

TRADING KNOWLEDGE LIBRARY

The Micro Entry

Price Action



Trading Papers

The Micro Entry

Small structure, distant targets, and the arithmetic that connects them

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	Ten to one is a stop problem	4
02	Structure below the timeframe you decide on	5
03	How small a stop is allowed to be	6
04	The arithmetic, before the chart	7
05	Staged entries	8
06	What a partial close actually does	10
07	Sizing when the stop is small	11
08	Costs decide this method	12
09	The fill you actually get	13
10	Where it works and where it cannot	14
11	Holding is the hard part	15
12	One trade, end to end	16
13	What has to be on paper first	18
14	Testing it honestly	20


01

Ten to one is a stop problem

A trade that returns ten times what it risked is almost never a trade that predicted ten times further than usual. It is a trade whose invalidation sat unusually close to its entry. The target was ordinary; the risk was small.

This is worth stating plainly because it reverses the order most people work in. The common approach is to find a direction, enter, and then look for somewhere to put the stop — which produces a stop placed where there is room, and a reward-to-risk ratio that is whatever happens to fall out. The asymmetric approach starts from the other end: find the price at which the idea is demonstrably wrong, and then look for an entry close enough to it that the distance is small.

Everything else in this paper follows from that inversion. Structure below your decision timeframe matters because it is where those close invalidations live. Staged entries matter because the first attempt at a close entry often fails. Partial closes matter because an asymmetric trade spends most of its life in a state you will want to escape.



The hit rate each payoff multiple must beat merely to break even, before costs. The curve flattens fast: most of what a large multiple buys you is bought by the time you reach 5R, and everything past it is bought with a longer hold.

One caution before the arithmetic: a small stop is not a safer stop. It is a stop that is hit more often. What it buys is not survival on the individual trade but the size of the payoff when the trade works — and those two are in direct tension for the rest of this paper.

02

Structure below the timeframe you decide on

Two timeframes are doing different jobs. The decision timeframe is where the idea comes from: the level being approached, the trend being followed, the range being left. The execution timeframe is where you find the entry, and it is lower — often much lower.

Micro structure is simply structure on that execution timeframe: the small sequence of higher lows inside a pullback, the narrow channel a retracement carves out, the point where that channel stops holding. None of it is visible on the decision timeframe, where the whole episode is one candle with a wick.

	Decision timeframe	Execution timeframe
Answers	Is there an idea here at all, and where is it going?	Where exactly is this attempt wrong, and by how little?
Never answers	Where the tight invalidation is — it cannot see one.	Whether the idea is any good — it sees only the last few minutes.

The two timeframes and the different questions they answer. Most of the damage done with this method comes from asking one of them the other's question — taking a signal from the execution chart, or a stop from the decision chart.

What it gives you is a specific, recent, objective price at which the small move you are trading has failed — and that price is usually far closer than the last swing on the decision timeframe. Trading the same idea with an invalidation five points away instead of fifty does not make the idea better. It makes it four or five times more profitable when it is right, and it makes it wrong more often.


The cost of going lower

Three costs come with the lower timeframe and none of them is optional. Noise rises: a five-point stop is inside the ordinary jitter of many instruments during active hours. Costs rise relative to the trade: a spread that is trivial against a fifty-point stop is a fifth of a ten-point one. And the number of decisions rises, which is a tax on attention rather than on the account, and is the reason this approach exhausts people who try to run it all day.

03

How small a stop is allowed to be

There is a floor below which a stop stops measuring anything, and it is not a matter of preference. Every instrument has an ordinary amount of back-and-forth at every hour of the day — the distance price travels while nothing is happening. A stop placed inside that distance is not tight. It is a coin flip with extra steps.



Schematic, for one liquid FX pair during London hours. The first bar is the ordinary five-minute jitter; the rest are candidate stop distances against it. A stop inside the jitter is not tight — it is random.

The floor is measurable and worth measuring once: take the average five-minute range of your instrument during the hours you trade it, and treat that as the minimum meaningful

invalidation. A stop at half of it will be hit close to at random, regardless of whether the idea was any good, and the record will show a method with an excellent reward-to-risk and a hit rate too low to survive.

This is the correction to the naive version of everything in section 01. Smaller stops do raise the multiple, but only until they fall inside the noise, at which point the multiple keeps rising and the expectancy collapses. The optimum sits just outside the ordinary jitter and no tighter — which, for most liquid instruments during active hours, is a stop of ten to twenty points rather than three.

A quick diagnostic: if more than a fifth of your stops are hit within two minutes of entry and price then continues in your direction, the stop is inside the noise. That pattern has one cause and one fix.

04

The arithmetic, before the chart

Expectancy for a fixed-risk approach is the only formula that matters here: the average outcome per trade is the win rate times the average win, minus the loss rate times the average loss. Stated in units of risk, where a loss is 1R:

Leak	Where it happens	Effect on the multiple
Spread	at entry, every attempt	-0.2R to -0.4R
Slippage on the stop	at exit, on the losers	loss becomes 1.1R
Partial close at 2R	management	10R becomes 6R
Early manual exit	the trader, at 4R	the tail is cut off

Where a nominal 10R method arrives at an effective 6R. Each leak is small and each is certain; taken together they roughly double the hit rate the method actually needs.

- At 5R per win, breakeven is a hit rate of 1 in 6, about 17%.
- At 10R per win, breakeven is 1 in 11, about 9%.
- At 3R per win, breakeven is 25% — which is why 3R systems feel comfortable and 10R systems feel like losing.

Two things follow immediately. The first is that an asymmetric approach can be wrong most of the time and still work, which is the entire attraction. The second is less often said: those breakeven rates assume the average win really is 10R and the average loss really is 1R. Both assumptions leak. Slippage makes losses slightly larger than 1R. Partial closes make wins substantially smaller than 10R. A system designed at 10R and executed at an effective 6R needs nearly twice the hit rate it was built for.

Before trading an asymmetric method, write down the hit rate it needs and the hit rate you have actually achieved on the last hundred attempts. If the second is not clearly above the first, the method does not have an edge problem — it has an arithmetic problem, and no amount of chart reading will fix it.

05


Staged entries



Schematic. The same idea, attempted three times as the move develops. Each attempt risks the distance to its own invalidation, and each later one is closer to the level that proves the idea wrong — which is why the third attempt can reach the same target at a fraction of the first one's risk.

A close entry fails often by construction. The response that turns this from a flaw into a feature is to treat one idea as several attempts rather than one trade.

The structure of it is simple. The first attempt is taken at the first place the execution timeframe offers an invalidation. If it is stopped, the idea on the decision timeframe has not necessarily changed — only the small structure that housed the entry. A second attempt is taken when new micro structure forms, with a new invalidation. And so on, usually to a limit of three.



The same idea attempted three times, with the numbers each attempt actually carries. Note that the risk per attempt never changes, so three failures cost 3R — but the achievable multiple grows, because the invalidation gets closer.

Two properties make this work. Each attempt risks the same fixed fraction of the account, so three failures cost 3R rather than an escalating amount. And each later attempt tends to sit closer to the level that would genuinely disprove the idea, so its stop is smaller and its reward-to-risk larger — the third attempt at a move can reach the same target at a fraction of the first attempt's risk.

This is not averaging down, and the difference is not semantic. Averaging down adds to a position whose invalidation has already been passed, increasing exposure as the evidence worsens. Staged entry closes each attempt at its own invalidation and re-enters only when new structure supplies a new one. The first has no defined loss; the second has three defined losses at most.

Where it goes wrong

The limit has to be written down in advance, because in the moment there is always one more level. A trader who allows a fourth, fifth and sixth attempt has stopped taking staged

entries and started refusing to be wrong — and the account cannot tell the difference between conviction and stubbornness, only between 3R and 9R.

06

What a partial close actually does

Closing half the position at a modest multiple and moving the stop to entry is the most common management rule in asymmetric trading, and it is almost always adopted for a reason that has nothing to do with expectancy: it converts an uncomfortable position into a comfortable one.

It is worth being precise about the trade it makes. Closing half at 2R and running the rest to 10R yields 6R on a winner instead of 10R — a 40% reduction in the payoff that justified the whole approach. In exchange, the distribution narrows: fewer full losses, more small wins, and a much smaller chance of watching a large open profit return to zero.

Rule	Result on a winner	What it needs to be right
Run the whole position	10.0R	the tail actually arrives often
Half at 2R, rest to target	6.0R	most 2R trades fail to continue
Half at 5R, rest to target	7.5R	a middling assumption either way
Stop to entry at 2R	10.0R or 0R	noise is small next to the stop

The same winning trade under four management rules. There is no correct row; there is a correct row for your own record, and it is decided by the number in the last column.

Whether that is a good trade depends entirely on a number most traders never measure: how often a position that reaches 2R goes on to reach the target. If it is frequent, the partial close is expensive. If it is rare, the partial close is the reason the method is profitable at all. The measurement takes an afternoon on your own record and settles an argument that otherwise runs for years.

Moving the stop to entry deserves the same scrutiny. It removes the loss and it also removes

the trade — a position stopped at breakeven on ordinary noise cannot participate in the move it was right about. On a small-stop approach, where ordinary noise is a large fraction of the original risk, this happens constantly.

07

Sizing when the stop is small

Fixed fractional risk means the position is whatever makes the stop cost one unit of risk. Nothing about that changes here — but the consequence of it is more uncomfortable on a small stop than most people expect the first time they see it.

Stop	Risk budget	Position	What it means in practice
50 points	1% of account	1 unit	ordinary, forgiving
20 points	1% of account	2.5 units	still comfortable
10 points	1% of account	5 units	a normal wick is now the whole loss
5 points	1% of account	10 units	size looks alarming; risk has not changed

Constant risk, varying stop. The position grows as the stop tightens, which is arithmetic everyone knows and a trap almost everyone falls into: the tightest stop produces the largest position at the moment the stop is least reliable.

Read the table honestly. The last row risks exactly as much as the first; the account cannot tell them apart. What differs is how the position looks on the screen and how it behaves in the seconds after entry, and both of those are handled by the trader rather than by the arithmetic.

Two real hazards hide in that row, and neither is psychological. The first is slippage: a large position filled against thin liquidity slips further, and a slipped stop on a five-point risk is a materially larger loss than 1R. The second is that many instruments have a minimum position increment, so the theoretical size is rounded — usually upward, quietly turning a one per cent risk into a one and a half per cent one.

The rule that survives both is to size from the stop, cap the position independently of the arithmetic, and take the smaller of the two. When the cap binds, the trade is simply smaller than the formula wanted; that is a cost worth paying to keep the loss distribution the shape the method assumed.

08

Costs decide this method

For most approaches, transaction costs are a drag. For this one, they are a design constraint, and they are the reason it works on some instruments and cannot work on others.

The arithmetic is unforgiving. If the spread is two points and the stop is ten, then every trade begins 20% of the way to its own invalidation, and the effective reward-to-risk on a nominal 10R target is closer to 8R. If the spread is two points and the stop is five, a quarter of the risk is gone before price moves. Three failed attempts have then paid six points of spread — more than half of a single attempt's risk — for nothing.



How much of the risk the spread has already taken at the moment of entry, for a two-point spread. This is the single number that decides whether the method is viable on an instrument, and it is why the smallest stop is not the best stop.

Three practical consequences. Trade the instrument whose spread is smallest relative to its typical micro structure, not the one whose chart looks cleanest. Prefer the active hours, when spreads are tightest, even though ranges are wider. And measure your own fills: the difference between the theoretical entry and the achieved one, averaged over fifty trades, is a number that either confirms this approach is viable for you or quietly explains why it is not.

09

The fill you actually get

Every number in this paper assumes you were filled at the price you meant. On a fifty-point stop that assumption is harmless. On a ten-point one it is the difference between a method and a story, because the gap between the intended price and the achieved one is measured in the same units as the risk.

Order	What it guarantees	What it costs	Where it fails
Backtest fill	the price on the chart	nothing	it does not exist
Limit at the level	the price or nothing	the trades you miss	fills only when wrong
Market on the signal	a fill	spread plus slippage	worst at the open
Stop entry above	a fill, on confirmation	a wider entry	widest stop of the three

The same order, filled four ways. Only the first row is what a backtest assumes, and the gap between it and the others is measured in the same units as the stop — which on this method means the gap is a large fraction of the risk.

The second row is where most people go, and it deserves the warning in its last column. A limit order at the level fills only when price reaches it — which is to say, it fills reliably on the attempts that carry on through, and misses the ones that turn at the level. The fills you get are therefore biased towards the losers, and a record built on limit entries flatters the method's average entry while quietly worsening its hit rate.

The remedy is not a cleverer order type. It is to measure: log the intended price and the achieved price on fifty consecutive attempts, take the average difference, and add it to the spread. If the sum is a fifth of your stop, the method is paying a fifth of its risk before price has moved, and the honest response is a wider stop and a smaller nominal multiple rather than a better order.

10

Where it works and where it cannot

This method is not universal, and the boundary is not a matter of taste. It needs an instrument whose spread is small relative to the invalidation its micro structure supplies, during the hours you are able to trade it. That is one ratio, and it either clears or it does not.

Instrument	Typical spread	Viable stop	Spread as share of risk
Major FX pair, London hours	0.6 pips	8–15 pips	4–8% — workable
Major FX pair, Asian hours	1.4 pips	5–8 pips	18–28% — marginal
Index future, cash open	1 tick	8–20 ticks	5–12% — workable
Gold	wide, variable	large	unstable — measure before assuming
Small-cap equity	wide	large	the method does not apply

Where the method has room and where it does not. The question is never whether the chart looks clean — it is the ratio in the last column, which can be computed in an afternoon for any instrument you are considering.

The Asian-hours row is the one that catches people, because the instrument is identical and only the hour has changed. The spread widens and the ranges compress at the same time, so the two halves of the ratio move against each other — a method that clears comfortably at nine in the morning in London can be structurally unprofitable at three.

	It works when	It fails when
Market	Range is expanding and structure forms cleanly.	Range is compressing: every tight stop sits inside the noise.
You	You can watch the execution chart and take three attempts calmly.	You trade it around a job, and attempt two is missed every time.

The conditions the method needs, against the conditions in which it quietly stops working. Nothing in the right column announces itself; each has to be checked, which is why the checking belongs in the routine rather than the moment.

The lower row of that matrix is the one worth taking seriously. Staged entries assume you are present for all three attempts; a trader who reliably catches the first and misses the second has not implemented this method at all, but a worse version of it that takes only the widest stop of the three.

11

Holding is the hard part

An approach that is wrong two thirds of the time and pays for it in one large win produces a psychological experience that most people have not designed for, and it is worth describing before it is experienced.

Losses cluster. With a 25% hit rate, runs of six or seven consecutive losses are not bad luck — they are the expected behaviour of the distribution, and they arrive regularly. Meanwhile the wins, which carry the entire result, are rare enough that missing two of them through an early exit can turn a profitable quarter into a flat one.

Hit rate	Longest expected losing run	How often five or more in a row
50%	about 7	several times
33%	about 11	very often
25%	about 15	constantly
15%	about 25	the normal state of affairs

What a hit rate actually feels like over two hundred trades. These are not bad patches to be endured occasionally — they are the ordinary output of the distribution, and a method abandoned during one was abandoned on schedule.

The practical defences are unglamorous and they work. Fix position size as a constant fraction of the account so that no single attempt can matter. Write the target, the partial rule and the maximum number of attempts before entering, when nothing is at stake. Record every attempt including the ones that were stopped, because a record of only the memorable trades will convince you the method is better and rarer than it is. And judge the method over a hundred trades rather than ten, because ten is inside the noise of a distribution this skewed.

12

One trade, end to end

Everything above is easier to hold in mind as a single sequence. This is one idea followed from the daily chart to the exit, with the rule that produced each decision named beside it.

Step	What was seen	What was done	Why
Idea	daily trend up, price at a weekly level	watch only	decision timeframe
Attempt 1	five-minute higher low forms	long, stop 18 below	first micro invalidation
Stopped	the low fails	out, -1R	the attempt was wrong, not the idea
Attempt 2	new higher low, tighter	long, stop 11 below	new structure, new stop
Managed	reaches 2R	held, no partial	record says 2R continues often
Exit	target at the prior high	closed, +16R	written before entry

One trade, logged as it happened, with the decision at each step and the rule that produced it. Every row was decided before the trade began — which is the only property of this table worth copying.

Three things in that table are the point. The first attempt lost and was not a mistake — it was the price of the tighter invalidation, and the second attempt only existed because the first one closed cleanly. The decision at 2R was made by consulting a record rather than a feeling. And the exit was at a price written down before entry, which is the only reason a sixteen-multiple was collected rather than half of it.

Trading Papers

Column	Why it has to be there
Attempt number	tests whether later attempts really are better
Stop distance	tests whether micro structure tightens anything
Planned multiple	separates a bad trade from a bad plan
Achieved multiple	the only honest measure of the method
Reached 2R?	settles the partial-close argument
Spread at entry	turns cost from an opinion into a number

The minimum log. Without the attempt number and the planned multiple, none of the questions this paper raises can be answered later — and answering them later is the whole point of keeping a record.

The log is what makes the next hundred trades cheaper than the last hundred. Without the attempt number, the claim that later attempts are better is an article of faith. Without the spread column, the cost argument cannot be applied to your own broker. The columns are chosen so that every disputed question in this paper is settled by a spreadsheet rather than by an opinion.

13


What has to be on paper first

Every argument in this paper resolves into a small number of written decisions. They are worth collecting in one place, because the difference between a method and an inclination is whether these six things existed in writing before the first trade.

Field	Where the number comes from	What happens without it
Instrument and hours	the spread-to-stop ratio	the method is tested on the wrong market
Minimum stop	measured five-minute range	stops land inside the noise
Risk per attempt	a fixed fraction of the account	size drifts upward after losses
Attempt limit	written, and it is three	a fourth attempt becomes a fifth
Target	structure, not hope	the exit is decided while in profit
Partial rule	your own record at 2R	the rule changes with the mood

The written plan, field by field. Each row is a number that has to exist before the first trade, and each has a section of this paper behind it — which is why a plan with a blank in it is a plan whose blank will be filled in at the worst possible moment.

The last two rows are the ones people leave blank, and they are the two that get decided under pressure. A target chosen while a position is in profit is not a target; it is relief. A partial rule invented at 2R is not a rule; it is a way of ending discomfort, and it will differ from the one used yesterday in a direction you cannot predict and cannot measure.



The order in which the plan is finished. Nothing here is done at the screen; all of it is done on a day when no position is open and nothing needs defending.

None of this is done during a session. The measurements need a spreadsheet and an hour, the decisions need a day when no position is open, and both need to be finished before they are needed rather than while they are needed. That is the entire mechanism: not discipline in the moment, but decisions made where discipline is cheap.

It is also what makes the honest test in the next section possible at all. A method with six written fields can be compared against the boring alternative, because both were specified in advance. A method that was improvised cannot be compared against anything, and the record it leaves behind measures the mood of its author rather than the behaviour of a market.

14

Testing it honestly

The method makes three separate claims, and they can be tested separately rather than as a bundle.

- That micro structure supplies a closer invalidation than the decision timeframe — measurable directly: the average stop distance under each rule, over the same set of ideas.
- That later attempts have better reward-to-risk than earlier ones — measurable by logging the attempt number alongside the achieved multiple.
- That the whole thing beats simply taking the decision-timeframe entry with a wider stop and a smaller target — the only comparison that matters, and the one almost nobody runs.



The comparison that decides whether the method deserves the attention it costs. Almost nobody runs the third step, which is why almost nobody knows whether the micro entry beat the boring alternative.

That last comparison is the honest test. An asymmetric method should be measured against the boring alternative it replaced, on the same instrument and the same period, after costs.

If it does not win that comparison, its appeal was the story rather than the arithmetic.

One last observation, offered without much comfort. The methods in this paper are mechanical, and the parts of them that fail are not. Every rule here exists because a specific decision becomes very difficult at a specific moment — the second attempt after the first one lost, the hold through 3R, the log entry on a day that went badly — and writing the rule down in advance is the only device that has ever reliably worked.

What it looks like	What it is	The rule that prevents it
A fourth and fifth attempt	refusing to be wrong	the attempt limit, written before the day
Tightening the stop further each week	chasing the multiple into the noise	a floor set from measured range
Taking profit at 3R “just in case”	collecting the frequent half of the distribution	the target, written before entry
Skipping the log on losing days	building a record of only the good trades	log the attempt, not the outcome
Judging it after fifteen trades	reading noise as evidence	one hundred trades, then decide

The five ways this method is most often abandoned, and what each one actually is. None of them is a chart-reading failure; every one is a rule that was never written down before it was needed.

If the method is going to be abandoned, it will be abandoned during a losing run of the



The whole paper as a sequence, in the order the decisions are actually made. Everything before this figure is the argument for one of these five steps; nothing in the method happens outside them.

ordinary length shown in section 09, for reasons that will feel like judgement at the time. Knowing the length of that run in advance does not make it pleasant. It does make it survivable, which is the whole of what this paper can offer.

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