

TRADING KNOWLEDGE LIBRARY



Confirmation and Its Price

Trading Strategies



Trading Papers

Confirmation and Its Price

Entering after the market shows its hand, and what the wait costs

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	The exchange every entry makes	4
02	What confirmation actually is	5
03	The evidence you never see	6
04	Putting numbers on the wait	7
05	The break-even that settles it	8
06	Where the order sits	9
07	Which way the surprise runs	11
08	Managing without interfering	12
09	Rules that make revision testable	13
10	The cost of not trading	14
11	Logging what you passed	15
12	Choosing your position on the scale	16
13	Splitting the difference	17
14	One idea, carried through both ways	18

01

The exchange every entry makes

There is one decision underneath every entry method ever taught, and most arguments about entries are really arguments about where to stand on it: the earlier you act, the better your price and the thinner your evidence; the later you act, the more you know and the more of the move you have already paid away.

Neither end is correct. They are different products. An anticipated entry buys a tight stop and a large multiple, and pays for it with a lower hit rate. A confirmed entry buys a higher hit rate, and pays for it with a worse price, a wider stop, and a smaller multiple on the same target. The mistake is not choosing one — it is choosing one and then measuring it with the other one's expectations.



Schematic. The same move, entered twice. The anticipated entry sits at the level with its stop beyond it: the best price available and no evidence yet. The confirmed entry waits for the close back through the level and enters above it: worse price, wider stop, and one fact in hand that the first entry did not have.

That framing does more work than it looks. It converts a question of temperament — am I patient or decisive — into a question with an answer: which product has the better expectancy on this instrument, at this cost structure, given the hit rate I actually achieve. Everything that follows is an attempt to make that computable.

	Anticipated	Confirmed
What you buy	Best available price. Tight stop. Large multiple on the same target.	One branch of the possible futures removed before you commit.
What you pay	A lower hit rate, and fills on levels that turned out to mean nothing.	Worse entry, wider stop, smaller multiple, and the moves that never paused.

The two entries are not better and worse versions of one thing; they are two products with opposite specifications. Read down a column to see what you are buying, across a row to see what it costs.

This paper does not argue for confirmation. It argues that the exchange is arithmetic rather than temperament, and that the arithmetic is computable on your own instrument in an afternoon. What it rules out is the position most traders actually hold, which is to take the anticipated entry and then feel betrayed by its hit rate.

02

What confirmation actually is

"Confirmation" is used loosely enough to mean nothing, so it is worth fixing. A confirmation is an event that could have failed to happen and did not. That is the whole definition, and it does two things: it removes a branch of the possible futures, and it costs you the distance price travelled while producing it.

The definition is stricter than the usage. An indicator turning up is not a confirmation, because it could not have failed to turn up once price rose — it is a restatement of the move, arriving late. A close beyond a level is a confirmation, because price could have touched and rejected instead. The test is whether the alternative was genuinely available.

- cheapest** ● A candle closes beyond the level rather than only touching it.
- A candle closes back through the level after breaking it.
- Structure shifts on the execution timeframe — the first lower high.
- dearest** ● A retest holds, by which time a third of the move is usually gone.

The four common confirmations, in rising order of both evidence and cost. Nothing here is free: each rung removes more doubt and takes more of the move with it.

Each rung buys something specific. The close beyond a level separates a wick from a decision. The close back through a broken level is the event most reversal methods are actually waiting for, whether or not they say so. The shift in structure is the first fact that the trend you were trading has stopped rather than paused. The held retest is the strongest and the dearest, because by the time it holds, the move has usually spent a third of itself.

Notice that all four are events on the chart rather than readings from a tool. That is not a stylistic preference. An event has a moment and a price, which means it can be tested and it can be turned into a resting order. A reading has neither.

03

The evidence you never see

Every filter has a false-negative rate: the cases it excludes that would have worked. For a confirmation filter these are the moves that ran without pausing to produce the event, and they are the single reason this subject is argued about instead of measured.

Confirmed, worked		in the journal
Confirmed, failed		in the journal
Never confirmed		invisible

Schematic, and the reason confirmation is argued about rather than measured. A filter sorts the same set of ideas into three outcomes, and only two of them reach the journal. The moves that would have worked but never produced the event leave no record at all.

They are invisible in a journal because they are not losses — they are absences. Nothing was entered, nothing was risked, nothing is recorded. The trader remembers a handful of them as "the one that got away" and forgets the rest, which means the memory of them is both vivid and unrepresentative.

The consequence is a specific, predictable distortion. Traders overestimate what confirmation costs them, because the misses that hurt are memorable. And they cannot say what it earns them, because the losses it prevented also never happened. Both halves of the ledger are missing, and the argument is conducted entirely on feeling.

The fix is mechanical and it is in section 09: log the setups you passed on. Until that log exists, any statement about whether waiting is worth it — yours or anyone else's — is a guess wearing the clothes of experience.

04

Putting numbers on the wait

The exchange can be written down. Take a setup with a target T , an anticipated entry at the level with stop distance S , and a confirmed entry that costs you C in price and widens the stop to S' . The anticipated version pays $(T - 0) \div S$. The confirmed version pays $(T - C) \div S'$.

	Anticipated	Confirmed
Entry	at the level	beyond the confirming close
Distance given up	none	$\approx 25\%$ of the way to target
Stop distance	S	$\approx 1.5 \times S$
Multiple on the same target	$4R$	$\approx 2R$

One setup, priced two ways. The target is unchanged; only the entry and the stop move. This is the whole exchange in four rows, and every number in it is measurable on your own instrument.

On a typical intraday setup where the confirmation costs a quarter of the distance to target

and widens the stop by half, a 4R anticipated entry becomes roughly a 2R confirmed one. That is the price of the evidence, and it is steep: half the payoff, gone before the trade has done anything.

It is worth being clear that neither number is a law. On an instrument that confirms cheaply — one where the close beyond a level tends to happen close to the level — the cost might be a tenth rather than a quarter, and the confirmed version might pay 3R. That is exactly the sort of thing that varies by instrument and by hour, and exactly the sort of thing that is measurable rather than arguable.

05

The break-even that settles it

So the question is precise: does the hit rate rise enough to pay for the reward given up? Break-even is where the two expectancies meet, and for a fixed-risk approach it is a short table rather than a debate.

Reward per win	Break-even hit rate	Roughly
2R	1 in 3	33%
3R	1 in 4	25%
4R	1 in 5	20%
5R	1 in 6	17%
10R	1 in 11	9%

Break-even hit rates for a fixed-risk approach, where a loss is 1R. The right-hand column is the one that decides an argument: a 4R method needs to be right one time in five, a 2R method one time in three.

Read it in the direction that matters. Halving your reward multiple does not require a slightly better hit rate — it requires a dramatically better one. Going from 4R to 2R moves break-even from one in five to one in three, which on a hundred trades is thirteen additional winners that the confirmation must produce out of nowhere.

4R at 30%		+0.50R
2R at 40%		+0.20R
2R at 50%		+0.50R — break-even with the first

The same setup, three ways, in expected R per trade. Confirmation has to lift the hit rate a long way to pay for the reward it gives up — and a lift from 30% to 40% does not come close.

That last comparison inverts the intuition the whole industry runs on. Being right more often is not the objective and never was. A method that is right 40% of the time and pays 2R is worse than one that is right 30% of the time and pays 4R, and no amount of comfort changes it.

The practical test: measure how much your confirmation actually costs in price, and how much it actually lifts the hit rate, on the same set of setups. Both numbers come from the same sample. If the lift does not cover the cost, you are paying for the feeling of certainty, which is a real thing to buy but should be bought knowingly.

06

Where the order sits

The mechanics follow from the choice, and they are worth stating plainly because half of execution error is an order type that does not match the intent.

	Fills without evidence	Fills only on evidence
Best price	Limit order resting at the level. Fills whether or not the level meant anything.	— (a better price on stronger evidence is not on offer)
Worse price	— (paying up for nothing is not a strategy)	Stop order beyond the confirming extreme. Cannot fill on a wick.

Order type is not a detail of execution, it is the decision itself expressed in the platform. The bottom-right cell is the one most traders occupy while believing they are in one of the others.

Anticipating

A resting limit order at the level, stop beyond it. It fills at the best price available and it fills whether or not the level meant anything, which is the entire trade-off in one sentence. Its discipline problem is that it fills while you are asleep, on a level you drew when you were tired.

Confirming

A stop order beyond the confirming candle's extreme, so entry happens only if price goes on to prove the event. It cannot fill on a wick, it cannot fill on the level alone, and it is why a confirmed method does not need you to be watching: the market either produces the event or it does not.

Neither

A market order taken because the chart looked ready. This is the one to name, because it is what most people actually do while believing they are doing one of the two above. A market entry has the price of a confirmed entry and the evidence of an anticipated one — the worst half of each.

07

Which way the surprise runs

A minor point that costs real money: the difference between the two order types is not only price but slippage direction.

Limit: fills at price or better  in your favour

Stop: fills at price or worse  against you

Schematic. The two order types do not merely differ in price, they differ in which direction the surprise runs. A limit order's slippage can only help you; a stop order's can only hurt, and in a fast market it is not small.

A limit order fills at your price or better. Its surprises are pleasant and small. A stop order fills at your price or worse, and its surprises arrive precisely when the market is moving fast — which is to say, precisely on the confirmations you most wanted to act on. The gap is larger than most stop-placement rules assume, and it is systematically in one direction.

This asymmetry is a hidden tax on confirmed methods, and it belongs in the arithmetic of section 05 rather than as an afterthought. If your average slippage on stop entries is a fifth of your stop distance, then the confirmed version of a setup is not paying 2R — it is paying closer to 1.7R, and the break-even hit rate it needs moves again.

Measuring it requires nothing but your own fills: the difference between the price your order named and the price you got, averaged over fifty entries. Most platforms will export it. Very few traders have looked.

08

Managing without interfering

Once in, there are two coherent positions and one incoherent one, and most journals contain all three without distinguishing them.

	While winning	While losing
Static	Plan holds. Testable.	Plan holds. Testable, and uncomfortable.
Dynamic	Revise on new structure. Defensible if the rule was written first.	Revise because it hurts. Not a method.

Two coherent positions and one that is not a position at all. The diagnosis is mechanical: if your revisions cluster in the lower-left cell, the method is not dynamic management, it is reluctance.

The static position is that the trade was decided before entry: target, stop, size, all fixed, and nothing that happens between now and the outcome is new information worth acting on. Its strength is that it is testable — every trade is the trade you designed — and its weakness is that it ignores evidence that genuinely arrives after entry.


The dynamic position is that new structure justifies revising the plan: extending a target when the level ahead is taken out, tightening a stop when the structure that supported the idea breaks. Its strength is that it uses information; its weakness is that it is nearly impossible to test, because no two trades are managed the same way and the record cannot tell you what the rule was.

The incoherent position is to hold the static plan while winning and to revise it while losing. This is not a method, it is an emotional response with a chart attached, and it is diagnosable in any journal: look at whether your management decisions cluster on the losing side.

09

Rules that make revision testable

Dynamic management is defensible, but only in one condition: that the revisions were specified before the trade rather than justified after it. The difference is not moral, it is statistical — a rule written in advance can be measured across a hundred trades, and a decision made in the moment cannot.



If you manage dynamically, this is the minimum that makes it testable later: the evidence, the action and the record, decided in advance and written where the rest of the plan lives.

The last item is the one that gets skipped and the one that matters most. If the journal records only entries and exits, then in three months there is no way to separate the trades where a rule fired from the trades where you simply felt uneasy. Logging which rule fired turns management from a story into data.

A useful constraint while you are building the habit: forbid revision on price alone. Price moving against you is not evidence about the idea, it is the ordinary noise the stop was sized for. Only structure — a level taken, a swing failing, a session ending — should be allowed to move anything.

10

The cost of not trading

Waiting has a second cost that never appears in a spreadsheet: the moves that ran without you. A method built entirely on confirmation will miss the fastest moves in the sample, because the fastest moves are exactly the ones that do not pause to produce evidence.



Schematic, and the hidden bill for patience. In most instruments a small number of days carry a disproportionate share of the range, and they are the days that gap, run and never retest — precisely the days a confirmation filter removes.

This matters more than it sounds. In most instruments, a small number of days carry a large share of the range, and they are disproportionately the days that gap, run and never retest. A filter that removes them removes the fat tail your expectancy was leaning on — and it does it invisibly, because a trade you never took leaves no record.

It is worth sitting with the shape of that. A method can be right about direction, right about level, and right about risk, and still underperform a cruder method that took the same idea earlier — purely because the crude method was present on the four days a year that paid for everything.

11

Logging what you passed

The honest correction is to log the setups you passed on and what they did. It is tedious and it is the only way to see the cost of your own patience.

Setup	Why passed	What it did	R missed
level held, no close	no confirmation	ran to target	+4R
break, no retest	waited for retest	never retested, ran	+6R
close beyond, thin hour	outside window	reversed	0R
retest, weak candle	trigger not met	stopped out anyway	-1R avoided

The log that makes the invisible cost visible. Four columns, filled in for the setups you did not take. A month of this settles whether your filter is protecting you or robbing you.

Four columns are enough: what the setup was, why you passed, what happened next, and the R you gained or missed by not taking it. The last column is the one that turns the log from a diary into a measurement — and note that it is signed, because a passed setup that would have lost is a saving, not a miss.

Traders who keep this log are frequently surprised in both directions. Some discover their filter is throwing away the best trades in the sample, and that a looser rule would have paid for its own losses several times over. Others discover the moves they mourn were mostly the ones that reversed, and that their patience is worth more than it feels like at the time.

Run it for a month before drawing conclusions, and count the entries. If the log has fewer than twenty rows, the filter is not the thing limiting your results — frequency is.

12

Choosing your position on the scale

There is no universally right point on the anticipation–confirmation scale, but there are three things that push a given trader toward one end, and they have nothing to do with conviction.

- costs** ● Spread large relative to your stop pushes toward anticipation.
- attention** ● Cannot watch the screen? Resting orders, which means anticipating.
- temperament** ● Cannot sit through a six-loss run? Take the higher hit rate.

Three inputs decide where on the scale you belong, and none of them is conviction. Work down the list: the first two are facts about your account, the third is a fact about you.

Cost structure comes first because it is arithmetic. If spread and commission are large relative to your stop, confirmation is expensive twice over — the wider stop and the worse fill — and anticipation is usually the cheaper product regardless of anything else.

Attention comes second because it is a fact rather than a preference. Anticipating means resting orders, which suits someone who cannot watch. Confirming on a lower timeframe demands presence at the moment the event forms, and a method that requires presence you do not have is not a method you own.

Temperament comes third, and honestly assessed it outranks the other two more often than anyone admits. A method with a 30% hit rate produces losing runs of six and seven routinely. If those runs will make you abandon the method, its higher expectancy is theoretical, and the lower-paying, higher-hit-rate version you will actually follow is worth more.

The best method you will not run is worth less than the mediocre one you will. This is not an argument for comfort — it is an argument for choosing a method you can survive, then improving it from the inside.

13

Splitting the difference

Sooner or later every trader tries the obvious compromise: half the position at the level, half on confirmation. It is worth pricing rather than dismissing, because it is a real answer to a real problem — and because what it actually buys is not what it is usually adopted for.

	All anticipated	Split 50/50	All confirmed
Average entry	at the level	between	beyond the close
Effective multiple	4.0R	≈ 3.0R	1.9R
Hit rate	31%	≈ 37%	44%
Expectancy	+0.55R	+0.48R	+0.28R
Trades where you feel wrong	many	most	fewer

Half at the level, half on confirmation, against the two pure versions. The split does exactly what it looks like it does: it lands between them on every measure, including the one people hope it will escape.

Read the table as a blend, because that is all it is. The split sits between the two pure versions on entry, on multiple, on hit rate and on expectancy, in roughly the proportion you split. There is no interaction effect, no free lunch hiding in the middle, and any account of the split that promises the multiple of one with the hit rate of the other is describing something that does not exist.

What it does buy is variance. Two half-positions entered at different prices produce a narrower distribution of outcomes than either pure version, which matters to an account that cannot survive a long run of full losses. That is a legitimate reason to choose it, and it is a different reason from the one most people give, which is that it feels better to have some position on when the move goes without them.

The cost is bookkeeping, and it is not trivial. A split entry is two trades with two entry prices, and unless the log records them separately, neither the confirmation cost of section 07 nor

the two hit rates of section 13 can be computed afterwards. A trader who splits without logging both halves has bought the middle of the scale and lost the ability to find out where on it they belong.

14

One idea, carried through both ways

Abstraction is where this subject usually dies, so here is the same idea taken end to end under both products. The level, the target and the day are identical; everything after the second step differs.



One idea, carried through both products end to end, so the arithmetic stops being abstract. Same level, same target, same day — everything after the second step differs.

The sizing step deserves a note. Position size follows from the stop, not from how sure you feel — which means the confirmed version, with its wider stop, takes a smaller position for the same risk. Traders who size by conviction rather than by distance quietly take more risk on the trades they like most, which is the opposite of what the arithmetic wants.

	Anticipated	Confirmed
Entry	1.0850	1.0863
Stop	1.0838 (12 pips)	1.0845 (18 pips)
Target	1.0898	1.0898
Multiple	4.0R	1.9R
Hit rate over 100 trades	31%	44%
Expectancy per trade	+0.55R	+0.28R

The same trade, both ways, with the numbers filled in. The confirmed version is right more often and earns less — which is the point of the whole paper, stated once in figures rather than prose.

Read the last two rows together. The confirmed version is right thirteen percentage points more often and earns roughly half as much per trade. Both facts are true at once, and only one of them shows up in how the trading feels.

15

Reading your own record

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


Measure	Where it comes from	What it decides
Average confirmation cost	entry price minus level, 50 trades	the reward you give up
Stop widening factor	confirmed stop ÷ anticipated stop	the multiple you keep
Hit rate, anticipated	your own record	half of the comparison
Hit rate, confirmed	your own record	the other half
Average slippage on stops	fill minus order price	the hidden tax
Spread in your window	platform export, by hour	whether either version works

The six numbers that decide this for you rather than in general. None of them requires data you do not already have: five come from your own fills, one from your platform's export.

Collect them in that order, because the first two can settle the question on their own. If your confirmation costs a third of the distance and widens the stop by two, no plausible lift in hit rate rescues it, and there is no need to argue about the rest.

The two hit rates are the slow part, and the reason is sample size rather than effort. A difference of ten percentage points is invisible in thirty trades and merely suggestive in a hundred.

That is an uncomfortable number for anyone who wants an answer this month, and it is the honest one. It also argues for measuring on historical setups rather than only on trades taken: the chart supplies hundreds of instances of both products, and it supplies them without costing anything to be wrong about.



How much of a difference a sample can actually see. With a 10-point gap in hit rate, thirty trades tell you nothing you can rely on; the answer only firms up past a hundred, which is why a month of trading rarely settles this question.

16

What this does not settle

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

First, the numbers here are illustrative. A quarter of the distance and a 1.5× stop are reasonable middle values, not measurements of your instrument. The method of section 05 is the transferable part; the figures in it are not.

Second, the hit rates you measure are contaminated by you. A confirmed method traded impatiently produces the statistics of neither product, and no amount of arithmetic separates the method's behaviour from the operator's.

Third, and most awkward: both products can be positive expectancy at once, and the better one is then a question of which you will execute consistently over a hundred trades. That is not a cop-out. It is the same reason two competent traders can hold opposite views on this and both be right about their own accounts.

Term	As used here
Confirmation	An event that could have failed to happen and did not.
Anticipation	Entering at a level before any such event.
R	One unit of planned risk: the distance from entry to stop.
Expectancy	Average R per trade over a sample, after costs.
False negative	A move that would have worked but never produced the event.
Static management	The trade is decided before entry and not revised.
Dynamic management	Revision permitted, by rules written before entry.

The terms this paper uses in a fixed sense. Where the industry uses one of them loosely, the definition here is the narrower one.

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