



- 01 Human-centred capability
- 02 Ethics as a lens
- 03 Designed for people
- 04 Building independence
- 05 Radical transparency

Our AI Principles

How we think about AI — and how we use it

David Kolb Consultancy Ltd

Strategic AI Advisory

AI enhances human capability — it doesn't replace human judgment



We believe AI is most powerful when it amplifies what people do well.

We believe AI is most powerful when it amplifies what people do well: recognising patterns across large amounts of information, generating options worth considering, handling the repetitive so people can focus on the consequential. But the decisions that matter — the ones with real consequences for real people — require human judgment, context, and accountability.

In our own work, every AI-assisted output is reviewed, validated, and owned by a human. We don't hand over final judgment to a model. In our advisory work, we help organisations design the same approach: AI systems with humans genuinely in the loop, not as a rubber stamp but as an active participant.

This isn't caution for its own sake. It's a practical recognition that probabilistic systems produce probabilistic outputs. The organisations getting lasting value from AI are the ones who have designed human verification into the workflow from the start — not bolted it on afterwards.

Ethics is not a constraint on AI — it's part of what makes it work



We take the ethics of AI seriously — because we've seen what happens when it isn't.

Bias embedded in training data. Outputs that feel authoritative but are factually wrong. Automation that removes the human accountability that customers and regulators depend on.

Ethical AI, in our view, means being honest about what a system can and can't do. It means designing for fairness and inclusion, not assuming neutrality. It means being transparent with clients, students, and partners about when AI is involved in producing work. And it means not recommending tools or approaches we wouldn't be comfortable operating under ourselves.

We actively bring ethics into every engagement — not as a separate workstream, but as a lens on the work itself. In financial services, this has meant defining clear boundaries around what AI can recommend and what requires human sign-off. In civic and public sector work, it has meant exploring how AI can serve inclusion and accountability rather than undermine it. The context changes; the commitment doesn't.

We believe in human-centred AI — designed around how people actually work



The organisations that get lasting value from AI design it to fit how people think, decide, and take responsibility.

Human-centred AI means starting with the human workflow, not the model capability. It means asking what decisions people need to make, what information they need to make them well, and where AI genuinely helps versus where it introduces noise or erodes the judgment it was supposed to support.

It also means being honest about impact. AI changes what work looks like — for individuals, for teams, and for the people on the receiving end of AI-informed decisions. We think those are conversations worth having directly, not managing around.

The organisations furthest along are the ones that had real conversations about what AI means for their people, early on, and designed with that in mind. We help facilitate those conversations — and then build the systems that follow from them.

The goal is capability — not dependency



We think about success differently to most consultants.

A successful engagement, for us, isn't one where the client needs us again next quarter. It's one where the client's own people can evaluate, govern, and deploy AI with confidence — without waiting for external help every time.

This means we build frameworks people can use independently. It means we teach rather than just advise. It means we're explicit about our reasoning, so clients can apply the thinking to the next challenge themselves.

The best outcome is an organisation that genuinely doesn't need us anymore — because we've helped them build what they needed internally. It's a longer-term bet, and it creates a different kind of relationship. The organisations that come back to us do so because they trust the thinking, not because they're dependent on it.

We are transparent about how we use AI — including in our own work



We use AI tools in our research, analysis, content development, and advisory preparation.

We think it is important to be open about this — and it would be inconsistent for a Strategic AI Advisory to be otherwise. We use Claude (Anthropic), accessed under enterprise terms with appropriate data protection agreements.

What we are clear about is the boundary: AI assists our thinking; it doesn't replace it. Client-sensitive information is handled in accordance with our Information Security and AI Acceptable Use policies. All outputs that reach clients have passed through human review and professional judgment. We take full responsibility for the quality and accuracy of our work, regardless of the tools involved in producing it.

We bring the same expectation of transparency to client work. When AI is involved in a process or output, we say so. We don't think that transparency diminishes the work — we think it's part of what makes it trustworthy.



*These principles aren't a destination. AI is moving fast,
and our thinking moves with it. What stays constant is the
underlying commitment: to use AI in ways that are honest,
human-centred, and worth trusting.*

David Kolb

Director, David Kolb Consultancy Ltd

1 July 2025