

TRADING KNOWLEDGE LIBRARY



The Rule You Can Actually Hand Over

Market Analysis



Trading Papers

The Rule You Can Actually Hand Over

Writing a rule, testing it honestly, and the three jobs that stay yours

Published by **Trading Papers**

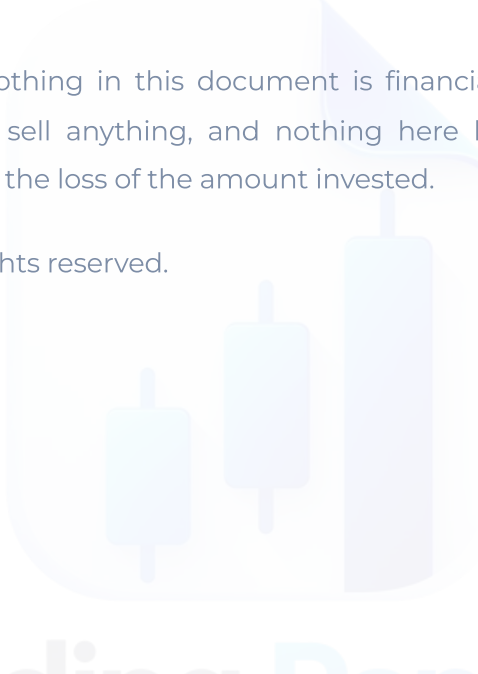
Series **The trading knowledge library**

Edition **First edition · 2026-08-27**

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 nothing here knows your circumstances. Trading carries risk, including the loss of the amount invested.

© 2026 Trading Papers. All rights reserved.



Trading Papers

Contents

01	An idea is not a rule	4
02	The sentence a stranger could run	5
03	Where the assistant genuinely helps	6
04	Counting the occurrences	7
05	Tallying what happened after	8
06	The rule that moved while you tested it	10
07	The number that decides nothing	12
08	What the test cannot tell you	13
09	The three jobs that stay yours	14
10	Sizing, which is arithmetic you still do	15
11	Living with a rule that is running	17
12	A session that works, end to end	18
13	Where each piece already has a paper	19
14	What this does not solve	20

01

An idea is not a rule

Everyone who trades has ideas about the market, and almost none of those ideas survive a test. That is not a criticism of the ideas. It is a statement about what they are made of.

The two things a trading idea can become, and the fork is narrower than it looks. Only one branch can be tested, only one branch can be repeated, and almost every idea people describe out loud is sitting on the other one.

An idea is a description of something you recognise. A rule is a sentence somebody else could follow without asking you a single question. The gap between the two is where this paper lives, and it is wider than it looks from the inside.

What people say	What is still undecided
Buy when the trend is up	what a trend is, in bars
Enter on a pullback to support	how far back, and how wide the level
Wait for confirmation	confirmation by what, within how long
Cut it if it goes wrong	wrong at what price, decided when

Four idea-shaped sentences and what each is missing before it can be run. The right column is not pedantry. Each gap is a decision, and if you do not make it somebody or something else will make it for you and never mention that it did.

Read the right column of that table slowly. Every gap is a decision, and every decision changes what the rule does. Leave one open and something fills it — you, differently on each pass, or an assistant, silently, which is the argument the second paper made at length.

So the first job is to close the gaps. It is a writing job, and writing is the thing a language model is actually for, which is why this is the part of trading where an assistant earns its place most clearly.

The third paper in AI Trading, and the last. It assumes the first two: that the assistant turns a picture into a description, and that each analysis job is a count that needs a rule from you.

02

The sentence a stranger could run

A rule reaches its finished state when nothing in it requires your judgement to execute. That is a high bar, and it comes down to six clauses.



The six clauses a rule needs before it can be counted. A rule missing any one of them cannot be tested, because the missing clause will be filled in differently on every pass, including by you.

Six is not a style preference. Drop any one and nobody can count the rule, because every pass through the data resolves that clause differently — a little more generously on the trades that were about to work, a little less on the ones that were not.

Before	After
Buy the pullback in an uptrend	on the hourly, after three higher swing highs
...at support	...when price returns inside a zone one percent wide
...with confirmation	...and one hourly bar closes above that zone
...and manage it	...stop below the zone, out at twice that distance or in two days

The stranger test, applied to one rule twice. The left column is the version people write; the right is the same idea after every clause has been decided. Only the right column can be handed to anyone, including to a future version of you who has forgotten what you meant.

The right column of that table is what an hour with an assistant should produce. It is not a better idea than the left column; it is the same idea with every hole filled, which is the only version that can be tested, repeated, or handed to anyone else.

It is also the version you will still understand in six months. A rule written in the left column's language is a memory, and memories drift towards whatever you have started believing since.

03

Where the assistant genuinely helps

Across a set of sessions this treatment worked through, the help was real and unevenly distributed.



Where an assistant actually earned its place across the same forty sessions, measured as the share of the work it removed. The two tall bars are language and arithmetic, which is the whole thesis of these three papers. The short bar is the one people ask for most.

The two tall bars are language and arithmetic. That is the whole thesis of these three papers, arriving for the third time: the machine turns pictures into descriptions, it counts what is in front of it, and everything it does well is one of those two things wearing different clothes.

The short bar deserves a sentence of its own. Asked whether a rule was any good, the answer agreed with whatever the question implied, nine times in ten — not because the model is dishonest, but because the fourth blind spot of the first paper applies here without modification. It cannot decline, so it produces a verdict, and a verdict has to lean somewhere.

- Useless** ● give me a profitable strategy
- Weak** ● improve this strategy
- Usable** ● list every clause this rule still leaves undecided
- Strong** ● write it so a stranger could run it, and mark my choices

Four ways to ask for a rule, and what each returns. The bottom rung is the request people actually make, and it is answered with a rule that sounds finished because sounding finished is what the machine is for.


The bottom rung of that ladder is the request most people actually make, and it is answered. What comes back is a rule assembled from the most common phrasings in the training data, with clauses that sound decided and were never tested together, delivered in the voice of something that has checked.

The top rung is different in one specific way: it asks the assistant to mark which parts came from you. That single instruction turns an opaque paragraph into a document you can audit.

04

Counting the occurrences

Once you have the rule in writing, count how often it happened at all.



The first count a rule needs: how often the condition happened at all. Three qualifying setups in this window, marked at the bar where the entry clause was met. A rule with four occurrences in a year is not a strategy, and this count is what tells you which one you have.

This sounds trivial and it is the step that ends most strategies within ten minutes, which is a good use of ten minutes. A rule that fires four times a year cannot be evaluated by anybody, and a rule that fires two hundred times a day is a different problem wearing the same words.

Ask for	Because
The number of occurrences	a rule with six is untestable
The bar number of each	you can check two by eye
The longest gap between them	tells you what waiting feels like
The count in the first half and second	a rule that stopped firing has died

What to ask for alongside the count, and why each one matters. The last row is the one that saves the most time: a rule that fires twice a year cannot be evaluated at all, and finding that out costs one question.

Ask for the bar numbers alongside the count, for the same reason the second paper gave: a count with positions can be checked in thirty seconds, and a count without them is a claim. Check two. Not all of them, and not none.

The last row of that table is quietly the most useful. Split the window in half and count each half separately. A rule whose occurrences all sit in the first half has stopped happening, and no amount of outcome arithmetic afterwards will tell you that as directly.

05

Tallying what happened after

The second count is what followed each occurrence, with the rule's own exit clauses resolving it and nothing else.



The second count: what happened after each occurrence, tallied under the rule's own exit clauses. Four buckets, no interpretation, and the fourth is the one people forget to ask for and the one that hurts most in practice.

Four buckets, because a trade under a written rule can only end four ways: it reached the target, it reached the stop, or the clock ran out with it slightly ahead or slightly behind. There is no fifth bucket for trades you would have managed differently, and adding one is how a test becomes a story.

Reading	What it emphasises	What it hides
Thirty-four percent hit target	the wins	the size of the losses
Forty-one percent stopped out	the losses	that the target was twice the stop
Half the trades went nowhere	the waiting	that going nowhere cost almost nothing

The same tally read three ways. A single set of counts supports three different sentences, and the assistant will write whichever one the question leans towards, in the same confident voice each time.

That table is the part worth pinning up. One set of counts, three sentences, all of them true, and an assistant will write whichever one the question leaned towards — because you asked whether the rule works, and the machine has to answer something.

The defence is mechanical rather than clever. Ask for the four counts as four numbers, before you ask for any sentence at all. Numbers first, prose second, and read the numbers before you read the prose.

06

The rule that moved while you tested it

This is the section the rest of the paper exists to set up, because it is where an afternoon of good work turns into nothing at all.

	The numbers looked good	The numbers looked bad
Rule written down first, unchanged	A result. Worth running forward on new data.	A result. You just saved yourself the money.
Rule adjusted while looking	Unknown. You may have fitted the rule to the sample.	Unknown, and you will keep adjusting until it flips.


What separates a test from a story. The row is whether the rule was frozen in writing before the count ran; the column is what the count said. Only the top row produces anything you can act on, and the bottom right is the most common outcome of an afternoon with an assistant.

The bottom row is not fraud. It is what happens naturally when the rule and the results share a screen and both invite editing, and an assistant makes it frictionless: you change a clause and re-run the whole count in the time it takes to read this sentence.

The edit	What it really did
Widened the zone slightly	added the trades that worked
Moved the stop a little further	removed the losses that stopped out
Skipped the quiet months	deleted the period it failed in
Added one more filter	removed the trades you already knew about
Changed the target after seeing the tally	chose the answer, then measured

Five edits that feel like refinement and are not. Each one moves a clause after seeing the outcome, which is the definition of fitting, and an assistant will make any of them instantly and describe the result as improved.

Every edit in that table feels like refinement while you are making it. Each one is a clause moved after seeing the outcome, and each produces numbers that look better and mean less than the numbers before it. Do it five times and you have a rule that describes your sample perfectly and knows nothing about the market.



How to make the count mean something, in the order that matters. Step two is the one people skip, and it is the only step that turns everything after it into evidence rather than illustration.

The remedy is the second step, and it is administrative rather than intellectual: put the written rule somewhere you cannot quietly edit before you run the first count. A message sent to yourself, a dated file, a photograph of a notebook page. What matters is that the version you tested is recoverable afterwards.

The fourth step is the other half of the same discipline. Deciding in advance what result would make you abandon the rule is what makes the count capable of telling you something you did not want to hear, which is the only kind of result worth running.

07

The number that decides nothing

Ask an assistant how a rule performed and the first number it offers will be the percentage of winners. That number, alone, decides nothing whatsoever.

The number	On its own it tells you
Percentage of winners	nothing at all
Average win against average loss	nothing without the first number
The two together	whether the rule made or lost money
The worst run of losses	whether you could have kept trading it

Four numbers a test returns and what each one is worth on its own. Read the right column downwards: the number people quote first is the number that decides least, and an assistant will quote it first because that is what everyone asks for.

It is the number everyone asks for because it is the number that sounds like a grade. Read the table's right column downwards and the ordering is almost the reverse of the way people quote them.

Rule A · winners		each worth about three units
Rule A · outcome		made money
Rule B · winners		each worth about half a unit
Rule B · outcome		lost money

Two rules with the same tally of winners and opposite results, which is why the first number decides nothing. The left pair wins less than half the time and makes money; the right pair wins more than half and loses it. Nothing but the size of the outcomes separates them.

Two rules, the same forty occurrences each, and opposite outcomes. The one that wins less than half the time makes money because its winners are larger than its losers; the one that

wins nearly two thirds of the time loses money for the mirror reason. Nothing but the size of the outcomes separates them, and the percentage of winners cannot see size at all.

The last row of the earlier table is the one that decides whether a rule is tradeable by a person rather than by a spreadsheet. A rule that makes money over forty trades but goes through eleven losses in a row is a rule almost nobody keeps trading, and abandoning it halfway is worse than never starting.

08

What the test cannot tell you

Suppose you froze the rule, checked the counts, and the numbers came out good. There are still four questions the whole exercise is silent on.

The question	Why the test is silent
Will it work next year?	the sample is the past, entirely
Is this luck?	forty trades cannot separate skill from noise
Would I have taken them all?	the record assumes you did
What about the costs?	only if you told it the spread and fees

Four questions a test cannot answer, however clean the count was. None of the four is a flaw in the arithmetic; each is a fact about what a record of the past is, and no amount of computing changes it.

The first is the one everyone knows and nobody acts on. A test is a record of the past, in its entirety, and the confidence it produces is about a period that has already finished. That is not a reason to skip it — an untested rule is worse — but it is a hard ceiling on what a good result means.

- Nothing** ● a dozen trades on one instrument
- A hint** ● forty trades, one instrument, one period
- Something** ● the same rule across three instruments
- A finding** ● and it survived a period you did not design it on

How much confidence a count actually buys, by how much of it there is. The top rung is where most people believe they are standing after an afternoon, and the bottom two are where they usually are.

The ladder is the practical version of the same point. Most people finish an afternoon believing they are on the top rung, and the arithmetic of small samples says they are on the first or second. Forty trades on one instrument in one period cannot separate a real pattern from an ordinary run of luck, and no assistant will volunteer that while it is writing you a summary.

The fourth row of the table is the cheapest to fix and the most often skipped. Costs are not a detail; on short-horizon rules they are frequently the entire difference between the two columns of the expectancy figure. An assistant will leave them out unless you hand it the spread and the fees, because it has no way to know them.

09

The three jobs that stay yours

Everything to this point has been work that can be handed over. Three things cannot, and the reason they cannot is the same reason in all three cases.

The decision	Why it cannot be handed over
How much to risk	it needs your account, which it has never seen
Whether to take this one	it needs what else you are holding
When to stop trading the rule	it needs what you can bear, not what is optimal

The three decisions that stay with you, and the reason each one stays. Notice that the reason is never that the machine is bad at arithmetic. It is that each of these needs a fact about you, and there is no fact about you anywhere in the conversation.


Read the right column. Not one of these says the machine is bad at arithmetic. Every one says the calculation needs a fact about you — your balance, your open positions, what a losing run does to your sleep — and there is no fact about you anywhere in the conversation unless you put it there.

That was the second blind spot in the first paper, and it is the one with a price attached. An assistant asked to size a position will size it, immediately and plausibly, from a balance you mentioned once or from a number it absorbed from an article, and the answer will look exactly like an answer computed from your account.

10

Sizing, which is arithmetic you still do

Position sizing is the sharpest illustration in the paper, because it is pure arithmetic and it still does not leave your side of the table.



Position size, as a sequence you run yourself. Every step is arithmetic an assistant could do in a second, and every step needs a number it has no way to know, which is exactly why the sequence stays on your side of the table.

Five steps, none of them hard, and each one needs a number that exists only in your account. The machine could run the sequence faster than you can read it and has access to none of the inputs.

What it answers	What it assumed
Risk one to two percent	a figure from an article, not your account
About this many units	a balance you mentioned once, in passing
You could size up here	that this setup is better, which nothing showed
That is a reasonable risk	that your other positions do not exist

What an assistant will produce if you ask it to size a position anyway. The left column is what it says; the right is the assumption underneath, which it will not state unless asked directly.


That table is what comes back when you ask anyway. The left column is the answer; the right is the assumption it rests on, and the assumption is never stated unless you ask for it directly — which, by now, is the recurring shape of everything in these three papers.

The fourth row is the one to watch. Sizing a position as though your other positions do not exist is the failure that turns three reasonable trades into one concentrated bet, and it is invisible in any single answer because each answer was reasonable on its own.

11

Living with a rule that is running

A written, counted, frozen rule now has to survive contact with a working week, and this is where the discipline usually slips.



Using an assistant while a rule is actually running, without letting it near the decisions. The shape is deliberately dull: it reads and counts, and you decide, and the order never reverses.

The shape is deliberately dull. The assistant reads and counts; you decide and size; the order never reverses. Nothing in those five steps asks the machine for an opinion, and that is the entire design.

Say	Instead of
Does this bar meet clause three?	should I take this trade?
What did the rule say about exits?	should I hold it longer?
Count the times this happened before	is this one different?
Read me the rule as written	what do you think now?

Four sentences worth keeping ready while the rule runs, and the request each one refuses. The refusals matter more than the questions: every one of them would have been answered, fluently, if it had been asked.

The right column of that table is the more instructive one, because every question in it would have been answered. Fluently, in a paragraph, with reasons. That is what makes the drift so easy: the wrong question does not

fail, it succeeds, and it hands you something that feels like counsel.

The last row is the one to reach for on a bad afternoon. Read me the rule as written is a request the machine cannot get wrong, and it is usually the only thing that needs to be said.

12


A session that works, end to end

Put together, the whole thing is about an hour, and it produces one rule with a count behind it and a written condition for giving it up.



One complete pass, from an idea in the morning to a written rule with a count behind it. It takes about an hour, and the first two steps are the ones that decide whether the other five were worth doing.

The first two steps decide whether the other five were worth doing. An idea stated vaguely and then interrogated for its undecided clauses turns into a rule; an idea handed over with instructions to make it profitable turns into a paragraph.



The other half of the same discipline, written before you need it. A rule with no stated end runs until something forces it to stop, and by then the decision has been made by an account balance rather than by you.

And a rule with no named end is not finished. Without the three numbers in that figure it runs until an account balance forces it to stop, which means the decision was made by the losses rather than by you, at the point where you were least able to make it well.

Write those three next to the rule, in the same file, on the same day. They cost nothing to decide while you are calm and cannot be decided at all while you are not.

13

Where each piece already has a paper

None of the underlying ideas here is new, and none of them was invented for use with an assistant.

- The rule** ● what a testable plan looks like — The Plan That Can Be Tested
- The test** ● what a result must show — What a Claim Has to Show
- The risk** ● the arithmetic of survival — The Account That Doubled
- The counts** ● the four analysis jobs — The Analysis That Is Really Arithmetic

Where each half of this paper has a longer treatment elsewhere in the library. This one is about the division of labour between you and an assistant; those are about whether the work is worth doing at all.

Each rung is a paper that spends twenty pages on whether the thing is worth doing and what it means when it is done. This paper only says which half of the labour can be delegated, and that is a narrower claim on purpose.

Hand over	Keep
Naming what is on the chart	deciding what it means
Counting anything countable	supplying the rule that defines the count
Writing the rule out in full	choosing which rule to trade
Tallying what happened after	deciding whether the tally is enough
Reading the rule back to you	how much to risk, and when to stop

The three papers of this collection in one table, which is also the summary of all three. Two columns, and the division between them has not moved once across sixty pages.

That table is the summary of the collection. Two columns, and the line between them has not moved once across three papers: naming and counting on the left, deciding and bearing on the right.

It is worth noticing what the left column is worth. It removes most of the tedium from a process that people abandon because of the tedium, and that is a real gain. It just is not the same thing as an edge, and nothing on the left column will ever become one.

14

What this does not solve

Three limits, and then the short version of all three papers.

The limit	What it means here
A testable rule is not a good rule	most well-written rules lose money
A count is not a forecast	the past is the only thing that was measured
Faster iteration is not better research	twenty rules an hour is twenty ways to fit

Three limits, stated plainly, in the spirit of the rest of the library. The first is the one that undoes the most optimism, and it applies just as hard to a rule written flawlessly.

The first is the one that undoes the most optimism. A rule you wrote flawlessly, counted honestly and froze properly still, most of the time, loses money. Testability belongs to the sentence rather than to the market, and getting the sentence right only means you now know which of the two you have.

The third limit is the newest and the least discussed. An assistant lets you write and test twenty rules in the time one used to take, and twenty attempts against the same sample is twenty chances to find something that fits it by accident. Speed makes the fitting problem worse rather than better, and the only defence is the same freezing step, applied every time.

Term	As used here
Clause	one decided part of a rule, stated as a price or a count.
Occurrence	a bar at which the rule's entry condition was met.
Frozen	written down before the count ran, and not edited after.
Fitting	changing a clause after seeing what the outcome would be.
Trial	an agreed number of trades, with an end named in advance.

The terms this paper uses in a fixed sense, carried forward from the first two and ending here.

The short version of the collection is three sentences. The assistant turns your chart into a description, so everything it tells you is downstream of what was inside the frame and nothing else. Every analysis job it does well is a count, every count needs a rule that comes from you, and an interpretation arrives attached whether you asked for one or not. And the work of writing rules and counting outcomes can be handed over entirely, while deciding what to risk, what to take and when to stop cannot be handed over at all — not because the machine is weak at them, but because it has never seen your account and will answer anyway.

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