
DAN MCSPIRIT, CFA · VERIFICATION PRACTICE

What It Actually Takes to Trust AI-Built Analysis

AI did not reduce the work of building an institutional-grade valuation model.
It relocated the work, from construction to verification.

NOT INVESTMENT ADVICE

SHARPLINK (SBET) · ABSOLUTE VALUATION

THE CALCULATOR COMPARISON

Determinism is what earned the calculator its trust

A CALCULATOR

It has never once returned a confident, plausible, wrong answer.

The same input returns the same output, every time, and the failure mode is a mistyped key — which is the user's error, visible to the user.

AN AI SYSTEM

Being confidently wrong is intrinsic to how it works.

For prose this barely matters: a wrong word is visible on the page. For a valuation model, a wrong number is invisible, arrives fully formed, and propagates.

NOT INVESTMENT ADVICE

CONTEXT

What was built

Absolute valuation of a digital asset treasury company holding roughly \$2.2bn of ETH.

Excel workbook generated by a Python script, never hand-edited.

An 18-check regression harness, re-run on every build.

A companion web app, live, with its own generated-constant build step.

Every figure tied to SEC filings. Unsourceable figures left blank and documented rather than estimated.

BUILD TIME

2 days

Model, harness, app and documentation.

THE ARTIFACT

sbet-absolute-valuation.vercel.app

NOT INVESTMENT ADVICE

Eight errors

Errors 6, 7 and 8 were challenged by the second AI in the moment. On error 5 there was no challenge: the second AI had used \$11.24 throughout, and the divergence only surfaced later, from a screenshot.

#	WHAT THE AI SAID	WHAT WAS TRUE	CAUGHT BY
1	Reported “0 failed, 17 passed” after a set of changes	The harness defaults to the workbook beside the script; the build writes elsewhere. It validated a file the changes never touched. Pointed at the real output: 3 blockers	Re-running against the explicit path
2	The scorecard covered “2 of 6” deals	It was 3 of 6. Three composition lines cannot be two deals; the live cell said 3. Relayed onward and built into the app before it was caught	Re-deriving the composition
3	Quoted the coherence check as \$0.1817	It was \$0.1808. The figure was correct two constant-sets earlier and went stale	Checking against the rebuilt workbook
4	The model was seeded with a stale \$6.42 price	It was \$7.59, the real NASDAQ close. Read a known-issues list and asserted the issue was still open without checking the generator	Reading the generator
5	Gave the justified price as \$11.27	The tile compares against a treasury-stock-method line, so the coherent figure is the economic-basis \$11.24. A basis mix — the exact discipline being policed elsewhere	Not caught in the moment
6	Asserted a digit transposition in a correct figure	There was no transposition. The verification script had an $n > 1000$ filter that silently dropped a 252-share tranche. A 4.5-share residual had been visible and was dismissed as “immaterial rounding”	The second AI, challenging in the moment
7	Predicted a test would move a value to ~\$275,403	It moves to ~\$316,985, and in the opposite direction. A pre-refresh number quoted as a post-refresh expectation	The second AI
8	Said net optionality was -0.03	It was -0.01. Confused the size of the cut with the resulting value	The second AI

THE PATTERN

The arithmetic was never wrong

8

ERRORS IN THE SESSION

0

CALCULATION ERRORS

8

WRONG CLAIMS ABOUT OUTPUT

THE THREE FAILURE MODES

- A cross-reference that was never re-derived.
- A stale figure carried from an earlier state.
- A mismatch of basis between two sides of a comparison.

The AI computed correctly and then described its own output incorrectly.

So checking the math is close to wasted effort. The risk lives in checking that the claims about the math still hold: that the cell says what the sentence says, that the figure is from this build and not the last one, that both sides of a comparison sit on the same basis.

ERRORS AS STATED BY THE SYSTEM THAT MADE THEM

A harness that validated a file it never read

REPORTED

0 failed, 17 passed

RE-RUN AGAINST THE EXPLICIT PATH

3 blockers

The harness defaults to the workbook sitting beside the script. The build writes somewhere else. Both facts were true for months and neither mattered until a build's output and the harness's default diverged.

WHY IT IS THE WORST KIND

Nothing in the report was false about the file it read. The file was the wrong one.

A passing check is a claim about a specific artifact. The claim gets scrutiny; the identity of the artifact almost never does. A failure announces itself. This one announced the opposite, and it would have kept announcing it on every subsequent build.

The fix was one argument on one command. Finding it required doubting a green result that had no reason to be doubted.

A defect asserted in correct code

WHAT WAS CLAIMED

A digit transposition in the share count, and a 4.5-share residual dismissed as “immaterial rounding.”

WHAT WAS TRUE

No transposition. The verification script carried an $n > 1000$ filter that silently dropped a 252-share tranche.

DROPPED TRANCHE

252 shares

The consequence, had it gone unchallenged: correct code edited to satisfy a defective check. The second-worst outcome in a model is a wrong number. The worst is a right number removed on purpose.

THE REVIEWING SYSTEM'S REASONING

“You confirmed net warrant shares of 533,763 tie to the workbook exactly. Only the inclusive count produces that figure, so the workbook is already counting the 252-share tranche at \$7.77 — which is in the money against the \$7.91 close. 5,528,280 is the count consistent with the net you validated.”

“I said nothing downstream moves either way. It does, by about 4 shares, and those 4 shares are what distinguishes 533,763 from 533,759. The two figures aren’t independent, which is what settles it.”

RECOMPUTED BOTH WAYS AT \$7.91

Including the 252 shares	net 533,763.245
Excluding it	net 533,758.785

PRE-EXISTING DEFECTS

What the model itself was hiding

DOCUMENTED AND INERT

\$41.8m

Protocol incentives were documented as decaying 50% a year, in two places. No formula read the input. The valuation capitalised them flat for ten years, overstating the yield present value by \$41.8m.

The note beside it called the treatment “deliberately conservative.” It was the opposite.

HARDCODED ADDRESSES

3 places

Hardcoded cell addresses in the harness. They fired correctly once, when rows shifted.

The dangerous case is the silent one: had a different tab shifted, they would have compared the wrong cell and passed.

UNCONDITIONAL PASS LINE

Both

A reporting bug printed the pass line whatever happened, so a single run could report a failure and a pass together.

All three predated the session and had survived multiple prior audits. The harness that found them was itself defective, and the defects were only exposed because a build finally moved some rows.

THREE LAYERS, RANKED

What caught things, and what stayed human

01 The regression harness

Cheap, automatic, the only layer that runs every time. It catches only what someone thought to encode. It ran clean for five builds over an inert input worth \$41.8m.

02 A second AI reviewing the first

Involved in four of the eight, including the two subtlest: three challenged in the moment, one a silent divergence. Not because it was smarter, but because it had no stake in the first system's conclusions and re-derived rather than accepted.

03 The human

Caught the rest, and made every call that mattered.

The most effective check was not a better model. It was two systems disagreeing, with a human adjudicating.

JUDGEMENTS, NOT OUTPUTS

Capitalising a decaying yield strip on its own schedule rather than flat.

Cutting a forward option value from 5% to 2% against a completed track record.

Leaving three cells blank rather than estimating them, and documenting which disclosure was missing for each.

Testing the company's own “premium to NAV” characterisation and finding it did not hold.

Every share-count basis convention in the model.

AI did construction, retrieval and arithmetic checking. It did not decide what any number meant.

THE TRADE

A calculator earned its trust by being deterministic. AI has not earned that and is not going to. What it offers instead is enormous leverage on construction, in exchange for a verification burden most people building with it have not yet priced.

You can't get somethin' for nothin'.

The checking is the job now, and it is a different job than the one you know how to do.

Provenance: this deck was drafted by the same AI that made all eight of these errors, working from the author's brief, and checked by the author.

sbet-absolute-valuation.vercel.app
danmcspirit.com · @danmcspirit
NOT INVESTMENT ADVICE