

TRADING KNOWLEDGE LIBRARY

Time as a Variable

Professional Trading



Trading Papers

Time as a Variable

Defining, testing and trading a session window

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	What this paper adds	4
02	Terms, stated precisely	5
03	Why the window is the testable unit	6
04	Naming the clock	8
05	What a window is allowed to claim	9
06	The protocol	11
07	Splitting the data	12
08	How hourly edges disappear	13
09	What the hour costs	15
10	Execution inside a validated window	16
11	One window, validated	17
12	The limits	18
13	Glossary	20

01

What this paper adds

Trading by the Clock argued that the hour changes the distribution of what happens next. That is a claim about the market. This paper is about what you can do with it: how to state a timing rule precisely enough to test, how to test it without discovering an edge that exists only in the sample you used to find it, and how to execute inside it once it survives.

The distinction matters because the first claim is nearly free to believe and the second is expensive to get wrong. Almost every trader agrees that some hours are better than others. Very few can say which hours, on which instrument, measured how, over what period, with what confidence — and the gap between those two states is where most of the money is lost.



The horizon decides how much of this paper applies to you. It is fixed before entry or it is not fixed at all — and the lower-right cell is where most timing research is quietly wasted.

A note on what is being claimed. Nothing here is a proprietary system, and none of it is unique to this library: the sessions are public knowledge, and the statistical cautions are standard practice borrowed from ordinary empirical work. What this paper offers is a way of putting them together that is precise enough to argue with.

02

Terms, stated precisely

Most disagreements about session trading are disagreements about words. Four definitions, used consistently for the rest of this paper:

Horizon

The intended lifetime of a position, decided before entry. Intraday means the position is closed within the trading day; multi-day means it is not. The boundary is not a matter of style — it determines which statistics apply to the trade, because a position held past its intended horizon is being managed by a rule you never tested.

The failure mode here is silent. A trade opened as intraday, left open because closing it would realise a loss, has become a multi-day trade by default rather than by decision. The position is the same; the plan governing it no longer exists.

Time point

A single instant with a known cause: a session open, a scheduled release, a fixing, a settlement. Time points are precise and often significant, but as a unit of analysis they are weak — each one produces a single observation per day, which is a slow way to accumulate evidence.




The three sessions on one 24-hour axis, in UTC. Boundaries are conventional rather than exact, and the shaded band is the London–New York overlap. Everything in this paper is an attempt to say something measurable about this picture.

Time window

A bounded interval, defined in one named clock, over which a measurable property of price is claimed to be stable. "07:00–09:00 UTC" is a window. "The London open" is not, until you say when it starts and ends and in whose clock.

Session

The working day of a financial centre, which is the underlying cause of most windows worth having. A session is why a window exists; the window is the part you can measure.



Schematic, and the first output of the protocol in section 07. This is the shape a liquid FX pair tends to take in UTC; your instrument will have its own, and computing it is the point at which this paper stops being reading.

03

Why the window is the testable unit

The argument for preferring windows over points is statistical rather than aesthetic, and it is worth spelling out because it also tells you how long a window should be.

A time point gives you one observation per session. Two years of data is roughly five hundred observations, which sounds like plenty until you split it by direction, by day of week, or by whether the previous day trended — after which any subgroup you actually care about has a few dozen members and a confidence interval wide enough to drive through.

A window gives you a distribution. Over the same two years, a two-hour window contains not one number per day but a range, a path, a maximum excursion in each direction, and a relationship to the levels around it. You can ask how often it expanded beyond the previous day's range, how often its first thirty minutes contained its extreme, how its behaviour changed after a trend day. Each of those is a testable statement built from the same underlying data.

Unit	Observations in two years	What each one gives you
A time point	about 500	one number: the move at that instant
A two-hour window	about 500	a range, a path, an extreme, a relationship to levels
Point, split by weekday	about 100 each	a confidence interval too wide to use
Window, split by weekday	about 100 each	still a distribution in each subgroup

Why the window rather than the moment. The same two years of data yields two very different amounts of evidence depending on which unit you choose, and the difference is not marginal.

There is a second, subtler reason. A time point is exactly the kind of variable that invites overfitting: with twenty-four hours and a few years of data, some specific minute will look extraordinary by chance alone. A window is harder to fit to noise because it aggregates — and when it does show an effect, you can usually name the participants who cause it.

The practical rule that follows: prefer the shortest window that still contains enough activity to measure, and be able to explain who is trading in it. An hour you cannot explain is a coincidence waiting to be discovered.

04

Naming the clock

A timing rule that does not name its clock is not a rule. There are at least four in play and they disagree with each other, sometimes by an hour, sometimes twice a year.

- UTC — never shifts, which is why it is the only sensible reference for research.
- Exchange local time — what the venue itself uses, and what its own schedule is published in.
- Broker server time — usually somewhere between UTC+1 and UTC+3, chosen for the broker's convenience, and what your chart actually displays.
- Your local time — the only one that determines whether you will really be at the screen.

The conversions are trivial and the errors are not. A cash equity session that runs 09:30–16:00 in New York is 14:30–21:00 UTC in winter and 13:30–20:00 UTC in summer, and it lands on different broker-clock hours depending on whose server you are on. Note also that it ends at 16:00 local — a rule that says "until the end of the day" and means midnight has quietly included four hours of a different market.

Clock	Moves?	What it is good for
UTC	never	research, storage, anything compared later
Exchange local	with local daylight saving	reading the venue's published schedule
Broker server	usually, unpredictably	nothing — but it is what your chart displays
Your local	with your daylight saving	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 — by you, never mind by anyone else.

Daylight saving deserves its own line because the northern hemisphere does not change together. Europe and the United States shift on different weekends, leaving a few weeks each spring and one each autumn when the gap between London and New York is an hour

off its usual value, and the overlap window moves with it. Japan observes no daylight saving at all, so the Asian session drifts relative to both twice a year.

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

The weeks each year when the usual gap between the centres is wrong. During these, a rule stored in broker time lands an hour away from the data it was validated on, and the results quietly stop describing the same thing.

The protocol is short: do all research in UTC, record the broker offset alongside every result, convert to local time only for display, and re-verify the mapping after each changeover. A stored result whose clock is unrecorded is not a result.

05

What a window is allowed to claim

Before testing anything, be clear about which kind of claim you are making, because they need different evidence and offer different value.

Distributional claims

"This window's average range is larger than the day's average." These are the easiest to establish and the most useful, because they change position sizing and stop distance rather than direction. They are also the most stable across regimes.

Structural claims

"The day's extreme forms inside this window more often than chance would predict." Harder, but testable: the null is the proportion of the day the window occupies. A two-hour window in a twenty-four-hour market covers 8.3% of the day; if it contains the daily high or low 25% of the time, that is a real concentration and worth knowing.

Claim	Example	Evidence needed	Survives out of sample?
Distributional	this window's range is wider	modest	usually
Structural	the day's extreme forms here	moderate	often
Conditional	after a trend day, it retraces	large — the sample halves	sometimes
Directional	this window is bullish	very large	almost never

The four kinds of claim, ordered by how much evidence each needs and how much it is worth once you have it. The order is almost exactly the reverse of how often each is made.

Conditional claims

"After a trend day, this window opens with a retracement." These are where most edges live and where most self-deception happens, because conditioning splits your sample and every additional condition halves it again.



What each additional condition costs. Starting from five hundred observations, three conditions leave you with a subgroup whose apparent edge is indistinguishable from chance — which is why conditional claims are where self-deception lives.

Directional claims


"This window is bullish." Treat these as false until proven otherwise, and expect them not to survive. Time concentrates activity; it does not choose sides. An hour that appears directional in-sample is the single most likely thing in this paper to be noise.

06

The protocol


What follows is a procedure, not a strategy. It produces a window you can defend, or it produces nothing — which is also a result.

- Fix the instrument and the horizon first. A window derived for one instrument does not transfer to another, and one derived for intraday work says nothing about a multi-day position.
- Take at least two years of hourly data, in UTC. Three months is one regime wearing a disguise; a single year contains one instance of each seasonal effect and cannot distinguish it from noise.
- Split the data before looking at it: the first two thirds for exploration, the last third sealed. Nothing you find in the first part is real until it survives the second.
- In the exploration set, compute for each hour: mean and median range, mean absolute return, the frequency with which the hour contained the day's extreme, and the mean maximum adverse excursion for a position opened at its start.



The protocol in section 07, compressed to the five steps it cannot be run without. The third is the one that is skipped, and skipping it is what makes the other four worthless.

- Look at the profile before forming a hypothesis. The shape of the day for your instrument will usually be obvious, and it will usually correspond to a session boundary. If it does not, be suspicious rather than pleased.
- State the window as a hypothesis with a mechanism: which participants, why that hour, what they are doing. If you cannot write the mechanism in one sentence, you have found a pattern rather than a cause.
- Test it once on the sealed third. Once. A second look after adjusting the window is a new experiment, and it needs new data.
- Re-measure everything your plan depends on inside the window: average range, hit rate, average time in trade, adverse excursion. A stop calibrated on the full-day average is wrong in both directions — too tight for the open, too wide for the quiet hours.



A window as a decision boundary rather than a preference: entries only inside it, positions opened in it managed afterwards, nothing new after that. The hours shown are an example. What matters is that the boundary exists before the day starts and is the same boundary your statistics were measured on.

07

Splitting the data

The protocol's third step carries all of its weight, and it is worth a section of its own because it is both the cheapest step to perform and the one most reliably skipped.

Portion	Share	Permitted use	Once used
Exploration	first two thirds	look freely, form hypotheses	reusable
Sealed	last third	one test of one stated hypothesis	spent — permanently
Live forward	everything after today	the only honest confirmation	accumulates slowly

How the data is divided and what each part is permitted to do. The rule that makes this work is the last column: a part used for the wrong purpose cannot be restored to usefulness by intending better next time.

The discipline it requires is unusual: you must decide how the data will be used before you have seen what is in it. That feels arbitrary, and it is exactly the point. Any division made after exploration is a division chosen — perhaps unconsciously — to contain the result you already found.

What makes the sealed third valuable is that it can be spent only once. A hypothesis that fails there is finished; adjusting the window and testing again on the same data is not a second attempt but a slower form of the first mistake, and the fact that it is done in good faith changes none of the arithmetic.

When a hypothesis does fail on sealed data, the correct response is to return to the exploration set with a genuinely different idea, and to accept that the next test will have to wait for genuinely new data. Windows are cheap to imagine and expensive to validate, and the expense is the reason a validated one is worth anything at all.

08


How hourly edges disappear

Four failure modes account for most of the timing research that works until it is traded. All four feel like diligence while you are committing them.

Multiple comparisons

Testing twenty-four hours at the conventional 5% threshold means expecting roughly one spurious result per pass by construction. Testing twenty-four hours across five instruments

means expecting six. The fix is not a cleverer threshold; it is holding out data and requiring a mechanism.




Expected spurious results at the conventional five per cent threshold, purely by construction. Nothing in this figure is about markets — it is what testing many things at once does, and it is why an unexplained hour should be treated as a coincidence already found.

Regime confusion

Volatility regimes persist for months. A window discovered in a period of central bank tightening may be describing that period rather than that hour. Requiring the effect to appear in at least two visibly different regimes costs you sample size and buys you the only kind of confidence that matters.

Cost blindness

Narrow windows produce fewer and shorter trades, and spread and commission are a fixed toll per trade regardless of duration. A rule that halves your average holding time doubles the share of your gross edge that costs consume. Worse, spreads are themselves a function of the hour: the widest spreads of the day sit exactly where the largest ranges do.



Spread as a share of the same hour's median range, schematic. The widest spreads sit exactly where the widest ranges do, so the two partly cancel — and the quiet hours, which look cheap on a chart, are the expensive ones.

Survivorship in your own record

A window that "works" in your journal may be one you only traded when conditions looked favourable. The record then measures your selectivity rather than the hour. This is why the protocol above tests on price data rather than on your executed trades.

09

What the hour costs

Every result in this paper is measured gross and traded net, and the gap between the two is itself a function of the hour. This is not a rounding correction — for narrow windows it is frequently the whole of the edge.

Failure	What it actually is	The remedy
Multiple comparisons	testing until something passes	hold out data; demand a mechanism
Regime confusion	describing a period, not an hour	require two visibly different regimes
Cost blindness	measuring gross, trading net	re-measure expectancy at that hour's spread
Survivorship in the journal	measuring your selectivity	test price data, not executed trades

The four ways a validated window turns out not to have been validated. Each one feels like diligence at the time, and each has a different remedy — which is why naming them separately is worth the space.

The third row of that table is the one worth expanding. Spread is a toll per trade, not per hour held, so a rule that narrows the window and shortens the average trade raises the share of gross profit that costs consume, even though nothing about the market has changed. A window that halves your holding time has doubled your cost ratio by construction.

It compounds with the hour itself. Spreads are widest where liquidity is thinnest, which is to say in exactly the quiet hours that look calm and tradeable on a chart. The correct measurement is therefore not the spread in points but the spread as a share of that hour's own median range, which is the only

form in which the numbers from different hours are comparable at all.

Compute it once for your instrument, hour by hour, on your own broker's data. It takes an afternoon, it settles the question of which hours can support your method, and it is the single most frequently omitted step between a validated window and a profitable one.

10

Execution inside a validated window

A tested window changes four concrete things. None of them is an entry signal.

It sets the expected distance. If the window's median range is known, a target beyond it is a low-probability request and a stop far inside it will be reached by ordinary noise. Sizing follows from the same number: risk a constant fraction of capital, and let the window's volatility determine position size rather than the other way round.

It sets the cost budget. Knowing the typical spread in that window tells you the minimum move that can be profitable, which in turn tells you whether the setup you are considering is worth taking at all at that hour.

It sets	From	What changes in the plan
Expected distance	the window's median range	targets and stop distance
Position size	the window's own volatility	risk stays constant across hours
Cost budget	the typical spread in that hour	which setups are worth taking at all
The boundary	the window's end	when you stop opening positions

What a validated window is actually for. None of these four is an entry signal, and a trader expecting a fifth row is expecting the framework to do a job it has never claimed.

It provides reference levels that are objective and shared: the window's own high and low, the previous session's extremes, the opening price. These are not predictions. They are the coordinates everyone else is looking at.

And it defines when you stop. A window is a boundary in both directions: entries inside it, management afterwards, nothing new later. Deciding that boundary before the day starts is the entire point — it is made when you hold no position and have nothing to justify.

What a window does not do is tell you what to do inside it. Timing narrows the field and improves the odds attached to a decision made on other grounds; it never makes the decision. A trader with a validated window and no entry method has bought a better seat at a game they still cannot play.

11

One window, validated

The protocol is easier to trust once it has been watched running. What follows is a single window taken from an untouched dataset to a tradeable rule, with the result at each stage and nothing omitted — including the stage at which the effect shrank.

Stage	What was done	Result
Fix	EUR pair, intraday horizon	stated before any data was opened
Data	two years hourly, UTC	sealed the final eight months
Profile	range and extreme frequency by hour	clear peak at 07:00-09:00
Mechanism	London desks open; overnight flow clears	one sentence, so the hypothesis stands
Hypothesis	the window holds the day's extreme	in-sample: 24% against 8.3% expected
Sealed test	one test, no adjustment	21% - survives
Costs	spread as share of that hour's range	5% - the method clears it

One window taken through the whole protocol, with the result at each stage. The hypothesis survived, narrowly, and the last row is the reason it was tradeable rather than merely true.

Two features of that table are the point. The in-sample concentration was 24 per cent against an expected 8.3, and the sealed test returned 21 — smaller, as it almost always is, and

still a real concentration. An effect that survives out of sample at a reduced size is the normal shape of a genuine finding; one that survives unchanged should be inspected for a leak between the two sets.

The other feature is the last row. A validated structural claim is worth nothing if the spread in that hour eats the move it predicts, so the cost measurement is part of the validation rather than a step afterwards. Had that row read twenty per cent instead of five, the finding would have been true and untradeable, and the correct response would have been to record it and stop.

Firm	the range in this window is wider than the day's average
Firm	size and stop distance should be computed from it, not the day
Held	the day's extreme forms here more often than chance predicts
Weak	after a trend day it opens with a retracement - sample too small
None	the window is bullish - never tested, never claimed

What the surviving window is now allowed to say, in descending order of confidence. Nothing on this ladder is an entry, and the bottom rung is the one that most often gets promoted to the top by wishful reading.

What the window now licenses is set out on that ladder, and the discipline is to trade only the rungs that were actually tested. The conditional claim near the bottom was interesting and its sample was too small; leaving it there, unpromoted, is the difference between a framework and a story.

12

The limits

Three limits are structural rather than incidental, and a framework that hides them is being sold rather than explained.

First, windows decay. They exist because particular participants are active at particular hours, and participants change: venues open, algorithms are retired, a central bank moves its announcement time. A window is a measurement


with a shelf life, and 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
An announcement time moves	the day's time point relocates	the old hour goes quiet
Volatility regime turns	range compresses everywhere	all hours shrink together
Your broker changes server clock	nothing real	everything appears to move at once

Why a window that worked stops working. Only the last row is a mistake; the others are the market changing who is in it, which is the thing the window was measuring all along.

Second, the effect is conditional on liquidity, not on the clock itself. On a holiday, in a crisis, or on the day of an unscheduled announcement, the hour tells you nothing because the population it describes is not there. The clock is a proxy for who is present; when the proxy breaks, so does everything built on it.

Third, and most often ignored: the largest single timing variable in most trading records is the trader. Losses concentrate in identifiable hours — late in a session, after a loss, in the last hour before a planned stop. That effect is usually larger than any market effect in this paper, it requires no data licence to measure, and it is the only one entirely within your control.



Schematic of the effect this paper cannot measure for you and which is usually the largest one in the record: where a trader's own losses concentrate. It costs nothing to measure and it is the only variable here entirely within your control.

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

13

Glossary

The vocabulary of this paper, gathered. Each of these is used in one sense throughout, and where a term is contested in ordinary usage the narrower reading is the one intended.

Term	As used in this paper
Adverse excursion	the furthest a position moves against you before closing.
Conditional claim	a claim holding only when a stated condition is met.
Horizon	the intended lifetime of a position, fixed before entry.
Out-of-sample	data withheld during exploration and used exactly once.
Overlap	hours in which two major sessions are open together.
Regime	a persistent market state, typically lasting weeks to months.
Time point	a single instant with a known cause.
Time window	a bounded interval, in a named clock, over which a property is claimed stable.

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

One term deserves a longer note than a table row allows. Out-of-sample is not a property of data; it is a property of how the data was used. A sealed third that has been looked at, even once, even briefly, even for a different question, is no longer out of sample — and no amount of care afterwards restores it. This is the single rule in the paper that cannot be recovered from once broken, which is why it is stated twice.