



Meet Your New Colleague

A 30-minute KS4 AI literacy half lesson centred on the question of responsibility. Students meet Riley, a fictional AI agent working at Meridian Group, watch Riley make a serious mistake, and decide who is responsible. Built on real 2026 workplace reality.

Year group	KS4 (Years 10–11)
Length	30 minutes
Subject	PSHE / Citizenship / Computing / Careers crossover
Resources	Slides or whiteboard, printed handout (the email)
Prep	3 minutes — read Riley's profile and the email below

Aligned to the OECD AILit Framework & PISA 2029

What this means. In 2029, the OECD's PISA assessment will test 15 year olds on Media and AI Literacy (MAIL) for the first time. The assessment is built on the AILit Framework, which organises AI literacy into four domains: Engage with AI, Create with AI, Manage AI, and Design AI.

AILit domain	Depth	Where in the lesson
Engage with AI	Developed	Reading Riley's email. Spotting where AI is at work.
Manage AI	Developed	Mistake responsibility debate. Approval-ceiling discussion.
Design AI	Introduced	Brief Cappelli vs Bersin framing in the plenary.

Background: what's actually happening in 2026

In 2026, enterprise software is being rebuilt around AI agents. **Microsoft Agent 365** launched on May 1, 2026 as a way for companies to observe, govern, and manage AI agents the same way they manage human users. **Korn Ferry** reports that more than half of talent leaders are planning to add autonomous AI agents to their teams in 2026.

The big question for KS4 students. They are the first cohort whose first job is being designed in this environment. Riley is fictional. Everything Riley does in this lesson is happening somewhere right now.

Riley's profile

Project this on screen as students walk in. Don't explain it. Let them read it.

MERIDIAN GROUP · EMPLOYEE PROFILE · MG-A0142

NAME	Riley	TYPE	AI agent (autonomous)
JOB TITLE	Junior Marketing Coordinator	REPORTS TO	Sarah Chen, Head of Marketing
START DATE	January 2026	SALARY COST	\$12,000 / yr (vs \$42,000 human)
WORKING HOURS	24 / 7	BUILT ON	Microsoft Agent 365



Lesson run

0–5 min**Hook: meet Riley**

- Project Riley's profile (previous page) on screen as students walk in. Don't introduce it.
- Once the room is settled: **"This is Riley. Riley started at Meridian Group in January. Riley is doing the job a junior marketing coordinator would have done. Riley's salary cost is a quarter of a human's. Riley works 24/7. Riley does not exist as a person."**
- Ask the room: **"What's your gut reaction?"** Take three hands. Don't comment.

5–15 min**The mistake**

- Reveal the email artefact (next page). Display it. Read it aloud, slowly.
- Riley sent a pricing email to 4,000 customers with the wrong figures. The error cost Meridian Group £180,000 in goodwill discounts before the company caught it.
- **Pair task:** rank four candidates from most to least responsible. **(a)** Sarah, Riley's manager. **(b)** Riley itself. **(c)** Meridian Group. **(d)** Microsoft, the vendor that built Riley.
- Each pair justifies their top and bottom with one sentence each.

15–25 min**Discussion: what does accountability mean?**

- Take feedback from three or four pairs. Probe the reasoning.
- Surface the key tension: a human junior who sent that email might be fired. Riley cannot be fired. Riley does not have a contract. Riley does not have a Disciplinary Committee.
- **The big question on the board:** *If the worker isn't a person, what does "accountability" actually mean?*
- If time allows, introduce the two positions briefly: **Bersin** ("AI as colleague — build HR systems around them") vs **Cappelli** ("The fallacy of treating AI agents as fellow employees. Agents need to be supervised, not staffed.").

25–30 min**Plenary: one sentence**

- On the board: **"When an AI tool makes a mistake, the person who is responsible is ____, because ____."**
- Each student writes their sentence in their book or on a sticky note.
- Take three out loud. End there.



Artefact: The email

Project this in full screen during the mistake section. Read it aloud, slowly. Then ask: who is responsible?

FROM	Sarah Chen (<i>Head of Marketing, Meridian Group</i>)
TO	marketing-leadership@meridiangroup.com
SUBJECT	URGENT — Pricing email error, 4,000 customers affected
SENT	Tuesday 14 March, 16:42

Team,

I need to flag a serious error from this morning. At 09:14, our agent Riley sent a promotional email to our full retail customer list (n=4,012) advertising our Spring package at £49 per month.

The correct price is £149 per month. The discount field in the source brief was misread. Riley generated, formatted, and sent the email autonomously — under the approval ceiling Riley is permitted to operate within.

By 11:30 we had 312 customers attempting to sign up at the £49 rate. Legal advised that under consumer protection rules we should honour the price for those who acted in good faith before the correction. Estimated cost of goodwill: £180,000.

I have paused all of Riley's outbound capabilities pending review. We need to discuss (1) the approval ceiling, (2) the audit trail, and (3) what we say publicly.

Sarah

For discussion (in pairs). Rank the following from *most* responsible to *least* responsible: **(a)** Sarah, Riley's manager. **(b)** Riley itself. **(c)** Meridian Group, the company. **(d)** Microsoft, the vendor. Justify your top and bottom answers with one sentence each.



Stretch and support

Stretch

- **Extension reading:** Peter Cappelli, *The fallacy of treating AI agents as fellow employees* (HR Executive, Feb 2026). 100-word summary plus response.
- **Approval ceiling design:** students draft three rules for what Riley should and shouldn't be allowed to send without human review.

Support

- Pre-teach two terms: **autonomous** (operating without human approval for each action) and **approval ceiling** (the size of decision a worker can make on their own).
- For the responsibility ranking task, reduce to two candidates (Sarah vs Riley) instead of four.

Assessment for learning

- **During the responsibility ranking:** can the student articulate why an AI agent cannot be "responsible" in the legal sense, even when it caused the error?
- **In the plenary sentence:** strong answers name a specific person and a specific reason. Weaker answers say "the AI" without naming the human behind it.

Safeguarding and framing notes

- **Cut both ways on the mistake.** Riley's pricing error was serious. But humans make those errors too — and Riley fixed it faster than a human could have. Make sure both sides surface.
- **The pronoun choice is pedagogical.** Throughout this lesson, Riley is referred to as "it." Anthropomorphic pronouns prejudice the question of whether AI agents should be treated like fellow employees. If students switch to "they" or "she," ask them why.
- **Don't catastrophise.** The honest framing is *this is happening in some industries in 2026; what does it mean for how we prepare?* — not *your career is doomed*.

Sources

- Microsoft: Agent 365 launch announcement (general availability May 1, 2026).
- Korn Ferry: HR and Talent Trends Shaping 2026.
- Cappelli, P. (February 2026). *The fallacy of treating AI agents as fellow employees*. HR Executive.
- Bersin, J. (April 2026). *Introducing HR 2030: A Vision For Agentic Human Resources*.
- OECD & European Commission: AILit Framework, draft for review (May 2025).
- OECD: PISA 2029 Media and AI Literacy (MAIL) assessment framework.