

TRADING KNOWLEDGE LIBRARY

The Spike and the Second Leg

Trading Strategies



Trading Papers

The Spike and the Second Leg

Trading the pullback after an impulsive move, and the price of a fixed target

Published by **Trading Papers**

Series **The trading knowledge library**

Edition **First edition · 2026-08-26**

Format **A4 · digital**

Educational content only. Nothing in this document is financial advice, no part of it is a recommendation to buy or sell anything, and it knows nothing about your circumstances. Trading carries risk, including the loss of the amount invested.

© 2026 Trading Papers. All rights reserved.



Trading Papers

Contents

01	What an impulse leaves behind	4
02	Why the pullback arrives in two legs	5
03	Defining the spike so it can be counted	6
04	The gap is the evidence, not the decoration	8
05	Where the second leg ends	9
06	The stop belongs to the origin	10
07	When the stop is not the same size twice	11
08	The one-to-one target, priced honestly	12
09	What a fixed target does to the distribution	13
10	Context decides more than the pattern	15
11	Timeframe, cost, and whether this survives	16
12	How it fails	17
13	One trade, end to end	18
14	Measuring it on your own record	20

01

What an impulse leaves behind

Price moves because orders meet. Most of the time they meet evenly: buyers and sellers trade at every price on the way, and the chart records a range that was fully populated. Occasionally they do not. One side arrives with more size than the book can absorb, resting orders are consumed faster than they are replaced, and price travels through a band of levels at which almost nothing changed hands.

That band is the subject of this paper. It is not a pattern in the decorative sense — it is a receipt. It says that at these prices there was demand and no supply, or supply and no demand, and that the market moved on before the imbalance could be settled. Everything that follows is an attempt to trade the settling.

	An impulse	Merely a big candle
What happened	Resting orders on one side were consumed faster than they were replaced.	A wide range filled in both directions, with buyers and sellers trading throughout.
What it leaves	A gap: prices at which almost nothing changed hands.	A full record of transactions at every price inside it.

The distinction the whole setup rests on. A large candle is not a spike; a spike is a move that outran the book and left a hole where trading should have happened. The right column is what most charts are full of.

The distinction in that figure is the one most versions of this setup lose. A large candle is common and means little: a wide range can be filled in both directions, with trading throughout, and it leaves no hole behind it. An impulse is a specific claim about what did not happen, and the gap is the only evidence for it that a chart actually carries.

A caution before the mechanics. This is a common setup taught under many names, and none of what follows is proprietary. What is worth having is not the shape — the shape is free — but the definition precise enough to count, and the arithmetic that says whether counting it is worth your time.

02

Why the pullback arrives in two legs

The retracement after an impulse is usually described as one movement, and it usually is not. It comes in two, with a pause between them, and the pause is the part that makes the shape usable rather than merely recognisable afterwards.



Why the pullback usually arrives in two pieces rather than one. Each step is a different group acting, and the pause between them is what makes the shape countable instead of merely felt.

The mechanism is ordinary. The first leg is largely those who were filled during the impulse taking money off the table; it is quick, shallow and self-limiting, because the people doing it run out. Then nothing happens for a while, which is the market waiting to see whether new interest arrives at the new price. If it does not, a second leg follows — slower, deeper, made of latecomers and of orders that were resting above the move and are now being reached.

First leg ends		shallow: around a third of the move
Pause		little travel, and it can be long
Second leg ends		deeper: often near the gap
Beyond the origin		no longer a pullback — the idea is dead

Schematic, and the reason a single number for 'the pullback' is useless. Measured from the end of the impulse, the two legs stop in different places often enough that a rule written for one of them misses the other entirely.

The two legs stop in different places, and that is the practical reason to distinguish them. A rule written for the first leg — enter on the first pullback — fires early and often, and is wrong in exactly the cases the pause would have warned about. A rule written for the second leg fires less often and has the useful property that the first leg has already told you where the sellers were.

None of this is a law. Retracements sometimes come in one leg, sometimes in three, and sometimes they simply do not stop. The claim is narrower: two is the modal shape, it is countable, and a method that waits for it trades a smaller and better-defined subset of what happens after an impulse.

03

Defining the spike so it can be counted

Everything above is words until the spike has a definition that a spreadsheet could apply without you. This is the part of the setup worth the most effort, because it is the part that decides what enters your sample — and a sample assembled by eye measures your taste rather than the market's behaviour.

Test	Written as	Why it is there
Direction	three or more candles, same sign	one candle is an event, not a move
Size	range $\geq 2 \times$ the last 20 candles' average	separates an impulse from ordinary travel
Gap	a price band with no overlap between candles	the evidence that the book was outrun
Cleanliness	closes in the last third of each range	a long wick means the move was already fought
Origin	the candle the sequence began from	this is the level the stop will belong to

The definition written so that a spreadsheet could apply it. Every row is a number you fix before looking at a chart, and a move that fails any of them is not the thing this paper is about.

The size test deserves a note. Two times the recent average range is a starting value, not a discovery; the right multiple for your instrument is whatever separates the moves that leave gaps from the ones that do not, and half an hour with your own data will find it. What matters is that the number is fixed before the day starts and applied to every candidate, including the ones you would rather take.


The cleanliness test is the one most often dropped, and dropping it is expensive. A candle that closes far from its extreme has already been fought: the aggressive side pushed, met resistance, and gave back part of the move within the same interval. That is not an outrun book; it is a contest, and the pullback that follows behaves like the end of a contest rather than like the settling of an imbalance.

Write the five rows down before you use them, and count how many candidates a week produces. If the answer is more than a handful on a five-minute chart, the definition is too loose and the rest of this paper will not save it.

04

The gap is the evidence, not the decoration

Of the five tests, the gap is the one that carries the argument. The other four describe how the move looked; the gap is the only one that says something about what happened in the order book, and it is the reason this setup has a mechanism rather than a following.



Not every hole is the same hole. The rungs run from the strongest evidence of an outrun book to the weakest, and the bottom rung is where most disappointing versions of this setup come from.

Read the ladder as a scale of evidence rather than a filter with one threshold. A true opening gap is the strongest and the rarest, and outside equity opens and weekend breaks you will not see many. The second rung — consecutive candles whose ranges do not overlap at all — is the workable definition for a continuously traded instrument, and it is strict enough to throw away most of what looks impulsive.

The bottom rung is where the disappointment lives. A large candle with full overlap has an entirely conventional explanation: a lot of trading happened, in both directions, at every price inside it. Treating that as an imbalance is the single most common way this method gets taught wrongly, and it converts a setup with a mechanism into a setup with a shape.

One practical consequence: grade the gap in your record rather than merely requiring one. If the strongest two rungs produce results the bottom two do not, you have learned something specific about your instrument that no general description of the pattern could have told you.

05


Where the second leg ends

The entry has to be a price, not an impression. What follows is deliberately mechanical: it does not ask whether the pullback looks finished, it asks whether a particular thing happened, and that is what makes a hundred of these comparable to each other later.



The entry, written mechanically. It is not a judgement about whether the pullback 'looks finished' — it is a price that either trades or does not, and the difference is the whole reason this can be measured later.

The trigger is a candle that reaches beyond the previous candle's extreme and then closes back through it. After an upward impulse, that means a candle that takes out the low of the candle before it and closes above that low. It is a small event and it is objective, and its virtue is not that it predicts anything — it is that it cannot be argued with after the fact.



The two places the trade can be taken, and what each costs. The lower rung is the one advertised as an improvement; it is better priced and worse filled, and the record settles which of those matters more on your instrument.

There is a second, better-priced entry: halfway between the trigger price and the stop. It improves the multiple by a meaningful amount and it fills less often, and which of those

dominates is a property of your instrument rather than of the method. Log which entry each trade used and the question answers itself in a few dozen trades — long before any argument about it would have finished.

What the trigger is not is a signal to be taken anywhere. Applied without the gap test and without the pause, the same candle appears several times an hour in any market, and the method degrades into an ordinary reversal entry with a good story attached.

06

The stop belongs to the origin

The invalidation is the part of a setup that has to correspond to something. The question a stop answers is not how much you are willing to lose — that is decided by position size — but what would have to happen for the idea to be finished.

Stop placement	What it means when hit	What it costs
Behind the origin candle	the impulse itself was undone	the widest stop, the smallest multiple
Behind the second leg's low	this attempt failed	tighter, but hit by ordinary noise
Behind the gap	the hole was filled	middle ground, and defensible
A fixed number of points	nothing in particular	unrelated to the structure it trades

Where the invalidation belongs, and what each alternative actually buys. The first row is the only one that corresponds to something happening in the market; the rest are ways of paying less for a stop that then means less.

For this setup the answer is unusually clean. The idea is that an impulse left an imbalance and the market is settling it. If price returns beyond the candle the impulse began from, the imbalance has been settled and then some: whatever was unfilled has now been filled, and there is nothing left for the trade to be right about. That is a real event, and it is why the origin is where the stop goes.

The tighter alternatives are tempting for the arithmetic and wrong for the reason. A stop

behind the second leg's low is smaller and produces a larger multiple, but being hit means only that this particular attempt failed — which happens constantly, and says nothing about the imbalance. Traders who move the stop there usually report the same thing: a better reward-to-risk on paper and a hit rate that fell by more than the multiple gained.

The middle row is defensible and worth testing. A stop behind the gap corresponds to a real event too — the hole has been filled — and it sits between the other two in both distance and meaning. If your record shows the origin stop is rarely touched before the gap stop resolves the trade anyway, the tighter version is the better instrument, and you will know because you measured rather than because you preferred it.

07

When the stop is not the same size twice


Because the stop belongs to the origin candle rather than to a fixed number of points, its distance changes with every setup. That is the correct behaviour — the invalidation follows the structure — but it has a consequence that a fixed-target method makes sharper than most.

Stop	Position	Target distance	What it feels like
8 points	2.5 units	8 points	large size, quick resolution
14 points	1.4 units	14 points	the ordinary case
25 points	0.8 units	25 points	small size, long wait
40 points	0.5 units	40 points	barely worth the screen time

Four trades from the same method on the same instrument, sized to the same risk. The stop is set by the origin candle, so it is a different distance every time — and the position has to move with it or the risk does not stay constant.

Read the last column rather than the first. All four trades risk exactly the same amount and the account cannot tell them apart, but they do not feel alike at all: the eight-point version resolves in minutes at a size that looks alarming, and the forty-point version sits there for an

hour at a size that looks pointless. Traders who size by feel take too much of the first and skip the last, which quietly turns a constant-risk method into a variable-risk one.



The order the numbers are computed in. Reversing the first two steps is how a trader ends up with a fixed position and a variable risk, which is the opposite of what the arithmetic in this paper assumes.


The rounding step matters more here than in most methods. On a small stop the theoretical position is large, so rounding to the nearest tradeable increment moves the risk by a percentage point or two; on a large stop the position may round down to almost nothing. Round down rather than up in both cases and accept the slightly smaller trade — the alternative is a risk budget that drifts upward exactly on the setups that resolve fastest.

There is one filter hidden in the table. A stop several times wider than your usual means the origin candle is far away, which usually means the impulse was long or the second leg went deep. Both are legitimate, and both give a trade that ties up attention for a small position. Deciding in advance the widest stop you will take is not risk management — the risk is constant either way — it is a decision about what your time is worth.

08

The one-to-one target, priced honestly

The default exit for this setup is a target at the same distance as the stop. It is a reasonable place to start and it is the most demanding rule in the paper, because of a fact about fixed targets that is easy to state and easy to forget.



What each reward multiple must be right about, before costs. The default one-to-one target is the demanding one; it needs to be right more often than a coin, every month, on an instrument that does not care.

A one-to-one target has to be right more than half the time to make money, and rather more than half once costs are subtracted. That is a high bar for any method, and it is a particularly high bar for one traded on short timeframes, where the cost is a large fraction of the risk. Nothing about waiting for a second leg guarantees a coin can be beaten; the record has to show it.

The attraction of the rule is not expectancy, it is variance. A near-even hit rate produces short losing runs and a smooth equity curve, and both make a method easy to keep executing. That is a legitimate thing to buy — a method you will actually run beats a better one you will abandon — but it should be bought knowingly, and it should be compared against the alternative rather than assumed superior to it.

The comparison that settles it takes one column in the log: did the trade run past twice the risk before it stopped? Count that over a hundred trades and the argument about targets is over, in whichever direction your own instrument decides.

09

What a fixed target does to the distribution

Exiting at a fixed multiple does something specific to the shape of your results, and it is worth seeing laid out rather than reasoned about in the abstract.

Exit rule	Typical hit rate	Average result	What it needs to be true
Fixed 1:1	55%	+0.10R	that the setup clears a coin, after costs
Fixed 2:1	38%	+0.14R	that price runs past the first target often
Half at 1:1, rest to 2:1	55%	+0.16R	that the runners cover the halved winners
Trail behind each leg	44%	+0.19R	that trends continue more than they reverse

The same set of trades under four exit rules. There is no best row in general; there is a best row for a measured record, and the last column names the measurement that decides it.

The rows are close together, which is itself the finding. Across four sensible exit rules the expectancy differs by a few hundredths of a unit of risk, and the differences that look decisive on a spreadsheet are well inside the noise of a hundred trades. The exit rule is not where this method is won or lost.

What differs sharply between the rows is not the average but the experience. The fixed one-to-one wins more than half its trades and pays little; the trailing version loses more often and occasionally pays a lot. Two traders running the same setup with those two rules will report entirely different things about it, and both will be describing the same underlying edge.

The practical advice is unglamorous: pick the rule you can execute for a hundred trades without changing it mid-sample, because a rule changed halfway through produces a record of neither version. If you genuinely cannot tell which one suits you, run the fixed target first — its shorter losing runs make the first hundred trades survivable, and surviving them is what makes the measurement possible at all.

10

Context decides more than the pattern

The shape can be identical in two places and be worth taking in one of them. That is not a hedge — it is the most measurable claim in this paper, and it belongs in the entry rule rather than in a paragraph about experience and feel.

	Worth taking	Worth skipping
Where the impulse began	From a level, a range edge, or against an exhausted move.	From the middle of a range, with nothing behind it.
What it did on the way	Broke a level that had held more than once.	Stopped short of every level anyone was watching.

The same pattern in two settings. The shape is identical and the expectancy is not, which is why context belongs in the entry rule rather than in a paragraph about discretion.

An impulse that begins from a level people were watching is a different event from one that begins in the middle of a range. In the first case there is a reason a queue existed at that price and a reason it was cleared; in the second there is a large candle with nothing behind it, and the pullback has no particular reason to stop where the method wants it to.

With the higher timeframe  the second leg tends to fail early

No clear higher trend  close to a coin

Against the higher timeframe  the pullback keeps going

Schematic, and the single filter that changes the numbers most. The pattern is the same in all three rows; only the direction of the timeframe above it differs, and that is enough to separate a method from a coin.


The higher-timeframe filter is the one that moves the numbers most, and it is close to free to apply. Taking the setup only in the direction of the trend above it removes roughly half the candidates and removes a disproportionate share of the losses, because a second leg that runs against a higher-timeframe move is not a pullback at all — it is the beginning of the next move down, seen too early.

Both filters cost you trades, and that is the point. A method that fires all day is not a method with more opportunity; it is a method with a looser definition, and the extra trades are drawn from exactly the population the definition was supposed to exclude.

11

Timeframe, cost, and whether this survives

This setup is usually taught on one-minute and five-minute charts, and that is where it is most often destroyed. The reason has nothing to do with the pattern and everything to do with arithmetic that applies to every method traded at that speed.



The spread as a share of the stop, for a setup traded on four timeframes. This is the number that decides whether a one-minute version of this method is a strategy or a subscription to your broker.

The spread is a toll paid once per trade regardless of duration, so shortening the trade raises the share of the edge it consumes. On a one-minute chart, a stop behind the origin candle is often ten or twelve points, and a spread of three or four is a third of the risk gone before price has moved. A one-to-one target that needs to win half the time now needs to win rather more than that, and the method has been asked to clear a bar that has quietly risen.

Two responses are honest. The first is to trade the same setup one or two timeframes higher, where the same structure produces stops several times larger and the same spread becomes a rounding error. The second is to stay short-term and restrict the method to the hours in which the spread on your instrument is tightest, which for most is the overlap of the major sessions.

The dishonest response is to tighten the stop so the multiple looks better. That does not reduce the cost; it raises the cost as a share of risk, and it moves the invalidation away from the event it was supposed to represent. It is worth naming because it is the most natural mistake to make and it looks, on the spreadsheet, like an improvement.

12

How it fails

Four failure modes account for most of the distance between this setup as described and this setup as traded. All four are rule failures rather than reading failures, and all four happen for the same reason: the strict version produces fewer trades than a person wants.

How it looks	What it is	The rule that prevents it
Taking spikes with no gap	trading size, not imbalance	the gap test, applied before the entry
Entering on the first leg	impatience wearing the setup's name	wait for the pause and the second leg
Moving the stop off the origin	buying a better multiple with a worse stop	the stop belongs to the structure
Trading it in the quiet hours	paying a third of the risk in spread	measure the spread against the stop

The four ways this setup stops working, and what each one actually is. None of them is a chart-reading failure; each is a rule that was loosened in the moment because loosening it produced a trade.

The first row is the fundamental one. Trading impulses without gaps is not a slightly weaker version of the method — it is a different method, one that trades candle size rather than

imbalance, and its results have no reason to resemble the ones the definition was built for. If a single rule from this paper survives into your trading, it should be that one.

	It works when	It fails when
The market	Volatility is expanding and impulses leave real gaps.	Range is compressing: every candle overlaps the last.
You	You can wait through the pause without taking the first leg.	You need a trade today, and the definition softens to provide one.


The conditions this method needs against the ones in which it quietly stops paying. Nothing in the right column announces itself on the chart, which is why both rows belong in a checklist rather than in judgement.

The lower row of that matrix is the one worth sitting with. Almost everyone can wait for the pause when nothing is at stake. The test arrives on a quiet afternoon after two losses, when a first leg is developing and the definition can be read generously, and no amount of understanding the mechanism helps at that moment. What helps is having written the rule down where it can be consulted rather than recalled.

13

One trade, end to end

The method is easier to hold as a single sequence than as a set of rules. What follows is one trade from the impulse to the exit, with each decision made at the point the method says to make it.



One trade in the order the decisions were made. Every step is a rule from an earlier section, and the point of the sequence is that nothing in it was decided while a position was open.

Two details in that sequence are the point. The stop level was fixed before the entry existed — it belongs to the origin candle, which was on the chart long before there was anything to be right or wrong about. And the eleven minutes of nothing were not a delay to be endured but the evidence the method requires: without a pause there is no second leg, and without a second leg this is a different trade.

	Value	Where it came from
Entry	1.0824	reversal candle close
Stop	1.0812 (12 pips)	behind the origin candle
Target	1.0836	1:1, fixed before entry
Gross result	+1.00R	target filled
Spread paid	1.4 pips	M5, London hours
Net result	+0.79R	what the account actually received

The same trade with the numbers filled in, and the two rows people leave out. The gross result and the net result differ by a fifth, which is the whole argument of section 10 stated once as arithmetic.

The last two rows are the ones usually left out of a worked example, and they are the reason section 10 exists. A gross result of one unit of risk arrived in the account as roughly four fifths of one, because the spread was paid at entry and is measured against a twelve-pip stop. Repeat that across a hundred trades and the difference between the gross and net records is not a detail — it is most of the argument about whether the method is worth running.

14

Measuring it on your own record


Everything above is general. The version that decides anything is computed on your own fills, and it needs six columns — none of which requires data you do not already have.

Column	Why it has to be there
Gap grade	tests whether the gap rule earns its place
Higher-timeframe direction	the filter that changes the numbers most
Entry: trigger or halfway	settles which of the two entries is better for you
Stop distance	turns the multiple into something comparable
Reached 1R before stop?	the raw hit rate, before any exit rule
Ran past 2R?	decides whether a fixed target is costing you

The minimum record. The first three columns describe the setup and the last three describe what it did, and without the middle one the most important question in this paper cannot be answered later.

The fifth column is the one that makes the rest useful. Recording whether the trade reached one unit of risk before it was stopped gives you the raw behaviour of the setup, independent of whatever exit rule you happened to be using — which means a change of exit rule does not invalidate the sample you already collected.

The sixth column settles the target argument by itself. If a large share of trades that reach 1R go on to reach 2R, the fixed target is expensive and the record will say so in numbers rather than in intuition. If they mostly turn around shortly after, the fixed target is the reason the method is profitable, and the temptation to let winners run is the thing to resist.



How many of these you need before the record says anything. At a hit rate near a coin, the honest answer is uncomfortable, and it is the reason a good week proves nothing at all.

The sample-size figure is the unwelcome one. At a hit rate near a coin, small samples are almost uninformative, and the difference between a method that wins 52% and one that wins 46% — which is the difference between a living and a slow loss — is invisible in twenty trades and barely visible in sixty. This is why a good week proves nothing, and why the only honest instruction is to keep the record before you need it.

15

What this does not settle


Three limits, stated plainly, because a paper that presents a setup as a solution has stopped being useful.

First, every number here is illustrative. Two times the average range, a third of the move, a one-to-four-pip spread — these are reasonable middle values chosen to make the arithmetic concrete, not measurements of your instrument. The method of measuring is the transferable part; the figures inside it are not.

Second, the pattern is common enough that it is being traded by many people, some of them faster than you. That is not a reason to avoid it — a widely known structure can still be worth trading if the reason it exists is real — but it does mean the edge is thin, and a thin edge is exactly the kind that costs and sloppiness erase.

Third, and most awkward: a near-even hit rate makes this setup unusually easy to misjudge in both directions. A run of eight wins will feel like mastery and a run of eight losses like a

broken method, and both are ordinary output from a coin-adjacent distribution. The discipline this asks for is not patience during the pause — it is patience across a hundred trades, which is a different and rarer thing.



The whole paper as a sequence. Everything before this figure argues for one of these five steps; nothing in the method happens outside them.

Term	As used in this paper
Impulse	A move that consumed resting orders faster than they were replaced.
Gap	A band of prices at which almost nothing traded.
Origin	The candle the impulse began from.
First leg	The initial retracement, usually taken by those already filled.
Pause	The interval between the legs, in which little travel occurs.
Second leg	The deeper retracement, and the one this method enters against.
Trigger	The candle that takes the previous extreme and closes back through it.
R	One unit of planned risk: the distance from entry to stop.

The terms used in a fixed sense throughout. Several of them are used loosely elsewhere; where that happens, the narrower reading here is the one intended.

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