News & Insights · 16 September 2026

From one spreadsheet to a working system: the early days of AI at Merlon

Hamish Carlisle · Portfolio Manager

In July, one of our portfolio managers asked a narrow question: how do we get our monthly performance attribution out of Excel?

Attribution is how we explain a month's result: how much came from the sectors we held, and how much from the individual stocks we chose. For years it had run through a combination of large spreadsheets, sets of macros and files from the fund administrator that often arrived weeks after month end in unwieldy formats. It worked. It was also the kind of process every firm has somewhere: important, inelegant, understood by very few people, and uncomfortable to change.

Two months later, that question has grown into a system that supports much of how the firm works: our reporting, our research library, our company engagement and the work of our Advisory Board. Hamish Carlisle drove the project, working with an AI coding tool, Anthropic's Claude Code, which was involved in more than three quarters of the changes. Hamish is not a software engineer. The rest of the investment team is now building the same skills.

It is early days. But the journey so far has shaped how we think about the next few years, and very little of it is specific to funds management.

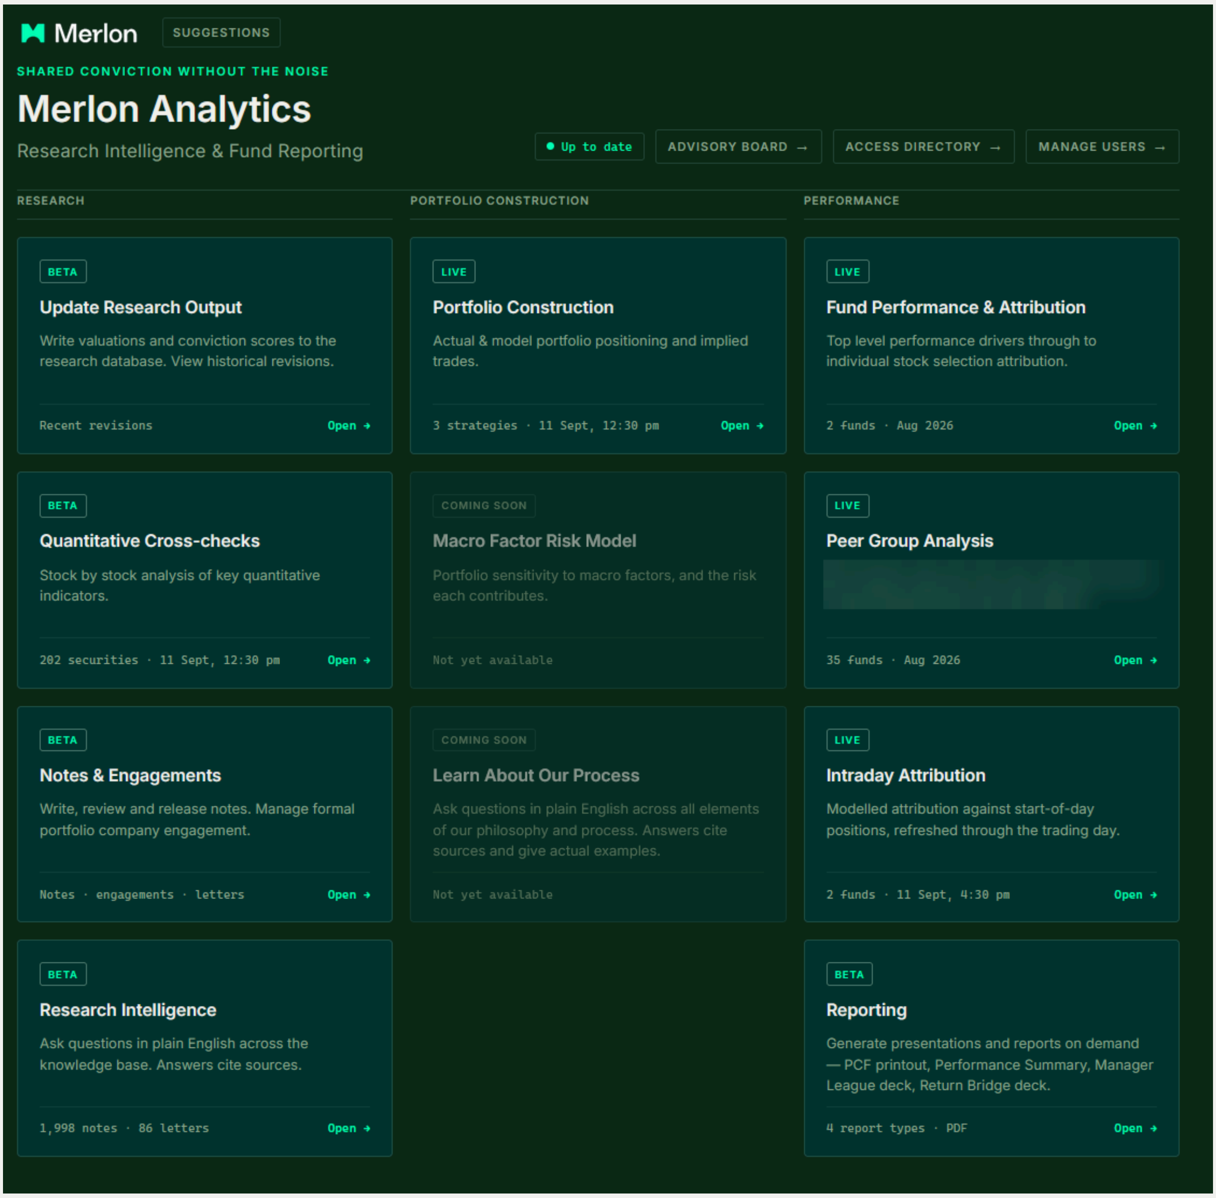

Where it lives: one internal home for research, reporting and engagement tools, built by the investment team.

Why a boutique would bother

For a small investment team, time is the scarcest resource. Every hour spent assembling a report, chasing a number between systems or searching for an old note is an hour not spent on what drives results for advisers and their clients: understanding companies, testing our views and making good investment decisions.

That was the point from the start. The aim was never to automate the investing. It was to automate the work around it, so more of the team's time goes into the investing itself.

It started by writing things down, not by writing code

The first change did not touch the spreadsheet's logic at all. It took decisions that were buried in the workbook and put them somewhere they could be read.

Take the table that matches our portfolio system's security codes to standard market identifiers. It had lived inside the process, so changing it meant opening the process. It moved into a plain file with a history, so every change can be seen, reviewed and reversed. The monthly output got the same treatment: one agreed description of what each month's file contains, which everything downstream relies on.

Two days after the first change, the results no longer lived only in the spreadsheet: they came out as clean data files that any other tool could read. The attribution calculation itself was left exactly as it was: it was correct and understood, and the AI was told not to rewrite it. The new output was checked against what we had already published, field by field. It matched.

Reporting was the smallest part

Once the output was clean, much of our regular reporting came to run on a schedule: monthly fact sheets and attribution packs, portfolios and index weights loaded into our portfolio system each morning, a private dashboard for clients and the performance tables on this website.

That is useful, but it turned out to be the least interesting part of the journey. The larger gains came where AI does more than move data around.

Research Intelligence: a research library we can ask questions of

In the last seven years our team have written about 2,000 research notes, and since 2021 we have sent more than 70 letters to company boards and received dozens of replies. Most of it sat in notebooks and folders, where finding what we thought about a company two years ago meant knowing where to look.

It is now one searchable library, and adding to it is as simple as sending an email. An analyst can send a note from their phone straight after a company meeting, including a photo of handwritten notes, which AI reads so the content can be searched.

The part the team uses most is **Research Intelligence**. A member of the investment team types a plain-English question, such as how our view on a company has changed or what we asked a board for and what happened, and gets an answer that cites the specific notes and letters it came from. Three rules keep it honest:

Every claim cites a source. If our research does not cover the question, it says so rather than guessing.

A company's words are never presented as our view. A reply from a board is labelled as the company's position.

The newest thinking comes first. When our view has changed, the answer leads with the latest note and dates it.

It answers from our own research, not the internet, and it does not decide anything. What it saves is the hour of searching before a meeting, and it puts the full history of our thinking in front of whoever is making the next decision.

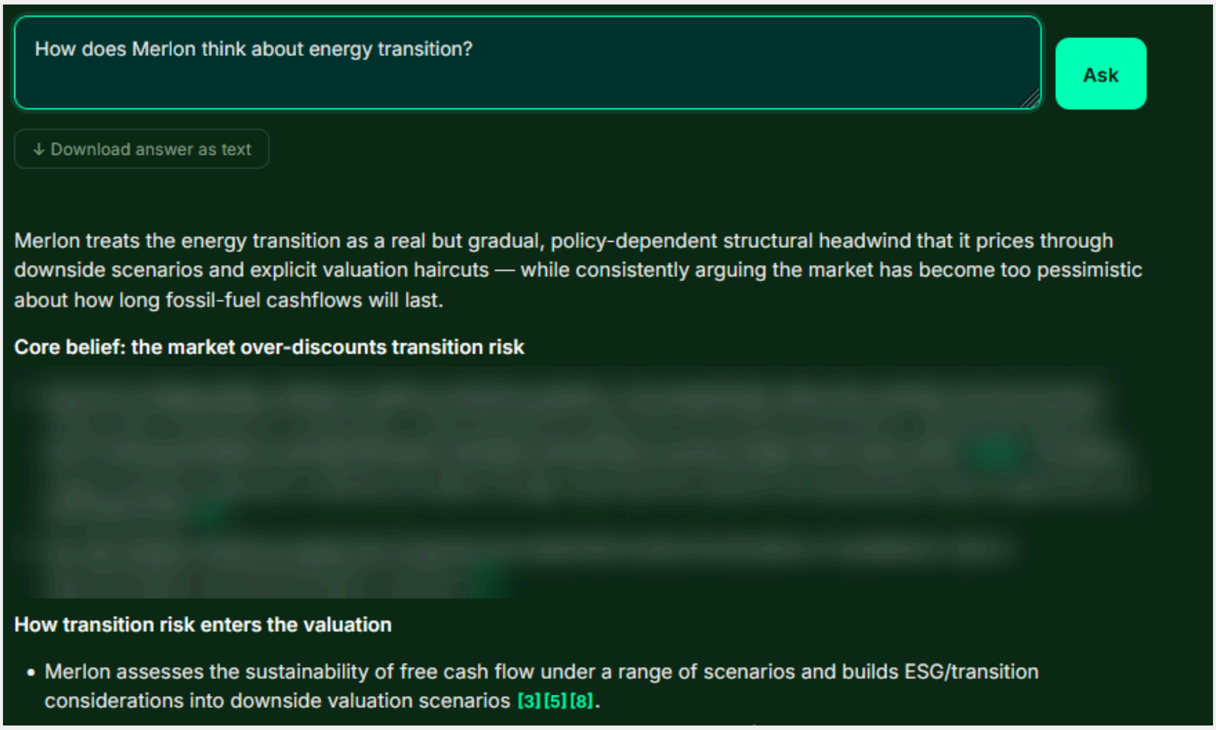

Research Intelligence answers from our own research and cites every source, so the reader can check the note behind each sentence.

Company engagement and the Advisory Board

Writing to company boards is central to how we invest. An introduction letter sets out how we see the company, the issues we think matter and where we believe the market has it wrong.

AI now prepares the first draft from our own research, setting out the market's view, the company's view and ours for each issue. From there it is a human process. The analyst holding the draft rewrites it. It passes through peer review, and our independent Advisory Board reviews and advises on it before it is sent. The system records who holds the draft at each stage, so nothing goes out without passing every review.

The Advisory Board's own work runs in the same place: meeting agendas, the papers tabled, and follow-up actions carried from one meeting to the next. Members can read the research behind a letter but cannot change it. Because engagement is now recorded as it happens, the counts we report to the Board, such as meetings held and letters sent, can be built from the underlying records rather than compiled by hand at year end.

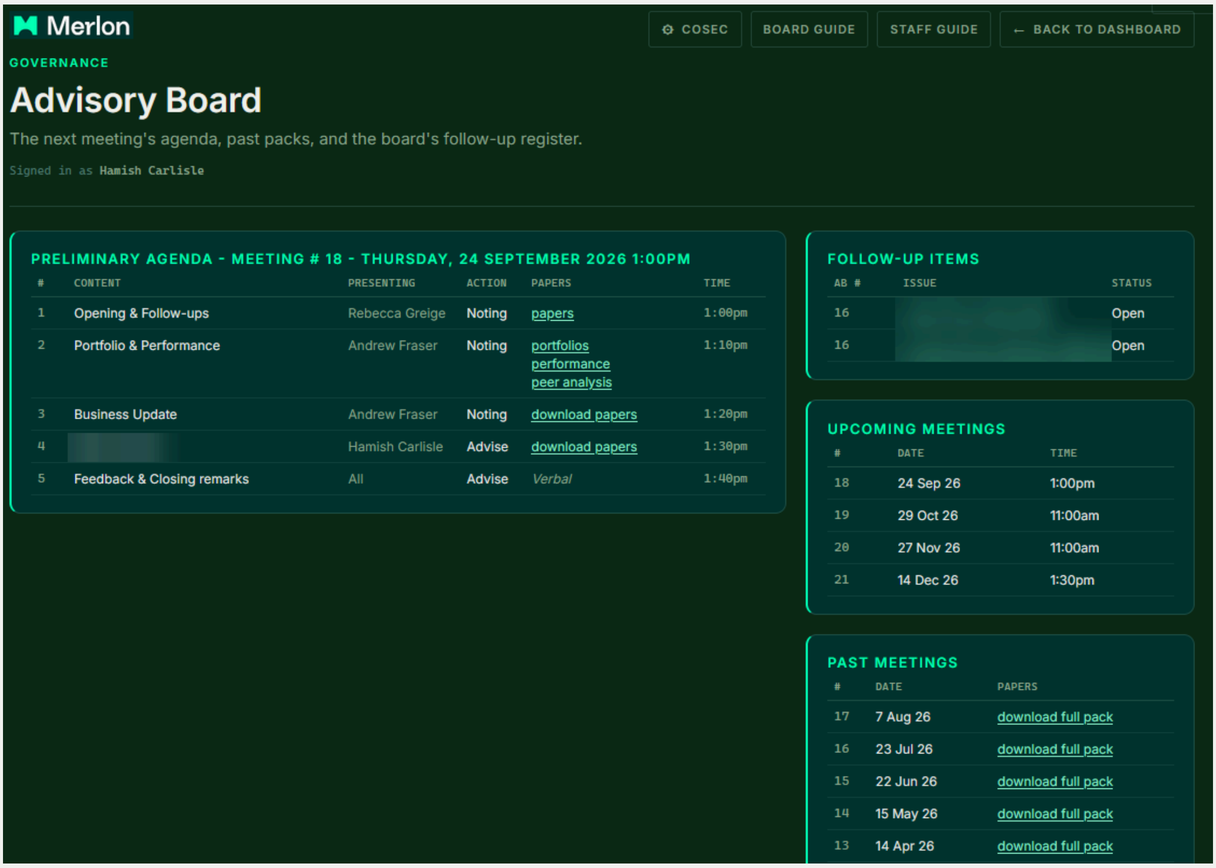

The Advisory Board's agendas, papers and follow-up actions, kept in one place.

What working with an AI coding tool actually looks like

It is less dramatic than it sounds. A person describes the problem in plain English. The AI reads the existing code and notes, proposes an approach, writes the change and tests it. The person reviews it, questions it and decides.

Three habits made it work.

Keep the written record next to the work. Every part of the system has notes explaining what it does, the decisions behind it and the mistakes that taught us something. The AI reads those notes before it starts, so each session begins from what is already known. That record is also what lets a colleague who is new to this pick up where someone else left off.

Isolate every change. Each piece of work happens in its own copy of the code, and nothing reaches the live system except through a reviewed request to merge it. Since late August nobody can bypass that step: not the team, and not the AI.

Measure rather than assume. An AI is confident whether or not it is right, and so are people. We test how a data provider or website actually behaves before relying on it. More than once, the confident answer was wrong.

The controls matter more than the tools

If there is one lesson for advisers, it is this. The tools will change. The controls around them are what make automation safe to rely on, and they carry over to any practice.

A person decides anything that reaches a client. We keep a written list of every way a change can reach a client or the public, and none is triggered by the AI's own judgement. A clean automated check is not sign-off: client-facing documents still go through compliance review.

Checks refuse rather than warn. A warning that appears every time trains everyone to click past it. Our checks stop the job when something looks wrong.

Every job reports that it ran. Twice this year a scheduled job stopped and nobody noticed for about a week, because data that stops updating looks exactly like data that has not changed. Now every scheduled job records each run, so silence is visible.

Decisions live in files, not in code. A judgement call written down where it can be reviewed tells the next person, or the next AI session, why it was made.

Early days, and why we are investing now

Two months is a short time, and we are at the very beginning of this. What exists today is a first version, built mostly by one person while the rest of the team learns. We expect it to look quite different in a year.

Three convictions have come out of it, and they shape what we do next.

Capability is the investment. The tools are improving quickly, and whatever we build this year will be rebuilt. What lasts is a team that knows how to use them well: how to describe a problem, check the answer and put the right controls around it. So we are investing in that skill across the investment team, not only in the systems.

The investment team should lead it. The people who understand a process best are the people who do it every day. When a portfolio manager can turn a frustration into a working tool in an afternoon, nobody has to write a specification, wait for a project queue, or explain the subtleties of attribution to someone who has never needed them. Technology specialists still matter for security, infrastructure and the controls around them, but the ideas and the judgement come from the people doing the investing.

This is an evolution, not a project. We are early in a change that we believe will completely transform the way we operate. It will not happen in one step, and it will not replace the judgement at the centre of what we do. It will steadily change how much of the team's time goes into that judgement.

The point is better decisions

None of this is AI making investment decisions. Our portfolio decisions and the judgement behind them remain with the investment team, and the attribution maths is exactly as it was.

What has changed is where the team's time goes. Less of it is spent gathering, reconciling, formatting and searching. More of it goes into meeting companies, testing our views against our own history and deciding what to own. That is where results for advisers and their clients are made, and it is where we want our time to go.

If you are thinking about doing the same

Start with one process everyone relies on and nobody wants to touch. For us it was a spreadsheet. For a practice it might be review preparation, fee reconciliation or client reporting.

Let the people who do the work lead it. They know where the time goes and what "right" looks like.

Write the process down before you automate it. Much of the value came from discovering what the process really was.

Think about where your own knowledge lives. Years of file notes and client conversations are worth far more when they can be searched and questioned.

Keep a person's sign-off on anything client-facing, and make that step impossible to skip.

Make every automated job report that it ran.

Think about client data first. Before any client information goes near an AI tool, consider your privacy obligations and your licensee's policies on AI use. Most of what we automated runs on market, portfolio and research data rather than client personal information, which made it a sensible place to start.

We began with one narrow question about one spreadsheet. We are a long way from finished, and that is the point: the most valuable thing we have built so far is the capability to keep going.

Disclaimer.

This document has been prepared by Merlon Capital Partners Pty Limited (ABN 94 140 833 683 / AFSL 343 753) (Merlon).

The advice contained in this document is general advice. It has been prepared without taking account of your objectives, financial situation or needs and because of that, you should, before acting on the advice, consider the appropriateness of the advice, having regard to your personal circumstances. You should also consult a licensed financial adviser before making an investment decision.

Any investment in a Merlon Fund (Fund) needs to be made in accordance with the Product Disclosure Statement (PDS), Additional Information Booklet and Target Market Determination (TMD). Investors should consider the PDS, Additional Information Booklet and TMD before deciding whether to invest in a Fund or continue to hold units in a Fund. Applications for units in a Fund can only be made pursuant to the application form relevant to that Fund. Distributors of our products must consider our TMD. You can access the relevant PDS, Additional Information Booklet and TMD at the Merlon website: https://www.merloncapital.com.au/

Any projections contained in this communication are estimates only. Such projections are subject to market influences and contingent upon matters outside our control and therefore may not be realised in the future. Investing involves risk including the risk of loss of principal. Please note that investments are subject to market and other risks and there is no assurance or guarantee that the intended investment objectives will be achieved. An investment's return and an investor's principal value will fluctuate so that, when redeemed, an investment may be worth more or less than their original cost. An investment's current performance may be lower or higher than the investment performance noted within the article. Past performance is not a reliable indicator of future performance. Any reference to past performance in this article is intended to be for general illustrative purposes only and should not be relied upon.

While we believe the information contained in this article is based on reliable information, no representation or warranty, express or implied, is made as to the accuracy or completeness of the information and opinions contained in the article. While every care has been taken in the preparation of this material, to the fullest extent permitted by law, Merlon its respective officers, directors, employees or agents, disclaim any and all liabilities for any loss or damage (including consequential loss or damage) howsoever arising in connection with the use of this article or its contents.