



Safe and Effective Generative AI in Regulated Debt Advice

Evidence Report

Exploring the role of generative AI-enabled self-serve tools in a complex,
vulnerability-rich and regulated environment

Evidence Report v1.0
March 2026

Contents

1	Executive Summary	3
2	Introduction and Context	8
3	Methodology	12
4	Debt Advice Sector characteristics and the implications for Gen AI self-serve tools	14
5	Additional risk factors	22
6	Opportunities for Gen AI in regulated debt advice	24
7	Safeguards for more robust deployment of Self-Serve	33
8	Strategic Implications	36
9	Conclusion	40
	Appendix 1: Sector Perspectives on AI Readiness	41
	Appendix 2: The consumer perspective: Public attitudes to AI and financial guidance	42
	Appendix 3: Glossary of Terms	44

1 Executive Summary

1.1 Purpose of the research

This research explores the potential role of generative AI-enabled (Gen AI) self-serve tools within the regulated debt advice sector in the United Kingdom. It examines the opportunities, limitations, and risks of introducing these tools into a high-stakes, regulated environment where decisions can have material consequences for highly vulnerable people, including housing stability, legal enforcement, and long-term financial well-being.

The research was commissioned by the Money Advice Liaison Group (MALG), funded by the Money and Pensions Service (MaPS) Debt Advice Transformation Fund 2025/26 and delivered by Wyser as the research and delivery partner. Fieldwork was conducted between January and March 2026.

1.2 Research approach

The study adopted a mixed-methods approach comprising 28 semi-structured interviews with sector experts across two phases (an exploratory phase of 15 interviews followed by a validation phase of 13 interviews), a nationally representative public survey of 2,060 UK adults conducted through YouGov Plc and a sector-wide professional survey. Interview participants included frontline debt advisers, supervisors and quality assurance staff, service managers, compliance and policy specialists and digital and product leads across national providers, local agencies, commercial firms and specialist services.

Where figures are from YouGov Plc - Total sample size was 2060. Fieldwork was undertaken between 11th-12th March 2026. The survey was carried out online. The figures have been weighted and are representative of all UK adults (aged 18+).

1.3 Key Findings

The research identified a series of structural, professional and systemic features of regulated debt advice that are highly relevant to the design and deployment of Gen AI self-serve tools.

1 Debt advice is structurally complex and reliant on professional judgement

Regulated debt advice involves layered assessment, iterative engagement and relational skills developed over years of practice. It operates across intersecting areas including welfare benefits, housing, creditor enforcement, safeguarding and others. Advisers routinely reconstruct financial information, challenge unrealistic budget entries and navigate competing priority debts requiring careful judgement that cannot be reduced to simple decision logic.

Commented [TK1]: is this an exhaustive list?

Commented [MG2R1]: Updated to include 'and others'

2 Vulnerability is widespread and often disclosed only once trust is established

Professionals surveyed indicated that it takes years to become a fully proficient debt adviser, particularly in the soft skills associated with the role. Key skills, including extrapolating from incomplete client information, identifying hidden vulnerability and managing very nuanced and often

unique personal circumstances, develop through experience. This suggests that while there may be scope for Gen AI to assist with lower-complexity scenarios, this should not be at the risk to the development of advisers.

Clients often present in crisis

3

Clients frequently seek help only after enforcement action, including council tax enforcement, rent arrears proceedings or court correspondence. Advisers frequently act to support clients in diffusing crises, and participants questioned the ability of a Gen AI self-serve tool to do so. The risk of confidently presenting a wrong answer was also highlighted, with advisers noting that some, but not all, individuals in crisis may accept any option, even an inappropriate one, which has implications for how Gen AI tools frame their outputs for users.

Commented [GM3]: What if clients could easily access advice before reaching a crisis point? Is there a reluctance to ask for help until it is too late? Smaller problems that are not a crisis might be easier for an AI service to support

Debt, housing and benefits advice are closely interconnected

4

Debt rarely exists in isolation, and clients often need support in multiple areas, often from different advisers and different organisations. In this finding, the high drop-off rate from housing or benefit advice to debt advice was noted: when clients perceive their primary issue as resolved, they often disengage before debt is fully addressed. Self-serve tools that focus too narrowly on debt may fail to capture broader stabilisation needs.

Commented [MG4R3]: There is a reluctance and advisers saw an opportunity for a tool for earlier engagement - covered in chapter 4

Commented [GM5]: Perhaps - but having an app in your the pocket that supports/ guides/ rewards healthy and consistent habits could help reduce this "drop off"

The financial statement is the core of debt advice

5

Across all organisations, the financial statement was described as the foundation of debt advice. Constructing a sustainable budget involves challenging unrealistic client details, identifying hidden expenditures, interpreting fluctuating income, and accounting for future risks. Automated bank statement analysis tools, used by some organisations, have had mixed results and have not eliminated the need for skilled interpretation. Initial information gathering is possible via a Gen AI tool, but the strategic and tactical elements of budget construction that involve professional judgement may not currently be replicable by Gen AI.

Commented [MG6R5]: I've added 'too' in front of narrowly - but absolutely

Advice relationships are rarely one-off

6

Debt advice relationships often continue for months and in some cases years, involving follow-up, reassurance and re-evaluation as circumstances change. Advice is iterative, not transactional. In designing debt advice tools, organisations should not underestimate the amount of follow-up with creditors and other organisations that individual debt advisers undertake for their clients.

Commented [GM7]: ...therefore an app in the pocket that supports and rewards consistent healthy habits over a sustained period of time...

Client use of Gen AI for debt advice is lower than anticipated, but growing

7

Based on media reporting, particularly reports on harms caused by Gen AI, it was anticipated that client use of Gen AI would be common. However, across 28 interviews, advisers reported limited evidence of clients using Gen AI tools, though emerging cases were identified, including clients presenting with expectations of particular debt solutions based on Gen AI output, at least one Gen AI-driven bankruptcy decision, and Gen AI-prepared creditor information citing irrelevant legislation. The YouGov survey conducted for this research found that 37% of UK adults have used an AI chatbot for personal advice, though only 21% would be likely to use one for advice about financial issues.

8

The public does not yet trust Gen AI for advice on financial issues

72% of the public is not confident in Gen AI financial advice. 81% would prefer human advice. Only 2% would follow Gen AI advice exactly as given. The strongest trust-building factor was knowing that Gen AI advice is regulated by financial authorities (32%), followed by human expert review (27%).

9

Sector attitudes to Gen AI are mixed but not uniformly resistant

Larger organisations exhibit more urgency regarding Gen AI implementation; smaller organisations tend toward caution. Many contributors stated that Gen AI tools must be developed with meaningful sector involvement and that debt expertise must shape design. There was notable openness to Gen AI supporting advisers, reducing administrative burden and improving consistency.

10

Adviser-assist tools have the strongest support and most immediate potential

Across all three data sources, Gen AI tools that reduce administrative burden for advisers received the strongest endorsement. In the sector survey, 82% rated adviser-support tools as suitable now. Several organisations are already piloting transcription and case note tools with positive results. Credit report analysis, creditor communication and legislation decoding were also identified as high-value applications.

11

Self-serve tools require careful scoping based on organisational context

A notable finding of the study was that the clients served by different organisations varies widely, even within organisations that operate in different regions, the clients served also vary. Therefore, the suitability of client-facing Gen AI self-serve tools depends heavily on the nature of the client base served. There is no one-size-fits-all answer. Organisations serving individuals in acute crisis with low digital literacy require fundamentally different approaches from those with digitally confident populations. For some organisations, self-serve may never be appropriate for any meaningful proportion of their clients.

12

Accuracy, jurisdictional sensitivity and timeliness of information are critical

Debt advice information changes frequently through case law, legislation and regulatory interpretation. Scotland operates an entirely separate legal and debt solutions framework. Interviews reported teams and individuals spending hours to days per month staying current with debt advice practice. Any Gen AI tools will require well-understood knowledge sources, robust update mechanisms and clear jurisdictional routing.

1.4 Implications for the sector

The research identified a graduated landscape of possible applications, each with distinct risk profiles and attracting varying levels of sector appetite. Adviser assistance Gen AI tools received near-universal support, and several organisations are already piloting them with positive results. Client-facing Gen AI self-serve tools attracted cautious openness, with confidence declining as influence over client decisions increases. Agentic artificial intelligence (Agentic AI) applications that take autonomous actions are largely forward-looking, and the sector expressed significant caution.

The public does not yet trust Gen AI for financial guidance, and the technology is not yet accurate enough for unsupervised use in complex or high-stakes scenarios. However, client use of Gen AI for debt information is growing, and the sector cannot afford to ignore it. The emerging Debt Advice AI Roadmap provides a structured basis for proportionate decision-making; organisations can identify which applications are appropriate for their specific context, client base and governance capacity rather than treating AI adoption as a single binary choice.

The central principle is proportionality. As the influence of an AI tool over client outcomes increases, so must governance, human oversight and safeguards. For some organisations, adviser assistance may be the only appropriate application. For others, carefully scoped client-facing tools may genuinely extend reach and improve access. The evidence supports both conclusions equally.

Level	Application types
Level 1 Adviser Assistance	1.1 Administrative assistance 1.2 Knowledge and compliance support 1.3 Operational intelligence
Level 2 Client Self-Serve: Information and Assistance	2.1 General information 2.2 Contextual guidance 2.3 Structured triage 2.4 Pre-appointment data gathering 2.5 Communication support 2.6 Documentation support 2.7 Ongoing engagement
Level 3 Agentic AI: Taking Action and Sustained Relationships	3.1 Personalised advisory 3.2 Adviser-directed actions 3.3 Semi-autonomous actions 3.4 Autonomous process management

The Debt Advice AI Roadmap. See Section 6 for full details

1.5 The structural challenge

While there is an appetite within the sector for Gen AI tools and recognition that Gen AI use by clients and organisations is inevitable, the sector’s structural challenges cannot be solved by technology alone. Several structural issues were raised during many of the interviews for this research, including the prevalence of chronic deficit budgets for clients, decreases in funding for organisations, an increasing number of clients requiring support for increasingly complex issues,

and the attrition of experienced advisers, Gen AI can reduce administrative burden and may enable some interactions to be handled without direct adviser involvement. Still, it cannot replicate the professional judgement, emotional intelligence and sustained relationship that the evidence shows are essential for effective debt advice to the most vulnerable. Gen AI is a tool that should be adopted where it genuinely improves outcomes, not deployed at any cost simply because the technology exists.

Commented [TK8]: needs reworking into two sentences

Commented [MG9R8]: Updated

2 Introduction and Context

2.1 Background

2.1.1 The state of UK debt in 2025/26

Over the past decade, demand for debt advice across the United Kingdom has been shaped by a combination of economic pressures, structural financial insecurity and rising living costs.

Citizens Advice helped more than 400,000 people with debt problems during 2025, up nearly 45% since the cost-of-living crisis took hold in 2021 (*Citizens Advice issues stark warning as 2,670 people seek debt help every day in first week of the year, CitizensAdvice, January 15th 2026*).

Periods of inflation, rising housing costs, energy price volatility, and wage stagnation have exerted sustained pressure on household finances. Advice providers have reported growing demand for support from individuals facing complex financial circumstances, often linked to broader issues such as housing instability, benefit changes, illness or relationship breakdown. Step Change reported that almost 18,000 clients completed full debt advice in January 2026, 9% higher than in January 2025 and 56% higher than in December 2025. The number of clients citing cost-of-living increases as their main reason for debt rose to 18% in January 2026, while nearly one in ten named needing credit to cover living costs (*Monthly Client Data Report: January 2026, StepChange*).

At the same time, debt advice itself has become increasingly complex. Changes in legislation, developments in case law, and evolving debt solutions require ongoing interpretation and expertise. Debt cases rarely involve a single isolated issue. Advisers frequently work across intersecting areas, including welfare benefits, housing stability, creditor enforcement and safeguarding concerns. Resolving financial distress often means stabilising multiple aspects of a client's circumstances simultaneously.

Perhaps the most concerning indicator is the prevalence of deficit budgets among debt advice clients, a finding of both this report and other sector reports. In 2024, 50% of Citizens Advice debt clients and 43% of National Debtline clients had essential living costs that exceeded their income (*Money Advice Trust and Citizens Advice publish new Insolvency Report, as debt remedies need 'major rethink' | Money Advice Trust, Money Advice Trust, July 2025*).

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. In the Institute of Money Advisers (IMA) report in January 2024, seven out of ten advisers say their workload has increased over the last two years (to 2023) (*Casework requirements and workloads in the money advice sector, IMA, January 2024*), and based on the interviews carried out in this research, this workload increase continues.

2.1.2 The rise of Generative AI

Parallel to these pressures, the past two years have seen rapid advances in Gen AI. Large Language Models (LLMs) capable of producing conversational responses, summarising information and answering complex queries have rapidly become accessible to the public. These tools are increasingly used by individuals seeking guidance on legal, financial and administrative questions, often before engaging with professional services.

The speed of adoption has led many organisations to explore potential Gen AI applications, often driven by both perceived opportunity and competitive pressure. However, the deployment of Gen AI in regulated environments raises important questions about reliability, governance and consumer protection that require careful consideration.

Gen AI is powered by LLMs trained on vast amounts of text data to learn statistical patterns and probabilistic relationships between tokens, words and concepts. When a user asks a question, the model generates a response by predicting the most likely words, producing text that reads as though a knowledgeable person wrote it. This is powerful but carries important limitations: Currently, LLMs do not verify facts, do not access live databases unless specifically connected to them, do not understand the legal or regulatory implications of what they produce and cannot detect vulnerability, distress or the unspoken context behind a question. They can produce information that sounds authoritative even when factually incorrect, a phenomenon known as hallucination.

More recently, the technology has advanced with Agentic AI, systems that can take autonomous actions, break complex tasks into steps, use external tools, 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 correspondence and schedule follow-up actions without a human 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.

The relevance to debt advice is immediate. People are already using publicly available Gen AI tools to ask questions about debt. These tools can provide quick, confident answers, but they are not regulated, often inaccurate and rarely recognise vulnerability. The question is not whether people will use Gen AI for debt information; they already are, but how the sector can harness the technology responsibly.

2.1.3 The regulatory landscape

The FCA has confirmed that it will not introduce AI-specific regulations, instead relying on a technology-neutral, principles-based, outcomes-focused approach that uses existing frameworks, including Consumer Duty, the Senior Managers and Certification Regime (SM&CR), and operational resilience rules. FCA Chief Executive Nikhil Rathi reaffirmed this position in December 2025, citing the rapid evolution of AI technology (["AI Regulation in Financial Services: Turning Principles into Practice"](#), BCLP (2026)).

Commented [TK10]: does this need some kind of reference that this is the current situation? In future, AI will be able to reason which will change the accuracy of what they produce

Commented [MG11R10]: Included the word 'currently'

The Bank of England and FCA's third survey of AI and machine learning in UK financial services, published in November 2024, found that 75% of firms were already using AI, with a further 10% planning to adopt within three years (*"Artificial Intelligence in UK Financial Services — 2024"*, 21 November 2024, Bank of England and Financial Conduct Authority (2024)).

Replying to the [Guardian in relation to AI chatbots providing inaccurate financial advice](#), the FCA said, "Unlike regulated advice provided by authorised firms, any advice provided by these general-purpose AI tools are not covered by the Financial Ombudsman Service and the Financial Services Compensation Scheme." Despite the growing deployment of general-purpose Gen AI chatbots by non-regulated organisations, which provide information which is regulated when provided by a regulated organisation, there does not appear to be a UK regulatory requirement for chatbot accuracy certification or hallucination testing. This matters because the public wants to see regulation in this space: the YouGov survey conducted for this research found that knowing AI is regulated by financial authorities was the single strongest factor in building public trust (32%), ahead of human expert review (27%) and up-to-date regulations (21%).

The self-regulation of social media companies provides a cautionary example. More than a decade after the widespread adoption of social media, regulators across the UK and Europe are still catching up with the harms caused to young people, with over 2,000 lawsuits pending against major platforms in 2026 and UK regulators publicly stating that the status quo is not working. The debt advice sector has an opportunity to engage with AI governance for regulated advice proactively, not reactively, ensuring frameworks are in place before harm occurs rather than after. The Safe-Use Framework accompanying this report is intended as a contribution to that effort.

2.2 Research objectives

This research initially set out to examine how Gen AI self-serve tools might interact with the practice of regulated debt advice. Rather than evaluating specific tools, it focused on understanding the real-world environment in which such tools would operate and the considerations for different uses of these tools.

Six core questions shaped the research:

1. How do individuals typically seek help and enter advice services?

Understanding pathways into advice is essential to assessing where self-serve tools might be encountered and what role they could play.

2. Where do misunderstandings, confusion, or risk most commonly arise?

Identifying where misunderstandings occur and clarifying where Gen AI tools might provide informational support and where they might reinforce incorrect assumptions.

3. How do advisers assess complexity and vulnerability in practice?

Professional debt advice relies heavily on judgement developed through experience. Understanding this is critical when evaluating what can or cannot be automated.

4. Where could Gen AI tools potentially add value?

Commented [GM12]: Is the risk/ opportunity wider than just debt? Debt is rarely the only issue - see earlier - the debt sector can lead the others to ensure there is alignment/ control/ regulation/ principles...

Commented [MG13R12]: Making this 'advice' rather than 'debt advice' - it was originally AI governance generally but Mark asked for it to be tighter so the report doesn't over reach

Gen AI tools may support information provision, triage, data gathering and adviser assistance. Identifying realistic areas of benefit enables proportionate consideration of use cases.

5. Where could Gen AI tools introduce harm or unintended consequences?

Automation can create risks where systems oversimplify, fail to recognise vulnerability or influence decision-making without sufficient context.

6. What governance and safeguards would be necessary for safe deployment?

Any Gen AI tools must operate within frameworks that protect consumers and ensure accountability.

2.3 Scope and definitions

Based on the progression of the research, particularly discussions on adviser use and concerns about serious harms from automatic decision-making, Gen AI tools for adviser assistance and some Agentic AI for self-serve are also discussed in the report and have been included in the Safe-Use Framework. There are some more general references to artificial intelligence (AI), intended to encompass both Gen AI and Agentic AI rather than the whole field of AI.

For this report, regulated debt advice refers to professional guidance provided to individuals on managing personal debt, including assessing financial circumstances and identifying appropriate solutions. In the UK, this activity operates within regulatory frameworks overseen by the Financial Conduct Authority (FCA), alongside the provision of funded advice and delivery in the charitable sector.

Gen AI refers to generative AI-enabled tools that allow individuals to obtain guidance, information or structured assistance without direct interaction with a human adviser at the point of use. Importantly, the term 'self-serve' does not imply the complete absence of human oversight. In many potential use cases, Gen AI tools may gather information or provide preliminary support before a professional reviews the case.

This report does not evaluate or endorse any specific Gen AI product or platform. The analysis focuses on general patterns of potential use and risk rather than the performance of particular technologies.

For a full set of terms, see Appendix 3: Glossary of Terms.

3 Methodology

3.1 Proposed research design

A mixed-methods approach was designed to explore the potential role of Gen AI self-serve tools within regulated debt advice. The original methodology comprised three components:

- Sector expert interviews: Semi-structured interviews with approximately 20 sector experts from frontline debt advice services, intended to capture professional perspectives on practice, risks and how AI might interact with existing processes
- Observation of Gen AI tool use: Observation of around 40 participants interacting with a Gen AI debt advice tool, to understand how individuals might use automated guidance and where misunderstanding or misuse might occur
- Sector survey: a survey targeting approximately 100 debt advice sector professionals to gather broader perspectives and validate emerging findings

3.2 Adaptation of the research approach

As the research progressed, the methodology was refined in response to emerging insights from initial interviews. Early discussions with sector experts revealed a significant gap between the anticipated user group for Gen AI self-serve debt advice tools and the individuals most commonly accessing in-person services. Advisers consistently reported that many of their clients present in crisis, have low digital literacy, rely heavily on in-person or telephone support and are referred through housing services, benefits advice or community organisations.

This changed the question from 'how are people using Gen AI self-serve tools?' to 'to what extent are people actually using Gen AI for advice on financial issues?' Based on this emerging question, the research team determined that observing a small group interacting with a Gen AI tool would provide limited insight; therefore, funding for observation was redirected toward a nationally representative YouGov survey of 2,060 UK adults, providing population-level data to complement the practitioner perspective.

3.3 Sector expert interviews

28 interviews were conducted across two phases. The exploratory phase comprised 15 semi-structured interviews to understand how regulated debt advice operates in practice and identify areas of complexity, ambiguity and risk relevant to Gen AI deployment. The validation phase comprised 13 further interviews to test emerging themes, sense-check possible Gen AI applications within the sector, and gather additional perspectives. Participants included frontline advisers, supervisors, QA staff, service managers, compliance specialists, digital leads, and creditor representatives from national providers, local offices, commercial firms, specialist charities, local authorities, and organisations across England, Wales, and Scotland. Collective experience among participants exceeded 250 years in the debt advice sector.

Commented [TK14]: such as what?

Commented [MG15R14]: Clarified

3.4 Sector survey

A sector survey was conducted targeting professionals in debt advice and related roles. 17 responses [were] received. While not statistically significant, the responses provide useful directional data when triangulated against interview evidence and the public survey. Respondents included debt advisers, QA professionals, managers and a creditor representative from national charities, local charities, local authorities and a housing association.

Commented [TK16]: needs updating

3.5 Public survey on Gen AI use

To address the identified evidence gap regarding public use of Gen AI, a nationally representative survey was conducted through YouGov, collecting responses from approximately 2,000 UK adults. The survey explored awareness and use of Gen AI tools for personal rather than professional purposes, willingness to use AI for advice about financial issues, trust levels and perceived advantages and risks.

3.6 Analytical approach

Interview outputs were analysed using affinity mapping and thematic analysis across the full corpus of 28 transcripts, identifying recurring patterns related to pathways into debt advice, sources of misunderstanding, vulnerability indicators, professional judgement in casework, and attitudes toward AI. Insights from the sector and public surveys were used to triangulate and contextualise these findings, allowing comparison between frontline advice practice and emerging patterns of AI engagement among the wider public.

Quotes have been taken from the original transcripts and attributed by participant code to maintain anonymity while enabling cross-referencing.

3.7 Limitations

Several limitations are acknowledged. First, sector expert interviews represent professionals' perspectives; the research does not include direct interviews with clients receiving debt advice. Second, the public survey provides insight into general attitudes toward Gen AI for advice on financial issues, but does not specifically target individuals experiencing debt problems. Finally, the research reflects the specific regulatory, legal and institutional frameworks of the [UK] debt advice system.

Commented [TK17]: does this include Scotland or just E&W?

Commented [MG18R17]: Both

Despite these limitations, the combination of practitioner insight, sector perspectives and population-level data provides a strong foundation for understanding how Gen AI self-serve tools may intersect with regulated debt advice practice.

4 Debt Advice Sector characteristics and the implications for Gen AI self-serve tools

Before considering Gen AI tools, it is essential to understand key elements of how debt advice operates in practice so that the implications for the implementation of Gen AI are clear. Debt advice is a highly specialised and experience-driven field. Advisers develop their expertise over many years, and much of that knowledge is tacit, developed through repeated exposure to complex cases and evolving regulatory environments. Several themes emerged consistently across organisations, suggesting structural characteristics of regulated debt advice. This report does not cover all aspects of regulated debt advice, but it highlights the key characteristics that have implications for the use of Gen AI in the sector.

4.1 Professional judgement and tacit knowledge

Professional judgement plays a central role throughout the advice process. Participants emphasised that becoming an effective debt adviser takes years of experience. Key skills include reading between the lines in client conversations, identifying subtle indicators of vulnerability, gathering accurate information from incomplete disclosures, challenging client assumptions about debts and priorities, and navigating creditor processes.

Advisers also assess whether potential solutions are sustainable over time, considering factors such as income stability, the client's motivation, and the likelihood of recurring financial difficulties. Even where individuals appear to have a workable financial plan, those operating with very narrow margins remain vulnerable to future financial shocks.

The strategic use of Breathing Space is a case in point, with some advisers holding it in reserve while using informal holds to ensure the best outcome for individuals. This was cited across at least four organisations as an example of tactical expertise that automated systems would struggle to replicate (AP01, AP20, AP22, AP24). One adviser noted a specific concern that existing automated systems may be recommending Breathing Space routinely, without the strategic judgement about whether to reserve it for a potential future emergency (AP20).

Experienced advisers described picking up clues from letterheads, reference numbers, letter formatting, tone of voice and body language that informed their understanding of a case in ways that would be invisible to any AI system (AP34). The knowledge required spans not only debt law and regulation but also welfare benefits, housing law, creditor practices, local authority procedures and the specific characteristics of hundreds of individual creditors.

4.1.1 The knowledge maintenance challenge

In addition to exercising professional judgement, remaining current with the debt advice knowledge base was identified as a significant and ongoing challenge. Larger organisations have teams dedicated to knowledge management and dissemination. One participant estimated that members of their knowledge team spent eight hours per person per week simply staying current (AP03). For cross-jurisdictional advisers, this effort is magnified as Scotland operates an entirely

separate legal system with distinct statutory solutions, different enforcement timescales and separate regulatory frameworks (AP09, AP12). For any Gen AI tool operating across the United Kingdom, jurisdictional accuracy is not an optional enhancement but a fundamental safety requirement.

The potential for AI tools to support knowledge maintenance is discussed in Section 6.

4.2 Crisis entry and late presentation

One of the most consistent findings across all 28 interviews was that clients overwhelmingly present to debt advice services at a crisis point rather than at an early stage when lighter-touch intervention might suffice. The most commonly cited triggers were council tax enforcement, bailiff action, possession proceedings and benefit disruptions. Advisers described clients arriving with months or years of unopened correspondence, having delayed engagement due to a combination of shame, denial, and the hope that the situation would resolve itself.

“Quite a few things need to happen before they actually seek advice. For them to actually come over the threshold and meet someone face to face is a big step. There’s a lot of embarrassment and guilt, and there’s tears” AP22

One participant from the commercial sector estimated a nine-month average delay between an individual recognising they had a financial problem and taking any action to address it (AP04). Another cited research that suggests the delay could extend to 18 months after reaching a crisis point (AP18). The pattern was universal across organisation types, though the severity at presentation varied. National telephone services tended to see clients with single issues and higher capability, while local face-to-face services retained only the most complex and vulnerable cases, having already signposted more capable individuals to providers with digital journeys (AP29).

4.2.1 No-show rates and disengagement

Several participants reported substantial dropout between referral and engagement. One organisation tracked a rate in which 50% of referrals never accepted a further appointment, with a subsequent 30% no-show rate among those who did book in (AP20). Another reported 50% no-shows in a single month (AP25). Seasonal patterns were noted, with engagement peaking after the new year and declining through summer months (AP25). A recurring cycle was described in which clients engage briefly during an acute crisis, receive help with the triggering issue, but disengage before addressing their broader debt picture, only to return six to twelve months later when a new crisis emerges (AP22).

This pattern has profound implications for Gen AI self-serve tools, which are typically designed for single-session, information-based interactions. If the majority of people presenting to debt services are already in crisis, tools designed for calm, methodical self-assessment may be poorly suited to their state of mind. Tools need to be designed for people who are distressed, overwhelmed and potentially in acute crisis rather than for people calmly researching their options. The research consistently found that crisis-driven presentation is the norm rather than the

Commented [GM19]: This is the opportunity for an app - help whilst the problem is small... no stigma/ no shame/ single issue...

Commented [MG20R19]: There is a paragraph at the end of the section referencing the opportunity

Commented [TK21]: what signifies a lighter touch provider from any other provider?

Commented [MG22R21]: Clarified to providers with digital journeys

exception. However, multiple participants observed a potential opportunity: a well-designed information tool might serve as a bridge to the earlier stage, when people are beginning to worry but are not yet ready to engage with a formal service (AP22, AP29, AP24).

4.3 Referral routes

People arrive at debt advice services through a range of referral routes. For many community-based advice organisations, referrals from partner services represent a major pathway. These often come from housing advice services, welfare benefits advisers, local authority support hubs, community support organisations and food banks.

Word-of-mouth referrals were widely cited, particularly among organisations embedded in local communities. Some contributors reported referrals from GPs, social prescribers or community networks, though these were less common. Larger national or commercial organisations reported greater reliance on online search and digital discovery. Self-referral also occurs, though it often happens only when individuals have already reached a crisis point.

The diversity of referral routes has significant implications for where Gen AI self-serve tools might be most effective. If the majority of clients reach debt advice through partner organisations, community hubs, food banks and word-of-mouth rather than through direct online search, then deploying self-serve tools on a debt advice organisation's own website may reach only a fraction of the people who need help. The tools may be sitting in the wrong place.

This raises a strategic question: which organisations are best placed to deploy client-facing Gen AI tools? A housing association, a food bank, a GP surgery or a local authority support hub may be the point at which someone first discloses financial difficulty. A well-designed general information or triage tool embedded at the point of first contact could identify debt issues earlier, provide immediate, basic guidance, and route individuals to appropriate advice services before their situation deteriorates further.

Additionally, the prevalence of word-of-mouth and community-based referrals suggests that trust and human connection play a significant role in whether someone seeks help at all. Gen AI tools that replace or bypass that human first contact risk losing the very mechanism that brings people to the service. This reinforces the case for Gen AI tools that support and enhance existing referral pathways rather than attempting to create an entirely new digital front door that many of the most vulnerable clients will never walk through.

4.4 Vulnerability and hidden circumstances

The research revealed that vulnerability in the debt advice context is rarely a static, declarable condition. It is dynamic, layered and frequently hidden. Across the interviews, advisers consistently described a pattern in which the most significant aspects of a client's situation only emerge after trust has been established through sustained human interaction.

“People do curate the amount of information they provide, and you have to draw it out of them as you go along”, AP13

Domestic abuse and financial coercion emerged as significant hidden factors across multiple interviews. One adviser noted disclosures of domestic abuse emerging only after an hour-long conversation when sufficient trust had been established (AP01). Silence in response to questions about household finances was identified as a key vulnerability indicator, invisible to any text-based Gen AI system (AP19). Participants described cases requiring careful professional judgement to distinguish between cultural norms around household financial management and patterns of coercive control (AP19).

Hidden spending patterns with material consequences for financial assessments were routinely described. These included illegal drug expenditure concealed through cash withdrawals, undisclosed gambling and spending on credit cards not visible on the primary bank account (specific interviews not cited for privacy). Mental health conditions were cited by virtually every participant as the most prevalent vulnerability, with impacts ranging from an inability to open letters or make telephone calls through to complete incapacity to engage with any aspect of financial management. Even clients who appeared capable could be so immobilised by their situation that the slightest additional burden would cause them to disengage entirely (AP13).

In addition to hidden vulnerabilities, clients often present with a specific issue that initially appears to be the primary problem, but which later proves to be only part of a broader set of financial difficulties. Individuals may initially focus on a credit card balance or parking fine, while deeper issues, such as rent arrears and council tax liabilities, emerge during the advice process.

Advisers frequently noted that clients misunderstand the relative seriousness of different debts, sometimes prioritising non-priority debts while underestimating the consequences of priority ones. Addressing the presenting issue often becomes the first step in building trust before advisers can explore the wider financial context.

The implications of hidden vulnerability and presenting-issue bias for AI tools are discussed in Section 5 (risk factors) and Section 6 (safeguards at each Roadmap level).

4.5 The financial statement: cornerstone and challenge

Every participant identified the income and expenditure assessment as the foundational element of regulated debt advice. Its accuracy is not merely desirable but determinative. *The arithmetic is comparatively simple; the psychology is comparatively complicated* (AP22). Getting the financial statement wrong means everything that follows is wrong.

Participants were unanimous in the view that clients cannot reliably self-report their financial circumstances without professional assistance. Clients routinely underestimate food expenditure, forget irregular costs such as school uniforms and vehicle maintenance, confuse weekly and monthly figures, omit expenditure they consider embarrassing and fail to account for spending made through cash or on credit cards that do not appear on their primary bank statement.

“The budgets that you get back are absolutely awful. In 9 out of 10 cases, they haven’t got a clue what they’re spending their money on, which is why they’re in debt” AP22

One participant noted clients claiming £50 per week on food for a family of four, a figure that could not possibly be accurate (AP12).

Advisers work to identify suppressed needs, ensuring that budgets reflect a sustainable level of expenditure rather than unsustainable sacrifice. This involves distinguishing between what a client believes they can currently afford and what should reasonably be included in a healthy budget, a skill that requires experience and professional judgement.

4.5.1 The limits of open banking

Several participants had experience with open banking tools and, while acknowledging their utility, identified significant limitations. Open banking captures what has been spent, but not what should have been spent. A client who has cancelled building insurance because they cannot afford it shows no expenditure in that category. Still, an experienced adviser would immediately identify the gap and insist on including it as a necessary cost (AP26). Other limitations included the inability to categorise family lending and informal financial transfers, the misclassification of internal transfers between accounts as income or expenditure and the absence of cash transactions (AP02, AP12). One adviser described the difference between the budget a client produces using an online calculator and the budget produced from bank statement analysis as being worlds apart (AP29).

4.5.2 Strategic budget construction

Perhaps the most striking finding was the degree to which budget construction involves professional judgement and strategic thinking rather than simple arithmetic. Advisers described actively managing budgets to achieve optimal outcomes for clients, knowing which figures creditors would challenge, understanding how different creditors and the court treated the same categories of income, and strategically structuring or presenting financial information to maximise the client's options (AP01, AP13). Disability benefits such as Personal Independence Payment could be treated differently by different creditors: some would disregard it as income, allowing it to be offset against care costs; others would count it in full. The court, the council, the insolvency service and the DRO intermediary all used different financial assessment methods (AP13).

4.5.3 The prevalence of deficit budgets

As noted in Section 2, the prevalence of deficit budgets has increased significantly, with approximately half of clients at some organisations now presenting with essential costs exceeding income. For these clients, the role of the debt adviser extends beyond options assessment to include income maximisation, benefit checks, charitable grant applications, and referrals to local support services.

Overall, the financial statement evidence presents perhaps the most fundamental challenge to Gen AI self-serve in debt advice. If self-reported financial information is unreliable in 9 out of 10 cases, any Gen AI tool that accepts client-provided income and expenditure figures at face value will build its analysis on a flawed foundation. Open banking can supplement self-reporting, but

captures spending history rather than actual need, misses cash transactions and cannot identify the insurance policy that was cancelled or the essential cost that has been suppressed.

Overall, the financial statement evidence presents perhaps the most fundamental challenge to Gen AI self-serve in debt advice

4.6 Complexity beyond debt

Debt problems rarely exist in isolation. Benefits issues, including delays, sanctions or incomplete claims, frequently exacerbate financial hardship. Housing instability can require parallel intervention. Participants also reported a noticeable drop-off between housing or benefits advice and debt advice referrals. When an immediate crisis is resolved, some clients disengage before addressing their underlying debts, often returning later when they feel more ready.

One adviser described navigating a landscape of hyper-local charitable trusts, each with different eligibility criteria and geographical boundaries, to secure practical help, including energy grants, food bank vouchers and household items for families fleeing domestic abuse (AP24).

The interconnections among debt, benefits, housing, and broader support create both a challenge and an opportunity for Gen AI tools. The challenge is that a tool designed to address debt in isolation will miss the broader context that determines whether a client's situation can actually be resolved. A client whose debt stems from a benefit sanction needs the sanction addressed before any debt solution becomes viable. A client in rent arrears may need housing intervention before debt advice can begin in earnest. Gen AI tools that focus narrowly on debt risk, providing technically correct but impractical guidance because they cannot see the whole picture.

The drop-off between crisis resolution and debt advice follow-up presents a further risk. If a Gen AI tool successfully helps a client navigate an immediate housing issue, the client may disengage before addressing the debts that will bring them back into crisis six months later. Without the human adviser's ability to maintain a relationship and gently re-engage, the cycle continues.

In many organisations, more experienced debt advisers often act as consultants to other debt advisers on very complex cases, again illustrating how, even with existing knowledge bases and guides, cases with very specific circumstances often require in-depth analysis and judgement.

The potential for AI tools to support advisers navigating multi-issue complexity is discussed in Section 6.

4.7 Iterative casework and unseen effort

Debt advice rarely (if ever) takes place in a single interaction. Cases involve multiple appointments, follow-up calls, reassurance and prompting for documentation, liaison with creditors and, in some cases, engagement with third parties such as family members or support workers. Case duration varies significantly, with some cases extending over many months or years. One adviser described maintaining hundreds of case notes for a single client, reflecting the evolving nature of complex long-term relationships.

The research also revealed that advice appointments represent only a fraction of the total work involved in managing a debt case. Behind each case lies a substantial volume of creditor correspondence, benefit applications, charitable grant submissions, housing referrals, safeguarding actions and administrative tasks. One adviser described spending an afternoon reviewing a single client's bank statements (AP22). Another recounted a case with 218 case notes requiring a full independent file review (AP29). A third described routinely working beyond contracted hours because the volume of follow-up could not be contained within the working day (AP24).

"If we don't do it, there's nobody else that's going to help that client do it" AP24

Several participants described how their role had evolved beyond debt advice into what amounted to social work: arranging food bank referrals, applying for charitable grants for essential household items, liaising with adult social care and supporting clients through benefit claims that other services should have handled. This scope of activity is invisible to any system that measures productivity by the number of advice sessions delivered or debt options recommended, and it represents the most significant gap between what Gen AI tools can currently do and what the job actually requires.

The volume of unseen effort that accompanies each advice case represents both the strongest argument for Gen AI assistance and the area requiring the greatest caution. Adviser assistance tools offer clear and immediate value here: automating correspondence, summarising lengthy case files, drafting standard creditor letters and reducing the administrative burden that drives advisers to work beyond their contracted hours. If Gen AI can give an adviser back even a fraction of the afternoon spent reviewing a single client's bank statements, that time goes directly to the clients waiting for help.

4.8 Clients cannot always articulate their situation accurately

In self-serve situations, without adviser support, clients must be able to articulate their situation accurately to receive suitable advice. However, participants report that clients do not always understand their circumstances, may not describe them accurately or in depth and may use incorrect terminology. The misunderstandings described by participants are not simply individual knowledge gaps. They arise partly from the way the debt advice ecosystem operates and the circumstances in which individuals seek help. As described earlier, many clients arrive in crisis with incomplete information and a limited understanding of how the debt system operates. Advisers therefore spend a significant portion of their work correcting misconceptions and reframing clients' understanding of their situation. When this client's misunderstanding of their own circumstances and the landscape meets Gen AI tools, it can lead to a greater risk of a detrimental outcome for clients.

Participants also described widespread misunderstanding about the nature and implications of different debt solutions. Clients frequently arrive with preconceived ideas about particular solutions, often shaped by advertising, online content or recommendations from friends and

family. Advisers reported hearing statements such as 'My friend has a Debt Relief Order, I want one too' or 'I saw an advertisement for an Individual Voluntary Arrangement (IVA). Isn't that the best option?'

These requests reflect a partial understanding of the solutions involved. Eligibility criteria, long-term consequences and suitability depend heavily on an individual's full financial circumstances. Advisers therefore begin by reconstructing the client's financial circumstances and explaining the range of available options before assessing which solution, if any, is appropriate.

A related dynamic is advice shopping, where individuals consult multiple sources in search of confirmation of a preferred course of action. It was identified as an existing pattern, with clients trying multiple organisations in search of the answer they preferred (AP17, AP21). Gen AI tools could amplify this: a confident chatbot response that aligns with a client's preferences may increase misplaced confidence, making the adviser's role in reframing harder. As one participant noted, the concern is not that AI adds a new problem but that it turbocharges an existing one.

5 Additional risk factors

In addition to the characteristics of the debt advice sector having specific implications for the implementation of Gen AI tools in general, several other factors were identified as problematic during the interview. These arise from the way chatbots 'behave' and from commercial behaviour that the use of Gen AI could amplify.

5.1 Chatbot 'behaviour'

During interviews, several participants reported experimenting with publicly available Gen AI tools to test how they responded to debt-related questions. While contributors acknowledged the speed and fluency with which these tools generate responses, they identified several risks closely connected to the structural characteristics of debt advice, particularly late presentation, incomplete information and hidden vulnerability.

Current Gen AI systems have no capacity to detect silence, read body language, sense evasion, or create the conditions of trust that enable disclosure. Any framework for safe Gen AI self-serve must include robust mechanisms to identify when circumstances are more complex than the inputs suggest and to ensure rapid escalation to a qualified human adviser. A self-serve system responding to a client's presenting question may probe beneath the surface. Still, without knowing all the nuances, there is no way to guarantee that sufficient information has been gathered to identify the main issue and recommend a specific solution correctly.

Several participants also described clients or their family members using ChatGPT to draft letters to creditors. One case involved a client's sister generating a letter requesting a debt write-off that cited FCA regulations inapplicable to the specific debt type (AP15). Another described an AI-rewritten Universal Credit write-off letter that was inappropriately aggressive and Americanised in tone when the situation called for a conciliatory approach (AP35). The letters were fluent and convincing, but made arguments that were legally irrelevant or factually incorrect for the specific circumstances.

5.1.1 Plausible but incorrect solutions

Gen AI tools can produce explanations that appear convincing and authoritative, even when the underlying reasoning is incomplete or incorrect. In regulated debt advice, relatively small details can significantly affect the suitability of particular solutions. Eligibility may depend on precise income and expenditure levels, the types of debts involved, asset ownership, recent financial behaviour and jurisdictional rules.

Gen AI systems that generate responses based solely on user-provided inputs may produce outputs that appear reasonable but are based on partial or inaccurate information. In situations where individuals are already uncertain about their financial circumstances, confidently presented but incorrect guidance may reinforce misunderstanding rather than correct it.

One participant working in quality and compliance described clients using ChatGPT to generate complaints citing the Human Rights Act, GDPR and the Equality Act, none of which were relevant

to their situation. In one case, a client lost their property because they followed ChatGPT's guidance rather than the professional advice they had been given (AP21).

5.1.2 Overconfident framing

Gen AI responses typically adopt a confident and authoritative tone. While helpful in some contexts, this creates potential risks in debt advice scenarios where users may interpret general information as personalised advice. This risk is amplified by the circumstances in which individuals often seek guidance. As described earlier, many clients seek help during moments of financial crisis, when stress and cognitive overload may limit their ability to critically evaluate information.

5.1.3 The reasoning deficit

"AI is really poor at putting down any kind of reasoning. Syllogistic reasoning of any sort, which is crucial to explaining your decisions as a debt adviser" AP

This observation was supported across multiple interviews. Gen AI tools were competent at recording factual information and listing tasks but could not explain why a particular option was chosen, how competing factors had been weighed, or what the rationale was for departing from the obvious solution. In regulated debt advice, where advisers must demonstrate that recommendations are suitable for the specific client, this reasoning deficit is not merely inconvenient; it represents a fundamental gap in required capability.

5.2 Commercial motivations and organisational bias

A recurring concern was the influence of commercial motivations on the quality and impartiality of debt advice and the potential for Gen AI self-serve tools to amplify existing problems.

"All roads lead to Rome. It looks like it's kind of an impartial process, but in fact it's not" AP

Individual voluntary arrangements were cited across virtually every organisation type as the debt solution most frequently mis-sold. Participants described a well-established pipeline: individuals searching online encounter commercially funded advertisements leading to IVA providers who conduct superficial financial assessments and recommend IVAs regardless of suitability (AP13, AP17, AP22, AP26, AP34). One participant identified specific lead-generation operations disguised as support communities on social media (AP26). IVAs remain the largest form of personal insolvency in the UK despite being unsuitable for many of those directed toward them (AP32). Several participants warned that if Gen AI tools draw from the broader internet, they will inevitably incorporate this commercial bias, presenting it with the additional veneer of apparent impartiality that Gen AI responses carry (AP17, AP22).

6 Opportunities for Gen AI in regulated debt advice

While preceding chapters have highlighted the complexity and risk inherent in Gen AI self-serve tools within the regulated debt advice sector, the research also identified substantial areas where Gen AI tools could provide genuine value. One of the clearest findings was that opportunities and risks are not separate conversations. They are two sides of the same coin: every application that offers a benefit also introduces considerations that must be understood and managed.

This chapter presents the opportunities identified through the research, organised around three levels of application that form the Debt Advice AI Roadmap. For each level, the chapter sets out what the research found in terms of both potential value and the considerations that accompany it.

Level	Application types	Risk and sector appetite
Level 1 Adviser Assistance	<ul style="list-style-type: none"> 1.1 Administrative assistance 1.2 Knowledge and compliance support 1.3 Operational intelligence 	<p>Risk: Low to moderate</p> <p>Appetite: Strong. 82% of sector survey respondents rated adviser-support tools as suitable now—strongest endorsement of any application. Several organisations are already piloting.</p> <p>Technology: Generative AI. The adviser reviews all output before use.</p>
Level 2 Client Self-Serve: Information and Assistance	<ul style="list-style-type: none"> 2.1 General information 2.2 Contextual guidance 2.3 Structured triage 2.4 Pre-appointment data gathering 2.5 Communication support 2.6 Documentation support 2.7 Ongoing engagement 	<p>Risk: Low to high, increasing with influence over client decisions</p> <p>Appetite: Graduated. 82% support general information tools; drops to 53% for triage and data gathering. 71% say AI should never make solution recommendations. 71% say AI should never assess vulnerability</p> <p>Technology: Generative AI interacting directly with clients. Current accuracy (~75%) falls short of the 85–90% threshold identified by the sector.</p>
Level 3 Agentic AI: Taking Action and Sustained Relationships	<ul style="list-style-type: none"> 3.1 Personalised advisory 3.2 Adviser-directed actions 3.3 Semi-autonomous actions 3.4 Autonomous process management 	<p>Risk: High to critical</p> <p>Appetite: Very cautious. Personalised advisory rated suitable by only 18% (non-priority) and 12% (all debts). No organisations are currently deploying agentic AI in debt advice.</p> <p>Technology: Agentic AI that takes actions, not just generates content. Largely forward-looking. 3.4 Autonomous process management may not be currently appropriate for regulated debt advice</p>

The Debt Advice AI Roadmap

The accompanying Safe-Use Framework provides detailed organisational readiness assessments, minimum standards and safeguards for each application type. The detail is included

here so that this report can stand alone, but organisations seeking to implement should refer to the framework for the full operational guidance.

The three levels reflect a fundamental distinction in the type of AI involved and the degree of influence over client outcomes. Level 1 covers adviser assistance using Gen AI. Level 2 covers client self-serve and support tools using Gen AI. Level 3 covers Agentic AI applications that take actions and maintain sustained relationships. Each level carries a different risk profile, sits at a different point of technological readiness and attracts a different level of appetite from the sector.

6.1 Level 1: Adviser Assistance

The most consistently cited opportunity across all data sources was Gen AI tools that reduce administrative burden for advisers. In the sector survey, 82% of respondents rated adviser-support tools as suitable for use now, the strongest endorsement of any Gen AI application. Interview participants from every organisation type identified administrative support as the most immediately valuable area. Several organisations are already piloting these tools with positive results.

Level 1 tools use Gen AI to produce content that the adviser reviews before use. The Gen AI does not interact with clients. If it produces an inaccurate output, the adviser catches it before it goes anywhere. This creates the lowest-risk environment for Gen AI deployment and a safe space for organisations to build experience, confidence and governance capability.

6.1.1 Administrative assistance (1.1)

The administrative burden of case documentation was described by virtually every participant as the single most time-consuming aspect of their work after direct client interaction. Participants reported spending hours writing up case notes, with one citing a requirement to produce approximately 50 pages of documentation per client (AP20). Organisations already using Gen AI transcription tools reported positive experiences, with one describing staff as loving the tool (AP05).

Credit report analysis was identified as a further high-value application. One participant described spending hours reviewing a 110-page credit report in which half the entries were Klarna buy-now, pay-later transactions (AP22). The task of extracting relevant debts, identifying defaults and transferring information to a financial statement is repetitive, rule-based and time-consuming, making it well-suited to automation with human oversight.

Creditor communication was raised independently by multiple participants. Currently, advisers must identify individual creditor contact details, send separate correspondence to each creditor and follow up when responses are not received. One adviser described sending 10 individual letters offering £1 per month token payments, two or three of which were lost in transit (AP22).

6.1.2 Knowledge and compliance support (1.2)

There was strong interest in AI tools capable of decoding complex legislation into accessible language for advisers who are not legally trained (AP18). Consumer credit law, insolvency legislation and benefits regulations are dense and technical. Tools that summarise policy

guidance or highlight relevant rules were seen as particularly useful in environments where staff turnover is relatively high, advisers are early in their careers or training periods are short.

However, participants expressed caution about relying on AI to interpret regulatory frameworks without human oversight. Most viewed such tools as supplementary knowledge support rather than decision-making systems. The key consideration is accuracy: debt advice operates across multiple legal jurisdictions with frequently changing legislation. Any Gen AI tool operating in this space must draw from approved, curated sources rather than the open internet and must be updated rapidly when case law or legislation changes.

6.1.3 Operational intelligence (1.3)

Quality assurance at scale was identified as a specific opportunity. Current sampling approaches review only a small fraction of cases. One participant suggested AI could potentially assess 100% of cases, identifying trends, flagging potential errors and directing human reviewers to the cases most likely to require intervention (AP21). An AI-versus-adviser comparison model was also proposed as a testing methodology (AP18).

The consideration at this level is the use of data. Trend analysis and caseload management involve processing information about individual clients at an aggregate level. Appropriate legal bases for this processing must be established, data should be anonymised or aggregated where possible, and findings must be reviewed by humans before action. Pattern matching across cases should not lead to automated decisions about individual clients.

Sector appetite: Level 1 attracted the strongest and most consistent support of any area. For many organisations, particularly those working with highly vulnerable populations, adviser assistance tools were viewed as significantly more appropriate than client-facing self-serve. By supporting the work of advisers rather than replacing it, these tools allow organisations to extend the reach of existing expertise while maintaining professional oversight. For some organisations, Level 1 may represent the full and appropriate extent of AI adoption.

6.2 Level 2: Client Self-Serve: Information and Assistance

Level 2 covers Gen AI tools that interact directly with individuals seeking help with debt and tools that support clients alongside the advice journey. These tools use Gen AI to provide information, answer questions, assist with practical tasks, and, at the higher end, gather pre-appointment information for adviser review. Risk increases progressively as the tool moves from providing general information toward influencing individual financial decisions.

Participants were cautiously open to client-facing tools but felt deployment must be carefully scoped based on the organisation's client base. Sector survey confidence mapped closely to the risk gradient: 82% rated general information tools (2.1) as suitable now, but this dropped to 53% for structured triage (2.3) and pre-appointment data gathering (2.4). The sector was clear about boundaries: 71% said AI should never make final solution recommendations, 71% said it should never assess vulnerability, and 65% said it should never determine insolvency eligibility.

There is some natural overlap between Levels 2 and 3. Tools at the upper end of Level 2, particularly those maintaining ongoing client engagement, share characteristics with agentic applications. The distinction is that Level 2 tools generate content and information; they do not take actions or make decisions autonomously.

6.2.1 General information (2.1)

Many people approaching debt advice services need basic information before they are ready for a full advice session. What is a DRO? How does an IVA work? What happens if I do not pay a council tax bill? These are factual questions that can be answered from curated, approved sources without personal data or individual assessment.

The consideration is accuracy and currency. Debt advice information changes frequently through case law, legislation and regulatory interpretation. Scotland operates an entirely separate legal and debt solutions framework. Any general information tool must have jurisdictional awareness, must draw from restricted knowledge sources and must be updated rapidly when changes occur. Content must be written at a suitable reading age to serve different clients and, if necessary, interactions must include clear signposting to professional advice.

6.2.2 Contextual guidance (2.2)

Some queries require the tool to use limited contextual information provided by the user. A person might describe a specific debt type and ask how to prioritise it, or request help understanding a creditor letter. This is more than general information because the response is shaped by the individual's circumstances, but it stops short of assessment or recommendation.

The consideration is that even without formal recommendations, contextual guidance can shape perceptions of urgency and priority. Conservative framing in crisis scenarios is essential. The tool should not collect sensitive data beyond immediate query needs and should not assess eligibility or recommend solutions.

6.2.3 Structured triage (2.3)

Several participants noted that self-serve tools could serve as a gateway for people who were not yet ready to engage with a formal service, potentially reducing the delay between problem recognition and help-seeking (AP22, AP29). A well-designed triage tool could identify urgency, route high-risk cases to immediate human contact and help lower-complexity cases prepare for their appointment.

This is the point at which the consequences of getting it wrong become more serious. A triage tool that fails to identify vulnerability, misclassifies urgency or gives false reassurance could cause real harm. The research found that vulnerability often emerges gradually through trust-building over multiple interactions. Clients may minimise or withhold aspects of their circumstances due to embarrassment or stigma (AP07, AP09). In many cases, advisers only become aware of underlying issues through careful questioning, observation and review of supporting documentation. Gen AI systems accepting responses at face value face significant limitations in this area.

71% of sector survey respondents said AI should never assess vulnerability. The interview evidence demonstrated that the most consequential vulnerabilities, domestic abuse, financial coercion, safeguarding concerns and capacity issues, emerge through sustained human interaction and are invisible to systems operating on self-reported information alone. Failure to detect vulnerability could lead to inappropriate insolvency decisions, undetected abuse or missed safeguarding referrals.

6.2.4 Pre-appointment data gathering (2.4)

Pre-appointment data gathering emerged as a specific opportunity: Gen AI guiding clients through initial budget preparation, flagging inconsistencies and gathering documents before an adviser appointment. 53% of sector survey respondents considered this suitable now.

The consideration is fundamental. The financial statement was widely described as the core foundation of debt advice, and all 17 sector survey respondents agreed that an accurate financial statement is essential. While automated systems may assist in gathering financial information, constructing a sustainable budget requires interpretation and professional challenge. Advisers frequently encounter clients who report unrealistically low expenditure because they underestimate costs, omit irregular expenses or suppress spending on essential needs. The interview evidence demonstrated that self-reported expenditure is unreliable in the vast majority of cases, that open banking captures spending but not need, and that strategic budget construction involves professional judgement about how to present truthful information to different creditors and institutions (AP03, AP22).

Automated budget tools accepting self-reported figures without professional challenge risk producing assessments that appear complete but are fundamentally unsound. Solutions built on unrealistic financial statements may fail over time, leading to recurring debt problems. Any tool at this level must position all information as preliminary and subject to mandatory adviser review.

6.2.5 Communication support (2.5)

Creditor letter coaching was seen as particularly promising, with multiple participants noting that clients already use Gen AI for correspondence, which they then send to advisers for review. While adequate for standard holding letters, this was not trusted for consequential communications such as statute-barred debt responses or write-off requests (AP29, AP35). Tone-checking for client correspondence was identified as a specific valuable feature (AP35).

Sector-specific tools could improve the quality and safety of these interactions compared with generally available Gen AI. The key distinction is between standard communications, which carry lower risk, and consequential communications that could affect a client's legal position and require adviser review.

6.2.6 Documentation support (2.6)

Participants identified tools that help clients organise paperwork, understand correspondence and prepare documents for appointments. Many clients arrive at advice appointments with disorganised documentation and a limited understanding of what they have received. A tool that

helps clients sort and categorise their documents could improve the efficiency of the advice session and reduce client anxiety.

The consideration is that these tools should help with organisation, not interpretation. Flagging enforcement or court-related documents as urgent is appropriate. Interpreting the legal or financial significance of correspondence may not be.

6.2.7 Ongoing engagement (2.7)

Appointment reminders, progress tracking, repayment prompts and re-engagement with clients who have dropped out of the advice process were identified as areas where Gen AI could add value. Debt advice is rarely a single interaction, and clients may need sustained support to maintain progress.

The consideration is tone and sensitivity. Clients in financial difficulty are often already experiencing stress. Contact must be sensitive in frequency and tone, must adapt to non-engagement rather than escalating pressure and must offer help rather than demand action. Engagement data should identify the need for support, not be used to penalise.

Sector appetite: graduated and conditional. The sector is broadly supportive of general information and communication support tools, but increasingly cautious as influence over client decisions grows. The suitability of Level 2 tools depends heavily on the nature of the client base served. Some advice providers primarily serve individuals who present in acute crisis with limited digital access. In such contexts, fully self-serve systems may have limited applicability. Organisations that already serve large numbers of clients through digital channels may find self-serve tools more appropriate in certain stages of the advice journey.

6.2.8 Technology readiness

The research found that current Gen AI accuracy is insufficient for some applications within Level 2. Purpose-built tools tested by participants achieved approximately 75% accuracy against a threshold of 85–90%. General-purpose Gen AI tools returned outdated information, US-sourced content for English queries and fabricated legal citations. Gen AI demonstrated a fundamental deficit in reasoning: competent at recording facts but unable to explain why a particular option was chosen or how competing factors had been weighed. Organisations considering Level 2 deployment should satisfy themselves that the specific tool they are evaluating meets the accuracy requirements for the application type in question.

6.3 Level 3: Agentic AI: Taking Action and Sustained Relationships

Level 3 addresses a fundamentally different category of AI. Where Levels 1 and 2 use Gen AI to produce content and information, Level 3 involves Agentic AI: systems that take autonomous actions, execute multi-step workflows, use external tools and maintain sustained evolving relationships with clients. This is AI that does things, not just says things.

The research found no organisations currently deploying Agentic AI in regulated debt advice. This level is largely forward-looking, as the technology landscape is moving rapidly toward agentic

capabilities and the sector needs a framework for assessing these developments as they emerge rather than responding reactively.

The inclusion of Level 3 draws on this research but also on Wyser's wider work in AI project implementation, including work with the Motor Ombudsman integrating summarisation and retrieval-augmented generation into case handling, and work with the Ministry of Justice's Private Family Law Programme. These projects, alongside the rapid advancement of Agentic AI platforms more broadly, indicate that the capabilities described in Level 3 are not theoretical. They are being developed and deployed in adjacent regulated sectors. The debt advice sector will encounter them.

6.3.1 Personalised advisory (3.1)

An Agentic AI tool operating at this level would maintain an ongoing relationship with the individual, store their information across sessions and provide personalised guidance that evolves over time, including scenario modelling and progress tracking. This moves beyond generating information into maintaining a sustained advisory influence.

The sector expressed significant caution. Personalised advisory tools were rated as suitable by only 18% of survey respondents for non-priority debts and 12% for all debts. The research identified specific cases of AI incorrectly ruling out DRO eligibility for deficit-budget clients (AP34), directing clients toward bankruptcy without identifying circumstances making it unsuitable (AP20, AP35), and confidently presenting IVAs as solutions when DROs would have been more appropriate and free. The risk is compounded by clients' tendency to grasp at any confidently presented option when in crisis.

The evidence suggests this application type may not be currently appropriate for most organisations without significant investment and testing.

6.3.2 Adviser-directed actions (3.2)

Agentic AI executing specific tasks under direct adviser instruction: filing forms, sending pre-approved correspondence and updating case management records. The adviser initiates every action, confirms before execution and can cancel or reverse.

This is the Level 3 application type most likely to emerge in the near future. It represents a natural extension of Level 1 tools: moving to AI drafting, formatting, and sending the letter, still under adviser supervision. The critical principle is that every consequential action, anything affecting a client's legal position, financial obligations or access to services, must have explicit human approval before it happens.

6.3.3 Semi-autonomous actions (3.3)

Agentic AI acting within pre-approved boundaries with approval gates for consequential actions. An example might be a system that automatically sends appointment reminders, generates routine correspondence or updates records within pre-approved rules, but requires an adviser's sign-off before sending anything to a creditor or filing a statutory form.

Commented [GM23]: this doesn't make sense "rather than the AI drafting a letter for the adviser to review, it drafts, formats and sends under the adviser's instruction" - I think you meant to say that the "AI drafts, formats and sends under the advisers instruction"

Commented [MG24R23]: Updated

This requires clearly defined and documented action boundaries, robust logging, override capability and mandatory limited-caseload piloting. Outcome monitoring against human-directed equivalents is essential to establish whether the Agentic AI is performing safely.

6.3.4 Autonomous process management (3.4)

Agentic AI independently manages multi-step processes. The findings of this research make it clear that this may not currently be appropriate for regulated debt advice. The combination of legal complexity, the prevalence of vulnerability, jurisdictional variation, and the potential for irreversible harm means that autonomous process management would require a level of reliability, accountability, and regulatory clarity that does not yet exist.

This application type is included to provide the sector with a structured way to assess fully autonomous capabilities if and when they become viable. It would require extensive piloting to demonstrate equivalent outcomes, full liability and accountability frameworks, automatic suspension if outcomes fall below thresholds, and sector-wide agreement on use cases and boundaries. It may also require regulatory approval or explicit guidance.

Sector appetite: very cautious. The sector has limited appetite for Level 3 applications at present, and for good reason. The technology is evolving faster than the governance, regulatory and accountability frameworks needed to support it in a regulated advice context. However, agentic capabilities are advancing at pace across the wider technology landscape, and the sector is likely to encounter Level 3 applications sooner than many expect. The inclusion of this level ensures that when those conversations arise, the sector has a structured basis for assessment rather than having to start from scratch.

6.4 Using the Roadmap

The Debt Advice AI Roadmap can be used at three levels. At the organisational level, providers can assess which levels and application types are appropriate given their client base, governance capacity and technical infrastructure. An organisation primarily serving highly vulnerable clients through face-to-face appointments may determine that Level 1 is appropriate and that client-facing tools would not serve their clients well. This is a legitimate and well-evidenced conclusion, not a failure to innovate.

At the tool level, when evaluating or developing a specific AI tool, the Roadmap provides a structured way to assess where it sits, what safeguards are required and what governance arrangements need to be in place before deployment.

At the regulatory level, regulators can develop proportionate oversight that calibrates expectations to the level and application type, rather than applying a single set of requirements to all AI tools regardless of function or risk profile.

The accompanying Safe-Use Framework provides detailed guidance for each application type, including organisational readiness requirements across six dimensions, minimum standards, escalation triggers and a possible implementation roadmap.

6.5 The role of general-purpose Gen AI tools

An important theme throughout the research is that many individuals already have access to general-purpose Gen AI tools. Even without sector-specific systems, members of the public consult these when seeking information about debt or financial problems. The YouGov survey found that 37% of UK adults have used an AI chatbot for personal purposes, though only 21% would be likely to use one for advice about financial issues.

Attempts to restrict AI use within the debt advice sector would not prevent individuals from consulting publicly available systems. The FCA has stated that the Financial Ombudsman Service or the Financial Services Compensation Scheme does not cover advice from general-purpose AI tools. Yet individuals in financial distress are consulting these tools and acting on the guidance. The research identified specific cases, including clients directed toward bankruptcy based on AI advice (AP20, AP35) and clients presenting with AI-informed complaints citing irrelevant legislation.

This reality underscores the case for sector-specific tools operating within the Roadmap's framework. Well-governed, accuracy-tested tools that draw on curated knowledge sources are likely to produce better outcomes than the uncontrolled use of general-purpose Gen AI. The sector's engagement with Gen AI is not simply about efficiency; it is also about harm reduction.

6.6 Sector conditions for effective AI use

Participants consistently felt that meaningful use of AI within debt advice would require close collaboration between technology developers and the advice sector. Several expressed concern that tools developed without a deep understanding of the sector's realities could easily miss critical nuances, particularly around complex regulatory frameworks, jurisdictional differences, evolving case law and diverse client vulnerabilities.

The strongest message from the interviews was clear: tools developed without sector involvement will fail. The complexities of jurisdictional variation, emergent vulnerability, tactical budget construction and interpretive debt solutions are not apparent from the outside. Technology developers seeking to create tools for the sector must embed sector expertise at every stage of design, development and testing.

Some of the technical requirements emerging from the research that technology developers should consider include jurisdictional routing, distinguishing England and Wales from Scotland, knowledge base updates within days for case law changes, handling of messy real-world financial data, including multiple accounts, family lending and irregular income, and client-facing content at an appropriate reading age. Pricing should consider organisation size to ensure that they are not locked out of innovation, particularly for Level 1 tools. Sector-level licensing, shared infrastructure models or tiered pricing should be considered to avoid widening the digital capability gap.

7 Safeguards for more robust deployment of Self-Serve

The evidence gathered from sector experts suggests that while Gen AI may offer potential opportunities, deployment requires careful design and governance, including escalation and oversight. The question is not simply whether Gen AI tools should be used, but how they should be deployed to maintain consumer protection and professional accountability.

7.1 Restricted and reliable knowledge sources

One of the most immediate safeguards concerns the sources of information used by Gen AI tools. Debt advice operates within a regulatory and legal environment that is both complex and subject to ongoing change. Gen AI tools intended to support debt advice should rely on clearly defined, authoritative knowledge bases rather than indiscriminately drawing from open internet sources. These knowledge bases should include version control, regular review, jurisdictional accuracy, and rapid update mechanisms for case-law changes.

Safeguards include limiting Gen AI systems to approved knowledge sources, maintaining clear version control for guidance materials, implementing regular review and update processes, and ensuring that jurisdiction-specific guidance is correctly applied. These measures may help reduce the risk of outdated or inaccurate information being presented to users.

The interview evidence demonstrated that even purpose-built tools trained on curated content can produce incorrect answers (AP26) and that the speed of change in the debt advice knowledge base means any system requires continuous maintenance at least as rigorous as that required for human advisers.

7.2 Clear escalation pathways

Gen AI tools must be designed to identify potential indicators that should trigger referral to a human adviser. Circumstances that may require escalation include evidence of enforcement action, such as bailiffs or court proceedings; rent arrears or housing instability; indications of safeguarding concerns or coercive control; signals of mental health distress or an acute financial crisis; and incomplete or inconsistent financial information.

Escalation thresholds may need to be conservative, recognising the limitations of automated systems in interpreting complex situations. The cost of over-escalation, routing someone unnecessarily to a human adviser, is lower than the cost of under-escalation.

Organisations are best placed to tailor their escalation paths based on their knowledge of their client base. Still, the following were escalation triggers highlighted by the research and ranked as follows in the sector survey:

Client confusion (76%), low digital confidence (76%), repeated contradictory responses (71%), bailiff involvement (59%), financial coercion indicators (59%), court proceedings (53%) and mental health disclosure (53%). The interview evidence adds priority debt of any kind, silence or evasion

in response to financial questions and any indication of safeguarding concerns. Thresholds should be conservative, with the default being human escalation when confidence is low.

7.3 Human-in-the-Loop

There was a strong consensus that decisions involving the selection of debt solutions should remain subject to professional oversight. Human review may be required for the validation of financial statements, the assessment of solution eligibility, the interpretation of complex financial circumstances, and the final recommendations regarding formal debt solutions. Maintaining professional involvement ensures that advice reflects the full context of a client's circumstances.

Individual organisations should consider the points at which they want to include human interaction. It may be that sufficient testing and piloting of a tool before deployment is sufficient for some organisations. Still, for others, a human may always be required to finalise options. Decisions on debt solutions must remain subject to professional oversight.

7.4 Transparency and auditability

Gen AI-assisted processes must remain transparent and auditable. Advice providers operate within regulated environments where decisions may need to be reviewed by supervisors, auditors or regulators. Safeguards include maintaining logs of Gen AI interactions and outputs, recording the sources of information used, ensuring advisers can review and override system outputs, and enabling audit trails for quality assurance. These measures help ensure that automated systems do not create opaque decision-making processes that are difficult to review.

Participants consistently expressed the view that accountability for AI-generated advice should rest with the organisation deploying the tool, with shared responsibility between the funder and the delivery organisation where funding relationships exist (AP21). Several participants called for MaPS certification of Gen AI tools used within funded services (AP36) and for named organisational (though not individual) accountability (AP12).

7.5 Clear scope and boundary setting

Users must understand when they are receiving general information and when professional advice is required. Systems should communicate clearly about the purpose of the tool, the limits of automated guidance, and when to seek human advice. In practice, this may involve designing tools that explicitly direct users toward professional advice services when complex or high-risk situations are identified.

7.6 Sector collaboration and governance

The successful deployment of Gen AI tools within the debt advice sector will require ongoing collaboration between technology developers, advice organisations and regulators. Technology solutions developed without meaningful sector involvement may fail to capture the complexities of real-world advice practice. Responsible deployment requires sector experts to be involved in

system design, ongoing testing with advice providers, clear governance arrangements for system oversight, and collaboration with regulators to clarify acceptable use cases.

7.7 Quality assurance

Gen AI tool outputs should be subject to quality checking analogous to that applied to human advisers. Multiple participants emphasised that the absence of complaints does not indicate the absence of errors (AP21, AP34). Periodic testing, including the adviser-versus-AI comparison model proposed during the research (AP18), should be embedded in organisational governance from the outset.

8 Strategic Implications

The evidence gathered suggests that while Gen AI may offer meaningful support in certain areas, the structure of debt advice work poses challenges for automated systems seeking to replicate professional judgement. This final section considers the implications for regulators, advice providers, technology developers and future research.

8.1 Implications for regulators

The growing availability of Gen AI technologies raises important questions for regulators overseeing financial guidance and debt advice. A key theme is that AI in debt advice should not be viewed as a single category of activity; different forms carry varying levels of risk depending on the extent to which they influence financial decision-making.

The Debt Advice AI Roadmap presented in this report offers a structured approach for distinguishing between these levels of influence. It may help regulators and sector bodies develop proportionate guidance that supports innovation while maintaining consumer protection.

The growing use of general-purpose Gen AI tools by the public also raises important considerations: individuals may increasingly consult Gen AI when seeking information on financial difficulties, even in the absence of sector-specific systems. Potential regulatory considerations include clarifying the boundaries between general information and regulated advice, encouraging transparency around the limitations of Gen AI-generated guidance, and establishing expectations for governance where AI tools are integrated into advice services.

8.1.1 Address the regulatory gap around general-purpose Gen AI tools

The FCA has stated that the FOS or FSCS does not cover general-purpose AI advice. Yet individuals in financial distress are consulting these tools and acting on guidance. At least two cases of bankruptcy based on AI advice were reported in this research (AP20, AP35), along with a case of property loss (AP21). The sector should press for protections before harm scales.

8.1.2 Clarify the information-advice boundary for AI systems

A Gen AI tool describing a DRO provides information; one that assesses circumstances and suggests a DRO is appropriate is functionally providing advice. Regulators should define this boundary explicitly for AI-delivered content.

8.1.3 Develop proportionate oversight using the roadmap

Level 1 requires different oversight from Level 2.4 pre-appointment data gathering, which requires different oversight from Level 3 agentic applications. This graduated approach supports innovation at lower levels while maintaining protections where AI touches the core regulated journey.

8.1.4 Require accuracy standards

No UK regulator certifies the accuracy of Gen AI chatbots or requires hallucination testing. In debt advice, where errors can lead to housing loss or inappropriate insolvency, minimum accuracy thresholds should be established for Gen AI tools that provide debt-related information.

8.2 Implication for sector leaders

8.2.1 Coordinate a collective sector response

Sector bodies should develop shared standards for safe AI use, facilitate knowledge sharing between organisations and represent sector interests in regulatory discussions. The Debt Advice AI Roadmap and the Safe-Use Framework provide a starting point.

8.2.2 Address workforce concerns proactively

Adviser anxiety about displacement is real. One participant described training many advisers just to experience high turnover driven by the emotional burden of the role (AP06). Organisations that positioned Gen AI as an efficiency tool for admin found staff became positive once they experienced time savings (AP05). Clear communication, investment in training and genuine engagement are essential.

8.3 Implications for advice providers

Gen AI adoption decisions should be guided by both the nature of the technology and the characteristics of the organisation's client base. The populations served by different advice providers vary significantly; some organisations support individuals presenting in acute crisis with complex vulnerabilities, while others engage large numbers of clients through digital channels.

Adviser-assist tools that support administrative workflows appear to have the broadest potential applicability across the sector. Fully automated self-serve tools may be more appropriate in limited contexts, particularly where users already engage digitally, and systems are carefully designed with appropriate escalation pathways. Effective implementation will require investment in governance, training and oversight.

8.3.1 Start with Level 1

Adviser assistance tools have the broadest applicability, the strongest evidence, and the lowest risk. Specific applications include Gen AI transcription and case note summarisation, correspondence drafting, credit report analysis, policy lookups and QA support. Always review Gen AI content before client or creditor use.

8.3.2 Assess your client base before considering Level 2

Before client-facing self-serve, undertake an honest assessment: what proportion could realistically engage digitally? Multiple participants estimated that approximately 50% of their clients were digitally excluded or would not engage with AI tools (AP06, AP28). What happens to excluded clients if resources shift toward digital? For some organisations, Level 1 may be the appropriate extent of AI adoption.

8.3.3 Where self-serve is appropriate, start with applications 2.1 & 2.2 to progress cautiously

Begin with general information and contextual query support. Build conservative escalation from the outset. Ensure every interaction includes a pathway to human advice. For some organisations, the most appropriate first step is not self-serve at all, but rather adviser-assistance tools that build organisational confidence and governance capability before any client-facing deployment.

8.3.4 Develop organisational AI policies

Multiple participants described staff using AI tools informally, without organisational guidance, which created data protection and quality risks (AP35). One organisation was publishing its AI usage policy within days of the interview. Clear policies covering approved tools, data protection boundaries, and quality expectations should be established before adoption, rather than after.

8.4 For technology developers

Technology developers seeking to create tools for the debt advice ecosystem may benefit from working closely with advice organisations during system design, incorporating sector expertise into knowledge bases and training data, testing tools within real-world advice environments, and designing systems that support rather than replace professional judgement. Collaboration helps ensure that technology solutions address genuine operational needs rather than attempting to automate processes that rely heavily on professional interpretation.

8.4.1 Build with the sector, not for the sector

The strongest message from interviews: tools developed without sector involvement will fail. The complexities of jurisdictional variation, emergent vulnerability, tactical budget construction, and interpretive debt solutions are not apparent from the outside. Embed sector expertise at every stage.

8.4.2 Ensure affordability

Consider sector-level licensing, shared infrastructure models or tiered pricing. Technology accessible only to large providers risks widening the digital capability gap.

8.5 For the Money and Pensions Service

8.5.1 Review documentation proportionality

The evidence from this research is unambiguous: current MaPS documentation requirements are perceived as disproportionate by the advisers delivering funded services. The administrative burden is the single strongest driver of interest in AI tools. MaPS should assess whether extensive documentation achieves its purpose or consumes capacity better directed to client work.

8.5.2 Commission shared AI infrastructure

Smaller organisations cannot afford to develop AI independently. MaPS could commission shared tools: knowledge bases, sector-specific LLM fine-tuning, common evaluation frameworks or centrally procured adviser-assist tools for funded organisations.

8.5.3 Fund further research

The AI arena is rapidly evolving, and while this research is timely, further research should be funded as the technology evolves. Possible areas for exploration are outlined in the Conclusion section of this document.

9 Conclusion

This research set out to explore the potential role of Gen AI self-serve tools within regulated debt advice. The findings suggest that debt advice is fundamentally relational, iterative and judgement-based. Advisers work with individuals whose financial circumstances are often complex, incomplete or evolving, and effective advice frequently depends on professional interpretation developed through experience.

AI technologies offer valuable support in certain areas, particularly reducing administrative workload, improving information access and extending expertise. However, systems that attempt to replicate or replace professional decision-making introduce significant risks. The evidence of harm, while not yet widespread, is sufficient to warrant proactive attention: clients being directed toward inappropriate insolvency solutions, clients losing property by following AI guidance over professional advice, and purpose-built tools achieving accuracy levels well below those required for a regulated environment.

The Debt Advice AI Roadmap provides a structured approach for the sector to assess where Gen and Agentic AI may be applied safely. The central principle is proportionality: as influence over client outcomes increases, so must governance, oversight and safeguards. The Roadmap's three levels provide a shared vocabulary that accommodates both current applications and future developments, while the Safe-Use Framework offers the detailed guidance needed for responsible implementation.

The sector's structural challenges, the prevalence of deficit budgets, decreasing funding of free advice services, the administrative burden of regulatory compliance and the attrition of experienced advisers cannot be solved by technology alone. As one adviser observed: *if we don't do it, there's nobody else that's going to help that client do it*. Gen AI can reduce the burden and may eventually enable some lower-complexity interactions to be handled without direct adviser involvement. But it cannot replicate the professional judgement, emotional intelligence and sustained relationship that the evidence shows are essential for effective debt advice to the most vulnerable.

The public is not waiting. People are using Gen AI for debt information, and some are acting on it. The question is not whether to engage, but how to shape the responsible use of Gen AI, ensuring that innovation strengthens rather than undermines the principles of free, impartial, high-quality advice.

9.1 Areas for Further Research

Several areas would benefit from further research, including: evaluating the effectiveness of AI-supported triage systems in improving access to advice; exploring how AI tools may support adviser training and knowledge management; assessing creditor perceptions and use of AI in relation to debt management; and assessing the long-term outcomes of AI-supported advice processes.

Commented [TK25]: These are different to the areas mentioned at the end of Ch 8

Commented [MG26R25]: Updated 8 to reference this section

Appendix 1: Sector Perspectives on AI Readiness

Findings from the sector survey (n=17). While not statistically significant, the responses provide useful directional data that can be triangulated with qualitative evidence.

Attitudes toward AI

Test in limited pilots: 7. Actively explore: 5. Proceed very cautiously: 3. Avoid AI: 2. Somewhat confident in safe deployment: 9. Not at all ready or slightly ready: 12 of 17. The dominant position was cautious openness combined with awareness of low readiness, a position captured by one interview participant as: we're going to have to roll with it. It's coming, so we need to think about how we can use it in the best way possible (AP12).

Where AI should not be used

Final solution recommendation: 12/17. Vulnerability assessment: 12/17. Insolvency eligibility: 11/17. Creditor communication: 7/17. These exclusions were strongly supported by the interview evidence, which demonstrated that each of these areas involves precisely the kind of contextual judgement, hidden complexity and potential for serious harm that current AI cannot reliably navigate.

What should come first?

Adviser-support tools were most frequently top-ranked, consistent with the Debt Advice AI Roadmap recommendation to begin with Level 1. General information and structured triage were also ranked highly. Personalised advisory tools were consistently ranked last. This ordering was identical to the pattern that emerged from the 28 interviews.

Appendix 2: The consumer perspective: Public attitudes to AI and financial guidance

Findings from the YouGov survey of 2,060 UK adults conducted in March 2026 establish a baseline for public attitudes to AI and financial guidance.

Current use of AI chatbots

37% of UK adults have used an AI chatbot for personal advice. 59% reported never having used one. The population currently accessing community-based debt advice tends to skew toward lower digital confidence, suggesting direct AI use among existing clients may be lower than these figures indicate. This was consistently confirmed across interviews, with the majority of participants (approximately 18 of 28) reporting no direct evidence that their clients used AI tools for debt advice.

Willingness to use AI for financial advice

Only 21% said they would be likely to use AI for financial advice (6% very likely, 14% fairly likely). 75% said unlikely, with 47% not at all likely.

Confidence in AI accuracy

Only 22% expressed confidence in the accuracy of AI financial advice (3% very confident, 19% fairly confident). 72% were not confident.

Preference for human advice

60% strongly preferred human advice, 21% somewhat preferred it. Only 3% expressed any AI preference. 12% would use both equally. This was echoed across the interview programme, where every participant, regardless of their overall attitude to AI, identified human interaction as essential for complex or vulnerable cases.

How people would respond to AI financial advice

The most common responses were to use it as a starting point for further research (37%), compare it with official sources (37%), research online to verify (37%), and check with a human expert (36%). Only 2% would follow AI advice exactly as given. This pattern is significant: while the public would not follow AI advice uncritically, those in financial crisis may behave differently, as 94% of sector survey respondents confirmed that clients tend to accept any confidently presented option when distressed.

What would increase trust?

Regulation by financial authorities (32%), human expert review (27%), up-to-date regulations (21%) and clear reasoning (18%) were the strongest factors. 28% said nothing would increase

their trust. Free-text responses revealed deep concerns about data security, commercial use and the impersonal nature of AI-driven guidance.

Implications

The data reinforces that AI self-serve tools deployed without robust safeguards and clear human oversight risk low adoption and potential harm. It also highlights a growing risk from unregulated AI use by individuals who do not understand its limitations, particularly those in financial distress. The gap between the general population's cautious scepticism and the crisis-driven behaviour of people in problem debt is a critical consideration for deployment strategy.

Appendix 3: Glossary of Terms

Artificial Intelligence

Artificial Intelligence (AI)	A broad field of computer science concerned with building systems capable of performing tasks that typically require human intelligence, such as learning, reasoning and problem-solving. In this report, AI is used as a general term encompassing both generative AI and agentic AI, rather than the whole field of AI
Generative AI (Gen AI)	AI systems powered by large language models that can generate text, hold conversations, summarise information and answer questions. Gen AI produces content based on statistical patterns learned from training data rather than by understanding or verifying facts. In this report, Gen AI refers to tools that allow individuals or advisers to obtain guidance, information or structured assistance
Agentic AI	AI systems that can take autonomous actions, break complex tasks into steps, use external tools and execute multi-step workflows with minimal human input. Unlike Gen AI, which generates content for a human to review, agentic AI acts by filing forms, sending correspondence, managing cases, and making decisions within defined boundaries. Agentic AI raises particularly significant questions for regulated environments because the more autonomous a system becomes, the greater the need for oversight and governance
Large Language Model (LLM)	The underlying technology powering Gen AI. LLMs are trained on vast amounts of text data to learn statistical patterns and probabilistic relationships between words and concepts. When a user asks a question, the model generates a response by predicting the most likely sequence of words. LLMs do not verify facts, do not access live databases unless specifically connected to them and do not understand the legal or regulatory implications of what they produce
Hallucination	A phenomenon in which an AI system generates information that sounds authoritative and plausible but is factually incorrect. In debt advice, hallucinations are particularly dangerous because small inaccuracies can have material consequences for clients, including directing them toward inappropriate debt solutions or providing incorrect information about legal rights and obligations
Retrieval-Augmented Generation (RAG)	A technique that connects a Gen AI system to a curated knowledge base so that responses are grounded in specific, verified source material rather than relying solely on the model's training data. RAG can improve accuracy by restricting the information the AI draws upon, though it does not eliminate the risk of hallucination entirely

Fine-tuning	The process of further training an existing AI model on a specific dataset to improve its performance in a particular domain. In debt advice, fine-tuning could involve training a model on curated case law, regulatory guidance and sector-specific content to improve the accuracy and relevance of its outputs
Token	The basic unit of text that a large language model processes. A token can be a word, part of a word or a punctuation mark. Models have limits on the number of tokens they can process in a single interaction, which constrains the length and complexity of conversations

Debt Advice Sector

Regulated debt advice	Professional guidance provided to individuals on managing personal debt, including assessing financial circumstances and identifying appropriate solutions. In the UK, this activity operates within regulatory frameworks overseen by the Financial Conduct Authority
Financial statement	A detailed assessment of an individual's income and expenditure that forms the foundation of all regulated debt advice. The financial statement determines what a person can afford to repay and which debt solutions are appropriate. The research found that financial statements involve professional judgement and strategic construction rather than simple arithmetic, and that self-reported financial data is unreliable in the majority of cases
Standard Financial Statement (SFS)	A commonly used format for recording income and expenditure in debt advice, maintained by the Money and Pensions Service. The SFS provides a standardised structure for financial assessments and is used across many advice organisations
Deficit budget	A financial position in which an individual's essential living costs exceed their income, leaving nothing available for debt repayment. The research found that an increasing proportion of debt advice clients present with deficit budgets
Priority debt	Debts where non-payment carries the most serious consequences, such as loss of housing, imprisonment, disconnection of essential services or loss of essential goods. Priority debts include mortgage or rent arrears, council tax, energy arrears, child maintenance and certain tax debts. Correctly identifying and prioritising these debts is a core element of professional debt advice
Non-priority debt	Debts where the consequences of non-payment, while serious, are generally less immediate than priority debts. Non-priority debts typically include credit cards, personal loans, catalogue debts and overdrafts.

	<p>Clients frequently misunderstand the distinction between priority and non-priority debts</p>
Breathing Space	<p>A statutory debt respite scheme in England and Wales that provides individuals with legal protection from creditor action for up to 60 days. During this period, most enforcement action, interest and charges are paused. The research identified the strategic use of Breathing Space as an example of professional judgement that automated systems would struggle to replicate, with some advisers deliberately holding it in reserve rather than applying it immediately</p>
Individual Voluntary Arrangement (IVA)	<p>A legally binding agreement between an individual and their creditors to repay debts over a fixed period, typically five to six years, overseen by an insolvency practitioner. IVAs were cited across the research as the debt solution most frequently subject to commercial mis-selling and inappropriate marketing</p>
Debt Relief Order (DRO)	<p>A form of insolvency available to individuals with low income, few assets and debts below a specified threshold. The research identified cases where AI tools had incorrectly ruled out DRO eligibility</p>
Sequestration	<p>The Scottish equivalent of bankruptcy. Scotland operates an entirely separate legal system with distinct statutory solutions, different enforcement timescales and separate regulatory frameworks, making jurisdictional accuracy a fundamental requirement for any AI tool operating across the UK</p>
Advice shopping	<p>A pattern in which individuals consult multiple advice sources seeking confirmation of a preferred course of action rather than objective guidance. The research identified this as an existing behaviour that Gen AI tools could potentially amplify if they provide responses that align with what a user wants to hear rather than what they need to know</p>
Tacit knowledge	<p>Professional expertise developed through years of practice that is difficult to articulate or codify. In debt advice, tacit knowledge includes the ability to read between the lines in client conversations, identify subtle indicators of vulnerability, interpret creditor behaviour and construct financial statements strategically. This knowledge represents a significant barrier to AI automation</p>

Regulation and Governance

Financial Conduct Authority (FCA)	The UK's financial services regulator. The FCA has confirmed that it will not introduce AI-specific regulations, instead relying on a technology-neutral, principles-based, outcomes-focused approach using existing frameworks
Consumer Duty	A regulatory framework introduced by the FCA that requires firms to deliver good outcomes for retail customers. Consumer Duty sets higher standards for the care that firms give to consumers and is relevant to organisations deploying AI tools that interact with or affect clients
Financial Ombudsman Service (FOS)	An independent body that resolves disputes between consumers and financial services firms. The FCA has stated that advice provided by general-purpose AI tools is not covered by the FOS
Financial Services Compensation Scheme (FSCS)	A scheme that compensates consumers when authorised financial services firms fail. As with the FOS, the FSCS does not cover advice provided by general-purpose AI tools
Senior Managers and Certification Regime (SM&CR)	A regulatory framework that sets standards of conduct for individuals working in financial services and holds senior managers accountable for the activities within their areas of responsibility

Organisations

Money Advice Liaison Group (MALG)	A cross-sector membership body that brings together creditors, debt advice providers, government and regulators to improve practice and outcomes in the debt advice sector. MALG commissioned this research
Money and Pensions Service (MaPS)	An arm's-length body sponsored by the Department for Work and Pensions. MaPS is responsible for improving the financial wellbeing of UK citizens and funds debt advice services across England. MaPS funded this research through its Debt Advice Transformation Fund 2025/26
Wyser	A specialist consultancy with expertise in AI, data and technology. Wyser delivered this research as the research and delivery partner

Research and Framework Terminology

Debt Advice AI Roadmap	A three-level framework developed through this research that provides a structured approach to understanding where AI may be applied safely in regulated debt advice. The three levels are: Level 1 (adviser assistance), Level 2 (client self-serve information and assistance) and Level 3 (agentic AI taking action and sustained relationships)
Safe-Use Framework	The practical companion document to the Debt Advice AI Roadmap, providing organisational readiness assessments, minimum standards, safeguards and escalation triggers for each application type within the Roadmap
Human-in-the-loop	A design principle in which a human professional reviews, approves or oversees AI-generated outputs before they are acted upon or shared with clients. The research found strong consensus that decisions involving the selection of debt solutions should remain subject to professional oversight
Escalation	The process of transferring a client from an automated interaction to a human adviser when certain triggers are detected. The research identified conservative escalation as a key safeguard: the cost of over-escalation (routing someone unnecessarily to a human) is lower than the cost of under-escalation (failing to identify a situation requiring professional intervention)
Self-serve	Tools that allow individuals to obtain guidance, information or structured assistance without direct interaction with a human adviser at the point of use. Importantly, self-serve does not imply the complete absence of human oversight. In many potential use cases, self-serve tools may gather information or provide preliminary support before a professional reviews the case
Triage	The process of assessing the nature and urgency of an individual's situation to determine the most appropriate level of support or referral. In debt advice, effective triage involves identifying vulnerability, assessing complexity and routing individuals to the right service
Open banking	A system that allows third-party providers to access bank transaction data with the account holder's consent. While useful for identifying spending patterns, the research found that open banking captures what has been spent but not what should have been spent, limiting its value for constructing accurate financial statements
Vulnerability	In the context of debt advice, vulnerability is rarely a static or declarable condition. It is dynamic, layered and frequently hidden. The research found that the most consequential vulnerabilities, including

domestic abuse, financial coercion, safeguarding concerns and capacity issues, often emerge only through sustained professional interaction rather than disclosure in response to direct questions
