop a smaller one, so that either way the loss equals your fixed risk (here 1%). Place the stop logically first, then size to it.

The maths of survival and growth

Combine fixed fractional risk with the reward-to-risk filter from the last chapter and the mathematics work in your favour. If you risk one percent per trade, take only setups offering two-to-one or better, and win even a modest share of them, your account grows over a large sample while never suffering a catastrophic loss. This is the entire financial engine of trading: small, controlled losses; larger, liquidity-targeted wins; and a position size that guarantees no single trade can end you. Analysis finds the trades; risk management makes them survivable and, in aggregate, profitable.

Leverage and its temptation

A final warning. Leverage lets you take positions far larger than your account, and it is the mechanism by which most beginners destroy themselves. Leverage is not inherently evil — it is just a multiplier — but combined with poor risk sizing it is fatal. If you fix your risk per trade as a fraction of your account and size positions from the stop, leverage becomes irrelevant to your risk: you are risking one percent whether or not leverage is available. The moment you start thinking in terms of how large a position you can take rather than how little you can afford to lose, you have inverted the entire discipline. Never do it.

KEY POINTS

- Risk management, not analysis, determines survival — the market can't take what you don't risk
- Risk a small, fixed fraction (often 1%) of your account on every trade, meaning the loss if the stop hits
- Position size follows the stop: size so the stop distance equals your fixed risk amount
- Place the stop logically first, size second — never widen a stop or oversize on conviction
- Fixed-fractional risk plus a reward-to-risk filter is the financial engine that grows an account safely

COMMON MISTAKES

Risking a large or variable amount per trade, exposing yourself to account-ending losses

Sizing by how large a position you can take rather than how little you can afford to lose

Widening stops to avoid losses, or oversizing because a setup feels certain

PRACTICE DRILLS

- ○ Decide your fixed risk percentage and write it into your plan.
- ○ For three setups with different stop distances, calculate the correct position size.
- ○ Model a ten-loss streak at your chosen risk and confirm you'd survive comfortably.

RECAP

Risk management is the most important skill in trading: risk a small fixed fraction of your account per trade, and size each position so the stop distance equals that fixed risk. This caps every loss, survives inevitable losing streaks, and — combined with a reward-to-risk filter — forms the engine that grows an account safely. Survival is prerequisite to every edge.

Builds on: Ch 26, Ch 5

CHAPTER 28 · EXECUTION, RISK & BECOMING A TRADER

Managing the Trade

Getting into a trade well is half the job. How you manage it once live — when to move your stop, when to take profit, when to do nothing — determines whether good entries become good results.

IN THIS CHAPTER YOU WILL

- → Manage a live trade without sabotaging it
- → Understand moving to breakeven and partial profits
- → Know when the best action is no action

The hardest part is doing nothing

Most trade management errors come from doing too much. Once you are in a well-defined trade with a logical stop and a liquidity target, the correct action is very often to leave it alone and let it resolve. Beginners, unable to tolerate the discomfort of an open position, meddle: they move stops, take profit early out of fear, or exit on a single candle that scares them — usually just before the trade would have worked. The first principle of trade management is that your plan was made with a clear head, and mid-trade emotion is a worse advisor than that plan.

Moving to breakeven

One legitimate management action is moving your stop to breakeven — the entry price — once the trade has moved meaningfully in your favour. This removes the risk of a winner turning into a loser. But it is a double-edged tool: move to breakeven too early and normal noise will stop you out of a trade that was still valid, denying you the win. A reasonable guideline is to move to breakeven only after price has reached a clear intermediate milestone, such as a minor liquidity level or a distance comparable to your risk. Do not move the stop simply because you are nervous.

> KEY IDEA
>
> Breakeven is a tool for removing risk after the trade has proven itself, not a reflex for soothing anxiety. Move to breakeven at a logical milestone, not the moment you feel uncomfortable — premature breakeven stops are one of the most common ways traders sabotage valid winners.

Partial profits

Scaling out — taking partial profit at an intermediate target and letting the rest run to the full draw — is a reasonable way to balance the psychological desire to bank something against the mathematical benefit of large winners. Take a portion at a nearer liquidity level, move the stop

to breakeven on the remainder, and let the rest target the full draw. This is optional; some traders prefer to take the whole position at the primary target for simplicity. What matters is that the plan is decided before the trade, not improvised in the heat of an open position.

Let the target or stop do its job

The cleanest management, especially while learning, is often the simplest: set your stop and target when you enter, then let one of them be hit without interference. This forces you to trust your analysis and removes the endless opportunities to sabotage yourself through meddling. As you gain experience you can add nuance — breakeven moves, partials, reading price action to exit early when the thesis clearly breaks — but the discipline of leaving a good trade alone is the foundation. Master doing nothing before you attempt doing something clever.

KEY POINTS

- Most management errors come from doing too much — often the right action is to leave the trade alone
- Move to breakeven only after a logical milestone, not to soothe anxiety — premature breakeven kills winners
- Scaling out (partial at an intermediate target, rest to the full draw) is optional but must be planned in advance
- The cleanest approach while learning: set stop and target, then let one be hit without interference
- Mid-trade emotion is a worse advisor than the plan you made with a clear head

COMMON MISTAKES

Meddling — moving stops, taking profit early on fear, exiting on a single scary candle

Moving to breakeven too early and getting stopped out of a valid trade by noise

Improvising management mid-trade instead of deciding it before entry

PRACTICE DRILLS

- ○ Take five paper trades where you set stop and target and do not touch them; observe the results.
- ○ Define in advance your rule for when to move to breakeven.
- ○ Journal one trade you were tempted to meddle with, and what happened when you didn't.

RECAP

Good trade management is mostly restraint: set a logical stop and a liquidity target, then let the trade resolve without meddling. Move to breakeven only at a logical milestone, treat scaling out as an optional plan made in advance, and — especially while learning — master leaving a good trade alone before attempting anything clever.

Builds on: Ch 26, Ch 27

CHAPTER 29 · EXECUTION, RISK & BECOMING A TRADER

Journaling & Backtesting

This is how you actually get good. Backtesting compresses months of experience into weeks, and journaling turns your trades into lessons. Skip these and you're not learning — you're just gambling repeatedly.

IN THIS CHAPTER YOU WILL

- → Backtest the method systematically to build pattern recognition
- → Keep a journal that produces real lessons
- → Use your records to prove (or disprove) your edge

Why backtesting is non-negotiable

Backtesting means going through historical charts, bar by bar, and marking where your setups occurred and how they resolved. It feels slow, but it is the single fastest way to build the pattern recognition that live trading requires — because you can review months of price in days, seeing hundreds of examples of sweeps, arrays, and model setups. A trader who has backtested five hundred setups recognises them instantly; a trader who has only read about them hesitates and misreads. There is no shortcut around this repetition, and backtesting is how you get it without risking money.

How to backtest well

Backtest deliberately, not randomly. Pick one instrument and one setup — say, the complete entry model in the New York AM kill zone. Go through historical data and mark every instance: the context, the sweep, the shift, the array, the entry, the stop, the target, and the outcome. Record each in your journal. Resist the temptation to cherry-pick winners or to move your rules to fit what you see — that is lying to yourself, and the market will collect on the lie later. Honest backtesting over a large sample tells you whether your setup actually has an edge, before a cent is at stake.

> KEY IDEA
>
> Backtesting answers the only question that matters before you risk money: does this setup, traded by these rules, actually make money over a large sample? Answer it honestly on historical data first. If it doesn't work in backtesting, it will not work live — and you have saved yourself the tuition.

Keeping a real journal

A trading journal records every trade — backtested and live — with enough detail to learn from. For each trade, note the setup and context, a screenshot, your entry, stop, and target, the

outcome in terms of your risk (won two-to-one, lost one), and — crucially — what you did well and what you would change. The journal is where undisciplined trades get caught: a pattern of moving stops, or entering outside kill zones, or oversizing, becomes visible only when written down. Most traders resist journaling because it makes their mistakes undeniable. That is exactly why it works.

Proving your edge

Over a large sample, your journal answers the essential question: do you have a documented edge? Tally your trades in terms of R — your fixed risk unit. If, over a hundred or more trades, you are net positive in R, you have evidence of an edge and can trade with justified confidence. If you are net negative, the journal tells you where the leak is — poor entries, bad time-of-day, breaking risk rules — so you can fix it. This is the difference between trading and gambling: the gambler has no idea whether they have an edge, while you have written proof either way. Do not risk meaningful money until that proof is positive.

KEY POINTS

- Backtesting — marking setups on historical charts — is the fastest way to build pattern recognition
- Backtest deliberately: one instrument, one setup, every instance recorded honestly, no cherry-picking
- A journal records each trade with context, screenshot, R-outcome, and what to do differently
- The journal makes undisciplined patterns (moved stops, wrong hours, oversizing) undeniable
- Over 100+ trades, your journal proves whether you have an edge — don't risk real money until it's positive

COMMON MISTAKES

Skipping backtesting and trying to learn on live money instead

Cherry-picking backtest results or bending rules to fit — lying to yourself

Not journaling, so recurring mistakes stay invisible and unfixed

PRACTICE DRILLS

- ○ Backtest fifty instances of one setup on one instrument, recording each honestly.
- ○ Set up a journal template with context, screenshot, R-outcome, and a lessons field.
- ○ After fifty trades, tally your R and identify your single biggest leak.

RECAP

Backtesting compresses experience by marking hundreds of historical setups, and journaling turns every trade into a lesson by recording context, outcome in R, and what to change. Together they answer the only question that matters before risking money — do you have a documented edge? — and separate disciplined trading from gambling.

Builds on: Ch 5, Ch 25

CHAPTER 30 · EXECUTION, RISK & BECOMING A TRADER

Building Your Trading Plan

A trading plan turns a method into a business. Without one, every decision is improvised and emotional. This chapter helps you write the document that governs exactly what you trade, when, and how much.

IN THIS CHAPTER YOU WILL

- → Understand what a trading plan must contain
- → Draft your own rules for setup, timing, and risk
- → Commit to following the plan mechanically

Why a written plan

A trading plan is a written document that specifies, in advance, exactly what you will and will not do. Its purpose is to remove in-the-moment discretion, because discretion under the pressure of an open market is where discipline dies. With a plan, a trade either meets your written criteria or it does not — the decision is made calmly, once, and simply applied. Without a plan, every chart is a fresh argument with yourself, and emotion wins those arguments more often than not. Professionals trade a plan; amateurs trade their feelings.

What the plan must contain

A complete plan answers concrete questions. What do you trade — which instrument, which setup (for example, the complete entry model only)? When do you trade — which kill zones or windows, and never outside them? How much do you risk — your fixed fractional risk per trade, and your maximum for a day or week? What are your entry, stop, and target rules — precisely defined so there is no ambiguity? And what are your hard 'no-trade' conditions — no clear draw, outside a window, reward-to-risk below your minimum, or after hitting a daily loss limit? Write each of these down explicitly.

> KEY IDEA
>
> If a trade is not specified by your written plan, it is not your trade to take. The plan is not a suggestion you consult when convenient; it is the sole authority on whether you act. A setup that violates the plan is a pass no matter how good it looks — because the plan, not the moment, holds your edge.

Rules that protect you from yourself

Beyond the setup, a good plan includes rules that guard against your own worst tendencies. A daily loss limit — stop trading after losing a set amount — prevents the tilt-driven spiral where one loss becomes ten. A maximum number of trades prevents overtrading. A rule to stop after

a big win can prevent giving it back through overconfidence. These are not arbitrary; each targets a specific, well-documented way traders self-destruct. Your journal (previous chapter) will reveal which tendencies are yours, and your plan is where you install the guardrails against them.

Following it mechanically

A plan only works if you follow it mechanically, especially when you do not feel like it. The trades you most want to take against your plan — the revenge trade after a loss, the oversized 'sure thing', the setup outside your window because you are bored — are precisely the ones the plan exists to prevent. Treat plan adherence as a separate skill from analysis, and track it in your journal: did I follow my plan on every trade, yes or no? A trader who follows a sound plan mechanically will outperform a more talented analyst who does not. Discipline is the edge that protects every other edge.

KEY POINTS

- A trading plan is a written document that removes in-the-moment discretion
- It specifies what you trade, when, how much you risk, your entry/stop/target rules, and no-trade conditions
- If a trade isn't specified by your plan, it isn't yours to take — the plan is the sole authority
- Include guardrails against yourself: daily loss limit, max trades, stop-after-big-win
- Follow the plan mechanically — the trades you most want to break it for are the ones it exists to prevent

COMMON MISTAKES

Trading without a written plan, so every decision is improvised and emotional

Having a plan but overriding it in the moment for 'special' setups

Omitting guardrails (loss limits, trade caps) that protect against tilt and overtrading

PRACTICE DRILLS

- ○ Draft your full trading plan covering what, when, how much, entry/stop/target, and no-trade rules.
- ○ Add three guardrail rules targeting your own worst tendencies from your journal.
- ○ Add a 'did I follow my plan?' yes/no field to your journal and track it for twenty trades.

RECAP

A trading plan is the written document that removes emotional discretion by specifying exactly what you trade, when, how much you risk, and your precise entry, stop, target, and no-trade rules — plus guardrails like daily loss limits. If a trade isn't in the plan, it isn't yours to take, and following the plan mechanically is the discipline that protects every other edge.

Builds on: Ch 27, Ch 29

CHAPTER 31 · EXECUTION, RISK & BECOMING A TRADER

Psychology & Discipline

The final skill is the one inside your own head. You can know every concept and hold a perfect plan and still fail if your psychology sabotages your execution. This chapter confronts that directly.

IN THIS CHAPTER YOU WILL

- → Recognise the main psychological traps in trading
- → Understand why discipline beats intelligence
- → Build habits that keep emotion from wrecking your edge

Why psychology is the final boss

By this point you have the analysis, the model, the risk rules, and the plan. What stands between you and results is yourself. Trading pits your money against your emotions in real time, and emotions are powerful, fast, and usually wrong in this context. Fear makes you exit winners early and skip valid setups; greed makes you oversize and overtrade; the pain of a loss makes you seek revenge, breaking every rule to 'get it back'. None of this is a character flaw — it is universal human wiring, poorly suited to markets. Acknowledging that is the first step to managing it.

The main traps

A few traps account for most damage. Revenge trading: taking an impulsive, oversized trade to recover a loss, usually losing more. Fear of missing out: chasing a move you missed, entering late at a bad price. Overconfidence after wins: sizing up and getting sloppy just when a run makes you feel invincible. Loss aversion: holding losers past their stop hoping they come back, and cutting winners early to lock in a small gain — the exact opposite of the edge. Each of these is predictable, which means each can be guarded against with rules made in advance, when you are calm.

> KEY IDEA
>
> Your plan and risk rules exist precisely because your in-the-moment psychology cannot be trusted. The disciplined trader is not someone who feels no fear or greed — they feel it and follow the rules anyway. Discipline is acting on your calm decisions rather than your live emotions.

Why discipline beats intelligence

It is a hard truth that in trading, discipline beats intelligence. A brilliant analyst who cannot follow their own rules will lose to a mediocre analyst who follows theirs mechanically. This is

because the market does not reward being clever; it rewards consistent, correct execution over many trades, and consistency is a matter of discipline, not intellect. If you have read this far you are clearly capable of understanding the method. Whether you succeed will come down to whether you can execute it with discipline when it is emotionally hard — which is most of the time.

Building the habits

Psychology is not fixed by willpower alone; it is managed by systems and habits. Follow a written plan so decisions are pre-made. Keep your risk small enough that no single trade triggers a strong emotional response — much fear and greed simply evaporate when the stakes per trade are modest. Journal your discipline, not just your trades, so lapses are visible. Take breaks after losses to avoid tilt. Detach your self-worth from any single trade's outcome; you are running a large sample, and any one result is noise. Over time, these habits build the calm, mechanical temperament that trading rewards — and that temperament, more than any setup, is what separates those who make it from those who do not.

KEY POINTS

- After analysis, model, risk, and plan, the final obstacle is your own psychology
- Main traps: revenge trading, fear of missing out, overconfidence after wins, loss aversion
- Your plan and risk rules exist because in-the-moment emotion can't be trusted
- Discipline beats intelligence — the market rewards consistent execution, not cleverness
- Manage psychology with systems: pre-made plans, small risk, journaling discipline, breaks, detachment from single outcomes

COMMON MISTAKES

Believing you'll master emotion through willpower rather than systems and small risk

Revenge trading, chasing moves, or sizing up after wins

Holding losers past the stop and cutting winners early — inverting your edge

PRACTICE DRILLS

- ○ Identify which psychological trap you're personally most prone to.
- ○ Reduce your risk per trade until no single trade provokes strong emotion, and note the effect.
- ○ Add a discipline score to your journal and review it weekly for lapses.

RECAP

Psychology is the final obstacle: fear, greed, revenge, and loss aversion sabotage even perfect analysis. Your plan and small, fixed risk exist because in-the-moment emotion can't be trusted, and discipline — acting on calm decisions rather than live feelings — beats intelligence. Manage it with systems and habits, not willpower, because temperament is what ultimately separates those who succeed.

Builds on: Ch 27, Ch 30

CHAPTER 32 · EXECUTION, RISK & BECOMING A TRADER

Capstone: A Complete Trade, Start to Finish

Everything you've learned, combined into one worked trade from blank chart to closed position. Read this as the demonstration that the pieces really do form a single, executable whole.

IN THIS CHAPTER YOU WILL

- → See the entire method applied to one trade end to end
- → Watch each concept play its specific role in sequence
- → Leave with a clear template for your own execution

Setting the scene

Let us trade one complete example, naming every concept as it appears. It is the start of the New York AM session. Before anything else, we prepare. On the higher timeframe, structure is bullish — a clean series of higher highs and higher lows. Price has pulled back into the discount half of the daily dealing range. Above sits an untapped pool of buy-side liquidity at an old high. We judge the draw on liquidity to be upward, toward that pool. Our thesis, stated before we look for an entry: 'the draw is up; we want longs from discount toward the buy-side liquidity above.'

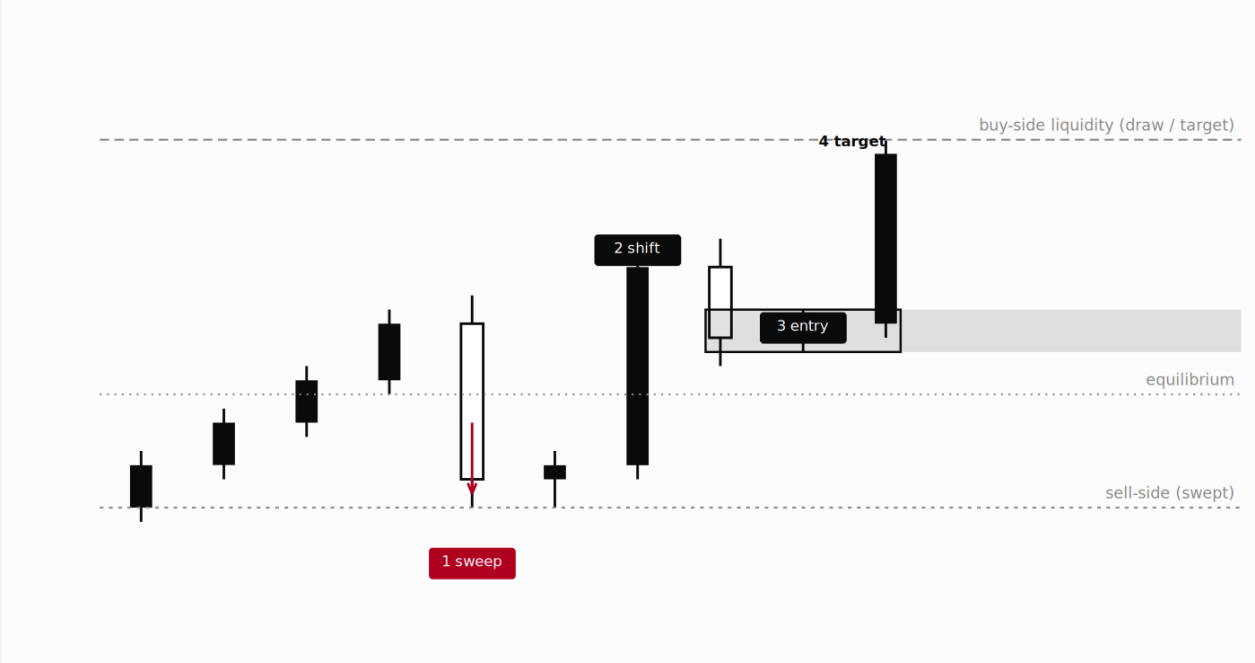

One complete trade: bullish context with price in discount and the draw above (prepare); sweep of sell-side liquidity in the kill zone (1, manipulation); structure shift up on displacement (2); entry into the OB+FVG array (3); delivery to the buy-side target (4). Stop sits beyond the sweep; risk fixed at 1%.

Waiting for the window and the sweep

We do not act until the kill zone opens — time is a filter, and a setup outside our window is not our trade. As the New York AM session begins, price does something that would frighten an untrained trader: it drops, sweeping the sell-side liquidity below a recent low. The crowd sees a breakdown and sells. We see the manipulation phase of AMD — a Judas move grabbing sell-side liquidity and trapping breakout sellers, exactly the fuel institutions need to drive price up toward our draw. The sweep is not a threat to our thesis; it is the signal we were waiting for.

Confirmation and entry

A sweep alone is not enough. We wait for confirmation, and it comes: price shifts structure to the upside on displacement — a strong, large-bodied move that breaks the recent minor high and signals delivery has flipped up. That displacement leaves a fair value gap behind it, and just below sits the order block — the last down-candle before the move. The two overlap, giving us a confluence entry zone in discount. We enter long on the retrace into that zone. Our stop goes just beyond the low of the sweep: if price trades back below it, our manipulation-then-distribution thesis is wrong and we are out for a small, fixed loss.

Sizing, target, and management

Before confirming the entry, we size the position so that the distance to our stop equals one percent of the account — no more, whatever the setup's allure. Our target is the buy-side liquidity above, the draw we identified at the very start. From entry to target is roughly three times our risk: a 3R trade that clears our reward-to-risk minimum. Now we manage by doing very little. Price delivers upward in the distribution phase. We hold, resisting the urge to meddle, and let price reach the buy-side liquidity. The target is hit; we close the trade. Three R gained, one R ever risked.

> KEY IDEA
>
> Look back at what that single trade used: structure and premium/discount (Part II), liquidity, the sweep, and the draw (Part III), the order block and fair value gap in confluence (Part IV), the kill zone, AMD, and the complete model (Part V), and the entry, stop, position sizing, target, and management discipline (Part VI). Every part of this course appeared, in order, in one trade. That is the whole point: the concepts are not separate — they are one method.

Your turn

This is the template for every trade you will take: prepare the context and draw before the session; wait for your time window; let price sweep liquidity against your direction; demand a structure shift on displacement; enter the resulting array in the correct half of the range; size from a logical stop; target the draw; and manage by mostly leaving it alone. It will not always work — no method wins every trade, and it does not need to. It needs only to be executed with discipline over a large sample, with your losses small and your winners targeting real liquidity.

You now have the complete method. What remains is the practice — backtest it, journal it, forward-test it, and let the edge reveal itself over time. Begin.

KEY POINTS

- A complete trade: prepare context and draw, wait for the window, let liquidity sweep, confirm a shift, enter the array, size from the stop, target the draw, manage minimally
- The sweep that frightens the crowd is the manipulation signal you were waiting for
- Entry confluence: displacement leaves an FVG overlapping the order block, in discount
- Stop beyond the sweep, size to fixed risk, target the pre-identified draw — a 3R trade
- Every part of the course appears, in order, in one trade — the concepts are one method

COMMON MISTAKES

Being frightened out by the manipulation sweep instead of recognising it as the signal

Entering before the structure-shift confirmation, or outside the time window

Meddling during distribution instead of letting the trade reach its target

PRACTICE DRILLS

- ○ Reverse-engineer one historical move into all these steps: context, draw, sweep, shift, array, entry, stop, target.
- ○ Do the same on ten examples until the full sequence is second nature.
- ○ Only then, forward-test the complete method on demo, journaling every step — and begin your real practice.

RECAP

One complete trade weaves the whole course together: prepare the context and draw, wait for the kill zone, let price sweep liquidity as the manipulation signal, confirm a structure shift on displacement, enter the overlapping order block and fair value gap in discount, size from a logical stop, and target the draw — then manage by leaving it alone. Every part of the method appears, in order, in a single trade. The knowledge is complete; the practice begins now.

Builds on: Ch 25, Ch 26, Ch 27, Ch 28

Begin.

You have the complete method. Knowledge becomes skill only through practice — backtest it, journal it, forward-test it, and let your edge reveal itself over a large sample. The market will still be here. Take your time, and trade with discipline.

theinnercircletraders.com