



Chad Cracknell, Account Executive, Microsoft UK Financial Services

Miriam Lloyd, Director, Microsoft UK Financial Services

Adam Richardson, Account Chief Technologist, Microsoft UK Financial Services

14:

The Building Society Workforce: People, Skills and the Mutual Advantage

Executive Summary

- **AI changes tasks, not the purpose of the mutual.**
In the next decade most routine financial services work is expected to be performed by AI agents. As Microsoft's Brad Smith puts it, a job is best understood not as a title but as a bundle of tasks, some AI can do, many you do better with AI, and some only people should do.¹ The roles colleagues do will change; the purpose of a mutual, serving members, does not.
- **The skills profile changes in two directions at once.** AI fluency becomes a near universal requirement, while the distinctively human skills, judgement under uncertainty, empathy, and ethical reasoning will rise in value. Most of the skills employers seek today are expected to remain relevant, applied alongside AI rather than displaced by it.
- **People centered service is the durable, though not automatic, differentiator.** As AI lowers the cost of standard products across the market, member trust, mutual ownership and genuine personal service will become harder for competitors to replicate. Realising that advantage depends on deliberate design, not on technology alone.
- **Trust and accountability are the foundations.** Governance, explainability and meaningful human oversight, set against a clear view of operational, model, data and conduct risk will allow boards to adopt AI with confidence and keep members protected.
- **Sector collaboration is the decisive enabler.** No single society can fund, govern and reskill around AI and agents alone. Acting together through the BSA on four fronts, shared AI foundations, common Responsible AI standards, pooled reskilling pathways, and a shared regulatory voice will be an important choice the sector can make through the next decade.

¹Brad Smith, AI, jobs, and the next generation, Microsoft On the Issues, 10 June 2026



Introduction — people at the core of a digital future

Building societies face two connected shifts. AI and agentic systems are rapidly becoming part of the operating fabric of financial services. As AI lowers the cost of delivery for products in every provider, and members increasingly use AI to find, compare and take products, the durable provider differentiation will narrow towards what is hardest to copy, trust, mutual ownership and community presence.

AI is not a substitute for the mutual model, nor a guaranteed rescue for it. It is a capability that can amplify what building societies are for if the sector invests in the right skills, governs the technology well and acts together.

This chapter addresses three practical questions: how roles will change, what skills will matter, and what choices the sector should make now.

1. The central recommendation: Collaborate

The challenge is not whether AI is useful, but the scale and cost of doing it safely. Deploying AI, governing it to regulatory standards and reskilling a workforce will require a significant and sustained investment. A sector ranging from large national mutuals to single-branch societies cannot absorb that evenly, and a society by society response risks leaving smaller mutuals behind.

The case for collaboration rests on scale, not on an absence of competition. The cost of safe AI foundations and regulatory evidence is largely fixed and largely shared, pooling effort lets every society, however small, reach a standard none could justify alone.

We therefore recommend sector action on four fronts:

- Shared AI foundations: reusable, pre-governed building blocks such as reference architectures, vetted models and common controls.
- Common Responsible AI standards: a shared framework for fairness, explainability and oversight, evidenced to regulators at sector level.
- Pooled reskilling pathways: shared curricula, accreditation and talent programmes so societies can develop AI-fluent colleagues to a common standard.
- A shared regulatory and policy voice: a single, credible channel through which the sector engages the FCA, PRA and government on AI, proportionality and the mutual model.

For a sector rooted in community, collaboration is not a side issue. It is how the sector modernises without narrowing its diversity. This also answers a wider call: past transformations left too many people and places behind, and doing better means shared responsibility across employers, government, educators, communities and workers.



2. From execution to orchestration: how roles will change

As AI agents absorb structured, repeatable work, roles will shift from execution to direction, judgement and member engagement.² Routine processing moves into agentic systems whilst people remain responsible for exceptions, conduct, vulnerability and decisions with real consequences.

The table below maps that shift onto sector specific roles. These are not distant end states, but directions of travel are already visible today.

Today's role	Evolves towards	Now focuses on
Mortgage advisor	Member outcomes advisor	Understanding the member's circumstances, sense checking AI, override it where needed, and guiding members through trade-offs.
Branch colleague	Community and digital coach	Coaching, financial literacy support and digital inclusion, while explaining how AI supported decisions are reached.
Customer service rep	AI quality controller / escalation specialist	Monitoring agent outputs, intervening on complex cases, retraining edge cases and escalating where AI is wrong or unfair.
Compliance and risk officer	Responsible AI lead	Testing agents for bias before deployment, monitoring live outputs against mutual principles and publishing transparency statements.
Vulnerability specialist	AI augmented vulnerability lead	Reviewing AI detected distress signals, deciding the intervention and conducting the human conversation.
(new)	AI orchestrator / agent operations lead	Owning the portfolio of production agents, monitoring performance, retraining cadence and human handover design.
(new)	Prompt and policy engineer	Translating policy, regulation and mutual principles into prompts, guardrails and rubrics that govern AI behaviour.
(new)	Member outcomes auditor	Auditing AI decisions against outcomes data to identify unfairness.

Colleagues become conductors of AI-assisted work. Systems handle volume; people remain accountable for outcomes.

²Cambon, A. et al. (Microsoft WorkLab) (2025–2026). Human–AI collaboration and task decomposition research.



3. The new skill set: AI fluency, judgement, and the skills only people bring

The workforce implication is straightforward in that AI fluency becomes foundational across almost every role, while the capabilities only people bring rise in value. Microsoft research shows that as AI is applied to more and more processes, expertise and empathy matter even more^{3,4}. Simply put, as automation expands via AI, the quality control and critical thinking people bring, increases the success^{3,4}. A simple three-layer framework helps define the new skills profile:

- **Foundational (everyone):** AI fluency in day-to-day work, digital and data literacy, regulatory and ethical awareness, and continuous learning.
- **Specialist (functional / AI-adjacent):** evaluating and overseeing AI outputs, agent design and orchestration, prompt and policy engineering, and domain expertise expressed through audit trails.
- **Distinctive (the skills only people bring):** empathy, deep listening, judgement under uncertainty, ethical reasoning, behavioural insight, vulnerability sensitivity and fluency in mutual values.

The requirement is easy to state and hard to deliver, widespread AI fluency, targeted specialist capability and stronger human judgement in member facing work.

4. Talent development, work redesign, and the scale question

Changing a workforce's skill profile and culture is a larger undertaking than a mere technology rollout, and cannot be outsourced. The path forward requires continuous talent development and deliberate work redesign.

- **Always-on reskilling:** the useful life of a skill is shortening. The World Economic Forum estimates that 39% of workers' core skills will change by 2030 and 59% of the workforce will need retraining over the same period⁴. Continuous learning has to be built into the operating model, not run as a one-off programme.
- **Time given back to members:** the payoff from AI fluency is redeployable time. Yorkshire Building Society is piloting AI agents to support internal risk and control testing. Early results suggest efficiency savings of around 40% – freeing up time for professional judgement rather than replacing it⁵. For a mutual, time that can be reinvested in member relationships.
- **Flexible workforce models:** colleagues start to move between advising members, overseeing agent outputs and community engagement, rather than delivering fixed single task roles.
- **Work redesign led from the top:** Societies that decide deliberately what to automate, and where a person must remain in control, will capture more value and do so safely.

Culture matters as much as tools. Organisational factors weigh more than twice as heavily in whether AI delivers^{3,4}. Given that less than 1 in 4 colleagues report clear leadership alignment, closing alignment gap is a board responsibility.

³Microsoft, 2026 Work Trend Index Annual Report – Agents, human agency, and the opportunity for every organization, 5 May 2026

⁴World Economic Forum (2023–2025). Future of Jobs / Jobs of Tomorrow reports (with Microsoft contribution).

⁵Microsoft (2026). Work Trend Index – Organizational performance and AI adoption insights.

⁶How Yorkshire Building Society is using AI to give colleagues more time for members, 3 July 2026



5. Trust, ethics and risk: people in the loop

Trust is the foundation of the mutual model, and AI raises the stakes on both sides. Used well, it can extend fair, consistent, member first decision making across all interactions. Used carelessly, it introduces new failure modes that boards must understand and govern.

Responsible AI should be treated as a workforce and governance discipline, not just a technical control set⁷. AI enabled services should be designed for explainability, human oversight and override.

- **Operational resilience:** greater reliance on AI and a small number of providers increases the impact of outages.
- **Model drift and risk:** without monitoring, challenge and retraining, a model that was fair and accurate at launch may become neither.
- **Third party and concentration risk:** shared foundations bring efficiency, but also common dependencies that require substitutability and exit planning.
- **Explainability and contestability:** members affected by an AI-influenced decision should be able to understand and challenge it.
- **Data privacy and security:** richer data and greater automation widen the attack surface and raise the stakes on consent, minimisation and protection.
- **Member exclusion and digital access:** automation must not leave behind members who are older, offline or less confident digitally.
- **Vulnerable customers:** AI may detect difficulty earlier, but the response should be a prompt handover to a qualified colleague, not an automated outcome.
- **Consumer Duty alignment:** AI mediated journeys must deliver good outcomes, avoid harm and support members' financial objectives with evidence a board can stand behind⁸.
- **Governance, accountability and human oversight:** accountability cannot be delegated to a system, there must be clear ownership and a named person answerable for outcomes.

Several of these issues are cheaper and safer to address collectively. Common Responsible AI standards and shared controls let the sector evidence safety to regulators at a higher standard than smaller societies could reach alone.

What this means for boards and regulators

For boards: Treat AI as a workforce and governance programme, not an IT project². Set the risk appetite, name the accountable owner, and insist on explainability and human override by design. Invest early in reskilling, the binding constraint on value is people and process, not tools. Where the economics are hard, pursue shared sector capability through the BSA rather than going it alone.

For regulators and policymakers: The mutual model's diversity and consumer focus is a good worth protecting. Proportionate expectations, and a willingness to engage with shared, sector level Responsible AI standards, can let smaller societies adopt AI safely without bearing large society compliance costs. Engaging the sector through a single BSA channel is the most efficient way to set expectations that are both safe and achievable. Consumer Duty provides a ready made outcomes lens for supervising AI in retail financial services⁵.

⁷Microsoft, Responsible AI Transparency Report 2025, June 2025

⁸Financial Conduct Authority, Consumer Duty (PRIN 2A), in force from 31 July 2023



6. How an AI augmented workforce can make mutuality a differentiator

As strong digital platforms and capable AI become common across retail finance, the question is not whether building societies can use AI, but why their use of it should produce different outcomes. The answer is structural, it lies in ownership, not technology.

A shareholder owned firm is expected to return efficiency gains to investors. A mutual has no external claim on that dividend. That difference allows a Society to direct more of AI's gain back to members.

- **AI can buy back the time personal service depends on:** every minute an agent saves on administration is a minute a colleague can spend on a member relationship.
- **AI can extend fair, member-first decisioning consistently:** colleagues can set the rules, audit the outputs and override the system where member outcomes demand it.
- **AI can let the human moments count more:** removing friction from everyday transactions frees colleagues for major life decisions, financial difficulty and questions of trust.

If AI drives standard products towards sameness, trusted advice, genuine care and verifiable member-first outcomes become more valuable, not less.

Conclusion — a call to action

The technology questions are increasingly answerable. The decisive questions are about skills, governance and whether the sector acts together.

- **For building societies and their boards:** commit to AI fluency for every colleague and work redesigned around people plus AI. Set risk appetite and accountability, and reinvest the time AI saves into members.
- **For the BSA and its members:** convene the four shared capabilities in Section 1: shared AI foundations, common Responsible AI standards, pooled reskilling pathways and a shared regulatory voice.
- **For regulators and policymakers:** protect the diversity of the mutual model through proportionate, outcomes based expectations and engage with the sector's shared Responsible AI standards as a route to safe, affordable adoption for societies of every size.

History offers grounds for optimism: when earlier technologies increased what was possible, human ambition generated new demand rather than less work, and AI used to accelerate expertise rather than replace it keeps human judgement as the differentiator. In a market where AI pushes standard products towards sameness, the value of genuine care, trusted advice and verifiable member-first outcomes may well rise, the advantage an AI augmented mutual workforce is well placed to deliver, if the sector makes the choices, and the investment, to earn it.

AI does not dilute the mutual difference. It will make it more visible. Competitors can match technology, what sets a mutual apart is the choice to send more of the benefit back to members.