



AI at Work

10 Questions for the CPO Agenda

A practical discussion guide to help senior People leaders explore where AI has already changed the work, and what operating assumptions now need to change with it.




Introduction to this Discussion Guide

In September 2025, OpenAI, Harvard and Duke published one of the most detailed studies of ChatGPT usage to date. The findings challenged some of the dominant assumptions: ChatGPT was not being used primarily as a coding tool, but as a broad knowledge-work layer, heavily applied to practical guidance, information-seeking, writing, synthesis and decision support.

UK-focused evidence has since added an important labour-market dimension. British Progress' April 2026 report, AI and the UK Labour Market: The Evidence So Far, found no sign that AI has replaced jobs at scale in the UK. The more immediate picture is one of uneven, task-level change. AI adoption is concentrated in specific tasks, and exposure to AI does not automatically mean work has been redesigned.

That pattern is consistent with what we are hearing directly from senior People and Technology leaders, including CPOs operating at the leading edge of AI adoption. The conversation is shifting. Tool approval and adoption metrics are giving way to harder questions: workforce redesign, capacity creation, reskilling, junior development, manager capability, governance and the future shape of roles.

This discussion guide is designed to help senior People leaders pursue those harder questions, moving beyond whether AI is being used, and into what it is doing to work, to capability and to the assumptions workforce plans are still built on.





Informal Redesign

1) What decisions about how work gets done have already been made, without anyone realising they were decisions?

AI adoption rarely waits for formal organisation design. Employees may already be changing how work gets drafted, researched, checked, summarised, escalated or delivered.

Why This Matters: The organisation may already have a de facto AI policy. It just has not been written down, reviewed or agreed. The question is whether that policy is one you would endorse.





Shadow AI as Workflow Signal

2) Where are employees using AI because the official process is too slow, too manual, or too poorly supported?

Shadow AI should not only be seen as poor discipline. It may also be a signal that employees are working around processes that no longer fit the pace or complexity of the work.

Why This Matters: The most interesting shadow use may reveal where the organisation needs workflow redesign, not just tighter policy.

AI




Headcount and Workflow Failure

3) Which headcount requests this year are solving a workflow problem that AI or automation could reduce, and does anyone in the approval chain know the difference?

Some teams may be adding people to compensate for manual processes, fragmented systems, poor handoffs, or weak self-service.

Why This Matters: Workforce plans built on current task economics may already be out of date. The risk is not overhiring in general. It is locking in cost for work that may already be changing elsewhere in the market.





Capacity Reallocation

4) Where has AI created capacity, and is that time being reinvested deliberately or just filled?

Freeing up time is not the same as creating value. Capacity can be reinvested in better management, stronger relationships, improved service, more strategic work, or simply absorbed into more activity.

Why This Matters: Most organisations will not struggle to fill recovered time. The harder question is whether that time is being redirected towards better outcomes, or simply sustaining more activity at lower cost. If it is the latter, the productivity story may be real, but the performance story is not.





Judgement Exposed

5) Which skills are becoming more important because AI has made output easier, but judgement more exposed?

AI may make output faster, but it does not automatically make work better. In many roles, the differentiating skills may become the ones AI cannot easily replicate: trust, empathy, judgement, psychological safety, influence, and the ability to navigate ambiguity.

Why This Matters: The reskilling agenda is not prompt training. It is strengthening the human capabilities that matter more when output gets faster: judgement, trust, context, challenge, empathy, and decision quality.





Junior Learning

6) Which junior tasks are we at risk of removing before we understand what they were teaching?

Routine work can be inefficient, but it often teaches context, pattern recognition, judgement, and professional confidence.

Why This Matters: If AI removes developmental work without replacing the learning it created, organisations may improve short-term efficiency while weakening future capability.





Quality Variance

7) If two teams use the same AI tool on the same task, how confident are you that the output would meet the same standard, and who is responsible for that gap?

The tool may be consistent, but the surrounding human system is not. Team norms, review discipline, manager confidence, and expectations around disclosure can all vary.

Why This Matters: Technology access is becoming equal faster than capability, judgement, and review discipline are. The risk is not the tool itself. It is that variance in human oversight becomes the main driver of variance in work quality.





Manager Control Point

8) Where are managers becoming the critical link between AI adoption and work quality?

Managers are increasingly responsible for setting norms, reviewing AI-assisted work, managing risk, and deciding whether AI use is disclosed, hidden, or trusted.

Why This Matters: The biggest source of AI risk in many organisations may not be the model. It may be the manager who does not know what to check, when to challenge, or where human judgement still needs to sit.






Decision Influence

9) Where has AI quietly moved from helping people produce work to shaping decisions?

The governance threshold changes when AI moves from drafting text to influencing recommendations, prioritisation, assessment, task allocation, or workflow actions.

Why This Matters: Decision support can become decision influence very quickly. CPOs need to know where human judgement still sits and where audit, transparency, and review are required.





True Cost Model

10) What is the true cost model of AI-enabled work?

AI may reduce manual work, but it creates new costs: systems ownership, workflow maintenance, quality assurance, governance, data protection, tool licences, change management, and manager enablement.

Why This Matters: Most AI business cases model the upside. Fewer account for the operating model that sits behind it: who owns the workflow, who checks the output, who maintains the process, and who is accountable when the AI gets it wrong. “The cost question is not just ‘what can we save?’. It is ‘what new capability, ownership, and control model does this create?’”.





Signals Worth Watching

1

AI use happening outside approved tools or formal policy.

2

Time saved that nobody has traced to an outcome.

3

Headcount models built on task volumes that AI is already reducing.

4

Managers setting different standards for AI-assisted work.

5

Routine work being removed before learning pathways are redesigned.

6

AI being used in hiring, assessment, or performance decisions without a governance framework.

7

AI-supported recommendations being treated as decisions.

8

Approved enterprise tools existing on paper, while employees still use easier consumer tools.

9

Managers being asked to absorb AI-driven change while their own capacity or spans are worsening.

10

Employees not engaging with AI, with no clear view on whether the issue is confidence, trust, permission, capability, tooling, or relevance.



The Central Question

Not:

**Are people using
AI?**

But:

**Where has AI already
changed the work, and
what operating
assumptions now need
to change with it?**



Evidence Base

This guide is grounded in UK and international AI-at-work evidence from 2025 and 2026, alongside live conversations with senior People and Technology leaders, including CPOs operating at the leading edge of AI adoption.

The consistent signal is that AI is not yet showing up primarily as mass job replacement. The more immediate shift is happening at task level: adoption is uneven, workflow integration matters, manager support strongly affects use, governance is lagging experimentation, and individual productivity gains do not automatically translate into organisation-level redesign.

Sources reviewed include OpenAI Signals and enterprise usage disclosures, Anthropic's Economic Index, worker-adoption research from Bick, Blandin, Deming, Fuchs-Schündeln and Jessen, Federal Reserve adoption monitoring, Gallup workplace research, European and UK-focused labour-market analysis, Danish labour-market research from Humlum and Vestergaard, and professional services research from Thomson Reuters and Deloitte.