



Safe and Effective AI in Regulated Debt Advice: Adviser Guide to AI

Understanding and working with artificial intelligence
A practical guide for frontline debt advice professionals

This guide is written for you: a debt advice professional. It is designed to help you feel informed and confident about AI, not anxious. A machine cannot replicate your professional judgement, relational skill and lived experience.

Adviser Guide to AI v1.0
March 2026

Contents

Contents

1	Introduction.....	3
2	Understanding Generative AI.....	5
3	Where AI fits into Debt Advice	9
4	Working alongside AI.....	13
5	When clients have used AI.....	16
6	What our survey reveals about consumer views on AI for financial advice.....	18
7	Ethical and professional boundaries	19
8	Conclusion.....	21

1 Introduction

AI is already part of our world

While there has been a lot of hype around AI, it is not new and not as mysterious as marketers would have you believe. If you have ever asked your phone for directions, unlocked it with your face or had Netflix choose what you watch, you have already used artificial intelligence. It is a set of technologies that have been quietly woven into everyday life over many years. What has changed recently is the arrival of generative AI, which can hold conversations, write text and answer questions in a way that feels (almost) human.

This guide exists because AI is becoming increasingly relevant to the work you do, the organisations you work within, and the people you support. Some of your clients may already be using AI tools before they reach you. Your organisation may be exploring how AI can support your work, and the wider debt sector is actively discussing how these technologies fit within a regulated advice environment.

None of this means your role is under threat. Extensive sector research conducted for this project, involving 28 interviews with debt advice professionals, a national public survey of over 2,000 people and a sector-wide professional survey, reached a clear and consistent finding: the professional judgement, interpersonal skill and lived experience that you bring to your work cannot be replicated by a machine. Debt advice is fundamentally human work.

Why does this guide matter now?

Since 2022, generative AI tools like ChatGPT, Microsoft Copilot and Google Gemini have become publicly available. The YouGov survey conducted for this project found that 37% of UK adults have used an AI chatbot for personal purposes. While client use of AI for debt advice remains relatively low, it is growing. At the same time, organisations across the sector are exploring how AI might support adviser work. Understanding these technologies is not about replacing what you do, it is about equipping you to work alongside them confidently.

AI you may already be using

Most people interact with AI dozens of times a day without being aware of it:

- Spam filtering: every time your inbox catches a phishing email, AI is making that decision
- Predictive text: your phone suggesting the next word as you type is AI predicting what you are likely to say
- Streaming recommendations: when Netflix suggests a programme, AI is analysing patterns across millions of users
- Banking fraud alerts: your bank texts about an unusual purchase uses AI to flag transactions that do not match your patterns
- Sat nav rerouting: when your app diverts around a traffic jam, AI is processing live data from thousands of drivers

- Voice assistants: Alexa, Siri and Google Assistant all use AI to understand and respond to spoken language

AI is not something coming over the horizon. It is already here, and you have been using it comfortably for years. What is new is generative AI, which can create content, hold conversations and respond to complex questions.

2 Understanding Generative AI

What is generative AI?

Traditional AI analyses data and makes predictions or classifications. It can tell you whether an email is spam or what the fastest route home is. It works with existing information.

Generative AI is different. It creates new content, text, images, code and audio, based on patterns learned from vast amounts of existing material. When you ask ChatGPT a question, and it responds with a fluent paragraph, it has not looked up that answer in a database. It has generated it, word by word, by predicting the most likely next word based on patterns in its training data.

This distinction matters. Generative AI does not know things the way a human does. It does not understand meaning, context or consequences. It produces statistically probable outputs. This is why it can sound confident and authoritative even when it is wrong, a phenomenon known as hallucination.

How large language models work

The technology behind most generative AI is called a large language model (LLM). These are trained on enormous amounts of text to learn statistical relationships between words. Think of it like a highly sophisticated autocomplete: your phone might suggest “you” after “thank” and an LLM does the same thing but at vastly greater complexity, generating entire paragraphs that read as though a knowledgeable person wrote them.

This is powerful, but LLMs do not verify facts, do not access live databases (unless specifically connected to them), do not understand legal or regulatory implications and cannot detect vulnerability, distress or the unspoken context behind a question.

How organisations make AI more reliable

Given these limitations, you might reasonably ask why anyone would deploy AI in a regulated environment. The answer is that organisations do not use raw, general-purpose AI for professional applications. They apply a range of techniques to make AI significantly more accurate and controlled. Understanding these techniques helps you ask better questions about the tools being introduced in your workplace and gives you a clearer sense of what “safe” AI deployment actually involves.

The most important technique is retrieval-augmented generation, commonly abbreviated as RAG. Instead of allowing the AI to draw on everything it learned during training (which may be outdated, inaccurate or irrelevant), RAG connects the AI to a specific, curated collection of documents. When someone asks a question, the system first searches the approved knowledge base, retrieves the most relevant content, and then uses it to generate its answer. Think of it as the difference between asking someone to answer a question from memory versus asking them to

look it up in an approved reference book first. The reference book approach is far more likely to produce an accurate, current and relevant answer.

For debt advice, this means an AI tool could be connected to your organisation's policy documents, current legislation, creditor protocols and procedural guides rather than drawing from the open internet. This does not eliminate the risk of error, but it substantially reduces it and makes the tool's responses auditable, as you can trace the source material it used.

Alongside RAG, organisations use guardrails: rules built into the system that control what the AI can and cannot do. These might prevent the tool from recommending specific debt solutions, require it to include disclaimers, force it to escalate certain topics to a human adviser or block it from processing particular types of sensitive information. Guardrails are how developers draw the line between what an AI tool handles on its own and what requires professional intervention.

System prompts are another layer. These are hidden instructions given to the AI by its developers that shape its behaviour, tone and scope. A system prompt for a debt advice tool might instruct the AI always to use UK-specific terminology, never present information as personal advice, always signpost to regulated services for complex queries, and respond with empathy when financial distress is indicated. You will not see the system prompt when you use a tool, but it significantly influences the quality and safety of the responses you receive.

Finally, some organisations choose to fine-tune AI models. This involves further training a general-purpose model on a specific dataset, such as thousands of case notes or adviser interactions, so that it becomes more attuned to the language, patterns and expectations of debt advice work. Fine-tuning is more resource-intensive than RAG and is typically used by larger organisations. Still, it can produce tools that feel significantly more natural and accurate within a specific professional context.

None of these techniques makes AI infallible. They make it significantly better than a raw, unconstrained model, but the combination of human oversight with technically grounded AI is what creates genuinely safe deployment. This is why your role in reviewing, checking and validating AI outputs remains essential regardless of how sophisticated the underlying technology becomes.

What is agentic AI?

More recently, the technology has moved further with agentic AI. Standard generative AI responds to prompts: you ask, it answers. Agentic AI can take actions. It can break a complex task into steps, use external tools, make decisions about what to do next and execute multi-step workflows with minimal human input.

In a debt advice context, agentic AI could theoretically gather financial information, cross-reference eligibility criteria, draft creditor correspondence and schedule follow-up actions, all without an adviser directing each step. This is why agentic AI raises particularly significant questions for regulated environments. The more autonomous a system becomes, the greater the need for oversight and governance.

These developments have occurred alongside capacity pressures within the advice sector. Many organisations rely on a combination of professional advisers, trainees, volunteers and triage staff. Frontline services frequently report high demand relative to available specialist capacity, leading organisations to explore ways to increase efficiency while maintaining quality and safety.

Key terminology

Artificial Intelligence (AI)	The broad field of computer systems that perform tasks that normally require human intelligence.
Large Language Model (LLM)	The type of AI model behind tools like ChatGPT and Gemini. Trained on vast text to generate human-like responses.
Generative AI	AI that creates new content rather than simply analysing existing information.
Agentic AI	AI that can take autonomous actions, use tools and complete multi-step tasks with limited human direction.
Hallucination	When AI generates information that sounds plausible but is factually incorrect.
Prompt	The question or instruction you give to an AI tool.
Training data	The information used to build an AI model has a cut-off date and may contain biases.
Knowledge base	A curated, approved set of information that an AI tool draws from, as distinct from open internet sources.
Retrieval-Augmented Generation (RAG)	A technique where an AI tool is connected to a specific, curated set of documents or data so that its answers are grounded in approved content rather than its general training. This is how organisations make AI more accurate and relevant for their particular context.
Guardrails	Rules and restrictions that are built into an AI system to prevent it from producing harmful, inappropriate or out-of-scope responses. In debt advice, guardrails might prevent the tool from recommending specific solutions or ensure it always directs users to a human adviser for complex situations.
Fine-tuning	The process of further training an existing AI model on a specific set of data to make it more specialised. An organisation might fine-tune a model on debt advice content so it responds more accurately to sector-specific questions.
Grounding	Connecting an AI system to verified, up-to-date sources of information so its responses are based on factual content rather than patterns from training data alone. A grounded system is less likely to hallucinate because it references real documents.
Token	The basic unit of text that an AI model processes. Roughly equivalent to three-quarters of a word. Tokens matter because AI tools have limits on how much text they can process in a single interaction, which affects how much context they can consider.

System prompt	Hidden instructions given to an AI tool by its developer that shape how it behaves, what tone it uses and what it will or will not do. You do not see these, but they significantly influence the responses the tool gives.
Human-in-the-loop	A design principle where a human must review, approve or intervene at critical points in an AI process. In regulated debt advice, human-in-the-loop means an adviser checks AI outputs before they are applied to a client's case.
Application Programming Interface (API)	A connection that allows different software systems to communicate with each other. When your organisation's case management system connects to an AI tool, it does so through an API
Natural language processing (NLP)	The branch of AI concerned with enabling computers to understand, interpret and generate human language. It is the foundation technology behind chatbots, voice assistants and text analysis tools.
Data processing agreement (DPA)	A legal contract between your organisation and a technology provider that governs how client data is handled, stored and protected. Essential for any AI tool that processes personal information

3 Where AI fits into Debt Advice

The sector research identified that AI in debt advice is not a single use case. It can be deployed for a wide range of purposes, each with varying levels of risk, safeguard requirements, and implications for your role.

The Debt Advice AI Roadmap, developed through this research and cross-checked with practitioners from across the sector, provides a structured way to understand these possible applications. It is organised into three levels, each building on the one before it. Understanding these levels matters because they show where AI is most likely to affect your work, where it carries the greatest risk and why certain applications need your professional oversight while others are relatively straightforward.

Level 1: Adviser Assistance

This is AI that helps you work more efficiently. It does not interact with clients directly. It uses generative AI to produce content that you review before use. It is the safest starting point and the area where the sector sees the most immediate value.

- **1.1 Administrative assistance:** transcription, case note summarisation, letter drafting. You review everything before it is used
- **1.2 Knowledge and compliance support:** policy lookups, legislation decoding, compliance checking. Particularly helpful for newer advisers or unfamiliar scenarios, but still requires professional understanding and knowledge to assess relevance to specific cases
- **1.3 Operational intelligence:** trend analysis, QA support, caseload management. Primarily used by management and quality teams

In the sector survey, adviser-support tools were rated as suitable for use now by 82% of respondents, the strongest endorsement of any AI application within the sector. Several organisations in the research are already piloting these tools with positive results.

“The time it saves is really good. The staff are loving it. Every subsequent communication goes through AI, and it works well.”

Adviser Supervisor

Level 2: Client Self-Serve: Information and Assistance

This is AI that interacts directly with people seeking help with debt, or helps clients with practical tasks alongside the advice journey. It uses generative AI to provide information, answer questions and assist with tasks. Risk increases as the tool moves from providing general information to influencing individual decisions.

- **2.1 General information:** factual Q&A, explainers, process walkthroughs. Lower risk with proper governance, including pre-implementation testing
- **2.2 Contextual guidance:** responds to individualised questions using limited context
- **2.3 Structured triage:** screens for urgency, routes users, flags vulnerability
- **2.4 Pre-appointment data gathering:** collects financial information for your review
- **2.5 Communication support:** helping clients draft creditor letters and prepare for conversations
- **2.6 Documentation support:** helping clients organise paperwork and understand correspondence
- **2.7 Ongoing engagement:** reminders, progress tracking, re-engagement prompts.

Your clients are already using general-purpose AI for some of these tasks. One adviser described clients sending them AI-drafted creditor letters for review. Sector-specific tools could improve the quality and safety of these use cases.

Level 3: Agentic AI: Taking action and sustained relationships

This is AI that takes actions or maintains an ongoing, evolving relationship with clients. It is fundamentally different from Levels 1 and 2 because it does things rather than just generating content. It is largely forward-looking and carries the highest risk. The tools listed here are possible applications of agentic AI. Individual organisational assessment of user needs and use cases will determine how suitable these are. This may change in the future as agentic AI advances at pace.

- **3.1 Personalised advisory:** ongoing support using stored information, scenario modelling, sustained advisory influence. Not currently appropriate for most organisations
- **3.2 Adviser-directed actions:** AI executes tasks under your instruction, such as filing forms or sending approved correspondence
- **3.3 Semi-autonomous actions:** AI acts within pre-approved boundaries with approval gates
- **3.4 Autonomous process management:** AI acts autonomously based on information supplied by the client. Findings of this research indicate that this is not currently appropriate for regulated debt advice

You may have already encountered level 3 tools, or may in the near future. The key principle, identified by the sector, is that every consequential action, anything affecting a client's legal position, financial obligations or access to services, must have human approval before it happens.

Here is where agentic or semi-agentic AI is already being used in other regulated sectors:

- **Healthcare:** NHS trusts across England are now deploying AI scribing tools that transcribe clinical consultations and draft summarised notes for clinicians to review

before they are entered into patient records. A major NHS England-sponsored trial across nine London sites evaluated over 17,000 patient encounters and found that clinicians spent nearly a quarter more time directly with patients when using the technology (GOSH/NHS England, 2025). NHS England published national guidance on AI scribing in April 2025 and in January 2026 launched a registry of 19 approved suppliers, actively encouraging adoption. In primary care, AI triage tools are also being piloted: an NHS-backed evaluation of the Rapid Health Smart Triage system found a 73% reduction in GP waiting times, with the tool assessing symptoms and routing patients to appropriate care within clinical guardrails (Integrated Care Journal, 2024)

- **Legal services:** AI-powered contract review, risk flagging and document drafting are now widespread across UK law firms. A 2025 industry analysis found that 93% of mid-sized UK firms use AI in at least one workflow, with platforms reporting 50-90% reductions in review time (Compare the Cloud/LexisNexis, 2025). The Law Society has published guidance noting the emergence of agentic AI systems in legal practice that can autonomously break tasks into components and interact with documents and other software. Tools like Harvey AI (used by A&O Shearman and PwC) and Luminance (developed in Cambridge with a proprietary legal language model) are in active use. In all cases, outputs operate under solicitor supervision and do not provide legal advice directly
- **Banking and financial services:** Major UK lenders are actively testing agentic AI for customer-facing services. NatWest is piloting AI agents to accelerate complaint handling, Lloyds has launched an employee-facing pilot to help consumers manage their finances, and Starling plans to offer personalised budgeting tools with predictive spending caps, all under FCA oversight (CeFPro, 2025). The FCA's Chief Data Officer confirmed that agentic AI could enter retail banking applications from early 2026, with the regulator applying existing rules, including Consumer Duty and the Senior Managers Regime, to hold firms accountable. Globally, a Capgemini survey found that the leading use cases for AI agents in banking are customer service (75% of firms), fraud detection (64%) and loan processing (61%) (Capgemini Research Institute, 2025)
- **Insurance:** Insurers are deploying agentic AI to process straightforward claims with minimal human intervention. AXA has more than 60 agentic use cases in testing or partial deployment, particularly in contact centres, underwriting and claims handling (Computing, 2026). Allianz launched an agentic AI solution in mid-2025 that processes low-complexity claims end-to-end, from notification through document verification to payment, with complex or disputed claims escalated to human handlers (Allianz, 2025). A 2026 industry survey found that UK insurers reported the highest AI-driven cost reductions of any sector at 21%, with 40% of companies already using AI in claims and risk management (Insurance Times/EXL, 2026)

- **Local government:** AI chatbots are now used by at least 59 councils across the UK and Ireland for routine enquiries on council tax, housing, waste collection and benefits (EBI, 2024). The London Borough of Barking and Dagenham's AI chatbot reduced phone calls by 1,000 per month while maintaining an 85% customer satisfaction score, freeing the telephone service to focus on vulnerable residents (Equality and Human Rights Commission, case study). Derby City Council was the first UK local authority to deploy a generative AI-powered 24/7 digital front door, handling 50,000 calls with a 56% deflection rate and including vulnerability and emotional cue detection (Government Technology, 2025). These typically operate at the general information and contextual guidance level, with handoff to human caseworkers for anything involving a decision or assessment of individual circumstances

The common thread across all these examples is that AI handles routine, structured, and lower-risk tasks, while human professionals retain oversight of complex judgements, vulnerable individuals, and consequential decisions. This is the model most likely to emerge in debt advice.

Sources for further reading

GOSH/NHS England AI scribe trial (2025): <https://www.gosh.nhs.uk/news/researchgosh-led-trial-of-ai-scribe-technology-shows-transformative-benefits-for-patients-and-clinicians-across-london/>

NHS England ambient scribing guidance (2025): <https://www.england.nhs.uk/long-read/guidance-on-the-use-of-ai-enabled-ambient-scribing-products-in-health-and-care-settings/>

Rapid Health Smart Triage evaluation (2024): <https://integratedcarejournal.com/nhs-backed-study-shows-73-reduction-gp-waiting-times-ai-triage/>

Law Society guidance on AI in contract review: <https://www.lawsociety.org.uk/topics/ai-and-lawtech/ai-use-cases-in-law-legal-contract-review>

Mid-market UK law firm AI adoption (2025): <https://www.comparethecloud.net/articles/uk-mid-market-legal-firms-ai-document-review-luminance-harvey-sra-compliance>

NatWest, Lloyds and Starling AI trials (2025): <https://connect.cefpro.com/article/view/fca-warns-banks-as-agentic-ai-nears-consumer-rollout>

Capgemini World Cloud Report for Financial Services (2025): <https://www.capgemini.com/us-en/news/press-releases/banks-and-insurers-deploy-ai-agents-to-fight-fraud-and-process-applications-with-plans-for-new-roles-to-supervise-the-ai/>

AXA agentic AI (2026): <https://www.computing.co.uk/feature/2026/how-axa-using-ai-rewrite-insurance-model>

Allianz Project Nemo (2025): <https://www.allianz.com/en/mediacenter/news/articles/251103-when-the-storm-clears-so-should-the-claim-queue.html>

UK insurance AI cost reductions (2026): <https://www.insurancetimes.co.uk/expert-views/ai-reality-in-the-uk-and-europe-moves-from-pilot-to-production-exl/1457555.article>

Local government AI assistants (2024): <https://ebi.ai/blog/local-government-ai/>

Barking and Dagenham chatbot: <https://www.equalityhumanrights.com/guidance/artificial-intelligence-case-studies-good-practice-local-authorities>

Derby City Council digital front door (2025): <https://www.governmenttechnology.co.uk/features/leading-ai-revolution-local-government-derby-city-council-story>

4 Working alongside AI

Where AI can help you

The strongest consensus from the research was around tools that reduce your administrative burden:

- Case note drafting: AI that listens to or reads your client interactions and produces structured case notes for you to review and edit
- Letter drafting: Producing first drafts of creditor correspondence that you adjust and send
- Policy lookups: Quickly finding relevant guidance, procedural requirements or specific rules
- Information gathering: Structured data collection during initial contact, improving the starting point (though you will still need to reconstruct the full picture)

The key principle: AI as your assistant, not your replacement

Think of AI as a very fast, very thorough administrative assistant that never gets tired and can draft documents quickly. But this assistant has no professional judgement, cannot read the room, does not understand vulnerability, and will sometimes confidently get things wrong. You remain a qualified professional. AI is a tool in your hands.

Where AI falls short and why you still matter

AI is increasingly capable in structured tasks, but it cannot yet perform reliably enough for the high-stakes, human-centred work that defines your role.

- Detecting hidden vulnerability: AI can spot some signals in data, but it cannot recognise vulnerability that emerges gradually through trust-building, understand the shame and fear behind incomplete disclosure, or interpret the subtle cues that alert an experienced adviser to underlying issues
- Replacing professional judgement: AI can apply structured rules to straightforward scenarios, but falls short on the interpretive, contextual, experience-informed judgement that accounts for the full complexity of a real person's life
- Building an advice relationship: the trust that develops between you and your clients over months or years, the trust that allows someone to disclose a gambling problem or admit to financial abuse, is something AI genuinely cannot replicate
- Guaranteeing accuracy: AI connected to curated knowledge bases can be very accurate, but in a regulated environment where incorrect guidance can cause real harm, the current margin for error remains too high for unsupervised use in complex client use cases

- Understanding what a client is not saying: AI can flag data gaps but cannot intuit the human reasons behind them, or draw on the professional intuition built through years of practice

Practical tips for working with AI tools

- Always review AI-generated content before it reaches a client or creditor. You are professionally accountable for everything that goes out under your name
- Treat AI outputs as a starting point to edit and improve, not a finished product
- Be cautious with client data. Only use AI tools that your organisation has approved. General-purpose tools like ChatGPT are not appropriate for client data unless specific safeguards are in place
- Ask questions about new tools. How does it work? What data does it use? What are your responsibilities?
- Report errors. Your professional knowledge is essential for identifying where AI gets things wrong
- Watch for over-reliance. If you are a newer adviser, make sure AI tools support your learning rather than substitute for it. The ability to construct a budget, identify vulnerability and navigate creditors develops through practice

Understanding how AI output quality works

Even if you are not writing prompts yourself, understanding how AI output quality is determined will help you make sense of why tools sometimes produce excellent results and sometimes fall short.

AI output quality depends on three things: the quality of the underlying model, the quality of the knowledge it has been connected to and the quality of the input it receives. In tools designed for professional use, the first two are handled by the developers. The third, the input, is where you and your clients have the most influence.

A prompt is simply the question or instruction given to an AI tool. The more specific and structured the input, the better the output. This is why structured data collection at the front end of a process matters so much. If a self-serve tool asks a client vague or poorly sequenced questions, the AI will have incomplete information to work with, and its responses will reflect that. Equally, if an adviser asks an AI tool to "summarise this case" without specifying what aspects matter, the summary may miss critical details.

In practice, this means that when you use AI tools in your work, being clear and specific in your instructions will produce noticeably better results. If you are reviewing AI-generated content that seems off, the issue may not be the AI itself but the information it was given to work with. Feeding back on these patterns helps your organisation improve how its tools are configured.

Where your data goes: understanding AI processing

When your organisation tells you a tool has been "approved", it helps to understand what that means in practice. There are broadly three ways AI tools process data, and each carries different implications for client confidentiality.

Cloud-based AI tools send data to external servers for processing. Most commercial AI tools, including ChatGPT in its free version, work this way. Data leaves your device, is processed on the provider's servers, and a response is returned. In many free-tier products, the data may also be used to improve the model, which means client information could become part of the AI's future training. This is why free, general-purpose tools are not appropriate for client data without explicit organisational approval and a data processing agreement.

Enterprise AI agreements are different. When an organisation purchases a business licence for an AI tool, the contract typically includes provisions that the provider will not use client data for model training, will store and process data in specific jurisdictions and will comply with UK data protection requirements. This is the level of agreement your organisation should have in place before any tool is used with client information.

On-premise or private deployment means the AI runs on your organisation's own servers or within a private cloud environment. Data never leaves the organisation's infrastructure. This is the most secure option but also the most expensive and technically demanding, and is typically only available to larger organisations.

If you are unsure which model your organisation uses for a particular tool, ask. Understanding whether client data is processed locally, within an enterprise agreement or sent to a public service helps you make informed decisions about what information you share with the tool.

5 When clients have used AI

Some clients will arrive having consulted AI tools about their financial situation. The YouGov survey found that while only 21% of the public would use AI for financial advice, those who do tend to treat it as a starting point. Only 1.7% said they would follow AI advice exactly. But that still means some people will act on what AI tells them and arrive at your door with expectations shaped by that interaction.

What clients might have been told

- Technically plausible but contextually wrong: An AI tool might correctly describe a Debt Relief Order (DRO) while failing to assess whether the client qualifies. Advisers have reported clients arriving saying the AI told them they could get a DRO when they were not eligible
- Confidently presented regardless of accuracy: Generative AI always sounds authoritative. It does not signal uncertainty. For a client in crisis looking for reassurance, a confident-sounding answer can feel compelling. The sector survey confirmed that 94% of professionals agreed that stressed clients accept any confidently presented option
- Based on incomplete information: AI can only work with what the client provides. Without full financial details, enforcement context and personal circumstances, the AI produces answers based on an incomplete picture
- Jurisdiction-nonspecific: General-purpose AI tools are unlikely to apply the correct jurisdiction-specific rules. Researchers testing AI tools found US-specific content returned for English queries, outdated DRO fee information and incorrect jurisdictional assumptions

“Someone went bankrupt because the AI bot told him to. And he’s now saying he wants to sue the AI.”

Team Leader

How to respond

The skill of managing a client who arrives with fixed expectations from AI is essentially the same skill you already use when clients arrive influenced by Individual Voluntary Arrangement (IVA) advertising or advice from friends. The approach that worked for experienced advisers in the research was:

- Acknowledge what they have done: Recognise that the client has tried to help themselves. This builds rapport rather than making them feel foolish
- Do not dismiss the AI outright: Saying the AI is rubbish will put the client on the defensive. Instead, explain that AI gives general information but cannot see the full picture

- Reframe gently: Explain that your role is to look at their specific situation in detail, which the AI could not do. Position your assessment as a more thorough version of what they have started
- Address specific misconceptions directly: If the AI has told them something incorrect, explain clearly why it does not apply to their situation. Use concrete examples
- Be aware of emotional investment: Some clients will have spent significant time with AI and feel invested in the outcome it suggested. Changing their mind may take patience

Emerging use cases to be aware of

The research identified several specific ways clients are beginning to use AI:

- Drafting creditor letters: Clients using ChatGPT to write holding letters to creditors, then sending them to advisers for review. Generally adequate for standard correspondence, but not suitable for consequential communications
- Researching debt solutions: Clients arriving with detailed but sometimes inaccurate knowledge about DROs, IVAs and bankruptcy, shaped by AI responses that may reflect outdated information or marketing-influenced content
- Filing applications: At least one reported case of a client filing for bankruptcy based on AI guidance, without understanding the full implications
- Emotional support: Some clients are turning to AI for reassurance about their financial situation. While not inherently harmful, this raises concerns about AI providing comfort that delays engagement with professional advice

Remember: you have already handled this dynamic

Clients arriving with preconceived ideas is not new. You have always dealt with people influenced by advertising, friends, or online content. AI is a new source of prior information, but the professional skill required to manage it, patient explanation, careful assessment and building trust, is something you already possess.

6 What our survey reveals about consumer views on AI for financial advice

Understanding how the general public views AI can help you anticipate client attitudes and respond appropriately. A nationally representative YouGov survey of 2,060 UK adults was conducted as part of this research. Here are the key findings that are most relevant to your work:

- 37% have used AI chatbots for personal advice (as opposed to only in a work context), with higher usage among younger adults. But 59% have never used one
- Only 21% would use AI for advice on financial issues (6% very likely, 14% fairly likely). 47% said not at all likely
- 72% do not trust AI accuracy for financial information. Only 3% expressed any preference for AI over human advice
- 80%+ prefer human advice. 60% strongly prefer it. This is consistent across age groups
- Only 1.7% would follow AI financial advice exactly. Most would use it as a starting point and check with official sources or human experts
- 28% said nothing would increase their trust in AI for financial matters
- The strongest trust factor was regulation by financial authorities (32%), followed by human expert review (27%)

What this means for you: most of your clients will either not have used AI or will have used it cautiously. The minority who have acted on AI guidance may need particularly careful handling. The data strongly reinforces the value of your role as the trusted human expert that the public says they want.

7 Ethical and professional boundaries

Data protection and client confidentiality

Debt advice involves handling deeply sensitive information. How this interacts with AI is a critical question:

- Never enter client data into unapproved tools. General-purpose AI tools in their free public versions are not designed for sensitive data. Unless your organisation has specifically approved a tool, do not enter client information
- Understand where data goes. Information entered into AI may be stored, used for model improvement or processed overseas. Your organisation should tell you whether tools have been assessed under UK GDPR
- Apply the same standards you always have. Data protection principles, purpose limitation, data minimisation, security, and accountability apply equally when AI is involved

Professional accountability

A fundamental principle: if AI helps you produce something, you are still responsible for it. Whether it is a case note, a creditor letter, a financial assessment, or a recommendation, professional accountability rests with you and your organisation, not with the AI.

The FCA has confirmed that existing regulatory frameworks, including Consumer Duty, apply fully to AI-involved services. Organisations cannot delegate accountability to technology. For you, this means continuing to check facts, question outputs and use your judgement before acting on anything AI produces.

Transparency with clients

- Be open about AI use. If your organisation uses AI in any part of its processes, clients should be informed clearly
- Make the human role visible. Clients should understand that a qualified adviser reviews their case. Where AI has contributed, human oversight should be communicated
- Avoid over-reliance on signals. If clients feel their advice comes from a computer, trust erodes. Transparency about AI's supporting role, combined with visibility of your professional involvement, maintains confidence

Bias and fairness

AI systems learn from historical data, which means they can inherit existing biases. You are not expected to audit AI for bias, but you should be aware that AI is not inherently neutral and should

trust your instincts if an output feels wrong, incomplete or unfair. Your role as a human check on AI is one of the most important safeguards in the system.

The boundary between information and advice

One of the most important distinctions in debt advice, and one that AI complicates, is the line between general information and regulated advice. AI can explain what a DRO is. It cannot assess whether a DRO is right for a specific person. The distinction is critical because clients may not recognise it. A confident AI response about debt solutions can feel like personalised advice even when it is generic information.

The regulatory position

The FCA has noted that the Financial Ombudsman Service or FOS does not cover advice from general-purpose AI tools. If a client acts on AI guidance and suffers harm, they may have no recourse. This is why professional debt advice, delivered by qualified advisers within regulated frameworks, remains essential.

Maintaining your professional development

There is a legitimate concern that over-reliance on AI could affect skill development, particularly for newer advisers. The ability to construct a sustainable budget, identify hidden vulnerability and navigate creditor negotiations develops through practice. If AI handles too much cognitive work too early, junior advisers may not develop the depth of expertise that experienced colleagues possess. AI should support your learning, not substitute for it.

8 Conclusion

If there is one message to take from this guide, it is this: AI is a tool, and you are the professional who decides how and when to use it.

The work you do - sitting with someone in financial crisis, building trust over weeks and months, reading between the lines, identifying hidden vulnerabilities, constructing realistic budgets, navigating complex regulations, guiding people toward sustainable solutions - is deeply, fundamentally human. It requires empathy, judgement, experience and relational skills that no AI system can replicate.

The sector research reached a clear conclusion: debt advice is structurally complex, iterative, relational and dependent on professional interpretation developed through years of practice. The question is not whether AI will replace advisers; it will not, but how these technologies can be integrated responsibly to support the work you already do.

AI may help you write case notes faster. It may help clients gather information before they reach you. It may help your organisation manage demand more efficiently. These are genuinely valuable contributions. But the heart of debt advice - the professional relationship between you and your client - remains yours.

Stay curious about AI. Ask questions. Engage with training. Provide feedback when tools get things wrong. And continue to bring the expertise, care and judgement that your clients depend on.

You are not being replaced. You are being supported.

A note from the research team

The research that informed this guide involved extensive interviews with experienced debt advisers and sector leaders. One of the strongest and most consistent findings was the depth of respect for the professional skill involved in debt advice. Every contributor emphasised that the expertise, relational ability and professional judgement of frontline advisers are the foundation on which effective debt advice rests, and that any use of AI must be built around that foundation, not in place of it.