COLDCODE SI LABS
LIBRARY · HARARE

Manufacturing the Number

Technical Note WP-03 · What a rule search produces on a feed with nothing in it



COURAGE GARIKAI

TECHNICAL NOTE · NOVEMBER 2026

Manufacturing the Number
What a rule search produces on a feed with nothing in it, and the threshold you must apply if you search

© 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-03 · 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

Manufacturing the Number

What a rule search produces on a feed with nothing in it, and the threshold you must apply if you search

Courage Garikai · WP-03 · Version 1.0

Abstract

The question this note answers is the one an operator asks immediately after reading WP-01: if the edge has to be created rather than found, can it be created?

Not by me, and not by anyone, on a feed generated by an algorithm with no exploitable serial dependence — and WP-01 measures that feed rather than asserting anything about it. What can be created, reliably and on demand, is the **number**. This note builds the machine that creates it, calibrates the machine against known ground truth, and states the acceptance threshold that any real search must clear.

The results, from 545 rules of the kind an operator actually writes, on captures with known generating process:

- On a **pure random walk**, selecting the best rule in-sample produces in-sample hit rates of **52.6–57.0%** (mean 54.4%) and out-of-sample hit rates averaging **50.5%**. Six of eight seeds are classifiable as manufactured artefacts; **none survives**.
- On a feed with **real structure** (AR(1), φ = 0.90), the search returns 65.7% out of sample, with the lower confidence bound clearing break-even in **3 of 3** seeds.
- On a feed with **structure that is real but marginal** (φ = 0.50, ceiling 52.5% by WP-01's economic map), the search **cannot tell the difference** in 3 of 3 seeds. That is the honest limitation and it is the practically important case.

The consequence is a threshold that most operators do not apply and every vendor implicitly relies on:

$$p_{\text{required}} \;=\; \frac{1}{1+b} \;+\; \sqrt{2\ln K}\cdot\sqrt{\frac{1}{4n}}$$

where K is the number of rules searched and n the number of settled contracts. Searching 545 rules on 5,000 settlements moves the bar from 52.08% to **54.6%**. If you search, your threshold rises, and the size of the rise is a function of how hard you searched.

---

1. The rule space

545 rules, each a signal of +1, −1 or 0 at the decision tick, using only information available at that tick:

- **momentum and reversion** across 127 lookbacks from 1 to 400 ticks (254 rules) — the backbone of every "M1/M5 confluence" claim;

- **moving-average crosses**, all 45 ordered pairs from 3 to 233 (90);
- **Bollinger z-scores** on 8 windows × 10 conviction thresholds × 2 directions (160) — the "only high-confidence signals" filter, which stands aside below threshold and therefore reports fewer trades;
- **range breakouts** on 10 horizons × 2 directions (20);
- **RSI-style oscillators**, 3 windows × 3 band pairs × 2 directions (18);
- **three pure-noise controls** with no information whatever.

The zero state matters. A conviction threshold mechanically raises the hit rate by shrinking the sample, and shrinking the sample widens the interval — so selectivity manufactures the appearance of skill while destroying the ability to verify it. That trade is not a bug in the space; it is the space.

---

2. What the search produces when there is nothing there

Eight independent random walks, 300,000 ticks each, contract horizon 60 ticks, so about 2,500 settlements per half. Select the best rule in-sample, score it out-of-sample.

seed	selected rule	in-sample	out-of-sample	95% CI	verdict
0	`rsi_mom_14_2080`	53.3%	54.6%	47.2–61.8%	unclear
1	`rsi_mom_14_2080`	54.5%	46.7%	39.4–54.3%	manufactured
2	`rsi_mom_21_3070`	54.3%	49.9%	45.1–54.7%	manufactured
3	`bz_rev_20_2.0`	52.6%	52.6%	46.6–58.4%	unclear
4	`rsi_rev_14_2080`	55.2%	49.7%	42.6–56.9%	manufactured
5	`bz_rev_10_1.75`	57.0%	50.0%	44.3–55.7%	manufactured
6	`bz_rev_15_2.0`	54.2%	50.9%	44.3–57.6%	manufactured
7	`bz_mom_10_1.25`	54.4%	49.8%	46.7–52.9%	manufactured

Mean in-sample 54.4%. Mean out-of-sample 50.5%. Zero of eight produced a lower confidence bound above break-even.

Note seed 0 in particular. A pure random walk produced a rule showing **54.6% out of sample** — above the 52.08% break-even — and its interval runs from 47.2% to 61.8% because the conviction threshold left only 174 settlements to judge it on. A vendor publishing that rule is publishing a true number about nothing. That is not a hypothetical: it is what 545 rules on a random walk will hand you, most weeks, for free.

---

3. The multiple-comparisons floor

For K worthless rules whose predictions are independent, each estimated on n settlements, the best one sits near

$$p_{\max} \approx \tfrac12 + \sqrt{2\ln K}\cdot\sqrt{\tfrac{1}{4n}}$$

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