Grey Matters
ExploreGreyMatters.ai

CASE STUDY • LEARNING LIBRARY

From Pensions Planning to a Living Decision Planner

How a financial forecast evolved into a practical way to explore real-life choices.

Human-led. AI-enabled.

Grey Matters • October 2026 • Published edition v1.2

CASE STUDY / FROM PENSIONS PLANNING TO A LIVING DECISION PLANNER

Some decisions are modest but still matter. Do we repair a broken appliance or replace it? Could a four-day week make family life easier? How do we deal with an unexpected bill or a debt that is becoming uncomfortable? Other choices, such as moving home or changing when we stop working, can affect years ahead. We do not all have the same resources or choices, but we can all benefit from being clearer about what a decision changes elsewhere. My route into this was pensions planning. It is an example, not an entrance requirement.

It began with a forecast, not a grand plan

I wasn't trying to invent a new approach to life planning. I wanted to make sense of choices that were becoming more important: what we might spend, what we wanted to enjoy, what we needed to protect, and how one decision might affect another later on.

There were professional financial forecasts to start from. They were useful. They showed how different assumptions about spending, investment returns and time could change the shape of the future. But I found that I wanted to ask questions around them — sometimes slightly awkward ones. What happens if we make a different choice? Which assumptions are doing most of the work? If a plan says we're 'on track', on track for what, exactly?

AI made those questions easier to pursue. It could read complex information, compare alternatives, develop explanations and help build models. Unlike a static report, the conversation could follow a new question the moment it occurred. That felt liberating, particularly because financial planning wasn't a strange foreign territory to me. The lived context was familiar. The figures mattered, but I already knew why we were asking.

Another person might begin with a weekly budget, a choice between two affordable second-hand cars or a conversation about sharing caring duties. They would not need our forecasts. The useful starting point is the same: what are we trying to improve, and what would we be giving up?

The starting point was never “How clever a model can we build?” It was “Can we make a better-informed choice?”

When possibility got out of hand

The trouble is that AI rarely runs out of suggestions. One scenario led to another: different spending patterns, ways of financing a purchase, changing tax rules, market shocks, property decisions, retained assets, other contingencies. Each addition could be justified. Each looked like another useful way to see the future.

Before long, the workbook had grown into an elaborate piece of machinery. At one stage there were roughly fifty worksheets, many hidden away to keep the visible surface manageable. It was impressive in its own way. It was also becoming difficult to understand what was driving a result, how one setting affected another and which version of the calculation was actually authoritative.

That's a slightly uncomfortable admission, because I had helped direct the expansion. AI made it easy to keep adding capability, and I kept finding interesting questions to ask. The outcome was not simply an AI mistake. It was a human–AI collaboration that had lost sight of the appropriate level of complexity.

We had to stop treating a larger workbook as progress. The model needed a different job description: preserve trusted source information, make the assumptions visible, test a specific choice, and show the consequences clearly enough for a human to understand and discuss. The untidy development history was useful learning, but it shouldn't become the permanent operating interface.

In ordinary life the warning sign might be smaller: asking AI to make a ten-sheet budget for a one-off household purchase, or to investigate every possible lender before writing down how much can realistically be repaid. More detail can be useful, but only when it answers a question that matters.

A number can be correct and still tell the wrong story

Some of the most important discoveries weren't about spreadsheets failing to calculate. They were about spreadsheets calculating a meaning we hadn't intended.

An asset moved from one account to another could appear to have vanished from the family's wealth. A financed purchase could be treated as money gone while the value of the asset itself was omitted. A repayment could be recorded as a cost without recognising the ownership it created. At one point a reassuring check was effectively comparing a reference against itself, rather than testing the calculation we were actually using.

AI could help identify and correct these problems, and it could build more systematic tests. But the challenge usually began with a human question: 'Hang on, that doesn't sound like what happens in real life.' The arithmetic mattered; so did the ordinary meaning of buying, owning, spending, transferring and keeping something.

The same distinction applied to investment risk. A scenario with a higher assumed return might produce a more attractive future value. That didn't mean it was automatically a better choice. The extra uncertainty, the timing of losses and the ability to continue living comfortably through a difficult period belonged in the discussion as well.

The same common-sense challenge applies if someone is comparing a loan, a repair and waiting. A monthly payment is not the whole cost, an attractive rate may come with conditions, and a bargain is no bargain if it removes the money needed for essentials. The figures need checking, and serious debt difficulties deserve competent independent support, but the underlying question belongs to the person living with the consequences.

Simplifying the model changed the conversation

We eventually separated things that had been getting tangled together. Original external forecasts were retained as a reference, not rewritten to make new scenarios fit. Actual information about the present had its own place. Agreed planning assumptions formed a working baseline. And a possible life choice sat above that baseline as an experiment until somebody actually decided to do it.

It sounds obvious written down, but this was a turning point. We could ask a question without pretending the question was already a decision. We could test a different approach without erasing the last accepted position. We could update real-world information without quietly changing the assumptions used to explain a previous conclusion.

Most importantly, we could distinguish an attractive output from a useful one. The measure of success stopped being how many scenarios the tool could calculate. It became whether the result gave us greater clarity about the choice in front of us.

THE PRACTICAL STRUCTURE WE SETTLED ON

What we know → What we assume → What we could do → What changes → What we choose. The model supports the first four. Only a human can settle the fifth.

The website now shows the fuller process we developed: personal circumstances and relevant outside changes feed into possible choices, scenario testing and trade-offs, followed by a human decision and later review. That is the mature end of our own journey, not the apparatus someone must acquire before they can begin.



Figure 1 — The full Living Decision Planner process from the Grey Matters website. This shows where our more sophisticated model evolved to; readers can begin with a much simpler decision and a written comparison.

Numbers needed an interpreter

Even the cleaned-up model wasn't the whole answer. A table could show that an option changed future resources or resilience. It couldn't tell us what that change meant for the kind of life we wanted.

That led to a separate written analysis — not simply a printout of the spreadsheet. It explained what we had examined, which assumptions were influential, what changed between the options, what might go wrong and which issues warranted further discussion or independent confirmation. Conclusions and recommendations could be tentative, including a recommendation to do nothing yet.

We discovered that the written interpretation was as important as the numbers. It turned a model's output into something we could understand, disagree with, refine and take into conversations with

other people. It also left a record of why a decision looked sensible at the time, rather than just the answer produced by the latest formula.

Someone considering an expensive home repair could use the same separation without a spreadsheet: one page comparing the urgent work, what can safely wait, available help and the strongest reasons for each choice; a second short paragraph saying what they intend to do and why. Understanding is the deliverable, not the number of cells.

The model asked, “What happens if…?” The written analysis asked, “So what does that mean for us?”

A living planner, not an annual event

The process didn't stop once a workbook was working. Personal circumstances evolve; outside events sometimes matter too. Inflation, policy, markets, employment and new technology can alter the assumptions behind a plan. But not every headline deserves a new decision.

A recent example was an investigation into whether apparently different investments might share exposure to the same underlying economic forces. We used AI to examine common factors, stress several possible outcomes and explore a less obvious question: could even a successful new technology create disruption elsewhere in the economy?

The analysis was developed in stages, with written conclusions retained separately. It sharpened questions about concentration, resilience and the evidence we would need next. It did not, simply by being interesting or dramatic, justify an immediate change. That was the living-decision benefit: the ability to absorb new information, test what it actually changed, and leave the existing decision alone when the case for movement wasn't established.

A decision planner worth keeping must be capable of returning 'nothing material has changed'. Otherwise it becomes a machine for reacting to news — and probably for giving its owner unnecessary work.

For a household purchase, a meaningful trigger might be a confirmed repair quote or a change in income. For a debt plan, it might be an actual creditor response or a change in essential outgoings. A new AI opinion, by itself, is not a changed fact.

Why human leadership felt different here

Building the Grey Matters website sometimes meant hanging on by my fingertips while AI moved confidently through unfamiliar technical language. The life-planning work felt different. I wasn't a specialist in every financial or tax detail, and that called for appropriate professional checks. But I did know the circumstances, the preferences, the things that worried us and the experiences worth protecting.

That familiarity changed the balance. Rather than trying to catch up with an AI explanation of unfamiliar machinery, I could challenge the model from a position of practical knowledge. If its answer didn't feel plausible, I could ask why. If it privileged the highest eventual number over the quality of life along the way, I could say that wasn't the question. I could use independent advice more effectively because I had already thought through what mattered and where I needed another person's expertise.

AI contributed structure, patience, comparison, calculation and fresh angles. The human contribution was not a ceremonial tick at the end. It was the purpose of the exercise, the understanding of real life and the ability to recognise when a technically neat answer was beside the point.

That human advantage is not reserved for financial expertise. Someone managing a tight household budget knows which costs genuinely cannot move. A carer understands the realities behind a timetable. A person weighing a job change knows how much a less stressful week might matter. AI can widen the options; it does not inherit that lived knowledge.

You do not need our model — or an adviser — to begin

The version built for our own circumstances was complex because those circumstances and the choices under consideration warranted it. That doesn't make it a template everyone should copy. A person deciding whether to change working hours, manage an unexpected household cost, replace an unreliable vehicle or support someone they care about can begin with a notebook, a simple table and a careful conversation with AI.

Some decisions are opportunities we welcome. Others arrive without an invitation: an income change, a caring responsibility, a health problem or an unexpected bill. Both deserve clarity. Depending on the stakes, someone might also want a trusted friend, employer, lender or qualified professional involved. Access to specialist advice can be useful; it is not the entrance ticket to asking better questions.

This isn't evidence that AI makes every personal decision faster, or that financial projections can predict anyone's future. What our experience demonstrated was more modest and more valuable: a way to keep facts, assumptions, alternatives and written conclusions connected, so choices can be revisited without starting again from scratch.

Where we landed

The most useful output wasn't a clever spreadsheet. It was a better conversation — one that could accommodate real choices, changing circumstances and disagreement without losing sight of the reason we were planning in the first place.

Sometimes the answer is to act. Sometimes it's to wait, to ask somebody else, or to accept that an enjoyable experience is worth more than a higher number in a distant forecast. AI can make the trade-offs more visible, but it can't tell a person what makes a life worthwhile.

If we do this, what else moves?

That's the question I would take from the whole experience. Not as a slogan for a complex financial model, but as a practical way to think before a decision — and to return to it when life changes.

COMPANION RESOURCE — Make Better Life Decisions with AI — the practical playbook. Start with one choice and finish with a Living Decision Card and a plain-English Decision Brief you can keep, discuss and revisit.

Grey Matters · Human-led. AI-enabled. This account reflects one ongoing collaboration; it is not financial, investment, legal or tax advice.