TRADING KNOWLEDGE LIBRARY

How Long a Trend Lasts

Price Action

Trading Papers

How Long a Trend Lasts

Persistence as a number rather than a slogan

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

Contents

01

A slogan with a number inside it

The oldest sentence in this business is that the trend is your friend, and it is usually said in a tone that closes the subject. It should not. Underneath the friendliness there is an ordinary factual claim, and the claim has a number attached to it.

What is usually said	What it would have to mean	Countable?
The trend is your friend	trends are pleasant to trade	no
Trends persist	a move underway tends to continue	not as stated
Don't fight the tape	trading against a move is unwise	no
A run of N bars tends to reach N+1	the next bar continues more often than chance	yes

The slogan taken apart. Read the right column: only the last row is something the market could contradict, and it is the only row anybody can act on. The other three are moods.

Read the right-hand column downward. Three of those rows are moods — agreeable, unfalsifiable, and useless. The fourth is a statement about the world that could turn out to be false, and it is the only one you can act on, because it is the only one that survives being written down.

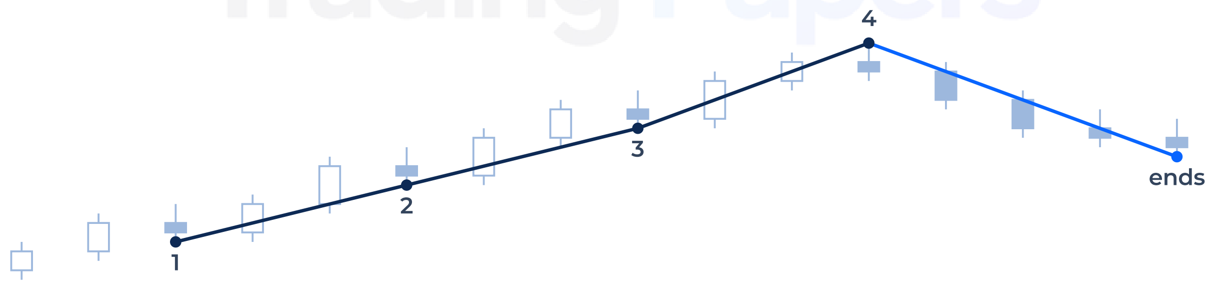

A run, marked by a written rule: consecutive higher swing lows on a stated swing definition, ending at the first lower one. Six swings, five of them higher, and then the one that stops it. The whole of this paper is about the length of that count and nothing else.

So here is a run, marked by a rule rather than by feel: consecutive higher swing lows on a stated swing definition, ending at the first lower one. This paper takes that definition entirely for granted. Which bars count as swings, and why that choice matters more than people think, is the subject of the structure paper in this library, and nothing here repeats it.

> Everything below assumes you have a written swing rule. If you do not, the numbers in this paper cannot be reproduced — not because the method is delicate, but because you would be counting a different thing each time you looked.

02

The shape of the answer

Count enough runs and a shape appears, and the shape is more useful than any single statistic taken from it.

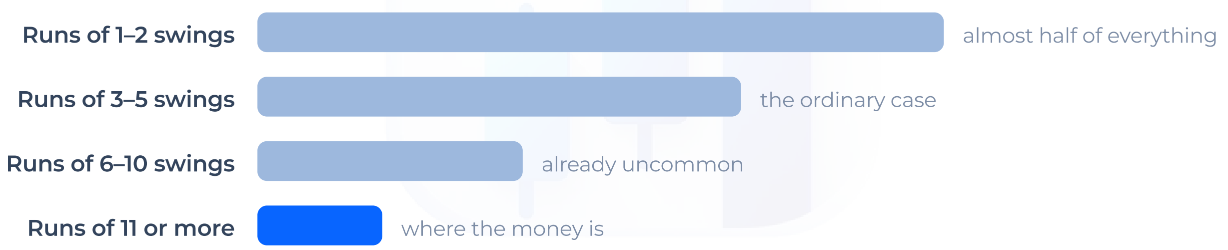

How the run lengths actually fall out. The shape is the point: short runs are common, long ones are rare, and there is no bulge anywhere in the middle for an average to sit inside honestly.

Nearly half of all runs are over almost as soon as they start. A third are ordinary. Fewer than one in five gets past six swings, and the runs long enough to be worth talking about are rare enough that most traders will go a whole quarter without holding one.

The statistic	Value	What it describes
Mean run length	about 4.1 swings	a run that is neither typical nor profitable
Median run length	2 swings	what you will usually be holding
Most common length	1 swing	the single likeliest outcome
Longest decile	12 or more	where the whole return sits

Why the average is the wrong number and quietly misleads. The mean is dragged upward by the rare long runs, so it describes a trend that is longer than most and shorter than the ones that matter — a length the market almost never actually produces.

Now look at what happens to the average. The mean is about four swings, and four swings is a length the market produces relatively seldom: it is longer than the typical run and far shorter than the ones that matter. The rare long runs drag the mean upward, past the crowd of short ones, without ever being represented by it.

This is not a statistical curiosity. A trader who has internalised four as the expected length will hold too long on the short runs and, much more expensively, will take profit far too early on the long ones — because at four swings a run that is going to reach thirty looks exactly like one that has already overstayed.

03

Identical from inside

Everything above is a fact about a large sample. The difficulty is that you never trade a sample; you trade one run, from inside it, without knowing which kind it is.

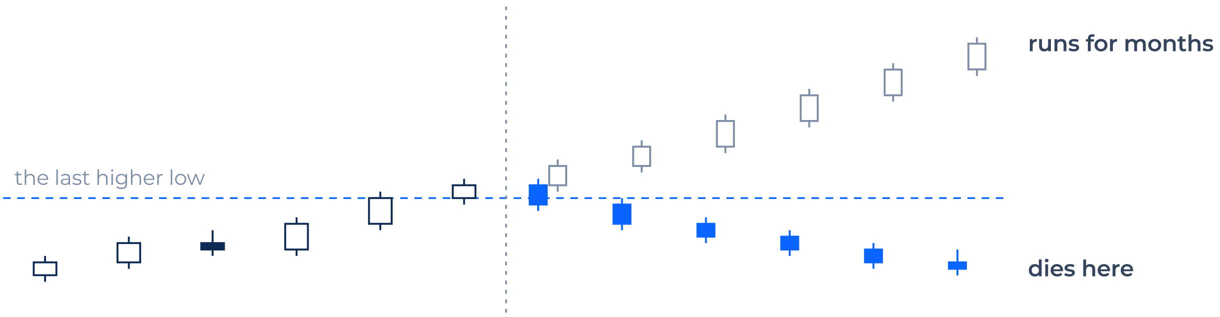

The reason none of this can be judged in the moment. Left of the divider is a run three swings old — identical bars, identical structure, identical everything. To the right it either dies immediately or goes on for a very long time, and standing at the divider there is nothing in the chart that separates the two.

To the left of the divider the two branches are the same bars. The same swings, the same higher lows, the same everything. To the right one of them is finished within a week and the other runs for months, and there is nothing at the divider — no candle, no pattern, no confirmation — that tells you which you are holding.

That is worth stating plainly because a great deal of teaching implicitly denies it. Every promise to identify the big move early is a promise to separate those two branches at the divider, and the branches are not separable there. What can be done is something less thrilling and considerably more useful: build a rule whose behaviour is acceptable in both branches.

A run that has reached	Chance it makes one more swing
1 swing	about 56%
3 swings	about 57%
6 swings	about 58%
10 swings	about 58%

The question people expect to matter, answered. If a run is already long, is it more likely to keep going? Read the right column downward: the chance of one more swing barely moves, which means age tells you almost nothing and the count has to be about size instead.

There is one obvious hope, and it is worth killing carefully because it is the one people reach for. If a run is already long, surely it is more likely to continue? Read the right-hand column: the chance of one more swing barely moves between a run of one and a run of ten. Age is close to uninformative, which means you cannot get an edge from waiting to see whether a trend is established.

And notice what that leaves. If you cannot tell short runs from long ones, and you cannot tell them apart later either, then the only lever remaining is how much you make when you happen to be in a long one and how much you lose when you are not. The whole method has to be built on the size of the outcomes rather than on the accuracy of the reading.

04

Where the money actually is

That conclusion sounds abstract until you look at where the return comes from in a sample of trend trades.

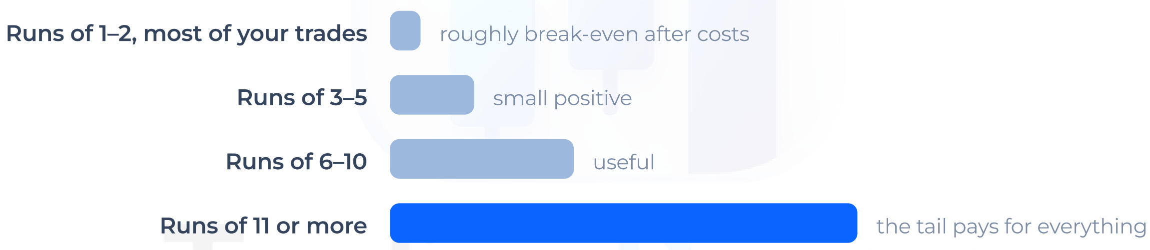

Where the return actually comes from, as a share of everything a trend method makes across a long sample. The top three bars are most of your trades and almost none of your money.

The top three bars are most of your trades and almost none of your money. Nearly two-thirds of everything a trend method earns across a long sample comes from the runs of eleven swings or more — which, by the histogram two sections ago, is fewer than one trade in twelve.

	Wins often	Wins rarely
Large average win	Rare, and usually a short sample or a hidden risk.	Trend following. Uncomfortable, and the shape that fits.
Small average win	Mean reversion. Works until the one that does not come back.	Nothing. This is a losing method with a comfortable feel.

The consequence people find hardest to accept, on one grid. A method built for this distribution is wrong most of the time on purpose, and a method that is right most of the time has usually been built by cutting the tail off.

That distribution forces an uncomfortable shape on the method, and the grid states it. A method fitted to this market is wrong most of the time, deliberately, because being right most of the time requires taking profits at a level the short runs can reach — and doing that amputates the tail that pays for everything.

The bottom-left cell is where most people end up, and it deserves sympathy rather than scorn. Mean reversion has a high win rate, a pleasant equity curve, and a long record of working right up until the move that does not come back. It feels like skill for years. The bottom-right cell is the one to actually fear, and it is where a trend method goes when its exits have been softened enough to feel comfortable.

05

Measured in what

Before going further it is worth settling a question that sounds pedantic and turns out to decide every number in this paper: what a run is counted in.

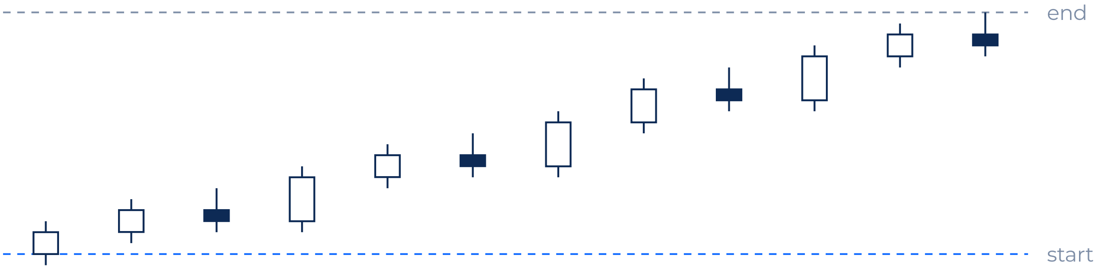

The same run measured two ways. Counted in swings it is four; counted in days it is thirty-one. Neither number is wrong, and they answer different questions — which is why the unit has to be stated before any of this paper's figures mean anything.

That figure is one run measured two ways. Counted in swings it is four. Counted in days it is thirty-one. Both are true, and they answer different questions, so a histogram that does not say which unit it used is not a histogram of anything.

Unit	What it measures	When it is the right one
Swings	how far the structure got	measuring persistence, as here
Days	how long you held	financing, overnight risk, patience
Points	how far price moved	sizing and expectancy
Bars	nothing stable	never — it changes with the timeframe

Three units, three different questions. The middle column is what this paper counts in, because it is the only one that does not change when the market's speed changes — and speed changes constantly.

This paper counts in swings, and the reason is in the third column. Days and points both move when the market's speed moves: a quiet month and a violent month produce very different day-counts for structurally identical runs. Swings hold still, because a swing is defined by shape rather than by pace.

That said, days are not a worse unit — they are the right one for a different question. If you are paying financing, or holding through weekends, or simply working out whether you can stand to be in something for two months, days is what you need. The mistake is only ever to mix them: to count persistence in swings and then form expectations about patience in days without ever converting between the two.

06

The entry barely matters

One consequence of the distribution deserves its own section, because it contradicts how almost everybody spends their time.

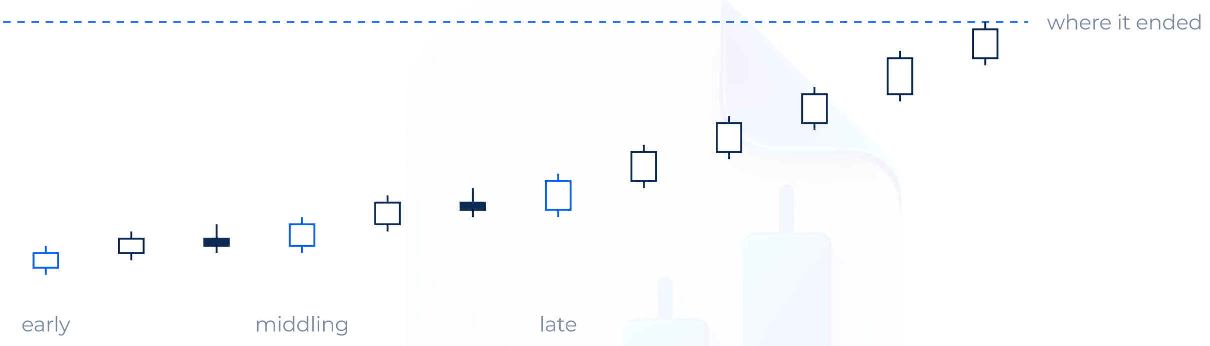

Three entries into the same run, taken by three different people with three different methods, spread across the first third of it. On a run of this length the difference between the best and worst of them is small compared with the part that comes afterwards.

Three entries into the same run, taken across its first third by three people with three different methods. One of them read it early, one waited for confirmation, one was late. On a run of this length, the distance between the best and the worst of them is a small fraction of what came afterwards.

What those three entries were actually worth across a hundred runs, in units of risk. The spread between the best and the worst entry is real and it is small — about one unit separates a careful reader from a careless one.

Across a hundred runs the spread is real and it is modest: about one unit of risk separates a careful reader from a careless one. That is worth having, and it is not nothing. But hold it in mind while you look at the next figure.

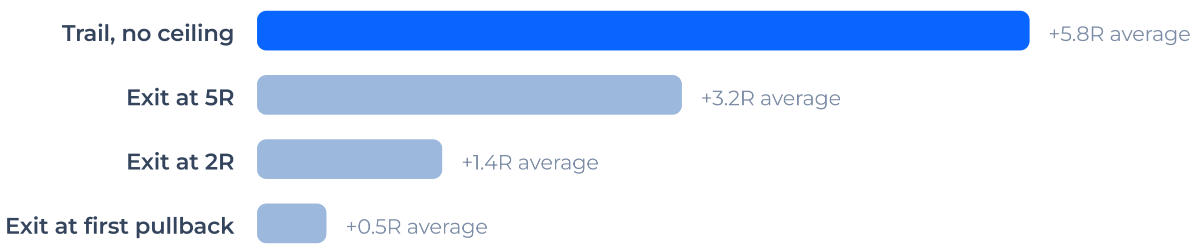

The same hundred runs, same entries, only the exit rule changed. Set this against the previous figure: the exit is worth several times what the entry is worth, and it is the half almost nobody spends their evenings on.

Same runs, same entries, only the exit rule changed — and the spread is several times larger. The difference between trailing without a ceiling and taking profit at the first pullback dwarfs every entry improvement available, and it does so on the same trades.

Set those two figures side by side and you have the most actionable thing in this paper. Entry technique is where nearly all study, all indicators and all screen time go; exit policy is where the returns are decided. That is not an argument for entering carelessly. It is an argument that if your evenings are finite, they are currently being spent on the smaller half of the problem.

07

Why the old advice is arithmetic

Cut your losses short and let your winners run is the other sentence everybody has heard and nobody derives. It is not a piece of temperament advice. It falls directly out of the shape in this paper.

1. Most runs end early, so most positions must be closed cheaply
2. A few runs are enormous, so those must not be closed early
3. You cannot tell which is which at the time
4. So the rule must be asymmetric: small fixed loss, unbounded gain
5. Any rule that caps the gain removes the only part that pays

Why cutting losses short is arithmetic rather than folklore, in five steps. Nothing here is about temperament: given the distribution, this is the only exit policy that keeps the tail attached to your account.

Read the five steps as a chain rather than as a list. Because most runs die early, most positions must be closeable cheaply. Because a few are enormous, those must not be closed early. Because the two are indistinguishable at the time, the rule cannot depend on telling them apart. What is left is a rule with a small fixed loss and no ceiling on the gain, and that is the only shape that fits.

	Fixed 2R target	Trail, no target
Win rate	51%	34%
Average win	2.0R	5.8R
Average loss	1.0R	1.0R
Longest run captured	2.0R	31R
Total over 100 runs	+53R	+131R

The same point as a ledger, on one hundred runs with a fixed one-unit risk. The fixed target looks better on every measure a person naturally checks, and loses.

The ledger makes it concrete, and it is worth reading the way a person naturally would: win rate first. On that measure the fixed target wins comfortably — fifty-one per cent against thirty-four. It also wins on how often the trader felt correct, which is not in the table but is doing a lot of work in real decisions.

And it loses by a wide margin, because the fourth row is where the sample's whole return lived and the fixed target was never able to reach it. This is the most expensive mistake available in trend trading, and it is made by people being careful rather than by people being reckless.

08

Size, and reaching the tail

Position size is normally treated as a separate subject from strategy. In a tail-driven method it is not separate at all, and the reason falls straight out of the histogram.

Because	most runs lose, and they arrive in streaks
Therefore	the size must survive a long run of small losses
Which means	a fixed fraction, never raised after a loss
And never	raised after a win either, which is the same error mirrored
Except	adding to a position already running, from its own profit

Why size is part of this argument and not a separate topic. Each rung follows from the distribution rather than from prudence: if the tail pays for everything, then surviving long enough to hold one is the whole job.

Read the rungs as a chain. Most runs lose. Losses arrive in streaks, because nothing about the distribution spaces them out politely. So the size has to be small enough that a long streak leaves the account intact — not because caution is a virtue, but because the return lives in a run that has not happened yet and you must still be present when it does.

Risk per trade	Exit at 2R	Exit at 5R	Trail, no ceiling
0.5% per trade	+26R	+58R	+124R
1% per trade	+53R	+112R	+131R
5% per trade	+61R	+40R	account gone in the streak

The same three exit policies from earlier, now run at three position sizes, in total return over the sample. Read across the bottom row: the aggressive size does not amplify the good method, it removes it, because the account does not survive to reach the tail.

The bottom row is the one to sit with. At five per cent a trade, the best exit policy in the table becomes the worst outcome in the table, because the account does not survive the streak that precedes the tail. The aggressive size did not amplify the good method. It deleted it.

The last rung of the ladder is the one genuine exception, and it is worth stating precisely because it is so easily abused. Adding to a position that is already running, out of profit the position itself has produced, does not raise the risk on the original decision — it is a second, separate bet on a run that has now demonstrated something. Adding to a position that is losing, or raising base size because the last trade won, is the same error wearing two different hats.

09

What it feels like to be right

Numbers on a page understate the difficulty, so it is worth drawing the experience the distribution guarantees.

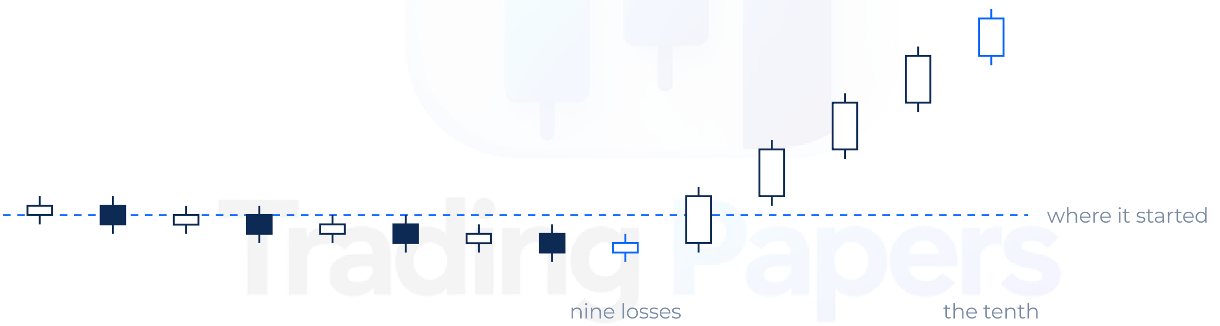

What that better method feels like from inside. Nine consecutive losses, none of them a mistake, before the run that pays for all of them. The account line is not drawn here because it is not the difficulty — the difficulty is that these nine bars arrive one at a time, with days between them.

Nine losses in a row, none of them an error, none of them avoidable, each one a correctly taken trade under a rule that is working. Then the tenth pays for all of them and more. Over a long sample this sequence is not bad luck — it is the ordinary consequence of a method that wins a third of the time.

The thing the drawing cannot convey is the spacing. Those nine losses do not arrive in an afternoon; they arrive over weeks, with quiet days between them in which there is nothing to do but wonder whether the method has stopped working. Every one of those gaps is an invitation to adjust something, and adjusting something is how the tail gets amputated.

This is the honest reason most people cannot trade trends, and it has nothing to do with intelligence or discipline in the abstract. It is that a correct method delivers a long, unbroken experience of being wrong, and the only thing that gets a person through it is having counted the distribution beforehand and knowing that nine is normal.

10

Two markets, two shapes

Everything so far has described a shape. The shape is common; the numbers filling it are not, and the difference between two instruments can be large enough to decide whether this method belongs on one of them at all.

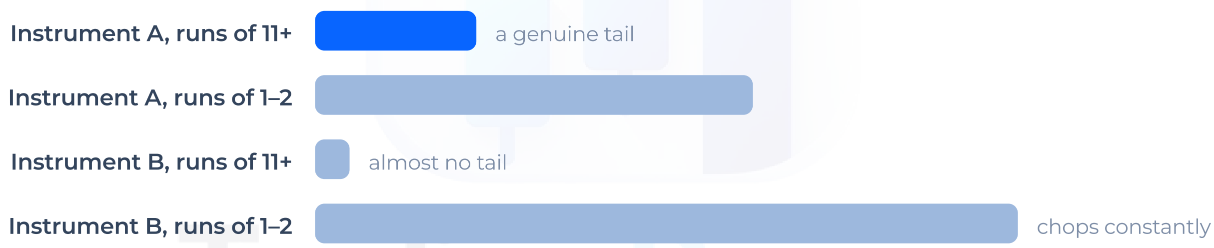

The same count run on two instruments over the same period. The shape is recognisably the same family — most runs short, a few long — but the tail is much fatter on one than the other, and that difference decides whether trend following is worth attempting there at all.

Two instruments, same period, same swing rule. Both produce the familiar family — many short runs, few long ones — but one has a genuine tail and the other barely has one. Instrument B chops: most of its runs die at the first swing and almost nothing survives to eleven.

	Instrument A	Instrument B
Median run length	3 swings	1 swing
Top decile share of return	64%	22%
Longest losing streak	11	19
Suits a trend method?	yes	no — count says mean reversion

The two instruments side by side on the numbers that matter. Instrument B is not a worse market; it is a different one, and applying a trend method to it is a category error rather than bad luck.

The bottom row is a judgement the count makes for you rather than one you make about the count. On Instrument B a trend method is not unlucky, it is misapplied: there is no tail for the losses to be paying for, so the arithmetic that justified the whole approach simply does not hold there.

This is the most common quiet failure in trend trading, and it looks nothing like a mistake while it is happening. The rule is followed correctly, the losses are taken properly, the discipline is real — and the instrument was never going to produce the run that made it worthwhile. Counting first is what separates a method that is being executed badly from one that is being executed faithfully in the wrong place.

11

When the shape changes

A distribution is a description of a period, and periods end. The practical question is how to notice that without inventing a regime change every time you lose four in a row.

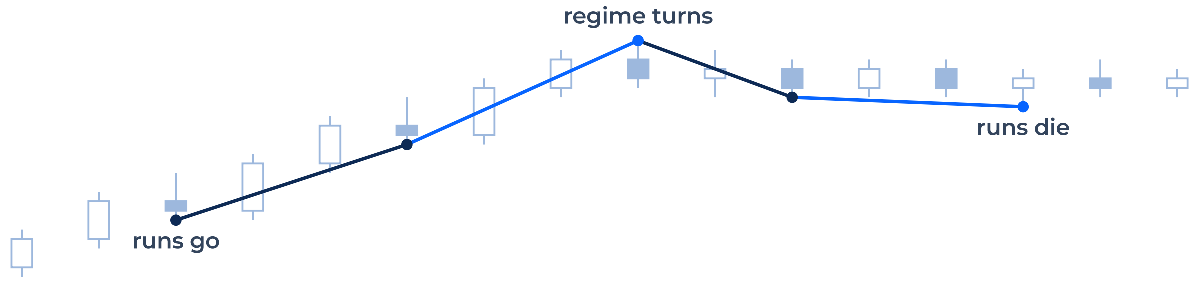

One instrument, one continuous stretch, and two distributions inside it. The left half produces runs that keep going; the right half produces runs that die at the first swing. Nothing about the method changed between them.

One instrument, one continuous stretch, two behaviours inside it. In the left half runs extend; in the right half they die at the first swing and price simply oscillates. Nothing about the swing rule changed between the two — the market did.

What you observe	What it means	What to do
Four losses in a row	ordinary; within any normal streak	nothing
Losing streak beyond the sample's worst	possible, still weak evidence	nothing yet, keep counting
Median run length halves over 50 runs	the distribution has moved	recount, then resize
Top decile share collapses	the tail has gone	stop; this method needs a tail

How to notice a shift without inventing one every time you lose. The middle column is what actually distinguishes a regime change from an ordinary drawdown, and the right column is what to do — which in three cases out of four is nothing.

The table is deliberately unexciting, because the right answer in three cases out of four is to do nothing. Four losses in a row is not evidence of anything; a streak longer than any in your sample is weak evidence at best. What actually constitutes evidence is a change in the distribution itself, measured over enough runs to be a measurement rather than a mood.

The last row is the only one that says stop, and it says it for a specific reason rather than out of caution. If the top decile's share of return has collapsed, the tail has gone — and a method whose entire justification was the tail has lost its justification, not merely its recent form. Everything else in the table is a request for patience; that row is a request for honesty.

12

Counting it yourself

Every number in this paper is one you can produce for your own instrument, and you should, because the shape is common but the values are not universal.

1. Choose one instrument and one timeframe, and do not change either
2. Fix the swing rule in writing, exactly as the structure paper sets out
3. Mark every run and record its length in swings, not in days
4. Build the histogram, then read the median and the top decile
5. Compare the top decile's share of return against everything else

How to produce these numbers for your own instrument in about a week of evenings. The second step is the one that must be fixed in writing first, because every number after it moves when the swing rule moves.

The second step decides everything after it. Run length is measured in swings, and swings are produced by a swing rule, so a different rule gives a different histogram from the same data. Fix it in writing before you start and never adjust it mid-count; the structure paper in this library explains why that instruction is less obvious than it sounds.

The number	What it decides
Median run length	how long a hold you must be able to tolerate
Share of runs that die immediately	how many small losses you must budget for
Top decile's share of total return	whether an exit rule may ever cap a winner
Longest losing streak in the sample	whether you can survive the method you chose

The four numbers worth having at the end, and the decision each one drives. The last row is the one that tells you whether trend following suits you rather than whether it works.

The four numbers each drive a decision. The median tells you what a normal hold feels like. The share of runs that die immediately tells you how many small losses to budget for. The top decile's share of return tells you whether you are ever allowed to cap a winner — and if that share is above about half, the answer is no, permanently.

The last row is different in kind, because it is not about the market. The longest losing streak in your sample is a fact about what you will have to sit through, and it is the number to check against yourself rather than against the method. A strategy you cannot hold through its normal drawdown is not a strategy you own.

13

What the method promises

It is worth writing the bargain out explicitly, because every part of it is knowable in advance and almost nobody reads it until they are already inside.

	It can give you	It cannot give you
About the outcome	A positive expectancy carried by rare large winners.	A high win rate, or a smooth line, ever.
About the experience	Long stretches of being wrong that are entirely normal.	Any way of knowing, at the time, that they are normal.

What a trend method can and cannot promise, as two honest sentences in each direction. The bottom-left cell is the whole bargain, and a person who has not accepted it will abandon the method during its first ordinary streak.

The two cells on the left are what a trend method genuinely offers: a positive expectancy carried by rare large winners, and long stretches of being wrong which are entirely normal. Both are real. Only the first of them is what anyone signs up for.

Term	What you are agreeing to
Win rate	roughly a third; two out of three trades will lose
Streaks	nine or more consecutive losses, repeatedly
Where returns come from	under a tenth of trades
What you may never do	cap a winner, or raise size after a loss
What decides success	whether you still hold the rule at loss number nine

The contract, written out. Read the right column before adopting the method rather than during its ninth loss, because that is the moment the terms start feeling negotiable.

Read the right-hand column as terms rather than as warnings. Two out of three trades will lose. Nine consecutive losses will happen more than once. Under a tenth of your trades will produce your returns, which means the great majority of your working life inside this method will feel like failure while being exactly correct.

The last row is the only one that decides anything, and it is not about the market at all. Whether you still hold the rule at loss number nine is the entire question, and it is settled long before loss number nine arrives — by whether you counted the distribution first, or adopted the method because someone showed you the tenth trade.

14

What this does not settle

Three limits, and then the short version.

The limit	What it means here
Run lengths change with regime	the distribution is not a constant; recount periodically
Swing rule decides the count	a different rule gives different lengths, so state it
Past shape need not persist	the count bounds your expectations, it does not guarantee them

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

The first is the most practical. Run-length distributions are not constants; they widen in trending regimes and collapse in quiet ones, which is exactly the subject of the next paper in this collection. A histogram built on two years is a description of those two years, and it should be rebuilt rather than trusted indefinitely.

The second is the dependency this paper has leaned on throughout. Every length here is measured in swings, and swings come from a rule you chose. Change the rule and every number moves — not because the market changed but because you did. That is not a flaw in the count, it is a requirement to state the rule alongside the result.

The third is the ordinary caution about any historical count. A shape that held for three years is evidence, not a law, and the honest use of it is to bound your expectations rather than to guarantee your returns. What the count genuinely buys you is knowing in advance which experiences are normal — and that is worth more than it sounds, because it is what stops you abandoning a working method during its ninth loss.

Term	As used in this paper
Run	a sequence of swings in one direction, by a stated rule.
Run length	the number of swings in that sequence, not its duration in days.
Median	the middle value: half of all runs are shorter than it.
Top decile	the longest tenth of runs in a sample.
R	one unit of planned risk: the distance from entry to stop.
Hazard	the chance a run continues, given that it has already reached some length.

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. Trends persist slightly more often than chance, but the length of a trend is a heavily skewed distribution rather than a typical value, so the average describes a run the market rarely produces. You cannot tell a short run from a long one while you are inside it, and a run's age tells you almost nothing about its future, so accuracy is not the lever — size is. Which means the whole method reduces to one obligation: keep the losses small and fixed, and never put a ceiling on the winner, because a small fraction of runs carries the entire return and a capped exit is a decision to give that fraction away.

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