GLOBAL FINANCIAL ENGINEERING, INC. · PUBLIC DOCTRINE SERIES · THE SELF-JUDGING MACHINE, PAPER I

The Martingale Ledger

Why a trading system's closed results flatter it, and the century-old theorem that tells the truth on day one

GFE-PUB-DOC-20261106-A · 6 November 2026

Dr. Glen Brown · Global Financial Engineering, Inc.

I. The flattering ledger

Every systematic trading operation keeps a ledger of closed positions, and almost every one of them, at some point, reads that ledger as a verdict. Sum the realised results, divide by the count, and the quotient is the edge. The procedure is so natural that it is rarely questioned. It should be. For a large class of trading designs, the closed ledger is not merely noisy in its early life; it is biased, and biased in the direction that management most wants to believe.

Consider a design in which every position rests a modest profit target close to its entry and a loss limit much further away. Such designs are common because they win often: the target is near, so most positions reach it. But the geometry has a consequence for measurement. Positions moving toward their near target arrive there quickly. Positions drifting toward their distant stop take far longer to get there. In the first weeks of any cohort, therefore, the closed ledger fills with wins. The losses that the cohort will eventually take are still open, carried as unrealised marks the closed ledger does not count. A system can show a comfortable edge on everything it has closed while being, in truth, slightly negative on everything it has done.

The usual remedy is patience: let the cohort run until enough positions have resolved for the ledger to settle. That remedy has a cost the textbooks rarely price. A system that admits and sizes positions every minute is making decisions throughout the waiting period, and it makes them on some estimate of its own edge, whether that estimate is stated or not. Waiting does not avoid the judgement; it merely defers the moment at which the judgement becomes honest. The question this paper answers is whether an honest judgement can be had on the first day.

II. The theorem that settles it

It can, and the instrument is one of the oldest results in the theory of random processes: the optional stopping theorem. Take as the null hypothesis that prices move as a fair walk, with no drift, between a position's target and its stop. Measure the position in units of its own risk, so that the stop sits at $-d$ and the target at $+b$, and let the position currently stand at a mark m between them. A fair walk started at m and stopped at the first barrier it reaches arrives at the target with probability $(m+d)/(b+d)$ and at the stop with probability $(b-m)/(b+d)$. Its expected final value is therefore

$$E[\,X \mid m\,] = b\cdot(m+d)/(b+d) - d\cdot(b-m)/(b+d) = m$$

Under a fair walk, the expected final result of an open position is exactly its present mark. It is neither a win waiting to happen nor a loss to be feared. It is, in expectation, precisely what it is worth now. Counting each open position at its mark, beside the realised results of the closed ones, gives an estimate of the cohort's expectancy that is unbiased from the first position onward.

The same two-point distribution supplies the uncertainty that remains. Given the mark, the variance of the final result is

$$Var[\,X \mid m\,] = (b - m)\cdot(m + d)$$

zero at either barrier and largest between them. A position near its stop adds little uncertainty because its fate is nearly sealed; one near entry adds the most. A ledger that counts marks without this term claims more certainty than it has.

The role of the fair walk deserves precision. It is not a belief about markets. It is the hypothesis to be rejected. A system with no edge is, by definition, one whose positions behave as though the walk were fair. Under that hypothesis, closed results plus open marks is exactly unbiased; under a real edge, the estimate moves in the direction of the edge. The instrument is exact precisely where exactness matters, at the boundary between edge and no edge.

III. An honest error: shocks are shared

A second correction is needed before the ledger can be trusted. Positions opened in the same asset class on the same day are not independent observations. They share a market, and when that market moves, they move together. A ledger that treats them as independent understates its own error, and an understated error is more dangerous than an admitted ignorance, because it invites action.

The remedy is the cluster-robust estimator familiar from econometrics. Positions of one class opened on one day form a cluster. Their deviations from the mean are summed within the cluster before they are squared, and the ratio of the clustered sum to the independent sum, floored at one, is the design effect:

$$deff = \left(\,\Sigma_g\,\left(\,\Sigma_{i\in g}\,(x_i - \hat{E})\,\right)^2 + \Sigma\, v_i\,\right) \Big/ \left(\,\Sigma\,(x_i - \hat{E})^2 + \Sigma\, v_i\,\right), \quad N_{eff} = N \,/\, deff$$

Here v_i is the conditional variance of each open position from Section II. The standard error is inflated by the square root of the design effect, and the effective number of observations falls accordingly. In practice the effect is rarely dramatic and frequently material: a design effect of 1.4 means that one hundred and forty positions carry the information of one hundred.

IV. The Case

This case is an illustration built from operating data; its figures have been altered, and it reports no account's results. It describes four multi-asset books run under one set of

rules, each judged on the positions opened since a common start date a week earlier, with older positions excluded from the judgement because they were opened under earlier rules.

Book	Positions	Closed only, R per close	Martingale ledger, R per position	Clustered error	Design effect	Verdict
A	206	+0.034	−0.006	0.016	1.00	Open
B	131	+0.045	−0.007	0.025	1.38	Open
C	72	+0.010	−0.017	0.035	1.21	Open
D	14	+0.061	+0.061	0.011	1.00	Thin

Illustrative figures, altered from operating data. Verdict: Thin below thirty positions; otherwise positive or negative only where the one-sided 95% bound excludes zero, Open where it does not.

Three readings follow. First, every book with open positions looked better on its closed ledger than it was, by between three and five hundredths of a unit of risk per position. In a design whose target is about six hundredths of a unit, that is the difference between a profitable system and an unprofitable one; the flattery is not a rounding error. Second, once open positions are counted at their marks and the error is clustered, every book with enough positions reads Open: none has yet proven an edge of either sign, and an honest instrument says so rather than guessing. Third, Book D shows why a minimum count belongs in the verdict. Fourteen closed positions, all winners, produce a bound above zero. A run of fourteen wins is entirely consistent with a loss rate of a few percent, and the ledger declines to be impressed.

The case contains one further lesson, about the unit itself. In one book, a single exchange-traded fund position closed at a trailed stop for a small profit in money but was recorded at more than twenty units of risk, while a second instrument in the same system recorded the same close at about five. That one row was enough to turn its book from slightly negative to apparently strongly positive. The theorem is exact about expectation; it says nothing about whether the unit was measured correctly. A ledger kept in units of risk needs a guard on its unit as well as a bound on its error.

V. What would falsify this reading

The martingale ledger rests on three assumptions. The first is the fair-walk null, which is not a weakness but the point: drift is what the ledger exists to detect, and it detects it without bias at zero. The second is that the barriers are where the ledger believes them to be; a stop that gaps, a target refused by the spread, or a stop moved after entry all change b or d, and a ledger that does not also track money will not notice. The third is the correctness of the unit of risk, which the case shows cannot be taken for granted.

The reading would be falsified if, as the open positions of a cohort resolve, their realised results diverged from the marks at which they were counted by more than their conditional variance allows. That is a test any system can run on itself, and one it should.

VI. The principle

The theorem in this paper is a century old. What is new, in the practice described here, is the decision to apply it to a live book on its first day and to print its verdict beside the error it deserves. Behind that decision sit three convictions. Evidence should belong to the system being judged, not to an average of its past versions. The absence of evidence is not permission to stop deciding, because a system that trades continuously acts on some estimate of itself whether or not that estimate is stated. And an instrument that flatters is more dangerous than one that is merely noisy, because flattery is acted upon. The closed ledger flatters. The martingale ledger does not.

The next paper in this series takes the ledger's output and asks the question that follows from it: once a system can say honestly what each of its strategies has earned, how much risk should each be allowed to carry?

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.