

TRADING KNOWLEDGE LIBRARY



The Stops Above the High

Market Analysis



Trading Papers

The Stops Above the High

Liquidity, sweeps and stop hunts — the pattern that is real and the story that is optional

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	A word doing two jobs	4
02	Why the orders are there	5
03	The pattern, drawn	7
04	One history, two futures	8
05	Selected on the outcome	10
06	Defining it before you know	11
07	Where the queues actually form	14
08	Inducement, taken seriously	15
09	Eight names, three claims	17
10	Counting it for yourself	19
11	The most expensive bar of the day	21
12	What to do with a queue	22
13	What is knowable here	24
14	What this does not settle	25


01

A word doing two jobs

Liquidity is one of the few words in this field that is both perfectly precise and hopelessly vague, depending on who is holding it. In its precise sense it is a measurable property of a market: how many orders are resting, at what prices, and therefore how much you can trade before you move the price against yourself. In its other sense it is a noun for the fuel above the highs and below the lows — something the market goes to get.

Sense	What it refers to	Can you count it?
Liquidity as depth	resting orders in the book; how much you can trade without moving price	yes, directly
Liquidity as a pool	stop orders clustered at an obvious price	yes, by inference
Liquidity as a target	the claim price was driven there to trigger them	no, not as stated
Liquidity as an excuse	the reason your stop was hit	no, and it is not meant to be

One word, two jobs. Almost every disagreement about liquidity is two people using different rows and assuming they are using the same one. The left row is measurable and dull; the right row is vivid and, as normally stated, cannot be checked.



The first sense, drawn. Beside the price is the book: orders resting at each level, thicker where more of them sit. Nothing here is hidden or sinister — it is a public queue, and it is thin in the places nobody has a reason to care about.

That first figure is the precise sense drawn. Beside the price is the book — the queue of orders waiting at each level — and it is thicker where more people have a reason to be waiting. There is nothing hidden about it. On many venues you can watch it change in real time, and on all of them its shape is a consequence of where people think the price matters.


This paper is about what happens when the second sense is layered on top of the first: when the observation that orders cluster becomes the story that price is driven to them. Both halves of that sentence are worth taking seriously, and they need to be taken seriously separately, because one of them can be counted and the other cannot.

The vocabulary here is the field's own and is used in its ordinary sense throughout. Nothing in this paper argues that the terms are meaningless. It argues that most of them are three claims wearing eight names, and that separating the countable part from the unfalsifiable part is what makes any of it usable.

02

Why the orders are there


Start with the part that is not in dispute, because it is arithmetic rather than interpretation.



The second sense begins here. Three attempts stop within a point of each other, and the result is a line that every single person watching this instrument can draw, in the same place, without discussing it. That agreement is the whole mechanism.

Three attempts stall within a point of each other. The result is a line that every person

watching this instrument can draw, in the same place, without discussing it with anybody. No one had to be told; the line is a function of the chart, and there is only one chart.



What that agreement does to the book. Nobody coordinated it and nobody had to: a stop is a function of the chart, and there is one chart. The bar just above the line is not evidence of a plan, it is evidence that arithmetic has no imagination.

Now look at what that agreement does to the book. Everyone short into the range has a stop just above the line, because just above the line is where being short is wrong. Everyone waiting for a breakout has a buy order in the same place, because a break of the high is their entry. Everyone long from the bottom has a stop under the low for the mirror-image reason. Nobody coordinated any of this and nobody had to.

Who	Their order	Why exactly there
Anyone short the range	buy stop	the high is where being short is wrong
Anyone waiting for a breakout	buy stop	a break of the high is the entry
Anyone long from the low	sell stop	the low is where being long is wrong
Any book with a bracket order	both	the bracket is drawn from the same two lines
Anyone who read the same course	same as above	the course taught the same level

Who is standing in that queue, and why each of them chose the same price without meeting. Read the right column downward: not one of these people needed to be told where to put the order, and not one of them is trying to trap anybody.

Read the right-hand column of that table downward. Not one of those people is trying to


trap anyone. Each of them independently solved a small problem — where is my idea wrong — and because they were all looking at the same picture, the independent answers landed on top of each other. That is the entire mechanism, and it is enough to produce everything the rest of this paper discusses.

It is worth pausing on how ordinary this is. Cars queue at the same junction without conspiring. Every one of them chose the shortest route independently, and the shortest route was the same route. The queue is real, the congestion is real, and there is no one arranging it.

03

The pattern, drawn

Here is what the clustering produces on a chart, drawn properly rather than described.




The canonical pattern, drawn properly. Price reaches through the equal highs by a point and a half, trades there briefly, and closes back underneath. Everything above the line was filled on the way through, and there is now nothing left up there to fill.

Price reaches through the equal highs by a point and a half. Everything resting above the line is filled on the way through — the shorts stopped out, the breakout buyers filled — and then there is nothing left up there to fill. The next bar closes back underneath, and the move that follows runs in the opposite direction to the push that started it.

This is a real phenomenon and it does not need defending. The mechanism is straightforward: a burst of buy orders arrives all at once because they were all conditional on

the same price; the burst is absorbed by whoever was willing to sell up there; and once the burst has passed there is no more of it, because it was finite. What was fuel a moment ago is now exhaust.




The same pattern as a sequence, so that each stage can be argued with separately. Only the first three are observable. The fourth is where a mechanism is being asserted, and the fifth is where a motive is being added on top of it.

Set out as a sequence, the pattern becomes easier to argue with. The first three stages are observable — the level is on the chart, the clustering follows from the level, the burst is visible in volume. The fourth is a mechanism being asserted, and it is a reasonable one. The fifth adds something the first four do not contain, and the rest of this paper is about what it adds and what it costs.

04

One history, two futures

Everything so far is compatible with the story that someone is doing this deliberately, and also compatible with nobody doing anything deliberately. To see why that matters, it helps to look at the case the screenshots never show.



One history, two futures, on the same axes. Every bar to the left of the divider is shared: the same level, the same push, the same close beyond it. To the right, one branch returns and one keeps going — and the name the whole picture is given depends entirely on which. This is the figure the rest of the paper is about.

Every bar to the left of the divider is shared. The same level, the same three attempts, the same close beyond it by a point. To the right, the picture splits: in one branch price comes back inside and runs the other way, and in the other it keeps going and never returns.

Both branches happen, regularly, on every instrument. And the name the whole picture receives depends entirely on which one you are looking at. The left branch is a liquidity sweep. The right branch is a breakout. Nobody calls the right branch a failed sweep, because by the time it has failed it has been renamed.

If it comes back	If it keeps going	What actually happened
a liquidity sweep	a breakout	a close beyond a level
a stop hunt	a genuine move	resting orders were filled
the trap was set	the level gave way	supply above the line ran out, or did not
smart money accumulated	smart money was already long	unobservable either way

The same six bars, named twice. Every term in the left column is applied after the outcome is known, and every one of them carries a claim about intent that the bars themselves do not contain.

That table is worth reading slowly, because it is the hinge of the argument. Each row is one pair of bars receiving two different descriptions on the basis of what happened afterwards.


The bars did not change. What changed was the outcome, and the outcome was allowed to reach back and rename the cause.

This is not dishonesty and it is not unique to trading. It is how language works when a phenomenon is named for its ending. But it has a consequence that is severe and specific here, and it is the subject of the next section.

05

Selected on the outcome

If a sweep is only called a sweep once price has returned, then every sweep ever recorded has, by definition, returned.



How a study of sweeps goes wrong without anybody lying. Only the top-left cell ever gets collected, because only the top-left cell gets the name — and a sample that is defined by its outcome will report that outcome every time.

Look at where the sample comes from. Only the top-left cell is ever collected, because only the top-left cell gets the name. The bottom-right cell — the identical push that kept going — is silently reclassified as a breakout and never appears in any tally of how often sweeps work. It is not hidden. It is filed elsewhere.

Sweeps as usually collected		true by construction, worth nothing
oses beyond a level that return, instrument A		counted in advance
Same, instrument B		different, and that matters
Same, but only inside one session window		where a real edge would hide

What that does to the number people quote. The top bar is what a sample selected on its outcome must report. The lower bars are what the same instrument gives when the candidate is defined before the outcome — which is the only version that carries information.

The top bar is what a sample defined by its outcome must report, and it is worth nothing at all. It is not a discovery about the market; it is a restatement of the definition. Ninety-three minutes of examples in which the pattern worked is ninety-three minutes of a tautology, however carefully the examples are drawn.


The lower bars are what the same instrument produces when the candidate is fixed before the outcome is known, and they are unremarkable numbers. That is exactly what makes them useful: an unremarkable number can be compared against a coin, against a cost, and against last quarter. A hundred per cent cannot be compared against anything.

The general form of this problem is treated in its own right by the paper on what a claim has to show, elsewhere in this library. It appears here in a particularly pure form, because the naming convention itself performs the selection, and it does so before anybody has decided to be careless.

06

Defining it before you know

The repair is mechanical and slightly tedious, which is why it is rarely done. It is to fix the candidate at a moment when the answer is not yet available.



A candidate defined before anyone knows the answer: any five-minute close more than one point beyond the level, marked at the close of that bar. Four of them here. Two came back, two did not, and nothing about the marked bars told you which was which.

Four candidates here, each marked at the close of the bar that produced it, using a rule stated in advance: any close more than one point beyond the level. Two of them came back inside. Two of them carried on. Standing at the close of each marked bar, nothing distinguished the ones that would return from the ones that would not.

Candidate	Closed beyond by	What followed	Kept in the count?
1	1.0 point	returned inside within 2 bars	yes
2	2.0 points	continued for 3 more points	yes
3	3.0 points	returned inside within 2 bars	yes
4	2.0 points	continued to a new high	yes


Those four candidates counted honestly, including the two nobody would have screenshotted. This is a sample of four and proves nothing; the point is the shape of the ledger, which has a column that the usual version does not have.

The ledger has a column the usual version does not have: whether it was kept. Every candidate is kept, including the two that would ordinarily have been quietly renamed. Four is not a sample and settles nothing — the shape of the ledger is the point, not its contents.

Field	This rule says
Universe	one instrument, chosen before the week starts
Level	a high touched twice, marked before the session opens
Candidate	a close more than one point beyond that level
Trigger	the next close back inside the level
Invalidation	any close a further two points beyond
Size	fixed fraction, from the invalidation distance
Record	every candidate, including the ones that never triggered

The claim written as a rule a stranger could run — the same seven fields the plan paper in this library argues for. The line that does the most work is invalidation, because it is the line that decides in advance what a failed sweep looks like.

Written as seven fields it becomes something a stranger could run, which is the standard this library applies everywhere. The line carrying the most weight is invalidation, because it is the line that decides in advance what a failure looks like. Without it, there are no failures, only breakouts.



What the rule calls a failure, drawn so it cannot be quietly reclassified later. Price closed beyond, never came back inside, and kept going. Under the naming convention this figure is a breakout; under the rule it is a sweep that did not work, and it stays in the count.

And here is one of those failures drawn, so that it cannot be reclassified after the fact. Under the naming convention this figure is simply a breakout and nobody would think to count it

against anything. Under the rule it is a candidate that was invalidated, and it stays in the tally with the same weight as the ones that worked.

07

Where the queues actually form

If the clustering is the real part, then the practical question is where it is thickest, and the answer follows directly from the mechanism rather than from anyone's theory.



A day's worth of obvious levels and the queue at each. The sizes are not measured, they are reasoned: a line more people can draw carries more orders. That is the only claim in this figure, and it is enough to be useful.

The sizes in that figure are not measured; they are reasoned. A level that more people can draw carries more orders, because the clustering was never about the level being special — it was about the level being shared. Yesterday's high is thick because nobody has to interpret it. A retracement level is thin because two people using different anchors get different lines and their stops land in different places.

The level	Who can draw it	How thick the queue
Yesterday's high and low	everyone, identically	thickest
Round numbers	everyone, identically	thick
Overnight high and low	everyone on the same session clock	thick
Equal highs in today's range	everyone watching today	moderate
A moving average	only those with the same setting	thin
A retracement level	only those anchoring identically	thinnest

The same levels ranked by the only property that drives the clustering: how many people can draw the line without argument. A level requiring a setting, a preference or an interpretation is a level fewer people share, and a thinner queue.


That ranking has a slightly uncomfortable implication for the more elaborate end of the subject. The more sophisticated a level is — the more settings, judgement or doctrine it requires — the fewer people share it, and the thinner the queue behind it. The crude levels carry the crowd. This is the opposite of the usual assumption that better analysis finds better levels, and it follows from the mechanism rather than from any opinion about the analysis.

It also gives the first genuinely actionable consequence in this paper, though not the one people expect. If you want to know where an overshoot is likely to have something behind it, look for the line the largest number of people can draw without argument, and be suspicious of levels that only your own chart produces.

08

Inducement, taken seriously

The strongest version of the intent story deserves a proper hearing rather than a dismissal, so here it is at its most reasonable.



The inducement claim, drawn as fairly as it can be drawn. A small high forms early and is taken; the argument is that it existed to attract the orders that were then used. Read the bars alone and this is a minor high being exceeded, which is what minor highs mostly do.

A minor high forms early in the session and is taken out shortly afterwards. The inducement claim is that this high existed in order to attract orders — that it was formed so that the orders resting above it could be used to fund a larger move. Read the bars on their own and this is a small high being exceeded, which is what most small highs do.

The claim	What would have to be true	How you would check
The high attracted orders	orders rested above it	book data, or the count of moves that stall there
Those orders were used	volume spiked as it broke	tape at the break, on a real sample
It was created for that	someone chose to form that high	no observation distinguishes this from chance
Therefore the next one repeats	the sequence recurs above chance	count every minor high, not the memorable ones

Inducement taken seriously and asked the ordinary question. The middle column is what would have to be true; the right column is what you would have to see to believe it. Row three is where the claim stops being about the market and becomes about the person telling it.

Take the claim apart into its parts and something clarifying happens. The first two rows are ordinary claims about the world and can be checked, roughly, with data most people can obtain. The third row is different in kind: no observation distinguishes a high that was

created on purpose from one that simply formed, because both produce identical bars.

That is not a small objection. A claim that no observation could contradict is not a strong claim that happens to be hard to test; it is a claim that has stopped making contact with the market and is now making contact only with the person telling it. And notice that the fourth row — the only one that would actually help you — does not depend on the third at all. If the sequence recurs above chance, you can trade it without believing anything about why.

There is one more reason the intent story is attractive, and it is worth naming without contempt. It explains a specific, painful experience — my stop was hit and then it turned around — in a way that supplies an antagonist and removes the sting. That is a real service. It is simply a psychological service rather than an analytical one, and the two are easy to confuse when the second would cost money to obtain.


09

Eight names, three claims

The liquidity vocabulary is large, and its size is often taken as evidence of depth. It is worth laying the terms out beside what each one actually asserts.



A fair value gap: three bars where the middle one runs far enough that the first bar's high and the third bar's low do not overlap. The shaded band is the gap. The claim attached to it is that price returns to fill the band — which is a claim about a price that mattered, wearing new clothes.



An order block: the last opposing bar before a decisive move, marked, with the claim that price returns to it and turns. Compare it with the gap above and with any support level anywhere — three names, three drawings, one underlying assertion about a price that will matter again.

A fair value gap is a three-bar arrangement where the middle bar runs far enough that the outer two do not overlap, and the claim attached is that price comes back to fill the band. An order block is the last opposing bar before a decisive move, and the claim attached is that price comes back to it and turns. Both drawings are perfectly clear. Both claims are the same claim.

The name	The claim underneath
Liquidity pool	a price that mattered before will matter again
Order block	a price that mattered before will matter again
Fair value gap	a price that mattered before will matter again
Breaker block	a price that mattered before will matter again
Liquidity sweep	a move that overshoots reverses
Stop hunt	a move that overshoots reverses
Turtle soup	a move that overshoots reverses
Inducement	obvious structure is created to be used

The liquidity vocabulary mapped onto the claims it makes, in the manner of the five-claims paper in this library. Read the right column down: three entries, and the left has three times as many. That ratio is not an accusation — it is a reason to test three things instead of nine.

Read the right-hand column downward: three distinct assertions. The left column has nearly three times as many entries. The five-claims paper in this library makes this argument across trading strategies generally, and the liquidity vocabulary is the most concentrated example of it anywhere in the field.

This is not an accusation of fraud. Vocabularies grow this way naturally, as each teacher names the version of the pattern they noticed. But the practical consequence is precise: if you are testing eight things you are doing eight times the work and dividing your evidence eight ways. If you are testing three, each of them gets three times the data, and the results will actually mean something.

10

Counting it for yourself

Everything above is preparation for something you can do in about a week of evenings, and which will tell you more about your own instrument than any course can.



How to settle this for yourself, on your own instrument, in about a week of evenings. The third step is the one that is skipped, and skipping it is what turns a study into a scrapbook.

The third step is the one that gets skipped, and skipping it is the difference between a study and a scrapbook. Keeping the continuations in the same list, with the same weight, is uncomfortable precisely because they are the cases that make the pattern look ordinary. That discomfort is the sensation of learning something.

The number	What it settles
Candidates found	whether the pattern is common enough to matter
Share that returned inside	whether the claim beats a coin at all
Average distance travelled after	whether the return is worth more than the spread
Result net of costs, at one size	whether you have anything

The four numbers worth having at the end of that week, and what each one settles. The last row is the one that decides whether any of it is tradeable, and it is the row a chart study never produces.

Four numbers come out of it. The first tells you whether the pattern is common enough to build anything on. The second tells you whether it beats a coin. The third tells you whether the return, when it happens, is large enough to be worth the attempt. The fourth is the only one that decides anything, and it is the one a chart study structurally cannot produce, because a chart study does not pay costs.




What a real edge looks like when it is measured rather than sold. Nothing in this range is exciting, and the top bar is the one being advertised. A method that returns fifty-five per cent of the time at a sensible reward is a good method; a method claiming ninety is a method that has been selected on its outcome.

It is worth knowing in advance what a good answer looks like, so that a good answer is not mistaken for a disappointing one. A pattern that resolves in your favour fifty-five per cent of the time at a sensible reward is a genuinely useful method and considerably better than most people ever find. The number implied by the promotional version of this subject corresponds to no sample that has ever been collected honestly.

11

The most expensive bar of the day

Suppose the counting comes back favourable. There is a cost problem waiting, and it is peculiar to this pattern rather than general.



The moment the pattern completes is the most expensive moment in the session to act. The stops firing are exactly what widens the quote, so the bar you most want to buy is the bar with the worst price on offer — a point the five-minute paper in this library makes in general and that applies here at its sharpest.

The moment the pattern completes is, mechanically, the worst moment in the session to transact. The burst of stop orders that makes the pattern is the same burst that empties the near side of the book, so the quote is at its widest and fills arrive furthest from where they were asked for. The bar you most want to act on is the bar with the worst price on offer, and this is not a coincidence — it is the same event seen from the other side.

	If the pattern works	After the cost of that bar
Move captured	8 points	8 points
Spread at a quiet moment	-1 point	—
Spread as the stops fire	—	-4 points
Slippage on a market order there	—	-2 points
Left over	7 points	2 points

The same moment in money. The middle column is what the pattern is worth if it works perfectly; the right column is what is left after paying to participate in it at the one moment participation is dearest.


The arithmetic is unforgiving. A pattern worth eight points that costs six to enter is a pattern worth two, and two points is inside the range where ordinary variation decides the outcome. The paper on the five-minute chart in this library makes the general version of this argument; here it applies at its sharpest, because the signal and the cost are produced by the same event.

The two responses are the familiar ones and both give something up. Wait for the return to be confirmed and enter into the quieter bar afterwards, accepting that you will miss the ones that never pause. Or work a limit order at the level in advance and accept that you will be filled on the ones that keep going, which are precisely the ones you did not want. There is no version in which you get the burst's price without paying the burst's spread.

12

What to do with a queue

The most useful consequence of everything above is not an entry, which is where every treatment of this subject puts it. It is a decision about where your own order goes.



The consequence that pays whether or not the story is true. Two stops for the same position: one at the obvious price where the queue is, one a little beyond it. The first is taken by the overshoot; the second is not, and the position that survives is the one that was never in the queue.

Two stops for the same position. One sits at the obvious price — the round number, the prior low, the place the queue forms. The other sits a little beyond it. The overshoot takes the first and misses the second, and the position that survives is the one that was never standing in the queue.

The elegance of this is that it does not depend on any of the contested part. You do not have to believe anyone is hunting anything. You only have to believe that orders cluster at obvious prices — which is arithmetic — and that a cluster is a place where fills are worse and overshoots are more likely to reach. Both follow from the mechanism in the second section of this paper.

But it is a trade and not a free lunch, and the middle rung is where it usually goes wrong. A wider stop is only defensible if the position size falls in proportion. Widen the stop and keep the size and you have not avoided the queue, you have simply arranged to lose more when you are eventually wrong — which is a way of converting a small, frequent annoyance into a large, rare disaster.

Buys	●	the position is not in the queue everyone else is standing in
Costs	●	a wider invalidation, and therefore a smaller position
Requires	●	the size to fall in proportion, every time, without exception
Fails when	●	the wider stop is used at the old size to feel safer
Testable by	●	counting how many of your stops were taken by less than a point

What that costs and what it buys, stated plainly, because it is a trade and not a free lunch. The middle rung is the one people get wrong: the wider stop is only defensible if the size comes down with it, and if it does not, this is a way of losing more slowly.

There is a cheap diagnostic for whether any of this applies to you, and it costs one evening. Go through your last hundred losing trades and count how many were stopped out by less than a point beyond your stop, and then went the way you had originally expected. If that number is small, this section is not your problem and you should ignore it. If it is not small, you have just found something worth more than any entry rule in this paper.

13

What is knowable here

It is worth setting out the whole subject on one grid, because the confusion it generates is almost entirely a confusion about which cell a statement belongs in.

The bottom-left cell holds the work: orders cluster at obvious prices, that clustering is inferable, and it has consequences for where you put your own orders and what you pay to

trade near them. It is dull, it is checkable, and almost nobody spends their time there.

The top-right cell holds the entertainment. Who was on the other side, what they intended, whether the move was engineered — these are genuinely interesting questions, they are unavailable from a chart, and answering them would not change a single decision you could make. That combination is what makes them so durable as content: they can be discussed indefinitely without ever being settled or acted upon.

	Observable	Not observable
Predicts something	Orders cluster at obvious prices. Count it, trade it, or trade around it.	Who was on the other side. Interesting, unavailable, and not needed.
Predicts nothing	Volume was elevated on that bar. True, and it tells you nothing on its own.	The market was hunting your stop personally. The market does not know your name.

The epistemics of the whole subject on one grid. The bottom-left cell is where the useful work is and where almost nobody spends their time; the top-right is where the ninety-three-minute courses live, because it is the only cell that is entertaining.

The bottom-right cell deserves a sentence of its own, because it is where the emotional weight of this subject sits. The market did not hunt your stop. It does not know your name, your position, or your account size. What happened is that you put your order where a hundred thousand other people put theirs, for exactly the same sensible reason, and the queue was visible from further away than you were.

That is a less satisfying story than the alternative and a considerably more useful one, because it names something you control. You cannot stop a market maker from existing. You can move your order two points.

14

What this does not settle

Three limits, in the spirit of the rest of this library.

The first is the most serious. Predatory flow is not a fantasy; there are venues and

The limit	What it means for this paper
Some venues really do have predatory flow	intent is sometimes real; it is still not observable from a chart
The counting is instrument-specific	a result on one market says little about another
Clustering is public knowledge	an edge everyone can see is an edge that erodes

Three limits on everything above, stated as plainly as the argument was. None of them rescues the intent claim, and none of them is a reason to ignore the clustering.

arrangements in which someone genuinely does have both the information and the incentive to push price into a cluster. Nothing in this paper says that never happens. What it says is that from a chart you cannot tell the difference, and that the trading decision is the same either way — which makes the question interesting rather than useful.

The second is a caution about your own results as much as anyone else's. A count that comes back favourable on one instrument in one session window is a fact about that instrument in that window. The mechanism generalises; the number does not, and the most common way a genuine result is destroyed is by being carried somewhere it was never measured.

The third is the one that eventually applies to everything in this field. Every person reading this paper can see the same clusters, and an edge that is public is an edge under pressure. That does not make the clustering go away — it is produced by the structure of stop placement itself, not by ignorance — but it does mean the reversal claim built on top of it may thin out over time while the cost consequence stays exactly where it is.

The useful summary is two sentences. Stops cluster at obvious prices because everyone reads the same chart and solves the same problem the same way, and that clustering is real, inferable, and worth acting on. The claim that someone drives price into those clusters on purpose adds a villain, adds nothing you can check, and changes no decision you could make — so keep the queue, drop the hunter, count your candidates before you know how

Term	As used in this paper
Liquidity	resting orders available to be filled at a price.
Pool	an unusually dense cluster of resting stop orders.
Sweep	a close beyond a level followed by a close back inside it.
Candidate	a close beyond a level, before the outcome is known.
Inducement	the claim that obvious structure exists in order to be used.
Selection on outcome	defining a category by the result it produced.

The terms this paper uses in a fixed sense. They carry the same meaning across the rest of the library.

they end, and move your stop off the line everyone else is standing on.

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