OPEN BANKING

The Standard for Open Innovation™

DO LLMS HAVE STANDARDS?

Findings from OBL's AI consultation and implications for Open Banking and Open Finance



Author: Nick Davey OBL Head of Strategy

Date: November 2026

Enquiries: nick.davey@openbanking.org.uk

CONTENTS

At a glance: what this report finds	3
Foreword	4
Henk Van Hulle	4
CEO, Open Banking Limited	4
Author's Preface	5
Executive Summary	6
What respondents told us	6
What we conclude	7
Green Paper and Scope	8
Responses	8
Q1 — Can AI reduce the cost and complexity of API data provision?	8
Q2 — Does the Open Banking Standard give AI what it needs? What improvements could make it more AI-ready?	12
Q3 — Can AI help build entirely new customer and business products from scratch?	13
Q4 — What have we missed?	14
Thematic Views and Respondent Snapshots	15
Affordability and risk in regulated sectors	15
Standards and agentic AI readiness	16
Implementation variance and privacy automation	16
Layered AI services	16
TPP perspectives	16
Cross-cutting Insights	16
Implications for Open Banking Ecosystem Participants	17
A Trust and Certification Model for AI-enabled Open Banking	18
Alignment with UK Policy Direction (Mills Review and Subsequent Work)	18
Final Takeaways.	20
Glossary	21

Do LLMs Have Standards?

Findings from OBL's AI consultation and implications for Open Banking and Open Finance

At a glance: what this report finds

- AI could help the next phase of Open Banking and Open Finance be delivered faster and more efficiently. Artificial Intelligence Large Language Models (LLMs) can help organisations map and clean legacy data, generate and test code, improve documentation and prepare data for APIs, reducing manual effort across development.
- The opportunity is significant, but savings are not yet proven for Open Banking. Research in comparable activities suggests productivity improvements of around 10–30%[1]. While any efficiency gains are untested for Open Banking, the report therefore uses up to 30% as a planning assumption for well-governed, repeatable activities, to be tested and validated in practice.
- AI works best with accurate, permissioned and well-structured data, while poor-quality or ambiguous data can amplify errors. Validation, traceability and proportionate assurance will therefore remain essential.
- Open Banking standards need to be easy to access and interpret by machines as well as people. A single authoritative, version-controlled source, richer machine-readable information and clearer documentation could help AI tools deliver more consistent code, testing and implementation.
- AI should accelerate standards-led development, not replace standards, governance or accountability. It could provide an "intelligent middleware" layer across data transformation, API development and assurance, while regulated firms remain accountable for security, resilience, data quality and customer outcomes.
- Agentic AI raises the stakes. If AI systems become more capable of acting on customers' behalf, Open Banking and Smart Data infrastructure will need robust controls for consent, permissions, auditability, liability and revocation.
- Trust and resilience will be critical. Inaccurate outputs, inconsistent performance, supplier concentration and cyber threats require safeguards including validation, monitoring, audit trails, fallbacks and appropriate human oversight.
- The potential benefit goes beyond efficiency. AI-ready Open Banking infrastructure could accelerate the development of new products and services while strengthening the foundations for the UK's wider Open Finance and Smart Data ambitions.

[1] https://www.scnsoft.com/finance/large-language-models

Foreword



Open Banking has shown what can happen when an industry comes together around common standards, trusted infrastructure and a shared ambition to deliver better outcomes for customers. The UK has built something of real value, and our responsibility now is to protect that foundation while helping the ecosystem prepare for what comes next.

At Open Banking Limited, our role is to maintain and support the standards that underpin the ecosystem. In doing so, we also need to understand how technology, markets and customer expectations are changing, so that our work remains relevant, robust and useful to those who rely on it.

AI is an important part of that changing landscape. Not because technology is an end in itself, but because it has the potential to change how quickly we innovate, how efficiently we develop and maintain our infrastructure, and ultimately how rapidly better products and services can reach people and businesses.

That matters as we look towards the next chapter of Open Banking, Open Finance and the UK's wider Smart Data ambitions. If the future is open, the standards that underpin it need to be ready for that future too. They must continue to support trust and security, while enabling innovation and new ways for technology to interact with financial services.

For me, that is the opportunity explored in this report: how we can make the most of AI to reduce complexity, improve the development experience and support more consistent implementation across the ecosystem. The aim is not AI for its own sake, but to make the infrastructure beneath Open Banking work better and help unlock the next wave of innovation.

We should approach that opportunity with ambition, but also with care. AI does not lessen the importance of high-quality data, security, accountability or trust. It makes them more important.

I hope this report encourages the industry to think boldly about what is possible, while remaining focused on what has always mattered most: building an open, trusted and innovative ecosystem that delivers better outcomes for people and businesses.

Henk Van Hulle

CEO, Open Banking Limited

Author's Preface

There are almost daily stories of banks and fintechs finding new uses for Large Language Models (LLMs). But innovation within a firm is one thing; changing the infrastructure an industry shares is another. Collective schemes require agreement and coordinated implementation; progress often depends as much on legislation and regulation as on technology.

Open Banking illustrates that challenge. It took a regulatory mandate and years of bespoke development to launch and reach scale. But must its next phase – and the development of Open Finance – follow the same path? Or have the tools changed enough to revisit our assumptions about time, cost and effort?

That is the question behind this paper. LLMs can help write code, interpret technical requirements and navigate documentation. 'Vibe coding' makes the shift tangible: people with little developer knowledge can describe an application and use an LLM to help build it. We explore what these capabilities could mean for updating Open Banking standards and creating new standards for Open Finance.

It is worth noting however, that AI may reduce the technical effort required to implement collective change, but it does not remove the need for agreement on what changes are needed, or how implementation is coordinated.

Two cautions before you read on. The efficiency gains discussed here are planning assumptions to validate, not measured savings; LLM use also carries costs, with prices changing frequently. And this paper deliberately looks beneath the customer experience: at standards, coding, APIs and supporting infrastructure. Consumer-facing agentic AI is a separate workstream, considered here only where today's infrastructure needs to anticipate it.

What follows brings together feedback on our AI Green Paper and our analysis. I hope you read it as an invitation to question whether the next phase of open data delivery must be built like the first. It could be faster and cheaper – provided we make the rules and standards AI-ready.

Nick Davey

Head of Strategy, Open Banking Limited

Executive Summary

Open Banking took a regulatory mandate on the nine largest banks and building societies in Great Britain and Northern Ireland and years of bespoke development. Must its next phase – and the move into Financial Smart Data through Open Finance – follow the same path? The consultation evidence challenges that assumption. Open Banking Limited's (OBL) AI Green Paper, published in Q4 2025, asked whether LLMs could speed up API delivery by converting legacy and unstructured data into structured, API-accessible formats, and help third parties bring new products to market faster. This paper brings together the responses and our analysis.

The context has moved on, too. HM Treasury's Financial Services AI Adoption Plan (July 2026) prioritises scaled, responsible AI adoption. The FCA does not currently plan AI-specific rules, instead it expects firms to manage AI through existing outcomes-based frameworks. Financial institutions have moved from experimentation with LLMs towards production deployment. The Evident Global AI Index assesses their maturity across Talent, Innovation, Leadership and Transparency, finding that many banks are embracing AI.

What respondents told us

This paper goes beyond a summary of consultation responses, bringing together those responses, subsequent stakeholder engagement, external research and OBL's standards-development work to inform its analysis and conclusions.

We had:

- Six formal responses from a mixture of industry respondents;
- Significant informal responses from interaction with a much larger set of firms and trade associations through meetings and calls;
- Research based on existing studies and use cases.

There is consensus that AI can act as an 'intelligent middleware' layer that maps, cleans and standardises data before it is exposed through APIs, especially in legacy and multi-technology environments. Proposed applications include automated schema mapping, anomaly detection, lineage tracking, synthetic test data and API documentation.

That support came with conditions. Respondents were clear that data quality, documentation, governance and explainability are the key enablers, and that assurance is required wherever AI has been involved in preparing data. Hallucination risk and the difficulty of reproducing AI outputs remain live concerns, and they shape a legitimate choice each firm will make between AI automation and AI augmentation of existing processes.

Respondents see the Open Banking Standard as a strong base, but not yet optimised for AI consumption. They called for a single, authoritative, version-controlled source for all specifications, richer machine-

readable artefacts and clearer governance of AI use. OBL has begun to improve the accessibility and machine readability of the standards. We now treat AI as one of the main customers of the Open Banking Standards.

What we conclude

AI is likely to reduce the cost and complexity of API data provision. No respondent provided substantiated, generalisable estimates of savings. Empirical studies of customer support and software development demonstrate productivity gains of roughly 10–30%. For financial services more generally employee productivity is considered to have increased by around 20-30%, with back-office efficiencies closer to 40% in some deployments. It is not just about efficiencies, however, the approach to use of LLMs has also changed, with Vibe Coding – using LLMs to code functional software based on descriptions, as well as using AIs in a more structured manner to assist development. We have seen banks move to a build rather than buy model because the time and effort required to build in-house has dramatically reduced; ING openly announced it is using Vibe Coding to build core electronic trading infrastructure and the Boston Consulting Group published a report in June 2026 highlighting that Vibe Coding is coming to finance and that CFOs need guardrails.

Applied to Open Banking activities, such as data preparation, mapping, documentation, and quality assurance, we use the information about efficiencies and the change in many banks' strategies as a planning assumption rather than a measured result. We consider up to 30% efficiency gain is reasonable for well-governed, repeatable activities, subject to validation, with higher gains possible. Gains vary with the task, the people, the data and the governance required, and are partly offset by the cost of the AI tooling itself.

Data quality is the biggest driver of AI value – and of AI risk. As AI is layered on top of Open Banking to deliver products and services, having the right data available, with the right permissions, at the right level of quality becomes essential.

The risk profile changes as autonomy increases. Assistive AI supports people without taking decisions; advisory AI shapes personalised guidance; agentic AI plans and acts within delegated authority – and the step from explaining data to acting on it is the material one. Consumer-facing agentic AI is the subject of a separate workstream. This paper starts from the assumption that more and better quality automated data will be required to build the next phase of AI and that like with Open Banking today transparency and consent are the foundations of trust – consent needs to be revocable and auditable, with clear liability allocation, and resilience against model monoculture and vendor concentration.

The Green Paper and a Summary of Responses

Green Paper and Scope

The Green Paper, published in Q4 2025, focused on the consumption of established Open Banking standards by AI, primarily large language models (LLMs), and on the application of AI to convert unstructured or semi-structured data into structured, API-accessible formats. This could significantly reduce the cost and time required to deliver Open Banking and Open Finance compared with earlier delivery models, such as those used during the initial implementation by ASPSPs and TPPs.

AI has continued to develop rapidly since publication, with the market moving from experimentation towards production deployment of AI-enabled workflows in financial services. The most material development for Open Banking is the increased focus on AI models that can plan, call tools, orchestrate multi-step workflows and deliver data through APIs – or ultimately access those APIs with customer consent to provide enhanced products and services. Respondents also raised agentic AI. Consumer-facing agentic services are the subject of separate workstreams and are not a focus of this paper; agentic developments appear here only where the supply side – standards, consent infrastructure and assurance – must anticipate them.

Responses

We analysed inputs from our online survey, launched alongside the Green Paper, together with direct written submissions, email exchanges and conversations following publication. These inputs inform the analysis alongside external research and OBL's standards-development work. We thank everyone who contributed.

Most written responses either gave a more general view covering AI in Open Banking broadly or answered some of the questions posed but not all. The vast majority of respondents were from the Open Banking ecosystem or the AI industry. The survey was anonymous, so comments are not attributed to named respondents.

Q1 — Can AI reduce the cost and complexity of API data provision?

Feedback:

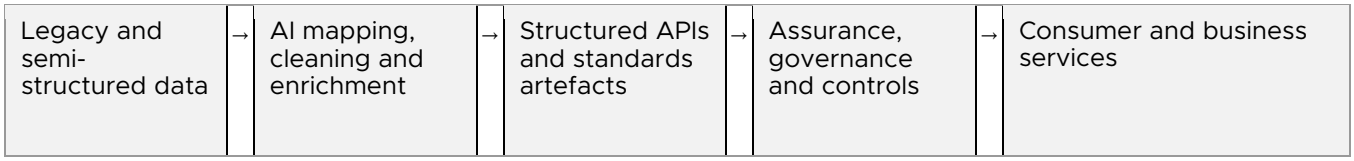
- Respondents gave a strong consensus that AI can act as an 'intelligent middleware' layer to map, clean and standardise data before it is exposed via APIs, especially in legacy or multi-technology environments.
- Proposed applications include automated schema mapping, anomaly detection, lineage tracking, synthetic test data and API documentation.
- Respondents cautioned against over-relying on AI for basic tasks such as simple data standardisation, suggesting instead that the LLM starts from a strong, high-quality foundation, with AI adding value on top.
- At least one respondent noted that issues remain with some LLMs, including hallucinations and shortened context windows, and flagged that assurance is required for any data where AI has been involved in its delivery.
- Another respondent noted that heavy delegation to LLMs can be risky because AI decisions, while recorded, are not always replicable – a consequence of continual model training and new version releases. For data automation, a replicable result is ideally required every time.

Since the consultation closed, debate has focused on two choices: commercial LLMs with enterprise controls – private deployments, restricted access and data isolation – or open-source models deployed within firms' own estates. Each involves trade-offs in accuracy, data protection and cost. Providers also differ in investment, development and model strengths, raising questions about reliance on a single supplier.

The newest model is not necessarily the best fit either. Some commentators suggest that older models can meet many financial-services needs at relatively low cost. That does not settle the separate question of whether to use one provider or several: the choice should reflect the task, the controls required and the risks of supplier reliance.

AI can also augment teams delivering data through APIs and TPPs building services on top, whether through direct mapping or by helping develop rules-based approaches. Figure 1 shows a full-automation route from legacy database to end services.

Figure 1: Where AI can create value in Open Banking delivery



AI acts as an accelerator across data transformation, tooling and assurance — not as a replacement for standards design, security or accountability.

However, some or all of this could be done using conventional development approaches, either augmented by AI throughout or using AI on specific components. The rationale for using AI less extensively is accuracy: hallucination rates vary materially by model and by task complexity. They are highest for complex reasoning and factual recall that is not grounded in source documents, and considerably lower for routine, well-bounded tasks. Therefore, basing outputs in verified and good quality source data, and validating the results remains essential wherever AI is involved in data delivery.

Accuracy limitations, together with the fact that accountability for outcomes remains with the ASPSP or TPP, will shape each institution's preference between AI automation and AI augmentation. These are legitimate preferences, and the right balance will differ by institution and by task.

Research box: Emerging tools to convert legacy and unstructured data into API-ready datasets

Since the publication of the Green Paper, a more developed market has emerged around AI-enabled data transformation. The common proposition of these platforms is that they can ingest data from legacy databases – and from documents, PDFs, emails, spreadsheets, operational systems and fragmented data silos – then extract, classify, normalise, enrich and consolidate that information into more structured, governed datasets.

The efficiency case is strongest where these tools replace manual discovery, mapping, document parsing, data-quality rule writing and reconciliation – repeatable work with accessible inputs and verifiable outputs. The scale of the gains, and the quality of the evidence behind them, is assessed under 'Savings: time and cost' below.

For Open Banking and Open Finance, this supports the proposition that AI can reduce the burden of preparing data for API-based access: these tools could help firms move from fragmented legacy estates to consolidated datasets that are sufficiently structured, validated and documented to be exposed through APIs that authorised TPPs can access. Their existence does not, however, remove the need for standards. AI can accelerate delivery, but the target data model, definitions, quality thresholds, lineage, permissions, assurance process and liability framework still need to be specified and governed.

Savings: time and cost

No respondent provided substantiated, generalisable estimates of time or cost savings (one vendor submission cited an acceleration figure for its own service; see example snapshots below). The wider evidence base has become more nuanced since the original Green Paper: recent empirical studies support material productivity gains in specific, well-defined tasks, but they do not support a simple assumption that AI will automatically reduce firm-level costs in all cases.

In customer support, a study of 5,172 customer agents, published in the Quarterly Journal of Economics in 2025, found that access to a generative AI assistant increased productivity by around 15% on average, with the largest gains for newer and less-experienced workers[2].

In software development, field experiments across Microsoft, Accenture and an anonymous Fortune 100 company, published in Management Science in 2026, found that developers with access to an AI coding

[2]Brynjolfsson, E., Li, D. and Raymond, L. (2025), *Generative AI at Work*, The Quarterly Journal of Economics, 140(2), pp. 889–942. https://academic.oup.com/qje/article/140/2/889/7990658

assistant completed around 26% more tasks, with stronger benefits for less experienced developers[3]. Earlier controlled coding experiments have identified larger time savings on specific tasks.

These studies demonstrate the potential for task-level gains of roughly 10–30%. They measure customer support and software development, not Open Banking delivery, so applying that range to activities such as data preparation, schema mapping, reconciliation and API assurance is an extrapolation and something to be tested over time. Up to 30% could therefore be seen as a reasonable planning assumption for well-governed, repeatable activities, subject to validation against firms' own experience, with higher gains plausible in narrower, structured tasks such as code drafting, testing, documentation and data mapping.

Some examples show higher efficiency gains of up to 40% on routine back office tasks, so there is potential for some elements of Open Banking to be subject to higher gains if these approximate routine back office tasks. It is also the case that data processing is much faster and more efficient under Augmented AI models, with estimates showing that data processing can be between 10 and 50 times faster. More complex data such as analysing multi page financial information is at the slower end, more automated data extraction exercises can be at the higher end.

For Open Banking and Open Finance API delivery, AI is therefore best viewed as an efficiency and assurance accelerator. It reduces the cost and effort of schema mapping, code generation, conformance tooling, documentation, anomaly detection and remediation. Regulated firms will still need to maintain accountability for data quality, security, operational resilience, and compliance functions. The gains are also partly offset by the costs of deploying and running AI itself; we have not assessed LLM pricing, which sits in a competitive market. We have received anecdotal input however that the models that financial institutions need in order to get these efficiency gains are not the latest frontier models; which puts downward pressures on the cost of development and delivery. Some also perceive that the cost of delivering AI LLMs will fall significantly in the future as fixed upfront costs are reduced[4], although this isn't a universally shared view.

Figure 2: Turning AI productivity into firm-level value

Task-level gains potential 10–30% improvement in well-governed, repeatable activities	→	Workflow redesign simpler processes, clearer roles and integrated tools	→	Operational change training, assurance, controls and benefit tracking	→	Realised value lower cost, faster delivery, higher quality or better outcomes

Productivity becomes value when it delivers measurable business or customer outcomes.

[3]Cui, Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S. and Salz, T. (2026), *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*, Management Science. https://pubsonline.informs.org/doi/10.1287/mnsc.2025.00535

[4] See: https://www.businessinsider.com/sam-altman-cost-using-ai-drop-10-times-every-year-2025-2

Q2 — Does the Open Banking Standard give AI what it needs? What improvements could make it more AI-ready?

Feedback:

- Respondents suggested the current standard is a strong base but not fully optimised for AI ingestion. Pain points include version drift across public sources, optional and variant fields, and limited machine-readable artefacts for workflows.
- Proposed enhancements: a single authoritative, version-controlled repository for all historic and current specifications; richer metadata and semantic precision; machine-readable supplements (e.g. knowledge graphs); and clearer governance on how AI may consume and act on the standards.
- Documentation uplift is essential for both humans and AI: more examples, consistent patterns and explicit field semantics.

Separately from the consultation responses, technical standards for AI-tool interaction have matured since the Green Paper. The Model Context Protocol (MCP) has become more relevant to agentic workflows, with the July 2026 specification[5]. This reinforces the need for Open Banking standards to be consumable not only by humans and LLMs, but potentially also by governed agents operating through structured, auditable tools.

OBL's standards-accessibility work is broader than AI alone: we are updating and centralising the standards on a new platform that includes adding the capability for MCP servers to connect to the platform, as well as having an AI agent (AskAI) tool embedded in the human-readable presentation layer. The centralisation of the standards has also been powered by AI; ensuring migration of information is both complete and accurate, and this is also delivering improvements in terms of consistency of terminology and structure. This includes capability that serves the API specification documents and API over MCP so AI clients can connect. We now treat AI as one of the main customers of the Open Banking standards. AI will need to understand and build to the standards to realise the efficiency gains highlighted, so where there are known deficiencies, we will work to remove the ambiguities that AI systems would otherwise inherit. We can also use AI itself to identify inconsistencies and possible gaps within the standards.

An LLM inherits the quality of the specification it consumes; ambiguous fields produce ambiguous code, version drift across sources produces inconsistent implementations, and a specification scattered across documents reproduces exactly the fragmented delivery the original implementation suffered. A single authoritative, version-controlled, machine-readable source inverts that. It enables AI tools to generate code, tests and documentation against the same ground truth as human implementers and allows conformance to be checked deterministically wherever requirements can be expressed as testable rules – with outputs still subject to validation. When the standard changes, regeneration against the updated source could reduce manual rework and make change management more efficient, with implementation

[6]FCA (2026), *The Mills Review: AI and the future of retail financial services*, July 2026. https://www.fca.org.uk/publication/corporate/the-mills-review.pdf

still flowing through the approved, tested and version-controlled release processes described under Question 3. Prioritising AI access to documentation, with clear signposts to the latest version will reduce the occurrences of out-of-date and stale information being consumed by AI and presented back to a user. This is not documentation housekeeping; it is the enabling condition for the efficiency case, and among the most promising investments in making the next phase of delivery faster and more cost-efficient than the first.

Q3 - Can AI help build entirely new customer and business products from scratch?

Feedback:

- Evidence covered a range of product areas, including budgeting and affordability analytics, fraud prevention and personalised assistants, alongside growing evidence of AI being used to create code and therefore software products.
- Respondents noted that data quality and standards maturity are prerequisites for safe automation.
- New opportunities include agentic workflows, proactive financial coaching, AI-assisted co-creation environments to design, test and certify new products, and agent-enabled payments or commerce journeys – pursued only within defined guardrails, with clear user intent, scoped permissions, escalation points and auditability.
- Identity and consent emerged as enablers: a unified digital identity and contextual, reusable consent could meaningfully reduce user friction and enable multi-source, AI-driven services.

While AI is not cost-free, there is a general consensus that the time and cost of delivering software-based products and services have fallen, enabling entrepreneurs to take ideas through to product delivery with smaller teams and in shorter timeframes. Partly offsetting this, organisations delivering financial services to consumers and businesses are regulated: they require authorisation, must submit regulatory returns, and carry other operational overheads. Some of these fixed costs AI will not materially affect.

Safe automation depends on high data quality and mature standards. This is addressed in part under Question 2 and through our work on the standards, but it also requires mechanisms by which changes to standards and maintenance releases can be automatically identified and interpreted by AI – with implementation into live APIs, products and services still flowing through approved, tested and version-controlled release processes. Sufficient maturity across that full lifecycle is what allows automation to operate safely and reliably.

AI-ready standards should be designed for use rather than just passive document ingestion alone. In practice, this means considering how Open Banking artefacts could support agent-accessible tools, explicit permissions, deterministic validation, structured error handling, test harnesses, benchmark datasets and auditable change monitoring. Standards such as MCP and Arazzo are relevant because they describe workflows, tools and interactions in formats agents can consume – but any adoption should be governed by Open Banking's existing requirements for security, consent, resilience and accountability.

Open Banking remains one of the few sources of authenticated data of its kind. Until other smart data

schemes provide comparable, trusted datasets, their availability is likely to remain an important constraint on the range of services that can be built – by AI tools or anything else.

Q4 - What have we missed?

Feedback:

- Agentic AI: delegated consent, entitlement and liability, and resilience concerns (model monoculture, vendor concentration, cyber contagion).
- Environmental and infrastructure considerations (compute, energy and water usage), and the need for model diversity and evaluation standards.

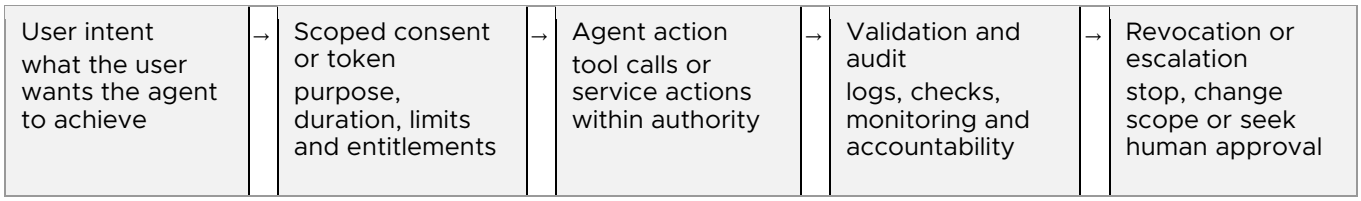
On bias, two distinct risk mechanisms matter, and their relevance varies by application. Upstream, where AI is used to classify, extract, enrich or map information before it becomes API-accessible – the use this paper examines under Question 1 – errors or systematic bias can be introduced into the transformations themselves. These risks are best controlled through the disciplines already discussed, including deterministic validation, reconciliation to source, lineage and data-quality thresholds. Downstream, where AI interprets structured financial data for consumers – generating explanations, recommendations or representations – bias arises in how information is framed, prioritised and presented, and is harder to constrain through rule-based checks alone. What has not yet been systematically tested is how model bias manifests in consumer-facing services built on Open Banking data; that is the specific gap requiring further investigation.

Model diversity and robust evaluation are relevant here. Diversity reduces the shared weaknesses that can arise from reliance on a narrow set of models, particularly models optimised for general purpose rather than financial services contexts; and models with similar headline accuracy scores may differ materially in explainability, consistency and handling of uncertainty. There is currently limited evidence of bias affecting Open Banking data flows or API processes directly, but these considerations become more relevant as AI is layered on top of Open Banking to deliver consumer-facing products and services.

Environmental and infrastructure impacts – compute, energy and water usage – are well documented. They sit outside OBL's control and economic influence, though we note industry initiatives such as Green AI that focus on more energy-efficient models and renewable energy, and we support efforts to minimise the negative environmental impacts of AI.

Finally, on agentic AI – raised by several respondents, and the subject of a separate OBL workstream on the user-facing side – the supply-side considerations are delegated consent, entitlements, liability, runtime supervision and operational resilience. Where agents act on behalf of users, consent should not be treated as a static, one-off permission. It should operate as an executable control surface that defines what an agent may do, for what purpose, against which account or service, over what period, subject to which limits, and with what revocation or escalation rights. This points towards policy-based authorisation, explicit delegation, scoped entitlements, and complete audit trails for agent activity. It also raises questions around liability, whether existing legislation is sufficient at assigning where responsibility sits if an agent acts incorrectly, causes consumer detriment, or relies on an inaccurate interpretation of data.

Figure 3: Agentic consent and control model



Consent becomes an active control layer defining what an agent may do, when it must stop, and how responsibility is evidenced.

There are also resilience considerations. Increased reliance on a small number of models, cloud platforms or AI infrastructure providers could lead to model monoculture, vendor concentration and correlated failure. Highly interconnected agentic systems may also increase cyber-contagion risk, where errors, prompt attacks, compromised tools or incorrect instructions propagate rapidly across services. These risks are not unique to Open Banking, but they become more significant as agentic AI is layered onto Open Banking infrastructure. A proportionate framework should therefore include fallbacks, kill switches, deterministic validation, runtime monitoring, incident reporting, model and tool version control, and clear boundaries for when human approval is required.

Thematic Views and Respondent Snapshots

The snapshots and bullets below summarise our representative's viewpoints. Quotes are anonymised where appropriate.

Affordability and risk in regulated sectors

"Open Banking transaction data varies widely in practice… AI can baseline semi-unstructured inputs and enrich them (e.g., infer MCCs), improving consistency and usability. Use cases include safer-gambling affordability assessments and proactive risk indicators."

- AI is effective as a last-mile standardiser and enrichment layer but should not replace well-maintained APIs.
- Support for enabling access to verified identity attributes (with consent) to improve onboarding and fraud controls.
- Emphasis on transparency and consumer awareness when sharing and processing data.

Standards and agentic AI readiness

A respondent told us that the standards should provide a single, authoritative source for LLM ingestion with version control; add benchmark datasets. We should plan for agentic AI (delegated consent, auditability, guardrails) to unlock future autonomous services within compliance boundaries.

- Single source of truth for specifications to reduce version drift across AI models.
- Machine-readable artefacts (knowledge graphs, event schemas) to support safe AI tooling.
- Explore delegated consent and VRP-aligned 'intent mandates' for agentic commerce.

Implementation variance and privacy automation

A respondent told us that standardised contracts and API gateways reduce cost; AI-powered mapping, automated testing and monitoring cut manual effort. Combine this with automated pseudonymisation and masking to meet GDPR/AI-Act obligations.

- Treat AI features as regulated products: end-to-end encryption, logging, auditability.
- Human-in-the-loop for borderline cases; model explainability and bias testing as defaults.
- Federated learning noted as a promising technique for privacy-preserving model improvement.

Layered AI services

An AI industry respondent told us that differential AI delivers Blueprint → Build: ingests legacy estates, generates requirements, diagrams, and tests; then orchestrates build with human-in-the-loop for verification – typically achieving ~80% acceleration.

TPP perspectives

TPPs noted that we use AI for insights today, but accuracy matters. Model updates can degrade reliability; firms should retain control over which AI to use and how it's governed. Hallucinations risk undermining consumer trust if outputs diverge from bank-facing channels.

Cross-cutting Insights

Four conditions run through everything respondents told us, and together they define what 'AI-ready' means in practice:

- Quality – upstream data consistency determines downstream value; remediation is typically costlier than preparation.
- Legibility – standards and documentation must be as consumable by machines as by people, through machine-readable workflows such as Arazzo and MCP.
- Permissioned authority – consent, identity, provenance and policy-based authorisation are what let AI act, not merely read, on a user's behalf.
- Resilience – model diversity, fallbacks, kill switches, version control and runtime monitoring are the price of depending on AI in regulated infrastructure.

Implications for Open Banking Ecosystem Participants

AI will affect different parts of the Open Banking ecosystem in different ways. The common theme is that AI may reduce the cost of data preparation, standards interpretation, testing and service development, but it will also increase the importance of governance, assurance, liability allocation and trust. Participants will therefore look for practical clarity on what AI-readiness means for their own obligations, opportunities and risks.

Participant	Potential opportunity	Key issues to manage
ASPSPs	Use AI to map, clean, reconcile and expose legacy or siloed data through more consistent API-accessible datasets; reduce manual work in documentation, testing and conformance.	Data quality and ability for reconciliation, schema validation, operational resilience, accountability for transformed data, supplier oversight and assurance evidence.
TPPs / Fintechs	Access richer, more consistent datasets; develop AI-enabled propositions more quickly; use AI for analytics, testing, customer support and personalised experiences.	Reliance on upstream data quality, explainability of customer-facing outputs, consent scope, liability for onward use, model monitoring and customer trust.
Consumers and SMEs	More personalised financial tools, faster onboarding, improved affordability insights, proactive support and eventually agent-enabled services.	Transparency, meaningful consent, ability to revoke permissions, clear redress, understandable explanations and protection from harmful automation or fraud.
Regulators	Use existing outcomes-based frameworks to support responsible AI adoption while encouraging innovation and competition.	Regulatory perimeter, Consumer Duty, accountability, operational resilience, systemic concentration risk, model governance and agentic payment risks.
Central Standards Body	Coordinate AI-ready standards and associated tools and ecosystem guidance.	Which standards body and regulatory remit will undertake work; then determining the appropriate changes in relation to certification, assurance, incident reporting, maintenance, delegated consent profiles and agent-accessible tooling.

The practical conclusion is that AI-readiness should not be framed as only a technical standards issue. It is an ecosystem operating-model issue: how data is prepared, tested, certified, accessed, explained, monitored and governed across the full chain from ASPSP source systems to TPP products and customer outcomes.

A Trust and Certification Model for AI-enabled Open Banking

If AI is used to prepare, transform, enrich or expose data for Open Banking and Open Finance, ecosystem trust will depend on more than the technical capability of the model or tool. Participants will need confidence that AI-transformed datasets, AI-readable standards artefacts and agent-accessible tools are accurate, repeatable, explainable, secure and governed. It is also likely that some flag would also need to be used to allow users to know the data has been transformed using an LLM model. This and a proportionate trust and certification model could therefore become an important part of the future operating model, covering three layers. The third is deliberately anticipatory: it prepares supply-side infrastructure for agent access ahead of the separate consumer-facing agentic workstream.

Certification layer	What could be tested	Why it matters
AI-ready standards artefacts	Version control, semantic precision, examples, workflow descriptions, endpoint metadata, test cases and consistency of interpretation.	Reduces ambiguity, implementation variance and model drift when AI systems consume standards materials.
AI-transformed datasets	Schema conformity, reconciliation to source systems, confidence scoring, exception handling, lineage, audit trail, completeness, timeliness and data-quality thresholds.	Gives TPPs, ASPSPs and regulators confidence that AI-assisted preparation has not introduced errors, hallucinations or undocumented transformations. This would also include the relevant flag to the dataset.
Agent-accessible tools	Consent scope, entitlement checks, deterministic validation, logging, tool permissions, refusal behaviour, escalation points, revocation and kill switches.	Ensures agents act only within delegated authority and that harmful or unauthorised actions can be detected, stopped and evidenced.

Certification should be proportionate and risk-based. For low-risk uses, like standards search, documentation assistance or developer support, lightweight checks and clear disclaimers may suffice. For higher-risk uses it is likely that stronger assurance would be needed e.g. providing data through APIs for affordability analytics, or in relation to anti-fraud controls.

The Future Entity could play a convening role in defining this trust model without becoming the sole certifier of every AI product or dataset: setting minimum assurance expectations, maintaining AI-ready standards artefacts, defining benchmark tests, coordinating ecosystem guidance and clarifying how certification interacts with existing Open Banking conformance, operational guidelines, consent requirements and regulatory obligations.

Alignment with UK Policy Direction (Mills Review and Subsequent Work)

Policy context box: The Mills Review and the role of Open Banking

The FCA's Mills Review, published in July 2026, examines how AI could reshape retail financial services by 2030 and beyond. Its central finding is that retail financial services are moving from human-led, episodic interactions towards AI-enabled, continuous and delegated services. The Review emphasises that AI could improve outcomes and support growth by reducing friction, addressing advice gaps, increasing switching and improving consumer support – but warns that it could also amplify risks around fraud, cyber security, market concentration, consumer harm, accountability and loss of consumer control.

The Review identifies agentic finance as a particularly important development and recommends, among other actions, that the FCA enable its foundations. Significantly for this paper, the route it recommends is to develop trusted agent standards within Open Finance – which is standards work, and sits on the supply side this paper addresses.

Open Banking is directly relevant to those foundations: the model for trusted data access, with consent, secure APIs, ecosystem standards, conformance and regulated participation. The next phase may be to test the application of AI to help provide Open Banking data; and may help newer participants or those that have developed Open Banking capabilities in the past to adopt more recent standards; and to make more information available to their customers. Beyond that there is a possible future that would see Open Banking infrastructure AI-innovation ready; and the products and services that ASPSPs and TPPs may provide could also become AI-innovative.

The themes above align with OBL's position submitted to the FCA's Mills Review[6] and with the subsequent UK policy direction on AI in financial services.

Near-term opportunities lie in:

- Efficiency,
- Standards accessibility,
- Data-quality tooling,
- AI-assisted development.

Mid-term priorities are:

- AI-ready standards,
- Evaluation datasets,
- Controlled tool access,

[6]FCA (2026), *The Mills Review: AI and the future of retail financial services*, July 2026. https://www.fca.org.uk/publication/corporate/the-mills-review.pdf

- Clearer governance across the AI lifecycle;

longer-term opportunities lie in agentic experiences, within clear consumer protections, scoped consent and existing regulatory obligations.

Final Takeaways.

- AI can change the delivery model. Legacy-data preparation, API delivery and product development could become faster and cheaper, supporting Open Banking updates and Open Finance.
- Savings seem real but the potential extent is a hypothesis to test. Up to 30% is a planning assumption for repeatable tasks, not a measured Open Banking saving; implementation, assurance and running costs affect net benefits.
- AI-ready standards unlock the opportunity. OBL's work towards a single, authoritative, version-controlled, machine-readable source can reduce ambiguity and enable more consistent code, testing and documentation.
- Quality and accountability remain essential. Permissioned data, source validation, traceability and proportionate assurance underpin trust. Automation or augmentation should reflect each task's risks; regulated firms remain accountable.
- Prepare for autonomy without losing control. Agent access requires scoped, revocable consent, clear liability, auditability and resilient fallbacks. OBL is looking to coordinate with industry on next steps.
- The evidence gives us reason to revisit assumptions about the technical effort required for the next phase of Open Banking and Open Finance. But it does not yet establish the scale of efficiency gains, or remove the need for collective decisions prior to technical change and implementation.

Glossary

Term	Meaning in this paper
Agentic AI	AI systems that can plan, call tools, execute multi-step workflows and act within delegated authority, rather than simply providing information or drafting support.
AI automation / AI augmentation	Automation uses AI to perform tasks with limited direct human intervention; augmentation uses AI to support people or conventional processes. Both require controls proportionate to the task and its risks.
AI-ready standards	Standards that are precise, version-controlled, machine-readable, testable and sufficiently documented for both human implementers and AI-enabled tools to interpret consistently.
API	Application Programming Interface: a defined technical interface that enables secure data exchange or service interaction between systems.
Arazzo	An OpenAPI Initiative specification for describing sequences of API calls and their dependencies to achieve defined outcomes, in a format readable by both people and machines.
ASPSP	Account Servicing Payment Service Provider: typically, a bank or account provider that holds payment accounts and exposes account or payment services through Open Banking APIs.
Assurance / certification	Assurance gathers and evaluates evidence to provide confidence that defined requirements are met. Certification is formal confirmation under a specified scheme that requirements have been met; it is one possible assurance mechanism, not a guarantee of error-free performance.
Audit trail	A reliable record of actions, decisions, data transformations, model outputs and approvals that can be reviewed for assurance, accountability and incident investigation.
Benchmark dataset	A curated dataset used to test whether AI systems or implementation tools perform consistently and accurately against defined tasks or expected outputs.
Conformance	The process of testing whether an implementation meets the relevant Open Banking standards, rules, technical requirements or assurance expectations.
Consumer Duty	The FCA framework requiring firms to act to deliver good outcomes for retail customers, including in relation to consumer understanding, support and foreseeable harm.
Delegated consent	A consent model in which a user authorises another service or agent to act on their behalf within defined purposes, limits, duration and revocation rights.
Deterministic validation	Rule-based checks that produce predictable results, used to confirm that data, tool calls or agent actions meet defined requirements before they are accepted or executed.
FCA	Financial Conduct Authority: the UK regulator for conduct in financial services and certain regulated markets.
Future Entity	The proposed future organisation or operating model intended to take forward Open Banking functions, standards and ecosystem coordination beyond current arrangements.
Hallucination	An AI output that is inaccurate, unsupported or fabricated, even though it may appear plausible or confident.
Human-in-the-loop	A control model in which a person reviews, approves, overrides or escalates AI outputs or actions, particularly where risks are higher or customer impact is material.

Term	Meaning in this paper
Knowledge graph	A structured representation of entities, relationships and meanings that can help AI systems interpret standards, datasets and workflows more consistently.
Large Language Model (LLM)	An AI model trained on large volumes of text and other data that can generate, summarise, classify, interpret and transform language-based information.
Lineage	Evidence showing where data came from, how it was transformed and how it moved through systems or processes.
MCC	Merchant Category Code: a four-digit code identifying a merchant's type of business. An AI-inferred category should be distinguished from a verified code supplied in source data.
MCP	Model Context Protocol: an emerging protocol for connecting AI systems to tools, data sources and services in a structured and governed way.
Model drift	A degradation or change in AI model performance over time, often caused by changes in data, model versions, user behaviour or the operating environment.
Model monoculture / vendor concentration	Model monoculture is widespread reliance on the same or similar AI models, potentially creating shared weaknesses and correlated failures. Vendor concentration is dependence on a small number of suppliers for models, platforms or infrastructure.
OBL	Open Banking Limited, the organisation currently responsible for maintaining and supporting the UK Open Banking standards and ecosystem functions.
Open Finance	The extension of consent-based data sharing beyond Open Banking payment accounts to a wider range of financial products and services, such as savings, investments, pensions and insurance.
Operational resilience	The ability of firms and ecosystem infrastructure to prevent, respond to, recover from and learn from disruptions affecting important services.
Policy-based authorisation	An approach where permissions are governed by explicit policies defining who or what may act, for which purpose, within what limits and under what conditions.
Schema mapping	Matching fields and meanings in a source dataset to a target data structure, specifying any transformations needed to make the data conform to that structure.
Schema validation	Checks that data conforms to a defined structure, field format, value type or technical specification before it is accepted or used.
Smart Data	A policy and technical approach to enabling secure, consent-based data sharing across sectors, potentially extending Open Banking-style models into areas such as energy, telecoms and other markets.
Tokenised consent	A consent approach in which permissions are represented through a portable or reusable token that records scope, purpose, duration, limits and revocation conditions.
TPP	Third Party Provider: an authorised provider that uses Open Banking APIs to deliver account information, payment initiation or related services to customers.
Vibe coding	An approach to software development in which a person describes the intended functionality in natural language and iteratively prompts an AI tool to generate or revise code. Generated software still requires appropriate testing, security checks and review.
VRP	Variable Recurring Payment: a payment model that allows authorised payments to be made within agreed parameters, often discussed in relation to future Open Banking payment use cases.



Nick Davey
OBL Head of Strategy
nick.davey@openbanking.org.uk

OPEN BANKING

The Standard for Open Innovation™

Open Banking Limited (OBL) - the Implementation Entity described in the CMA Order - built the UK's world-leading Open Banking Standard and industry guidelines to drive competition, innovation and transparency in UK retail banking.

There are now over 19 million active user connections - consumers and SMEs -of open banking-powered financial management apps and payment tools in the UK.

www.openbanking.org.uk
www.linkedin.com/company/openbanking