ACADEMY OF LIFE PLANNING

From Agentic AI to Human Agency

Designing Financial Services Around the Sovereign Consumer

How much authority can artificial intelligence help us return to human beings?

Steve Conley, Founder, Academy of Life Planning

24 November 2026

Version 1.0 | Draft for consultation

Public policy and strategic discussion paper

Contents

Executive summary 3

1. The regulatory moment 3

2. The historical problem: asymmetric capability 4

3. From agentic AI to human agency 5

4. Defining the sovereign consumer 6

5. Three competing futures for financial services 6

6. The AI Financial Agency Spectrum 7

7. Human Agency by Design 9

8. Measuring human agency: The Human Agency Test 10

9. Beyond Consumer Duty: agency as an observable outcome 11

10. The independent professional in an AI-enabled world 12

11. Interoperability: the infrastructure of sovereignty 13

12. The sovereign life record 14

13. Liability, accountability and delegated authority 14

14. The right to delegate — and the right to take back control 15

15. Vulnerability, accessibility and the limits of autonomy 16

16. Financial crime prevention and citizen capability 17

17. The Sovereign Consumer Charter 18

18. The commercial architecture of agency 18

19. A practical reference architecture 19

20. Policy recommendations 20

21. A proposed contribution to the FCA Agentic Policy Sprint 21

22. Research and implementation roadmap 22

23. Conclusion: from artificial intelligence to human flourishing 23

Invitation to dialogue 24

References and regulatory context 24

Executive summary

The emergence of agentic artificial intelligence represents a structural turning point for financial services.

For decades, consumers have relied on financial institutions and professional intermediaries to interpret complex information, assess alternatives and navigate financial decisions. This dependency has arisen, in part, from substantial asymmetries of information, expertise, technology and bargaining power.

Artificial intelligence has the potential to reduce those asymmetries.

However, two fundamentally different futures are emerging. In one, AI enables institutions and technology platforms to make increasingly sophisticated decisions on behalf of consumers. Human decision-making is progressively delegated to automated systems. In the other, AI equips individuals with the knowledge, analytical capability and practical tools necessary to exercise greater authority over their own lives.

Both futures involve technological advancement. Only the second necessarily involves the restoration of human agency.

The Academy of Life Planning proposes that the development of AI-enabled financial services should be assessed not only against conventional measures of safety, efficiency, innovation and consumer outcomes, but also against its contribution to human agency.

This paper introduces:

- Human Agency by Design: a proposed principle for financial services innovation.
- The Sovereign Consumer: an individual equipped to understand, choose and act with meaningful independence.
- The AI Financial Agency Spectrum: a framework distinguishing capability-building from delegated decision-making.
- The Human Agency Test: a proposed method for evaluating whether services increase or diminish consumer autonomy.
- The Sovereign Consumer Charter: a proposed set of consumer entitlements for the age of AI.
- A policy and implementation roadmap for regulators, financial institutions, technology providers and professional communities.

The central proposition is simple:

> **Artificial intelligence should expand the capacity of human beings to exercise authority over their lives, rather than merely expand the capacity of organisations and machines to exercise authority on their behalf.**

1. The regulatory moment

On 24 November 2026, the Financial Conduct Authority published an update on its approach to agentic AI. [1]

The FCA recognises the potential benefits of AI agents in widening access to financial information and tools. It also highlights questions concerning regulatory permissions, agent identity, consumer consent, delegated authority and accountability.

Significantly, it asks financial institutions to prepare for interactions with AI agents operating on behalf of their customers.

The FCA has also announced plans for a dedicated Agentic Policy Sprint, bringing together policymakers, technology providers and financial services firms. [1]

This follows the Mills Review into the long-term impact of AI on retail financial services, which identified risks including declining financial literacy, diminished consumer agency and manipulation as financial decisions become increasingly automated. [3–4]

These developments signal an important change. The regulatory conversation is moving beyond how financial institutions use artificial intelligence towards how consumers themselves participate in an AI-enabled financial system.

This creates an opportunity to ask a more fundamental question.

Not simply: How can we make agentic finance safe?

But: What kind of financial system should agentic finance help us create?

2. The historical problem: asymmetric capability

The traditional financial services system has developed around an imbalance of capability.

Financial institutions typically possess greater access to specialist knowledge, analytical infrastructure, product information, legal resources and professional expertise than individual consumers.

Consumers frequently encounter complex products, unfamiliar terminology, lengthy contracts and decisions whose consequences may not become apparent for decades.

Professional intermediation has helped address these challenges and continues to provide valuable expertise. However, dependency on intermediaries can also create structural vulnerabilities, particularly where the intermediary benefits financially from the consumer's eventual choice.

The problem is not necessarily the character or intentions of individual professionals. It is the architecture of the relationship.

When the person explaining the options also benefits from the selection of a particular option, a potential conflict exists. Regulation seeks to manage such conflicts. Alternative service architectures can sometimes remove them altogether.

AI creates the possibility of a different architecture.

Traditional architecture	Agency-led architecture
Institution → Intermediary → Consumer	Sovereign consumer + personal AI ↔ Financial institution

Conceptual illustration of a potential redistribution of capability, not a claim that existing asymmetries have disappeared.

A consumer equipped with appropriate AI tools can increasingly:

- Interpret complex financial documents.

- Explore the consequences of different decisions.
- Examine assumptions underlying professional recommendations.
- Understand contractual rights and obligations.
- Identify questions requiring specialist expertise.
- Challenge inconsistencies and request supporting evidence.
- Maintain an independent record of their financial affairs.

The consumer is no longer necessarily dependent on the institution to explain the institution's own proposition. This represents a potential redistribution of power.

> **The democratisation of intelligence creates an opportunity for the democratisation of financial authority.**

3. From agentic AI to human agency

The terms agentic AI and human agency sound related. They describe different things.

Agentic AI concerns the ability of a machine to pursue objectives, perform tasks and potentially act with varying degrees of autonomy.

Human agency concerns the ability of a person to understand their circumstances, establish objectives, make informed choices and act upon them.

An increase in machine agency does not necessarily produce an increase in human agency. Indeed, it can produce the opposite.

Dimension	Agentic AI	Human agency
Central question	What can the machine do?	What can the person do?
Primary capability	Autonomous task completion	Informed self-direction
Measure of success	Successful execution	Meaningful human control
Decision authority	May be delegated	Ultimately retained by the individual
Principal risk	Unsafe or incorrect execution	Dependency, exclusion or manipulation
Desired relationship	Machine performs tasks	Technology serves human purposes

A financial agent might automatically select products, move money, renew contracts and optimise a portfolio. These activities could create substantial convenience.

But if the individual cannot understand, challenge, amend or revoke the decisions being made, convenience may have been purchased at the expense of sovereignty.

Conversely, an appropriately designed AI system could make complex information accessible, help a person articulate their priorities and reveal alternatives previously unavailable to them.

The distinction is not between good and bad technology. It is between different objectives for technology.

> **The autonomy of the machine should be subordinate to the autonomy of the person.**

4. Defining the sovereign consumer

The Academy proposes the following working definition:

> **A sovereign consumer is an individual who possesses, or can access, the information, capability, resources and support necessary to make meaningful decisions in accordance with their own values, circumstances and objectives, while retaining ultimate authority over those decisions.**

Sovereignty does not mean that individuals must become experts in everything. It does not require them to manage every administrative task themselves. Nor does it prohibit the voluntary use of advisers, specialists or automated agents.

A person exercises agency when they knowingly choose to delegate a task within boundaries they understand and can control.

The distinction is between chosen delegation and imposed dependency.

Financial sovereignty has five dimensions.

Dimension	Meaning
Epistemic sovereignty — the ability to know	Access to understandable information and the ability to distinguish evidence from assertion.
Decisional sovereignty — the ability to choose	Meaningful alternatives and freedom from undue influence.
Operational sovereignty — the ability to act	Practical access to the tools and resources needed to implement decisions.
Data sovereignty — the ability to control information	Meaningful control over access, use, sharing and portability of personal records.
Remedial sovereignty — the ability to challenge	Accessible mechanisms to question decisions, correct errors and pursue appropriate redress.

These five dimensions provide a broader understanding of consumer empowerment than financial literacy alone.

Knowledge without the ability to act is insufficient. Choice without meaningful alternatives is illusory. Rights without accessible remedies are incomplete.

5. Three competing futures for financial services

AI is not inherently decentralising. It can distribute capability or concentrate it.

Three broad models may emerge.

Model	Location of control
Institutional dependency	Financial institution → Consumer
Algorithmic dependency	Consumer → Technology platform → Providers
Human sovereignty	Consumer → Consumer-controlled AI → Specialists, institutions and other tools

Model A: Institutional dependency

Institutions use AI to optimise distribution, servicing, compliance and commercial relationships. Consumers receive more sophisticated services but may remain dependent on institutional infrastructure and judgement.

Risk: existing information and power asymmetries become technologically reinforced.

Model B: Algorithmic dependency

Consumers delegate financial decisions to technology platforms, which become powerful gateways between individuals and financial institutions.

Risk: traditional intermediation is replaced by algorithmic intermediation.

Model C: Human sovereignty

Consumers retain meaningful control over their information, objectives and decisions, using AI and professional expertise as complementary resources.

Objective: continuous human agency, supported by episodic expertise.

The third model does not require the disappearance of financial institutions, professional advisers or large technology companies. It requires a different relationship between them and the individual.

Institutions become resources within the consumer's financial ecosystem rather than necessarily becoming the organising centre of that ecosystem.

> **Replacing a human gatekeeper with an algorithmic gatekeeper is not the democratisation of financial services. It is the digitisation of dependency.**

6. The AI Financial Agency Spectrum

Regulatory and commercial discussions frequently treat AI applications as a single category. A more useful approach is to distinguish systems according to their functions and the authority they exercise.

The Academy proposes an illustrative six-level spectrum.



Levels 1–3: capability building. Levels 4–6: increasing intervention and activity-specific safeguards. Illustrative functional framework, not statutory categories.

Level	Function	Illustrative application
1. Inform	Explain	Describe how a pension works
2. Interpret	Analyse	Explain the terms of a pension contract
3. Explore	Model	Compare alternative retirement scenarios
4. Recommend	Direct towards a specific action	Recommend a particular investment
5. Arrange	Facilitate	Arrange the acquisition of a financial product
6. Execute	Perform	Automatically purchase or sell investments

The distinction has important implications for governance, consumer understanding and regulatory analysis.

An application that explains financial concepts is functionally different from one that selects investments and executes transactions.

Nevertheless, even apparently educational systems can cross regulatory boundaries depending on how their outputs are generated, personalised and presented. No particular level should be treated as a regulatory safe harbour.

The applicable legal and regulatory requirements must follow the actual activity and circumstances, rather than the label attached to the technology.

The Academy's preferred architecture concentrates on building capability through the first three functions. This is a deliberate design choice.

> **The objective is not necessarily to progress towards greater machine autonomy. It is to progress towards greater human autonomy.**

7. Human Agency by Design

The Academy proposes Human Agency by Design as an overarching principle for the development of consumer-facing financial services and technologies.

It would complement established approaches to privacy, safety, accessibility and consumer protection.

Its purpose would be to ensure that the preservation and development of human decision-making capability becomes an explicit design objective rather than an accidental by-product.

Seven principles are proposed.

Principle	Design requirement
Human primacy	Technology serves the person's objectives rather than defining them.
Comprehensibility	People can understand the basis and limitations of material outputs.
Meaningful choice	Alternatives, uncertainties and trade-offs are made visible.
Structural independence	Commercial incentives and potential conflicts are disclosed and, where possible, removed.
Controllability	People can challenge, amend, suspend or revoke delegated authority.
Portability	People can retain and transfer their information without unreasonable barriers.
Accountability	Responsibility for errors, service quality and harmful conduct remains identifiable.

Human Agency by Design would require organisations to consider whether their products leave consumers more or less capable after using them.

This is a different measure of success from engagement, retention, conversion, assets under management or transaction completion. A service might achieve excellent commercial performance while increasing consumer dependency.

An agency-led service should be capable of succeeding commercially while reducing it.

> **The test of an empowering service is whether the person becomes more capable, not whether the provider becomes more indispensable.**

8. Measuring human agency: The Human Agency Test

Consumer empowerment is frequently described as an aspiration but less frequently measured as an outcome.

The Academy proposes a practical assessment framework organised around its five capability outcomes: Clarity, Control, Confidence, Care and Course. These are collectively referred to as the 5Cs of Human Agency.



Five complementary dimensions of consumer capability.

Clarity — can the person understand?

Can the individual explain their circumstances, available options, material risks and uncertainties in their own words? Has the service reduced confusion or simply provided an answer?

Control — can the person choose?

Can the individual reject a proposed course of action, explore alternatives, change providers and retain access to their information? Has the service increased or diminished the person's practical decision authority?

Confidence — can the person act?

Does the individual feel appropriately equipped to take the next step? Confidence must be calibrated against actual understanding. False reassurance and overconfidence should not be mistaken for empowerment.

Care — is the person appropriately supported?

Does the service recognise vulnerability, accessibility requirements, emotional circumstances and the limits of its own competence? Can the person access appropriate human expertise when required?

Course — can the person sustain progress?

Does the individual have an actionable plan, a record of their decisions and a means of adapting to changing circumstances? Can they continue making progress without unnecessary dependency on the original service provider?

Proposed assessment instrument

The following prototype illustrates how these dimensions could be incorporated into service evaluation. Score each dimension from 0 (absent) to 5 (strong).

Dimension and assessment statement	Before	After
Clarity: The consumer understands the decision and its consequences.	2	4
Control: The consumer retains meaningful authority and alternatives.	2	4
Confidence: The consumer can act with appropriately informed confidence.	2	4
Care: The consumer can access appropriate support and safeguards.	2	4
Course: The consumer can sustain progress and adapt independently.	2	4
Total (out of 25)	10	20

Agency change: +10 points. The scores are hypothetical; the instrument is a proposed research tool, not a validated psychometric measure or regulatory compliance assessment.

The framework could be developed through consumer research, longitudinal studies and comparative testing.

Self-reported confidence should be tested against demonstrated understanding. Aggregate scores should not conceal serious deficiencies in individual dimensions. For example, a high clarity score cannot compensate for the inability to revoke an agent's financial authority.

The underlying question remains:

> **What can the consumer understand, decide or do after receiving the service that they could not understand, decide or do before?**

9. Beyond Consumer Duty: agency as an observable outcome

The FCA's Consumer Duty establishes important requirements relating to products and services, price and value, consumer understanding and consumer support. [7–8]

Its emphasis on understanding and support provides a natural foundation for exploring human agency as an additional analytical lens.

The Academy is not necessarily proposing a new regulatory duty at this stage. Rather, it proposes research into whether measures of consumer agency could strengthen the evaluation of existing consumer outcomes.

Existing Consumer Duty outcome	Complementary agency question
Products and services	Can the consumer understand their needs and explore meaningful alternatives?

Existing Consumer Duty outcome	Complementary agency question
Price and value	Can the consumer evaluate what they are paying for and the value received?
Consumer understanding	Can the consumer demonstrate understanding independently of the provider?
Consumer support	Does the support increase capability while remaining available when needed?

Consider an ongoing financial service. A consumer may receive regular communications, annual reviews and access to professional support. These may be valuable.

An additional question is whether the consumer is becoming increasingly capable of understanding and navigating their circumstances.

Could they confidently challenge a recommendation? Could they distinguish the value of financial planning from the value of investment administration? Could they leave the service without losing access to their financial history or an understanding of their financial plan?

The objective is not to eliminate ongoing professional relationships.

It is to distinguish continuity of care from continuity of dependency.

10. The independent professional in an AI-enabled world

The restoration of human agency does not diminish the importance of human expertise. It changes when, why and how that expertise is deployed.

Historically, professionals have often combined knowledge provision, interpretation, judgement, administration and implementation within a single service relationship. AI enables some of these functions to be separated.

Routine explanation, document interpretation and scenario exploration may increasingly be accessible directly to consumers.

Human professionals can concentrate on matters where their expertise, judgement, experience and relational skills are particularly valuable. These include uncertainty, conflicting priorities, complex family circumstances, vulnerability, ethical dilemmas and situations requiring specialist competence.

The Academy describes this emerging architecture as:

Continuous agency with episodic expertise.

Life progression	Career • Family • Retirement • Later life
Continuous foundation	Consumer capability and AI support
Episodic layer	Specialist expertise accessed when circumstances require it

The individual's plan remains continuous. Specialist involvement is determined by need rather than necessarily by the commercial requirement for an ongoing intermediation relationship.

This also supports a wider understanding of financial planning. A person's financial future is shaped by their health, relationships, career, skills, enterprise, housing, family obligations and purpose, alongside their financial assets and liabilities.

The consumer should be able to plan the whole person before considering individual financial products.

Where specialist advice or implementation is needed, the appropriate professional can be engaged.

Self-directed by default. Supported when it matters.

11. Interoperability: the infrastructure of sovereignty

Human sovereignty requires more than the availability of AI applications. It requires a financial ecosystem in which consumer-controlled technologies can communicate effectively with service providers.

The FCA's November 2026 update recognises the emerging need for financial firms to identify agents, establish permissions and provide appropriate support when problems occur. [1]

The Academy proposes that interoperability should extend beyond the safe execution of transactions. It should also support the consumer's right and practical ability to understand their financial circumstances.

Priority areas could include:

Infrastructure	Potential consumer benefit
Machine-readable financial statements	Independent analysis of financial circumstances
Standardised fee disclosures	Transparent comparison of costs
Portable financial plans	Continuity when changing providers
Accessible product and contract information	Improved understanding before commitment
Structured transaction histories	Better record-keeping and reconciliation
Verifiable document provenance	Stronger protection against misinformation
Secure, permissioned data access	Reduced administrative burden
Traceable consent and authority records	Greater accountability and control

Financial institutions should be encouraged to accommodate legitimate, securely authenticated consumer-directed tools.

This does not mean granting unrestricted access to institutional systems or sensitive information. Appropriate authentication, data minimisation, consent, cybersecurity and operational safeguards remain

essential. The ICO has identified several such risks in its work on agentic AI. [10]

A particularly important distinction is between permission to read information and permission to act.

A consumer's AI should not automatically acquire the power to move money simply because it has permission to examine a financial statement. Read access and execution authority should be separately governed.

The future-ready financial institution may be one that can withstand independent scrutiny, communicate transparently and cooperate with technologies chosen by the consumer.

Transparency becomes a technical requirement, not merely a marketing promise.

12. The sovereign life record

One of the structural weaknesses of existing financial services is that the individual's financial history is often fragmented across institutions.

Pension providers, banks, employers, insurers, advisers, solicitors and other professionals each retain different parts of the person's story. The consumer may not possess a coherent record of their own life.

When providers merge, systems change, businesses close or records become inaccessible, individuals can struggle to establish historical facts or exercise longstanding rights.

A sovereign consumer architecture should therefore explore the development of a lifelong, portable, consumer-controlled record.

Such a record could contain financial plans, pension histories, contractual documents, professional recommendations, decision rationales, consent records and correspondence.

The individual would determine what information is shared, with whom, for what purpose and for how long, subject to applicable legal requirements. AI could make these records intelligible, searchable and practically useful.

However, such a record would require strong privacy and security protections, recovery mechanisms, export facilities and safeguards against inappropriate access. It should not create a new form of platform dependency.

The Academy's My Life Record™ proposition represents an initial practical exploration of this concept.

The broader policy opportunity is to establish principles and interoperability standards that allow individuals to retain continuity of their own financial histories across their lifetimes.

The record should follow the person, not merely the institution.

13. Liability, accountability and delegated authority

Agentic financial services create difficult questions of legal responsibility.

Consider an AI agent that receives a fraudulent invoice and authorises payment within an apparently valid delegation of authority. The consumer may never have seen the invoice. The financial institution may have received an apparently legitimate instruction. The AI provider may have generated an erroneous assessment.

The allocation of responsibility cannot safely be assumed in advance.

This illustrates why governance should distinguish decision support from delegated execution.

The Academy proposes the following accountability framework for further legal and regulatory examination.

Participant	Illustrative responsibility
Consumer	Establish objectives and make decisions within their chosen sphere of control
AI provider	Exercise appropriate care in system design, testing, disclosure, security and service delivery
Financial institution	Meet applicable regulatory, contractual and operational obligations
Professional specialist	Exercise competence and reasonable care within the agreed scope of engagement
Agent operator	Respect the limits of delegated authority and maintain appropriate records
Policymakers and regulators	Establish clear expectations and address gaps in consumer protection and redress

These proposed responsibilities are not a statement of current legal liability. The applicable obligations will depend on the relevant law, contracts, regulatory status and circumstances.

Importantly, preserving human authority must not become a mechanism for transferring all risk to the consumer.

A technology provider cannot reasonably claim to empower a person while disclaiming responsibility for foreseeable defects in its own service. Similarly, a professional who provides materially inaccurate information should not escape accountability simply because the consumer ultimately makes the decision.

Authority and accountability are distinct.

The individual should retain authority over their life. The service provider should remain accountable for the service it delivers.

14. The right to delegate — and the right to take back control

Some consumers will welcome extensive automation. Others will prefer to make decisions personally. Many will move between these approaches according to their circumstances.

Human sovereignty must accommodate this diversity. The objective is not compulsory self-management.

A consumer should be able to delegate routine or complex tasks where this serves their interests, provided that the delegation is meaningful, appropriately safeguarded and understood.

A proposed delegation framework would incorporate graduated permissions.

Permission level	Example
Read	Access specified financial information
Analyse	Interpret information and identify patterns
Prepare	Draft a possible action for consideration
Request approval	Present a proposed action for explicit authorisation
Execute within limits	Perform an authorised task within agreed boundaries
Suspend or revoke	Withdraw authority and prevent further actions where practicable

Delegation should be purpose-specific, proportionate, time-bounded where appropriate, traceable and revocable.

For consequential transactions, additional safeguards may be necessary, including limits, independent verification, confirmation requirements and appropriate human oversight.

The right to delegate should be accompanied by the practical ability to reclaim authority. Otherwise, delegation risks becoming dependency.

15. Vulnerability, accessibility and the limits of autonomy

A human agency framework must recognise that people differ in their circumstances, capabilities and support requirements.

Financial distress, bereavement, coercion, disability, cognitive impairment, poor digital access and language barriers can affect a person's ability to exercise meaningful agency. AI may reduce some barriers while creating new ones.

A system designed only for digitally confident, financially secure and highly literate individuals would risk reinforcing existing inequalities.

The agency-led model should therefore incorporate accessible communication, appropriate escalation pathways, supported decision-making, non-digital alternatives and specialist referral where necessary.

It should recognise that the appropriate level of support can change throughout a person's life.

The Academy's proposed six-stage pathway provides a useful model.

Stage	Purpose
1. Stabilise	Address immediate harm, distress and urgent needs.
2. Clarify	Build an understandable picture of the circumstances.
3. Orient	Identify priorities, resources and available support.

Stage	Purpose
4. Plan	Explore alternatives and develop a workable course.
5. Activate	Implement informed decisions with appropriate assistance.
6. Sustain	Maintain capability and adapt as circumstances change.

The Human Agency Restoration Pathway. Individuals may revisit stages as circumstances change.

This approach is particularly relevant to people recovering from financial exploitation. The first objective may not be financial optimisation. It may be restoring safety, confidence, access to information and the ability to make decisions.

Technology should support that restoration rather than demand autonomy before a person is ready to exercise it.

> **Human agency is not the absence of support. It is the presence of meaningful control, with appropriate support.**

16. Financial crime prevention and citizen capability

Financial crime prevention has traditionally relied heavily on institutions, regulators, law enforcement and specialist investigators. These remain essential.

However, consumer-controlled AI creates opportunities for upstream prevention.

Individuals could use appropriate tools to examine contractual terms, scrutinise unsolicited offers, identify inconsistencies, understand regulatory permissions and formulate questions before transferring money or entering an agreement.

This could complement institutional detection and intervention.

The Academy's Get SAFE initiative explores how individuals can develop practical capabilities to recognise and respond to potential financial exploitation.

Such systems must avoid creating false certainty. An AI application cannot reliably establish that an investment is legitimate simply because it finds no adverse information. Nor should it conclude that criminal conduct has occurred without sufficient evidence.

Its role should be to support critical enquiry, evidence preservation, informed decisions and appropriate escalation.

A citizen equipped to challenge information before committing is potentially less vulnerable than one who must rely entirely on retrospective protection.

The policy opportunity is to strengthen both institutional safeguards and citizen capability.

> **Consumer protection should begin before the consumer becomes a victim.**

17. The Sovereign Consumer Charter

The Academy proposes a Sovereign Consumer Charter as a basis for consultation with consumers, policymakers, regulators and service providers.

The following are proposed entitlements and design commitments. They are not presented as a statement that each is already an enforceable legal right.

Proposed entitlement	Design commitment
1. The right to understand	Access information in a form that supports meaningful comprehension.
2. The right to choose	Explore alternatives and make decisions without undue influence.
3. The right to independent scrutiny	Use appropriate tools and expertise to examine the claims of service providers.
4. The right to control personal information	Exercise meaningful control over the collection, use, access and sharing of personal records.
5. The right to portability	Retain continuity of information and planning when changing providers.
6. The right to informed delegation	Understand and determine the authority granted to professionals and automated systems.
7. The right to reclaim authority	Modify or revoke delegated permissions, subject to legitimate legal and contractual constraints.
8. The right to human support	Access appropriate human assistance for consequential decisions and complex circumstances.
9. The right to challenge and redress	Question outcomes, correct errors and pursue accessible remedies.
10. The right to develop capability	Access services and technologies that support rather than unnecessarily diminish human agency.

> **Technology serves the person. The person remains the principal.**

The charter is intended to stimulate discussion across the boundaries of financial regulation, consumer protection, competition policy, digital rights and professional ethics.

Its implementation would require careful consideration of legitimate limitations, competing rights, proportionality and institutional responsibilities.

It is a direction of travel rather than a completed legal instrument.

18. The commercial architecture of agency

The restoration of human agency must be economically sustainable.

A system cannot depend exclusively on the goodwill of professionals and technology providers whose commercial incentives favour a different outcome. Business models matter.

A firm remunerated principally through product distribution may have an incentive to increase product acquisition. A firm remunerated through assets under management may have an incentive to retain assets within its management arrangements. A technology platform remunerated through advertising, referrals or transactions may have an incentive to influence consumer choices.

These incentives do not establish that any individual firm is acting improperly. They do establish the importance of examining the relationship between commercial success and consumer outcomes.

An agency-led business model seeks to align its income with the delivery of capability, education, planning, decision support and specialist expertise. The consumer pays for the service they need, rather than necessarily for the implementation of a financial product.

Several models can coexist, but transparency about their economic architecture is essential.

Commercial model	Principal source of remuneration	Structural question
Product distribution	Sales or product-related revenue	Does the provider benefit from a particular purchase?
Asset-based services	Assets under management or administration	Does the provider benefit from retaining the assets?
Transactional platforms	Activity, referrals or transaction revenue	Does the provider benefit from increased activity?
Subscription services	Continued participation	Does the service create recurring value or unnecessary dependency?
Agency-led planning	Planning, capability and expertise	Does revenue reflect independently valuable services?

No payment model is inherently immune from conflicts.

Even an educational subscription may create incentives to discourage independence if its profitability depends on perpetual engagement.

The appropriate test is whether commercial arrangements are transparent and whether the provider can prosper when the consumer becomes more capable.

The economics of empowerment must reward empowerment.

19. A practical reference architecture

The Academy of Life Planning is developing an ecosystem that explores how the proposed principles might operate in practice. Its intention is to separate the consumer's ongoing capability from the episodic involvement of professional expertise.



Illustrative reference architecture, not independent validation or regulatory endorsement.

The Academy of Life Planning acts as the educational and public-interest movement.

Total Wealth Plans provides AI-enabled tools intended to support understanding, planning and independent decision-making.

Total Wealth Planner provides complementary human support and specialist planning expertise.

My Life Record™ provides a proposed foundation for continuity of consumer-controlled information.

Get SAFE addresses the prevention of financial exploitation and the restoration of agency following harm.

The wider professional ecosystem supplies specialist services according to the person's circumstances and requirements.

This is not intended to establish a closed ecosystem. Indeed, a genuinely sovereign architecture should be compatible with other providers, technologies and professional communities.

The Academy's own arrangements should be subject to the principles proposed in this paper, including transparency, accountability, portability and evaluation of consumer outcomes.

The Academy should not ask of others what it is unwilling to demonstrate itself.

20. Policy recommendations

The Academy proposes the following areas for consideration by the FCA, other regulators, government, standards-setting bodies and industry participants.

Recommendation	Proposed action
1. Recognise human agency	Include the preservation and development of consumer agency among the objectives considered in agentic AI policy development.

Recommendation	Proposed action
2. Differentiate AI functions	Distinguish information provision, interpretation and exploration from recommendation, arrangement and execution.
3. Research agency outcomes	Develop and validate measures of consumer understanding, control, appropriate confidence and practical capability.
4. Promote consumer-directed interoperability	Explore secure standards for legitimate consumer-controlled tools to access and interpret financial information.
5. Separate information access from execution authority	Develop granular permission frameworks reflecting the different risks of reading, analysing and acting.
6. Strengthen transparency of incentives	Examine commercial arrangements, referral relationships and ranking mechanisms that may influence consumer-facing AI outputs.
7. Clarify accountability and redress	Identify gaps in the allocation of responsibility across technology providers, institutions, professionals and agent operators.
8. Support lifelong information portability	Explore practical standards for continuity of consumer records and financial plans across providers.
9. Protect inclusive participation	Preserve appropriate human support, accessible alternatives and safeguards for consumers in vulnerable circumstances.
10. Test agency-led business models	Include independent, non-intermediating planning and capability-building models in regulatory research and innovation initiatives.

These recommendations do not assume that additional regulation is necessarily required in every area. Some may be addressed through existing law, regulatory guidance, technical standards, competition measures, professional practices or consumer education.

Others may reveal gaps requiring legislative or regulatory reform. The Academy encourages proportionate, evidence-led development rather than premature prescription.

21. A proposed contribution to the FCA Agentic Policy Sprint

The Academy would welcome the opportunity to contribute to the FCA's proposed Agentic Policy Sprint.

Its distinctive contribution would be to bring a consumer-agency perspective to a conversation that could otherwise concentrate primarily on institutional innovation and autonomous execution.

A proposed workstream could be entitled:

> **Human Agency in Agentic Finance: Consumer Capability, Sovereignty and Decision Authority.**

The workstream could explore five research questions.

- How can consumer-facing AI be designed to increase rather than diminish financial capability?
- How can consumers meaningfully understand and control the authority delegated to AI systems?
- How can consumer-controlled AI interact securely with financial institutions without creating additional forms of dependency?
- How can the quality of AI-supported decisions and the development of consumer capability be measured?
- What commercial, technical and governance arrangements best support human sovereignty?

Potential participants could include consumer representatives, independent financial planners, financial institutions, AI developers, behavioural scientists, accessibility specialists, data protection experts and professionals working with victims of financial exploitation.

The desired output would be a practical, testable framework for Human Agency by Design.

22. Research and implementation roadmap

The propositions in this paper should be tested rather than merely asserted.

The Academy proposes a phased programme of development.

Phase	Activity and output
1. Consultation and conceptual development	Publish the discussion paper, engage stakeholders, challenge definitions and identify existing research relevant to human agency. Output: revised framework and research agenda.
2. Measurement and testing	Develop and validate the Human Agency Test using appropriate consumer research and comparative service evaluations. Output: evidence-based assessment methodology.
3. Governance and technical standards	Develop practical design requirements covering permissions, privacy, transparency, portability, safety and accountability. Output: Human Agency by Design governance standard.
4. Cross-industry pilots	Test interoperability and agency-led service models with consumers, professionals, technology providers and financial institutions. Output: documented evidence, case studies and lessons.
5. Public-interest adoption	Share findings, develop educational resources and contribute to professional and regulatory standards. Output: open principles and practical implementation guidance.

Research should explicitly include negative findings and unintended consequences.

It should examine AI-generated misinformation, inappropriate confidence, privacy failures, digital exclusion, manipulation, excessive reliance and situations where professional or institutional intervention produces better outcomes.

The purpose is not to prove that an agency-led model is always preferable. It is to establish when, how and under what conditions it genuinely improves people's lives.

23. Conclusion: from artificial intelligence to human flourishing

Artificial intelligence is often described as a technology that will transform the way organisations operate. Its more profound potential may be to transform the relationship between organisations and human beings.

For much of the history of modern financial services, individuals have depended on institutions and intermediaries because they could not independently access the knowledge, infrastructure and capabilities required to navigate the system. That constraint may be beginning to loosen.

For the first time, consumers have the prospect of carrying sophisticated analytical capabilities into their interactions with financial institutions. They may be able to interrogate information, challenge professional assertions, explore alternatives and participate in decisions on a more informed basis.

But this outcome is not inevitable.

We could build a future in which institutions possess even greater technological power over consumers. We could replace traditional intermediaries with algorithmic intermediaries. We could create increasingly convenient systems in which individuals understand less, choose less and gradually lose the ability to act independently.

Or we could deliberately design a different future.

One in which technology expands access to knowledge.

One in which professional expertise remains available without requiring unnecessary dependency.

One in which individuals retain meaningful control over their information, decisions and lives.

One in which financial services exist to support human flourishing rather than becoming the organising purpose of human life.

The Academy of Life Planning believes that the future of financial services should be evaluated not only by what its machines can accomplish, but also by what its people become capable of accomplishing.

The objective is not simply more intelligent institutions. Nor is it necessarily more autonomous machines. It is more capable human beings.

The individual is not merely a customer, a risk profile, a source of data, a portfolio or a collection of transactions.

The individual is the principal. Everything else should serve that principle.

As policymakers, professionals and technology developers consider the future of agentic finance, we invite them to consider a question that may ultimately matter more than the sophistication of any algorithm:

> **How much authority can artificial intelligence help us return to human beings?**

Invitation to dialogue

The Academy of Life Planning welcomes constructive challenge to any of the principles, frameworks and proposals contained in this paper.

We invite regulators, policymakers, financial services professionals, technologists, academics, consumer representatives and members of the public to participate in the conversation.

Our intention is not to establish a closed doctrine or claim exclusive ownership of an emerging field.

It is to contribute to the development of financial services in which human agency becomes an explicit, observable and valued outcome.

The future is not something that happens to us. It is something we can participate in designing.

Agency in. Dependency out.

References and regulatory context

[1] FCA. Our approach to Agentic AI, 24 November 2026.

[2] FCA. Using AI agents, 24 November 2026.

[3] FCA. Mills Review announcement, 2026.

[4] FCA. Mills Review engagement paper, 2026.

[5] FCA. Building the future of finance in the age of AI, 19 October 2026.

[6] FCA. AI in financial services: good practice and areas for improvement, October 2026.

[7] FCA. About the Consumer Duty.

[8] FCA. Consumer understanding: good practice and areas for improvement.

[9] FCA. Money talks: Lessons from two LLM pilots on consumer guidance.

[10] ICO. Agentic AI: Data protection and privacy risks.

References carried forward from the discussion draft. The links and publication details have not been independently re-verified during this formatting exercise.