The Measurement Canon

Sixteen instruments for reading a trading system as a population of positions: the question each answers, the method beneath it, and the doubt that remains

GFE-PUB-CAN-20261112-A · 12 November 2026
Dr. Glen Brown · Global Financial Engineering, Inc.

I. A system is a population, not a stream

The measures in common use for judging a trading operation read it as a single stream of profit and loss. Return, volatility, the Sharpe ratio, maximum drawdown and the share of winning trades all take the account's equity, or its list of closed trades, as one series and summarise it. They answer a real question, which is how the account has fared, and they answer it well. They do not answer a different question, which is whether the machinery producing those results is sound, and why.

Many automated systems are not well described as a stream. They open many positions, each with its own stop and its own profit target, each living through states the system itself defines, and each ending in one of a few ways: it reaches its target, it reaches its stop, or it is closed near where it began. Such a system is a population, and the questions that decide whether it is sound are questions about that population. Where do its losses come from? How much of its stop do its winners use? How long does a deep position take to be decided? How much of its floating loss was inherited from rules it has since retired? How much of its intended size does its venue allow it to express? None of these can be asked of a single stream.

This paper sets out sixteen measures for asking them. Together they form what we call the Measurement Canon. They were developed for, and are computed live by, the proprietary system of the house; this public edition publishes the instruments themselves, stated generally, without the system's parameters and without any account's results.

II. Four principles

The unit belongs to the system. Every measure is expressed in a unit the system itself defines: R, the distance from a position's entry to its stop, valued in money; the bank, the gain a position takes at its target; and the day of opening. A measure stated in the system's own unit can be compared with the system's own design. A measure stated in percentage return cannot.

The method is borrowed; the question is original. Almost every measure stands on an established method from another discipline: survival analysis from medicine, cohort analysis from actuarial and credit practice, the adverse-excursion studies of the trading

floor, the bootstrap, reserving from insurance. What is new is the question each puts to a trading system and the unit in which it puts it. The borrowed method is what makes the new measure trustworthy, and each entry names it.

Every measure states its doubt. No measure removes uncertainty. Each converts guesswork into stated uncertainty, and each entry closes on what remains uncertain: the size of the sample, the censoring that young evidence suffers, the gap between what the system observed and what it could not. A reading is acted upon only within its uncertainty.

Measures inform; they do not act. Every measure is evidence only. None grants authority to close a position, none treats the passage of time as a reason to act, and none is read across accounts: each account is judged on its own evidence, and another account may be cited only as reference. A measure that is allowed to act is no longer an instrument; it is a rule, and it must then be designed and tested as one.

III. The vocabulary

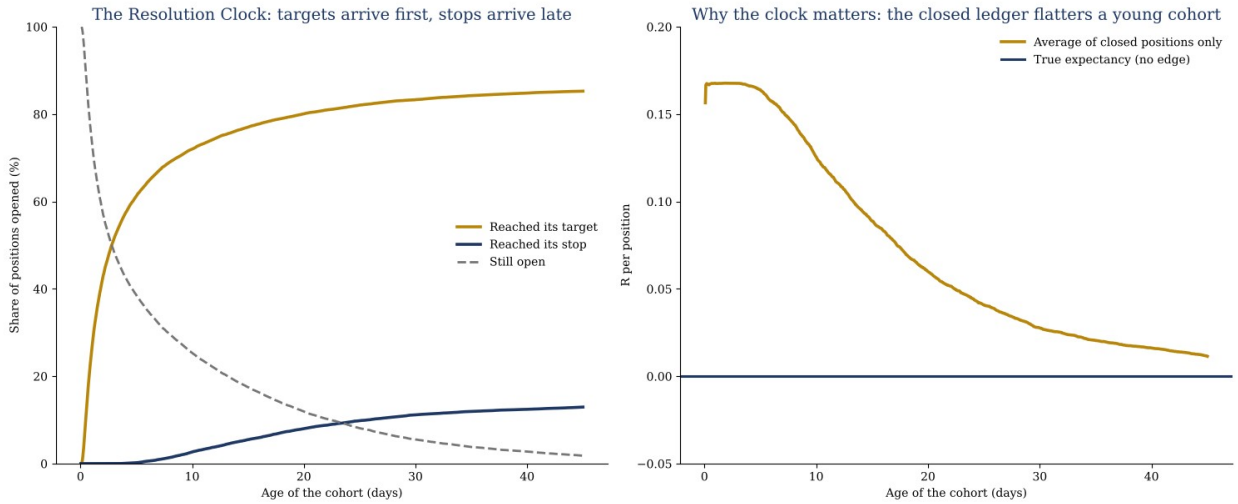
A **position** is opened with a stop and a profit target. When it reaches its target it **banks**; when it reaches its stop it is **stopped**; it may also be closed near its entry. Its **depth** is its deepest excursion against entry, in R. A **vintage** is the set of positions opened on one trading day. The **open mark** is the current unrealised result of the positions not yet resolved. The **current rules** are the rules in force today; positions opened under earlier rules form the **inheritance**. Where a venue refuses the full intended size of an order, a system may complete the position with further orders, and the group of orders is a **batch**.

The design to which these measures are addressed is a common one: a modest target close to entry, and a stop much further away. Such designs win often. Their losses are rarer, larger and slower to arrive, and that is the source of most of the measurement problems this paper addresses. Figure 1 illustrates the point with a simulated system that has no edge at all.

IV. Why the clock comes first

In the simulation the positions have, by construction, no edge. Yet for the first week almost every closed position is a winner, and the average of closed positions stands near its maximum. The losses arrive later, one at a time, and the closed average falls toward its true value only as the cohort matures. A system judged on its closed trades in its first weeks will look profitable whether it is or not.

Figure 1. A system with no edge, read by age



Simulated driftless price paths, 20,000 positions opened together, each with a target at 0.15 R and a stop at 1 R. Left: cumulative incidence of banks, stops and positions still open by age. Right: the average result of the positions closed so far, against the true expectancy of zero. Illustrative simulation; it describes no account.

This is why the Canon begins with the Resolution Clock and the Vintage Resolution Table. They do not judge the system; they say when the system can be judged, and which of its vintages are still too young to say anything. The companion paper The Martingale Ledger gives the estimator that removes the bias from the first day; the measures that follow ask what lies beneath it.

V. The sixteen measures

Each measure is set out in five parts: the question it exists to answer, the established method on which it stands, the exact definition of what is computed, a table of the readings it can give and what each means, and what remains uncertain. They are numbered in the order in which they were defined. By the question they answer, they fall into five families.

Family	Measures
When positions are decided	I, II, IV, XIII
Where losses come from	III, V, X, XI
Whether the stop is right	IX, XII (read only when XIII allows)
What the open book owes	VI, VII, VIII, XVI
What the venue allows	XIV, XV

Measure I. The Vintage Resolution Table

The question. Which days of opening are still undecided, and what each day of opening has earned or still owes.

The established method. Vintage, or cohort, analysis from actuarial and credit practice, in which each year of underwriting or each month of lending is followed to maturity on its own, so that a young book is never mistaken for a profitable one.

Definition. Positions grouped by the day on which they were opened. For each vintage: the number opened, the numbers that banked, were stopped or were closed near entry, the number still open, the realised result in R, the open mark in R, the share resolved, the median time to bank and the vintage's age.

Reading	What it means
A vintage resolved above about 85% with a positive realised result	The vintage has done its work; what remains open is its tail.
A young vintage with a low resolved share	Not yet readable; its open mark is a floor, not a result.
An old vintage still carrying a large open mark	The book's debt is concentrated there; read its concentration (Measure VIII) and its depth classes (Measure III).

What remains uncertain. Open marks are taken at a moment and can understate the depth a position has reached. A vintage of a handful of positions is an anecdote, not a reading.

Measure II. The Resolution Clock

The question. How fast the system converts openings into banks, and when its stops arrive.

The established method. Survival analysis with competing risks, and in particular the cumulative incidence functions of the Aalen–Johansen estimator, developed in medical statistics for patients who may leave a study by one of several causes. Here a position leaves by its target, by its stop, or by a close near entry.

Definition. Of the positions at least h old, the share that has banked and the share that has been stopped by age h, read at one hour, four hours, twelve hours, one day, three days and ten days. Only positions old enough to have reached a horizon are counted at it, so no open position is treated as a winner or as a loser. A horizon that no position has yet reached is reported as unreached, never as zero.

Reading	What it means
The banked share rising steeply within the first day	The near target is doing its work; the system turns its openings quickly.
Stops appearing only at three days and beyond	Losses are slow; they are governed by deep excursions rather than by noise, and the closed ledger of a young book will flatter it.
The banked share flattening while the open share stays high at ten days	A tail of long-lived positions is forming; read the Line of Fate and the cost of carrying them.

What remains uncertain. The longest horizons rest on few positions in a young book. The clock describes; it never acts. No position is closed because of its age.

Measure III. The Line of Fate

The question. Where the system's losses come from.

The established method. Stratified outcome analysis by a state variable, the reasoning a clinician applies to a severity score: patients are grouped by the severity they reached, and outcomes are read within each group.

Definition. Positions divided by their deepest excursion against entry, in R: shallow (never a quarter of R), middle (a quarter to half of R) and deep (half of R or more). For each class: banks, stops, other closes and positions still open.

Reading	What it means
No stops among shallow positions	The line holds: the system's loss risk lives in its deep excursions.
A stop in the shallow class	The line is breached and that loss must be explained: a price gap through the stop, a close made while the system was not watching, or a depth the record failed to see.
A rising stopped share among resolved deep positions	The deep class is turning against the book; read its vintages and the market factor they share.

What remains uncertain. A position's class is fixed by the deepest point it has reached so far; an open position can still move to a deeper class.

Measure IV. The Deep Resolution Horizon

The question. How long a deep position takes to be decided.

The established method. Time-to-event analysis on a conditional population, the reasoning behind the expected remaining life of a patient who has already reached a given stage.

Definition. For positions that reached half of R against entry, the time from opening to resolution, read separately for those that went on to bank and those that were stopped.

Reading	What it means
Deep resolutions measured in weeks	A deep position is undecided for weeks; its mark is not its fate.
Deep positions younger than the horizon	Nothing about them is decided yet; any value placed on them is provisional.

What remains uncertain. Read from few resolutions in a young system, and itself subject to the censoring it measures.

Measure V. Heat to Bank

The question. How far a system's winners go against themselves before they bank.

The established method. Maximum Adverse Excursion, a trade-study measure in use for decades, here expressed in the system's own unit of gain: the bank.

Definition. For each banked position, its deepest excursion against entry divided by its own bank, read as the median, the ninetieth percentile, and the shares beyond one and five banks.

Reading	What it means
Many winners more than one bank under water before they bank	The stop is earning its width: a narrow stop would have killed them.
Winners rarely under water	The book’s instruments or its youth keep winners shallow; the stop’s width is protecting something else.

What remains uncertain. A property of the instruments traded and of the book’s age: winners that bank only after a long excursion appear late.

Measure VI. Open Book Fair Value

The question. What the open book is worth, as distinct from what it is marked at.

The established method. Conditional expectation by state, the reasoning of insurance reserving: an open claim is valued at what claims of its kind have historically cost, not at its last payment.

Definition. Each open position valued at the historical payoff of positions of its depth class, summed for the book and added to the realised result to give an expectancy per position on a fair-value basis.

Reading	What it means
Fair value well above the mark	The mark overstates the debt; the book is closer to break-even than it looks.
Fair value near the mark	The open positions are expected to resolve roughly where they stand.

What remains uncertain. Wide where the deep class has few resolutions; it is always stated with its range and never alone.

Measure VII. Recovery Debt Days

The question. How long the system’s own banking would take to pay its open mark.

The established method. The payback period, the oldest measure in capital budgeting, turned on the system’s own flow of gains.

Definition. The open mark divided by the realised banking rate under the current rules, where the rate is measured from the first opening under those rules, in R and in money.

Reading	What it means
A debt of a day or less	The open book is small against the system’s earning power.
A debt of a week or more	The open book is large against the flow; read where it sits by vintage and by depth.

What remains uncertain. The banking rate is an average over the life of the current rules and moves with it; measured over too long a window it overstates the debt.

Measure VIII. The Vintage Concentration Index

The question. Whether a day's openings were one bet written as many positions.

The established method. Concentration measurement, of the same family as the Herfindahl–Hirschman index of market structure, reduced here to the share held on one side.

Definition. For each day of opening, the share of its positions on its larger side, read on the open book and naming the day that holds most of it.

Reading	What it means
An open vintage of several positions, three quarters or more on one side	The open book is one factor bet: its fate is shared, and its depth will move together.
Balanced vintages	Independent bets; their resolution is spread.

What remains uncertain. A mechanism rather than a law: a concentrated vintage may reflect the session in which positions were opened as much as the system's design.

Measure IX. The Stop Frontier

The question. Whether a different stop distance would make the system better.

The established method. Counterfactual replay on the book's own price paths, with bootstrap intervals; the reasoning of an efficient frontier, read along a single dial.

Definition. Expectancy per position at k times today's stop, with the bank and the money at risk held fixed. Narrower stops can be computed exactly from recorded excursions, open positions included. Wider stops are bounded between the rescues actually observed and the most favourable case.

Reading	What it means
A plateau around today's stop, falling on both sides	The stop is where it should be; widening has not earned its case.
The wide side rising on measured evidence	Stopped positions are being rescued beyond the stop, and the question of widening reopens.

What remains uncertain. Read only on matured vintages (Measure XIII); where a venue limits position size, read in money rather than in R.

Measure X. Stop Utilisation

The question. How much of the stop distance winners actually use.

The established method. Order statistics: the quantiles of the winners' deepest excursion in R.

Definition. The median, ninetieth and ninety-ninth percentiles and maximum of the banked positions' deepest excursion, in R.

Reading	What it means
The ninety-ninth percentile well inside the stop	The outer part of the stop is not where winners live; it is where deep positions are decided.
A bank recorded near the stop	The stop is protecting winners at its edge; any narrowing would cost them.

What remains uncertain. The maximum is a single position and moves with one event; the high quantiles are read together.

Measure XI. The Afterlife Depth

The question. What a stopped position's price did after the stop.

The established method. Post-event path analysis on the minute record of each instrument.

Definition. For each stopped position, the deepest point its price reached after the stop and before any return to the level of its bank.

Reading	What it means
Stopped positions going far beyond the stop without returning	A wider stop would only have made those losses dearer.
Stopped positions returning to the bank soon after the stop	A wider stop would have rescued them; count them against the Rescue Threshold.

What remains uncertain. Where the price record does not cover a stopped position, it cannot be followed, and it is reported as such.

Measure XII. The Rescue Threshold

The question. How many stopped positions a wider stop must rescue to pay for itself.

The established method. Break-even arithmetic on the book's own frontier.

Definition. The number of stopped positions that would have had to bank before reaching k times the stop for expectancy at k to equal today's.

Reading	What it means
Rescues observed below the threshold	Widening stays closed on evidence.
Rescues reaching the threshold	The question of widening reopens by itself.

What remains uncertain. Depends on the open book at the moment it is read, and is restated with each reading.

Measure XIII. Frontier Maturity

The question. When a frontier may be read at all.

The established method. Right-censoring, the central problem of survival analysis: an observation that has not yet ended cannot be counted as ended.

Definition. A frontier is read only on vintages older than the Deep Resolution Horizon, because a young vintage flatters both widening and narrowing.

Reading	What it means
A young vintage leaning one way	Not a reading; wait for the vintage to mature.
Mature vintages agreeing	The frontier may be read.

What remains uncertain. The horizon is itself measured, and moves with the evidence.

Measure XIV. The Venue Expression Ratio

The question. How much of its intended size the venue allows the book to express.

The established method. Fill-ratio and capacity measurement from execution analysis, in the tradition of implementation shortfall.

Definition. At each witnessed cap on order size, counted once per instrument per trading day, the share of the requested volume that the venue accepted, read as a mean over witnessed caps and as the number below a quarter.

Reading	What it means
High expression	The venue is not the book’s constraint.
Low expression in a class of instruments	The venue cannot carry the book’s size there: additional orders restore what they can, and the rest is refused.

What remains uncertain. Counts witnessed caps, not every order; it is a map of the venue, not a measure of the system.

Measure XV. The Batch Multiplier

The question. What additional orders, placed where a venue caps size, add to the money of the positions they extend.

The established method. Attribution analysis: the money of the added orders set against the money of the positions they complete.

Definition. One plus the money of the closed additional orders divided by the money of the closed positions entered at the venue’s maximum size.

Reading	What it means
Above one on winning batches	The batch is restoring size the venue removed, and at better prices.

Reading	What it means
Below one where batches lose	The batch carries the full intended size into a losing position as well, as it was designed to.

What remains uncertain. Mixes winning and losing batches in one figure; it is read beside the Line of Fate of the batched positions, and reported as unavailable until some additional orders have closed.

Measure XVI. The Inheritance Load

The question. How much of the floating loss belongs to positions opened under rules the system has since retired.

The established method. Legacy-book separation, as a bank separates its run-off portfolio from its new lending.

Definition. The floating result of positions opened before the current rules, its share of the book's floating loss, and its size in days of the current rules' banking.

Reading	What it means
A large share of the floating loss in the inheritance	The headline figures judge the past, not the current system.
A small inheritance	The book's floating result belongs to the current system.

What remains uncertain. Exact in money; the figure in days moves with the banking rate.

VI. Reading the instruments in a live system

A measure is only useful if its owner can read it while the system runs. In the house's own system, the measures are computed by the system itself from its own records, refreshed with each evidence report, hourly by default, and printed both on a dashboard and in the system's journal, so that every figure on the screen can be traced to a line of record. A reading is marked for attention when it crosses a threshold stated in advance: a stop in the shallow class, a bank recorded near the stop, an inheritance carrying half the floating loss, a concentrated open vintage, a recovery debt of a week or more. A mark for attention is a call to read, never an instruction to act.

Building the instruments into the system also exposes their own defects quickly. In its first day of operation, one instrument measured the banking rate over too long a window and overstated the recovery debt many times over; another reported a horizon that no position had yet reached as a share of zero. Both were corrected within the day. Each correction is reflected in the definitions above, which is why Measure II reports unreached horizons as unreached and Measure VII measures its rate from the first opening under the current rules.

VII. What would falsify the Canon

The Canon makes three claims that can fail. The first is that the population is the right object: if a system’s positions are not in fact independent bets but one bet written many times, measures computed position by position will claim more evidence than exists, which is why Measure VIII exists and why errors elsewhere in the house’s practice are clustered by day and class. The second is that the borrowed methods carry over: survival analysis assumes that censoring is unrelated to outcome, and a system that closes positions for reasons connected to their fate would break that assumption. The third is that the record is complete: every measure is computed from what the system observed, and a position that moved while the system was not watching carries a depth the record cannot see. Each claim can be checked against the system’s own record, and the checks belong in the system rather than in an occasional review.

VIII. What this edition does not contain

This edition publishes the instruments, not their readings. Readings depend on the age of the evidence, and most of the measures above are, by their own definitions, not yet readable on young books. Later papers in the series will report readings only where the evidence has matured, with figures altered so that no account’s results are disclosed, and will record the measures that the evidence calls for. The questions this paper leaves open are those the measures themselves pose: whether the Line of Fate holds as deep classes resolve, whether a Rescue Threshold is ever reached, and whether the concentration of a vintage proves to be a property of a design or of the sessions in which its positions were opened.

References

Aalen, O. O. and Johansen, S. (1978). An empirical transition matrix for non-homogeneous Markov chains based on censored observations. Scandinavian Journal of Statistics 5(3), 141–150.

Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. Journal of the American Statistical Association 53(282), 457–481.

Fine, J. P. and Gray, R. J. (1999). A proportional hazards model for the subdistribution of a competing risk. Journal of the American Statistical Association 94(446), 496–509.

Sweeney, J. (1997). Maximum Adverse Excursion: Analyzing Price Fluctuations for Trading Management. Wiley.

Efron, B. (1979). Bootstrap methods: another look at the jackknife. Annals of Statistics 7(1), 1–26.

Mack, T. (1993). Distribution-free calculation of the standard error of chain ladder reserve estimates. ASTIN Bulletin 23(2), 213–225.

Perold, A. F. (1988). The implementation shortfall: paper versus reality. Journal of Portfolio Management 14(3), 4–9.

Dr. Glen Brown (2026). The Martingale Ledger. The Self-Judging Machine, Paper I. Public Doctrine Series.

Dr. Glen Brown (2026). Engineered Defaults. The Self-Judging Machine, Paper IV. Public Doctrine Series.

About the Author

Dr. Glen Brown is President & CEO of Global Financial Engineering, Inc. and Global Accountancy Institute, Inc. He is the architect of the GCPIAUT framework and the broader GATS-native proprietary capital doctrine. His work focuses on sovereign capital governance, algorithmic trading architecture, reserve-first compounding systems, and institutional financial engineering within closed proprietary environments.

Business Model Clarification

Global Financial Engineering, Inc. and Global Accountancy Institute, Inc. operate as closed-loop proprietary institutions. They do not offer public investment products through the doctrines described herein, do not invite retail participation into their proprietary capital architecture, and do not present the GCPIAUT framework as a public collective investment scheme. The concepts discussed in this article are part of the firms' internal intellectual and operating doctrine.

Risk Disclaimer

Trading and investment activity across foreign exchange, equities, futures, commodities, and digital assets involves substantial risk. Market conditions may change rapidly, and losses may occur. This article is provided for intellectual, institutional, and educational discussion only and does not constitute investment advice, an offer, a solicitation, or a recommendation to buy or sell any financial instrument or to participate in any investment structure.

Dr. Glen Brown
President & CEO, Global Financial Engineering, Inc. and Global Accountancy Institute, Inc.