Institute *for the* Future of Work

Report

Good Work Design and AI Adoption

Action Research in Eight Cross-Sector Case Studies

May 2026

Funded by

UKRI Innovate UK

BridgeAI

In partnership with

CIPD

wbs WARWICK BUSINESS SCHOOL THE UNIVERSITY OF WARWICK

Contents

Project Team

IFOW
Semeli Hadjiloizou
Danielle Soskin
Mia Leslie
Anna Thomas MBE
Professor James Hayton
Dr Magdalena Soffia
Jo Marriott

CIPD
Ben Willmott
Lizzie Crowley
Peter Cheese
Hayfa Mohdzaini
David D'Souza
Was Rahman

Citation

Good Work Design and AI Adoption: Action Research in Eight Cross-Sector Case Studies, London: Institute for the Future of Work.
DOI: 10.5281/zenodo.20084164

1. Introduction to the work

Anna Thomas MBE

Artificial Intelligence (AI) is already reshaping work and working lives. Politicians and practitioners alike are lured by bold promises of significant gains in efficiency and productivity. Much less is known about how AI innovation can translate into a wider range of good outcomes - including new job creation and improved work quality - and where the process of adoption can go wrong.

To meet this challenge, we need a reorientation from just the technical towards the human and social dimensions of this transformation. We know that AI adoption is mediated by organisational capabilities and practice, decision-making and resources. But too little is known about the interaction between individual, firm and 'system' level factors that are engaged as AI is adopted.

Our case studies have taken a 'work design' lens to understand this messy interaction. They surface many of the firm and human capabilities - and the choices made - that decide which tasks are done, which skills are developed, how autonomy is protected and how work systems can be better designed in the age of AI.

To do this, we have used a framework for responsible AI and 'work design' produced as part of my own doctorate and developed by the IFOW team through the eight case studies summarised in this report.

This approach has allowed us to shed light on and share new insights about how barriers or 'frictions' can be overcome in practice. It strongly suggests that the design of 'good work systems' is the missing link in effective and responsible AI adoption.

The race to understand and shape better outcomes is on, as reports of the 'wild west' of AI adoption abound. Rushed experimentation, unclear goals, patchy governance, and the absence of shared learning or rigorous evaluation are common. So these case studies are a timely contribution that should help people, employers, and government develop a shared understanding, language and tools to meet new challenges.

Good work design is associated with higher motivation, engagement, satisfaction and wellbeing, and with better performance and resilience. Good work design increases creativity, innovation and productivity and is associated with augmentation and human-computer complementarity, and good quality work too.

Using our new approach and this 'good work design' lens, the IFOW team has been able to conduct action research at pace and share eight journeys of AI adoption across private and public sectors, surfacing their aims, those involved, those affected and whether new 'value' was actually created. They uncover new insights about the process and options for reshaping and directing AI implementation to meet pressing challenges, address risks, and open up new opportunities. The organisations involved are:

- **MemberOrg** — exploring co-created approaches to generative AI adoption in a not-for-profit organisation
- **ConstructionCo** — developing more purposeful approaches to AI adoption in construction
- **FinanceCo** — supporting AI-enabled software developer performance through collaborative job design
- **HealthTrust** — engaging NHS staff on the rollout of an HR chatbot

- **HolidayCo** — developing people-centred approaches to AI adoption in retail and tourism
- **InnovationCo** — building foundations for responsible AI adoption in a safety-critical and unionised sector
- **LargeCo** — examining how limited employee participation threatened return on AI investment
- **LegalCo** — introducing an AI legal assistant within a law firm

This publication shares summaries of these case studies. For IFOW, it is the first in a series of work conducted under the new 'AI Adoption' challenge of the new IFOW strategy. It aims to share deep insight in an accessible way, highlight possible interventions that can be made at a firm level and signpost how these powerful new technologies can be adopted in ways that are both efficient and responsible, transforming working lives for good.

I am hugely grateful to my supervisor, Professor James Hayton, and to Professor Sharon Parker and our project Steering Group, Dr Magdalena Soffia and the brilliant IFOW team, our partners - the CIPD and Innovate UK - for making this work possible so that we can continue to transform work and working lives for good.

Case Study 1: MemberOrg

A more participatory approach: a not-for-profit looks to deploy generative AI to build capacity

This case study details steps taken by a not-for-profit membership organisation to move from a self-led approach to generative AI to a shared organisational position on one general-purpose AI tool.

Profile

This not-for-profit membership organisation, referred to here MemberOrg, assembled a cross-functional working group of five people, the majority being senior staff, including roles responsible for HR, legal, governance, IT and data.

Operational context

MemberOrg is a values-driven organisation and is guided by its core mission to manage the rights of its members effectively. While it is committed to using new technologies to deliver better quality services and improve its processes, MemberOrg is conscious that its actions must align with its values. These values were the outcome of what they described as a 'big process of working collectively' towards 'greater transparency and co-creation', which began in 2023.

Co-creation has been enacted through the development of a staff consultation forum with one staff representative from each department. Through this, employees are involved in organisational decision-making, as their HR director explained:

> *"Staff members work directly with senior leadership and the Board through consultation and regular staff forums to ensure that the whole organisation comes together for decision-making across projects or workstreams where appropriate. This results in regular refinement on activity and objectives to ensure they are agile and manageable across teams, maintaining the co-creation process."*

Members of the working group reported that employees have high levels of autonomy and agency. This was echoed in the results of the survey completed by four-fifths of the workforce, which was deployed by IFOW.

It was also emphasised that the organisation offers significant room for growth and

development. Employees (both inside and outside of the working group) reported that they had the opportunity to undertake different roles at the organisation. This was said to have contributed to employees remaining with the organisation for several years.

The tasks conducted at MemberOrg spanned several different teams; some were very small, and others were relatively large. The tasks conducted by the different teams were described as diverse and complex.

This operational context produced what was described as a 'self-led' approach to AI. IFOW's survey revealed that two-thirds of the workforce were using general-purpose generative AI tools in their day-to-day work, most commonly for searching and retrieving information. Some were also engaging in more advanced usage, such as creating agents to execute discrete tasks and entire workflows. A member of the legal team described creating two agents for capability testing. One was tested to estimate fees for services. The second was tested to assess where an email might have come from, move it into the correct inbox and send a re-direction reply if the sender was from abroad.

Though some employees were reported to be using generative AI extensively, no tools had been formally integrated into organisational processes, nor had MemberOrg provided paid licenses or detailed guidance regarding how particular tools could be used for different tasks. Moreover, the use of AI was not formally reflected in organisational policy or strategy and responsibility for oversight and key decisions had not been assigned.

Challenge

When the action research project began, the self-led approach was governed by a set of guidelines created by the organisation (the 'AI Use Guidelines'). Implemented in November 2023, the guidelines were intended as a framework for employees to make self-assessments about the acceptability of their usage. The guidelines set out a checklist with requirements, along with examples of permitted and prohibited uses. They required employees to take the following actions:

- Ensure the accuracy of AI outputs.
- Not use any organisational or third-party confidential data as prompts for publicly available generative AI tools.
- Ensure that AI use does not result in discriminatory or biased outcomes.
- Ensure that usage does not produce material that may infringe third-party rights.
- Use AI only in ways they would be comfortable publicising.
- Be transparent with colleagues about AI use.
- Provide human oversight over input, processing and output and not rely on automated decision making.
- Comply with company values, policies and the law.

The examples were broad and did not always link directly to typical organisational tasks or processes that employees might perform. They were not linked to a specific generative AI tool and there was no clearly defined space or mechanism to surface, discuss and document different uses across teams. As such, there was not a shared understanding of how individual jobs and organisational workflows were changing because of the use of generative AI tools.

This led to concerns that certain uses could potentially be undermining organisational values, as well as lack of trust between employees about who was using generative AI tools and how. There was also a concern about adoption gaps emerging between different teams and individuals, with some using AI extensively and others fearing to use it at all.

The working group wanted to use generative AI tools to maximise resources in the context of capacity constraints and deliver efficiency gains. Certain teams within the organisation had significant amounts of administrative work, and it was felt that generative AI tools could potentially be used to reduce some of the time devoted to that. However, the working group emphasised that the complexity and diversity of tasks meant that the suitability of AI usage would vary across the organisation.

Finally, there was a concern among the working group that the AI Use Guidelines were placing an undue amount of compliance responsibility on employees, as well as creating anxiety around generative AI usage that was preventing experimentation. Described by a working group member as a 'panic response' to the proliferation of generative AI, the guidelines were drafted by the legal department and circulated among the senior leadership team. There was a sense that some employees were consulting them, as the legal department was receiving requests about whether certain uses of AI were acceptable in line with the guidelines.

However, employees had not been consulted or involved in their creation. Extending the commitment to co-creation to the development of an approach to generative AI and seeking employee feedback to update the AI Use Guidelines were underlined as key goals.

What they did

To facilitate a critical reflection about the use of generative AI tools in the organisation, the working group co-designed a half-day workshop with IFOW researchers. The workshop was hosted at MemberOrg's offices in September 2025. Half of the workforce attended, with representatives from every team, spanning different levels of seniority.

MemberOrg had initiated a dialogue about the use of generative AI tools in 2024 through a workshop and an employee survey. This surfaced a consensus that the use of these tools could be compatible with MemberOrg's values, with the caveat that careful attention is needed in relation to transparency, data privacy and the impact on human interactions. Data privacy was regarded as particularly important as MemberOrg manages the personal data of its members. However, the workshop had not led to any formal updates to either organisation's policy or practice. The 2025 workshop with IFOW as part of the action research was designed to reignite dialogue. The objectives were threefold:

- identify potential use cases and associated risks and impacts, both positive and negative,
- critically reflect on the existing approach to AI usage and the AI Use Guidelines and,
- co-design interventions to allow MemberOrg to maximise opportunities and to attempt to mitigate challenges.

The workshop surfaced a broad desire to explore the use of generative AI tools for 'low-risk' tasks and reduce the time spent on those that are mundane and repetitive. However, the use of generative AI tools for certain tasks was thought of as too low value to justify the environmental impact of the tools' energy use. Other use cases were thought of as too 'high risk' because of the employees' awareness that generative AI tools are inherently biased and likely to produce outputs with inaccuracies.

The negative impact of some use cases on human skills and relationships was highlighted too: reduced communication, negative impacts on critical thinking skills and the disruption of skill development pipelines for junior employees and future entrants to the workforce. It was emphasised that generative AI use could increase monotony for employees due to the need to check the accuracy of outputs. This is associated with disengagement and reduced vigilance, as well as potential negative effects for productivity.

Potentially valuable use cases highlighted were around data wrangling, document analysis, and the use of AI agents to determine the contents of emails and reply with pre-existing templates, which are currently copied and pasted by employees. However, it was emphasised that it would be important to determine which processes could benefit from generative AI specifically, as opposed to other non-AI-based automations. A workshop participant noted: "It's a huge bit of work, as to what's just automation and what's AI, because I find a lot of the time, that's mixed up. And it's like, actually, we don't need AI, that is just automating a process."

The importance of assigning responsibility for managing agents and maintaining an oversight of AI usage across the organisation was also discussed. It was emphasised that everyone did not necessarily have the skills or desire to develop agents and that processes would need to be clearly documented to prevent 'chaos' and organisational knowledge loss, if the employee that originally created the agent left the organisation.

This was discussed in relation to training requirements, with employees requesting structured guidance based on company-wide ways of doing things and 'safe spaces' to experiment and share knowledge. Reflecting on the current approach, one workshop participant noted:

> *"I think the problem has been, 'AI is here, go and use it', but also, don't put any sensitive data in it [...] So, it would need to be, here's your agent, if there are changes, these are the changes. And now we'll run through and do a few goes, and a manager will oversee you using it a few times, so that you're happy."*

Relatedly, the workshop confirmed that the AI Use Guidelines were not fit for purpose. An anonymous poll revealed that only 47% of participants were aware of and using them. Ultimately, the guidelines were criticised as abstract, subjective and dependent on high levels of knowledge about generative AI. For example, some employees were unclear what was meant by 'publicly available generative AI tools'. By design, it was also emphasised that they lacked consideration of environmental sustainability, one of MemberOrg's values.

This discussion surfaced a shared desire for a more practical set of guidelines. Those guidelines would be tied to a specific generative AI tool or tools, and relevant to organisational processes to reduce the possibility for vastly different interpretations about acceptability. A series of other interventions were also surfaced, which the working group deliberated. In the short-term, the working group opted to:

- Assign Microsoft Copilot as the formally vetted organisational tool and provide paid licenses to all employees.
- Undertake MemberOrg's 'process mapping' practice to identify how Copilot can and should be used within existing workflows.
- Provide a whole staff training on Copilot alongside the rollout of the paid licences (though use of Copilot would not be compulsory).
- Run weekly surgeries for a month after paid licenses are rolled out to provide a safe space for employees to experiment, learn and troubleshoot collaboratively.
- Develop a checklist of questions (eg, impact on job roles, existing company policies and ethical risks relating to bias and fairness) that will be used to vet proposed use cases for Copilot and AI agents.
- Update the AI Use Guidelines based on the discussion at the workshop and circulate for employee feedback.
- Maintain the working group established for the action research as a forum for deliberation and decision making.

Outcomes

Actively including employees in dialogue and decisions about the organisational approach to generative AI has allowed MemberOrg to surface potential use cases, recommendations for company policy and practice and important risks and harms, as seen by employees.

Reflecting on the process, the HR director noted:

> *"We know autonomy is important to our workforce, we know skill variety is important, we know people want to hear the authentic voice, not the voice of AI, all of these things have come out in a way where perhaps we hadn't consciously thought about it."*

Through this, MemberOrg has developed a foundation for a shared understanding and a higher level of engagement. A working group member noted:

> *"I can see the voice of different people across the organisation coming through, and the fact that we've co-created this, I feel very much like we're getting the momentum, and I think that's part of doing this as a co-creative process. So, we have got that enthusiasm within the team. People are keen to see what's going to happen next."*

In practice, it has initiated a move away from a purely 'self-led' approach to generative AI tools towards assessing of the utility of a vetted tool (Copilot) for different organisational tasks. This task-by-task approach is underpinned by a "desire to solve problems rather than kind of create new ones or doing it AI for AI's sake," as one working group member noted.

Recognising that employees know their jobs best and that they can play an important role in deciding how Copilot should be used to deliver efficiency gains that are good for both them and the organisation, MemberOrg will conduct process mapping in collaboration with employees. Changes to processes, as a result of the use of Copilot will be implemented with the agreement of all affected parties.

It was emphasised, however, that the team responsible for process mapping would need to be sufficiently aware of the functionalities of Copilot to highlight potential solutions. It was also emphasised that employees would need to be provided with the time to experiment for themselves. This was highlighted as a potential challenge considering that capacity constraints had initially motivated interest in the use of generative AI tools.

Ultimately, MemberOrg is at an early stage in relation to its use of generative AI and the actual impacts remain speculative. This highlights the importance of continually evaluating the effect that generative AI is having on work and workers, following the adoption of AI. While the working group is positive about the potential efficiency gains, it is keen to ensure that they are not secured at the expense of employee motivation and job satisfaction, with a recognition that increased efficiency may itself deliver negative effects that require careful monitoring and mitigation.

Reflecting on their own usage at the close of the action research project, a member of the working group cautioned about the "risks of burn out" that could arise from:

> *"Having a tool that – whilst it can do things quickly – may lead to, for example, compressing time frames in your mind of how long a job or task might take, or how many tasks I might be able to achieve. For example, previously, I might push a task into the next quarter. But now schedule it for next week because I know I've got this digital assistant that helps me get things done much more quickly. However, I've now also got an expectation on myself about how much more I can get done, and potentially this will detrimentally alter my view on how others, who have access to the same tools, can get done."*

Learning Points

Teams and individuals within organisations will exhibit different levels of tolerance and openness to generative AI tools in relation to varying risks and benefits.
A task-specific approach to general-purpose generative AI tools is necessary to identify and realise genuinely valuable innovations.

AI guidelines that depend on subjective assessments may create anxieties for employees and prevent experimentation.
Clear documentation and communication regarding how generative AI tools are and can be used at organisations is necessary to ensure shared understanding and trust. Responsibility for overseeing systems and mitigating risks should also be clearly assigned to management and signposted to employees.

Employees are best placed to evaluate the utility of generative AI tools for their own tasks, as well as risks and impacts that require attention and mitigation.
For employees to be able to do this effectively, they should be provided with protected time and support to experiment, and governance structures should create mechanisms for meaningful participation.

Training should extend beyond guidance about how to use a particular tool.
Creating safe spaces for peer-to-peer experimentation can help to foster a relational approach to learning and development, which can facilitate knowledge exchange.

Though efficiency is a goal for many organisations, it is neither a given nor does it come without wider implications, some of which may be negative.
Organisations should take a holistic approach to AI impact assessments which look beyond time savings and task volumes to consider impacts relating to impacts to autonomy, skill requirements, tasks, job demands, relationality and employee wellbeing.

Case Study 2: ConstructionCo

Using AI with purpose: a construction company's journey

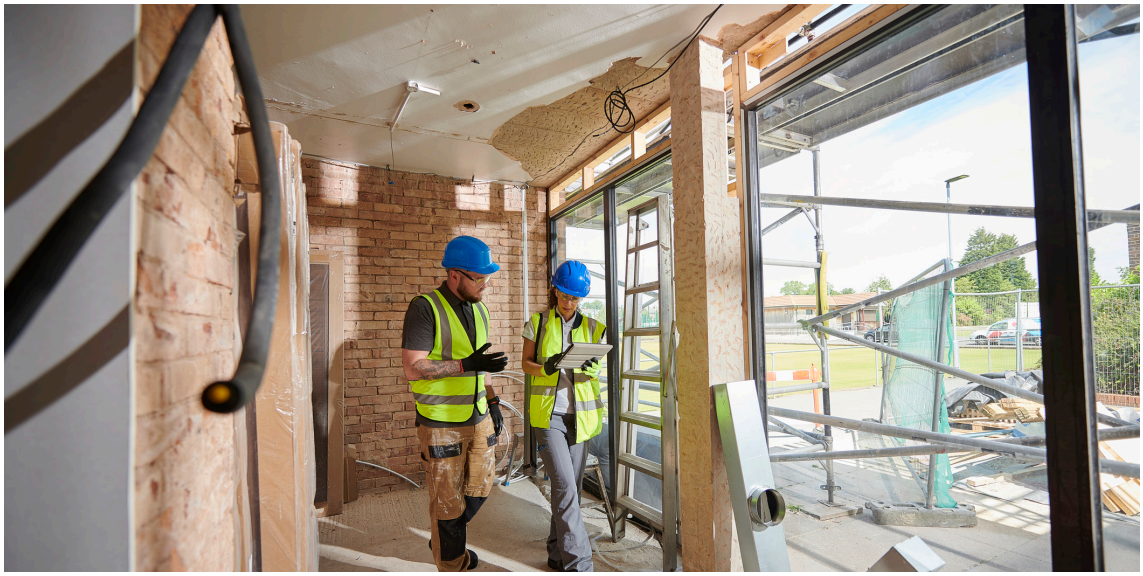

Profile

This case study focused on a construction company, referred to here as ConstructionCo. It looked at how they learnt from the past to become more deliberate in how they buy and use AI and other technologies.

ConstructionCo handles every step of building a home, from identifying and buying land to construction and sales. It employs several thousand people across its regional divisions in the UK. In the past two years, its HR team grew from one to 25 people.

Even though interviewees reported believing that the construction industry as a whole is slow to use new technologies, the people at ConstructionCo felt they were falling behind competitors. When IFOW started engaging with ConstructionCo, its HR department had procured an AI-enabled HR system, which was now in a testing phase. Meanwhile, selected departments across ConstructionCo were trialling a general-purpose generative AI tool.

This case study explored how ConstructionCo learned from the past and became more thoughtful in how they choose and use new technologies across the organisation. For example, how a prioritisation technique guided the HR department's procurement and how the legal department assessed a general-purpose generative AI tool.

A cross-functional working group got together to support the action research process with IFOW, including representatives from HR, IT and risk departments.

Operational context

Employees were wary about using new tools because previous launches had not gone well. In one case, a new system was introduced with limited employee consultation. Because the system did not meet their needs, hardly anyone used it, and ConstructionCo had ended up cancelling the contract.

Despite these past negative experiences, regular feedback from quarterly townhalls and employee surveys showed high levels of trust and engagement. During the townhall meetings, management shared their plans for new technology and invited employees to

share their thoughts about the new HR system and on using generative AI. This renewed interest in AI was partly sparked by a new pro-technology senior leader. With this fresh perspective, employees brainstormed practical ways to save time and share knowledge effectively.

However, as ConstructionCo is at an early stage of digital maturity, it lacked a clear strategy for leveraging AI, and this meant that use on the ground was fragmented. Employees were trying AI tools independently and did not have many opportunities to share their learnings with others.

Challenge

The biggest challenge for ConstructionCo was that employees had different levels of interest and confidence in using AI.

Employees who did not use general-purpose generative AI tools much were hesitant and anxious about the risk of data breaches. They were curious about how AI could help them streamline processes and do advanced data analysis. But because there was no clear plan or structured training, they did not feel they had the skills to figure out how AI could help them with specific tasks.

'Super users', on the other hand, were confident and willing to experiment with AI tools. But their efforts were ad hoc, and there were no structured processes to enable them to share learning more widely.

Because of this, fewer employees used AI tools and people missed out on chances to learn from one another. Without clear rules, structured training or a way to share ideas, ConstructionCo risked getting stuck in ad hoc experimentation instead of making real progress.

What they did

As part of the action research process, IFOW reviewed internal documents, interviewed employees across corporate departments and co-designed a workshop with the working group. IFOW also looked at how the legal department assessed a general-purpose generative AI tool.

One of the documents reviewed was a Request For Proposal (RFP) to HR system vendors outlining the HR department's requirements by priority. This RFP set the tone for a thorough nine-month HR system procurement process. Requirements in the RFP were organised by the MoSCoW prioritisation technique: 'must-have', 'should-have', 'could-have' and 'won't-have'. Vendors were asked to show how their solution addressed each item in the list. This helped the HR department systematically think about how vendor proposals for the new HR system could help employees, managers and the HR department in their daily work.

ConstructionCo's working group co-designed a workshop with IFOW to discuss with employees ways to overcome the identified challenges. The working group agreed that it would be useful to identify ground rules and structured learning opportunities to enable effective AI use at work.

Meanwhile, the head of legal independently led a pilot to assess whether using a general-purpose AI assistant could reduce their use of external legal panels.

Outcomes

The co-designed workshop plan included live demonstrations, breakout discussions, and activities to map AI's impact on job design. While there wasn't time to run the workshop during the action research process, the workshop co-design meeting helped the working group surface gaps to address:

- AI training should not just be about how the general-purpose generative AI tool worked. It needed to give employees the confidence to try new things safely and learn where the AI tool can help them do their jobs better.
- The meeting also confirmed the working group's belief that a clear set of rules (governance) is essential. Employees need these safety guardrails to prevent confidential data being leaked through the AI tool.
- The working group realised that asking employees to help each other informally was not enough to teach everyone how to use the AI tool effectively. They suggested creating structured learning pathways for everyone and creating a group of 'AI champions' who can give expert help.
- The HR department's use of MoSCoW to prioritise requirements during the procurement process was highlighted as good practice.

Furthermore, the legal department's AI tool pilot uncovered significant time and cost savings. The time spent reviewing a document fell from around an hour to 15 minutes. This highlighted an opportunity for ConstructionCo to bring routine work back to its own legal department and save money. Besides making work faster, the pilot created space for open conversation on how employees in the legal department were using AI tools. This allowed ConstructionCo to formally check the AI tool's performance while listening to employees.

Learning Points

Rebuild trust through transparency.
Past technology failures can make employees anxious or sceptical. To move forward, lead with open communication. By involving employees early you ensure AI tools actually solve their problems, turning fear of the new into trust in the process.

Create safety guardrails for learning.
Employees want to try AI tools but are anxious about making a mistake and leaking data. To fix this, provide clear rules (governance) and step-by-step training. When employees know where the boundaries are, they feel empowered to experiment safely, and it helps the organisation stay compliant.

Move beyond asking a friend.
Informal learning only works if you know the right people, which can leave some employees behind. To spread AI skills across the organisation, offer structured learning pathways and create a formal 'AI champions' network. This breaks down silos and makes sure everyone has access to the same learning opportunities and experts.

Case Study 3: FinanceCo

Collaborative job design: a bank's holistic view of performance and learning needs

Profile

This case study focused on a bank, referred to here as FinanceCo. It looked at how a bank explored the impact of an AI coding assistant on software developer daily work and lived experience. The findings show how a people-centred approach to AI creates holistic management strategies to support employees through change.

This bank offers both digital and in-person banking services. Headquartered in the United Kingdom, it has a workforce of over 26,000, including a team of 370 people professionals. Employee engagement scores are consistently in the upper decile.

In terms of AI maturity, FinanceCo reported itself as ahead of the curve in the sector. It has implemented mechanisms to support responsible AI practices, such as establishing an AI strategy, council and centre of expertise. FinanceCo's AI strategy was underpinned by its responsible AI principles, which include accountability, contestability and redress, fairness, empowerment, and societal and environmental wellbeing.

Among the range of AI systems being procured, fine-tuned, or developed at FinanceCo, this case study focused on an AI coding assistant that was made available to their software developers. To support the action research process with IFOW, FinanceCo composed a working group. The working group comprised senior professionals from AI delivery and people teams, with the latter including specialisms in AI and people, skills, strategic workforce planning, and colleague enablement.

Operational context

FinanceCo's approach to implementing technology - including AI systems - was cautious and defensive. It followed clear governance frameworks aligned with its risk appetite. AI use cases were tracked and monitored end-to-end to ensure risk and compliance frameworks were maintained. FinanceCo's AI strategy was underpinned by its responsible AI principles, which included accountability, contestability and redress, fairness, empowerment, and societal and environmental wellbeing.

As part of its strategic vision for AI, FinanceCo committed a people and communications workstream to support its workforce through change. This workstream included tailored skill development and insights sharing across the organisations. FinanceCo's senior people professionals told IFOW that there was a strong focus on ensuring that AI was deployed responsibly to improve work processes. For example, employees were encouraged to engage critically with AI tools as 'thinking partners'.

Fostering a hybrid approach in using the AI coding assistant presented a 'fundamental shift' in the software developers' job design, according to the AI delivery senior leaders. Software developers were encouraged to use the assistant as a boilerplate code generator, so they could focus more on the overarching architecture of the code. They were also encouraged to use the AI coding assistant to brainstorm and test the best way to code a solution, which the leaders felt allowed 'people's creativity to shine'.

Challenge

A key challenge for FinanceCo was understanding how software developers' work may have changed because of using the AI coding assistant. The working group hoped that understanding how AI changed role content would help them apply the learnings from here to other AI implementation projects in the organisation. It would shape their approach to learning pathways and strategic workforce planning.

In particular, the working group defined six areas that required further attention to ensure responsible and efficient AI implementation:

- Understanding how AI tools were being used, perceived and experienced by FinanceCo's software developers in practice.
- Understanding how meaningful engagement and involvement of the software developers around increasing AI use at work could be facilitated.
- Understanding how new AI capabilities could have an impact on existing workflows.
- Understanding what senior leaders and managers need to be aware of and communicate to the workforce around AI adoption.
- Understanding how the opportunities and risks presented by AI on existing capabilities and plans for future capabilities could be evaluated.
- Understanding requirements for effective learning pathways.

Achieving good job design would require a cyclical and reflective approach, as one working group member said:

> *"There's constant adaptation. You're not just using a software tool that has some limits because the software is improving all the time, and your ability to understand how to get the most from it is improving all the time as well. So, your role and the way in which you interact with that AI is dynamic."*

As FinanceCo's people team was building the organisation's wider strategy for learning and skills development, it was essential that they understood the on-the-ground experience of software developer work. This would enable the people team to develop learning pathways that were rooted in the realities of their workforce and ensure that AI implementation led to positive outcomes. This final point was highlighted as being especially important as working group members reported that the clear benefits of AI across the workforce were anecdotal.

What they did

To capture how software developers experienced changes to their work alongside the AI

coding assistant, the working group co-designed an in-person workshop session with IFOW. The workshop was attended by seven software developers of different seniorities and one senior people professional. The workshop was facilitated by researchers from IFOW with the aim that participants would feel comfortable sharing their thoughts and experiences with individuals external to FinanceCo. It offered an inclusive space for software developers of varying seniorities to share their opinions and experiences of how their work was potentially shifting with the introduction of the AI-powered coding assistant.

Workshop approach

The structure of the workshop itself mimicked a condensed version of the '6Rs Framework' by taking participants through the stages of 'Reveal', 'Reflect' and 'Reimagine'. It started with an icebreaker activity to surface what participants thought a responsible approach to AI entailed. Participants highlighted requirements around data privacy and security, human oversight and robust governance in addition to well-rounded training opportunities, informed consultation with employees and pro-worker employment protections.

The next section of the workshop was focused on gaining a baseline understanding of how the AI-powered coding assistant was currently being used and experienced by the software developers. First, participants were asked to pinpoint where they feel they currently stand along a learning journey, spanning from 'awareness' of the AI-powered coding assistant to 'confidence' in using it. Initial responses from senior software developers highlighted that their use of the AI-powered coding assistant was limited, preferring to use other AI tools instead.

They also reported that the stages of the learning journey set out were too broad to capture the range of use, which could vary in frequency or intensity. This highlighted an area of improvement that could be addressed in FinanceCo's future AI learning pathway to ensure it is responsive to the processes and workflows inherent to this high-skilled role.

After that, FinanceCo's roadmap of rolling out the AI coding assistant was presented by a senior lead, providing wider context to the organisation's adoption journey that had not been previously shared with the software developers. This, along with creating space for open discussion, prompted critical questions around performance measurement, potential efficiency traps and required changes to learning and development.

For the next segment of the workshop, IFOW researchers facilitated interactive breakout group discussions that homed in on how the software developers' specific work characteristics are being positively or negatively impacted by the uptake of the AI-powered coding assistant. Job design discussion points can be summarised as:

Job autonomy and control

1. Positive impacts:

More experienced software developers are able to take more significant decisions around the architecture and design of the code at a higher level.

2. Negative impacts:

Although junior software developers have access to greater information via the tool, their understanding of the wider context is reduced, creating a discrepancy between understanding the problem, but not how a solution is reached.

Increased expectation that junior software developers should be able to 'just do it' using the tool and with less access to managerial support that would provide greater learning.

Task variety and significance

1. Positive impacts:

Efficiency in task completion frees up time to dedicate to higher-value tasks required by customers.

2. Negative impacts:

Awareness of how employees' tasks fit within broader workflows is being lost. This is exacerbated by the increase in reactive tasks that allow for less agency, such as simply accepting the outputs generated by the AI tool.

Responsibilities could get diffused with increased difficulties to retain accountability if a software developer does not understand what they have built due to over-reliance on the AI tool.

Skill variety and use

1. Positive impacts:

AI tools enhance or create learning opportunities in different skill areas.

2. Negative impacts:

Impacts to critical thinking skills as the tool is designed to appease users rather than challenge them, leading to further over-reliance.

Creativity is reduced to writing or ideating a prompt for the tool, leading to potential difficulties in standing out as a high performer as employees tend to use the same tools in the same way.

Job feedback

1. Positive impacts:

More varied feedback on work conducted by employees is offered as the tool provides a 'second pair of eyes' on outputs.

2. Negative impacts:

Potential disruption of performance measurement metrics could lead to prioritising quantity of work rather than quality.

Performance feedback may be less accurate if individual workers take credit for code produced by the tool, which could also further impact job satisfaction or significance.

Social and relational

1. Positive impacts:

Managers viewed performance feedback that could be relayed via the tool as potentially less biased than what they would deliver.

2. Negative impacts:

Junior software developers reported a reduction in interactions with their team; they are told to find answers from the AI tool.

Peer-to-peer learning could decrease significantly as software developers rely on the tool for information and collaboration.

Job demands

1. Positive impacts:

May potentially lead to less admin-based tasks.

2. Negative impacts:

Junior software developers felt increased pressure to deliver work faster with the assistance of the tool.

Senior developers expected that their code review workload that is produced from junior developers would increase as tool leads to greater mistakes.

Workshop outputs

For the final segment of the workshop, participants focused on how work could be designed to maximise positive, people-centred outcomes by mind mapping a range of practical solutions, interventions or concrete actions that could be taken by FinanceCo. The suggested co-developed solutions by participants built around five key themes.

1. Meaningful learning and training

Moving beyond training exercises based on 'box ticking' and training offered by third-party technology providers; the software developers reported that they felt that these do not deliver value.

Increase software developer-led content and training pathways that include hands-on learning and impactful demonstrations.

Foster independence to allow deep dives into more complex learning, if desired.

Widen L&D focus on how to use certain AI tools as well as how to strengthen in broader aspects of the role.

Ensure the 10% L&D time included in FinanceCo's employment contracts is upheld and encouraged to be used.

2. Improved organisation-wide information and communication flows

Develop a clearer organisational stance on AI tools and how these interact with potential changes to work, with a focus on supporting employees' discretion over their work.

Strengthen communication flows between the organisation's innovation leaders and those who are 'on-the-ground' to ensure proper feedback around specific uses and experiences of emerging technologies.

3. Junior-senior feedback loops

Senior managers should engage more with juniors to understand how work processes are being enriched rather than focusing exclusively on task efficiency.

Seniors should facilitate meaningful learning and mentorship to allow juniors to understand wider context of code design and architecture.

Management to clarify expectations around use of AI tools as access to the tool does not necessarily mean that it will be used appropriately.

Renew 180 reviews or similar to allow juniors to provide constructive feedback to their seniors.

4. Reconsideration of metrics

Review how value creation is understood and measured in the organisation, particularly in relation to how the increased use of AI tools is impacting these metrics.

Ensure metrics capture quality of work over merely speed or quantity of work completed (eg, current metrics may include AI uptake levels but do not capture whether meaningful value was generated for the customer).

Consider integration of qualitative metrics alongside quantitative metrics to provide a more indicative picture of changes to performance, both organisational and individual.

5. Updates to job and role descriptions

Examine whether certain job and role descriptions need to be updated to reflect realities of the work (eg, writing code may have reduced, whilst understanding of systems structure and engineering of the code should receive greater focus).

Embed flexibility should be embedded and growth in job and role descriptions to prevent lock-in to a current list of skills or tasks.

The software developers discussed whether the increased rollout of the AI-powered coding assistant could lead to the introduction of new metrics for quantifying completed tasks that would then inform how employees' performance would be determined. They said this approach could result in performance metrics emphasising the quantity of tasks completed (eg, lines of code written) rather than the quality of outputs. This would also fail to capture the nature of senior developers' roles that involve less code writing, with more time dedicated to fixing code and mentoring direct reports instead.

The concern around new metrics shifting understandings of value creation at work is closely connected with the second concern around potential 'efficiency traps' being created by third-party technology providers promoting their products. Despite the high costs involved in purchasing the AI-powered coding assistant, no framework for measuring changes in productivity generated by the tool was provided by the technology provider.

This could lead to the organisation being locked into an artificial metric that is not truly indicative of productivity gains but is nevertheless utilised to justify the displacement of employees to achieve the return on investment that was initially promised by the technology provider.

In addition, the software developers discussed how learning and development needs to be approached more holistically with improved employee-manager support structures that allow for proactive coaching and mentorship.

This was highlighted as being especially pertinent for new starters and software developers in entry-level roles that are seeking to build their expertise but may not have the know-how to challenge the outputs generated by the AI-powered coding assistant. This could lead to over-reliance on the tool without fostering the necessary foundation to understand the underlying reasons behind and wider context of specific coding issues that may arise.

The software developers underlined how this foundational knowledge is especially important as changes to the required tasks and skills of the role means that understanding the infrastructure and design of the code is becoming more essential than knowing how to write the code itself.

The conversations facilitated in the workshop expanded beyond the AI tool itself to illuminate considerations at the levels of individuals, teams and the wider organisation that the software developers highlighted as being important to shaping positive experiences of work. The value of this people-centred approach is further demonstrated by one working group member who shared: "This is probably the best conversation on AI I've had in a long time."

Outcomes

Spearheaded by workers in their people team, FinanceCo's working group resulted in defining two critical programmes of work to be driven forward for ensuring a people-centred approach to AI adoption: fostering social learning and bolstering management capabilities. These are both areas that FinanceCo's people team already had capability to support and deliver on.

1. Fostering social learning

- Insights from the workshop revealed that the software developers, particularly those in more junior roles, were asking for a more holistic approach to learning that expands beyond the training materials provided by third-party technology providers. Feedback from the software developers included:
- Increasing dedicated face-to-face coaching with more experienced developers, providing them with deeper contextual awareness and criticality around challenges with coding architecture and design.
- Establishing a supportive peer culture for social learning in technical teams. This could include supporting communities of practice within developer teams to facilitate a co-learning culture. Additionally, team spaces for discussion, participation and feedback to support the integration of the AI-enabled coding assistant were highlighted in order to give teams the autonomy to pinpoint value-add tasks and learning pathways for achieving them.
- Establishing an AI champions network to support social learning engagement and collective decision-making around emerging AI capabilities.
- Integrating the AI champions model within developer teams to ensure that AI adoption practices are effectively aligned and integrated with learning and career pathways.

2. Bolstering management capabilities

- The working group identified the need for greater support for managers to ensure consistency of AI adoption practices and drive a wider cultural shift for high-involvement management practices. Feedback from the software developers included:
- Re-examining existing leadership programmes at FinanceCo, with particular emphasis on developer teams that have had low manager capability scores in past engagement surveys, strengthening the coaching skills of managers to support the change in software developers' roles from coding to greater collaboration with peers to review code.
- Managers to support the development of their team's communication, problem-solving and collaboration skills as requirements for effective peer-to-peer review will increase.

Learning Points

People teams may already have the relevant skills and experience to provide immediate value-add and improve return on investment on AI projects.

Job design can serve as an invaluable lens for understanding how new AI-enabled technologies are impacting the experience of work both at the individual- and team-level.

The deployment of an AI-enabled technology at work should be underpinned by a strategy that includes holistic learning approaches that extend beyond standard training provided by third parties.
Importantly, senior management and manager capabilities must be strengthened to enable a positive culture of peer-to-peer support and social learning across the organisation.

Thoughtful consideration of revised productivity measures are required.
Measuring changes in productivity or efficiency as a result of AI adoption is not a straightforward task. Meaningfully engaging directly with the workforce can be a powerful way of discerning potential shifts in actual value-add activities, avoiding lock-in to superficial metrics.

Case Study 4: HealthTrust

Engaging staff: an NHS trust plans to roll out an HR chatbot

Profile

This case study focused on an NHS Foundation Trust in England, referred to here as HealthTrust. HealthTrust employed over 5,000 people, including 80 in the HR and organisational development (OD) team. It looked at how an NHS Foundation Trust engaged with staff on plans to deploy an AI-enabled HR chatbot.

A working group of 15 HR and OD professionals was formed to support the action research process. The working group was led by the senior workforce delivery professional who led on the procurement of the HR chatbot.

Operational context

HealthTrust had an AI use policy in place as well as a cross-functional AI steering group. The people in the AI steering group came from a range of departments, including digital, information and medical ethics. The AI steering group's role was to make sure any AI solutions they picked solved real problems without overspending.

HealthTrust's plan to roll out an HR chatbot fits with the UK government's 10 Year Health Plan for England plan to "liberate staff from their current burden of bureaucracy and administration". AI and other technologies will enable the NHS to "personalise care, improve outcomes, increase productivity and boost economic growth". "There will be fewer staff in the NHS in 2035 … [but] those staff will be better treated, have better training, [have] more exciting roles … and so they will each achieve much more".

The plan to roll out came after a study to make sure the HR chatbot was the right fit for Health Trust. The team already knew that it could work from a similar past project. Staff were happy with the quick and reliable answers from another virtual assistant that was deployed.

Challenge

HealthTrust's HR and OD team was stretched thin by thousands of repetitive questions every month and the numbers were growing. Staff mainly asked about annual leave, occupational health appointments, policy details, and recruitment processes. As HealthTrust operated round the clock, these enquiries would come in at any time.

Although staff could access help guides and resources through HealthTrust's intranet, they preferred to contact the HR and OD team directly by email or phone. But the HR and OD team had limited capacity to respond immediately, especially at night and at weekends.

What they did

To lower the workload, HealthTrust decided to find an HR chatbot that could answer routine questions automatically. HealthTrust teamed up with other regional NHS Trusts to buy the HR chatbot together to get a better deal and save money.

A four-part roadmap was created so that each trust could choose to launch the HR chatbot in stages. At first, the HR chatbot would simply answer basic questions about annual leave and policies. Eventually, it would be able to handle more advanced tasks like pulling data from other systems automatically and guiding staff to complete forms.

During the procurement process, HealthTrust knew they also needed to understand how the HR chatbot would change the team's daily work. In interviews, working group members raised big questions on what staff would do with the free time saved by the HR chatbot, as well as what new skills they'd need to learn. They wanted to make sure no one was left behind.

Through the action research with IFOW, HealthTrust captured how the HR and OD team felt about the new HR chatbot and how it might affect their daily work. HealthTrust also met with different staff to discuss how the chatbot would change the way the team runs.

Senior managers that IFOW interviewed emphasised the importance of being honest with staff about new technology. Information needs to be shared, and the benefits clearly communicated to help with any worries. Interviews with non-frontline staff in the operations, HR, and OD teams revealed mixed feelings. While some were eager to see how the HR chatbot could help them, others were sceptical about whether it would work in a complex organisation like HealthTrust.

Findings from IFOW's survey highlighted the HR and OD team's top priorities for bringing in new technology responsibly:

- Recognise when AI shouldn't be used, for example, in highly sensitive or complex situations.
- Make sure all staff are aware of the AI tools used by the HR and OD team.
- Address concerns about job security early by having meaningful staff consultation. AI shouldn't be used just to cut jobs. Instead, any money saved should be put back into training the team with new skills.
- There needs to be robust data governance.
- Provide comprehensive training on using new technologies. 64% of respondents reported not receiving any relevant staff-led or self-led training.

Outcomes

After speaking with 21 staff from the HR and OD team, HealthTrust identified three ways to better prepare staff for the upcoming shift to using the HR chatbot.

Firstly, there needs to be better training on digital tools in general across HealthTrust. Even though the HR chatbot will work through familiar channels like the phone, many people still struggle with basic digital skills. Improving these basics is the first step toward helping staff understand and use AI confidently. The HR and OD team noted that HealthTrust needs to map out everyone's current skills, and that the digital education and people development teams should work more closely together to spot any skills gaps. This teamwork will make sure HealthTrust has the right skills to meet its main digital goals.

Secondly, the team emphasised the need to plan for changes in their daily work. As the HR chatbot takes over simple queries, staff will shift toward checking and improving its answers. This raises some big questions. Will staff still feel as connected to their work if they spend less time talking directly to colleagues? There is also a worry that junior staff might lose out on learning by doing if the chatbot takes over entry-level tasks. To stay on top of this, the Trust plans to use staff surveys more often to check how staff are adapting to these changes and make sure they still feel supported.

Finally, the team highlighted that success depends on clear and constant communication. Introducing new technology must be people-centred and not just about the software.

As one member of staff said: “This [chatbot] isn’t necessarily a pure IT project...Actually, that’s the main reason why AI and digital projects fail … because of the people side … Often the bulk of the resource goes into the tech side and less on the people side. But we really want to make this work, … we could be doing some work now to start to get the managers of our admin teams ready.” Staying open and sharing information builds trust, and it allows managers to address any worries or fears staff might have.

Learning Points

Build AI literacy and confidence early.
For an HR chatbot to work well, the team needs to have the skills to oversee it. Staff must feel confident checking the HR chatbot’s answers and fixing any mistakes to ensure the information is accurate.

Watch how jobs change.
As the HR chatbot is rolled out, keep a close eye on how daily tasks are shifting. A people-centred approach means regularly checking in on how staff feel about their new responsibilities and what new skills they might need.

Listen before you implement.
The time before the HR chatbot goes live is the most important. Use this phase to listen to a wide range of staff voices. Their views will help you build a solid plan for training, communication and other support that works for staff.

Case Study 5: HolidayCo

Cross-functional deliberation: a holiday retailer's journey towards responsible AI use

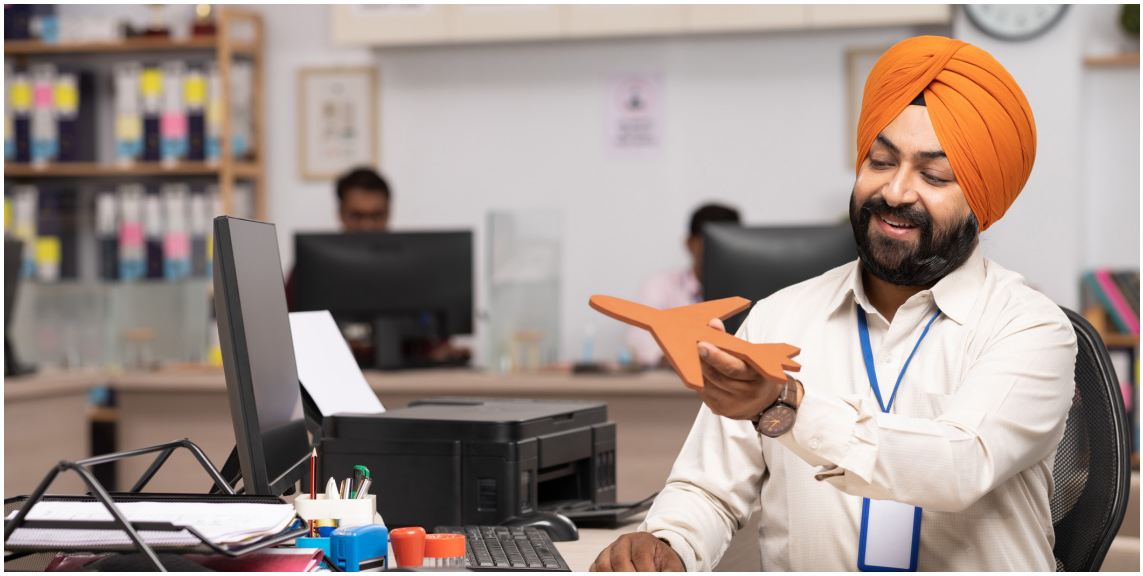

Profile

Headquartered in the United Kingdom, this organisation, referred to here as HolidayCo, specialises in providing package holidays and employs over 500 people.

This case study examined how the firm navigated the implementation of generative AI tools across the organisation and took steps towards tackling siloed working.

HolidayCo's centralised HR team comprises 15 people and covers the remits of culture and engagement, operations, and talent acquisition. HolidayCo has rolled out several AI tools across the business, some of which are targeted to specific functions and others that are general-purpose generative AI tools available to all employees. This case study focuses on the latter, which consist of three off-the-shelf generative AI tools that produce outputs such as text, image or code.

For this case study, HolidayCo composed a working group of senior professionals from people, legal, platform engineering, and digital transformation teams. By collaborating with the working group, this research project was able to surface the key challenges the organisation was facing around AI implementation, co-develop possible solutions, and home in on refining a specific intervention.

Operational context

At the organisational level, the use of generative AI is guided by HolidayCo's 'AI Use' policy. The policy was drafted in March 2025 in response to the proliferation of generative AI tools readily available to employees. In-depth interviews and a co-designed workshop revealed, however, that individuals and teams across the business lacked cohesive awareness and understanding of the capabilities of specific AI tools, how they are currently being used by different teams and how they may be best utilised to offer potential improvements to work processes.

As such, members of the organisation's working group, as well as workshop participants, described HolidayCo's current approach to AI enablement as being patchy, with big differences in attitudes and practices around AI adoption across teams. The organisation

currently does not have an overarching AI strategy and has enabled immediate access to the three off-the-shelf generative AI tools for the entire workforce, as opposed to a phased or piloting approach. One of the main reasons for this immediate rollout was to minimise the 'shadow' use of AI tools by employees, who were using personal generative AI accounts to perform work tasks, posing risks to company data.

Overall, the working group was positive about the perceived benefits of AI tools, even though no formal evaluation or cost-benefit analysis had been conducted. The overarching reasons for adopting the AI tools were largely attributed to ease of integration, clarity of data privacy and security policies, and the influence of external industry narratives.

AI governance at HolidayCo is largely guided by a data privacy and security risk management approach. Although there is an interest in expanding this to support wider people-centred outcomes within the context of AI adoption, such as enhanced literacy, learning and attention to work design changes, there is little clarity on how to achieve this. The people team isn't currently involved in decision-making processes around AI adoption.

Challenge

Keeping pace with advancements in AI technologies, whilst pursuing a clear and unified approach to AI adoption within HolidayCo, has surfaced challenges. The working group repeatedly referred to the sense of speed that seemed to dominate AI development, characterising the landscape as the 'wild west' or as a 'minefield'. Within this context, the working group acknowledged that there was an increase in anxiety and concern from the workforce around what the rise of AI tools may mean for skills development and job security.

Some members of the working group also expressed uneasiness around how increased AI usage at work may affect human connection, with fears around employee interactions being reduced or replaced with 'AI-to-AI' in the future. The need to hear directly from the workforce to understand their current perceptions, attitudes, and experiences of AI adoption at HolidayCo was underlined, as this insight was lacking.

Underpinning these challenges for HolidayCo was the lack of an organisational programme of work to drive AI strategy and governance activities in a clear and consistent way. Efforts to shape strategic decisions around AI adoption were siloed, leaving business leaders with information gaps in how these activities were being formulated and implemented. Moreover, current AI governance mechanisms focus narrowly on data privacy and security risks, lacking direction on how other relevant governance concerns would be approached. The working group highlighted the need for taking a more structured approach to AI governance across the business, but was unclear on how to tackle this.

As part of this ambition, the working group stressed the importance of fostering cross-functional collaboration and learning, both to strengthen workforce engagement and to cultivate innovation around pressing business challenges. In response to the prevailing information gaps, the working group underscored the importance of strengthening organisational communication channels and developing a coherent narrative around AI adoption at HolidayCo. Without a structured programme to bridge these silos, HolidayCo had no clear direction on how to manage the broader opportunities or risks.

What they did

In response to these challenges, the working group identified a series of areas that would need to be addressed. This included the lack of foundational learning and training on new generative AI tools and functionalities for the workforce to be able to integrate practices that would augment their current workflows. Ongoing avenues for coaching and experimentation were also identified as blockers. The working group felt these could potentially be addressed in the short- to medium-term through new organisational initiatives such as hackathons,

competitions or a champions network. To verify the relevance and value of these envisioned activities for the workforce, the working group agreed to break silos across functions by listening to stakeholders across the organisation.

The working group co-designed an in-person workshop with IFOW researchers. Fourteen participants of different seniorities attended from the people, finance, legal, operations, platform engineering, content, and digital transformation teams. It was the first time HolidayCo had brought colleagues across a range of seniorities and teams to explore the implications of generative AI tools at work.

The structure of the workshop itself mimicked a condensed version of the 6Rs Framework by taking participants through the stages of 'Reveal', 'Reflect', 'Reimagine', and 'Realise'. To set the scene, two lightning demos were held by HolidayCo colleagues to share how AI was currently being used to address business challenges. This opportunity for peer-to-peer learning that was tailored to the HolidayCo context allowed colleagues to gain insight into how other functions across the organisation navigate the opportunities and challenges presented by certain AI tools.

Through the workshop, employees participated in activities to identify their goals, enablers and blockers, and possible short- to medium-term solutions around AI.

Workshop Goals

IFOW researchers facilitated an activity around articulating specific goals around AI adoption at HolidayCo, with a focus on understanding how these goals intersect with employees' work, the work of their teams, and wider function. The value of this exercise was that HolidayCo colleagues were able to clearly pin down what mattered most to them at work and reflect on whether the introduction of AI tools would enhance or diminish their overarching goals. The workshop participants set out the following goals, which can be grouped into six broad themes:

Colleague support

- Ensuring positive working relationships with colleagues are supported, preventing human interactions from being reduced or compromised from increased AI usage.
- Increasing support for colleagues through coaching.

Creative and fulfilling work

- Ensuring that employees' expertise, skills and experience are utilised effectively.
- Enabling opportunities for greater creativity and autonomy at work.

Knowledge sharing

- Supporting testing, learning and knowledge exchange across the organisation.
- Ensuring organisational-wide information sharing and communication is robust and consistent.

Worker protections

- Ensuring employees feel valued for their work.
- Maintaining worker protections and due process.
- Offering opportunities for reskilling and upskilling.

Efficiency

- Reducing time spent on repetitive, admin-based tasks.
- Speeding up production of analyses and insights.
- Ensuring efficiency gains align with high-quality work outputs.

Governance

Ensuring appropriate safeguards for data privacy and security.

Ensuring interoperability of AI tools with the organisation's wider IT system.

Enablers and blockers

After that, the participants focused on identifying the range of enablers and blockers to their goals around AI adoption at HolidayCo.

Enablers

Participants felt the organisation's positive and encouraging attitude toward technological transformation encouraged innovation in AI adoption.

It was felt that the organisation is generally supportive of its workforce and has a people-centred culture, which supported their goals around AI.

Participants highlighted the organisation's robust frameworks for safe and secure IT use, putting them in a good place to approach AI adoption.

Blockers

Low levels of knowledge and literacy around use of generative AI tools across teams was highlighted as a blocker to AI adoption.

Flow of communication and information was restricted across the business, hindering individuals and teams in participating actively in key decisions around AI adoption.

Non-managerial employees highlighted general uncertainty and fear around long-term impacts of AI adoption on job quality and job security.

Possible solutions

Finally, participants worked together to mind-map practical solutions to the blockers they had previously identified, paying close attention to the potential impact and feasibility of their chosen solutions. A total of 11 practical solutions were raised, ranging from drop-in sessions, lightning demos, and a champions network to strengthening HolidayCo's existing AI working groups approach to implementing shadowing or role swap opportunities.

Following the co-designed workshop, the working group reviewed the range of solutions that were developed by participants. They prioritised the establishment of a more structured and joined-up approach to AI adoption at HolidayCo that would tie together workstreams around training and development, change management, information and communication flows, and inclusive participation of the workforce in the organisation's AI roadmap.

The working group identified that the most effective way of achieving this would be by evolving HolidayCo's existing AI working group. This existing group was initially established as a response to the increasing use of AI tools in the workplace and offered an informal forum for technical employees to explore different AI applications. As such, it was run in an ad-hoc fashion, lacked strategic priorities, and did not formally feed into HolidayCo's wider AI Risk and Governance Committee.

Outcomes

By convening colleagues from a range of seniorities and functions to reflect and discuss on the trajectory of AI implementation at their organisation, HolidayCo made progress towards driving cross-functional collaboration on key issues of technological adoption.

The co-designed workshop shed light on the wider considerations of work and organisational design that need to be in place to enable readiness, expanding beyond the narrow focus of solely technical access and functionality. This is further demonstrated by one working

group member who shared: “I think this process has hugely helped us cross-functionally collaborate, progress the agenda together, and knowing it’s not just a tech agenda, I feel like we’ve done so much just through getting together and talking about this, that we may not have done before.”

By evolving and expanding the current AI working group, HolidayCo would be able to move towards developing a comprehensive AI adoption roadmap for the organisation. The working group indicated that this refreshed approach would aim to encompass the voices of cross-functional representatives, offer spaces for training and peer-to-peer learning, unify the business’ internal communications around change management, and gradually open up mechanisms for understanding how the integration of AI tools may potentially be changing the nature of work performed by employees.

In underlining the importance of this intervention, one of the working group members stated: “There’s nothing that’s aligned at the moment. We might have lots of different workstreams running separately, but no way of linking them, no way of thinking about the change, the impacts, the risks. And almost measuring the success of all the things that we’re doing.”

In practice, it was agreed that the renewal of the AI working group would require the establishment of core members that would feed actions into the AI Risk and Governance Committee, the creation of clear objectives and key results for the group to deliver, and the creation of three main workstreams that would tailor content and activities around AI adoption for employees across varying competency levels, from basic, intermediate, to advanced. It was envisioned that each of the three workstreams would be supported by its own community of practice through the establishment of a champions network that would be employee-led and offer opportunities for personal development.

Learning Points

Organisational plans to adopt AI systems and tools need to take a holistic approach to understanding advantages and disadvantages.
These need to move beyond solely data privacy and security concerns to encompass improvements to governance, learning and internal communications programmes across the business.

Decision-making around AI should strive for cross-functional representation, including people professionals.
This will ensure alignment across the organisation and the integration of multiple perspectives and sets of expertise.

Cultivate interactive spaces for varying seniorities of the workforce to meaningfully inform and shape decisions around AI adoption within the organisation.
This important for building trust and fostering socially licensed innovation.

A systemic process for continuous reflection around AI adoption, such as the 6Rs Framework, offers a structured yet adaptable way of diagnosing, developing and evaluating organisational- and individual-level changes.

Case Study 6: InnovationCo

Building the AI groundwork: an innovation company looks to modernise a safety-critical and unionised sector

Profile

This case study focuses on a public innovation hub with fewer than 50 people, referred to here as InnovationCo, working inside a larger public body. It describes how a public innovation company introduced an AI tool for maintenance planning.

InnovationCo's mission is to use technology to modernise a sector that has been traditionally slow to change, partly due to the high importance given to safety and the need to ensure that changes do not put people at risk. Without an internal HR team, InnovationCo relies on the parent organisation's central HR for people support.

To support the action research, a working group of mostly engineers and an ergonomics specialist was formed. Their goal was to expand the use of cameras and AI video reviews for maintenance planning in a specific region, in line with national policy. While HR was not part of the working group, IFOW spoke with a national HR lead to gain broader insights into the parent organisation's culture and constraints.

Operational context

InnovationCo operates in a safety-critical, hierarchical sector that has a strong union presence. Small changes to processes or how people work could have major ripple effects on safety standards and management practices, and trigger formal engagement with the trade union. These factors created an uncertain environment for innovation. As AI use and its impacts grew, many employees experienced high levels of anxiety.

Challenge

The main hurdle was the parent organisation's fragmented structure and outdated processes. Because the project was primarily built for technical safety, there was no space for the team to formally flag – early on – how AI would change their daily work. As a result, complex problems went unaddressed.

This created a ripple effect of challenges:

- Fragmented decision-making. Without a clear way to 'own' decisions, teams worked

in silos. For example, the parent organisation's national HR lead saw the risk of 'tissue rejection' from the workforce and recommended early, frequent engagement with union colleagues. However, the national HR lead was not involved in the pilot phase.

- Outsourced strategy. In the absence of clear internal governance, the technology provider began to steer the project. This risked aligning the AI tool with the technology provider's goals rather than InnovationCo's mission or the needs of the engineers.
- Operational paralysis. Teams felt stuck. They wavered between focusing on their own tasks and worrying about how broader organisational changes would be perceived. Information wasn't shared across regions, which stifled any chance of peer learning or a shared vision.

Ultimately, the lack of coordinated oversight meant that the AI maintenance planner was being imposed on employees, rather than being designed deliberately to augment their skills. The team could see that redesigning workflows would be a helpful way to break through these organisational hurdles. But this remained a thought exercise rather than a practical shift because they felt they lacked the authority to change them.

What they did

To help InnovationCo navigate these hurdles, IFOW led a series of interviews and workshops to create space for dialogue. The goal was to move beyond the technical and surface human opportunities and risks.

Through interviews and observing a safety workshop and team meetings, IFOW identified that while the lead engineers were highly engaged, other critical functions such as HR remained at a distance. This confirmed that the challenge wasn't just technical. The deeply rooted organisational hierarchy made cross-functional collaboration difficult.

Engineers, parent organisation leads, and the technology provider came together to do the three targeted activities at a workshop run by IFOW:

System mapping

Participants mapped out how information flowed. This helped everyone see the system and identify who really needed to be in the room for decisions on the AI tool. It also gave them a chance to discuss openly the challenges they perceived.

Augmentation journey

Participants looked at the evolution of the AI tool, from an individual engineer's resource for fault detection toward a coordinated system for planning preventative maintenance. They discussed the ways an engineer's skills and judgment could be augmented over time. This helped identify decisions on how the AI tool would impact work and workflows.

Defining success

Participants began identifying measurable indicators of success to capture the people impact alongside technical ones.

After the workshop, IFOW met with the working group to translate the workshop themes into actionable insights. The working group realised that a single regional use case could not be viewed in isolation, so the discussion shifted to process and governance gaps at national level:

Gap in responsibility

Current innovation processes didn't allow for a systematic, cross-functional assessment

of the AI tool. This meant that the people impact was a major strategic risk, but the responsibility for assessing them was left to individual project teams.

Risk of escalation

Given the sector's strong trade union presence, leaving the people impact unaddressed at organisational level risked brewing future conflict. If significant risks around changing responsibilities in safety-critical areas surfaced too late, disputes could escalate.

Barriers to innovation

Outdated processes and rigid efficiency targets worked against the project because they didn't capture social or secondary impacts. Instead, they drove up a sense of pace and anxiety that restricted the team from reflecting on how best to use the time freed up by the AI tool.

The value of job design

Considering the job design of those affected by the AI tool should be built into the procurement and pilot phases and revisited as the tool evolved. It shouldn't be a one-off exercise.

Outcomes

The collaboration between IFOW and InnovationCo elevated the project from a technical implementation to a broader, strategic organisational evolution. The action research process surfaced four key outcomes:

Secured leadership buy-in

Senior management has been engaged to trigger and require a formal change management process. They recognised that these systemic challenges must be addressed at a national level, rather than being left to local teams.

Proactive union engagement

The working group committed to bring trade union representatives into the conversation before formal consultation begins. By establishing joint working groups in neglected areas like job design, human control over decisions and procurement, the organisation is moving toward a more transparent, less conflict-prone relationship.

A strategic pause for better design

Internal innovation processes were intentionally paused to allow for a deeper look at job design and workflows. This led to a proposal for a new work design panel within the ergonomics function so that the people impact was mapped properly before moving forward.

New holistic metrics

InnovationCo is developing an augmentation policy and reviewing its efficiency metrics. New success indicators have been set to move beyond simple technical output, which includes better information sharing and tracking the people impact of AI.

This progress provided a vital starting point. To ensure a responsible and effective rollout, the senior leadership must translate these early wins into a long-term strategy that puts job design at the heart of the project's success.

Learning Points

Prioritise cross-functional engagement.
Integrating AI is not a purely technical project. Early involvement from HR and operational teams avoids fragmented decision-making and ensures systemic impacts are understood across the entire organisation.

Formalise change management to break silos.
Ad-hoc testing without a clear roadmap risks reinforcing information silos and limiting what the organisation can learn. To ensure a responsible rollout, the organisation must establish formal governance and clear channels for sharing knowledge.

Take ownership of job design.
While technology providers engage with end-users, their commercial interests don't guarantee a focus on people impacts. Organisations must take the lead in assessing how AI will change work and augment people's skills and judgment.

Set strategy early to prevent drift.
Strong internal governance prevents external interests from taking the lead. Setting clear objectives from the start ensures the technology remains a tool for the organisation's priorities, not the other way round.

Case Study 7: LegalCo

A tech-positive law firm introduces an AI legal assistant

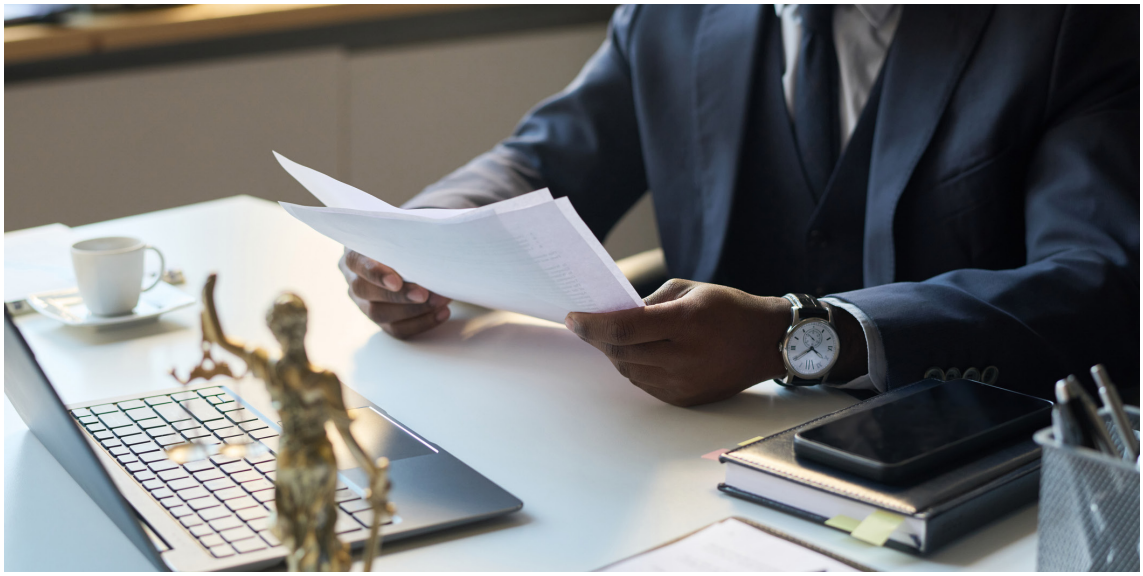

Profile

This case study focused on an international law firm headquartered in the UK, referred to here as LegalCo. It looks at how a they introduced an AI legal assistant for their employees. At the time of IFOW's investigation, employees in legal and business functions had access to the organisation's AI legal assistant for a year.

Operational context

LegalCo had consistently shown a strong pro-technology culture compared to peers by being open to trialling emerging technologies, often initiated by their own lawyers. AI implementation was viewed as part of a broader strategy to explore legal technologies to support both lawyers and clients. At LegalCo, lawyers work closely with the technology team to identify, test and implement technology solutions. Employees said their AI legal assistant was 'more sophisticated' and 'better suited for legal work' than tools used in other firms.

The AI legal assistant was implemented as part of a top-down directive from senior leadership. Senior leadership addressed concerns about job displacement directly, and clearly stated AI would not be used to reduce headcount in roles that were not revenue-generating. This senior-level buy-in created a supportive environment for employees to experiment with the tool and share their learnings with each other. LegalCo's operational context created an environment where AI was positively encouraged, openly discussed, and aligned with professional and firm-level identity.

What they did

While LegalCo did not share a formal AI strategy with IFOW, the interviews showed there was a clear strategic direction, and practical measures to create a safe environment for experimentation. There was strong senior leadership consensus on the goals and outcomes they hoped to achieve from rolling out the AI legal assistant.

This top-down directive created early alignment and eased many of the usual challenges associated with learning and development, change management and internal communications. Internal communication was regular and in multiple formats. Concerns

about redundancy were directly addressed, and employees were encouraged to try the AI legal assistant, helping reduce anxiety and normalise experimentation across levels. This multi-pronged approach to AI adoption included:

- **Creating a cross-functional steering group which included senior leadership from across legal and business functions.** This steering group met fortnightly to decide features, rollout timing, internal communications and training approaches, and adapt plans as new insights emerged. This created a coherent strategy across the firm and wraparound support for employees.
- **Gated access to the AI legal assistant, where employees had to complete mandatory live and recorded training before receiving access.** Part of the training sessions were delivered in collaboration with the service provider.
- **LegalCo also worked closely with the service provider on feature feedback.** The relationship remained collaborative and led by LegalCo's internal strategy rather than the service provider's priorities. This empowered both senior managers and junior lawyers to experiment openly, reflect on how tasks were changing, and understand how AI was augmenting their roles.
- **Senior lawyers created a culture of experimentation by instructing junior lawyers to try applying the AI legal assistant to different stated use cases.** Senior lawyers also experimented with the AI legal assistant and openly shared their experiences with the team.
- **LegalCo's HR department regularly communicated tips and best practice to employees in weekly emails.** The HR team also developed a 'champions' network, innovation awards and featured 'spotlights'.
- **Drilling the mantra, 'verify, verify, verify' to all employees.** Human verification of AI outputs was non-negotiable.

Challenges surfaced

While LegalCo's approach to implementing the AI legal assistant was well executed, two challenges were identified.

Firstly, newly qualified and junior lawyers using the AI legal assistant gained increased autonomy, real-time feedback on their work, and support with drafting and reviewing documents. However, concerns were raised that professional learning and development may need to be adapted to ensure legal skills continue to be developed by junior lawyers. These were addressed informally through on-the-job manager feedback rather than changing formal training plans.

Secondly, lawyers said their workloads increased, and wellbeing was negatively impacted, because of responding to a higher volume of AI-drafted communications from opposing parties. This resulted in some cases becoming more contentious or parties becoming more entrenched in their position. It also removed a key part of a lawyer's role of engaging with real people.

Learning Points

Make AI adoption explicitly people-centred by pairing encouragement to experiment with clear non-negotiables (eg verification, strategy, red lines).
LegalCo encouraged experimentation by having senior lawyers model use of the AI legal assistant. Then adjust their instructions to junior lawyers to draft, test or refine their work with the AI legal assistant. This normalised day-to-day use. Alongside this, LegalCo set firm red lines to protect quality and authenticity, such as a non-negotiable requirement to 'verify, verify, verify' every AI output. This combination reassured staff that experimentation was expected and safe, while maintaining human judgement as central to legal work.

Steer adoption using a cross-functional team that includes senior leadership, as this supports alignment on rollout, training and internal communication.
LegalCo's cross-functional steering group comprised senior leaders from legal, business services, HR, technology and practice teams. The group met fortnightly to decide on features, training requirements, rollout timing and internal communications, ensuring that all functions moved in step. This alignment reduced fragmentation and meant that staff received consistent messaging, coordinated training, and had a unified sense of the purpose of the tool.

Protect authenticity by identifying domains where human voice, trust or credibility are crucial and communicate these clearly to employees.
LegalCo's lawyers consistently identified witness statements as a task where authenticity and human voice were essential. They emphasised that the AI legal assistant could help structure documents, but the words must remain "the voice of the witness," and overly polished AI-style drafting was viewed as undermining credibility. Several lawyers worried that documents that "sound like AI" lose the personal texture required in litigation. By designating witness statements and strategic advice as strictly human-first domains, LegalCo protected trust and preserved the relational elements of legal work.

Redesign learning and development pathways to ensure junior employees continue to build the skills necessary for service delivery as they progress through their professional careers.
LegalCo's senior lawyers observed that the AI legal assistant shifted junior work away from producing first drafts. The time saved enabled their junior lawyers to do more editing, verification and higher-order legal reasoning. While this accelerated exposure to complex tasks, it also poses a risk to junior lawyers' professional development by missing opportunities to develop core drafting and legal reasoning skills. Managers had already begun informally adapting feedback and supervision to counter this. For example, discussing legal reasoning more explicitly during review meetings and updating formal training. This demonstrated the need for structured pathways that safeguard foundational skill development as AI reshapes early-career work.

Case Study 8: LargeCo

Questioning return on investment: a B2B operations firm reflects on deprecating an AI tool

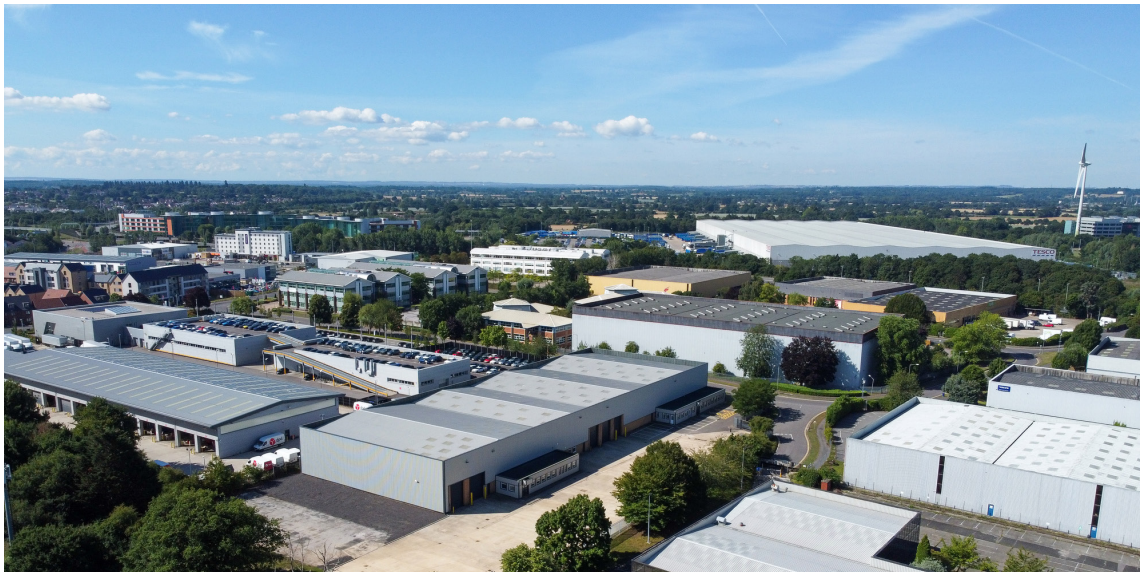

Profile

LargeCo has thousands of employees in corporate functions and frontline operations. Some HR activities are outsourced to a third-party provider offshore. This case study details the approach of an operations management company to AI adoption, specifically how a generative AI tool was introduced into the bid writing function

For this research, LargeCo assembled a cross-functional working group with senior leadership team members involved in various aspects of AI adoption across the organisation. This spanned enterprise technology, HR, people and sales functions.

Engagement with IFOW was initially led by an AI people lead. They described their role as engaging in work relating to 'people, capacities and skills' and 'bringing the responsible dimension' to business decisions about which tasks to automate. This individual left, and their role was subsumed into the remit of the chief of staff to the head of the people function. IFOW was not able to establish meaningful contact with the chief of staff. This case study is therefore observational.

Operational context

LargeCo operates business-to-business (B2B). They have a strong focus on risk mitigation with regard to cybersecurity and data privacy, and an impetus to communicate with clients that any AI usage was safe and compliant with relevant law and regulation.

In comparison to industry-level trends, AI adoption at LargeCo seemed fast-paced. IFOW identified at least 10 different AI tools in operation or under consideration across the organisation, mostly generative AI. LargeCo had begun to deploy AI agents to automate some administrative tasks.

The AI tools identified by IFOW were not necessarily reflective of AI adoption across the entire organisation. This was due to research limitations as well as an absence of comprehensive oversight within the company itself. During the research, the organisation was undertaking an audit to identify all AI tools and applications in use.

LargeCo's AI adoption is part of a broader digital transformation process. The fast-paced nature is connected to the perception that AI can be used to 'enable growth' and improve the capacity and capabilities of workers. A senior leader within the technology function noted:

> *"The way we improve our profit margin is by growing our top-level growth, but without increasing the operating cost of our business. And we see the way that we get there is through the efficiencies that AI-based tooling can have. So, we can double the size of our business without doubling the size of the team."*

In addition to driving efficiencies, members of the working group expressed the target of becoming industry-leading with regard to the use of 'secure' and 'trustworthy' AI systems.

To this end, LargeCo has developed an AI policy centring on principles including data privacy, transparency and fairness, and ethical oversight. A further iteration of the policy was in development when this research took place. LargeCo had also developed a risk-based framework for AI adoption that is aligned with the EU AI Act. This prohibits the adoption of certain AI applications, such as social credit scoring systems and emotion recognition for workplace hiring.

Decisions to adopt other applications–categorised as minimal, low or high risk–are taken by senior leaders situated across legal, digital, people, data privacy, business services and information protection functions who sit within a multi-level AI governance board (the board). High-risk use cases are added to a risk register and tracked by a separate team that reports into legal, and back up to the board. Representatives from the HR and learning and development are represented on the board in an advisory capacity.

Challenge

The cross-functional nature of the board was described as important to prevent the 'siloing of AI' from existing processes and ensure that impacts of AI adoption across the organisation were 'visible and understood'. The board was praised by those interviewed for this project. However, it was emphasised by a member of the working group that discussions relating to risks and impacts were skewed towards technical considerations at the expense of 'people impacts'.

Furthermore, although extensive efforts were said to be underway in the learning and development function to facilitate upskilling and prepare the workforce for impacts of AI adoption, it was acknowledged that a more proactive approach to considering the impacts of AI on job roles and responsibilities was required, along with role redesign and strategic workforce planning. This was informed by assumptions that AI would have a significant impact on jobs, as well as past experiences.

A member of the working group noted:

> *"So, we've had a couple [of automations] where we've gone live, and it has led to some redundancy consultations. What that is kind of informing is, as we start to look at the next area of focus, how do we kind of get ahead of the inevitable conversations that will happen so that we're thinking about the people impact earlier in the journey, because what we learned from the first couple of experiences was that it was too late in the journey that we started to talk to people about, actually, what does this mean for their job".*

In addition, organisational data viewed by IFOW highlighted a lack of trust among the workforce regarding the ethical and responsible use of AI by the organisation. This may be connected to the absence of mechanisms for employee participation in decisions about AI use and adoption, or a lack of transparency regarding strategy. The data viewed by IFOW simultaneously demonstrated that employees did not feel involved in decisions relating

to AI usage, with only a very small percentage aware of the organisational strategy. This is reflected in the structure of the board designed to govern AI use and adoption, which lacks any form of employee representation.

What they did

Due to the extent of AI use at LargeCo, the decision was taken by members of the working group to focus on a specific tool for the purpose of the research. This was a generative AI bid writing tool. The tool was procured from a third-party technology vendor based in the UK and introduced in the bid function in early 2024 via a phased rollout of licenses. The basing of the vendor in the UK was cited as a key reason for its selection, in comparison with other international providers.

The stated managerial intention of the rollout of the tool was threefold. To reduce time spent by employees to produce a first draft, to increase bid quality and to mitigate against the perceived risks of shadow AI usage by employees. The head of the bid function noted: "It was partly looking at an efficiency saving [...] but also that idea of having an AI solution that we could present to technologically savvy and kind of early adopter type bid people and say if you want to use AI, you can, but please use this one, don't use an arbitrary off-the-shelf open AI solution. Use this one because we've got all the kind of checks and balances in place. It's a secure environment. It draws from our own content."

The vendor provided training to employees and supported the bid function to deploy repeated surveys to assess employee attitudes to the tool. The survey addressed comfort with AI, confidence with the tool and time-savings, with only minimal reflections on impacts to the experience of work via questions relating to work-life balance and stress reduction. The questions were framed positively, limiting space for reflections on negative impacts. For example, employees were asked about their agreement with statements such as: "The [generative AI bid writing tool] reduces stress and makes my work more enjoyable", and "I would be disappointed if I could no longer use [the generative bid writing tool]." IFOW sought to deploy an additional survey with employees using the tool. This was optional for the organisation. The aim of the survey was to gain an understanding of employee attitudes towards the use of AI tools at work, the extent of information, and involvement in adoption and different aspects of current experiences of work, including levels of stimulation, mastery, autonomy, relationality and job demands.

Outcomes

LargeCo facilitated the deployment of the survey, which was voluntarily completed by less than a quarter of those with access to the tool. However, questions relating to current experiences of work were removed by the organisation, as they were deemed not 'relevant' to the AI tool or AI adoption at the organisation. This perception is suggestive of a broader absence of sufficient attention to how changes brought about by AI are impacting different characteristics of work, as experienced by employees. Nevertheless, the survey surfaced a range of attitudes towards the generative AI bid writing tool among employees.

On the one hand, there was enthusiasm about the broadening of capabilities. A survey respondent wrote:

> *"It has broadened my ability to complete certain written tasks that I would have otherwise struggled with. For example, assisting me in drafting case studies or personal biographies for use in tenders, from a set of technical answers that I have sourced internally."*

On the other hand, there were concerns about the creation of extra administrative work and skill erosion, as another explained:

> *"It has created a lot of administrative tasks and more support for teams asking questions about it [...] I'm wary of using it too much and impacting my ability to think for myself and use my existing skills."*

A final survey respondent reflected on differing attitudes among their fellow employees:

> *"[the generative AI tool] allows me to produce a first draft very quickly […] When I talk to colleagues, I feel like I find this easier than some, and having thought about this, I think this is due to my vast experience of drafting documents from scratch over many years. This has given me the experience to know what the answer should look like, and importantly, to remove the inconsistencies and inaccuracies that AI will undoubtedly produce."*

This emphasises the significance of foundational skill development, which may be hindered by AI adoption, if it impacts junior roles. Within the bid function, the stated intention was not to reduce headcount. However, an external working group member involved in higher-level organisational governance of AI emphasised that the tool was introduced in part to reduce the need to hire in 'specialist technical knowledge' and that impact on junior roles was expected 'because actually [the generative AI bid writing tool] can do some of the work that would have been done by junior sales resource.'

There was also an external expectation that the tool was saving significant amounts of time, with another working group member describing its impact as such:

> *"What used to take about, you know, seven to eight days by two, three bid writers to produce, [the generative AI tool] can do that in half a day, right? So that is the speed, that is the extent of the solution."*

It was reported that a small group of 'superusers' had seen notable time savings. However, the overall impact of the tool appeared to be more limited than expectations suggest. At the close of this research, it was noted that the tool may be dropped due to an absence of a sufficient return on investment.

This was attributed to unenthusiastic uptake among employees, as well as the inability to observe or capture improvements to bid quality due to tool usage. This was assumed to be connected to the frictions built into the tool and the rollout of an alternative tool across the whole of LargeCo, which was embedded into the existing information architecture. Reflecting on the experience, a member of LargeCo involved in the attempted rollout of the tool noted:

> *"What this whole process with [the generative AI bid writing tool] has taught me is that you cannot force people to use a tool that they are not getting value out of".*

Learning Points

Cross-functional representation in AI governance mechanisms will help to ensure that a diversity of impacts are surfaced.
LegalCo encouraged experimentation by having senior lawyers model use of the AI However, the representation of senior leaders from HR and people functions is not sufficient to ensure that workplace impacts are front and centre of company strategy or effectively attended to. Managers and people professionals need to be trained on how to steer organisational change and job redesign effectively.

Failing to meaningfully involve employees in AI governance will create distrust and concern, irrespective of organisational narratives regarding responsible usage.
Surveys about attitudes towards AI are not an adequate replacement for robust employee consultation and participation.

Vendors have a commercial interest in building an evidence base pertaining to the success of their tool.
Defaulting to vendors to assess employee attitudes towards will fail to offer a holistic picture.

No amount of tool-specific AI training will lead to buy-in if employees do not genuinely feel that the use of an AI tool adds value to the quality or experience of their work.

Institute *for the* Future of Work

IFOW is an independent research and development institute dedicated to transforming working lives for the better, co-founded by former employment barrister Anna Thomas MBE, Nobel prize-winning economist Sir Christopher Pissarides, and technologist Naomi Climer CBE.

Our core team at Somerset House works with a growing network of strategic partners striving for systems change.

Our vision is a future in which everyone flourishes in work they shape.

Our mission is to understand together how to transform working lives for good.

Our theory of change is that creating and sustaining good work is the best way to achieve this goal and ensure that innovation and social good advance together.

CIPD

The CIPD has been championing better work and working lives for over 100 years.

We help organisations thrive by focusing on their people, supporting our economies and societies.

We are the professional body for HR, L&D, OD and all people professionals – experts in people, work and change.

wbs
WARWICK BUSINESS SCHOOL
THE UNIVERSITY OF WARWICK

With bases in Warwick and at The Shard in London, WBS is consistently ranked as one of the world's leading business schools. WBS was a leading partner in IFOW's Pissarides Review into the Future of Work and Wellbeing.

UKRI Innovate UK

BridgeAI

Innovate UK BridgeAI empowers UK businesses in high-growth sectors, driving productivity and economic growth through the adoption of Artificial Intelligence. We bridge the gap between developers and end-users, fostering user-driven AI technologies.

With a focus on ethics, transparency and data privacy, we aim to build trust and confidence in the development of AI solutions. Strengthening AI leadership, supporting workforces, and promoting responsible innovation, BridgeAI shapes a collaborative and AI-enabled future.

BridgeAI is an Innovate UK funded programme, delivered by a consortium including Innovate UK, Digital Catapult, The Alan Turing Institute, STFC Hartree and BSI.

About our Reports

Our expert Reports present new work from across IFOW's research team, research fellows, practitioners and partners, offering policymakers, firms and academics access to the latest research in the Future of Work space.

If you would like to further information, please contact team@ifow.org

Citation

Good Work Design and AI Adoption: Action Research in Eight Cross-Sector Case Studies, London: Institute for the Future of Work.
DOI: 10.5281/zenodo.20084164

Permission to share

Institute for the Future of Work
Somerset House
Strand
London
WC2R 1LA
ifow.org / @ifow.org