


KUTTY STAUDER



The Spinozan Trader

A Geometric Critique of Strategies, Signals, and Evidence

EXPANDED REFERENCE EDITION · SEPTEMBER 2026 · SAMPLE

OPENING SAMPLE

The Spinozan Trader

A Geometric Critique of Strategies, Signals, and Evidence

Definitions, conditional proofs, counterexamples, historical research, and reproducible tests. Mathematical properties do not establish market profitability; evidence is labeled by its actual scope.

Kutty Stauder


Expanded reference edition · September 2026

Copyright © 2026 Kutty Stauder.

All rights reserved.

Philosophical inspiration: Baruch Spinoza. Contemporary prose and applications are by the author, not quotations from Spinoza unless expressly identified.

kuttystauder.com/books



PART I

Foundations

How to Use This Book

A screen can contain twenty indicators and only one idea. The lines have different colors, the names come from different decades, and the trader feels surrounded by confirmation. Yet each line may be another transformation of the same recent prices. The screen has multiplied the appearance of evidence faster than it has multiplied evidence itself.

This book asks a harder question than whether an indicator is good. Good for what decision, in which market, over what horizon, using what information, at what cost, and compared with what alternative? A moving average can reduce the complexity of a decision without predicting a return. A volatility estimate can improve position sizing without forecasting direction. An execution algorithm can save money while having no view of tomorrow's price. A pattern can predict a conditional distribution and still fail to pay its spread.

The geometric method gives these distinctions a durable form. Definitions prevent an attractive word from doing the work of a calculation. Assumptions disclose the world in which an argument holds. Propositions state results narrowly enough to be demonstrated. Counterexamples expose conclusions that outrun their premises. Historical studies then confront the premises with a world that does not owe them continued truth.

Choose a reading route

On a first reading, take the method of proof, definitions, expectancy, sizing, and information together. Then read The Calculus of Usefulness. The line-by-line proof example and symbol guide explain the notation without assuming prior mathematical training. You can follow the prose and worked examples first, returning to a demonstration when you want to inspect its assumptions.

Choose the route that matches the decision you want to make:

Your question	Read next	Produce before moving on
How do I specify my first systematic rule?	<u>Where edge comes from</u> , <u>operation</u> , and if the purchase is a challenge <u>prop firms</u>	A completed strategy worksheet that names the supposed payer, and a paper ledger
Where does a trading edge come from?	<u>Edge chapter</u> , <u>usefulness</u> , <u>market making</u>	Which of the named families, if any, your rule could access
How do prop-firm challenges work?	<u>Prop firms</u> , <u>risk</u> , <u>livelihood</u>	A written knockout sheet; passing is not expectancy
Which popular systems are being sold to me?	<u>Atlas</u> , the linked proofs, <u>laboratory results</u>	Each system's actual object: recoding, path constraint, or untested lore
Can I just watch candles and be successful over time?	<u>Watching candles</u> , <u>livelihood</u>	Which study, sample, exit, and cost you are borrowing; none is a discretionary livelihood
What is anchored VWAP, exactly?	<u>The anchored VWAP</u> , <u>execution</u>	A written object, volume, anchor clock, and role: benchmark, description, or trigger

Your question	Read next	Produce before moving on
Can I algo-trade like Jim Simons / Renaissance?	<u>Algorithmic trading and its myths, edge</u>	Which Renaissance book; Medallion is not RIEF and not a 13F copy
Is the COT report useful? Intraday or swing?	<u>Commitments of Traders, futures</u>	Tuesday snapshot, Friday 15:30; not an intraday increment; swing only after the print
Which tools on my chart do useful work?	<u>Atlas</u> , the relevant indicator chapter, then <u>execution</u>	Each tool's input, decision role, and a comparison with it removed
Does a backtest establish an edge?	<u>Data, inference, laboratory protocol, results</u>	A list of what was tested, what was selected, and what remains uncertain
Why can a “losing” rule still grow money, or a “winning” table still fail?	<u>Murky backtesting, monthly results, opening-drive, random controls</u>	Which meaning of <i>lose</i> you are using: expectancy, benchmark, path, or convention
How does the intraday ORB research translate into a test?	<u>Closing ORB study</u> , then <u>operation</u>	Separate specifications for stocks and futures; no transfer of return claims
How did the monthly SPY/cash filter behave?	<u>Monthly construction, results</u>	An investment-contrast worksheet; not the book's first trading method
How should I revise an existing process?	<u>Blueprints, worked cases, exercises</u>	A dated change proposal and a rule for rejecting it
What did the Turtle traders actually do?	<u>Turtle chapter</u> , then <u>risk</u>	A written unit, stop, and correlation-cap sheet; no transfer of Faith's 80% figure
How do I write a risk plan?	<u>Risk management, sizing, operation</u>	A completed risk worksheet with heat, gap, and pause rules

Your question	Read next	Produce before moving on
I trade forex or futures. Where could a repeatable edge come from?	<u>Anatomy of a repeatable edge</u> , <u>instruments</u> , <u>edge chapter</u>	A named payer, a net margin after costs, and the contract arithmetic that decides whether the size is even available
Which forex and futures strategies have public evidence?	<u>Trend</u> , <u>carry</u> , <u>currencies</u> , <u>commodities</u> , <u>time and flow</u> , <u>public records</u>	Each claim's sample, cost treatment, and status: gross or net, in sample, out of sample, or live
How do I test a published strategy, and how would I run the six-month momentum rule?	<u>A paper examined</u> , with <u>momentum</u> and <u>inference</u>	A written specification, a cost model, the months needed to see the edge, and a retirement rule
What exactly would I trade, and when would I stop?	<u>Blueprints</u> , <u>operation</u> , <u>edge factory</u>	A written specification with a sizing formula, a cost limit, and a retirement rule fixed in advance

The historical survey supplies context. The indicator and strategy chapters are a reference: read the family you need rather than treating every formula as a prerequisite to a first experiment.

Keep the studies separate

Several kinds of evidence appear in this book. This map prevents a favorable number from becoming a claim about a different strategy.

Construction	Evidence in this edition	What the reader can conclude
Edge taxonomy	Geometric identities EG1–EG3; <u>edge chapter</u>	Recodings of H are not a distinct payer; a challenge fee is not an expectancy proof

Construction	Evidence in this edition	What the reader can conclude
Prop-firm evaluations	Dated official rules plus PF1–PF3; prop firms	Passing, underlying expectancy, and payout are different objects; no industry pass rate
Random-rule SPY controls	Frozen coin-flip and shuffled channel, same tape as the laboratory; results	These random policies lost the full sample to buy-and-hold; a slice can still beat a named rule
Candlestick / “price action” rules	Sourced tests from 1975–2012 Japan and U.S. names, 5-minute DJIA, Taiwan windows, and 1973–1994 FX head-and-shoulders; watching candles	No inspected study follows discretionary candle-watchers as a school; isolated favorable tables depend on the exit and do not establish a livelihood
Anchored VWAP	Geometry V4–V8, Berkowitz–Logue–Noser 1988, TradingView formula, Shannon’s public definition; complete reference	A declared average of a declared history; no inspected journal test of reclaim or fade as a livelihood
Algorithmic trading / “like Simons”	ADV four-family split, 13F rule, Zuckerman/Cornell table as journalism, 2020 Medallion-versus-outside-fund year; myths	Automation is not an edge; Medallion is not a retail experiment; RIEF is the same firm’s other book
Commitments of Traders	CFTC clock and classifications, N3–N6, Sanders–Irwin–Merrin and energy Granger tests; COT	A Tuesday book released Friday; not an intraday signal; inspected panels find little lead from positions to returns
Monthly SPY/cash filter	The book’s retrospective simulation; full protocol and results	An investment-contrast risk tradeoff, with substantial benchmark shortfalls; not the first trading method

Construction	Evidence in this edition	What the reader can conclude
Daily SPY indicator rules	Thirteen constructions under common assumptions; comparison	How these exact rules differ in exposure, return, drawdown, and costs
Stock ORB with daily RVOL ranking	An externally reported 2016–2023 study; stock section	A promising published claim requiring independent replication
Fixed-basket FAANG ORB	This edition's long/short 2016–2023 test and separate 2024–2026 extension; opening-range study	Outcomes for five retrospectively selected survivors under declared costs and data restrictions
Long/short NQ ORB	Two limited 2026 public-data replays; pilot	Exploratory outcomes with unresolved contract and execution limitations
Donchian-gated SPY opening-drive	This edition's daily analogue, 1994–July 2026; closing study	A failed declared-edge test; not a 5-minute stock ORB and not the Turtle system
Cross-sectional factors and portfolios	Defined research factors and allocation identities; factors , portfolios	Investment-contrast literature and mathematics, not a recommended trading book
Other strategy templates	Conditional mathematics and identified family literature; blueprints	Questions and specifications for research, not new backtested recommendations
Turtle / Donchian futures system	Faith's published rules plus later trend-following papers; Turtle chapter	A complete-system specification and long-horizon family evidence, not a reconstructed Turtle book
Risk plan	Geometric identities and worksheets; risk management	Explicit heat, caps, and gap rules; not a source of expected profit

Construction	Evidence in this edition	What the reader can conclude
Forex and futures strategy families	Sourced public studies, with no backtest run by this edition; <u>anatomy of an edge</u> to <u>blueprints</u>	Every figure belongs to someone else's sample, labelled gross or net of costs and in sample or live; the blueprints are specifications rather than tested results
Carry and insurance-like premia	The cross-asset carry literature, its drawdown record, and dated crashes; CY1–CY4, <u>carry</u>	A real premium with a measured negative skew; a smooth sample is evidence about the states it did not contain
Contests and long public records	Operator-published standings and managers' own factsheets; PR1–PR2, <u>public records</u>	Self-reported and unaudited, sometimes restated across different accounts and leverage levels

The daily SPY laboratory uses February 1994 onward; the monthly study has its own start and initialization. Compare each strategy with the benchmarks calculated inside its own experiment. Dollar profit from one fixed futures contract, a stock portfolio's cumulative percentage return, and a fund strategy's annual growth rate are not a common ranking scale.

The evidence key

A **definition** is a convention used in this edition. Different software can implement the same named indicator differently. The edition states look-back, smoothing, seed, equality rule, and timing where these affect a result.

A **proposition** is an original demonstration in this book unless a source is explicitly credited. The first twenty-nine propositions retain their numbering from the shorter edition. Later propositions have stable family identifiers, such as F1 for filters. A label is a way to find an argument, not a claim that the author has discovered a previously unknown theorem.

An **empirical finding** reports an identified study or this edition's own calculation. Dates, market, comparison, and limitations travel with the finding. Results from a diversified futures study are not silently transferred to a single stock ETF. An intraday effect is not treated as evidence for a monthly rule.

A **counterexample** disproves a universal statement by constructing a permitted case in which it fails. It does not establish the frequency of that case in actual markets. One steadily rising price path disproves “RSI above 70 must be followed by a decline.” It does not establish that an RSI-based strategy always loses.

A **working hypothesis** is a proposition about markets that remains to be evaluated. A **usefulness verdict** identifies which roles are supported by the argument, which need empirical evidence, and which stronger claims have failed. The verdict can be favorable about measurement and unfavorable about a promotional promise at the same time.

Terms to keep beside the results

Term	Meaning in use
Expected value or expectancy	Mean payoff under a stated probability model; a sample average estimates it
Net profit	Profit after the costs included in that calculation; check which costs were omitted
CAGR	The constant annual compound rate connecting starting and ending wealth, under the stated cash-flow convention
Drawdown	Decline from a previous equity peak; daily-close measurement can miss intraday losses
Basis point (bp)	0.01 percentage point; 5 bp of \$10,000 is \$5
One risk unit (R)	The strategy's defined initial risk amount; it is not a guarantee of maximum realized loss

Term	Meaning in use
Out-of-sample evaluation	Evaluation on observations excluded from the specified fitting process; a repeatedly inspected holdout is no longer fresh
Trading edge	A claimed favorable net decision advantage against a stated alternative, requiring evidence beyond the indicator formula

Scope and date

The historical spine begins in the 1960s and follows the principal families into contemporary electronic and data-intensive trading. Earlier charting traditions appear where they explain an ancestor. This is a broad reference, not an enumeration of every proprietary strategy, every local-market convention, or every indicator variant. Contemporary source notes are dated to the September 2026 edition; exchange rules and product details must be rechecked when a strategy is implemented.

The monthly SPY study, daily indicator laboratory, random-rule controls, and closing ORB study each retain their own assumptions and evidence status. The forex and futures part, chapters 45 to 56, is a review of publicly documented strategy families and contract mechanics rather than an experiment of this edition's own: it reports other people's samples and states that no backtest was run here, so its blueprints are specifications to be tested rather than results to be trusted. The monthly filter and the factor and portfolio chapters are investment contrasts: they are retained as honest evidence, not as the first trading route. The [downloadable companion](#) contains protocols, code, tests, results, and input hashes. Reproduction requires the documented market inputs; the companion does not redistribute raw vendor data.

A reader's first experiment

Take one proposed trade. Write five sentences before looking for another indicator:

1. The instrument and exposure I would hold are these.
2. The information available before my order is this.
3. The mechanism that could pay me is this.
4. The observation that would make me reject the mechanism is this.
5. After costs and adverse execution, the loss I can survive is this.

If the third sentence is merely the name of an indicator, the argument has not begun. Name the family in the edge chapter, or admit that none applies. If the fourth sentence is impossible to write, the idea may be protected from evidence. If the fifth sentence depends on a fill the market has not promised, the account is carrying an assumption the trader has not acknowledged.

Carry those five sentences into the strategy worksheet when you are ready to make the experiment executable. A good trading book should leave the reader harder to deceive, including by the book itself.

The Method of Proof

The trader enters a world of numbers and too often leaves it with a superstition. A pattern becomes a promise. A winning month becomes a theory. A loss becomes an insult. The account records a result, but the mind supplies a story about what the result must mean.

The geometric method begins elsewhere. It asks what the terms mean, which premises are being granted, and what follows from them. It does not reward a conclusion for being forceful. The trader who writes “therefore” must show the bridge.

This book uses the architecture familiar from Euclid and Spinoza: definitions, common assumptions, propositions, demonstrations, corollaries, and explanatory notes. It borrows a discipline of reasoning, not the subject matter of Euclidean geometry. A market is not a triangle, and an empirical premise does not become self-evident because it has been called an axiom.

We will establish truths about the arithmetic of trading: expectancy, costs, compounding, drawdowns, position sizing, information, and the conditions of long-run growth. These results are conditional on explicit models. We will then construct a particular method and examine its historical record. The distinction between a theorem and a tested hypothesis is the spine of the book.

Three kinds of statement

A **definition** fixes how a term is used. Defining an edge as positive expected profit does not demonstrate that a trader possesses one. Defining a trend as a price above an average does not demonstrate that the price will continue upward.

A **deduction** establishes a conclusion from premises. If a position loses half its value, it must double to recover. Nothing about tomorrow's market is needed to prove that statement. If a proposed argument requires tomorrow's distribution, that distribution must appear among its premises.

An **empirical claim** concerns what occurred or is likely to occur in the world. Whether a trend filter has reduced drawdowns after realistic costs is such a claim. Data, execution conventions, alternative explanations, and uncertainty belong beside it. A historical result must carry its dates.

The error we are trying to remove is the quiet substitution of one category for another. "I can calculate the optimal stake if I know the odds" becomes "I know the optimal stake." "This rule succeeded in these observations" becomes "this rule works." "Every event has causes" becomes "I can predict the next event." Each substitution removes the most difficult premise from view.

Read a proof as a short argument

You do not need to know geometry to begin. The Euclidean method used here means stating the starting points and showing each step to a conclusion. It is a way to organize an argument. The word proof does not make an assumption about markets true.

Part of the argument	Plain-language question
Definition	What exactly does this word or symbol mean here?
Assumption or postulate	What are we granting for the argument, and could it fail in practice?
Proposition	What precise statement are we trying to establish?
Proof or demonstration	How does each step follow from the starting points or an earlier result?
Corollary	What further conclusion follows directly?
Scholium	What explanatory note helps me interpret or use the result?

Part of the argument	Plain-language question
Counterexample	Can one allowed case show that an “always” claim is false?
Q.E.D.	The argument is finished; the letters abbreviate “which was to be demonstrated”

A proof translated line by line

Take the first proposition in the next group: net profit equals gross profit less costs. The letters are temporary names for amounts. Let G mean gross profit, C mean costs, and X mean net profit, all measured in dollars.

$X = G - C$ reads: “Net profit equals gross profit minus costs.” If gross profit is \$100 and costs are \$7, net profit is \$93. The letters let the same statement work for other amounts without rewriting the sentence.

Now consider possible trades rather than one completed trade. $E[X]$ reads “the expected value of X ”: the probability-weighted mean of its possible outcomes under the assumed model. It is not a promise about the next trade, and the letter E is not another amount being multiplied by X .

$E[X] = E[G] - E[C]$ reads: “Expected net profit equals expected gross profit minus expected costs.” Taking a probability-weighted average preserves subtraction, provided the averages exist. That is the demonstration. It remains valid when profitable trades and costly trades tend to occur together; independence is not required for this accounting identity.

What has been proved? How the quantities relate. What remains unknown? The actual future distribution of gross profits and costs. The formula cannot supply those inputs. This separation is the habit to carry into the rest of the book.

Follow the labels without memorizing them

“Proposition 3” is an address, not the third step of a trading system. “Depends on D_4 , A_2 ” sends you to Definition 4 and Assumption 2 in the foundational definitions. Read those starting points before deciding

whether the conclusion applies to your account.

Later propositions use family labels: U1 is the first usefulness proposition, F1 the first filter proposition, and ORB1 the first opening-range proposition. Their numbers identify arguments, not confidence scores or recommended parameter values. The [linked index](#) lists every proposition.

Some letters recur in different chapters. “Definition D1” and “Proposition D1” are different addresses; the latter belongs to the data family. The same distinction applies to an assumption labeled A1 and a factor proposition labeled A1. Local headings and the index resolve the meaning. A range such as D2–D3 means both definitions, not a subtraction.

When a proof uses unfamiliar mathematics, first read its proposition, assumptions, example, and explanatory note. The [symbol guide](#) will help you return to the formal steps. Skipping those steps is a reading choice, not independent verification of the proof.

The Spinozan connection

Spinoza distinguishes activity from passivity through adequate causation and adequate ideas (Ethics III, definitions 1–3 and proposition 1). He describes conatus as a thing’s striving to persevere (III, propositions 6–7). Those concepts suggest a useful orientation: preserve the capacity to act, examine the causes of your conduct, and do not confuse the intensity of an affect with evidence.

They do not establish that prices have a biological desire to persist, that every frightened person makes a bad trade, or that a trading loss proves inadequate understanding. A well-specified strategy can lose on an outcome its model allows. A poor strategy can win by chance. The mathematical distinctions below prevent the philosophy from becoming another instrument of self-deception.

A statistical regularity can be useful before its entire causal mechanism is known. A compelling causal story can be useless if the relevant effect is already reflected in price. The standard is explanatory honesty joined to observable performance, not a requirement that a finite trader comprehend the whole market.

What this edition promises

Forty-five chapters and one hundred thirty-nine propositions examine the principal trading families from the 1960s to the present. The original twenty-nine propositions establish the mathematical foundation; later demonstrations analyze indicators, strategies, execution, and inference. A fully specified equity/cash rule and a comparative indicator laboratory provide reproducible historical evidence, including costs and adverse comparisons. The atlas, worked cases, strategy blueprints, and operating protocol turn the arguments into a method of evaluation.

The aim is not a method that cannot lose. Proposition 23 explains why history cannot supply that guarantee. The aim is a method clear enough to execute, limited enough to survive mistakes, and explicit enough to be rejected when it fails its declared purpose.

To read geometrically is to ask of each page: What is defined? What is assumed? What has actually been shown? This is the first discipline of the Spinozan trader.