TRADING KNOWLEDGE LIBRARY

The Market That Isn't Trending

Price Action

Trading Papers

The Market That Isn't Trending

The range as a distribution, and the one break that pays for it

Published by	**Trading Papers**
Series	**The trading knowledge library**
Edition	**First edition · 2026-08-26**
Format	**A4 · digital**

Contents

01

A state you can only name afterwards

The first thing to settle about ranges is not how to trade them but when they begin to exist, because the answer is later than anyone draws it.

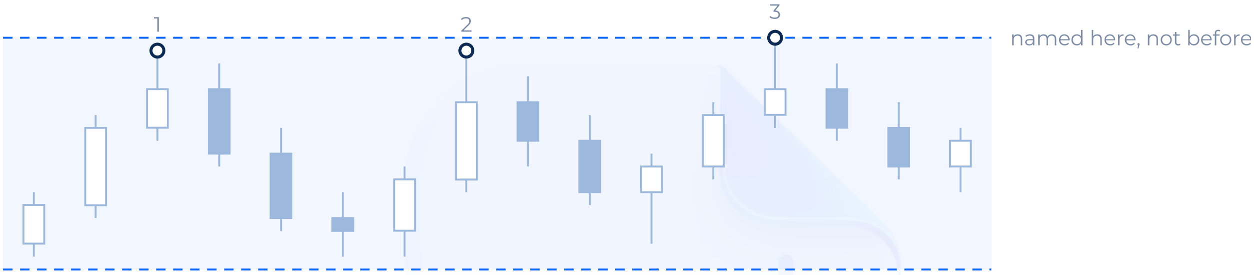

The moment a range becomes nameable, and how much of it has already happened by then. The first touch is a level. The second is a coincidence. Only at the third does a reader call it a range, and the shaded stretch to its left is the part that was already over.

A level touched once is a level. Touched twice it is a coincidence with a good story attached. It becomes a range at the third touch, when the edge has turned price back often enough that a reader will name the pattern, and by then the shaded stretch on the left of that figure has already happened.

After	What you know	What is left
one touch	a level was reached once	everything, but nothing is confirmed
two touches	it turned there twice	most of it, and the second turn may be chance
three touches	the edge has held	about half the range's life
four or more	the edge is well established	the part nearest the break

What the three touches cost to wait for. Read the right column: by the time the pattern is nameable, most of the move that defined it has been given away, and the trader who acted earlier was acting on something that was not yet a range.

Read the right-hand column downward. Every row buys certainty and spends opportunity, and the exchange rate is fixed by the pattern rather than by the trader. This is not a flaw in anyone's method. A range is a claim about repetition, and repetition cannot be observed the first time.

The consequence runs through the rest of this paper. Every range you trade is one you recognised late, which means the trades available to you are the later touches, and the later touches are the ones nearest the break. The method does not get to choose the good half.

> This paper assumes a swing rule, which the structure paper in this library settles, and the run-length count, which the first paper of this collection settles. Neither is repeated here.

02

What most of the time counts

Markets are said to range most of the time. The number attached to that claim is usually seventy per cent, sometimes eighty, and it is quoted as though it were a property of markets.

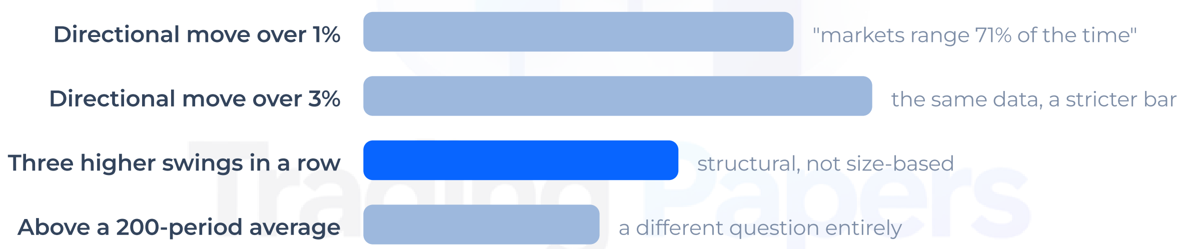

The famous share, recounted four ways on identical data. The only thing that changed between these rows is the rule for what counts as trending, and the answer moves by more than thirty points.

Those four rows are the same data. The only thing that changed is the rule for what counts as trending, and the answer moves from thirty-nine to eighty-four. A number that swings by forty-five points on a definition is not describing the market.

The rule	What it actually asks	Sensitive to
Move exceeds x%	did price travel far	the threshold, entirely
Swings in one direction	did structure advance	the swing rule
Price against an average	is price above its own past	the lookback length
Volatility below its median	is the market quiet	the sample the median came from

Why those four disagree. Each row is a defensible rule, and each measures a different thing, so the widely quoted number is not a fact about markets. It is a fact about whichever rule the person quoting it happened to use.

The third column is where each rule keeps its assumption. A percentage threshold hides in the threshold; a structural rule hides in the swing definition; a moving-average rule hides in the lookback. None of them is wrong, and no two of them answer the same question.

So the useful version of the claim is narrower and worth more. On your instrument, under your written rule, some share of bars sit inside a range you could have named at the time. That number you can produce, and this paper is about what to do once you have it.

03

The edge is a zone

A trader draws the edge as a line because a line is what fits on a chart. The touches do not land on it.

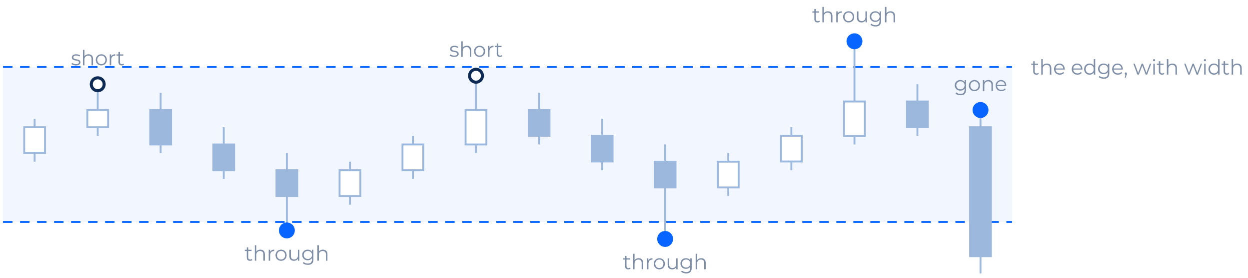

Six touches at one edge, drawn at the price each bar actually reached. Two stopped short of the line, three went through it and came back, one kept going. A drawing that reduced this to a line would erase the only thing worth knowing about it.

Six touches at one edge, each drawn at the price its bar actually reached. Two turned before arriving, three went through and came back, and the last kept going. The line is somewhere in the middle of that scatter, and it was never the boundary.

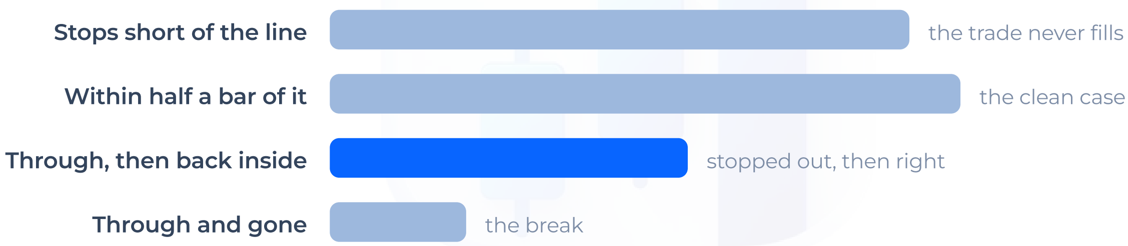

Where the touches actually land, across two hundred of them at one instrument. The line a trader draws is the mode of this distribution, not its boundary, and nearly a third of touches finish on the far side of it.

Across two hundred touches the shape is stable. About a third never reach the line, about a third arrive cleanly, and about a third finish on the far side of it. The third group is where the method is decided, because those are the trades that stop out and then turn out to have been right.

That is the practical content of calling the edge a zone. A stop at the line is inside the distribution of ordinary behaviour, so it will be hit by touches that were about to work. The zone's width is a number you can measure, and measuring it is what the stop section later depends on.

04

How often the edge holds

Everything so far has been about recognising the pattern. The method rests on one number, and it is the share of touches that turn back.

Touch number	Held	What it means
1st	—	not yet an edge
2nd	about 78%	the level is doing something
3rd	about 74%	the best odds you will see
4th	about 71%	still good
5th and beyond	about 66%	decaying, and the break is closer

The number the method rests on, and it is not the one people quote. An edge holding four times in five sounds like a business; the right column is what makes it one or not.

Roughly three touches in four hold. That is a high number and it is the source of both the method's appeal and its danger, because a trader who has seen four hold in a row has seen the ordinary case and will treat it as evidence about the fifth.

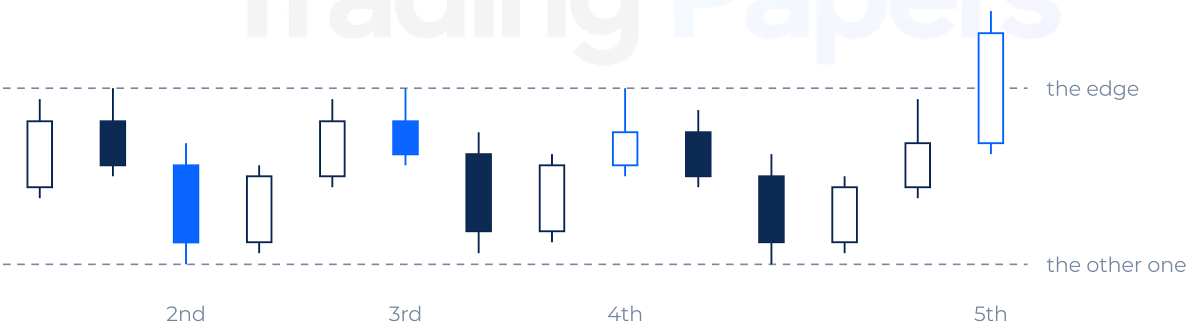

The decay drawn out. Each touch is slightly likelier to be the last one than the touch before it, which is the exact opposite of the trend case, where age told you almost nothing. Here age is information, and it points one way.

Notice the direction of the decay. The second touch is the strongest and each one after it is slightly weaker, which is the exact opposite of what the first paper of this collection found for trends. There, a run's age told you almost nothing about its future. Here age is informative, and everything it says is bad news.

The asymmetry between those two findings is not a coincidence, and it is the subject of the next section. A trend that has run a long way has demonstrated something about its own persistence. A range that has held a long time has spent some of the demonstration it began with.

05

The mirror

Range trading is trend following reflected. Once that is seen, most of the advice about both becomes derivable rather than memorable.

	Trend following	Range trading
Win rate	about a third	about three quarters
Average win	large, and unbounded	small, and capped by the range
Average loss	small and fixed	small until it is not
Where the tail is	on the right, in the winners	on the left, in the breaks
What kills it	capping a winner	not capping a loser
How it feels wrong	long stretches of losing	one afternoon

The two methods side by side on the numbers that describe them. Read down: every line is reversed. This is one distribution seen from two sides, and it is the reason the two methods fail in opposite ways.

Read down the two columns. The win rates invert, the win sizes invert, and the tails swap sides. The first paper described a method that is wrong most of the time and paid by rare enormous winners. This one describes a method that is right most of the time and charged by rare enormous losers.

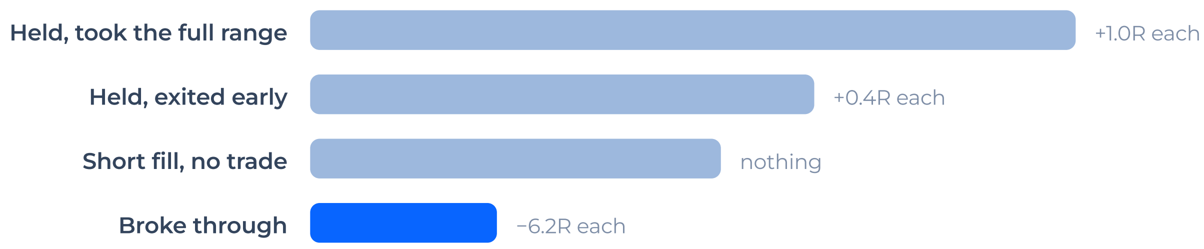

A hundred range trades, sorted by outcome. The first three bars are the business, and they are pleasant. The last one is every one of them handed back, and it arrives without notice.

The bars make the arithmetic concrete. Sixty-eight trades in a hundred make money, twenty-two never fill, and ten break through. Those ten cost more than the sixty-eight made, which is why the win rate is the wrong place to look and the size of the tenth outcome is the right one.

The mirror explains the shape of the mistakes too. Trend traders fail by capping the winner, because the winner is uncomfortable to hold. Range traders fail by not capping the loser, because the loser is uncomfortable to accept. Both are the same instinct applied to opposite distributions, and both remove exactly the part that pays.

06

Height against cost

Before any judgement about a particular range, there is an arithmetic filter that removes most of them.

Start with	the distance from one edge to the other
Subtract	the overshoot you must place the stop beyond
Subtract	the part nearest the far edge, where you exit early
Subtract	the spread, twice, and the cost of holding
What is left	is the trade, and it is often nothing

Why most ranges are untradeable before any judgement is applied. Each rung subtracts something from the distance between the edges, and what survives at the bottom is the trade.

Read the rungs as subtraction from one quantity. The distance between the edges is what the trade could pay. The stop must sit beyond the overshoot, which takes some. The exit happens before the far edge, because waiting for it means waiting through the touches that fail, which takes more. Costs take the rest.

Edge to edge	After stop room	After early exit	After costs
a narrow range	0.6×	0.4×	below break-even
a medium range	0.7×	0.5×	a thin positive
a wide range	0.8×	0.65×	the trade

The same subtraction on three real-sized ranges. Only the widest survives, which is why a method that takes every range it finds is not the same method as one that takes the ranges worth taking.

On three real-sized ranges only the widest survives the subtraction. This is worth doing before deciding whether a range looks good, because looking good is not a property the arithmetic recognises, and a narrow range can look every bit as clean as a wide one.

The filter also explains a common and expensive experience. A trader who takes every range that appears will find that the method works in backtest and loses in life, because the ranges that were too narrow to pay were also the most numerous and the easiest to spot.

07

The trade that pays for nothing

There is one trade the method must never take, and the method's own logic recommends it.

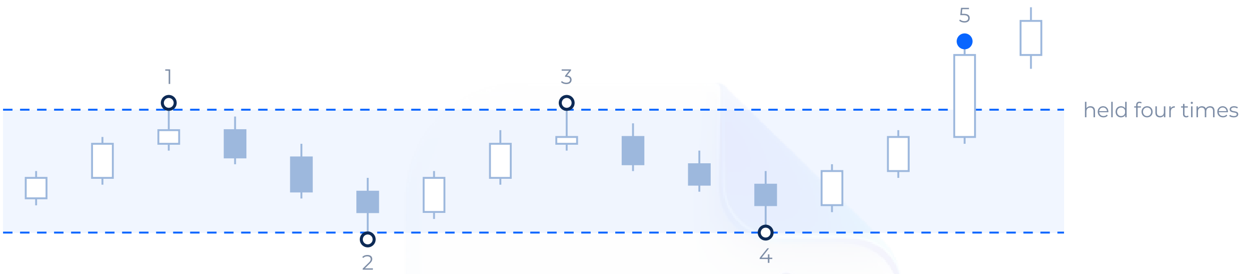

The one trade the method must never take, drawn at the moment it looks most attractive. The edge has held four times, so the fifth touch is the most confident entry the pattern has produced, and it is the entry standing in front of the break.

The edge has held four times. The fifth touch is therefore the most confident entry the pattern has produced, arrived at by exactly the reasoning the method teaches, and it is standing in front of the break.

Why the trap is built into the method and cannot be reasoned away. Each step is correct on its own, and following all five in order puts the largest position at the worst moment.

Each of those five steps is correct in isolation. An edge with more holds behind it has more evidence behind it. Confidence should track evidence. Size tends to track confidence. And a range ends exactly once, at some touch, which by construction is a late one.

So the trap is structural. It is not carelessness, it is the method executed faithfully, and it cannot be reasoned away by being more careful at the fifth touch. What breaks the chain is the third step: refusing to let size track confidence, so that the touch that breaks costs the same as the touch that holds.

08

Where the stop goes

The overshoot distribution from earlier now does its work, because it settles where the stop belongs.

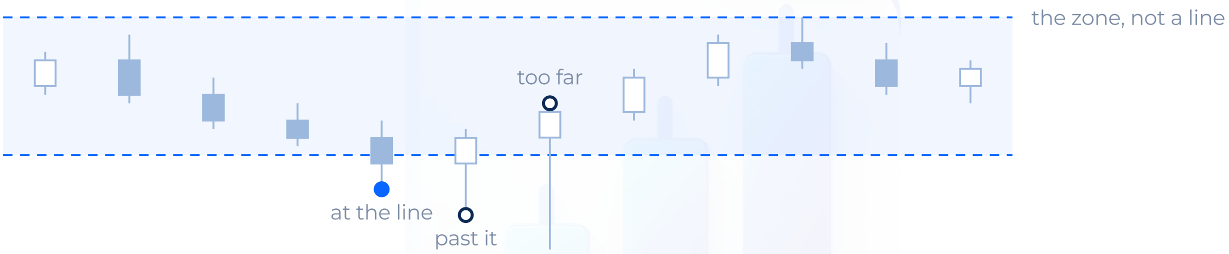

Three stops on the same trade. The first sits at the line and is taken out by an ordinary overshoot that resolves in its favour. The second sits past the overshoot distribution and survives it. The third is so far out that the trade it protects can no longer pay for it.

Three stops on one trade. At the line, an ordinary overshoot takes it out and the trade then resolves in the direction it was placed. Past the overshoot, the same bar leaves it alone. Far outside, it survives everything and the loss it eventually takes is bigger than the range could ever have paid.

Stop placed	Stopped out	Average loss	Total
At the line	44%	0.5R	+9R
Past the overshoot	19%	1.0R	+38R
Far outside	9%	2.8R	−4R

The same hundred trades under three stop placements. The middle row is not a compromise between the other two; it is the only one where the losses and the wins are the sizes the method assumed.

The ledger is unambiguous. A stop at the line converts winning trades into small losses, which is why its stop-out rate is more than twice the middle row's. A stop far outside converts a manageable loss into an unmanageable one. The middle row is not a compromise; it is the only placement where wins and losses are the sizes the method assumed.

The rule that follows is short. Put the stop beyond the measured overshoot, not beyond a comfortable-looking distance, and if the resulting stop is too wide for the range to pay for, the answer is to skip the range rather than to move the stop.

09

Size, and the left tail

In a method with a fat left tail, position size is the method. Nothing else decides whether the break is survivable.

	Small fixed size	Size rising with confidence
While the edge holds	Steady small gains. Unremarkable and durable.	Faster gains, and the record looks better.
When it breaks	One loss the account absorbs.	The loss lands at the largest size ever used.

Why size is not a separate topic in a method with a left tail. The bottom-right cell is where a long run of held edges quietly puts a trader, and it is the cell where the break arrives.

The grid has one dangerous cell and the danger is in how a trader arrives there. Nobody chooses maximum size on a whim. They arrive after a long run of correct trades, each one adding evidence, and the size that meets the break is the largest they have ever used precisely because the evidence was strongest just before it.

Sizing	Over 40 holds	The break costs	Net
Fixed 1%	+40R	−6R	+34R
Fixed 2%	+80R	−12R	+68R
Raised after each hold	+126R	−94R	+32R

The same edge, the same break, three sizing policies. The last row is not recklessness; it is the ordinary human response to a run of correct trades, and it is what turns a working method into a single bad afternoon.

The third row is the arithmetic of that arrival. Raising size after each hold earns half as much again over forty holds, and then hands back nearly all of it in one afternoon. The trader who did this was not reckless. They were responding to evidence, in the one method where responding to evidence is the error.

So the sizing rule for ranges inverts the intuition the pattern builds. Size must be flat across touches, and if it varies at all it should fall as the touch count rises, because the hold rate falls with it. A method that sizes up on the fifth touch has arranged to be largest exactly where it is weakest.

10

What it feels like from inside

The experience of range trading is the reason its risk is carried for so long without being noticed.

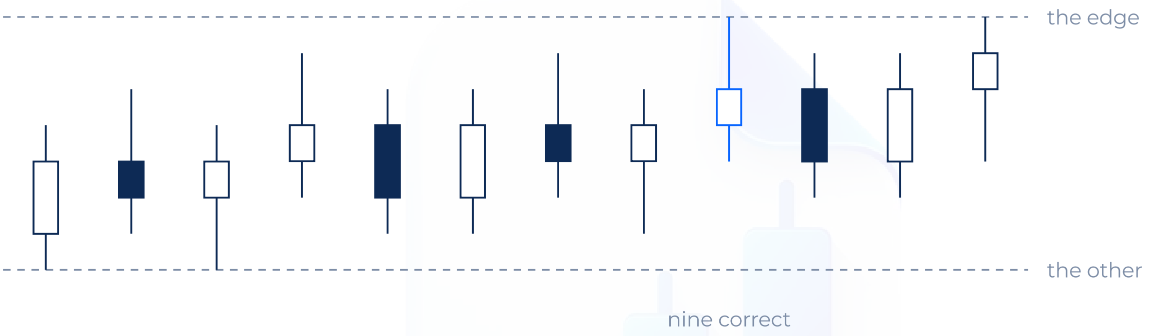

What the method delivers between the breaks. Eleven trades, nine of them correct, and none of them interesting. This stretch is the reason range trading is easy to keep doing and hard to keep sizing correctly.

Eleven trades, nine of them correct, none of them memorable. This is what the method delivers between breaks, and it goes on for weeks. Nothing in the stretch is exciting and nothing in it is alarming.

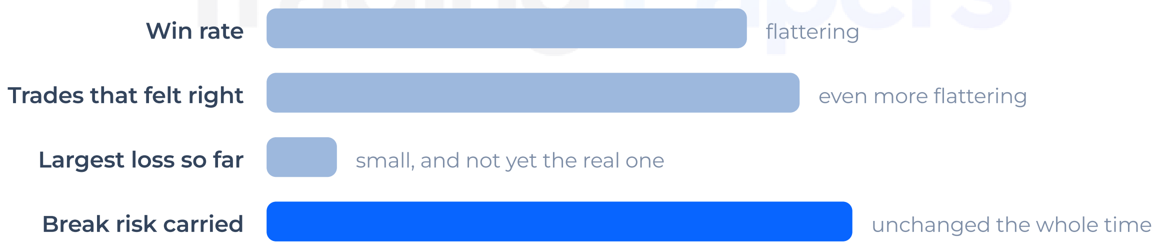

The trader's own record over that stretch, and why it misleads. Every measure a person naturally checks says the method is working, and every one of them is silent about the only number that decides it.

The record from that stretch says the method is working, and every measure it offers is true.

The win rate is high. The trades felt right. The largest loss so far is small. None of those numbers has anything to say about the break, because the break has not happened yet and its size does not depend on how the previous forty trades went.

Compare this with the first paper's account of trend following, where the difficulty was a long unbroken experience of being wrong. Here the difficulty is the opposite and harder to defend against: a long unbroken experience of being right, which is indistinguishable from a method that has stopped carrying risk.

11

Which markets actually range

The shape described so far is common. The numbers filling it are not, and two instruments can look identical on a chart while only one of them supports the method.

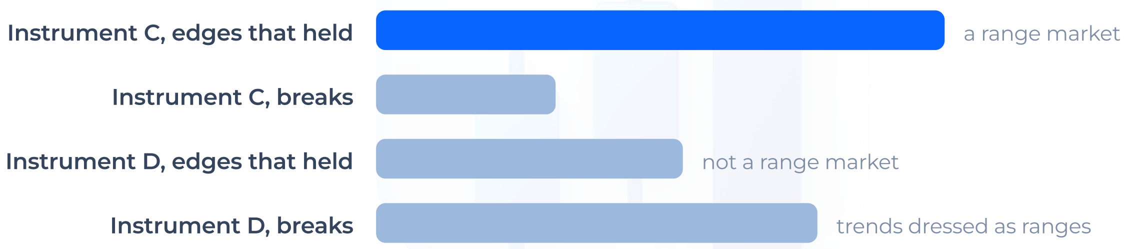

Two instruments counted over one period. The first spends most of its life between edges that hold; the second looks similar on a chart and produces half as many holds, because its edges are crossed rather than respected.

Instrument C holds three quarters of its touches. Instrument D holds fewer than half, which means its edges are being crossed rather than respected, and what a trader draws as a range there is a trend pausing.

	Instrument C	Instrument D
Edges that held	76%	41%
Median touches before break	6	2
Overshoot beyond the line	small	large and uneven
Suits range trading?	yes	no — the count says trend

The two instruments on the numbers that decide the method. Instrument D is the one that ruins range traders, because on a chart it is indistinguishable from C until the count is run.

The middle row is the tell that a chart cannot show. On instrument C the median edge survives six touches; on D it survives two. A trader who has drawn a clean-looking range on D is looking at a pattern that will typically end before the third entry.

This is the quiet failure that matches the one the first paper described for trends. The rule is followed correctly, the entries are taken at the edge as taught, the stops are placed properly, and the instrument was never going to hold. Counting first is what separates a method executed badly from one executed faithfully in the wrong place.

12

When the range ends

A range ends once, and it gives some warning. The warning is weak, and treating it as strong is its own mistake.

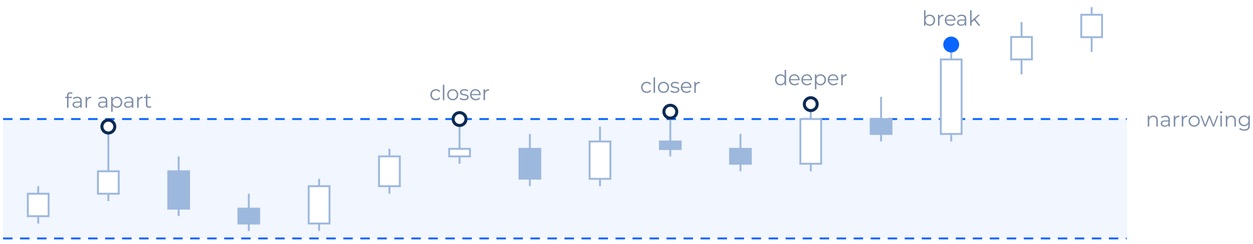

The tells that precede a break, in the order they appear. The touches arrive closer together, each one reaches further into the edge, and the returns toward the middle get shorter. None of the three is decisive, and all three together are the most warning the pattern gives.

Three things change before the break. The touches arrive closer together, each one reaches further into the edge, and the returns toward the middle fall shorter. All three are visible in that figure and all three are also visible before plenty of ranges that go on for months.

What you see	What it means	What to do
One deep overshoot	ordinary; the zone has width	nothing
Touches arriving closer together	weak evidence, common before nothing	nothing
Each return falls shorter	the range is compressing	expect less, exit earlier
All three at once	the best warning available	halve the size, or stand aside

What to do about each tell, which in two cases out of four is nothing. The last row is the only one that says act, and what it asks for is smaller size rather than a prediction.

The table is deliberately unexciting, because in two cases out of four the right answer is to do nothing. One deep overshoot proves nothing; the zone has width and deep touches are part of it. Touches arriving closer together happens constantly before nothing at all.

The last row is the only one that asks for action, and what it asks for is smaller size rather than a prediction. That distinction matters. Predicting the break requires being right about timing; halving the size requires only noticing that the evidence has weakened, which is a much easier thing to be right about.

13

Counting it yourself

Every number in this paper is one you can produce for your own instrument, and the values here are worth less to you than your own.

1. Choose one instrument and one timeframe, and change neither
2. Write down what counts as a touch and what counts as a break
3. Mark every touch and record whether it held
4. Measure how far each touch went past the line
5. Divide the holds by the touches, by touch number

How to produce the hold rate for your own instrument. The second step is the one that decides every number after it, and writing it down is what makes the count repeatable rather than an impression.

The second step is the one that decides all the others. What counts as a touch, and what counts as a break rather than an overshoot, are choices, and a different pair of choices produces a different hold rate from the same data. Write both down before you start.

The number	What it settles
Hold rate by touch number	whether to take the fifth touch
Overshoot distribution	where the stop goes
Median range height	whether the trade clears its costs
Worst single break	what size you can carry

The four numbers the count produces and the decision each one settles. The last is the one that decides whether the method belongs on this instrument at all.

The four numbers settle four decisions. The hold rate by touch number tells you whether the fifth touch is a trade. The overshoot distribution tells you where the stop goes. The median height tells you whether the trade clears its costs at all.

The last number is different in kind, because it is not about the method. The worst single break in your sample is a fact about what you will have to absorb, and it should be checked against your size rather than against your confidence. A break you cannot carry at your current size is a break that ends the account rather than the trade.

14

What this does not settle

Three limits, and then the short version.

The limit	What it means here
A range is named late	you never trade its first half
Hold rates move with regime	recount; the number is not a constant
The definition decides the count	another rule for touch and break gives other numbers

Three limits, in the spirit of the rest of this library. None of them rescues the advice, and none is a reason to skip the count.

The first is structural and has no remedy. A range is named after it has repeated, so the half of it that was cheapest to trade is always already gone. Any method that claims to catch ranges early is claiming to identify repetition before it repeats.

The second is the ordinary caution about any count. Hold rates move with regime, and a rate measured over two quiet years describes those two years. Recount periodically, and treat a falling hold rate as information about the instrument rather than about your recent luck.

The third is the dependency this paper leaned on throughout. What counts as a touch and what counts as a break are rules you chose, and every number here moves when they move. That is not a flaw in the count; it is the requirement to state the rule beside the result, which is the same requirement the structure paper made of the swing rule.

Term	As used in this paper
Range	a stretch where an edge has turned price back at least twice.
Edge	the zone at which those turns happened, not the line drawn through them.
Touch	a bar that reaches the edge zone, whether or not it goes through.
Overshoot	how far past the line a touch went before turning.
Hold	a touch that turned back inside; a break is one that did not.
R	one unit of planned risk: the distance from entry to stop.

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

The useful summary is three sentences. A range is a description applied after an edge has held twice, so you always trade its later half, and the later touches are the ones nearest the break. The method's edge is a high hold rate, roughly three touches in four, decaying slowly as the touch count rises — the mirror of trend following, whose win rate was low and whose tail was on the other side. Which means the whole method reduces to one obligation: keep the size flat as confidence rises, because the range ends exactly once and it ends at the touch you were most sure of.

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