

TRADING KNOWLEDGE LIBRARY

The Five Claims

Trading Strategies



Trading Papers

The Five Claims

Every named strategy, grouped by what it actually asserts about price

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	Twelve names, five claims	4
02	Claim one: a price that mattered before matters again	5
03	Claim two: a move in progress continues	7
04	Claim three: a move that went too far comes back	8
05	Claim four: the market moves in proportion	9
06	Claim five: price returns to what it skipped	11
07	Why one chart supports three claims	12
08	What it takes to test one	13
09	Indicators are not a sixth claim	14
10	Which claim suits which market	16
11	Stacking claims makes signals rarer, not safer	17
12	Choosing one	18
13	What this does not settle	19

01

Twelve names, five claims

There is a genre of trading content that promises every strategy in twelve minutes, and it always delivers a list: support and resistance, supply and demand, breakouts, trend following, momentum, mean reversion, Fibonacci, Elliott wave, harmonics, fair value gaps. The list is accurate. It is also the least useful possible arrangement of the material, because it presents as a dozen alternatives what is actually a much smaller number of ideas.

The name	The claim underneath
Support and resistance	a price that mattered before matters again
Supply and demand zones	a price that mattered before matters again
Dynamic support (moving averages)	a price that mattered before matters again
Breakout trading	a move in progress continues
Trend following	a move in progress continues
Momentum	a move in progress continues
Mean reversion	a move that went too far comes back
Overbought and oversold	a move that went too far comes back
Fibonacci retracement	the market moves in proportion
Elliott wave, harmonics	the market moves in proportion
Fair value gaps, imbalance	price returns to what it skipped

The names you will meet, and the claim each one is making. Read down the right column: there are five entries, and the left column has more than twice as many. That ratio is the whole argument of this paper.

Read that table down the right-hand column. Eleven names, five distinct claims — and three of the names are making exactly the same assertion in different drawing styles. Once the collapsing is done, the field stops looking like a menu and starts looking like

a small set of incompatible beliefs about why price moves at all.



The five claims, in order of how much they assume. The first assumes almost nothing about the market and the last assumes a great deal; that ordering is also, roughly, the order in which they get harder to test.


The ordering on that ladder is by how much each claim assumes. That a price which mattered before matters again assumes almost nothing: people remember, orders rest, and the same numbers are visible to everyone. That the market moves in fixed proportions assumes a great deal, and the ordering is roughly also the order in which the claims get harder to test — which is the subject of the second half of this paper.

The named techniques here are the common vocabulary of the field and belong to nobody. The grouping into five claims, the drawings, the arithmetic and the argument are ours.

02

Claim one: a price that mattered before matters again

This is the oldest claim and the one with the most costumes. It says that a particular price has a history, and that the history will show up again in behaviour.



Claim one, drawn. A price is reached, rejected, reached again and rejected again — and the third approach is the one every version of this claim is about. Support and resistance draws it as a line; supply and demand draws it as a band; a moving average draws it as a line that moves. All three assert the same thing about the same place.

The drawing is the whole claim. Price reached an area, was rejected, came back, was rejected again — and the third approach is what every version of this method is waiting for. Nothing about the mechanism is mysterious: orders that were not filled the first time may still be resting, traders who were wrong there may want out at break-even, and everybody can see the same number.

The costume	How it is drawn	The one real difference
Support / resistance	a single horizontal line	a price
Supply / demand zone	a band with a width	a tolerance around that price
Dynamic support	a moving average line	the price moves with time
Round numbers	a line at 1.2000, 100, 50000	chosen before the chart, not from it

The three costumes of claim one, and what actually differs between them. Only the last column is a real difference; the rest is drawing style, and arguing about which is correct is arguing about a pencil.


The costumes differ in drawing style and almost nothing else. Support and resistance draws a line; supply and demand draws a band; a moving average draws a line that moves with time. The only genuine difference is in that last column — a zone is a level with a stated tolerance, and a moving average is a level whose price changes — and arguing about which is correct is arguing about a pencil.

Round numbers are worth separating out, because they are the one member of this family chosen before the chart rather than from it. Nobody has to identify 1.2000 or 50,000; it was always there, everybody sees the same one, and that removes the largest source of disagreement in the whole family.

03

Claim two: a move in progress continues

The second claim is the engine of breakout trading, trend following and momentum alike. It says that what is already happening is more likely to keep happening than to stop.



Claim two, drawn. A range, a close beyond it, and continuation — the shape every breakout, trend and momentum method is describing. The accent candle is the one the strategies disagree about: breakout enters there, trend following waits for the next higher low, momentum waits for a reading to confirm what already happened.

The picture is a range, a close beyond it, and continuation. The three costumes disagree only about where to act. Breakout methods enter at the accent candle, on the close beyond the range. Trend followers decline to act there and wait for the first higher low afterwards. Momentum methods wait for a calculation to confirm what already happened, which is later still.


That disagreement is not cosmetic — it is the entire difference in their results. Entering at the break gets the best price and the highest rate of false starts. Waiting for the higher low gets a worse price and skips most of the failures. Neither is right; they are two positions on one trade-off, and a method that has not stated which it takes has not been written down.

Note also what this claim requires that the chart does not supply: a definition of the range, and a definition of beyond. A close one tick past the edge and a close two per cent past it are both breaks under most written descriptions, and they are not the same event.

04

Claim three: a move that went too far comes back

The third claim is the exact negation of the second, which is worth stating plainly because the two are usually taught in separate chapters as if they could both be true at once.



Claim three, drawn, and it is claim two's exact opposite. The same extension that a continuation method reads as strength, a reversion method reads as stretch. Both are looking at the identical candles; what differs is which of two incompatible things they believe about them.

Same kind of chart, opposite reading. A long extension away from an average is, to a continuation method, evidence of commitment; to a reversion method it is evidence that everybody who wanted in is already in. The candles are identical. What differs is which of two incompatible beliefs the reader brings.

	A long extended move	What follows from it
Continuation reads	Strength: participants are committed and keep committing.	Buy in the direction of the move.
Reversion reads	Stretch: everybody who wanted in is already in.	Sell against the move, or wait for it to snap back.

Claims two and three read the same candle in opposite directions, and neither is being unreasonable. This is the single most useful thing to understand about the whole field: the strategies are not competing on quality, they are asserting incompatible things about the same picture.


This is the most useful single thing to understand about the whole field, and it dissolves most arguments about which strategy is best. The strategies are not competing on execution quality. They are asserting contradictory things about the same picture, and both have long stretches of history in which they were correct.

The practical consequence is that a trader who holds both claims at once holds no claim at all. Any chart can be read as extended or as trending, so a method that permits either reading permits every trade, and permits it after the fact. If you cannot say which of these two you believe, that is the first thing to settle — before any question about entries.

05

Claim four: the market moves in proportion

The fourth claim is the most elaborate and the most contested. It says the size of the next move is related to the size of the last one by particular ratios.



Claim four, drawn. A move, then a pullback that stops at a particular fraction of it. Fibonacci names the fractions in advance; Elliott and harmonic patterns extend the same idea to sequences of them. What the drawing cannot show is the part that matters — whether that fraction was chosen before the pullback or after it.

Drawn, it is straightforward: a move, then a pullback that stops at a fraction of it. Fibonacci retracement names the fractions in advance — thirty-eight point two, fifty, sixty-one point eight. Elliott wave extends the idea into a sequence of moves with rules about which may be longest. Harmonic patterns specify several ratios that must hold at once.

The difficulty	What it does to a test
The move must be chosen first	two analysts draw different moves on one chart
Several levels are named at once	one of 38.2, 50, 61.8 is usually near a turn
The count can be revised	a wave that failed becomes a different wave
Tolerance is unstated	near the level does the work the level cannot

Why claim four is the hardest of the five to test, one difficulty per row. None of these makes it false; all of them make it slippery, and slipperiness is what lets a method survive without ever being checked.

The difficulties in that table are what make this claim hard rather than false. The swing has to be chosen before the ratios mean anything, and two analysts will choose differently on the same chart. Several ratios are named simultaneously, so a turn near any of them counts. A wave count that fails can be relabelled, and a relabelled count is a description rather than a prediction.


None of that proves the claim wrong. It means that testing it requires far more discipline

than the others: the swing selection has to be mechanical, the ratio has to be one ratio, and the tolerance has to be a number decided before you look. Most published enthusiasm for this family skips all three, which is why the argument about it never resolves.

06

Claim five: price returns to what it skipped

The fifth claim is the newest in popular usage and, usefully, the narrowest.



Claim five, drawn. Price moves so fast that a range is skipped — no trading happened between the first candle's high and the third candle's low — and the claim is that the skipped range gets revisited. It is the narrowest of the five claims and the easiest to define, which also makes it the easiest to check.

When price moves fast enough, a range gets skipped: no trading happened between one candle's extreme and the next but one's, and the claim is that the market comes back to trade there. It goes under fair value gap, imbalance, and inefficiency, and unlike the previous claim its object is unambiguous — the skipped range has exact boundaries you can read off the chart.

That precision is its real advantage. Where a Fibonacci level requires four judgements before it exists, a skipped range requires none: three candles define it arithmetically, two people will identify the same one, and whether price later traded back through it is a yes or a no. Of the five claims, this is the one a beginner can most easily count.


It has one specific weakness that its enthusiasts underplay, and it is a deadline. Given enough time, price returns almost everywhere; a claim that a gap will be filled eventually is

close to unfalsifiable. The claim only becomes testable when it says by when — within ten bars, within the session, before the next gap forms — and that number is almost never stated.

07

Why one chart supports three claims

Having separated the claims, it is time to admit the problem that makes the whole field so hard to argue about.



One chart, three claims, all present at once. There is a level that has held twice, a range that has just broken, and a move extended far enough for a reversion method to call it stretched. Nothing here is contrived — this is what an ordinary chart looks like, and it is why every strategy has screenshots that work.

That chart is not contrived. It contains a level that has held twice, a range that has just broken, and a move extended far enough that a reversion method would call it stretched. All three claims are present simultaneously, and each could be used afterwards to explain whatever happens next.

This is why every strategy has convincing screenshots, and why the screenshots are worth nothing. On any chart of reasonable length, an analyst holding any of the five claims will find instances of it — not by dishonesty, but because the claims describe things that genuinely occur. The claims are not rare enough for their presence to be evidence.

It also explains a common experience: a trader learns a method, sees it everywhere for a fortnight, and concludes it works. Seeing it everywhere is the expected outcome whether or

not it works, because that is what looking for a common pattern does. The only thing that distinguishes the two possibilities is counting, which is the subject of the next section.

08


What it takes to test one

All five claims are testable in the same way. What differs is how much has to be nailed down before the test can start.

Claim	Must be fixed in advance	Usually left vague
A price matters again	how a level is chosen, and its tolerance	which of many touches counts
A move continues	what counts as the range, and as a break	how far past the edge is a break
A move comes back	what counts as too far, from what mean	the period of the mean
Price returns to a gap	the exact skipped range, and a deadline	how long you wait for the fill
Proportions hold	which swing, which ratios, what tolerance	all four of those

What each claim needs before it can be tested, stated as the thing you must fix in advance. The right column is where each family of methods usually stays vague, and vagueness is not a small flaw here — it is what makes a count impossible.

The middle column is the price of admission for each claim, and the right column is where its literature usually stays vague. Notice that the vagueness is never about the interesting part — it is about the boring part. How far past the edge counts as a break, how long you wait for a fill, which of four touches makes a level: these are exactly the details that decide the result, and exactly the details a promotional explanation omits.



How to test any one of the five, in the same procedure. The fourth step is what separates a test from a collection of screenshots, and it is the step no promotional material ever performs.

The procedure is the same for all five, and the fourth step is the one that matters. Counting what happened after your signal tells you very little on its own; you need to know what happened after every other bar in the same data. A breakout that is followed by a rise sixty per cent of the time, in a market that rose sixty per cent of the time anyway, has told you about the market rather than about breakouts.

One realistic expectation, stated so that a negative result is not read as failure. What survives a test of this kind is usually a small edge, often smaller than the spread that pays for it, and often larger once the claim is required to occur in specific conditions rather than anywhere. Finding a small edge is the normal outcome for a claim that has any merit at all.

09

Indicators are not a sixth claim

A survey of strategies usually devotes a section to indicators, as though they were a separate family. They are not, and seeing why removes a great deal of confusion.

	What an indicator is	What it is treated as
In principle	A calculation on price or volume you could do by hand.	An independent source of information about the market.
In practice	A measurement of how far, how fast, or how much traded.	A reason to enter, with the claim left unstated.

Why indicators are not a sixth claim. Momentum and volume tools are instruments for measuring the claims already listed — they answer how much, not whether. Treating a reading as a reason is how a measurement gets promoted into a strategy.

An indicator is a calculation on price or volume. It cannot know anything the price does not already contain, because it is made of the price. What it does is measure — how far a move has gone, how quickly, how much traded — and a measurement is an input to a claim rather than a claim of its own.

The indicator	What it measures	Whose claim
Moving average	the average of recent closes	a price matters again
RSI, stochastic	how far a move has extended	a move comes back
MACD	the gap between two averages	a move continues
ATR	the typical size of a bar	none — it sizes the others
Volume	how much traded at these prices	any of them, as corroboration

Common indicators, and which of the five claims each one is measuring. Once the claim is named, the question stops being which indicator is best and becomes whether that claim holds on your instrument.

Read the last column of that table and the indicators stop competing. RSI measures extension, so it belongs to claim three; a trader using RSI is asserting that moves which went far come back, and the indicator only says how far. MACD measures the gap between two

averages, so it belongs to claim two. A moving average is a level, so it belongs to claim one.

ATR is the exception and worth its own sentence: it makes no claim at all. It measures how large a bar typically is, and its use is to set stop distances and position sizes in units the instrument understands. It is the only item in this paper that improves every method equally.

Volume is the one genuine outside input, since it is not computed from price. That makes it more interesting than the others and does not make it a claim: volume is used to corroborate one of the five, and a trader who cannot say which is using it as decoration.

10

Which claim suits which market

The reason the argument between methods never ends is that each one is nearly right some of the time and systematically wrong at other times, and the market does not announce which regime it is in.

	Trending market	Ranging market
Continuation methods	Work, and produce the screenshots.	Break out, reverse, and pay the spread repeatedly.
Reversion methods	Sell strength too early, repeatedly.	Work, and produce their own screenshots.

Which claim suits which market, and the reason nobody wins the argument. Each claim has conditions in which it is nearly right and conditions in which it is systematically wrong, and no market announces which it is currently in.

In a trending market, continuation methods work and produce the screenshots; reversion methods sell strength too early, repeatedly, and their users conclude the method is broken. In a ranging market the results invert exactly. Neither

user is doing anything wrong, and both have a period of evidence.

The tempting response is a filter that identifies the regime, and it is worth being honest about it: regime identification is itself a claim, usually claim two in disguise, and it is identified after the fact by the same tools. That does not make it useless — a long-run average does separate trending from ranging often enough to help — but it is not a way out of the problem, it is another instance of it.

The more workable response is narrower and less satisfying. Pick the claim whose failure mode you can live with, size so that its bad regime is survivable, and accept that a portion of every year will be spent being wrong in a way you predicted in advance.

11

Stacking claims makes signals rarer, not safer

The natural instinct on discovering that each claim fails sometimes is to combine several and act only when they agree. It is a reasonable instinct with an arithmetic problem.



Schematic of what happens when claims are stacked to make a signal safer. Each added condition genuinely raises the quality of what remains and cuts how much remains, and past two or three the sample collapses faster than the accuracy improves.

Each additional condition genuinely raises the quality of what remains — that part of the intuition is correct. What it also does is cut how much remains, and the cutting compounds. Two claims agreeing might leave forty per cent of the signals; three leaves a seventh; four leaves a handful a year on one instrument.

The problem is not that rare signals are bad. It is that a method producing seven signals a


year cannot be evaluated. You would need a decade to reach a sample worth arguing about, and long before then the market, the instrument or your own rules will have changed. A confluence method is often not a better method but an unfalsifiable one.

The practical limit for a beginner is two: one claim as the reason, and at most one other as a filter. That keeps the sample large enough to produce a number within a quarter or two, which is the only thing that will ever tell you whether any of this works for you.

12

Choosing one

The point of collapsing a dozen names into five claims is to make a choice possible. A menu of twelve invites collecting; a set of five incompatible assertions invites deciding.



How to choose one claim rather than collect five. The second step is the one people skip, and skipping it is why so many traders own an opinion about every method and a record of none.

The first step is more important than it sounds. Believing the logic is not sentimentality — it is what makes you keep following the rule through the losing stretch that every claim has. A trader running a mean-reversion method they find unconvincing will abandon it in exactly the drawdown it was designed to survive.

The second step is the one people skip. A claim that occurs four times a year on your instrument during your hours cannot be tested by you, however sound it is. Frequency is a property of the claim and your circumstances together, and it should be checked before any of the effort goes in.

One claim, tested	Five claims, collected
A number, and a reason to trust it	An opinion about each
A definition somebody else could apply	A feel for what looks right
A losing period you can interpret	A losing period that means switch
One thing to improve next quarter	Five things to read more about

What you get from one claim tested, against five claims collected. The right column is the more common position by a wide margin, and it feels like knowledge while producing none.

That table is the difference the choice makes. One claim tested gives you a number, a definition somebody else could apply, and a losing month you can interpret. Five claims collected give you an opinion about each, a feel for what looks right, and a losing month that always suggests switching to one of the other four.

13

What this does not settle

Three limits, in the spirit of the rest of this library.

First, the grouping into five is ours and it is a tool rather than a discovery. Someone could reasonably split continuation into trend and breakout, or fold gaps into levels, and end with four or six. The value is not in the number; it is in noticing that the field is smaller than its vocabulary and that several of its arguments are between people using different words for one idea.

Second, nothing here says which claim is true. Each has conditions where it holds and conditions where it does not, and this paper deliberately declines to recommend one — partly because the answer differs by instrument and horizon, and mostly because a recommendation you did not test would be exactly the kind of thing this library exists to argue against.

Third, this is a paper about direction, and direction is the smaller half of trading. A correct

claim, sized carelessly, still empties an account; the papers on risk in this library are not a supplement to this one but the part that decides whether any of it survives contact with a losing month.

Term	As used in this paper
Claim	An assertion about why price moves, which a method acts on.
Level	A price treated as significant because of earlier trading there.
Zone	A level with a stated width, rather than a single price.
The mean	A stated average a reversion method measures distance from.
Fill	Price returning through a range it previously skipped.
Horizon	The fixed number of bars over which an outcome is measured.

The terms this paper uses in a fixed sense. Several are used loosely across the field; where that happens, the narrower reading here is meant.

The useful summary is short. There are not twelve strategies. There are five claims about why price moves, most of the names are costumes, and the honest question is not which is best but which one you are prepared to define, count, and be wrong about on purpose for a quarter. Everything else in the twelve-minute list is vocabulary.

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