

TRADING KNOWLEDGE LIBRARY



Trading by the Clock

Market Analysis



Trading Papers

Trading by the Clock

Sessions, windows, and the shape of the market's day

Published by **Trading Papers**

Series **The trading knowledge library**

Edition **First edition · 2026-08-26**

Format **A4 · digital**

Educational content only. Nothing in this document is financial advice, no part of it is a recommendation to buy or sell anything, and it knows nothing about your circumstances. Trading carries risk, including the loss of the amount invested.

© 2026 Trading Papers. All rights reserved.



Trading Papers

Contents

01	The market does not move evenly	4
02	Which horizon the clock is for	5
03	The three sessions	6
04	Getting the clock right	7
05	The shape of a session	9
06	Session levels, and why price returns to them	11
07	Building a window you can actually trade	13
08	A week, written down	14
09	Scheduled events are part of the clock	15
10	The clock is not the same for every instrument	16
11	Testing a timing edge without fooling yourself	18
12	Windows decay	20
13	What this gets you, and what it does not	21

01

The market does not move evenly

Put a day of any liquid market on a chart and strip the times off the axis, and the day looks like one continuous thing: a line that goes up, or down, or sideways. Put the times back and it stops being continuous. Long stretches barely move. Then, at hours you can predict on a calendar rather than from the chart, the range of a whole quiet stretch is covered in twenty minutes.

That unevenness is not noise, and it is not a pattern hidden in the price. It is a consequence of who is awake. Markets are made of institutions, and institutions keep office hours. A currency pair is quiet at three in the morning in London for the same reason a motorway is quiet at three in the morning: the people who use it are asleep. When Tokyo opens, then Frankfurt and London, then New York, the population of the market changes, and with it the size of the orders, the tolerance for risk, and the reasons anyone is trading at all.

This is the single most useful thing a chart cannot show you, and it is why time deserves to be treated as an input rather than an axis. Two identical setups — the same structure, the same distance to the same level, the same trigger — are not the same trade if one appears at 02:40 and the other at 09:35. The first has to move through a market with almost nobody in it. The second has the whole of London behind it.



Schematic. The reason an hour behaves the way it does is who is in the market during it, and the mix changes three times a day. Everything else in this paper is downstream of this one picture.

The claim in this paper is narrow on purpose. Time is not a signal. Nothing here says an hour makes a trade good. What it says is that the hour changes the distribution of what happens next — how far price tends to travel, how cleanly, and how often a level holds — and that a plan which ignores this is being tested against an average that describes no actual hour of the day.

02

Which horizon the clock is for

Before any of this is useful, one thing has to be fixed: how long the position is meant to live. The clock matters enormously to an intraday trader, moderately to a multi-day one, and almost not at all to an investor — and a paper about sessions is worth nothing to someone holding for months.

Horizon	Position closed	What the clock is for
Intraday	within the session	choosing the window and the levels
Multi-day	days later	choosing the session to enter and exit in
Investment	months or longer	almost nothing — the clock stops mattering

Three horizons and the boundary between them. The boundary matters because it is rarely crossed by decision — a day trade held overnight has become a multi-day trade by default, governed by a plan that was never written.


The boundary is worth stating because of how it is usually crossed. A trade opened as intraday and left open because closing it would realise a loss has not been reclassified — it has escaped its plan. The position is the same; the rules that governed it no longer exist, and the statistics that justified it no longer apply.

So the first use of the clock is not choosing an hour to trade. It is deciding, before entry, which hour the position must be closed by — and treating that as part of the setup rather than as a preference.


03

The three sessions

The twenty-four hour market is conventionally divided into three sessions, named for the financial centres whose working day defines them. The boundaries are not exact and no exchange rings a bell, but the character of each is consistent enough to plan around.



The three sessions on one 24-hour axis, in UTC. Boundaries are conventional rather than exact — no exchange rings a bell — and the shaded band is the London–New York overlap, the deepest and fastest hours of the day.



Schematic, and the single most useful thing to compute for yourself — section 09 says how. This is the shape a liquid FX pair tends to take: a quiet Asian stretch, a violent London open, a deep overlap, and an afternoon that empties out.

Asia — roughly 23:00 to 08:00 UTC

Tokyo, Hong Kong, Singapore and Sydney. Volume is the lowest of the three and ranges are usually the narrowest. Moves tend to be orderly and to respect levels, because there is not

enough size in the book to push through anything substantial. The exception is the currencies and commodities the region actually trades — the yen, the Australian and New Zealand dollars, gold — which can move properly on regional data.

For everything else, the practical value of the Asian session is not the trading in it. It is that it builds a compact, well-defined range that the next session will use.

London — roughly 07:00 to 16:00 UTC

The largest FX session in the world by volume, and the one that most often sets the day's direction. The first hour is typically the most volatile of the day: overnight orders are worked, the Asian range is frequently taken out on one side or both, and the day's high or low is set surprisingly often within the first two hours of it.

New York — roughly 12:00 to 21:00 UTC

The second heavyweight, and the one that carries the majority of scheduled economic releases for the world's reserve currency. Its first ninety minutes overlap with London's afternoon, and that overlap — roughly 12:00 to 16:00 UTC — is the deepest, fastest window of the twenty-four hours. After London goes home, volume thins noticeably; the late New York afternoon is a different market from its morning, and trades that assumed morning liquidity often stop working in it.

Clock	Moves?	What it is good for
UTC	never	research, storage, anything you want to compare later
Exchange local	with local DST	reading the venue's own published schedule
Broker server	usually, unpredictably	nothing — but it is what your chart displays
Your local	with your DST	deciding whether you will actually be at the screen

Four clocks, all in play at once. Every timing rule you write is expressed in one of them, and a rule that does not name which one is unusable six months later.

04

Getting the clock right

Everything above is written in UTC, and that is deliberate. Session work fails more often on timekeeping than on analysis, and the failures are dull and avoidable.

Your broker's server clock is not UTC. It is usually somewhere between UTC+1 and UTC+3, it may shift with daylight saving, and it is chosen for the broker's convenience rather than yours. Whatever you conclude about hours will be expressed in whatever clock your chart is using, so the first task is to know which clock that is and to write it down next to your notes. A rule that says "the London window is 10:00 to 12:00" is worthless six months later if nobody recorded whose ten o'clock it was.

Daylight saving is the second trap, and it is worse than it sounds because the northern hemisphere does not change on one date. Europe and the United States shift on different weekends, which leaves two stretches each year — a few weeks in spring, a week in autumn — when the gap between London and New York is one hour different from usual. During those weeks the overlap window moves. Sydney and Tokyo add their own complications: Japan does not observe daylight saving at all, so the Asian session drifts relative to the others twice a year.

The practical answer is to anchor your analysis to one reference clock — UTC is the obvious choice because it never changes — and to convert for display, rather than doing the reverse. If your platform allows a UTC chart, use it for research even if you trade on a local one.

Period	What changes	Effect on the overlap
Mid-March to late March	US shifts, Europe has not	overlap starts an hour earlier
Late October	Europe shifts, US has not	overlap starts an hour later
All year	Japan never shifts	Asian session drifts twice annually

The weeks each year when the usual gap between the centres is wrong. During these the overlap window moves, and any rule written in broker time lands an hour off what it was tested on.

A short exercise worth doing once: mark the same session boundary on your chart on a date in January and a date in July, and confirm it lands where you expect. If it does not, you have found a clock problem that would otherwise have quietly corrupted every timing conclusion you draw for a year.

05

The shape of a session

Sessions have an internal structure that repeats often enough to be worth describing, and it is roughly the same shape in every one of them, scaled to the session's size.



Schematic, not measured: the shape a session tends to take. A fast open as overnight orders are worked, a slower middle where levels behave best, a smaller burst into the close, then the drift into the next session's quiet.

The opening is the most active part. Orders that accumulated while the centre was closed are worked at once, and the resulting move is fast and often one-directional. It is also the least forgiving: spreads are wider, slippage is real, and a level that would hold at midday is frequently cut straight through.

The middle is slower. Once the opening flow is absorbed, price tends to settle into a range, retrace part of the opening move, or drift with whatever the day's larger theme is. This is where levels behave best — where a support that has been touched twice is more likely to hold a third time — because there is enough participation for the level to be defended and not so much that it is simply run over.

The close brings a second, smaller burst as positions are squared, followed by a decline into the next session's quiet. In the New York afternoon this decline is pronounced, and it is a common source of frustration: a setup that looked identical to a morning one fails to travel, not because the analysis was wrong but because there is no longer anyone on the other side to move it.

	Opening hour	Middle hours
What it offers	Range, direction, the day's extreme surprisingly often.	Levels behave: a twice-touched support often holds again.
What it costs	Wide spreads, real slippage, levels cut straight through.	Less travel: the same target needs patience or a closer entry.

The two halves of a session behave differently enough to be traded differently. What is a virtue in one column is a hazard in the other, which is why one set of rules applied to the whole session performs worse than either half deserves.

Expansion and contraction


Underneath this is a simpler rhythm that any market shows: periods of contraction, where range narrows and price coils, are followed by periods of expansion, where it travels. Sessions organise that rhythm in time. The Asian session is usually contraction. The London open is usually expansion. Neither is a rule, and both fail regularly, but the sequence is common enough that seeing a very tight Asian range should raise, not lower, your expectation of movement in the hours that follow.

06

Session levels, and why price returns to them

A session does not only produce movement. It produces reference points, and those reference points are the most useful thing time gives you. The important ones are simple:

- The high and low of the Asian range.
- The high and low of the previous day, and of the previous week.
- The opening price of the London and New York sessions.
- The high and low of the first hour after a session opens.



Schematic. Protective stops rest just beyond an obvious extreme, so a cluster of them above the range is buying at a known price. Price reaches them, they fill, and the reason to continue is gone. Common — not reliable: plenty of opens break the high and keep going.

Level	Why it exists	Shelf life
Asian high / low	protective stops rest just beyond it	one session
Previous day high / low	the most widely watched extreme	one to two days
Session open	reference for everyone who started there	that session
First-hour high / low	the day's opening argument, settled	that session

The reference levels a session produces, and why each one attracts price. None of this is mystical: every row describes orders resting at a known price, visible to anyone reading the same chart.

These levels matter for a mechanical reason rather than a mystical one. Traders place protective stops just beyond obvious extremes, and resting stop orders are, by construction, orders that will execute if price reaches them. A cluster of them above the Asian high is a pocket of guaranteed buying, sitting at a known price, visible to anyone who can read a chart. A large participant who needs to fill a sell order of size has an obvious place to find the other side of it.

This is why the London open so often pushes through the Asian high or low and then reverses. Described without narrative: price moved to where the resting orders were, those orders filled, and having filled them the market no longer had a reason to continue. Whether you call this a stop hunt or simply a market seeking liquidity, the observable sequence is the same, and it happens at predictable times because that is when the size that needs filling shows up.




The sequence described without narrative. Whether it is called a stop hunt or a market seeking liquidity, the observable steps are the same — and they happen at predictable hours because that is when the size that needs filling arrives.

Two cautions. First, this sequence is common, not reliable: plenty of London opens break the Asian high and keep going, and a plan that assumes reversal every time will lose money on trend days. Second, the levels only carry weight while they are recent. Yesterday's Asian high matters this morning; last Tuesday's does not.

07

Building a window you can actually trade

Knowing that some hours are better than others is not yet a plan. Turning it into one means narrowing rather than widening, and the discipline is mostly in what you exclude.



A window written into the plan, laid over the day: entries only inside it, positions opened in it managed afterwards, nothing new after that. The hours are an example — the point is that the boundary is decided before the day starts.

Choose one window and learn it

Two hours is enough. A trader who knows one window deeply — how it usually opens, what a normal range looks like in it, which levels it respects, how it behaves the day after a trend day — has a far better sample to reason from than one who watches fourteen hours and remembers the memorable parts of each. The two hours after the London open and the first two hours of the New York session are the conventional choices because they carry the most volume, but the argument for them is participation, not tradition: pick the window your instrument actually moves in.

Write the window into the plan

A timing rule that lives in your head is not a rule. It should be written where the rest of the plan is: entries are taken only between these hours, positions opened in the window are closed or reduced by this hour, and nothing new is opened after it. The value is not that

these particular hours are magic. It is that the boundary is decided before the day starts, when you have no position and nothing to justify.

Expect the window to change what you measure



The window written as a procedure rather than an intention. Each step is a decision made before the day starts, when you hold no position and have nothing to justify.

A strategy tested across the whole day and then traded in one window is a different strategy, and its statistics do not carry over. If you narrow the hours, re-measure: average range, win rate, average time in a trade and typical adverse excursion all shift, sometimes substantially. This is not a formality. A stop that was comfortable against a full-day average range can be far too tight for the London open and far too wide for the Asian session.

08

A week, written down

A window becomes real at the point where it is recorded rather than intended. What follows is one week of a two-hour London window, logged as it happened, and it is worth reading for its shape rather than its result.

Day	Asian range	What the window did	Outcome
Mon	narrow	swept the high, reversed	+2.1R
Tue	wide	no setup inside the window	flat
Wed	narrow	broke and continued	-1.0R
Thu	narrow	swept the low, reversed	+1.8R
Fri	wide, news at 12:30	stood aside	flat

One week of a two-hour London window, as a record should look. The point is not the result — it is that every row can be checked against a rule written before the week began.

Two of the five days produced nothing, and both were correct. Tuesday offered no setup inside the window and the window is a boundary, not an obligation; Friday carried a scheduled release inside the hours and standing aside was the rule rather than a mood. A record in which every day contains a trade is a record of a trader who has not written a window at all.

The losing day is the one worth keeping. Wednesday looked identical to Monday and Thursday at the moment of entry — the same narrow Asian range, the same sweep — and went the other way. That is what a tested window buys: not days that work, but days that are alike enough to be counted, so that a month of them says something a month of improvisation never can.


09

Scheduled events are part of the clock

The economic calendar is a timing tool, and it is the one piece of genuinely forward-looking information every participant has for free. It does not tell you what will happen. It tells you exactly when the market will be told something, which is enough to plan around.

The largest scheduled movers for most instruments are central bank rate decisions and the press conferences that follow them, inflation releases, and employment data. Around these,

three things change at once: spreads widen, liquidity thins as market makers step back, and the range of possible outcomes over the next few minutes becomes much wider than the chart's recent history suggests. A stop placed on a technical level a minute before a rate decision is not protecting a position. It is a lottery ticket.



A scheduled release is a time point that changes what the windows around it mean. These are the three defensible responses; what is not defensible is the fourth, which is not knowing it was coming.

There are three defensible positions on this and they are all respectable: stand aside entirely for a defined period around the release; reduce size for it; or trade the aftermath once the initial spike has been absorbed and spreads have normalised, typically fifteen to thirty minutes later. What is not defensible is not knowing the release was coming. That is the only part of this the calendar removes entirely, and it takes a minute a day.

One further point, easy to miss: the release matters for the currency it belongs to, wherever you happen to be trading. United States data at 12:30 UTC moves gold, oil and the American indices as surely as it moves the dollar, and a trader working an unrelated pair can still be caught by the liquidity drain around it.

10

The clock is not the same for every instrument

Sessions describe the FX market well because FX is genuinely global and continuous. Other instruments have their own clocks, and applying the FX map to them produces confident nonsense.

- Currency pairs move most when at least one of their home centres is open. Euro and sterling pairs come alive at the London open; yen and Australasian pairs have their best hours during Asia; dollar pairs respond to New York across the board.
- Gold trades around the clock but takes most of its direction from the London and New York hours, and is unusually sensitive to United States data and to real yields.
- Stock indices have a hard boundary that FX does not: the cash open. An index future traded before the cash market opens is a thinner, different instrument, and the open itself is routinely the most volatile minutes of its day.
- Individual equities are the strictest case. Almost all of their volume occurs in the first and last hour of the cash session, and the middle of the day can be close to unusable for a short-term approach.
- Crypto never closes, which is often read as meaning it has no clock. It has one anyway: activity concentrates around the traditional session hours because the largest participants are the same institutions, and weekends are consistently thinner, which makes weekend moves both easier to start and less reliable.

Instrument	Liveliest hours	The thing that catches people
EUR, GBP pairs	London open, overlap	die off after London closes
JPY, AUD, NZD	Asian session	move on regional data nobody watches
Gold	London and New York	unusually sensitive to US data and real yields
Stock indices	cash open, first 30 min	futures before the cash open are a thinner instrument
Single equities	first and last hour	midday is close to unusable short-term
Crypto	traditional session hours	never closes, but weekends are consistently thinner

The session map is an FX map. Other instruments have their own clocks, and applying this one to them produces confident nonsense — the last row most of all.

The general rule behind all of these: an instrument's active hours are the hours in which the people who most need to trade it are at their desks. Work that out for what you trade, from its own data, and you will not need anyone else's session map.

Testing a timing edge without fooling yourself

Everything in this paper is a claim about tendencies, and tendencies should be checked against the instrument you actually trade rather than accepted because they are widely repeated. The work is not difficult and it does not require any special software.

Take at least a year of hourly data — not three months, which is one market regime wearing a disguise — and compute, for each hour of the day, the average range, the average absolute return, and how often that hour contained the day's high or low. The result is a profile of the day for your instrument, in your broker's clock, and it will make the general session map either concrete or irrelevant. This alone is worth more than any list of recommended hours, including the one in this paper.



The procedure that replaces every general claim in this paper with a measurement of your own instrument. It needs a spreadsheet and an afternoon, and nothing else.

Two traps are worth naming, because both are easy to fall into and both feel like research while you are doing them.

- Curve-fitting the clock. Given twenty-four hours and a few years of data, some hour will look excellent by chance alone. If an hour's edge cannot be explained by who is trading in it, treat it as coincidence until it survives on data you did not use to find it.
- Ignoring costs. Narrow windows mean fewer, and often shorter, trades. Spread and commission are a fixed toll per trade, so a timing rule that halves your average trade duration can quietly double the share of your profit that costs consume.

Curve-fitting the clock
Ignoring costs
How it looks

One hour tests beautifully, for no reason anyone can name.

The window works, but a narrower one means shorter trades.

The fix

Require a mechanism, and hold out data you did not use to find it.

Re-measure expectancy after spread, at that hour specifically.

The two ways an hourly edge disappears after you find it. Both feel like research while you are committing them, and the fix for each is different.

First hour of the window



prepared, rested

After a loss



revenge risk

Late in a session



tired, forcing

Last hour before stopping



trying to finish green

Schematic, and the timing edge nobody sells: when your own worst decisions happen. Traders keeping honest records commonly find losses concentrate in identifiable hours, and unlike the market's hours, these are entirely within your control.

One more thing worth measuring, though it is about you rather than the market: when do you make your worst decisions? Traders keeping honest records commonly find their losses concentrate in particular hours — late in a session, after a loss, in the last hour before they planned to stop. That is a timing edge too, and unlike the others it is entirely within your control.

12

Windows decay

A window is a measurement with a shelf life. It exists because particular participants are active at particular hours, and participants change — which means re-testing annually is maintenance rather than doubt.

Cause	What changes	Sign it has happened
A venue changes hours	the participants arrive elsewhere	the peak moves an hour
Central bank shifts announcement time	the day's time point moves	the old hour goes quiet
Volatility regime turns	range compresses everywhere	all hours shrink together
Your broker changes server clock	nothing real changes	everything appears to move at once

Why a window that worked stops working. None of these are failures of the method; they are changes in who is present, which is the thing the window was measuring all along.

The last row is the one that wastes the most time. A broker moving its server clock changes nothing about the market and everything about your notes, and the symptom — every hour appearing to shift at once — is diagnostic. If the whole profile moves by exactly one hour, look at the clock before looking at the market.

The others are real. When a venue changes its hours or a central bank moves its announcement, the participants genuinely relocate, and the correct response is to re-run the profile rather than to argue with it. A method that cannot survive its own measurement being redone was never measuring anything.

13

What this gets you, and what it does not

Used properly, the clock does three things. It tells you when to be at the screen, which is a claim about your attention rather than about price. It sets expectations for how far a move should travel, so that the same setup can be sized and stopped differently at different hours. And it produces reference levels — session highs and lows, opens, first-hour ranges — that are objective, recent, and shared by everyone else looking at the same chart.

	What the clock does	What it does not
Before the trade	Says when to be at the screen, and what a normal range looks like then.	Says nothing about direction. No hour is bullish.
During it	Sets size and stop from the hour's own volatility, not the day's average.	Cannot rescue an entry method you do not have.

The whole paper in four cells. The left column is what a tested window genuinely provides; the right is what it is routinely asked for and cannot give.

It does not tell you direction. No hour is bullish. Timing narrows the field and improves the odds attached to a decision you have already made on other grounds; it never makes the decision. A trader who has learned the sessions but has no method for what to do inside them has bought a better seat at a game they still cannot play.

And every tendency described here fails regularly. The London open sometimes does nothing at all. The Asian session occasionally produces the day's decisive move. The overlap is not always the deepest liquidity of the day, and on the days when a central bank surprises everyone, none of the usual shape survives. These are distributions, not laws, and the correct way to hold them is loosely: as a reason to expect one thing more than another, and never as a reason to skip the question of what you will do when it does not happen.

Term	As used here
Session	The working day of a financial centre.
Time point	A single instant with a known cause.
Time window	A bounded interval, named in one clock.
Overlap	Hours when two major centres are open together.
Session level	A price produced by a session: its extreme or its open.
Expansion	A period in which range widens and price travels.
Contraction	A period in which range narrows and price coils.

The terms used in a fixed sense throughout. Where the industry uses one of them loosely, the narrower definition is the one meant here.

Educational content only. Nothing in this document is financial advice, and no part of it is a recommendation to buy or sell anything.

Trading Papers