COLDCODE SI LABS
LIBRARY · HARARE

Separating a Mechanism from an Edge

Working Paper WP-01 · A calibrated instrument for broker-synthetic price feeds

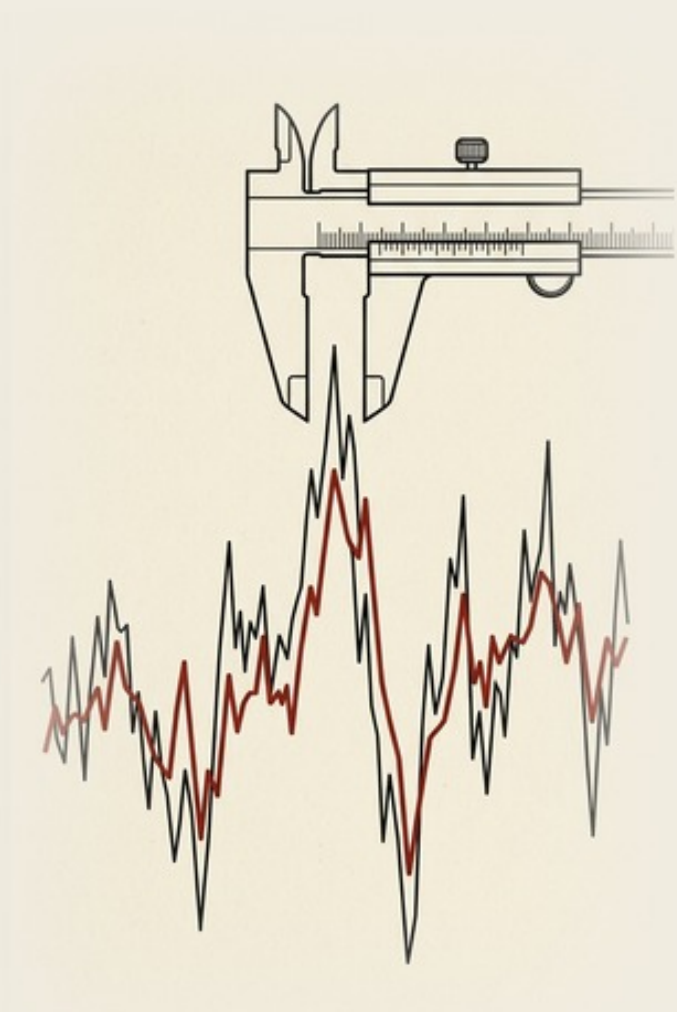

COURAGE GARIKAI

WORKING PAPER · NOVEMBER 2026

Separating a Mechanism from an Edge
A calibrated instrument for broker-synthetic price feeds, and the arithmetic that makes most of what it detects uninteresting

© 2026 Courage Garikai. All rights reserved.

No part of this working paper may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in reviews and as permitted by applicable law.

Coldcode Si Labs Working Paper WP-01 · First edition · November 2026
Published by Coldcode Si Labs, Harare, Zimbabwe

Correspondence: couragegarikai@atomicmail.io
Free samples and the full catalogue: library.coldcode-si-labs.workers.dev

Coldcode Si Labs Working Paper · Harare · November 2026

Separating a Mechanism from an Edge

A calibrated instrument for broker-synthetic price feeds, and the arithmetic that makes most of what it detects uninteresting

Courage Garikai · WP-01 · Version 1.0

Abstract

Everyone who has looked at a synthetic OTC feed has had the same two thoughts in sequence: there is structure in this chart, and therefore it is tradeable. Both can be true. In the class of feeds considered here they are almost never true together, and the gap between them is quantitative rather than philosophical.

This paper does three things. It states five mutually exclusive hypotheses for how a broker-synthetic price feed might be generated and gives each one an executable test with an explicit null. It calibrates that battery against known ground truth — five generators, sixty replicates — so that its error rates are known before it is pointed at a feed nobody controls: a perfect 5 × 5 confusion matrix in 12 replicates per class, and zero false positives on the efficient random walk. It then states the battery's detection floor honestly, at **$|\varphi| \approx 0.05$ in tick increments at 20,000 observations**, meaning a null result licenses "no serial dependence larger than that" and nothing stronger.

The fourth thing is the one that costs money. Using the same generators, we measure the highest hit rate any linear predictor achieves **out of sample at the venue's own contract horizon**, at a common horizon volatility. The result is stark: at a 60-tick expiry and a 92% payout, a feed must carry a tick-level autoregressive coefficient of **$\varphi \approx 0.43$** before its predictability reaches break-even. That is nearly an order of magnitude above the coefficient the battery detects easily. Structure that is unambiguously present, statistically significant, and detectable in minutes of data is economically worthless at the horizon the contract actually settles on.

Which leaves exactly one mechanism of the many available that produces the sessions everyone argues about: not prediction of the market, but the ordering of operations at settlement. A trade log can separate the two, and this paper specifies the discriminator and calibrates it.

---

1. The question, stated so that it can fail

The claim under test is not "the venue is fair" or "the venue is fraudulent." Both are unfalsifiable as stated, and both are somebody else's argument. The claim is narrower:

> **Given a capture of ticks from a synthetic feed, and the payout the venue is posting, is there a decision rule whose hit rate clears break-even with a confidence interval that excludes break-even?**

Everything in this paper is apparatus for answering that sentence and for saying honestly when the answer is we cannot tell yet.

Three design commitments follow.

A verdict names a generating process, not a trade. The battery does not output "profitable" or "rigged." It names one of five classes. Attribution and economics are kept in separate modules because conflating them is the central error in the public argument about these feeds.

Every test states its null. A test without a stated null is an invitation to read noise as structure, which is what a chart does to a human being at 3 a.m.

The instrument's error rates are known before use. We calibrate on generators whose truth we set, and we report what the calibration cannot rule out. Twelve replicates prove less than zero replicates in fifteen, and it is stated that way.

---

2. Five hypotheses, made executable

The generators are the hypotheses. If a feed is of class H, a sample from H's generator should be statistically indistinguishable from the feed.

Class	Generative story	What a reasonable operator would do
H0 efficient random walk	i.i.d. increments. A martingale in sign.	Nothing. There is no edge, and every win is variance.
H1 serial dependence	AR(1) in tick increments, $\varphi \neq 0$.	Maybe something — if the persistence survives to the contract horizon (see §5).
H2 replay	The feed is stitched from contiguous slices of a real historical tape.	Nothing legitimate. The structure is a copy, and the copy is detectable.
H3 mean reversion	Ornstein–Uhlenbeck on the price level.	Mean-reversion bets at horizons longer than the relaxation time.
H4 quantisation artefact	Structure in the tick grid, not in the returns.	Nothing. The structure is in the ruler, not the thing being measured.

The distinction that matters most in practice is **H0 versus H4**. A feed whose increments all sit on a grid of five ticks produces a spectacular autocorrelation function, an integer-valued ACF, visible "levels" on a chart, and exactly zero predictability. It is the single most common way an operator convinces themselves they have found something.

The battery: eight tests, eight nulls

#	Test	Null it rejects	Failure mode it catches
1		signs are independent	

Sample provided free for review. The complete work is available at https://library.coldcode-si-labs.workers.dev/