

PREPARED BY WE COMPLEMENT

Knowledge is cheap. Judgement is not.

What financial advice firms are
paid for in the age of AI

For advice firm owners, compliance
leaders, networks and advisers



AI has made knowledge cheap.
Judgement is what clients pay for next.



About this paper

This paper sets out why professional judgement, not knowledge, is now the asset advice firms need to govern and develop. It is written for advice firm owners, compliance leaders and networks, and for advisers.

It is a discussion of professional judgement in financial advice. The examples are illustrative. It is not legal or regulatory advice.

Prepared by We Complement. Written by Tony Slimmings. September 2026.

Contents

1	Introduction
2	Knowledge has become accessible. Judgement has not.
3	What clients pay for
4	Good advisers reach different conclusions
5	Why the variation stays hidden
6	Calibrated, not replaced
7	Growing judgement, not only knowledge
8	Where AI fits
9	What to do now
10	Sources and scope

1. Introduction

Most conversations about AI in financial advice start with the wrong question: which parts of the adviser's job will the machine take? This paper starts with a different one. If AI can now explain almost anything an adviser once spent years learning, what is left that clients, firms and regulators still depend on?

The answer is judgement: deciding what matters in a client's circumstances, weighing evidence that points in different directions, and standing behind the recommendation. That part of the job is becoming increasingly valuable. It is also more variable than most firms assume.

The author wrote this paper for two readers. For firm owners, compliance leaders and networks, it sets out why judgement, not knowledge, is now the asset to govern. For advisers, it explains why their judgement is the part AI makes more important, and how to strengthen it.

The question running through this paper

If AI can now explain almost anything an adviser once spent years learning, what is left that clients, firms and regulators still depend on?

2. Knowledge has become accessible. Judgement has not.

For most of the profession's history, knowledge was the scarce resource. Tax rules, product features, pension legislation and investment theory took years to learn and effort to keep current. Clients paid, in part, for access to that knowledge.

AI has changed the economics. A client can now ask a general-purpose AI tool to explain how pension tax-free cash works, the difference between a pension transfer and a consolidation, or how an investment bond is taxed, and receive a straightforward answer in seconds. The answer will not always be right, but it will often be good enough to make information feel free.

AI can also compare evidence and propose a conclusion, and it will get better at both. That does not settle the question. A plausible AI conclusion does not relieve the adviser of understanding it, testing it against what they know about this client and owning the decision that follows. The client's circumstances, the reliability of the reasoning and the accountability for the outcome all stay with the professional. That argument holds however capable the models become.

3. What clients pay for

Clients pay for someone to decide what matters in their circumstances and to stand behind that decision. Authority, in a profession, is accountable judgement: the right to decide, joined to the obligation to answer for the decision.

In advice, that accountability is named and personal. A recommendation carries an adviser's name, a firm's permissions and, behind them, senior managers answerable under the Senior Managers and Certification

Regime. When a complaint reaches the Financial Ombudsman Service, nobody asks what the adviser knew. They ask what the adviser decided, on what evidence, and why.

That is why AI does not make advisers less valuable. As knowledge becomes cheap, the relative value of judgement rises. The adviser who can interpret evidence, reconcile competing goals, explain a trade-off in plain words and accept responsibility for the result is doing the part of the work that clients cannot get elsewhere.

4. Good advisers reach different conclusions

The remarkable thing about professional judgement is not that it varies. It is how dependable it is, given the task. Advisers decide under uncertainty, with incomplete information, competing priorities and a client's future at stake. What follows is not a criticism of advisers. It is a description of what any human being faces when deciding under those conditions.

Daniel Kahneman, Olivier Sibony and Cass Sunstein separate two kinds of error in judgement. [1]

Bias is a consistent lean in one direction: the same scale always reading ten grams heavy. It is at least predictable, and so it can be corrected.

Noise is unwanted variation: the same scale giving a different wrong reading each time. Two competent professionals look at the same facts and reach different conclusions, and the facts do not explain the difference.

The authors break noise down further, and each part will be familiar to anyone who has worked in an advice firm:

Type of noise	What it looks like in advice
Level noise	One adviser is more cautious than another across every client, so the same client gets a lower-risk portfolio from one than from the other.
Pattern noise	An adviser's own history shapes calls: the one who lived through a transfer complaint treats every transfer differently.
Occasion noise	The same adviser reaches a different conclusion on a Friday afternoon, after a difficult meeting, or at the end of a long day.

The size of the effect surprises people. In one insurance company the authors studied, executives expected their underwriters' quotes on the same cases to differ by around 10 per cent. The typical difference was 55 per cent. [1]

Different recommendations do not on their own prove noise or unreliable judgement. Two advisers can reach different conclusions for good reasons: one noticed something the other missed, or weighed a client's priorities differently and can explain why. Differences in communication style, empathy and practical wisdom are part of good advice too. The question is whether a difference can be explained by relevant evidence and defensible reasoning. Variation that can be explained that way is legitimate. Variation that cannot is noise.

5. Why the variation stays hidden

Noise is hard to see because each file, read on its own, looks defensible. The cautious adviser's recommendation has a rationale. So does the bolder adviser's. A reviewer checking one file at a time against a standard finds two compliant files and moves on. The difference between them only appears when similar cases are placed side by side, and even then comparison only shows that a difference exists. Whether it is justified depends on the evidence and reasoning behind each recommendation.

Training does not remove it either. Kahneman spent much of his career studying cognitive bias and concluded that his own intuitive judgement was no less prone to it than before. What improved was his ability to recognise situations in which errors were likely. [2] If one of the world's leading authorities on judgement could not catch his own biases as they formed, no firm can expect its advisers to do so unaided.

Hindsight makes the problem look smaller than it is. Once an outcome is known, the overlooked assumption is easy to spot. Recognising it before the recommendation is made is the hard part, and that is where the effort needs to go.

6. Calibrated, not replaced

Other professions have faced the same problem and reached the same answer: strengthen professional judgement with better evidence and better systems, rather than replace it.

For advice, that answer can be called Calibrated Judgement. It works in two steps. Comparing similar cases reveals where advisers' judgements differ. Evidence then decides whether each difference is justified. The firm's norms are provisional summaries of the evidence so far, open to challenge and updated as the evidence improves. They are not a house view to conform to.

One principle holds the approach together: **agreement is not proof of quality, and departure from the majority is not proof of error**. A firm can be consistently wrong. An adviser whose judgement differs from colleagues' may be mistaken, or may have seen something the rest missed. Calibration exists to improve how evidence and uncertainty are handled, while preserving justified dissent. A well-designed firm can tell the difference between the two, and learns from the second.

Calibration is therefore not standardisation. The aim is not identical recommendations for clients who look alike on paper. It is recommendations whose differences can be explained. Calibration also works best as small, timely prompts before a recommendation becomes final, rather than a verdict after it:

- 'Have you considered...?'
- 'Comparable cases suggest...'
- 'What evidence explains this difference?'

None of these tells the adviser what to decide. Each one improves the thinking before the decision is made.

7. Growing judgement, not only knowledge

Years of experience do not on their own produce reliable judgement. Long service shows how many decisions someone has made, not how well they were made. Judgement develops when people examine

real decisions without blame and receive specific feedback that corrects a real weakness. Counting CPD hours measures time spent, not judgement gained.

Reflection has to be designed. Looking back at actual cases, including the ones that went well, needs a regular slot, a simple format and protection from blame. Without that, it drifts back into either nothing or a hunt for someone to fault, and it rarely survives the moment a senior person's case is on the table.

A second opinion has to be ordinary. Judgement is at its most fragile at the moment it is exercised: under time pressure, with incomplete information, alone. The firms that handle this well make asking for a second view so quick and normal that it carries no stigma. The test is simple. Who in your firm is least likely to ask for a second opinion? If the answer is your most experienced adviser, the mechanism exists on paper but not in practice.

The future delivery of judgement AI is quietly removing

Much of the work through which trainee advisers and paraplanners learn their craft is exactly the work AI now does fastest: researching products and providers, preparing information, and drafting reports.

That work was never only production. It was also training:

- **Research builds technical understanding.** Working out why one product suits a client and not another teaches the trade-offs that later become instinct.
- **Preparation teaches how to read a client.** Gathering and organising someone's circumstances is how a junior learns which details matter and which do not.
- **Drafting forces reasoning into the open.** Writing a recommendation down exposes the gap between what an adviser thinks and what they can justify.

If AI absorbs those activities and nothing replaces the learning they provided, firms will produce advisers with access to more knowledge than any generation before them and less practised judgement. The efficiency gain arrives now; the cost arrives in the next ten years, when those people are not available to make the calls.

Preserving the learning

The answer is not to refuse the efficiency. It is to preserve the learning deliberately.

1. **Make juniors test the draft, not just accept it.** An AI-drafted recommendation becomes a training exercise when a junior has to check it against the source documents, find what it missed and explain why the conclusion holds or does not.
2. **Ask for the reasoning back.** Before a case moves on, the person who prepared it explains in their own words why the recommendation fits this client. If they cannot, the learning has not happened, whoever wrote the report.
3. **Keep some cases hand-built.** Deliberately working a proportion of cases from scratch, then comparing the result with the AI's version, turns the difference into feedback.
4. **Let juniors see decisions being made.** Sitting in on the moment an experienced adviser weighs a close call teaches more about judgement than any amount of finished paperwork.

Developing judgement is now a design task for the firm, not a by-product of time served. Firms that plan for it will have the advisers they need in ten years. Firms that do not will discover the gap when it is too late to close quickly.

8. Where AI fits

Kahneman once wished for a warning bell that would sound before a serious error of judgement. [2] Unaided human thinking has no such bell. For the first time, firms have the means to build something close to it: systems that compare similar cases, bring current evidence to the adviser's attention, notice an omitted consideration and prompt reflection at the moment it can still change the outcome.

The division of labour is clear. AI brings evidence, comparison and prompts. The adviser weighs them, decides and is accountable for the decision. A tool that quietly makes the recommendation, or that approves a file because a sentence has been added, has crossed a line that clients and regulators rely on. The value of AI in advice lies in making the professional's judgement better informed at the moment of decision, not in taking the decision away.

The warning bell needs testing too. It would be inconsistent to ask suppliers for evidence that their tools work, as our paper *How do you know your AI advice tool works on your files?* does, and then assume that prompts improve judgement. They have to earn that claim. The test is whether the prompts identify relevant omissions, reduce unjustified variation and help advisers reason better, without raising distracting challenges or steering everyone towards the same flawed conclusion. A prompt that makes advisers agree more has not succeeded unless it also makes them right more often.

Consistency is also a matter of fairness

The Consumer Duty requires firms to act to deliver good outcomes for retail customers, and the FCA's emphasis is on firms monitoring the outcomes their customers actually receive. [3] Our argument, not a regulatory requirement, is that unexplained variation belongs in that monitoring.

If the quality of a client's advice depends on which adviser they happened to see, or what time of day the file was written, that is worth knowing and acting on. None of this means standardising recommendations. It means being able to explain the differences between them.

9. What to do now

None of this needs new technology to start. It needs a decision to look.

For firm owners, compliance leaders and networks

1. **Run a simple noise check.** Take three or four anonymised past cases and give each to several advisers separately, with the same information. Compare how they interpret the evidence as well as what they recommend. Kahneman and his co-authors call this a noise audit. [1] The aim is not to find who is right. It is to see how wide the range is, and whether each difference can be explained by evidence and reasoning.

2. **Record the reasoning, not only the recommendation.** Many disagreements between reviewers and advisers turn out to be reasoning the adviser held but never wrote down. A file that shows why is easier to review, easier to defend and easier to learn from.
3. **Make a second opinion routine.** Build a quick, blame-free way to ask for a second view on close calls, and check who uses it. If experienced advisers never do, find out why.
4. **Look for patterns across cases.** Review similar cases side by side as well as one at a time, so that consistent leans and unexplained spread become visible.

For advisers

1. **Write down why before you decide.** Setting out the two or three reasons that drive a recommendation, before committing to it, exposes an assumption while it can still be changed.
2. **Ask for a second view on close calls.** Treat it as what careful professionals do, not as an admission of doubt.
3. **Notice the conditions you decide in.** Fatigue, time pressure and a difficult previous meeting all move judgement. If a recommendation was finalised late on a Friday, look at it again on Monday.
4. **Know your own pattern.** Ask where your recommendations tend to sit compared with colleagues on similar clients. A consistent lean is not wrong in itself, but it needs to rest on evidence rather than habit.

The aim is not a firm of identical advisers. It is a firm that knows how its advisers' judgements vary, can explain that variation, and keeps improving it.

Closing thought

AI has made knowledge cheap. Judgement, exercised well and accountably, is what clients will pay for next.

10. Sources and scope

1. Daniel Kahneman, Olivier Sibony and Cass R. Sunstein, *Noise: A Flaw in Human Judgment* (William Collins, 2021). System noise is introduced in chapter 1, including the insurance underwriting study; level, pattern and occasion noise are set out in Part II; the noise audit is described in Appendix A.
2. Daniel Kahneman, *Thinking, Fast and Slow* (Allen Lane, 2011), 'Conclusions', on his judgement remaining prone to bias, recognising situations in which errors are likely, and the absence of a warning bell before serious errors.
3. Financial Conduct Authority, Consumer Duty, Principle 12 and PRIN 2A: firms must act to deliver good outcomes for retail customers, and monitor the outcomes customers actually receive (PRIN 2A.9).

This paper is a discussion of professional judgement in financial advice. The examples are illustrative. It is not legal or regulatory advice.