

S-07, Anonymized sample from a real engagement

QA report, return card data validation, sprint 12

A clean room recompute of the return cards for 24 funds, with zero fabricated values, plus a check of 12 claims in a sprint report.

Anonymized sample from a real sprint for a US investment analytics platform, 2026

What this is. An independent check of whether every figure on the platform's return cards is real and correctly computed. I recomputed each figure from raw prices with separate code, compared the two number by number, and had reviewers try to refute each conclusion. Fund tickers are replaced with Fund A to Fund X, and the price vendor and the public cross check site aren't named. The appendix is the claims check from the sprint before, which shows the same method applied to my own report.

Return card data validation

Funds audited	Fabricated values	10Y mean difference	Method
24	0	0.035 points	Clean room recompute

The headline is simple. Nothing on the return cards is fabricated. Every displayed figure was reconstructed on its own from raw dividend adjusted prices, and the annualized returns reproduce to about a thousandth of a percentage point. The audit did surface real notes on the one year figure and on disclosure, and two of them are already fixed. This report shows the raw data behind each number, so nothing here rests on trust.

1. How the audit was done

Two methods, run against each other. First, every stored card number was pulled straight from production. Second, a clean room script fetched raw daily adjusted closes for all 24 funds from a separate source and recomputed each return from scratch. It anchored the

multi year figures on the close nearest each calendar anchor. The two were then compared number by number. A panel of adversarial reviewers then tried to refute each conclusion and to catch any invented citation.

2. The core result, the numbers are real

Confirmed by both methods. Across all 24 funds, the platform figures match the recompute to a fraction of a percentage point on average. The match is tightest on the ten year horizon, which is the least sensitive to a day of anchor drift.

Horizon	Mean difference	Worst case	Reading
10 year annualized	0.035 points	0.111 points	Exact for display
5 year annualized	0.070 points	0.253 points	Exact for display
2 year annualized	0.194 points	0.689 points	Snapshot timing on volatile funds

Worked example, checkable by hand

Fund W, ten years. The compound annual rate is the latest adjusted close divided by the adjusted close ten years earlier, to the power one tenth, minus one. That gives 14.79 percent. The platform shows 14.81 percent. The small gap is the platform snapshot sitting a fraction above the independent close. It isn't a computation error.

3. A public source agrees

A third confirmation. The long run figures were cross checked against a public total return site, which publishes trailing annualized returns with dividends reinvested for these funds. The card figures sit within a few tenths of a point of the public record. The small gap is the end point, because the card stands one trading day later than the public page.

Fund	Card 10Y annualized	Public source 10Y	Difference
Fund W	14.81%	14.91%	-0.10 points
Fund V	20.05%	20.22%	-0.17 points

Both public figures were read directly from the source on the day of the audit. Three separate measurements now agree on each number, the platform, the clean room recompute and the public database.

4. The findings, ranked

All checked against raw data. None fabricated.

Finding	Verdict	Action
2, 5 and 10 year annualized returns reconstruct exactly from raw prices	Real and correct	None
Prices and the total return basis are correct, dividend and split adjusted	Real and correct	None
The one year figure used a fixed 252 trading day lookback rather than a true calendar year, which drifts slightly on the international funds. The drift was small, typically well under one point	Real, method	Fixed, now calendar anchored like the multi year columns
The figures are total return but nothing on the page said so. A user checking a broker price chart would see a smaller number	Real, disclosure	Fixed, source and basis line added
The price source is an unofficial endpoint, and an adjusted close is rescaled each time a new dividend posts	Real, low impact	Cited on the page
One commodity fund carries structural roll decay, so a long run buy and hold reading misleads by the fund's nature, not the data	Real, fund nature	Noted
Stored values carry four decimals, more precision than the data supports	Mitigated	The card and the export already show two decimals
No survivorship gap and no bad split. One fund's reverse split is carried correctly	Clean	None

5. What was fixed as a result

Deployed to production. The one year return now anchors on the close nearest one calendar year ago, the same way the multi year columns already anchor. So it's a true calendar year figure across every fund, rather than a trading day count that drifts. And the fund detail page now carries a source and basis line. It says the figures are total return, dividend and split adjusted, that the multi year columns are annualized, and where the prices come from. That line is the citation the data was missing.

6. Every fund, claim against recompute

The ten year figure is the compound annual rate from the ten year anchor to the last close. Claim is what the platform shows. Independent is the clean room recompute.

Fund	Asset class	10Y claim	10Y independent
Fund A	Developed market equity	7.7852%	7.7336%
Fund B	Developed market equity	4.4384%	4.4247%
Fund C	Developed market equity	8.5312%	8.5045%

Fund	Asset class	10Y claim	10Y independent
Fund D	Developed market equity	9.0853%	8.9932%
Fund E	Bonds	1.3787%	1.3704%
Fund F	Bonds	1.4261%	1.4192%
Fund G	Bonds	2.0318%	2.0203%
Fund H	Bonds	-2.3747%	-2.3898%
Fund I	Commodities	9.2563%	9.2899%
Fund J	Commodities	11.0481%	11.0330%
Fund K	Commodities	10.3289%	10.2765%
Fund L	Commodities	5.1286%	5.2391%
Fund M	Emerging market equity	7.7867%	7.7112%
Fund N	Emerging market equity	2.6068%	2.6011%
Fund O	Emerging market equity	8.3464%	8.2739%
Fund P	Emerging market equity	7.4589%	7.4240%
Fund Q	Developed market equity	7.5781%	7.5576%
Fund R	Developed market equity	8.4337%	8.4031%
Fund S	Developed market equity	10.0385%	10.0068%
Fund T	Developed market equity	9.7687%	9.7378%
Fund U	US equity	13.1102%	13.0898%
Fund X	US equity	10.5237%	10.4959%
Fund V	US equity	20.0531%	20.0301%
Fund W	US equity	14.8055%	14.7891%

In the report as sent, this table also listed the raw adjusted close at each anchor date, two, five and ten years back, for every fund. The prices are left out of this sample because they would name the funds.

Appendix, claims check of my own sprint 11 report

Every claim in the sprint 11 report was checked adversarially against raw live data before I sent it. Each check ran a real production shell, an API request, a cloud query or a code read. Each checker had the instruction to mark a claim unverified if real data couldn't confirm it.

Claims checked	Verified on live data	Overstatement caught	Fabricated
12	11	1	0

The one caught overstatement, and the fix

The draft said the report rebuild makes the PDF without running scripts again, stated as an absolute. The code shows an option that does run every attached script. So the absolute is false for that path. It's true only for the path the email sender uses, and I bound the wording to that path.

Claim	Verdict	Observed evidence
The new source endpoint returns a failing script's code	Verified	HTTP 200, 2,260 bytes of real Python, no error page
149 failing script sources exported	Verified	149 files, all .py, the first one opens with real imports
Naming convention code present	Verified	The error class, the source parser and the raise all present, with line numbers
Naming convention live on production	Verified	A production shell printed the raised error for an unloaded ticker and 500 rows for loaded ones
The sender builds a missing PDF before sending	Verified	The send path rebuilds, waits, refreshes, checks the latest PDF again, then attaches and sends
The report rebuild skips the script rerun	Corrected	True only for the default path the sender uses
The sign up page shows the current brand	Verified	No old logo import and no old brand text, screenshot on file
Deployed on main at the stated commit	Verified	GitHub head matches, app and API both answer 200
Daily pipeline scheduled and enabled	Verified	Both report rules enabled, with the right task targets
The time statement bills 10 hours honestly	Verified	Hour cells sum to 10.0, every line maps to a file on disk
Analysis sandbox kit built	Verified	Dockerfile, client, README and a working welcome notebook all present

Claim	Verdict	Observed evidence
The source endpoint does not run code	Verified	The failing script's code came back cleanly. The endpoint reads the file and runs nothing

Axdale, Michael Lip. Warsaw, Poland (CET). michael@axdale.com. www.axdale.com. Code record github.com/theluckystrike. Names, tickers and amounts are replaced.