北京大学光华管理学院 · 光华思想力

2024 — 2026 三轮追踪调查

人工智能驱动的人力资源管理

中国企业三轮追踪调查报告

Artificial Intelligence-Enabled Human Resource Management:
A Three-Wave Longitudinal Survey of Chinese Firms

北京大学光华管理学院

“人工智能驱动的人力资源管理”思想力课题组

2026年7月

引用本报告

中文引用（APA 第7版）

> 姜铠丰. (2026). 人工智能驱动的人力资源管理：中国企业三轮追踪调查报告（2024—2026）. 北京大学光华管理学院光华思想力课题组.

English citation (APA 7th)

> Jiang, K. (2026). Artificial intelligence-enabled human resource management: A three-wave longitudinal survey of Chinese firms (2024–2026). Guanghua Thought Leadership Platform, Guanghua School of Management, Peking University.

本报告基于2024年10月至2026年5月实施的三轮企业追踪调查。报告全文所载数据与图表版权归课题组所有， 引用请注明出处。

目录

CONTENTS

摘要 内容摘要 4

01 调研背景：为什么是人工智能，为什么是中国企业 5

02 方法论与样本：三轮追踪的设计与实施 8

03 采纳总体图景：从三分之一到接近半数 13

04 九大模块的应用格局 20

05 采纳的驱动因素：组织条件与人力资源负责人 28

06 动机归因与企业画像：分化在准备度，而非动机 37

07 挑战与治理：应用跑在了制度前面 41

08 组织后果：回报、差距与重构 45

09 三年全程追踪面板数据的发现 52

10 结论、启示与未来图景 57

附 致谢 64

内容摘要

自2024年10月至2026年5月，本课题组围绕人工智能在中国企业人力资源管理中的应用（以下简称AIHRM），对企业人力资源负责人开展了三轮追踪问卷调查，累计获得2,241个企业观测，其中487家企业至少参与两轮调查，154家企业三轮全程参与。本报告基于三轮数据的系统比较与匹配面板分析，刻画了中国企业AIHRM从起步到扩散的完整轨迹。

调查的核心发现可以概括为五点。第一，采纳率稳步攀升。在人力资源管理中已实际部署人工智能技术的企业比例从2024年11月的37.3%上升至2026年5月的47.9%；在第一轮和第三轮均参与调查的288家匹配企业中，这一比例从38.5%上升至51.0%，表明增长主要来自企业的真实转化而非样本构成变化。第二，应用呈现面广而不深的格局。招聘、人力资源行政与工作设计构成应用的第一梯队，但86.1%的企业仍处于探索或试点阶段，进入规模化及以上阶段的企业仅占2.7%。第三，采纳的分化根源在于组织准备度而非认知态度。潜在剖面分析识别出四类企业画像，各画像在采纳动机上的差异很小，在资源与人员准备上的差距却十分显著；人力资源负责人的年龄与专业背景对采纳具有稳定的预测作用，其中仅具备人力资源单一专业背景的负责人所在企业，在两个独立时间窗口的滞后检验中均表现出显著更低的采纳增量。第四，治理建设明显滞后于应用扩张。数据隐私与安全的关注度大幅上升并跃居各项挑战之首，但51%的企业尚未建立任何人工智能治理措施，覆盖率最高的单项措施也不足三成。第五，回报开始显形但分布不均。可以评估回报的企业中约六成感知到中等以上的投资回报，人力资本涌现效应逐轮增强。

基于上述发现，报告在结论章节提出了面向企业管理者的行动建议，包括将投入重心从模块铺开转向单点做深、把治理框架建设提前到规模化之前、将人力资源负责人的数据素养与跨界经验纳入高管配置考量，以及正视并规范员工层面自发的人工智能使用。

01

第 一 章

调研背景

为什么是人工智能，为什么是中国企业

1.1 人力资源管理正在成为人工智能落地的关键场景

过去三年，生成式人工智能完成了从技术突破到组织应用的快速跨越。大语言模型在文本理解、内容生成与多轮交互上的能力跃升，使人工智能第一次能够直接介入以语言和判断为核心的管理工作，而人力资源管理恰恰是这样一类工作的集中地带。从简历筛选、面试评估、培训内容生成，到员工咨询、绩效反馈与组织诊断，人力资源管理的几乎每一个环节都存在可以被人工智能重塑的空间。Gartner在2023年的调查显示，76%的人力资源负责人认为，如果未能在未来一到两年内引入人工智能解决方案，其所在企业将在竞争中落后于同行。这一判断在随后的市场演化中得到了验证：创业与投资数据库Crunchbase的记录显示，人工智能驱动的人力资源管理初创企业在经历调整后于近两年出现明显反弹，资本与产品供给的双重扩张为企业采纳提供了现实条件。

与技术供给的快速成熟相对照的，是学术界与实践界对采纳过程本身的认识仍然有限。已有研究多以单次横截面调查为主，能够回答某一时点上有多少企业在使用人工智能，却难以回答三个更重要的问题：企业的采纳决策如何随时间演化，哪些组织因素与个体因素在驱动或阻碍采纳，以及采纳之后组织发生了什么。回答这些问题需要对同一批企业进行跨年度的持续追踪，这正是本项调查的出发点。

1.2 中国企业为何需要专门的观察

将观察对象聚焦于中国企业，首先源于一个基准数据缺失的现实。目前管理者与研究者在讨论人工智能采纳水平，尤其是人工智能在人力资源管理方面的采纳水平时，引用的百分比大多来自国际咨询机构发布的全球调查，这些调查的受访企业以北美与欧洲市场为主体，来自中国企业的样本通常数量有限、行业与规模覆盖单薄，且很少公布中国子样本的独立结果。以此类数据推断中国企业的采纳进度，既可能高估，也可能低估，参照价值存在疑问。中国企业需要一套基于本土大样本、测量口径一致、可以跨年比较的基准数据，这是本项调查最直接的动因。

在此之外，聚焦中国企业还有三方面的理由。其一，中国拥有全球规模最大的劳动力市场和数量庞大的企业群体，人工智能对就业结构与管理方式的影响在中国具有最大的体量效应；国内基础大模型在2025年前后的快速普及大幅降低了企业的使用门槛，使得人工智能的组织扩散在中国呈现出与欧美市场不同的节奏与路径。其二，中国企业的人力资源管理正处于数字化基础参差不齐的阶段，先进制造企业与中小民营企业在数据基础、系统建设与专业能力上差距悬殊，人工智能的引入是叠加在这一异质性之上的，其扩散规律难以从发达市场的经验直接外推。其三，中国的制度环境具有独特性，包括数据安全与个人信息保护的监管框架、国有企业与民营企业并存的所有制结构，以及各级政府对人工智能产业的积极推动，这些因素共同塑造了一个值得独立研究的制度场景。

基于以上考虑，本课题组在光华思想力课题的支持下，于2024年秋季启动了这项追踪调查，计划以每年一到两轮的频率持续观察中国企业AIHRM的采纳进程。本报告呈现前三轮调查的完整结果，时间跨度为2024年10月至2026年5月，覆盖了生成式人工智能在中国企业界从概念热议走向常规部署的关键时期。

02

第 二 章

方法论与样本

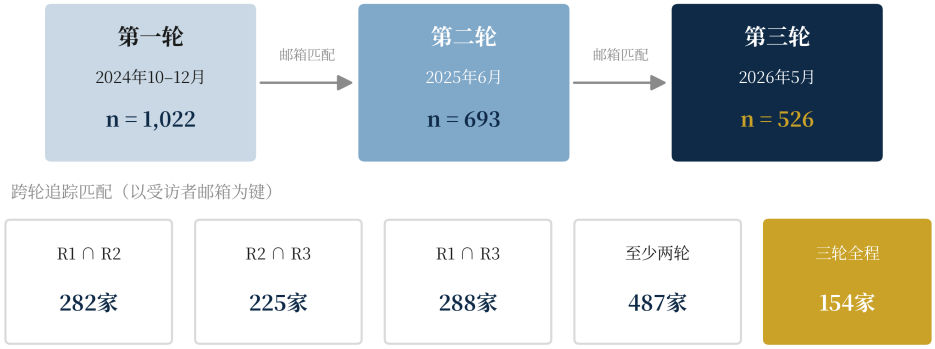
三轮追踪的设计与实施

2.1 调查设计与实施过程

本项调查采取重复测量与固定样本追踪相结合的设计。第一轮调查于2024年10月至12月实施，问卷主体回收集中在11月，共获得1,022份来自企业人力资源负责人的有效问卷。第二轮调查于2025年6月实施，回收有效问卷693份。第三轮调查于2026年5月实施，回收有效问卷526份。三轮调查的受访者均为企业人力资源工作的实际负责人，职位覆盖首席人力资源官、人力资源副总裁、人力资源总监与人力资源经理，每家企业由一位负责人代表作答，问卷以企业为分析单位。

为实现跨轮追踪，调查以受访者邮箱为匹配键，在数据清理中对邮箱进行标准化处理后建立企业层面的纵向链接。匹配结果显示，第一轮与第二轮共同参与的企业为282家，第二轮与第三轮为225家，第一轮与第三轮为288家；至少参与两轮的企业合计487家，三轮全程参与的企业为154家。这一结构使本报告能够同时给出两类证据：基于各轮全样本的总体图景，以及基于匹配面板的组内变化。凡涉及跨期比较的关键结论，报告均以匹配面板进行了复核，以排除样本构成变化的干扰。图2-1概括了整体设计与样本流转情况。

图 2-1
三轮追踪调查设计：2,241个企业观测，154家企业三年全程可追踪
受访者均为企业人力资源负责人（CHRO/HR副总裁/总监/经理）；三轮核心量表保持一致



数据来源：北京大学光华管理学院·光华思想力课题组，AIHRM追踪调查（2024–2026）。

2.2 测量内容

三轮问卷保持了核心测量的一致性。采纳状态通过单选题测量，区分自主研发部署、与外部服务商合作部署、正在考虑自研、正在考虑合作与暂无计划五种状态；第三轮将前两类进一步细分为九类具体实现路径，用于刻画2026年的采纳生态，在跨轮比较时折算回五类可比口径。应用广度通

过九个人力资源模块的多选题测量，每个模块列出七类具体应用；第三轮针对生成式与智能体技术新增了六个选项，跨轮比较时同样只统计三轮均出现的七类应用。组织层面的解释变量包括对AIHRM的重视程度、人工智能伦理与治理关注、技术契合度、资源与人员准备度、高绩效工作系统、行业竞争强度与行业人工智能普遍性，均采用五点量表的多题项合成。采纳归因区分增强逻辑、替代逻辑与外部压力三类。此外，问卷测量了实施挑战、组织绩效、人员流动，以及生成式人工智能促进个人知识向团队能力转化的涌现效应量表。第三轮新增了应用成熟度阶段、治理措施清单、投资回报感知、绩效差距感知、人力资源团队编制变化与新增角色等测量。

人力资源负责人的个体特征包括性别、年龄、受教育程度与专业背景，其中专业背景区分仅有人力资源背景与兼具工程技术、财务、法律、营销、战略或运营等跨界背景两类，三轮均有测量。考虑到年龄、专业背景等属于相对稳定的特征，本报告在匹配样本内将这些信息跨轮沿用，涉及此处理的图表均在注释中作了标明。第二轮还对负责人的八项胜任力进行了专项测量。

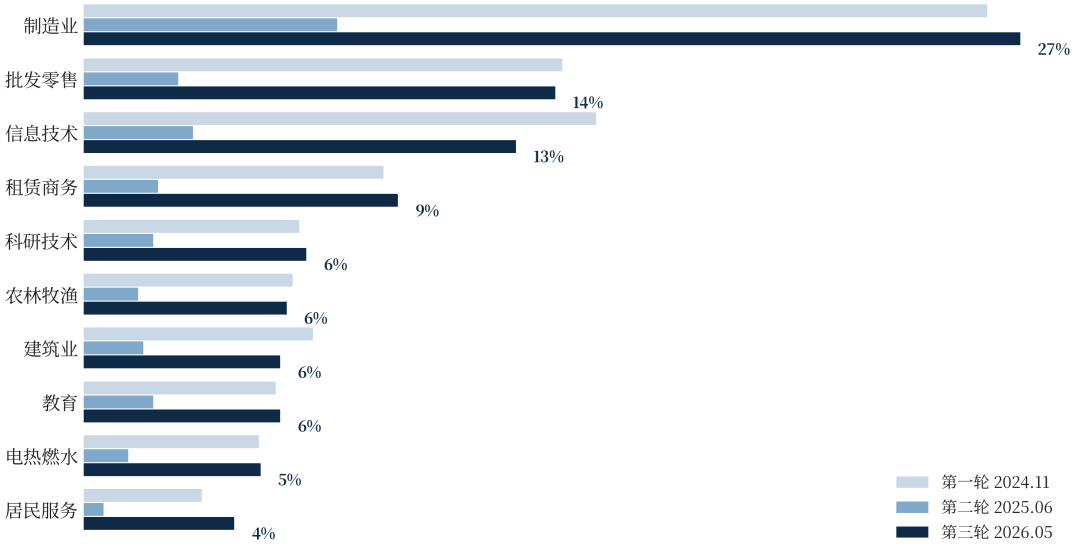
2.3 样本构成

三轮样本在行业、所有制与规模结构上保持了较高的一致性，为跨轮比较提供了基础。行业分布方面，制造业与信息技术业构成样本的两大主体，批发零售、金融、科研技术服务等行业依次跟进，三轮之间未出现明显的结构漂移。所有制方面，民营企业约占六成，国有及国有控股、股份制与外资合资企业构成其余部分。规模方面，样本以中小企业为主体，员工规模在两千人及以上的大型企业占比稳定在两成左右。图2-2与图2-3展示了具体分布。

图 2-2

样本行业结构三轮保持稳定：制造业与信息技术业为两大主体

各行业占比（多选题，占各轮样本比例，展示第三轮占比前十行业）

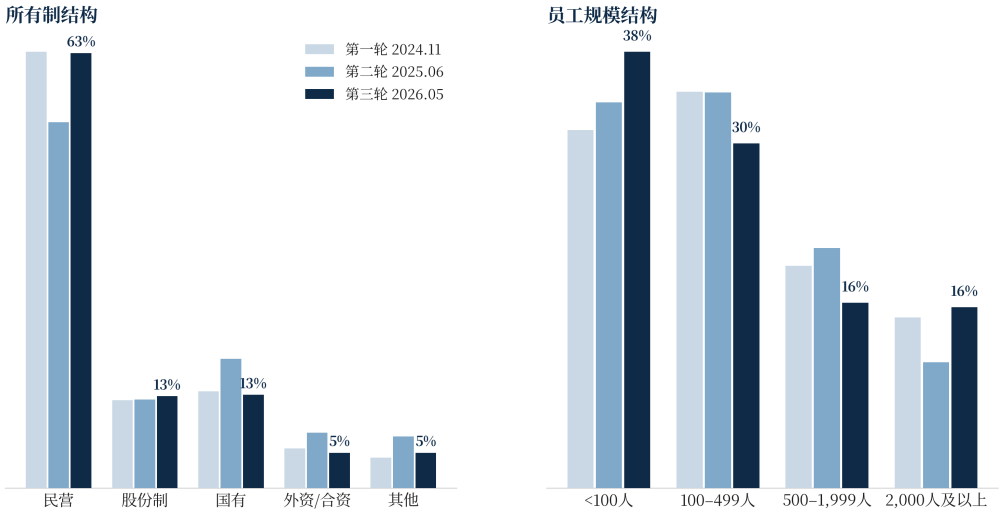


数据来源：AIHRM三轮追踪调查。行业为多选题，各行业占比之和大于100%。

图 2-3

民营企业约占六成，中小规模企业为样本主体，三轮结构高度可比

占各轮有效样本比例



数据来源：AIHRM三轮追踪调查。

2.4 解读与启示

三轮问卷保持同一套测量口径、同一类受访者、同一个分析单位，是为了提供一条可以对标的时间序列，而不是三张彼此无法比较的快照。这一点在实践中常被忽视：许多企业内部的人工智能盘点每年更换问法与口径，结果是数据积累了三年，却无法回答自己究竟前进了多少。建议企业在内部评估中借鉴这一原则，确定一套核心指标后至少锁定三年不变，把变化留给新增题目而非既有题目。

样本结构的稳定性也提供了一个使用提示。本报告的样本以民营企业与中小规模企业为主体，制造业与信息技术业占比较高，这决定了报告中的各项百分比更适合作为同类企业的参照，而非全国普查值。人力资源负责人在引用本报告数据时，宜优先对照与自身规模、行业相近的分组结果，而非仅看总体均值；第三章之后各章提供的分组数据正是为此准备。

03

第 三 章

采纳总体图景

从三分之一到接近半数

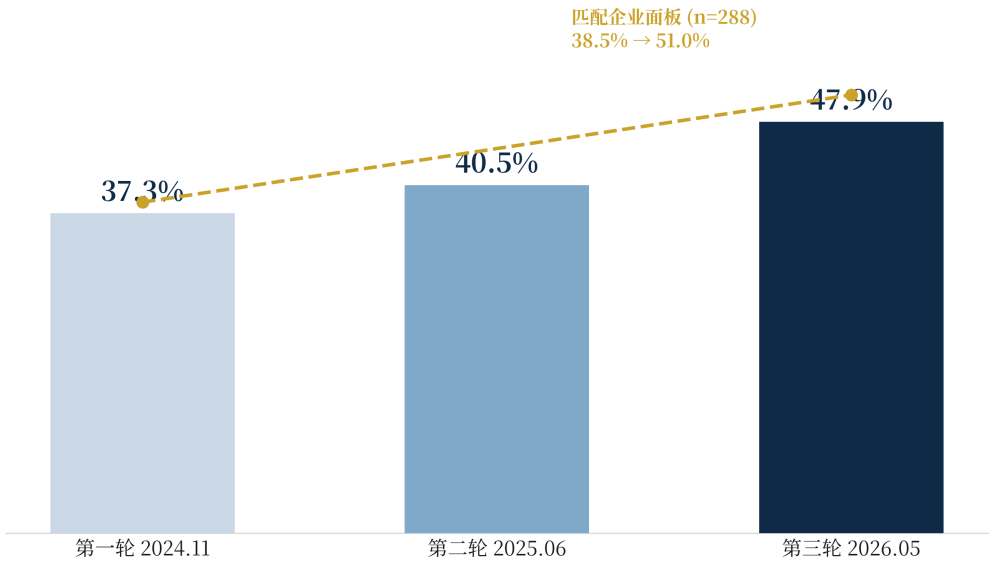
3.1 采纳率的三年轨迹

三轮调查记录了一条清晰的上升曲线。在人力资源管理中已实际部署人工智能技术的企业比例，从2024年11月的37.3%上升至2025年6月的40.5%，再升至2026年5月的47.9%。为排除样本构成变化的影响，我们对第一轮与第三轮均参与调查的288家企业进行了组内比较，其采纳率从38.5%上升至51.0%，涨幅超过全样本口径，说明总体上升主要由企业的真实转化驱动。按照当前速度，已采纳企业过半将在2026年内成为现实。

图 3-1

AIHRM采纳率三年持续攀升，接近半数企业已落地应用

已采纳＝自主研发或与外部服务商合作部署AI于HR管理（占各轮全样本比例）



数据来源：AIHRM三轮追踪调查（R1 n=1,022；R2 n=693；R3 n=526）。金色虚线为R1与R3均参与的288家匹配企业的组内变化。

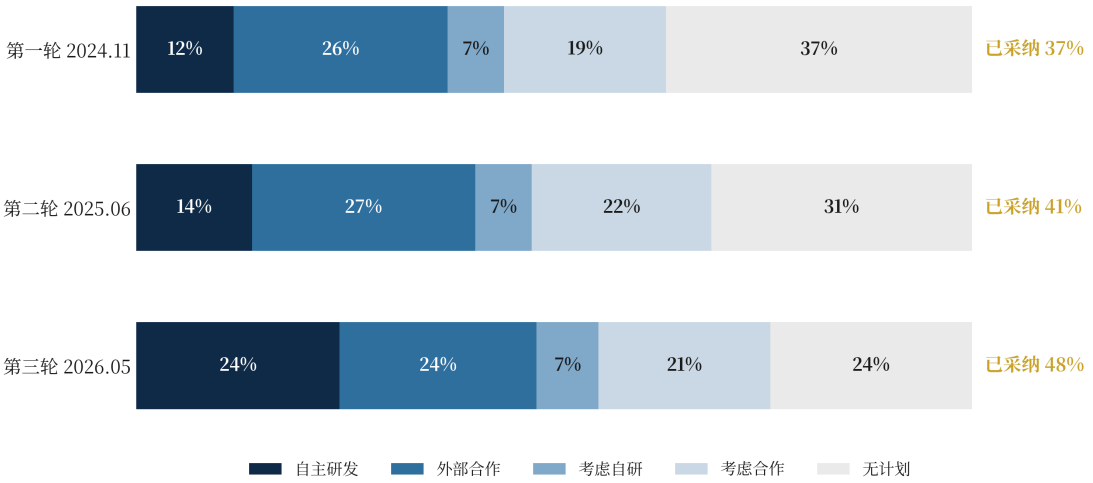
3.2 采纳方式的结构变化

采纳方式的构成揭示了增长的来源。2024年第一轮调查中，明确表示暂无计划的企业占36.6%，到第三轮已收缩至24.1%；同期自主研发部署的比例从11.6%上升至24.3%，与外部服务商合作的比例稳定在四分之一上下。观望阵营的持续收缩与自研阵营的扩大同时发生，表明企业不仅在更多地采纳，也在更深地介入技术实现本身。处于考虑状态的企业合计仍占约28%，构成未来一到两年采纳率继续上行的储备。

图 3-2

“无计划”阵营持续收缩，外部合作是采纳的主通道

AIHRM采纳方式构成（第三轮九类细分方式折算为可比五类）



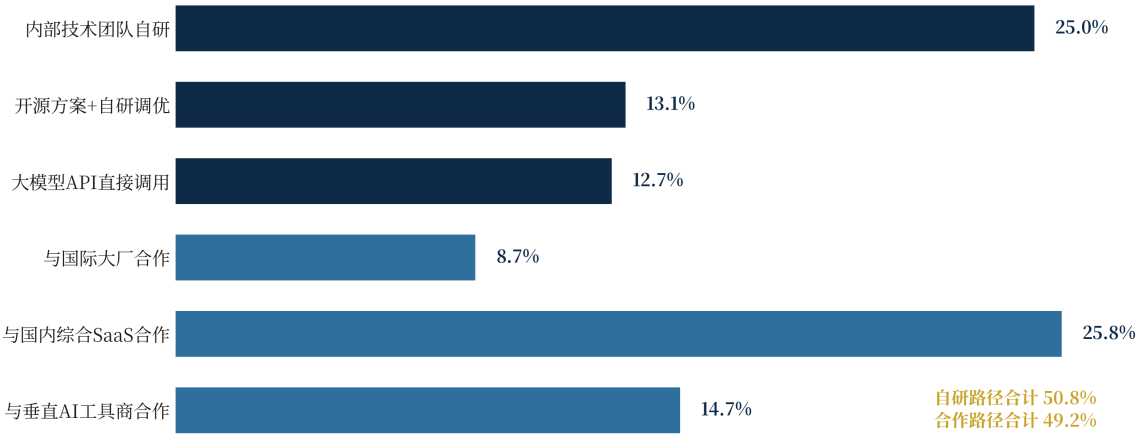
数据来源：AIHRM三轮追踪调查。

第三轮对采纳路径的九类细分提供了2026年采纳生态的截面。在252家已采纳企业中，自研类路径合计占50.8%，包括依托内部技术团队开发、基于开源方案调优以及直接调用大模型应用程序接口三种方式；合作类路径合计占49.2%，其中与国内综合软件服务商的合作最为普遍，与垂直领域人工智能工具商及国际厂商的合作依次减少。大模型接口的直接调用作为一条低门槛路径的兴起，是本轮调查观察到的显著变化，它使不具备大规模技术团队的中型企业也能够以较低成本进入自研阵营。

图 3-3

采纳生态走向多元：国内SaaS与内部自研并驾齐驱，大模型API成为新入口

已采纳企业的具体实现路径（第三轮，n=252，占已采纳企业比例）



数据来源：AIHRM第三轮调查。深色为自研类路径，浅色为外部合作类路径。

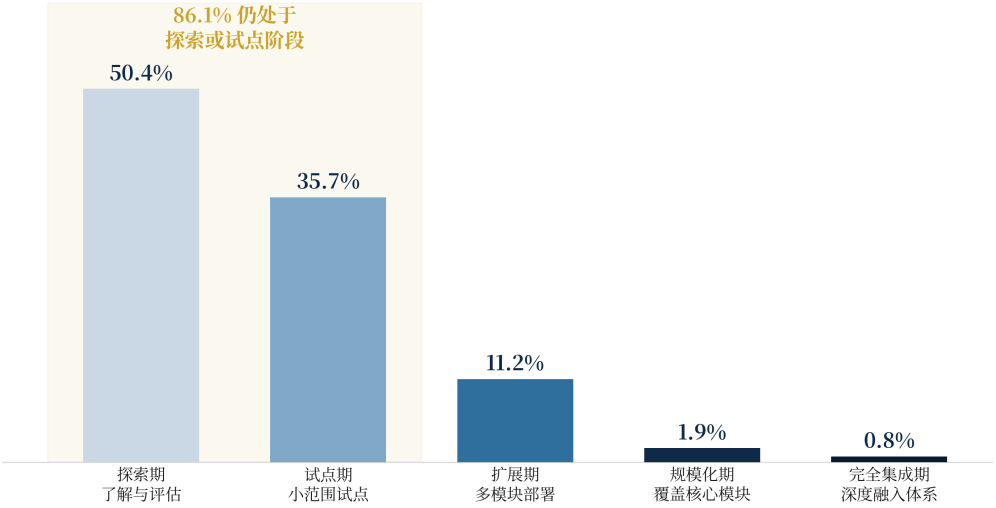
3.3 成熟度：面广而不深

采纳率的量变尚未带来成熟度的质变。第三轮调查引入了应用成熟度的五阶段自评，结果显示50.4%的企业处于了解与评估技术方案的探索阶段，35.7%处于小范围试点阶段，两者合计86.1%；进入多模块部署的扩展阶段的企业占11.2%，而达到规模化应用与深度集成阶段的企业合计仅占2.7%。将这一分布与接近半数的采纳率并置，可以更准确地理解当前阶段的特征：多数企业已经迈过从无到有的门槛，但距离人工智能真正融入人力资源管理体系仍有相当距离。第四章对模块层面广度与深度的分析将进一步支持这一判断。

图 3-4

面广而不深：采纳率近半，但仅2.7%的企业进入规模化及以上阶段

AIHRM应用成熟度五阶段分布（第三轮，n=526）



数据来源：AIHRM第三轮调查。

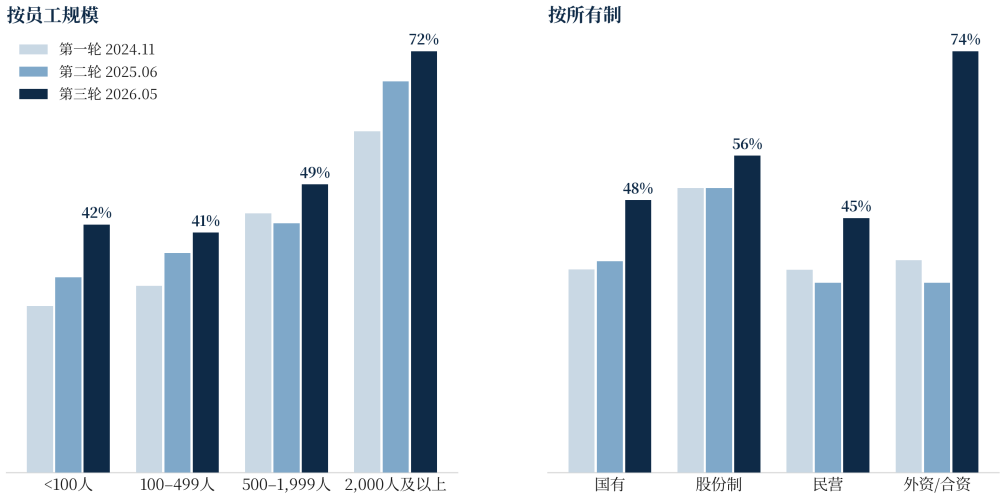
3.4 规模与所有制的分化

采纳率在企业规模上呈现稳定的梯度。员工规模两千人及以上的企业采纳率在三轮调查中均显著领先，且领先幅度未见收窄；百人以下小微企业的采纳率虽然也在上升，但与大型企业的差距依然明显。所有制维度的差异相对温和，民营企业与股份制企业的采纳率略高于国有企业，外资合资企业居中，且各类企业均呈上升趋势。规模梯度的持续存在提示，资源与数据基础仍是采纳的硬约束，这一线索将在第五章的驱动因素分析中得到更系统的检验。

图 3-5

规模优势明显：2,000人以上企业采纳率显著领先，且增速未减

AIHRM采纳率（已采纳占比）



数据来源：AIHRM三轮追踪调查。

3.5 解读与启示

三年数据传递的第一个判断是，观望的窗口正在关闭。采纳率接近半数，明确表示暂无计划的企业已从三分之一以上收缩到不足四分之一，这意味着人工智能在人力资源管理中的应用正在从少数先行者的选择变为多数企业的常规配置。对尚未启动的企业而言，问题已经不是是否要做，而是以何种方式进入、如何避免重复他人已经踩过的坑。值得注意的是，仍有约28%的企业处于考虑状态，这一群体构成了未来一到两年的转化储备，也意味着当前的先发优势窗口最多还有一到两年。

第二个判断是，采纳率这个数字本身容易造成误导。将47.9%的采纳率与86.1%的企业仍处探索或试点阶段并置，更接近真实的图景是：多数企业已经进场，但真正完成规模化推进的仍是少数。对管理者而言，这提示了两件事。其一，如果贵司尚未采纳，追赶的难度低于采纳率数字给人的压迫感，因为领先者的实际身位并不遥远；其二，如果贵司已经采纳，真正的竞争不在于是否用了人工智能，而在于能否率先从试点跨入规模化，那2.7%的企业才是真正的先行者。管理层在设定目标时，应当把成熟度阶段而非采纳与否作为衡量进度的标尺。

第三个判断关于路径选择。第三轮数据显示自研与合作大体各占一半，而大模型接口的直接调用作为一条新路径快速兴起，使中型企业得以在不组建大规模技术团队的前提下进入自研阵营。这为资源有限的企业提供了一个现实选项：不必在自建平台与全盘外包之间二选一，可以先以接口调用的方式在单一场景验证价值，再决定是否深度投入。规模梯度的持续存在同样值得注意，两千人以上企业采纳率已达72%，与百人以下企业的四十余个百分点差距主要源于数据基础与专业能力，

而非认知差异；中小企业的现实策略是借助成熟的外部服务快速补齐基础，而非试图复制大型企业的自建路线。

04

第 四 章

九大模块的应用格局

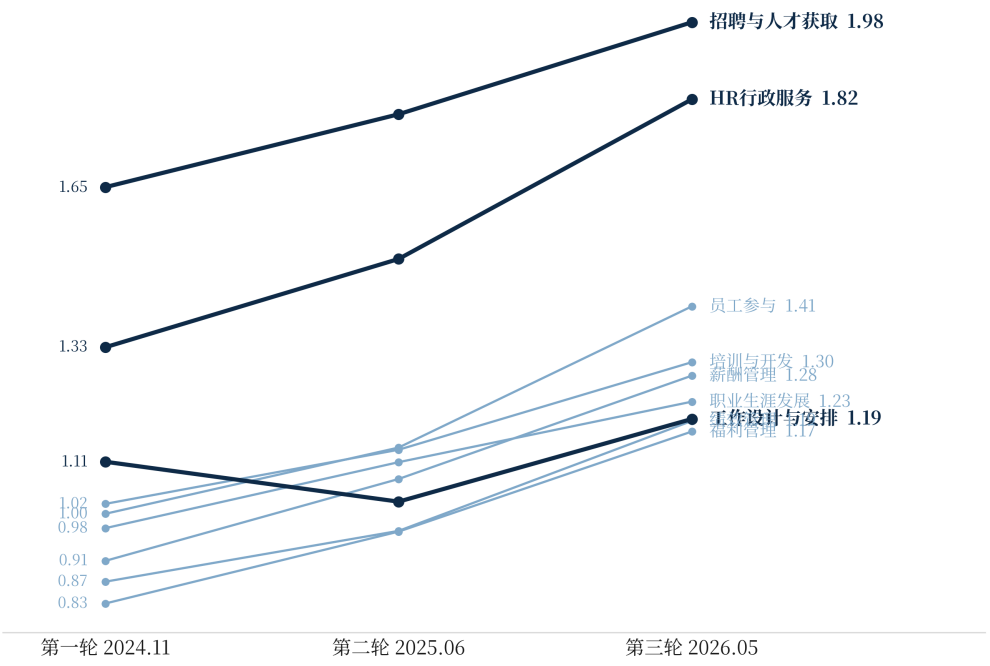
4.1 模块梯队与三年演进

按九个人力资源模块统计的平均应用数量显示，三年间所有模块均呈上升态势，且梯队格局保持稳定。招聘与人才获取、人力资源行政服务与工作设计构成第一梯队，培训与开发、员工参与紧随其后，绩效管理、薪酬管理与福利管理的应用数量相对最低。第一梯队的三个模块具有共同特征：任务标准化程度高、文本处理密集、决策风险相对可控，这使其成为人工智能进入组织的自然入口。相比之下，薪酬与绩效直接关联员工切身利益与合规风险，企业在这些模块的推进明显更为审慎。

图 4-1

九大模块全线上行：招聘、HR行政与工作设计构成应用第一梯队

各模块平均采用的AI应用数量（0–7项可比口径，剔除第三轮新增选项）



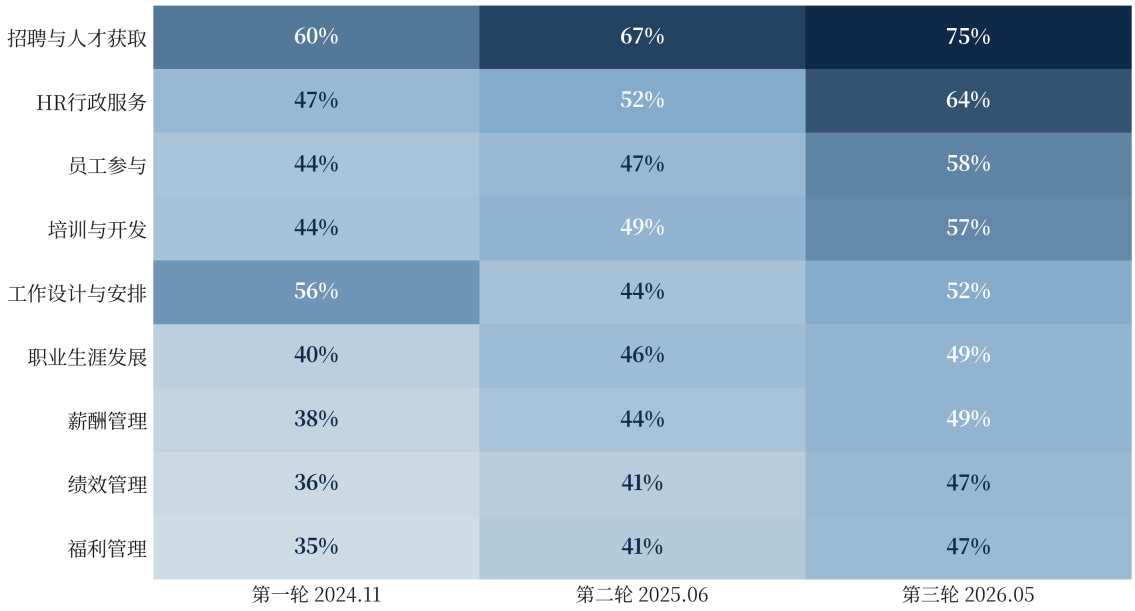
数据来源：AIHRM三轮追踪调查。为保证跨轮可比，仅统计三轮均出现的7类具体应用。

从至少采用一项具体应用的企业比例看，招聘环节在第三轮已超过半数，成为渗透率最高的模块；行政服务与工作设计的渗透率也已接近半数。薪酬与福利管理在三轮调查中始终处于渗透率的末端，构成应用版图上相对的洼地。图4-2的热力图完整呈现了模块与轮次交叉的渗透结构。

图 4-2

过半数企业已在招聘环节使用至少一项AI应用，薪酬福利仍是洼地

在该模块采用至少一项具体AI应用的企业比例



数据来源：AIHRM三轮追踪调查（可比口径）。

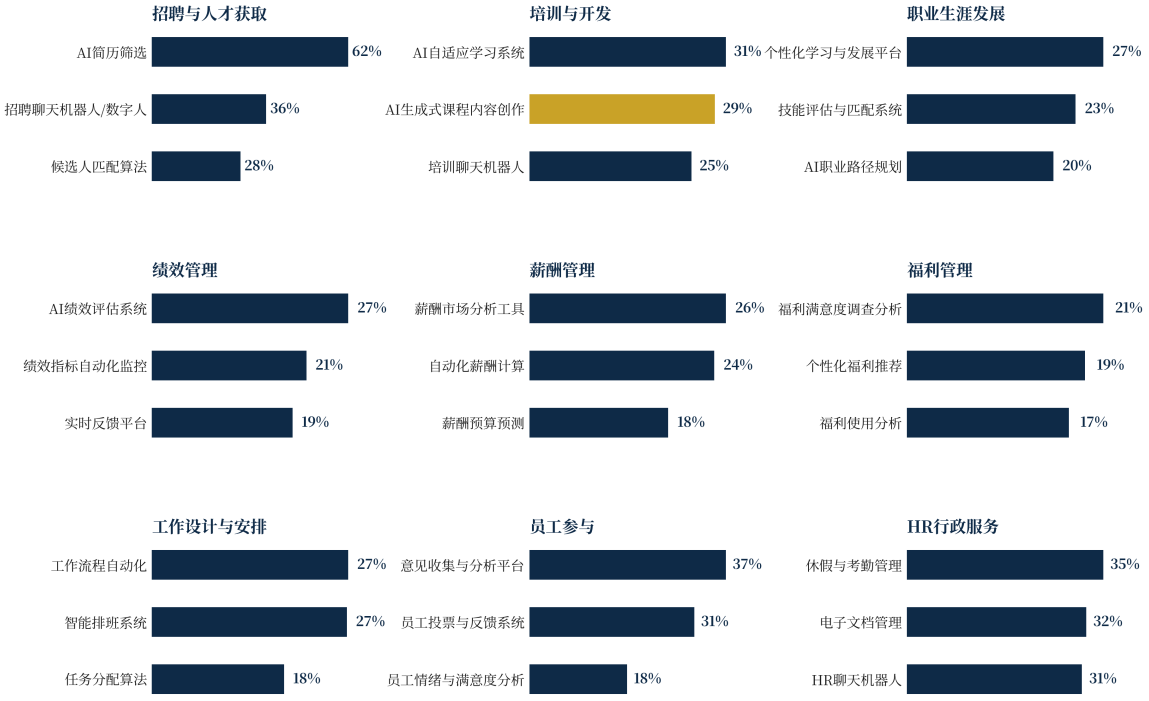
4.2 各模块的主流实践与前沿应用

第三轮调查为九个模块分别列出了具体应用清单，各模块使用率最高的三项实践共同勾勒出2026年中国企业AIHRM应用的真实构成。招聘模块由简历筛选绝对主导，62.4%的企业已使用人工智能算法自动筛选简历，这是全部清单中渗透率最高的单项应用，招聘聊天机器人与数字人（36.3%）、候选人匹配算法（28.1%）分列其后。人力资源行政服务的头部实践高度集中于流程自动化：休假与考勤管理（35.0%）、电子文档管理（31.9%）与全天候应答的聊天机器人（31.2%）三项并驾齐驱。员工参与模块中，意见收集与分析平台（36.7%）和员工投票与反馈系统（30.8%）的使用率超出预期，说明人工智能已经进入组织倾听员工声音的环节。工作设计模块以工作流程自动化（26.9%）与智能排班（26.7%）为主；培训模块中，自适应学习系统（30.6%）居首，而2026年新增测量的生成式课程内容创作以28.9%直接跃居第二，是全部新增选项中渗透最快的一项。

图 4-3

九大模块的主流实践：筛选、监控与流程自动化类应用占据各模块头部

第三轮各模块使用率最高的三项具体AI应用（占全部526家企业的比例；金色=2026年新增测量选项）



数据来源：AIHRM第三轮调查（n=526，多选题）。

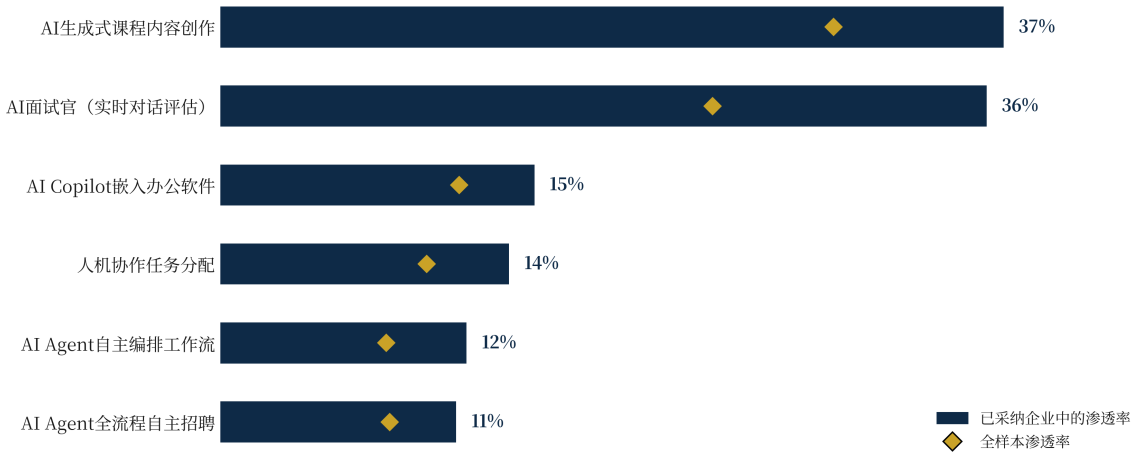
纵览各模块的头部实践，可以发现一个共同特征：占据首位的几乎全部是数据分析、筛选匹配与流程自动化类应用，包括绩效模块的评估系统与指标监控、薪酬模块的市场分析与自动计算、职业发展模块的个性化学习平台与技能匹配、福利模块的满意度分析。换言之，当前应用版图的主体是让机器处理数据与流程，而由人工智能直接生成内容或自主作出决策的应用仍居次要位置。这一格局与技术采纳的风险逻辑一致：企业优先把人工智能放在错误可纠正、责任可追溯的环节。

为追踪技术前沿的渗透进度，第三轮问卷针对2025年以来的技术演进新增了六类前沿应用的测量。在已采纳企业中，生成式课程内容创作与承担实时对话评估的AI面试官渗透最快，后者在全样本中的使用率已达23.2%，接近招聘模块传统实践的水平；嵌入办公软件的智能辅助工具与人机协作任务分配处于中段；由人工智能自主串联招聘全流程或自主编排工作流程的智能体应用仍在早期，全样本渗透率均在8%上下。总体而言，生成内容与辅助增强类的前沿应用扩散明显快于自主决策类应用，企业愿意让人工智能生成材料、参与评估，但尚不放心让它独立作出决定，这与第六章将要讨论的增强归因主导的动机结构相互印证。

图 4-4

生成式与智能体应用崭露头角：生成式课程与AI面试官渗透最快，Agent仍在早期

2026年新增测量的六类前沿AI应用采用比例（第三轮）



数据来源：AIHRM第三轮调查（n=526；已采纳企业n=252）。

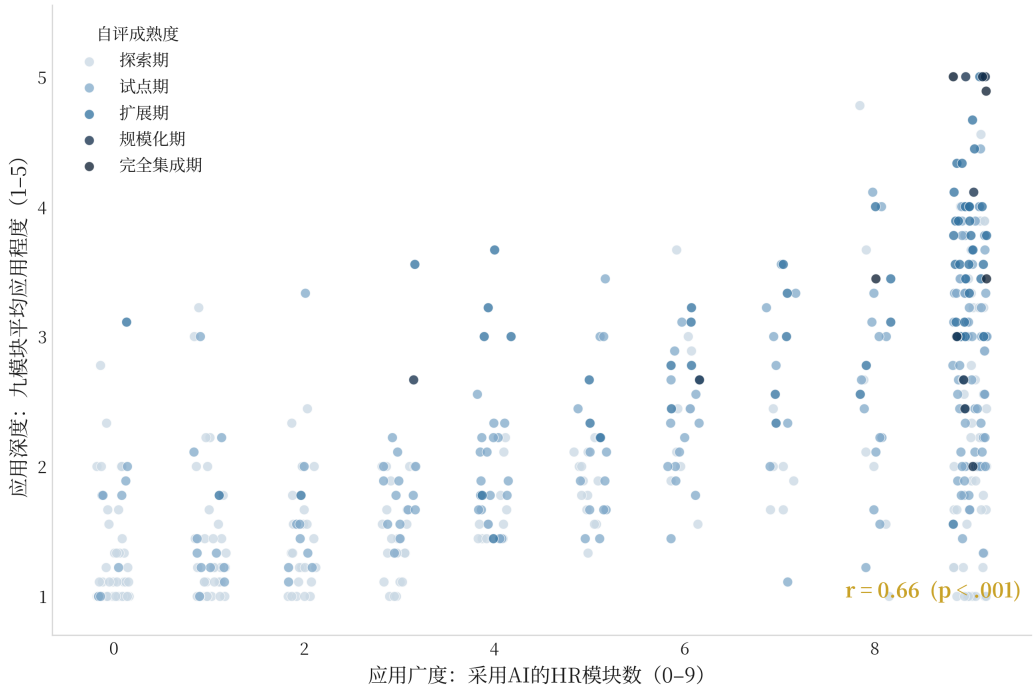
4.3 广度与深度的关系

为刻画企业层面的应用形态，我们以采用人工智能的模块数量衡量广度，以九模块应用程度评分的均值衡量深度。两者显著正相关，但联合分布显示多数企业集中在广度较高而深度居中偏下的区域：模块的铺开走在了单模块深耕的前面。结合成熟度自评的分层着色可以看到，真正进入扩展及以上阶段的企业同时具备高广度与高深度，而大量处于试点阶段的企业虽已触及多个模块，单个模块的应用程度仍停留在初步探索层次。

图 4-5

广度先行、深度滞后：多数企业铺开了模块，但单模块应用仍停留在浅层

第三轮企业层面应用广度与深度的联合分布（n=526）



数据来源：AIHRM第三轮调查。深度=Q17九模块1–5级应用程度评分均值。

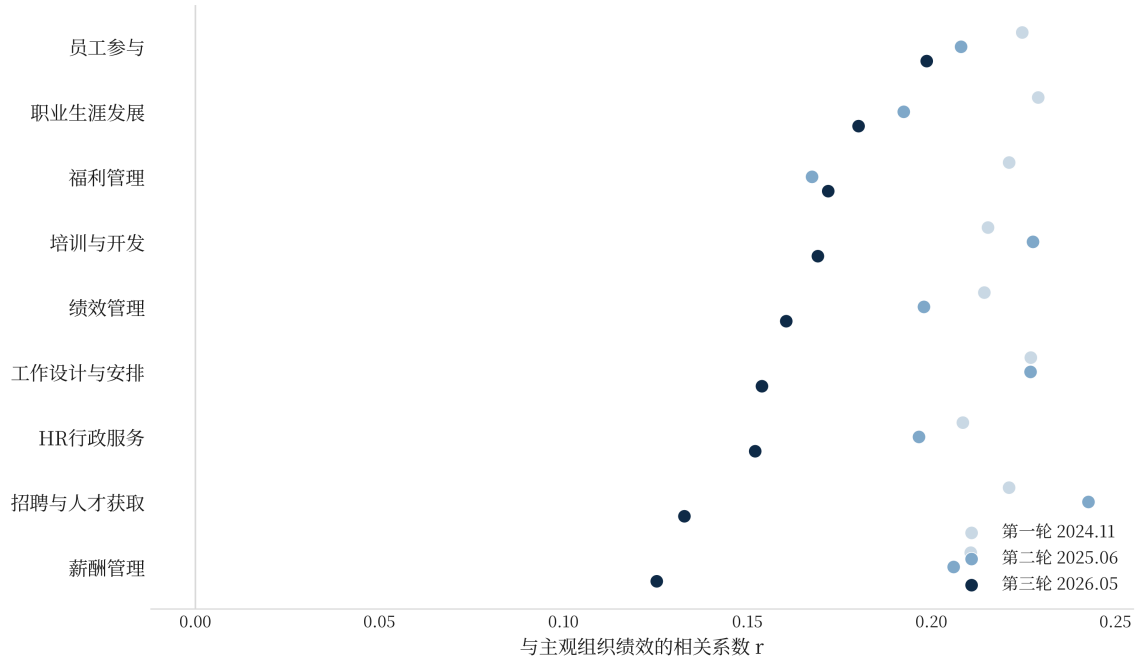
4.4 模块应用与绩效感知的关联

将各模块的应用数量与企业主观绩效评价进行相关分析，三轮结果呈现一致的正向关联，其中招聘与培训模块的相关最为稳定，在三轮中均达到统计显著。这一结果不宜作因果解读，绩效更好的企业可能本就更有余力投入人工智能；但模块间相关强度的差异仍具参考价值，它与第八章投资回报感知中招聘与行政位居回报前列的结果相互呼应，提示当前阶段人工智能的价值兑现集中在流程效率类模块。

图 4-6

AI应用与绩效感知普遍正相关，招聘与培训模块的关联最为稳定

各模块AI应用数量与主观组织绩效的相关（大点＝p<.05）



数据来源：AIHRM三轮追踪调查。相关不代表因果；绩效为与同行相比的主观评价。

4.5 解读与启示

模块层面的数据为投入排序提供了一条经验规则：人工智能的价值优先在任务标准化、文本密集、决策风险可控的环节兑现。招聘、人力资源行政与工作设计三个模块同时占据应用渗透率与回报感知的前列，并非偶然，它们的共同特征是流程清晰、样本量大、错误可纠正。相反，薪酬与绩效直接关联利益分配与合规风险，企业的审慎不是保守，而是理性的风险定价。管理者在制定路线图时可以据此排序：先在前一类模块建立能力与信任，再逐步向后一类模块延伸，而不是从最敏感的环节开始证明技术价值。

广度与深度的分布则揭示了一个普遍的资源配置误区。数据显示多数企业集中在广度较高而深度不足的区域，这往往源于一种可以理解的压力：管理层希望看到人工智能在各个模块都有动作，于是资源被摊薄到九个方向，每个方向都停留在演示与试用层次。真正进入扩展阶段的企业则表现出相反的形态，即在少数模块上同时具备广度与深度。对人力资源负责人的具体建议是，重新审视当前的项目清单，如果九个模块都有在建项目而没有任何一个模块完成了流程重构，那么合理的动作是收缩战线，将资源集中到一到两个已验证回报的场景，把它做到能够替代原有流程的程度，形成内部标杆之后再谈复制。

前沿应用的渗透格局提供了第三个提示。嵌入办公软件的辅助类工具扩散最快，而由人工智能自主编排流程的智能体应用仍处于早期。这一顺序符合技术采纳的一般规律：增强人的工具比替代人的工具更容易被组织接受，因为前者不改变责任归属。管理者在引入智能体类应用时应当预期到更长的接受周期与更复杂的流程改造，并将责任边界的界定作为方案设计的组成部分，而非上线之后的补充说明。

05

第 五 章

采纳的驱动因素

组织条件与人力资源负责人

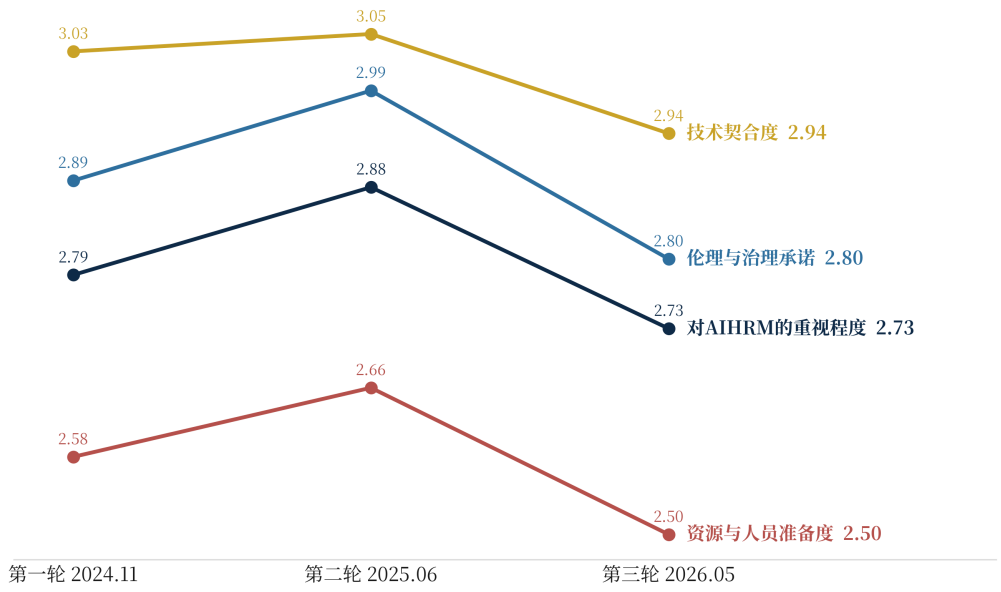
5.1 组织心态与准备度的三年曲线

技术—组织—环境框架下的四项核心指标在三轮调查中走出了先升后落的曲线。对AIHRM的重视程度从第一轮的2.79分升至第二轮的2.88分，第三轮回落至2.73分；伦理与治理关注、技术契合度与资源准备度呈现相同形态，其中资源准备度在三轮中始终处于四项指标的最低位，第三轮为2.50分。这一集体回调发生在采纳率持续上升的背景之下，两者的背离值得认真解读：随着实际部署的展开，企业对自身准备状况的评估从概念期的乐观转向落地期的审慎，行业整体进入了一个更为冷静的阶段。第三轮样本构成的变化可能贡献了部分回落，但匹配面板的复核显示方向一致。

图 5-1

热度先升后落：第三轮各项准备度指标集体回调，行业进入冷静期

TOE框架各维度均值（1–5分；第三轮样本口径变化亦可能贡献部分回落）



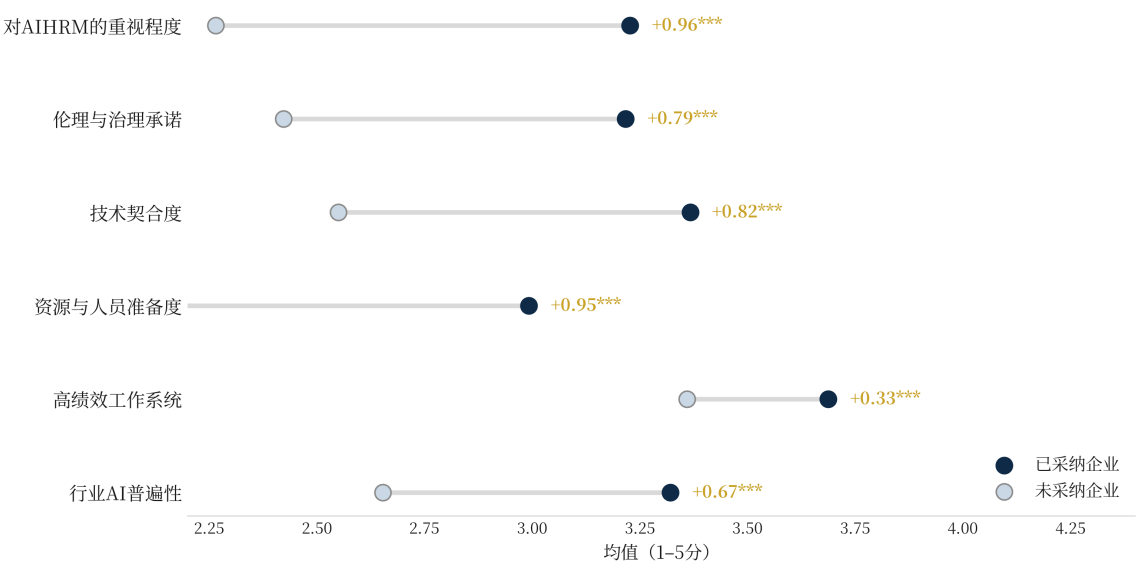
数据来源：AIHRM三轮追踪调查。

已采纳与未采纳企业的对比进一步定位了分化的关键。第三轮数据显示，两类企业在全部六项组织指标上均存在显著差距，而差距最大的两项是对AIHRM的重视程度和资源与人员准备度；技术契合度与伦理关注的差距相对较小。换言之，将采纳者与观望者区分开来的，与其说是认知与态度，不如说是实际的资源投入与人员能力储备。这一发现与第六章潜在剖面分析的结论相互印证，也构成了本报告管理建议的重要依据。

图 5-2

采纳者与未采纳者最大的落差在资源准备度，而非认知与态度

第三轮已采纳与未采纳企业在各维度上的均值差（* p<.05，** p<.01，*** p<.001）



数据来源：AIHRM第三轮调查。

5.2 高绩效工作系统与采纳：一个需要区分口径的关系

高绩效工作系统与采纳之间的关系，不能用简单的正向或负向来概括，因为两种统计口径给出的答案方向相反。就直接的相关而言，三轮数据均为正：高绩效工作系统得分较高的企业，整体上采纳比例也略高，图5-2中已采纳企业在该维度上的均值优势即是这一事实的体现。这并不难理解，体系成熟的企业往往规模更大、资源更充裕，而规模与资源本身就与采纳正相关。但课题组前两轮的回归分析显示，在控制企业规模、重视程度与准备度等因素之后，高绩效工作系统的系数转为显著为负。两个口径的背离意味着，负向关系是一个条件效应：在其他条件相当的企业之间，人力资源体系越成熟，采纳的边际倾向反而更低。对此的一种解释是路径依赖，成熟体系内含既定流程与既得能力，降低了变革的相对收益。第三轮的综合模型中，这一条件效应方向仍为负但已不再显著（见图5-8），提示随着工具可用性提高与同行示范效应累积，成熟体系的抑制作用正在减弱。

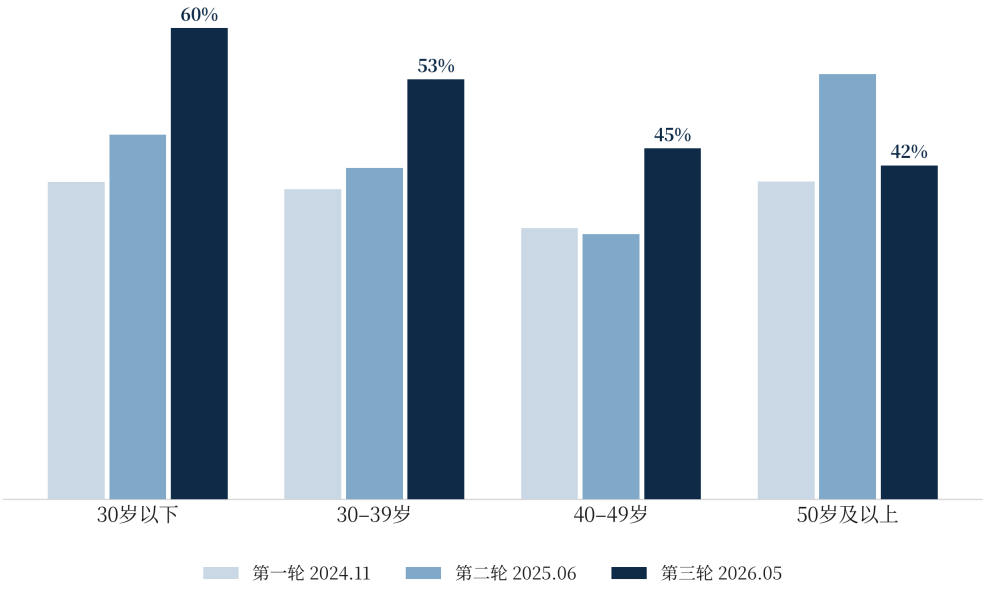
5.3 人力资源负责人特征：年龄与专业背景

企业人力资源负责人的个体特征与采纳率之间存在稳定的关联。按年龄分组的统计显示，负责人越年长，所在企业的采纳率越低，这一梯度在三轮调查中方向一致；按专业背景分组的统计显示，兼具工程技术、财务、战略等跨界背景的负责人所在企业，采纳率持续高于仅有人力资源单一背景的负责人所在企业，差距大体保持在十个百分点上下。

图 5-4

掌舵人年龄效应三轮稳健：HR负责人越年长，企业采纳率越低

按HR负责人年龄组的AIHRM采纳率（人口学特征在匹配样本内跨轮沿用）

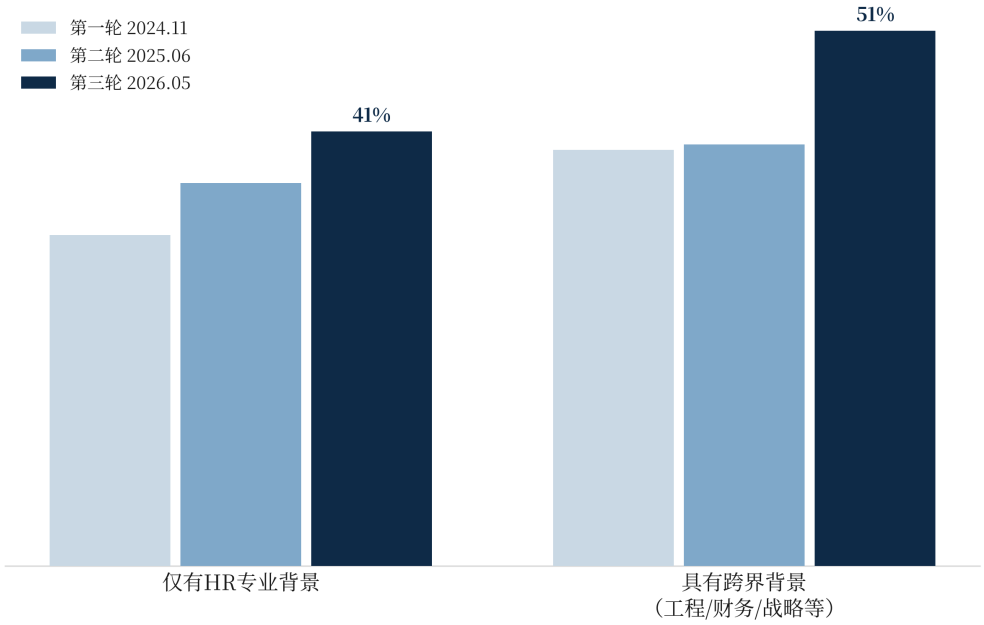


数据来源：AIHRM三轮追踪调查。第二轮问卷仅向新参与者采集人口学信息，重复参与者信息由匹配轮次补齐。

图 5-5

跨界背景的HR负责人所在企业采纳率持续领先，差距在10个百分点上下

按HR负责人专业背景的AIHRM采纳率



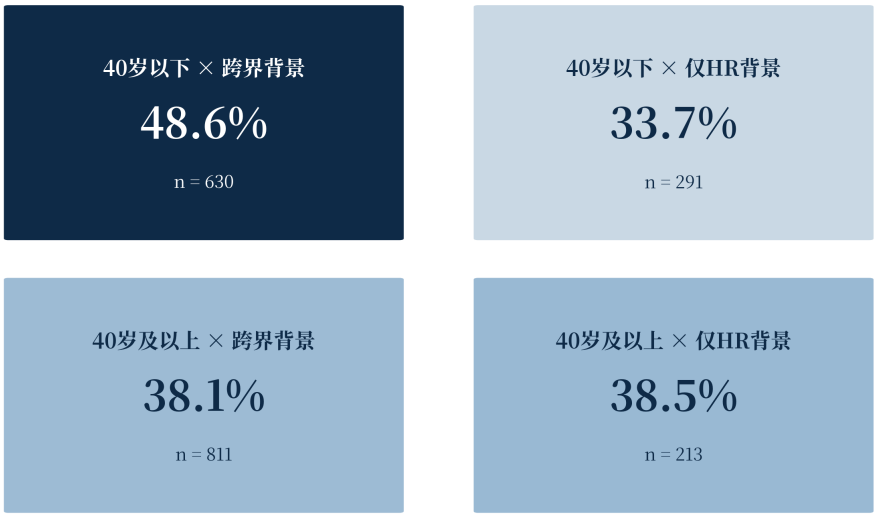
数据来源：AIHRM三轮追踪调查。跨界背景=拥有工程技术/财务/法律/营销/战略/运营等任一HR之外背景。

将年龄与背景交叉后，图景变得更加立体。三轮合并样本的四组比较显示，四十岁以下且具备跨界背景的负责人所在企业采纳率达48.6%，为四组最高；四十岁以下但仅有人力资源背景的组别采纳率为33.7%，为四组最低；四十岁及以上的两个组别居中且彼此接近。Logistic回归确认年龄、背景的主效应及两者的交互项均达统计显著，交互的含义是：跨界背景带来的采纳优势集中体现在年轻一代负责人身上，而在年长群体中背景差异基本消失。对这一格局的一种解读是，年轻且跨界的负责人兼具变革意愿与跨域知识，构成了组织内推动人工智能议程的最积极力量；单一专业成长路径的局限，则在最需要跨界判断的年轻管理者群体中表现得最为明显。

图 5-6

“年轻且跨界”的HR负责人是AI采纳的最强引擎，与最低组相差逾15个百分点

HR负责人年龄 × 专业背景四组企业的AIHRM采纳率（三轮合并）



数据来源：AIHRM三轮追踪调查合并样本。Logistic回归显示年龄与专业背景效应显著（p<.05）。

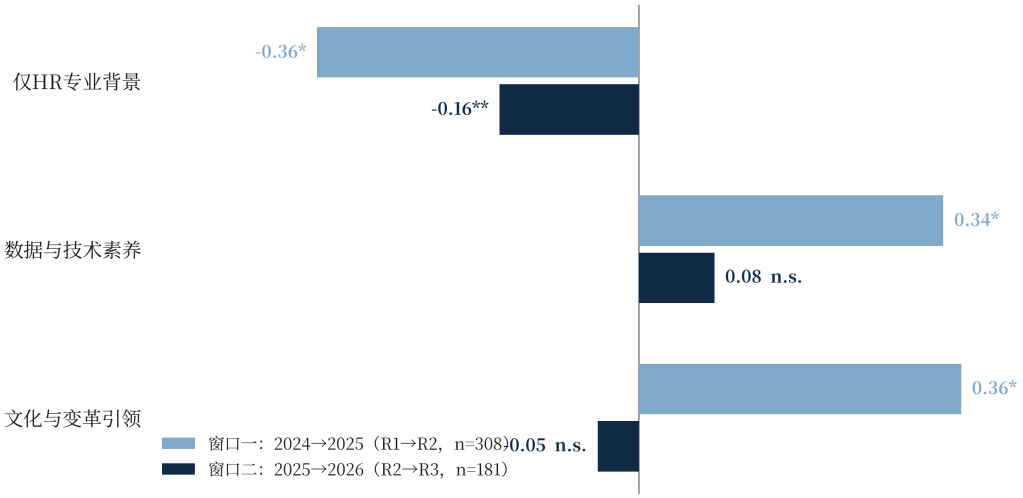
5.4 滞后检验：两个时间窗口的重复验证

横截面关联难以排除反向因果与遗漏变量的干扰，为此我们利用追踪设计进行了滞后检验。第一个时间窗口基于第一轮与第二轮的匹配企业，在控制第一轮采纳水平后考察负责人特征对第二轮采纳的预测；课题组前期分析显示，仅有人力资源背景的系数为-0.36，负责人的数据与技术素养系数为0.34，文化与变革引领能力系数为0.36，三者均达显著。本报告利用第二轮与第三轮的匹配数据构建了第二个独立窗口，在控制第二轮采纳水平、企业规模与高绩效工作系统后，仅有人力资源背景对第三轮采纳的标准化系数为-0.16，在千分之一水平上显著，与第一个窗口的结论完全一致；数据与技术素养的系数方向为正但未达显著，文化与变革引领接近于零。综合两个窗口，专业背景的抑制效应是掌舵人特征中最为稳健的一项，跨越两个年度、两组不同企业样本均得到重复；胜任力维度的预测力则有待更长的观察窗口与更大的匹配样本加以确认。

图 5-7

“仅HR背景”的抑制效应在两个独立时间窗口均显著，是最稳健的掌舵人效应

控制上一轮采纳水平后，CHRO特征对下一轮采纳的标准化系数（滞后设计）



数据来源：AIHRM追踪调查匹配面板。窗口一引自课题组前期分析；窗口二为本报告新增，控制企业规模与HPWS。
能力维度在窗口二方向为正但未达显著，可能与样本量和区间截断有关；† p<.10，* p<.05，** p<.01，*** p<.001。

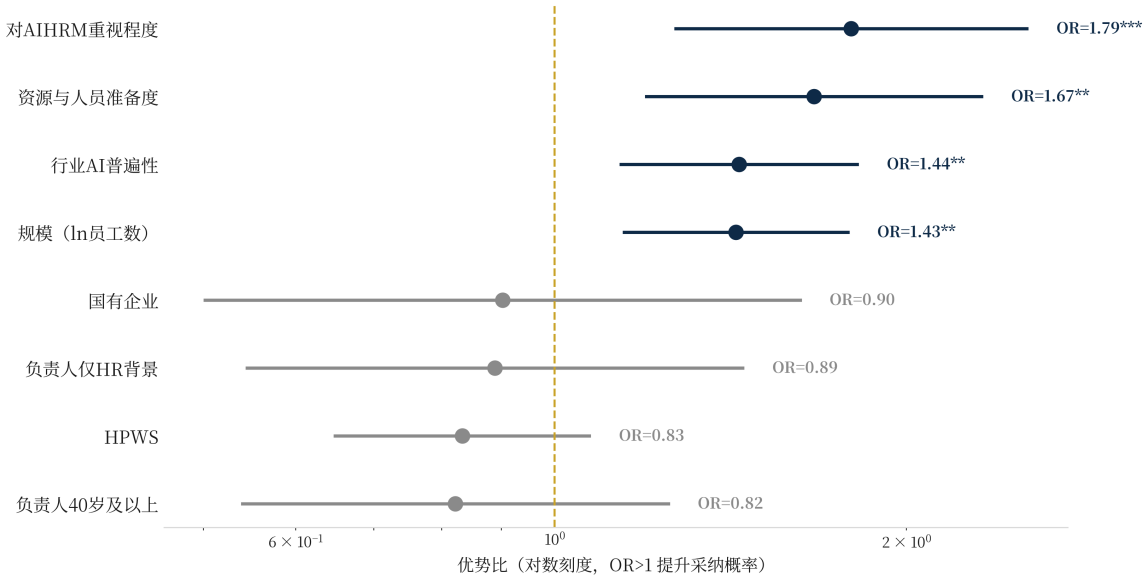
5.5 综合模型：直接效应与传导路径

将组织因素与负责人特征同时纳入第三轮的Logistic回归，可以观察各因素的相对强度。结果显示，对AIHRM的重视程度与资源准备度的优势比最大，每提高一个标准差分别使采纳概率的优势提升约79%与67%；企业规模与行业人工智能普遍性次之，均达显著。负责人的年龄与专业背景在纳入组织变量后不再显著，结合前述滞后检验中背景效应的稳健存在，一个与数据相容的解释是：负责人特征对采纳的作用主要经由组织状态传导，年轻且跨界的负责人通过提升组织对人工智能的重视程度与资源投入，间接推动了采纳，横截面模型中的组织变量吸收了这一间接路径。这一传导机制的严格检验需要更长的面板，但现有证据已足以支持一个管理判断：改变组织的准备状态是负责人影响采纳进程的现实杠杆。

图 5-8

重视程度、准备度、规模与行业环境构成采纳的直接引擎；掌舵人特征主要经由组织准备传导

AIHRM采纳的Logistic回归优势比（第三轮，n=515；连续变量已标准化）



数据来源：AIHRM第三轮调查。灰色=不显著。负责人年龄与背景在纳入组织变量后不再显著，提示其横截面效应主要通过重视程度与准备度等组织状态间接实现；滞后面板中的直接效应见图5-7。

5.6 解读与启示

本章的核心发现，是采纳者与观望者之间的差距主要落在资源准备度而非认知态度上。这意味着人力资源负责人在向管理层争取支持时，论证的重点需要调整：在多数企业里，高管并不需要被反复论证人工智能重要，重视程度的分歧远小于资源投入的分歧。更有效的提案不是再讲一遍趋势，而是把准备度上的具体不足说清楚，包括数据的可用性、系统的可对接程度、团队的技术理解力与专职人力的配备，并为每一项不足标注补齐的成本与时间。第三轮准备度均值仅2.50分、在四项指标中垫底，说明这不是个别企业的问题，而是行业的共性短板，谁先补齐，谁就获得了实质性的先发身位。

高绩效工作系统的条件性负向效应在第三轮明显减弱，这一变化对体系成熟的企业尤其重要。前两轮的回归分析曾提示，在规模与资源相当的企业之间，人力资源体系越成熟，推进人工智能的边际倾向反而越低，路径依赖是合理的解释；第三轮该效应已不再显著。对体系成熟的企业而言，这是一个可以重新入场的信号：既有流程的规范化程度本身正在从阻力转为资产，因为标准化的流程与积累的数据恰恰是人工智能应用的前提。这类企业过去的迟疑有其道理，继续迟疑则难有理由。

人力资源负责人个体特征的相关发现需要谨慎而诚实地使用。数据显示年轻且具备跨界背景的负责人所在企业采纳率显著领先，仅有单一人力资源专业背景的负责人所在企业在两个独立时间窗

口中均表现出更低的采纳增量。这一结论的正确用法不是把它变成用人上的年龄标签，年龄本身不是能力，四十岁以上群体内部的背景差异也已基本消失。更合理的解读是，这些个体特征标记的是跨领域知识与变革推动经验的可得性，而这两者是可以通过配置与培养获得的。对企业的建议是双轨并行：在选配层面，将数据素养与跨界经历纳入人力资源负责人的任职考量；在培养层面，为在任负责人提供系统的技术理解与数据分析训练，并在人力资源团队中引入具备技术背景的成员，以团队的知识结构弥补个人的背景局限。对人力资源负责人本人而言，这组数据是一份直接的提示：技术与数据能力已经不再是加分项，而是这一职位的基础配置。

综合模型中负责人特征在控制组织变量后不再显著，这一结果不应被读作负责人的作用不重要，而应被读作其发挥作用的路径已经清晰。负责人不是通过个人偏好直接决定采纳，而是通过塑造组织对人工智能的重视程度、争取资源与推动准备工作来影响结果。这恰恰是一条可操作的路径：改变组织的准备状态，是负责人手中最现实的杠杆。

06

第六章

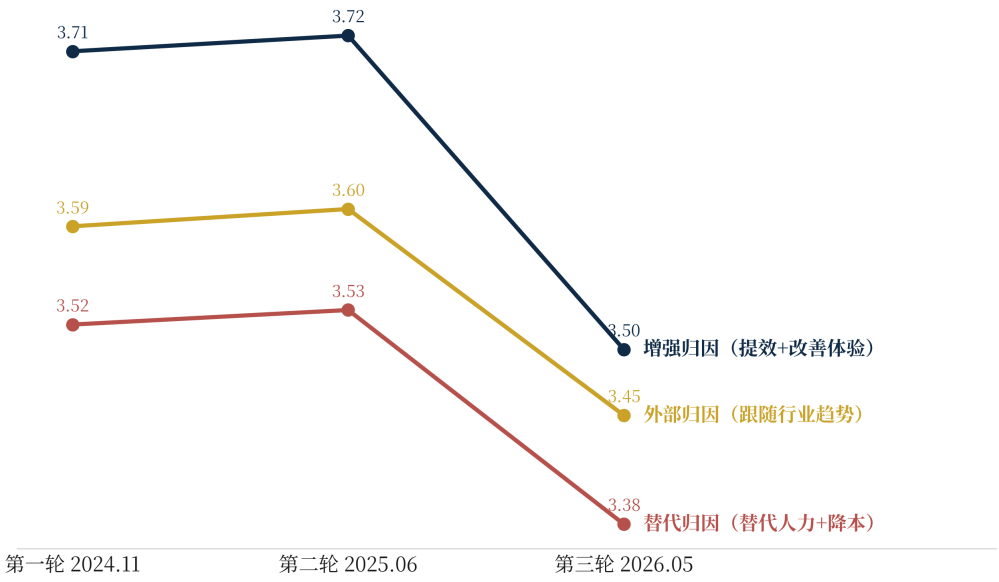
动机归因与企业画像

分化在准备度，而非动机

6.1 三类归因的同步降温

企业采用人工智能的动机可以归为三类：以提升效率与改善员工体验为核心的增强逻辑，以替代人力与压缩成本为核心的替代逻辑，以及跟随行业趋势的外部压力。三轮数据显示，三类归因的均值同步回落，排序保持稳定，增强归因始终居首。归因结构的整体降温与第五章观察到的准备度回调同属一个现象：概念期的高预期正在让位于落地期的审慎评估，企业对为什么使用人工智能的回答变得更加克制，也更加务实。替代归因未见相对上升，说明至少在受访负责人的表述层面，成本替代尚未成为主导动机，这与第八章观察到的编制收缩现象之间的张力，将在结论章节讨论。

图 6-1
三类归因同步降温：增强逻辑始终占主导，但企业整体进入理性评估期
企业采用AIHRM的动机归因均值（1–5分）



数据来源：AIHRM三轮追踪调查。

6.2 四类企业画像的识别

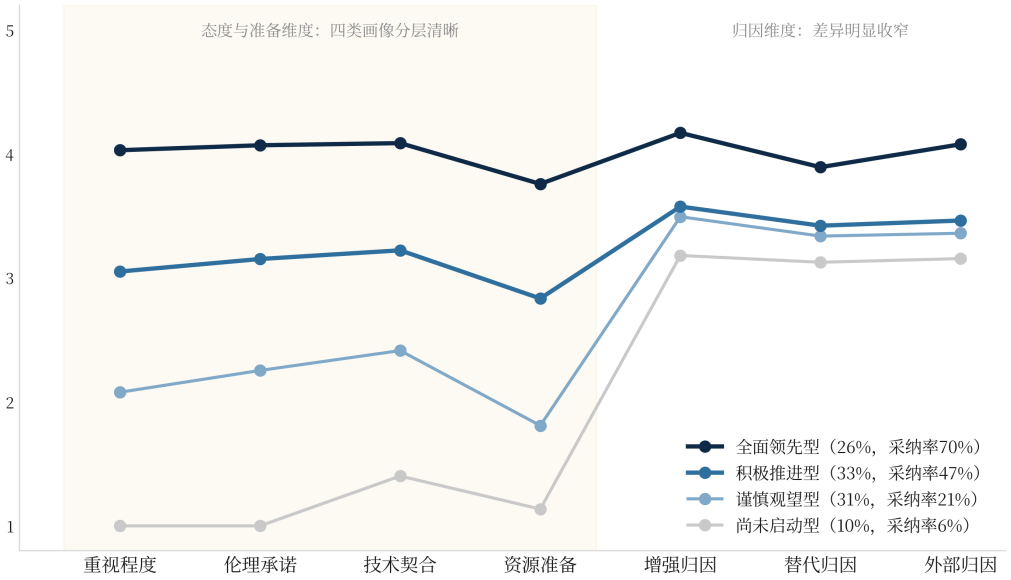
为超越单变量描述，我们对三轮合并的2,238个企业观测进行了潜在剖面分析，输入变量包括重视程度、伦理关注、技术契合、资源准备四项状态指标与三类归因指标，采用高斯混合模型估计。综合模型拟合与可解释性，四剖面解被确定为最终方案。第一类企业在全部状态指标上均处高位，占样本的26%，其采纳率达70%，可称为全面领先型；第二类各项指标中等偏上，占33%，采纳率47%，为积极推进型；第三类指标中等偏下，占31%，采纳率21%，为谨慎观望型；第四类在状态指标上接近量表底部，占10%，采纳率仅6%，为尚未启动型。

画像的判别结构揭示了一个关键事实：四类企业在归因维度上的差异远小于在状态维度上的差异。谨慎观望型企业同样认可增强逻辑，其归因得分与领先型企业相去不远，真正拉开距离的是资源准备与重视程度。这意味着中国企业在是否应当拥抱人工智能的问题上已接近共识，分化发生在能否组织起相应的资源与能力，采纳的瓶颈是执行问题而非认知问题。

图 6-2

潜在剖面分析识别出四类AIHRM心智画像：分化在于准备度，而非动机

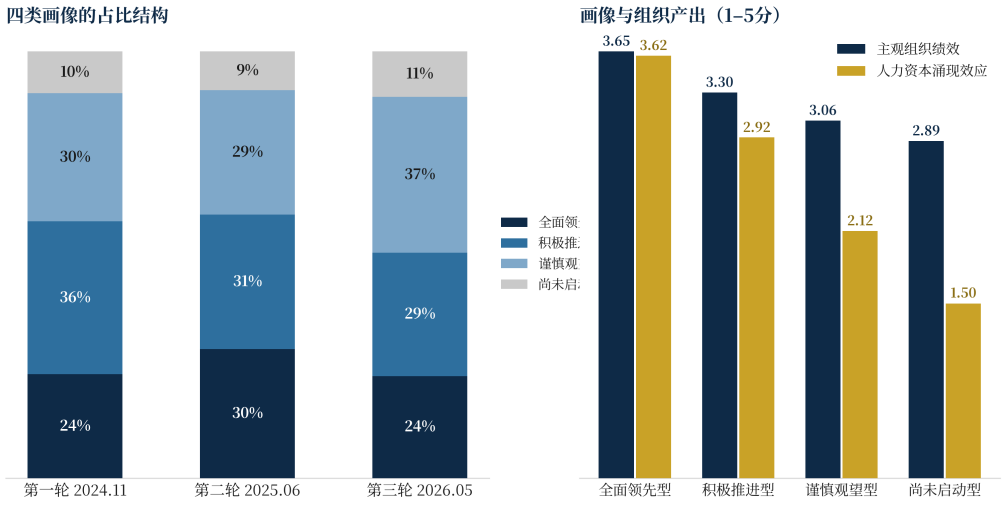
各画像在七个维度上的均值特征（高斯混合模型，三轮合并 n=2,238）



数据来源：AIHRM三轮追踪调查。模型基于标准化变量估计，图中展示原始量表均值；四剖面解兼顾拟合与可解释性。

画像结构在三轮之间保持了较高的稳定性，各类型占比的变动不超过三个百分点，表明这一分类刻画的是企业间相对持久的位置差异。产出维度的梯度则十分陡峭：从尚未启动型到全面领先型，主观绩效与人力资本涌现效应逐级抬升，领先型企业在两项产出上均显著占优。画像因此可以作为一个简明的诊断框架：企业可以对照四类特征定位自身位置，并将追赶的重点放在状态指标的短板上。

图 6-3
画像结构三年基本稳定，但产出梯度陡峭：领先型的绩效与涌现效应全面占优



数据来源：AIHRM三轮追踪调查。

6.3 解读与启示

四类画像的判别结构给出了一个反直觉但重要的结论：中国企业在为什么要用人工智能这个问题上已经接近共识，谨慎观望型企业对增强逻辑的认可程度与全面领先型企业相去不远。分化发生在能否组织起资源与能力。这意味着企业内部关于人工智能的争论，如果仍停留在理念层面，多半是在解决一个已经不存在的问题。管理层的会议议题应当从要不要做转向缺什么、谁来补、多久补齐。

画像框架可以直接用作自我诊断工具。人力资源负责人可以对照四类特征的维度得分定位本企业的位置：如果各项状态指标均处中等偏下而归因得分不低，属于典型的谨慎观望型，短板在准备而非意愿，对策是集中资源补齐数据与人员基础，而非继续做可行性论证；如果状态指标中等偏上，属于积极推进型，其风险是止步于试点，对策是选定场景做深并建立可量化的评估机制；如果已处全面领先型，真正的挑战转向治理与组织公平，这正是第七章与第八章讨论的议题。

画像结构三年高度稳定，这一事实本身值得管理者警觉。它说明企业在这一分布中的相对位置具有粘性，落后者并不会因为行业整体上行而自动追上，位置的改变需要主动的、结构性的投入。与之相对，产出维度的梯度十分陡峭，从尚未启动型到全面领先型，绩效与涌现效应逐级抬升。位置的粘性与产出的梯度结合起来，指向一个务实的判断：追赶的窗口存在，但不会自动打开。

07

第 七 章

挑战与治理

应用跑在了制度前面

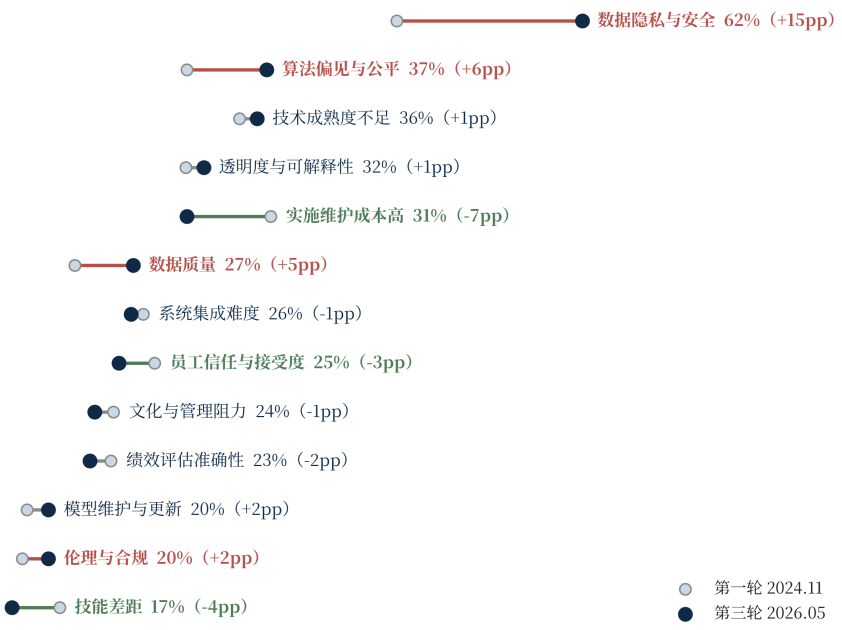
7.1 实施挑战的变化

对实施挑战的追踪显示，企业面临的困难清单正在发生结构性位移。数据隐私与安全的提及率从第一轮的47.3%升至第三轮的61.8%，以接近十五个百分点的涨幅居各项挑战之首；算法偏见与公平的提及率从30.8%升至37.1%，数据质量从22.0%升至26.6%。与此同时，实施与维护成本的提及率从37.4%回落至30.8%，技能差距的提及率也有所下降。这一升一降勾勒出清晰的阶段转换：随着工具价格下降与使用经验积累，成本与技能类的入门障碍正在缓解，而隐私、偏见与数据质量这类随应用深入而放大的内生风险，正在成为企业的主要顾虑。

图 7-1

数据隐私跃升为最突出挑战，算法偏见与数据质量的关注上升，成本顾虑回落

提及各项挑战的企业比例：2024年第一轮 vs 2026年第三轮（红=上升逾2pp，绿=回落逾2pp）



数据来源：AIHRM追踪调查第一轮与第三轮（多选题）。pp=百分点。

7.2 治理建设的滞后

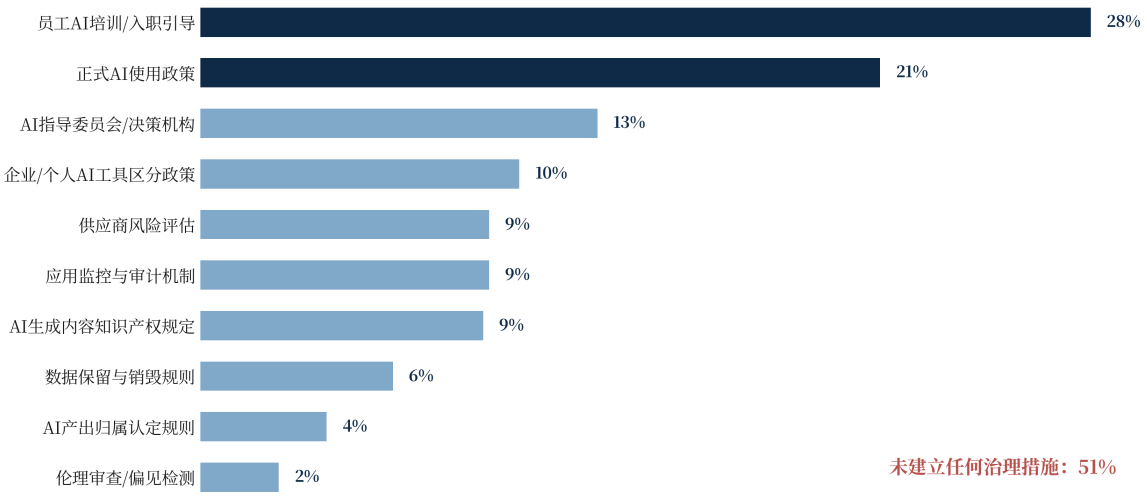
与风险感知的上升形成鲜明对照的，是治理体系建设的普遍缺位。第三轮调查列出了十类常见的人工智能治理措施，结果显示51%的企业尚未建立其中任何一项；覆盖率最高的员工人工智能培训也仅有28.1%的企业实施，正式的使用政策覆盖21.5%，设立专门指导委员会或决策机构的企业占

12.5%，伦理审查与偏见检测的覆盖率不足3%。将本章两项发现并置可以看到一个明显的缺口：企业清楚地知道风险在哪里，却大多尚未将这种认知转化为制度安排。考虑到86%的企业仍处于探索与试点阶段，治理建设目前尚有时间窗口，但窗口会随着应用规模化而快速关闭。

图 7-2

治理真空：过半企业尚未建立任何AI治理措施，单项覆盖率最高不足三成

已建立各项AI治理措施的企业比例（第三轮，n=526，多选）



数据来源：AIHRM第三轮调查。

7.3 解读与启示

实施AIHRM的挑战的变化为管理者提供了一份更新过的风险清单。成本与技能这类入门障碍的提及率在回落，说明进入门槛确实在下降；而数据隐私、算法偏见与数据质量的关注度在上升，这三项的共同特点是随应用深入而放大，属于内生风险。这一转换的管理含义是，人工智能项目的风险评估不能沿用立项时的清单。企业在试点阶段关心的是买不买得起、会不会用，进入规模化阶段后真正决定成败的是数据能不能用、结果能不能解释、责任由谁承担。人力资源负责人应当在项目从试点转向推广的节点上，主动发起一次风险清单的重估。

治理数据揭示的缺口更为紧迫。超过六成的企业把数据隐私列为首要挑战，却有51%的企业没有建立任何一项治理措施，覆盖率最高的员工培训也仅有28.1%，伦理审查与偏见检测不足3%。风险认知与制度安排之间的这道裂缝，是本报告观察到的最明显的管理短板。需要强调的是，这不是一个可以拖延的议题：治理措施的建设成本在试点阶段远低于规模化之后，一旦人工智能嵌入到招

聘筛选或绩效评估的正式流程，再回头补建审查机制，涉及的将是既成流程的推翻重来，以及可能已经发生的合规后果。

对人力资源负责人的具体建议是，把治理建设作为规模化的前置条件而非并行选项，并从成本最低、争议最小的几项入手：一是明确的使用政策，界定哪些场景可以使用人工智能、哪些决策必须保留人工复核；二是数据边界的界定，明确哪些员工数据可以进入模型、留存多久、如何销毁；三是面向使用者的培训，这是当前覆盖率最高也最容易推进的一项；四是针对招聘与绩效等高风险场景的偏见检测机制。当前86%的企业仍处于探索与试点阶段，这个时间窗口恰恰是建立治理框架的最佳时机，也是成本最低的时机，但窗口会随着应用规模化而快速关闭。

08

第八章

组织后果

回报、差距与重构

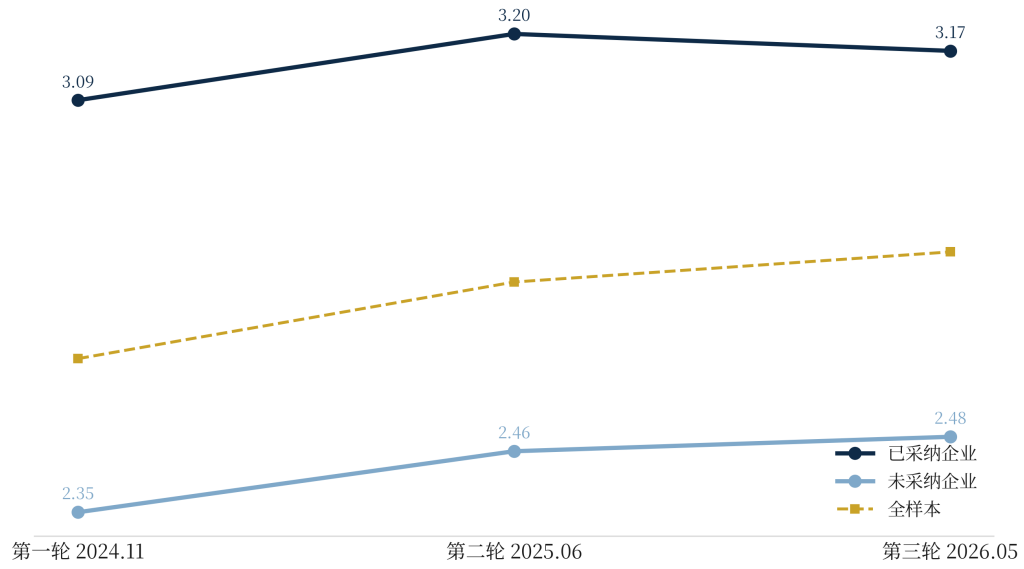
8.1 人力资本涌现效应的持续增强

在多数态度指标回调的背景下，衡量生成式人工智能促进个人知识与技能向团队综合能力转化的涌现效应量表，是少数逐轮上行的指标，全样本均值从第一轮的2.62分升至第三轮的2.81分。已采纳企业的得分在三轮中均显著高于未采纳企业，且两者差距未见收窄。涌现效应的稳步抬升提供了一个与情绪指标不同的信号：在实际使用人工智能的组织中，个体层面的能力增益正在向团队层面沉淀，这正是人力资本理论视角下技术投入转化为组织能力的关键环节。

图 8-1

生成式AI的人力资本涌现效应逐轮增强，是全书少数持续上行的产出指标

“个人知识转化为团队能力”八项量表均值（1–5分）



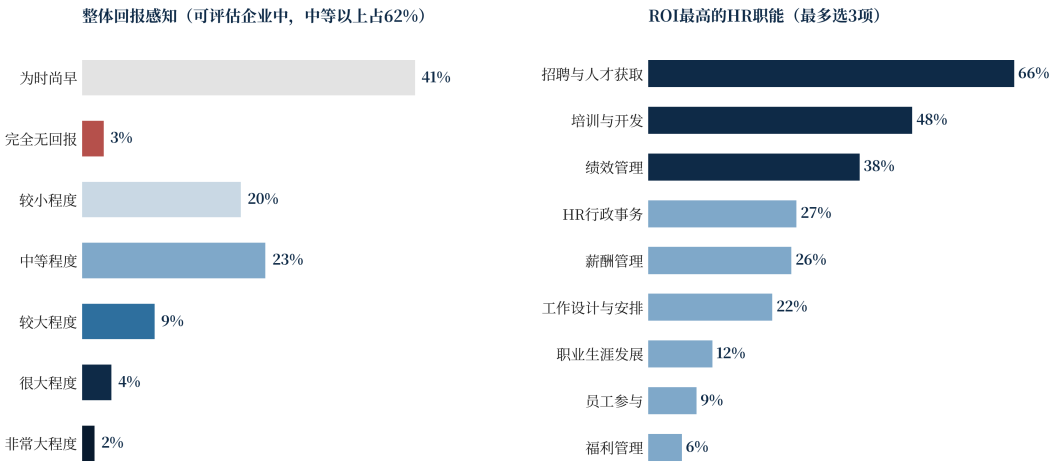
数据来源：AIHRM三轮追踪调查。量表衡量生成式AI促进个人知识技能向团队综合能力转化的程度。

8.2 投资回报：观察期中的初步兑现

第三轮对投资回报的直接测量显示，回报评估仍处于观察期：41.1%的企业认为下结论为时尚早。在能够作出评估的企业中，感知到中等及以上回报的合计约六成，认为完全没有回报的企业占比很小。分职能看，招聘与人才获取和人力资源行政事务等被最多企业列为回报最高的领域，与第四章模块渗透率的第一梯队高度重合。回报的兑现路径由此清晰可见：人工智能率先在流程密集、任务标准化的环节产生可感知的效率收益，而在绩效、薪酬等判断密集的环节，价值验证仍需时间。

图 8-2

回报进入观察期：四成企业认为评估为时尚早，已可评估者中六成感知中等以上回报

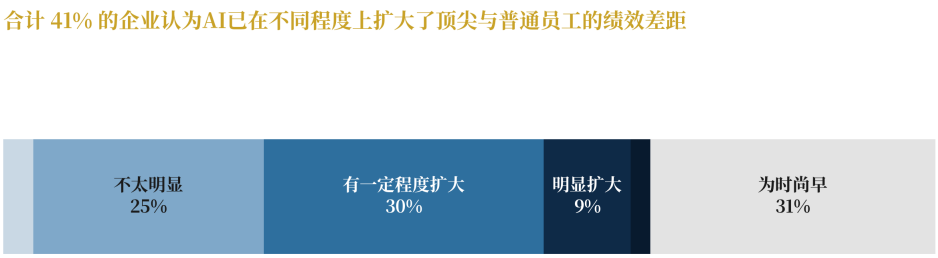


数据来源：AIHRM第三轮调查（n=526）。

8.3 绩效差距的扩大

人工智能的收益在员工之间的分布并不均匀。第三轮调查询问企业是否观察到人工智能应用扩大了高绩效员工与普通员工之间的差距，41.4%的企业给出了不同程度的肯定回答，其中约一成认为差距明显或大幅扩大。这一观察与近期学术研究中关于生成式工具放大能力回报的讨论方向一致：善于驾驭工具的员工获得了更大的产出杠杆。对管理者而言，这既是甄别高潜力人才的新信号，也提出了公平性管理的新课题，若放任差距自发扩大，可能侵蚀团队协作与薪酬体系的内部公平感。

图 8-3
AI正在放大“明星员工溢价”：近半企业观察到绩效差距扩大
“AI应用是否扩大了顶尖与普通绩效员工之间的差距”（第三轮，n=526）



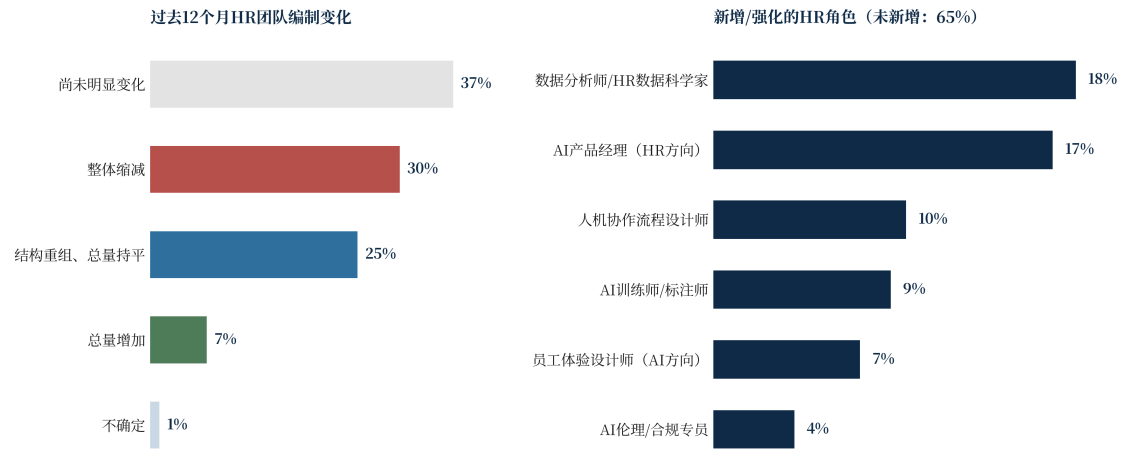
数据来源：AIHRM第三轮调查。

8.4 人力资源职能自身的重构

人工智能改变的不仅是人力资源管理的对象，也包括这一职能自身。过去十二个月中，30.2%的企业人力资源团队编制整体缩减，25.1%在总量持平下进行了结构重组，总量增加的企业仅占6.8%。与编制收缩并行的是角色更新：数据分析师是最多企业新增或强化的角色，人工智能产品经理、人机协作流程设计师等新型岗位开始出现。第二轮与第三轮对结构调整的同题比较显示，取消原有岗位、调整岗位职责、启动再培训与转岗计划的企业比例全面上升，组织结构层面的调整正在提速。收缩与更新的并行表明，人力资源职能正在经历的不是简单的规模削减，而是技能构成的置换。

图 8-4

编制收缩与角色更新并行：三成企业HR团队缩减，数据分析师是最常见新角色

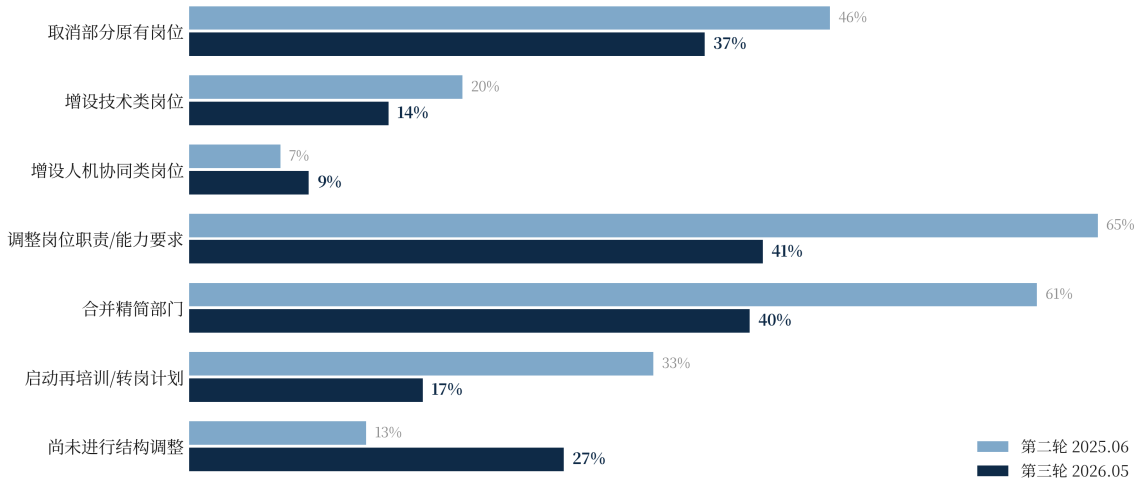


数据来源：AIHRM第三轮调查（n=526）。

图 8-5

组织结构调整全面提速：岗位职责重塑与再培训成为主要抓手

过去6个月进行过各类人力资源结构调整的企业比例（同题跨轮对比）

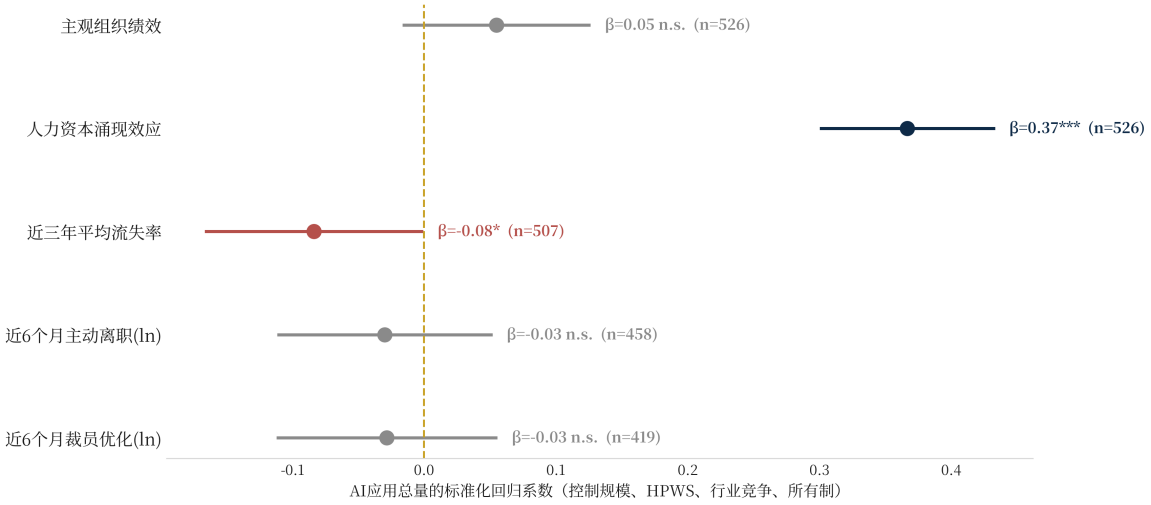


数据来源：AIHRM追踪调查第二轮与第三轮（多选题）。

8.5 应用深度与组织产出的关联

以人工智能应用总量为解释变量、控制企业规模、高绩效工作系统、行业竞争与所有制的回归分析显示，应用越深的企业，人力资本涌现效应与主观绩效评价越高，两项系数均达显著；与此同时，应用总量与员工流失率、主动离职人数及裁员优化人数也呈显著正相关。正向产出与人员流动的并存，与课题组前期基于两轮匹配数据结论一致：人工智能对减少人员的作用显现较快，对绩效提升的作用则需要更长时间验证。横截面数据不足以确定因果方向，流动加快既可能是应用的结果，也可能与同期的组织调整互为表里，但两类指标的同步上行足以提醒管理者，人工智能的落地过程同时是一个需要主动管理的人员转换过程。

图 8-6
AI应用越深，能力涌现与绩效感知越强，但人员流动同样更高
AIHRM应用总量对组织产出的关联（第三轮横截面，标准化系数与95%置信区间）



数据来源：AIHRM第三轮调查。横截面相关不代表因果；离职与裁员为对数化人数。

8.6 解读与启示

本章数据显示：人工智能的收益与代价是同时发生的。一方面，涌现效应逐轮增强，可评估回报的企业中约六成感知到中等以上回报，效率类模块的价值正在兑现；另一方面，应用越深的企业人员流动越高，三成企业的人力资源团队编制在收缩，四成企业观察到员工之间的绩效差距在扩大。管理者若只准备承接前者而未准备管理后者，转型过程中的组织成本会以人员流失与士气受损的形式显现。

回报评估的现实状况提示了另一个管理要点。41.1%的企业认为评估回报为时尚早，这一比例本身是一个信号：多数企业在投入人工智能时并未建立衡量标准。对人力资源负责人而言，与其在项目结束后追问回报几何，不如在立项时就确定基线与口径，例如单次招聘的周期与成本、行政事务

的人工工时、员工咨询的一次解决率。回报感知集中在招聘与人力资源行政这一事实也提供了参照：这两个环节的指标最容易量化，因而最容易被感知为有回报。选择可衡量的场景先行，本身就是提高转型可持续性的策略。

绩效差距的扩大是本章对管理者最直接的挑战。41.4%的企业观察到人工智能放大了高绩效员工与普通员工之间的差距，这既是识别高潜力人才的新信号，也是内部公平的新压力源。如果差距源于工具可得性的不均，解决方案是普及；如果源于使用能力的差异，解决方案是培训；如果源于任务性质本身对人工智能的敏感度不同，则需要在评价体系中作出校正，避免把技术红利误读为个人能力。放任差距自发扩大的风险在于，它会同时侵蚀协作意愿与薪酬体系的内部公平感。

人力资源职能自身的重构则要求负责人正视一个身份问题。编制收缩与角色更新并行发生，数据分析师成为最常见的新增角色，这说明这一职能经历的不是简单的规模削减，而是技能构成的置换。对在任团队的建议是尽早启动技能盘点，明确哪些工作会被工具吸收、哪些能力需要补强、现有成员的转型路径在哪里，并把再培训与转岗安排作为部署方案的组成部分与技术投入同步规划，而非在裁减决定之后才开始考虑。人力资源部门在这场转型中的特殊位置在于，它既是变革的设计者，也是变革的对象，这一双重身份决定了它的做法会被全公司当作样板来观察。

09

第 九 章

三年全程追踪面板数据的发现

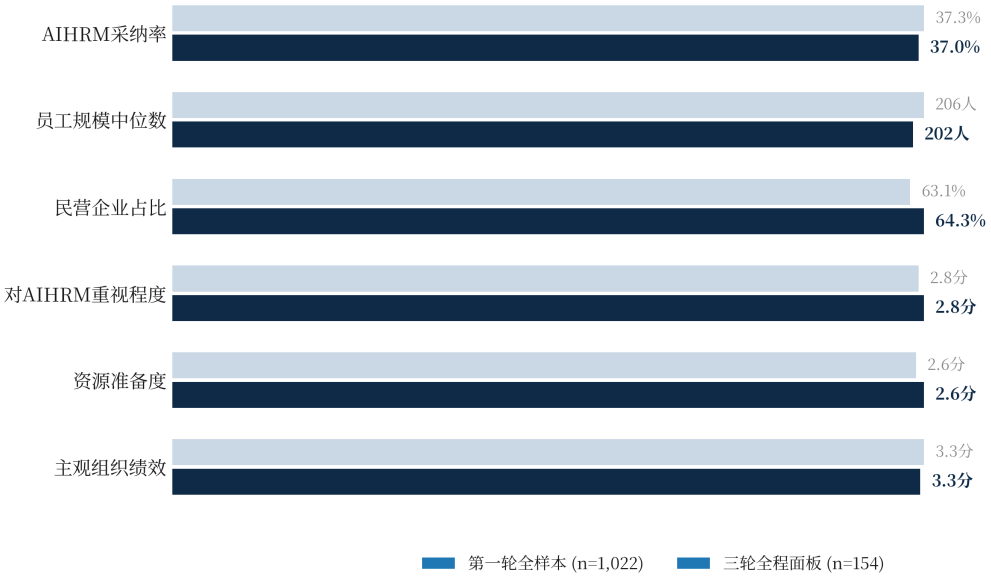
9.1 面板的代表性

三轮全程参与的154家企业构成了本项调查最珍贵的资产：对同一批企业的连续三次观测，可以将时代的变化与企业的变化区分开来。使用面板作推断的前提是其代表性，为此我们将面板企业与第一轮全样本在基线特征上进行了逐项对比。结果显示，两者在采纳率、员工规模、所有制构成、对AIHRM的重视程度、资源准备度与绩效水平上均高度接近，未发现统计显著的差异，面板不存在明显的自选择偏差，其结论具备向总体外推的基础。

图 9-1

154家全程追踪企业与总体样本高度相似，面板结论具备外推基础

面板企业与第一轮全样本的基线特征对比（以第一轮测量为准；条形按组内最大值归一）



数据来源：AIHRM三轮追踪调查。各指标均未见统计显著差异。

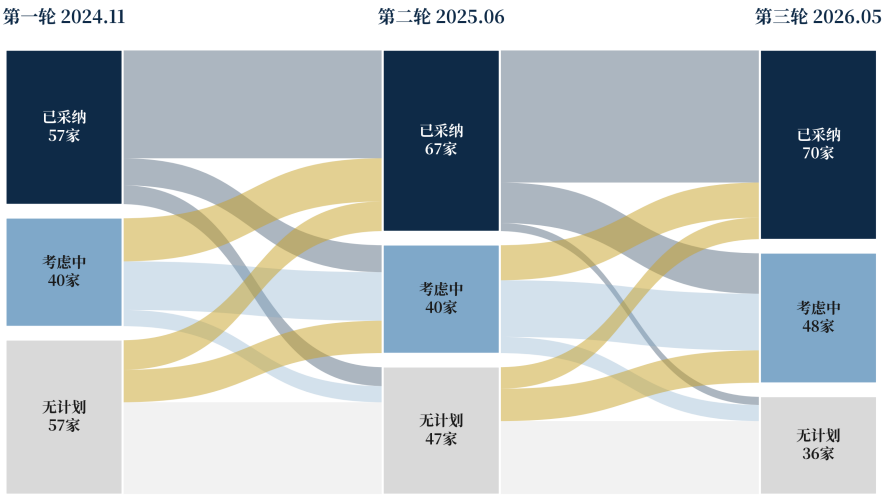
9.2 采纳状态的逐轮迁移

将面板企业的采纳状态区分为已采纳、考虑中与暂无计划三类，逐轮追踪其迁移路径。已采纳企业从第一轮的57家增至第二轮的67家、第三轮的70家；暂无计划的企业从57家降至36家。两个观察间隔中，共发生48次由考虑或无计划状态转入已采纳的跃迁，同期有35次反向退出。状态迁移呈现明确的净流入，但退出并非罕见，约每三次转入伴随两次退出，说明部分企业的早期采纳带有试探性质，在实践检验后选择了收缩。这一细节修正了扩散过程单向平滑的想象：总量曲线的平稳上升之下，是企业层面进进出出的活跃试错。

图 9-2

154家企业的三年迁徙：单向奔赴采纳，几乎无人回头

同批企业AIHRM采纳状态的逐轮转移（桑基图）



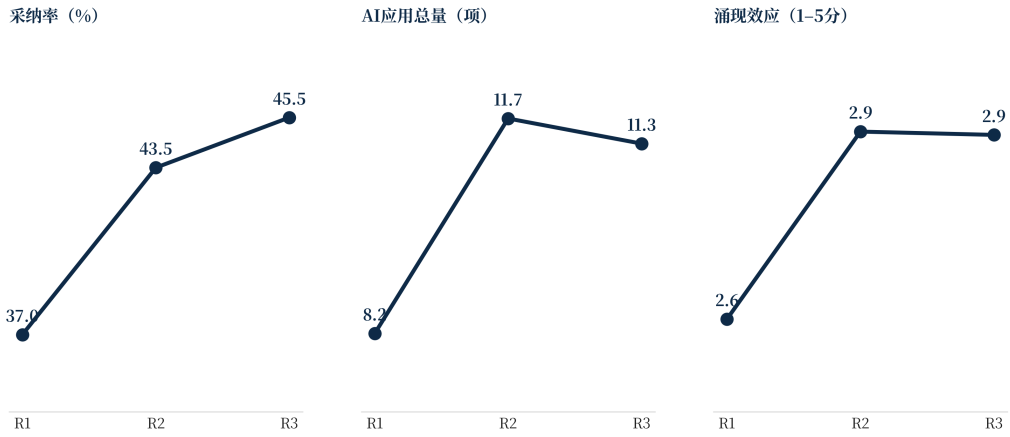
金色流=状态跃迁（向采纳靠拢）。两个间隔中共发生48次“转入采纳”，仅35次“退出采纳”，转型基本不可逆。

数据来源：AIHRM三轮全程追踪面板（n=154）。

9.3 同批企业的核心指标轨迹

组内配对检验为总体趋势提供了最严格的确认。154家企业的人工智能应用总量从第一轮的平均8.2项增至第三轮的11.3项，涌现效应均值从2.57分升至2.91分，两项提升均在千分之一水平上显著；采纳率从37%升至45%，受面板样本量限制，该提升为边际显著。应用总量在第二轮到第三轮之间的高位企稳同样值得注意，它与第五章的准备度回调相互印证：扩张最快的阶段可能已经过去，企业正在从增加应用数量转向消化与深化既有应用。

图 9-3
同一批企业的真实三年：采纳率、应用总量与涌现效应同步爬升
三轮全程面板企业（n=154）核心指标组内变化



数据来源：AIHRM三轮全程追踪面板。配对检验：AI应用总量与涌现效应R1→R3提升显著（p<.001），采纳率提升为边际显著（p=.069，受面板样本量限制）。

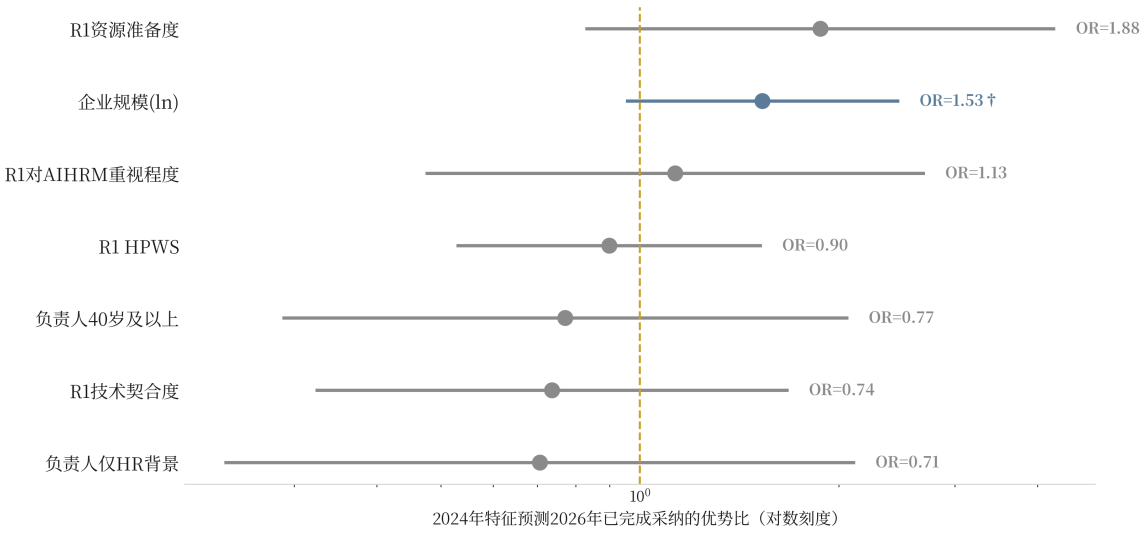
9.4 从观望到落地：谁完成了转化

面板设计允许我们回答一个前瞻性问题：在2024年尚未采纳的企业中，哪些特征预示了两年后的落地。对第一轮93家未采纳面板企业的Logistic回归显示，两年内的转化率为32.3%；第一轮的资源准备度与企业规模显示出方向性优势，准备度每高一个标准差，转化的优势比为1.88，规模的优势比为1.53，但受样本量限制，两者分别处于统计显著性的边缘之外与边缘水平，其余变量未见预测力。这一结果的稳妥解读是：早期的资源准备可能构成日后转化的先行条件，但基于93家企业的检验力有限，确证有待面板扩容后的第四轮数据。

图 9-4

从观望到落地：早期资源准备度与企业规模显示方向性优势，但受样本量限制未达显著

第一轮未采纳面板企业中的转化预测（Logistic回归，n=93，两年转化率32%）



数据来源：AIHRM三轮全程追踪面板。† p<.10；准备度OR=1.88（p=.13）、规模OR=1.53（p=.08）。

9.5 解读与启示

面板数据的价值在于它排除了样本更替的干扰，让管理者看到同一批企业的真实变化。三个结论具有直接的参考意义。其一，转型大趋势是明显的：三个采纳状态之间的迁移呈现明确的净流入，无计划阵营从57家收缩至36家。其二，转型并不平滑：48次转入伴随35次退出，大约每三次进入就有两次退出，说明相当一部分企业的早期采纳带有试探性质，在实践检验后收缩了战线。这一细节对管理者是一种提醒，也是一种宽慰：撤回一个不成立的试点不是失败，而是扩散过程中的常态；真正的失败是在没有评估机制的情况下既不深入也不退出，让项目悬置消耗资源。

其三，同批企业的应用总量在第二轮到第三轮之间高位企稳，而非继续快速扩张。结合第五章观察到的准备度回调，一个合理的判断是，铺开阶段的红利已经收割得差不多，下一阶段的竞争将发生在消化与深化上。对已经完成初步部署的企业而言，继续增加应用数量的边际收益在下降，把已有应用做深、做到能替代原流程的程度，才是当前更有价值的投入方向。

转化预测的结果虽受样本量限制未达统计显著，其方向性仍值得参考：2024年资源准备度较高、规模较大的企业，两年后完成采纳的可能性更大，而当时的重视程度与技术契合度并未显示预测力。若这一方向在后续轮次中得到确认，其管理含义相当明确——决定企业能否从观望走向落地的，是两年前是否在悄悄打基础，而不是当时表态有多积极。对尚在观望的企业而言，这提示了一个具体的动作：即便暂不启动应用，数据治理、系统对接与人员能力的准备工作也应当先行，它们才是日后转化的真实前提。

10

第十章

结论、启示与未来图景

10.1 三年观察的基本结论

三轮追踪调查支持以下基本判断。中国企业的AIHRM采纳已越过早期尝鲜阶段，接近半数的采纳率、观望阵营的持续收缩与匹配面板中的真实转化共同表明，人工智能进入人力资源管理已成为不可逆的趋势；但成熟度分布同时显示，这仍是一场刚刚进入深水区的转型，规模化与深度集成的企业不足3%，应用的广度显著领先于深度。驱动因素方面，认知与动机层面的分歧已经很小，资源准备度成为区分企业的核心变量，人力资源负责人的专业背景在两个独立时间窗口中表现出对采纳增量的稳健影响，其作用很可能经由组织准备状态传导。后果方面，效率类模块的回报开始兑现，人力资本涌现效应持续增强，但绩效差距扩大、人员流动加快与人力资源职能自身的收缩重构同步发生，收益与代价并行。制度方面，治理建设显著滞后于应用扩张，风险感知已明显上升，过半企业却仍无任何治理安排。

10.2 对企业管理者的建议

第一，将投入重心从铺开转向做深。数据显示广度的扩张已近高位而深度普遍不足，回报又集中在少数模块，企业与其在九个模块上均匀用力，不如选择招聘、行政等已验证回报的场景，将单点应用推进至流程重构的深度，形成可复制的内部标杆后再行扩展。第二，把治理框架的建设提前到规模化之前。使用政策、数据边界、伦理审查与培训体系的建设成本，在试点阶段远低于规模化之后的补课成本；数据隐私已跃居挑战之首而过半企业尚无任何治理安排，这一缺口应在应用规模化之前补齐。第三，将负责人配置纳入转型方案。数据表明年轻且具备跨界背景的人力资源负责人所在企业采纳进度领先，仅有单一专业背景的负责人所在企业在两个时间窗口中均表现出更低的采纳增量；企业在选配人力资源负责人时应将数据素养与跨界经验作为重要考量，对在任团队则应提供系统的技术与数据能力训练。第四，主动管理转型的人员面。应用深化与人员流动、编制调整的同步出现提示，再培训与转岗安排不应是裁减之后的补救，而应作为部署方案的组成部分与技术投入同步规划；对工具使用带来的绩效分化，应通过培训普及与评价体系调整加以引导，避免其固化为组织内部新的不平等。

10.3 技术前沿：从辅助工具到数字劳动力

在给出终局判断之前，有必要先看清技术本身正走向何处。过去三年，人工智能在组织中的形态经历了三级跳：先是以单点工具的方式进入，完成筛选、匹配、计算等界定清晰的任务；随后以副驾的方式嵌入日常工作软件，与使用者实时协作生成内容、处理信息；当前的前沿则是智能体，即能够围绕目标自主串联多个环节、调用多种工具并动态决策的系统。本报告第三轮的数据恰好为这一阶梯提供了截面刻度：属于第一级的简历筛选已被62.4%的企业常态化使用，属于第二级的生成

式课程内容创作与嵌入办公软件的辅助工具分别达到28.9%与11.3%，而属于第三级的全流程自主招聘智能体与工作流编排智能体的渗透率均在8%上下。三级形态并存且渗透率依次递减，说明中国企业正处在从第一级向第二级大规模迁移、向第三级零星试探的位置。

国际前沿实践显示了阶梯的上端可能是什么样子。部分领先企业已将招聘的寻访、初筛、沟通与面试安排交由智能体端到端执行，人只保留终审与关键岗位的介入；一些企业开始为智能体建立类似员工的管理档案，将其纳入编制与权限体系，数字劳动力由此从比喻变为管理对象。在人才管理的底层逻辑上，以技能而非岗位为单元的管理系统与内部人才市场正在扩散，企业通过持续解析员工的技能构成来完成配置、发展与继任决策。组织架构层面也出现了标志性的重组，美国生物技术企业莫德纳将人力资源与信息技术部门合并为一个统一的人力与数字技术部门，并由原首席人力资源官出任负责人，其背后的逻辑正如业界评论所言：人力资源部门理解工作，技术部门理解工具，当工作与工具深度缠绕时，二者必须在同一张桌子上被设计。

技术能力与组织应用之间存在明显的时间差，这是判断未来节奏的关键变量。Anthropic在2026年初发布的报告指出，人工智能在多数职业任务上的能力已经具备，但实际应用显著滞后于能力边界。本报告的成熟度数据与此呼应：接近半数的企业已经采纳，真正实现规模化与深度集成的不足3%，能力与应用之间的落差构成了未来数年变化的主要势能。课题组另一项基于中国上市公司招聘记录的研究提供了同方向的证据：全部在招岗位的平均人工智能暴露度在2019年见顶后持续回落，这一回落并非影响减弱的信号，恰恰相反，它意味着高暴露度岗位的招聘需求正在收缩，替代效应已开始在劳动力市场的结构层面显形。

10.4 终局推演：AIHRM的可能形态

基于三轮数据与前沿实践的延长线，我们对AIHRM的中长期形态作出五点推演。需要说明的是，这些判断是证据的外推而非预言，其实现速度取决于技术、成本与制度的共同演化。

推演一，事务层将趋于完全的系统承接。本报告第四章显示，九个模块当前的头部实践几乎全部是筛选、监控与流程自动化类应用，这类任务边界清楚、规则稳定、错误可纠正，正是自动化最先穷尽的地带。沿此延长，考勤休假、薪酬计算、政策问答、档案管理、入职离职等事务性工作的绝大部分将在未来数年内由平台层完整承接，员工与人力资源体系的日常接口将从人变为系统，人的介入收缩至例外处理与申诉环节。事务层的归零不是人力资源职能的消亡，而是其成本结构与人员结构的根本改写。

推演二，招聘将成为第一个被完整重构的模块。简历筛选已经常态化，实时对话评估的面试应用以23.2%的渗透率快速爬升，全流程自主招聘智能体已在8%的企业出现。三级形态在招聘模块内的完整在场，使其最有可能率先走完从工具到智能体的全程：常规岗位的寻访、筛选、沟通、评估由智能体端到端编排，人的角色收缩为标准的制定者、例外的裁决者与关键岗位的终审者。届时招

聘的核心竞争力将从执行效率转向两项人的能力：为智能体设定正确标准的判断力，以及在高价值候选人身上建立信任的关系能力。

推演三，管理的对象将从人扩展为人机混合的团队。当智能体承担越来越完整的工作单元，组织事实上多出了一类新的劳动力，其配置、授权、评估与退出都需要规则。业界流传过一种说法，认为人力资源部门将不复存在，但其原始表述的方向恰恰相反：信息技术部门需要像人力资源部门管理员工那样管理智能体。这一表述指出了一个更可能的走向，即人力资源管理的方法论将外溢为数字劳动力治理的模板，岗位说明、权限边界、绩效标准、行为规范这些为人建立的制度工具，将被改写后用于管理机器。为人机协作设计分工、为智能体设定行为边界，可能成为人力资源职能在下一阶段最具增量价值的新任务。

推演四，人才管理的基本单元将从岗位转向技能与任务。课题组对标普500企业岗位数据的分析显示，大量岗位名称未变而技能内核已被置换，同一职位名称之下正在并存两种不同的人。当岗位标签不再可靠，以技能为单元的人才账本将成为配置与发展决策的基础，内部人才市场将取代静态的岗位序列成为人才流动的主渠道。与此同时，本报告观察到的绩效差距扩大意味着绩效分布本身正在被技术重塑，善用工具者获得放大的产出杠杆，人才识别与评价体系需要为此重写，以区分工具红利与真实能力。

推演五，通往终局的路径不会是直线。本报告的面板数据显示，约每三次采纳转入就伴随两次退出，试错与回摆是扩散过程的常态而非例外。国际经验同样提供了节制乐观的理由：瑞典金融科技企业克拉纳曾以人工智能替代数百名客服人员，其后因服务质量下滑重新招聘人类员工；对2025年美国上百万裁员的统计显示，企业自己明确归因于人工智能的比例不足百分之五，以技术为名的人员调整中混杂着大量其他动因。经济学的研究早已指出，技术对就业的净效应取决于企业在增强与替代之间的选择，而这一选择最终由制度与管理决定，不由技术本身决定。终局的方向是清楚的，但抵达的方式仍然掌握在组织自己手中。

10.5 AI时代的人力资源从业者：市场的选择与能力的重构

人工智能重塑人力资源管理之后，由谁来做人力资源工作，是本次调研过程中受访者提问最多的问题。对这一问题，与其依赖观点，不如让招聘市场自己投票。课题组建立了覆盖2014年至2025年、共计71万余条人力资源类岗位招聘广告的数据库，企业在职位描述中真实写下的要求，构成了对这一职业变迁最诚实的记录。分析显示，十二年间人力资源岗位的硬技能要求发生了三个方向的位移：数据化、绩效化与战略化，市场对从业者的期待正在从熟悉系统操作与政策条文，转向能够分析数据、理解业务战略。软技能方面，沟通与协作作为基础要求在高位企稳，真正的增量集中在承担责任与影响他人两类要求上。下行最陡的是标准化内容交付类技能，在培训与开发领域，约七成此类技能条目的市场需求在收缩，这是生成式人工智能替代最彻底的地带。

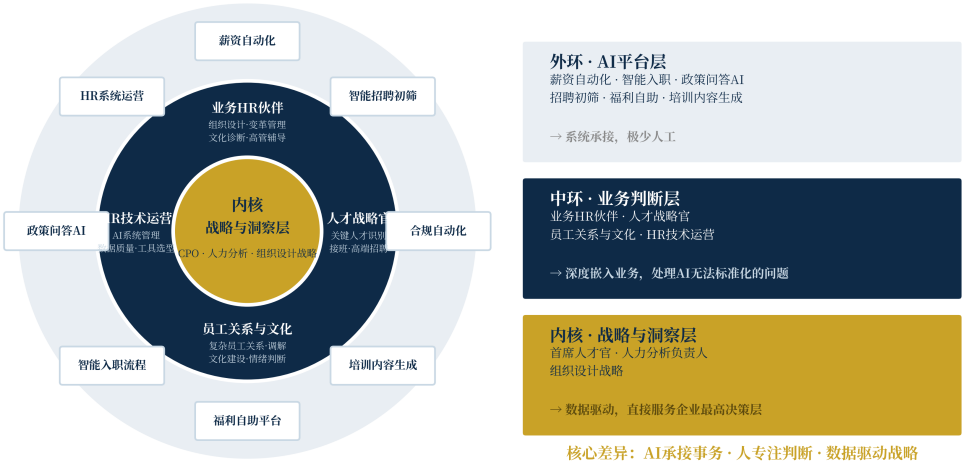
更值得注意的是职业内部的分化。将同一类人力资源职位按薪资分组比较后发现，高薪组的技能要求在全面加码，数据分析、战略理解与团队领导的权重同步上升；低薪组的要求则基本停留在原地。同一个职业名称，正在走向两种不同的命运，分化不发生在职业之间，而发生在职业内部。这一格局与本报告第八章观察到的员工绩效差距扩大在逻辑上同构：技术放大了能力结构的回报差异。综合市场十二年的投票，结论可以表述得很直接：被挤压的不是人力资源职业本身，而是其中以标准化流程执行为主的部分；懂业务、懂数据、懂人工智能的复合型从业者，其市场价值在持续上升。

落到组织形态上，还需要回答一个结构问题。过去二十年里，全球企业人力资源部门的组织设计几乎都以戴维·尤里奇（Dave Ulrich）提出的三支柱模型为蓝本，即专家中心、业务伙伴与共享服务中心的分工。人工智能的进入正在动摇这一架构的根基——共享服务中心所承载的标准化交付，恰恰是系统最先完整吸收的部分。也许是时候让三支柱演化为另一种形式了：课题组将未来人力资源部门的结构归纳为一个三环模型。外环是人工智能平台层，薪资自动化、智能入职、政策问答、招聘初筛、福利自助与培训内容生成在此由系统承接，仅保留极少人工；中环是业务判断层，由业务伙伴、人才战略、员工关系与文化、人力资源技术运营等角色构成，深度嵌入业务，处理人工智能无法标准化的问题；内核是战略与洞察层，由首席人才官、人力分析负责人与组织设计战略角色构成，以数据驱动，直接服务企业最高决策层。三环的分工可以概括为一句话：人工智能承接事务，人专注判断，数据驱动战略。需要说明的是，三环之间的边界并非一成不变。随着人工智能能力的持续演进，今天由中环人工处理的部分判断类工作，明天可能被系统吸收进外环；环与环的分界线会不断向中心收缩，人所专注的地带将随之向更高阶的判断与战略层面迁移。因此，三环模型刻画的不是一个静态的岗位分布，而是一个动态的收缩过程。三环之间的人员配比将因企业规模与行业而异，但价值重心由外向内收拢、机器边界由外向内推进的方向是一致的。

图 10-1

未来HR组织的三环结构：事务由系统承接，人的价值向判断与战略收拢

外环由AI平台承接事务；中环处理业务判断；内核以数据驱动战略



资料来源：课题组。该模型为基于现有证据的方向性推演，各环人员配比因企业规模与行业而异。

对应到个人，未来从业者的能力图谱呈现自下而上的层级结构。最底层是事务性人力资源管理能力，包括标准化的专业知识、流程执行与事务管理、标准化内容的制作与培训交付等；这一层的工作边界清晰、规则稳定，正在让渡给系统，市场需求持续下行。在其之上是战略性人力资源管理能力，包括接班计划与关键人才管理、需要公平性判断的薪酬与激励设计、高端招聘与候选人关系经营，以及组织设计与变革落地；这些属于传统人力资源管理功能划分中的高阶部分，依赖经验与专业深度，机器难以整体接管。在这两层传统能力之上，未来的从业者需要叠加一层新的能力，即人工智能与数字化能力，包括以数据讲清问题并对算法结论保持质疑的数据素养、判断人工智能输出质量与识别例外的工具把关能力、界定人机分工的流程设计能力，以及隐私边界与算法问责层面的伦理合规意识。而位于图谱顶端、也最为关键的，是两项核心能力：一是情境判断力与人际信任建立的能力，在复杂员工关系与高管信任场景中读懂情绪、作出判断、赢得信任，其立足点是以人为本、以员工为本；二是对业务的理解力，读懂损益表、将人才结构与经营战略连接、预判变革时机，其立足点是以公司为本、以绩效为导向。两项核心能力的共同特征是依赖情境、涉及信任、难以标准化，恰好落在机器最难复制的区间，也恰好分别回应了人力资源工作服务于人与服务于组织的双重使命。对从业者个人而言，可行的动作是对照这张图谱盘点自身的能力结构：工作内容越集中在底层事务性能力，转型的紧迫性越强；人工智能与数字化能力和顶端两项核心能力中的任何空白，都应作为未来两到三年的补课重点。

图 10-2

未来HR从业者的能力图谱：双核并立，事务性能力让渡给系统

基于课题组71万余条HR招聘广告（2014—2025）技能需求分析归纳



资料来源：课题组人力资源岗位招聘广告数据库。技能需求的量化位移详见正文。

莫德纳的案例为这一章提供了一个具体的注脚：当人力资源与数字技术合并为一个部门时，被选来掌舵的是原首席人力资源官而非技术负责人。这一人事安排背后的判断值得每一位从业者记取。当企业开始认真回答哪些工作交给人、哪些交给机器时，这个问题的圆心是组织、是工作、是人，而这正是人力资源专业积累最深的领域。技术从来不会自动向善，这场变革由效率逻辑单独定义，还是由既懂技术又对人保有温度的专业者引导，取决于谁坐在驾驶座上。市场十二年的投票已经表明，坐上驾驶座的门票是复合的能力结构；而本报告三轮数据反复出现的另一个事实是，多数企业的转型瓶颈恰恰是懂业务又懂技术的人力资源人才的稀缺。对从业者而言，这既是要求，也是机会。

10.6 调查的局限与后续计划

本报告的结论受若干条件约束。各轮样本存在构成差异，尽管关键结论均经匹配面板复核，总体统计仍应理解为参与调查企业的图景而非严格意义的全国代表性估计；主要变量来自人力资源负责人的自我报告，可能受共同方法偏差影响；横截面关联的因果解释保持谨慎，滞后检验虽提供了更强的证据，其样本量仍然有限。本章第三至五节的前瞻内容为基于现有证据的推演，不构成对具体时间表的预测。课题组计划于2027年实施第四轮调查，重点扩充三轮以上的追踪面板，引入客观经营数据的匹配，并对治理建设、智能体应用与员工层面使用行为进行更细粒度的测量，以持续记录中国企业这场组织转型的进程。

致谢

致谢

本研究于2024年8月正式启动。项目前期，在北京外企人力资源服务有限公司（FESCO）旗下人力资源协会（HRA）的协助下（特别感谢张淑艳女士的支持），课题组对近十家已经应用人工智能驱动人力资源管理的企业进行了深度访谈，这些访谈让我们得以在设计问卷之前先行理解真实的应用场景、困难与顾虑。在访谈与文献梳理的基础上，项目团队完成了问卷设计，并在中国HRD俱乐部（特别感谢蔡元启先生的支持）、人力资源协会（HRA）以及北京大学MBA人力资源协会的支持下，于2024年10月启动线上问卷调查，第一阶段调查于2024年12月底结束，形成了本报告的初步研究成果。其后两轮调查的样本招募同样主要依托上述渠道完成。没有这些机构在企业联络、渠道开放与持续动员上的投入，一项跨越三轮、历时近两年的追踪调查难以实施，课题组谨此致以诚挚谢意。

本项调查自启动之日起，得到北京大学光华管理学院光华思想力课题的全程资助。这项课题设立的初衷，是鼓励教员去做一些贴近企业实践的探索性工作，为未来的研究打下基础，而不必然以发表为单一目的。正是这一定位，给了课题组从容跟踪一个快速变化议题的底气：三轮调查、两年追踪，以及大量无法立即转化为论文的访谈与数据积累，都是在这样的支持下完成的。这项工作也带来了计划之外的收获，课题组在调研过程中孵化出一门面向MBA学生的课程《人工智能的商业场景应用》，将一线企业的真实实践带回了课堂。希望这份报告与这门课程，没有辜负光华思想力基金的资助初心。

感谢课题组的两位博士生同学。李航涛同学承担了前期线上问卷的设计与三轮数据的大量分析工作，许多细致的技术性劳动都凝结在本报告的数字背后；侯文昕同学则负责了三轮问卷的发放组织、参与企业的联络跟进以及纪念品寄送的登记等大量琐碎而具体的事务，正是这些环环相扣的执行工作保证了近两年追踪调查的顺利运转。值得一提的是，这些工作恰恰是人工智能技术暂时无法完全取代的——问卷背后的专业判断，以及与数百家企业保持联系所需要的耐心与责任心，仍然只能由人来完成。没有他们的努力，这项调研无法顺利进行。

我们尤其要感谢三轮调查中每一位拨冗作答的人力资源负责人。这份问卷并不轻松，题量大、涉及面广，许多题目需要回答者调取企业内部的具体信息才能作答，完成一次需要相当的时间与耐心。诸位在繁忙的工作之余认真填答，并且在第二轮、第三轮邀请发出时再度参与，这份坚持正是本报告最重要的基础。三轮全程参与的154家企业，使我们得以观察同一批组织在两年间的真实位移，这类证据在管理研究中并不多得。各位付出的时间，最终转化为可以被同行参照、被后来者使用的公共知识，这是对管理学术与管理实践共同的贡献。

在三轮调查陆续完成的过程中，阶段性研究结论曾在多个学界与业界场合分享，包括2024年浙江大学管理学院举办的数智组织研讨会、中国管理研究国际学会（IACMR）2025年年会暨《管理视野》杂志创刊十周年活动、FESCO旗下人力资源协会连续两年的超A大会、2026年HR智享会的年度论坛，以及课题组与瑞华智策合作开展的关于新质组织调研启动会。每一次分享收获的提问、质疑

与建议，都在不同程度上进入了后续轮次的问卷修订与本报告的分析框架，谨向提供这些交流机会的主办方和给予反馈的学界与业界同行致以谢意。

调查过程中，课题组还收到许多来信。不少来信者表达了对这一议题的高度关注，同时也坦陈了一种焦灼：明知人工智能正在改变人力资源管理，自己所在的企业却因数据基础、资源条件或组织认识的限制，一时还无法真正推进。我们愿意在此作一个回应：不必着急。三轮调查的数据显示，尚未采用的企业至今仍是多数，绝大多数已经采用的企业也还停留在探索与试点阶段，真正进入规模化应用的不足百分之三。当前的差距远没有行业叙事所渲染的那样悬殊。随着技术走向成熟、使用成本持续下降，工具的可得性将不再是少数企业的专利，普及是可以期待的。届时我们更希望看到的是，人力资源管理者能够从繁重的事务性工作中腾出手来，去做那些更有价值、更具战略意义的事情，更好地服务员工，更充分地体现人力资源工作应有的人文关怀。这既是技术进步应有的方向，也是本项调查持续开展的初衷。

最后作一点说明。本报告的写作得到了Claude Fable 5的大力支持，从数据整理、图表制作到文稿修订，人工智能深度参与了报告的生产过程，全部内容经课题组认真审核与反复调整后定稿。在此作出说明，既是课题组对AI工具使用保持透明的一份责任——我们相信，如实披露人工智能在知识生产中的参与方式，应当成为这个时代学术工作的基本规范；对一份以人工智能为研究对象的报告而言，这也是一次效率上的受益，是课题组积极拥抱人工智能技术的一次亲身实践与尝试。

本报告的观点与结论由课题组负责，文中若有疏漏，责任由课题组承担。

GUANGHUA SCHOOL OF MANAGEMENT · PEKING UNIVERSITY

A THREE-WAVE LONGITUDINAL SURVEY · 2024–2026

Artificial Intelligence-Enabled Human Resource Management

Evidence from Three Waves of a Longitudinal Survey of Chinese Firms

Guanghua School of Management, Peking University

Guanghua Thought Leadership Research Team on AI-Enabled HRM

JULY 2026

SUGGESTED CITATION

English citation (APA 7th)

> Jiang, K. (2026). *Artificial intelligence-enabled human resource management: A three-wave longitudinal survey of Chinese firms (2024–2026)*. Guanghua Thought Leadership Platform, Guanghua School of Management, Peking University.

Chinese edition

> 姜铠丰. (2026). 人工智能驱动的人力资源管理：中国企业三轮追踪调查报告（2024—2026）. 北京大学光华管理学院光华思想力课题组.

This report draws on three waves of a firm-level survey conducted between October 2024 and May 2026. Data and exhibits are the property of the research team; please cite the source when quoting. This English edition is a translation of the Chinese original; where wording differs, the Chinese edition governs.

Contents

A THREE-WAVE STUDY

— Executive Summary 4

1 Why Artificial Intelligence, and Why Chinese Firms 5

2 Method and Sample: Designing and Running a Three-Wave Study 8

3 The Overall Adoption Picture: From One Third to Nearly Half 13

4 The Application Landscape Across Nine HR Modules 20

5 Drivers of Adoption: Organisational Conditions and the HR Leader 28

6 Motives and Firm Profiles: Divergence in Readiness, Not in Motivation 37

7 Challenges and Governance: Application Has Outrun the Rules 41

8 Organisational Consequences: Returns, Gaps and Restructuring 45

9 Findings from the Three-Wave Panel 52

10 Conclusions, Implications and the Road Ahead 58

— Acknowledgements 67

Executive Summary

Between October 2024 and May 2026 our research team conducted three waves of a longitudinal survey on the use of artificial intelligence in human resource management (AIHRM) in Chinese firms, surveying the executives responsible for HR. The three waves yielded 2,241 firm observations; 487 firms participated in at least two waves and 154 firms took part in all three. Drawing on systematic cross-wave comparison and matched-panel analysis, this report traces the full trajectory of AIHRM in Chinese enterprises from first steps to broad diffusion.

Five findings stand out. First, adoption has risen steadily. The share of firms that have actually deployed AI in HR rose from 37.3% in November 2024 to 47.9% in May 2026; among the 288 firms that participated in both the first and third waves, the share rose from 38.5% to 51.0%, indicating that the increase reflects genuine within-firm conversion rather than shifts in sample composition. Second, application is broad but shallow. Recruitment, HR administration and work design form the leading tier, yet 86.1% of firms remain in the exploration or pilot stage and only 2.7% have reached scaled deployment or beyond. Third, what divides firms is organisational readiness rather than attitude. Latent profile analysis identifies four firm profiles that differ little in their stated motives but differ sharply in resource and staffing readiness; the age and professional background of the HR leader predict adoption consistently, and firms whose HR leader holds an HR-only background show significantly lower gains in adoption across two independent lagged windows. Fourth, governance lags well behind application. Concern about data privacy and security has risen sharply to become the single most cited challenge, yet 51% of firms have established no AI governance measure at all, and even the most widely adopted single measure covers less than three in ten firms. Fifth, returns are becoming visible but are unevenly distributed. Among firms able to assess returns, roughly six in ten perceive a moderate or better return on investment, and the human-capital emergence effect strengthens with each wave.

Building on these findings, the concluding chapter sets out recommendations for executives: shift investment from spreading across modules to deepening a few, build the governance framework before scaling rather than after, weigh data literacy and cross-functional experience when appointing HR leaders, and confront and regulate employee-initiated use of AI.

1

CHAPTER 1

Why Artificial Intelligence, and Why Chinese Firms

1.1 Human resource management is becoming a decisive site for AI deployment

Over the past three years generative artificial intelligence has moved rapidly from technical breakthrough to organisational application. The leap in large language models' capacity for text comprehension, content generation and multi-turn interaction has, for the first time, allowed AI to intervene directly in managerial work built on language and judgement — and human resource management is precisely where such work concentrates. From résumé screening, interview assessment and training content generation to employee advisory services, performance feedback and organisational diagnosis, almost every element of HR contains scope for reshaping by AI. A 2023 Gartner survey reported that 76% of HR leaders believed their organisations would fall behind competitively if they failed to adopt AI solutions within the following one to two years. Subsequent market developments bear this out: records from the venture database Crunchbase show that AI-driven HR technology start-ups rebounded markedly over the past two years after an earlier correction, and the parallel expansion of capital and product supply has created the practical conditions for firm-level adoption.

Set against this rapid maturing of technological supply, scholarly and practitioner understanding of the adoption process itself remains limited. Most existing studies rely on single cross-sectional surveys. They can tell us how many firms were using AI at one point in time, but they cannot answer three more important questions: how adoption decisions evolve over time, which organisational and individual factors drive or impede adoption, and what happens inside organisations after adoption. Answering these questions requires following the same firms across years, and that is the starting point of this survey programme.

1.2 Why Chinese firms warrant dedicated observation

The first reason for focusing on Chinese firms is the absence of an appropriate benchmark. When managers and researchers discuss levels of AI adoption — and particularly adoption within human resource management — the percentages they cite come mostly from global surveys published by international consultancies. Those surveys draw predominantly on North American and European respondents; Chinese firms typically appear in small numbers, with thin coverage across industries and firm sizes, and the Chinese sub-sample results are rarely reported separately. Inferring the progress of Chinese firms from such data may either overstate or understate the true picture, and its value as a reference is questionable. Chinese firms need benchmark data built on a large domestic sample, measured consistently, and comparable across years. That need is the most immediate motivation for this survey.

Beyond the benchmark problem, three further considerations argue for a China-specific study. First, China has the world's largest labour market and an enormous population of firms, so the effects of AI on employment structure and management practice carry their greatest aggregate weight here; the rapid diffusion of domestic foundation models around 2025 sharply lowered the barrier to use, giving organisational diffusion in China a rhythm and a set of pathways distinct from Western markets. Second, Chinese firms differ widely in their digital foundations: advanced manufacturers and small private firms are worlds apart in data infrastructure, systems and professional capability, and the introduction of AI is

layered onto this heterogeneity, so its diffusion cannot be extrapolated directly from developed-market experience. Third, the institutional environment is distinctive — the regulatory framework for data security and personal information protection, an ownership structure in which state-owned and private firms coexist, and active promotion of the AI industry by governments at every level together form an institutional setting that merits study in its own right.

On these grounds, and with the support of the Guanghua Thought Leadership programme, our team launched this longitudinal survey in autumn 2024, with the intention of observing AIHRM adoption in Chinese firms at a frequency of one to two waves per year. This report presents the complete results of the first three waves, spanning October 2024 to May 2026 — the critical period in which generative AI moved, among Chinese enterprises, from a subject of conceptual excitement to routine deployment.

1.3 Interpretation and implications

For senior executives and HR leaders, the practical message of this chapter concerns caution in citing data. When a meeting hears that “an international survey shows X% of firms have adopted AI,” the questions worth asking are how many Chinese firms were in that sample, how those firms were distributed across industries and sizes, and whether Chinese results were reported separately. Benchmarks drawn largely from Western markets can produce either unwarranted anxiety about falling behind or unwarranted complacency; both distort the allocation of resources. The three time series presented here are intended to give Chinese firms a domestic reference point of consistent construction and transparent provenance when setting their own AI agenda.

A second implication concerns urgency. HR has become a dense site of AI application because of the nature of the work: it runs on language, text and judgement, precisely the region where large language models are strongest. This means the HR function is unlikely to sit downstream of this wave of technology, waiting for other departments to demonstrate the way, as it often has before. The more realistic expectation is that HR will be among the earliest domains where AI is applied and among the earliest to be reshaped by it. That dual position makes the agenda one that cannot be outsourced and cannot be deferred.

2

CHAPTER 2

Method and Sample: Designing and Running a Three-Wave Study

2.1 Survey design and fieldwork

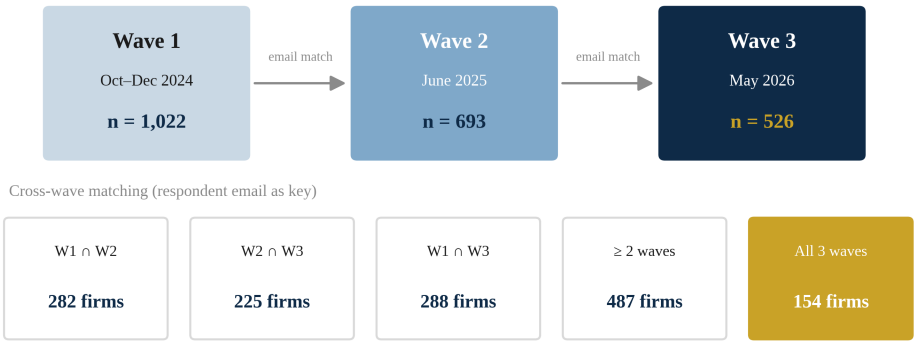
The study combines repeated measurement with fixed-sample tracking. Wave 1 ran from October to December 2024, with the bulk of responses collected in November, yielding 1,022 valid questionnaires from executives responsible for HR. Wave 2 was fielded in June 2025 and returned 693 valid questionnaires. Wave 3 was fielded in May 2026 and returned 526. In all three waves, respondents were the executives actually responsible for the HR function, spanning chief human resources officers, vice-presidents of HR, HR directors and HR managers. Each firm was represented by a single respondent, and the firm is the unit of analysis.

To enable cross-wave tracking, the survey used the respondent's email address as the matching key, establishing firm-level longitudinal links after standardising addresses during data cleaning. Matching identified 282 firms common to Waves 1 and 2, 225 common to Waves 2 and 3, and 288 common to Waves 1 and 3; 487 firms participated in at least two waves and 154 firms participated in all three. This structure allows the report to present two kinds of evidence in parallel: the overall picture based on each wave's full sample, and within-firm change based on the matched panel. Every key conclusion involving comparison over time has been checked against the matched panel in order to rule out shifts in sample composition. Exhibit 2-1 summarises the design and the flow of the sample.

Exhibit 2-1

Three survey waves yield 2,241 firm observations; 154 firms tracked across all three

Respondents are the executives responsible for HR (CHRO, VP HR, HR Director, HR Manager)



Source: Guanghua School of Management, Peking University — AIHRM Longitudinal Survey (2024–2026).

2.2 Measures

The three questionnaires kept their core measures consistent. Adoption status was measured with a single-choice item distinguishing five states: deployed through in-house development, deployed in partnership

with an external provider, considering in-house development, considering external partnership, and no current plans. Wave 3 subdivided the first two states into nine specific implementation routes in order to characterise the 2026 adoption ecosystem, and these were mapped back onto the five comparable categories for cross-wave comparison. Breadth of application was measured with multiple-response items across nine HR modules, each listing seven specific applications; Wave 3 added six options covering generative and agentic technologies, and cross-wave comparisons again count only the seven applications present in all three waves. Organisational-level explanatory variables include the strategic importance placed on AIHRM, attention to AI ethics and governance, technology fit, resource and staffing readiness, the high-performance work system, industry competitive intensity and the prevalence of AI within the industry, each formed as a multi-item composite on a five-point scale. Attributed motives distinguish augmentation logic, substitution logic and external pressure. The questionnaire further measured implementation challenges, organisational performance, workforce turnover, and an emergence scale capturing the conversion of individual knowledge into team capability through generative AI. Wave 3 added measures of maturity stage, governance measures in place, perceived return on investment, perceived performance dispersion, changes in HR team headcount and newly added roles.

Individual characteristics of the HR leader include gender, age, education and professional background, with background distinguishing those holding an HR-only background from those who also hold a cross-functional background in engineering, finance, law, marketing, strategy or operations; all were measured in every wave. Because age and professional background are relatively stable attributes, this report carries these data across waves within the matched sample, and every exhibit affected by that treatment is annotated accordingly. Wave 2 additionally measured eight competencies of the HR leader.

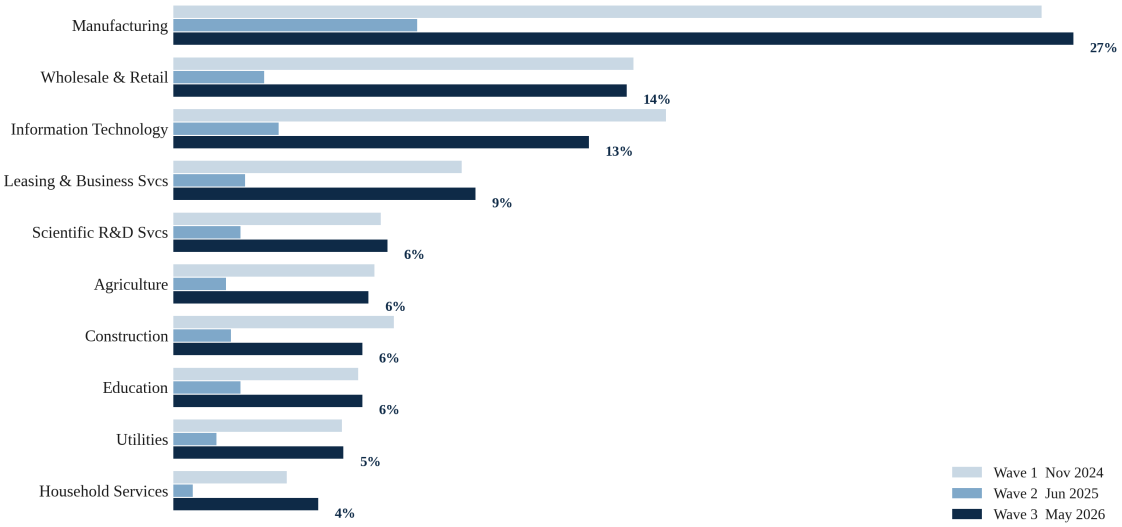
2.3 Sample composition

The three waves are highly consistent in industry, ownership and size structure, which provides the basis for cross-wave comparison. By industry, manufacturing and information technology form the two principal blocks, followed by wholesale and retail, finance and scientific and technical services; no marked structural drift appears across waves. By ownership, private firms account for roughly six in ten, with state-owned and state-controlled, shareholding, and foreign-invested and joint-venture firms making up the remainder. By size, the sample is dominated by small and medium-sized firms, with large firms of 2,000 or more employees holding steady at about one in five. Exhibits 2-2 and 2-3 present the distributions.

Exhibit 2-2

Industry composition is stable across waves; manufacturing and IT dominate

Share of firms in each industry (multiple response; top ten industries by Wave 3 share)

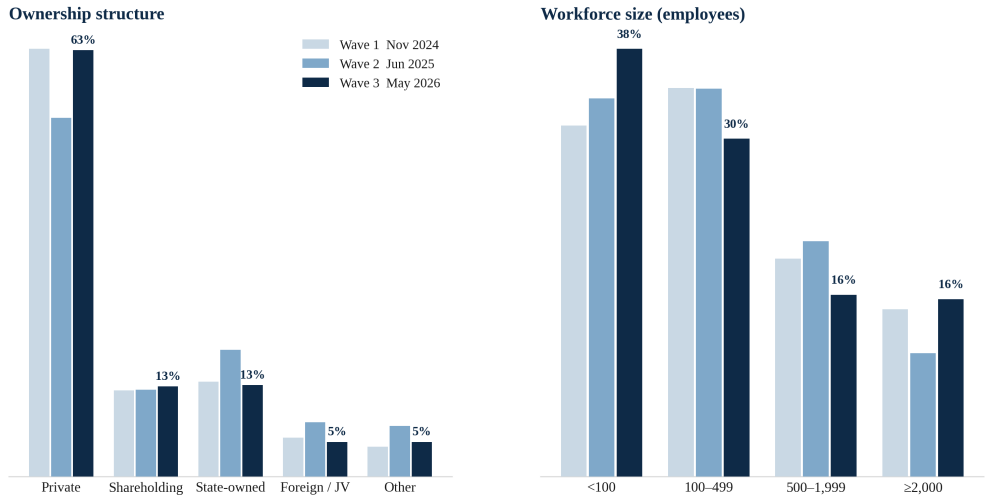


Source: AIHRM Longitudinal Survey. Industry is a multiple-response item, so shares sum to more than 100%.

Exhibit 2-3

Private firms make up about six in ten; small and mid-sized firms form the sample core

Share of valid responses in each wave



Source: AIHRM Longitudinal Survey.

2.4 Interpretation and implications

Holding the same measures, the same class of respondent and the same unit of analysis across three waves was intended to produce a time series that can be benchmarked against, rather than three snapshots that cannot be compared with one another. This point is often neglected in practice: many internal AI stocktakes change their wording and definitions each year, so that after three years of data collection a firm still cannot say how far it has moved. We would encourage firms to apply the same principle internally — fix a set of core indicators for at least three years and confine change to newly added items rather than existing ones.

The stability of the sample structure also carries a caveat about use. This sample is composed largely of private and small to medium-sized firms, with manufacturing and information technology well represented, which means the percentages reported here serve better as a reference for comparable firms than as national census values. HR leaders citing these data should look first at the subgroup closest to their own size and industry rather than at the overall mean; the subgroup breakdowns provided from Chapter 3 onward exist for that purpose.

3

CHAPTER 3

The Overall Adoption Picture: From One Third to Nearly Half

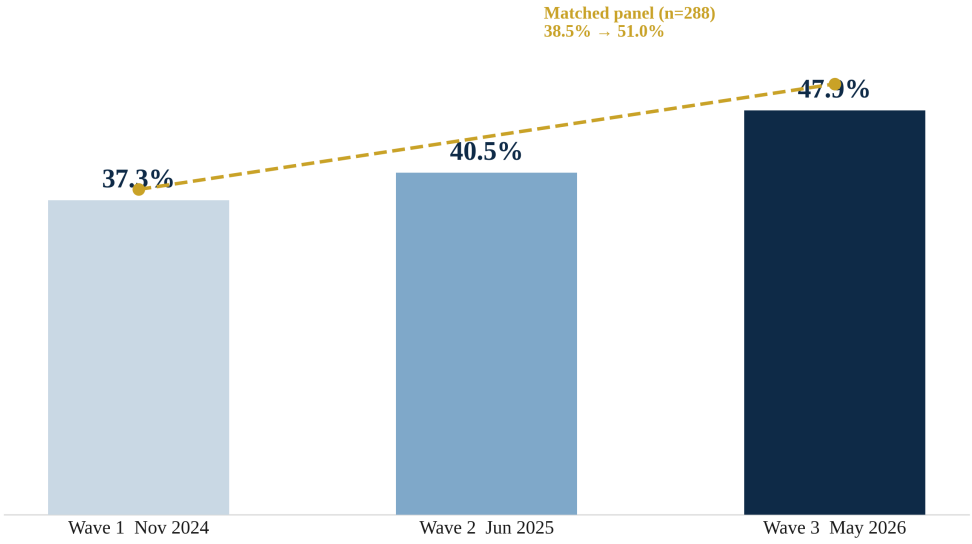
3.1 The three-year trajectory of adoption

The three waves record a clear upward curve. The share of firms that have actually deployed AI in human resource management rose from 37.3% in November 2024 to 40.5% in June 2025, and then to 47.9% in May 2026. To rule out the influence of changing sample composition, we compared within the 288 firms that participated in both the first and third waves: their adoption rate rose from 38.5% to 51.0%, a larger increase than in the full sample, indicating that the overall rise is driven mainly by genuine conversion among firms. At the current pace, adopters will constitute a majority within 2026.

Exhibit 3-1

AI adoption in HRM has climbed for three years, approaching half of all firms

Adopted = AI deployed in HRM, either developed in-house or with an external provider



Source: AIHRM Longitudinal Survey (W1 n=1,022; W2 n=693; W3 n=526). Gold line: within-firm change for the 288 firms in both W1 and W3.

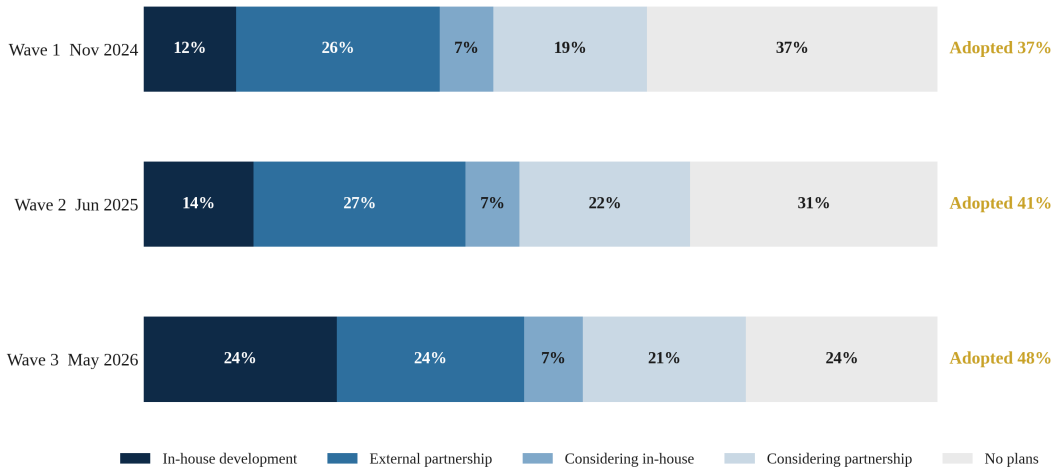
3.2 Structural change in modes of adoption

The composition of adoption modes reveals where the growth comes from. In the first wave in 2024, 36.6% of firms reported no plans; by the third wave that share had contracted to 24.1%. Over the same period the share deploying through in-house development rose from 11.6% to 24.3%, while partnership with external providers held steady at around a quarter. The continuing contraction of the wait-and-see camp and the expansion of the in-house camp occurred together, indicating that firms are not only adopting more but engaging more deeply with technical implementation itself. Firms still in a considering state account for roughly 28% in total, forming the reserve from which adoption will continue to rise over the next one to two years.

Exhibit 3-2

The “no plans” camp keeps shrinking; external partnership is the main route in

Composition of AI adoption modes (Wave 3's nine detailed routes mapped to five comparable categories)



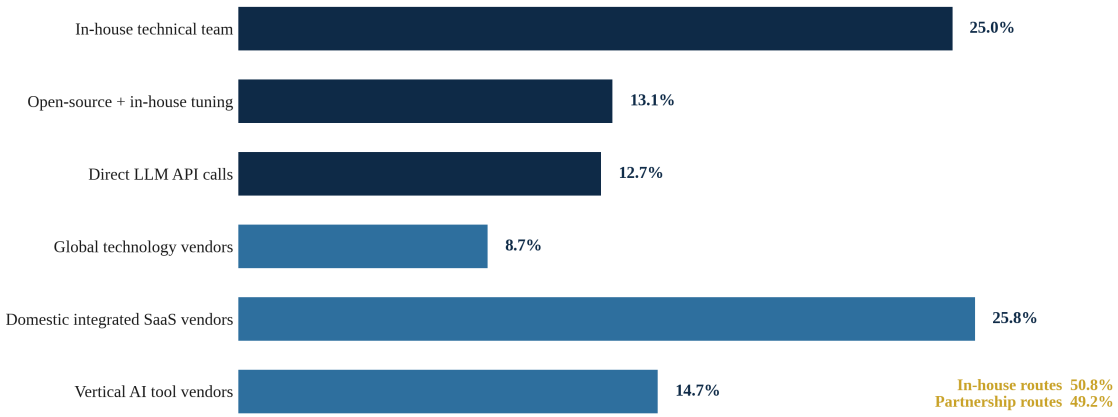
Source: AIHRM Longitudinal Survey.

The nine-way breakdown of implementation routes in Wave 3 provides a cross-section of the 2026 adoption ecosystem. Among the 252 adopting firms, in-house routes account for 50.8% in total, comprising development by an internal technical team, tuning of open-source solutions, and direct calls to large-model application programming interfaces. Partnership routes account for 49.2%, within which collaboration with domestic integrated software vendors is the most common, followed by vertical AI tool providers and then international vendors. The rise of direct API access as a low-threshold route is a notable change observed in this wave: it allows mid-sized firms without large technical teams to enter the in-house camp at relatively low cost.

Exhibit 3-3

A diversifying ecosystem: domestic SaaS and in-house build run neck and neck

Implementation route among adopters (Wave 3, n=252, share of adopting firms)



Source: AIHRM Survey Wave 3. Dark bars: in-house routes. Light bars: external partnership routes.

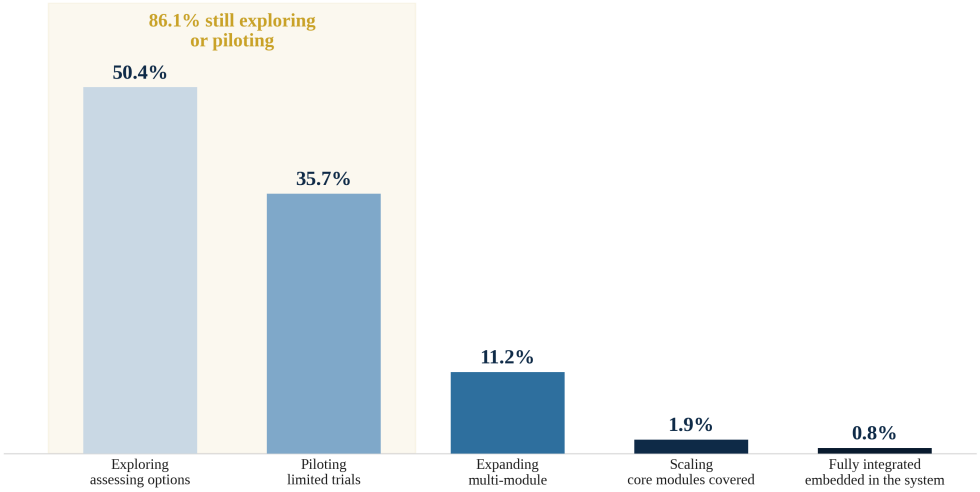
3.3 Maturity: broad but shallow

The quantitative rise in adoption has not yet produced a qualitative change in maturity. Wave 3 introduced a five-stage self-assessment of application maturity. The results show 50.4% of firms in the exploration stage of understanding and evaluating technical options and 35.7% in the stage of small-scale piloting, together 86.1%; 11.2% have reached the expansion stage of deployment across multiple modules, while firms at the scaled-application and deep-integration stages together account for only 2.7%. Placing this distribution alongside an adoption rate approaching half yields a more accurate reading of the current moment: most firms have crossed the threshold from nothing to something, but they remain a considerable distance from having AI genuinely embedded in the HR system. The module-level analysis of breadth and depth in Chapter 4 reinforces this judgement.

Exhibit 3-4

Broad but shallow: adoption nears half, yet only 2.7% have reached scale or beyond

Self-assessed maturity of AI application in HRM (Wave 3, n=526)



Source: AIHRM Survey Wave 3.

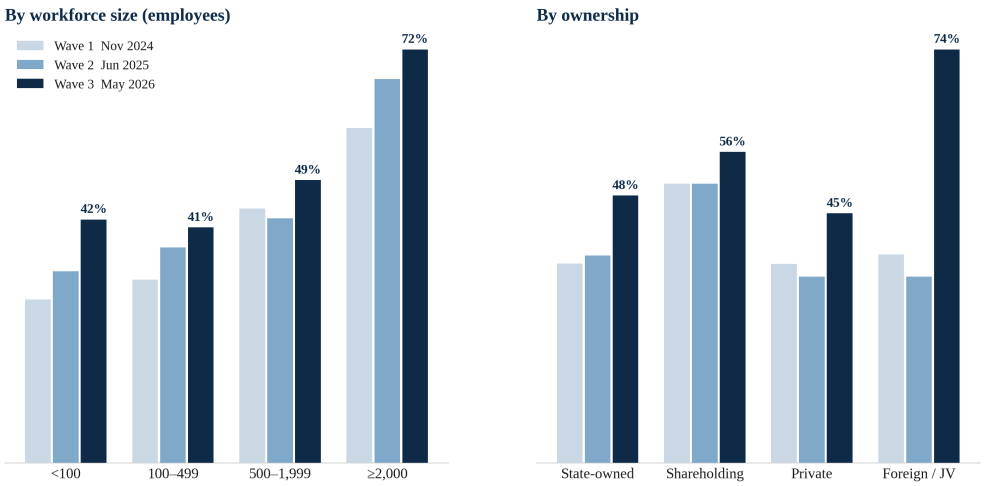
3.4 Divergence by size and ownership

Adoption shows a stable gradient by firm size. Firms with 2,000 or more employees led significantly in all three waves, and their lead has not narrowed; adoption among firms with fewer than 100 employees is also rising, but the gap with large firms remains clear. Differences by ownership are more moderate: private and shareholding firms adopt at slightly higher rates than state-owned firms, with foreign-invested and joint-venture firms in between, and all categories trending upward. The persistence of the size gradient suggests that resources and data foundations remain a hard constraint on adoption — a thread taken up more systematically in the analysis of drivers in Chapter 5.

Exhibit 3-5

Scale advantage persists: firms with 2,000+ employees lead and are still pulling ahead

AI adoption rate (share of firms that have deployed AI in HRM)



Source: AIHRM Longitudinal Survey.

3.5 Interpretation and implications

The first judgement conveyed by three years of data is that the window for waiting is closing. With adoption approaching half and the share of firms reporting no plans having contracted from more than a third to less than a quarter, AI in HR is moving from the choice of a few front-runners to standard equipment for the majority. For firms that have not yet started, the question is no longer whether to act but how to enter and how to avoid repeating mistakes others have already made. Notably, some 28% of firms remain in a considering state; this group forms the conversion reserve for the next one to two years, which also means that the current window of first-mover advantage may last no more than that.

The second judgement is that the adoption rate itself is easily misread. Setting the 47.9% adoption rate beside the 86.1% of firms still exploring or piloting gives a truer picture: most firms have entered the field, but few have completed a scaled rollout. Two things follow for executives. If your firm has not yet adopted, catching up is less daunting than the headline rate suggests, because the leaders are not far ahead in substance. If your firm has adopted, the real competition is not whether you use AI but whether you can be among the first to move from pilot to scale — the 2.7% are the genuine front-runners. In setting targets, management should use the maturity stage rather than the binary of adopted or not as the yardstick of progress.

The third judgement concerns the choice of route. Wave 3 shows in-house and partnership routes split roughly evenly, with direct API access rising quickly as a new path that lets mid-sized firms join the in-house camp without building a large technical team. This offers a practical option for resource-constrained firms: rather than choosing between building a platform and outsourcing wholesale, they can validate value in a single use case through API access first and then decide whether to invest more deeply. The persistence of the size gradient deserves equal attention. Adoption among firms with more than 2,000 employees has

reached 72%, and the gap of more than forty percentage points against firms under 100 employees stems mainly from data foundations and professional capability rather than differences in awareness. The realistic strategy for smaller firms is to close that foundational gap quickly through mature external services rather than attempting to replicate the in-house route of large enterprises.

4

CHAPTER 4

The Application Landscape Across Nine HR Modules

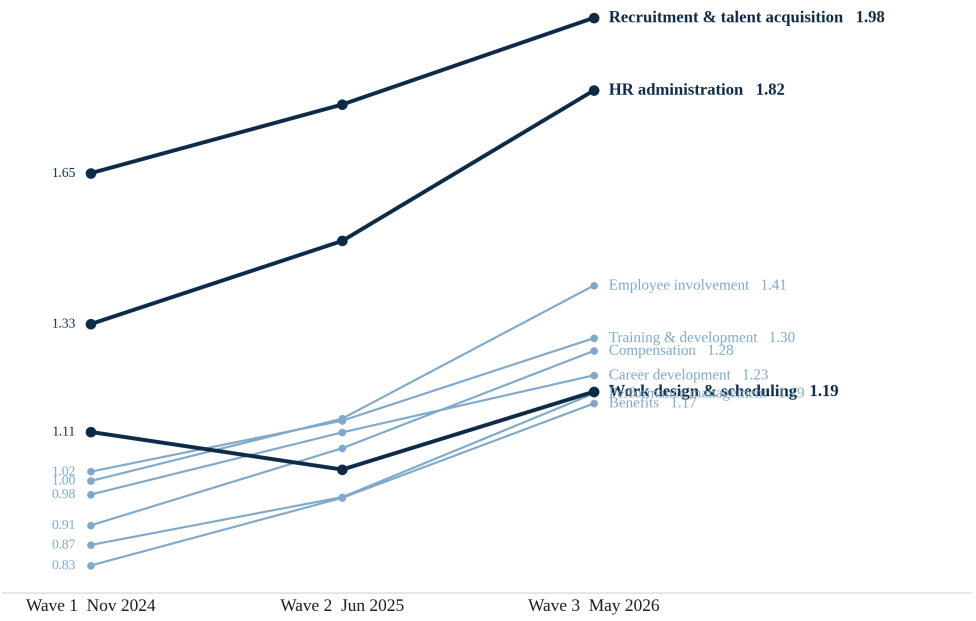
4.1 Module tiers and their three-year evolution

The average number of applications adopted across the nine HR modules rose in every module over the three years, while the tier structure remained stable. Recruitment and talent acquisition, HR administration and work design form the leading tier, followed by training and development and employee involvement, with performance management, compensation and benefits showing the lowest counts. The three modules in the leading tier share common features: highly standardised tasks, text-intensive processing and relatively contained decision risk, which makes them the natural entry point for AI into the organisation. Compensation and performance, by contrast, bear directly on employees' material interests and on compliance risk, and firms are visibly more cautious there.

Exhibit 4-1

All nine modules advance; recruitment, HR administration and work design lead

Average number of AI applications adopted per module (comparable 0–7 basis across waves)



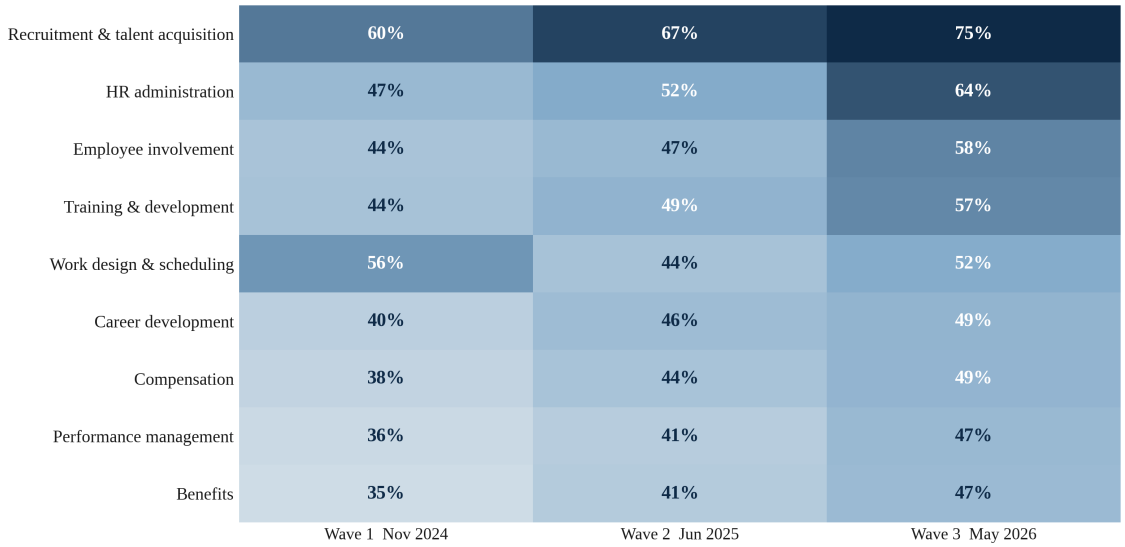
Source: AIHRM Longitudinal Survey. To keep waves comparable, only the seven applications measured in all three waves are counted.

Measured by the share of firms using at least one specific application, recruitment passed the halfway mark in Wave 3 and is the most penetrated module; administration and work design are also approaching half. Compensation and benefits have remained at the bottom of the penetration ranking in all three waves, forming a relative low ground on the application map. The heat map in Exhibit 4-2 presents the full structure of penetration by module and wave.

Exhibit 4-2

More than half of firms now use at least one AI application in recruitment

Share of firms using at least one specific AI application in the module



Source: AIHRM Longitudinal Survey (comparable basis).

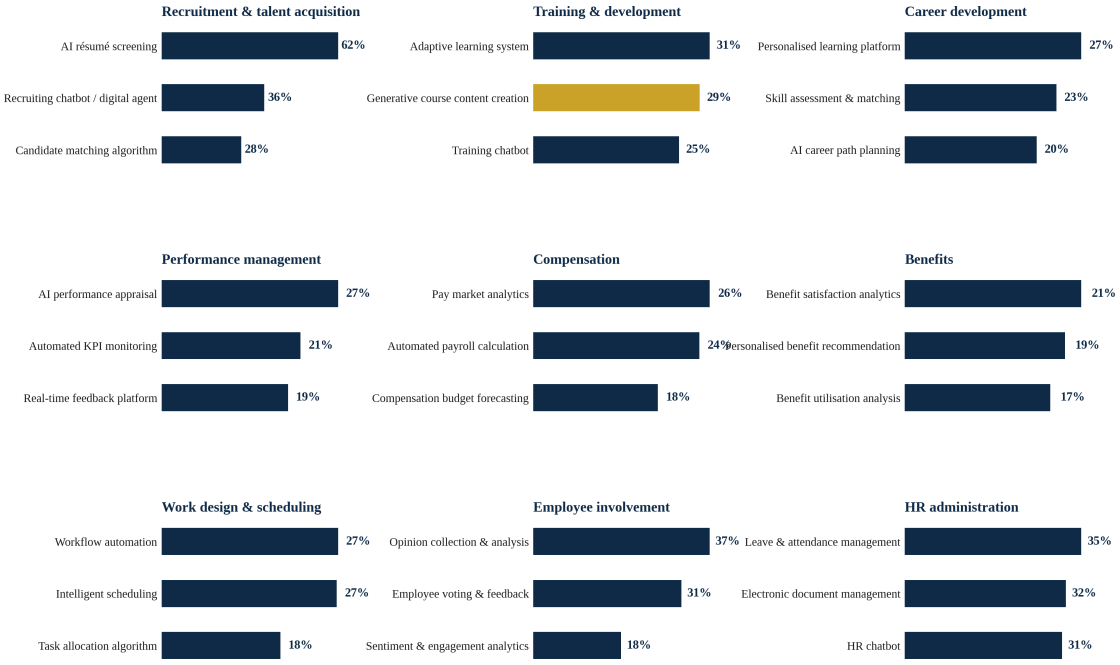
4.2 Leading practices and frontier applications by module

Wave 3 listed specific applications for each of the nine modules, and the three most-used practices in each module together sketch the real composition of AIHRM in Chinese firms in 2026. Recruitment is dominated overwhelmingly by résumé screening: 62.4% of firms use AI algorithms to screen applications automatically, the single most penetrated application on the entire list, followed by recruiting chatbots and digital agents (36.3%) and candidate matching algorithms (28.1%). Leading practices in HR administration cluster tightly around process automation: leave and attendance management (35.0%), electronic document management (31.9%) and round-the-clock chatbots (31.2%) run close together. In the employee involvement module, opinion collection and analysis platforms (36.7%) and employee voting and feedback systems (30.8%) are used more widely than expected, indicating that AI has entered the channels through which organisations listen to their employees. Work design is led by workflow automation (26.9%) and intelligent scheduling (26.7%). In training, adaptive learning systems rank first at 30.6%, while generative course content creation — newly measured in 2026 — moves straight into second place at 28.9%, the fastest-spreading of all newly added options.

Exhibit 4-3

Screening, monitoring and process automation head the list in every module

Three most-used AI applications per module (Wave 3, share of all 526 firms; gold = newly measured in 2026)



Source: AIHRM Survey Wave 3 (n=526, multiple response).

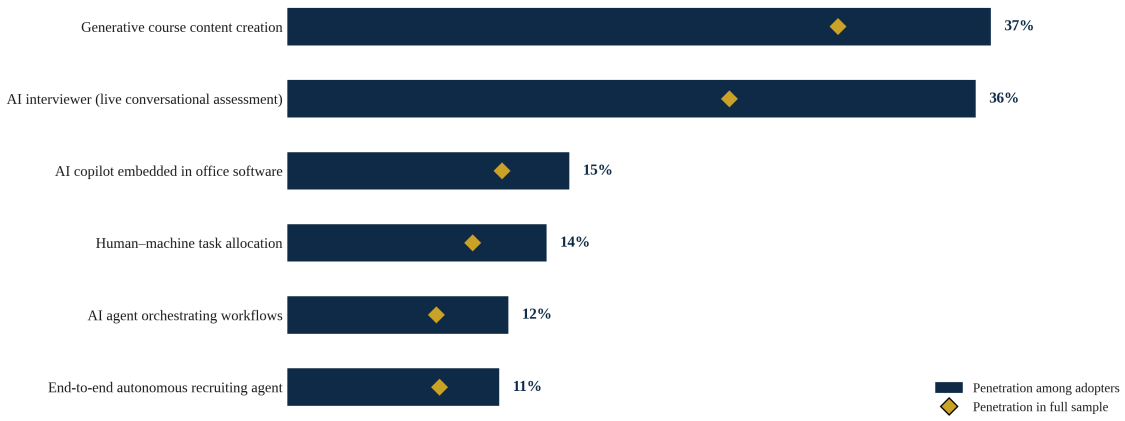
Viewed across all nine modules, the leading practices share a common feature: almost every top-ranked application is a matter of data analysis, screening and matching, or process automation — appraisal systems and KPI monitoring in performance, market analytics and automated calculation in compensation, personalised learning platforms and skill matching in career development, satisfaction analytics in benefits. In other words, the bulk of the current application map consists of having machines handle data and process, while applications in which AI generates content directly or makes decisions autonomously remain secondary. This pattern is consistent with the risk logic of technology adoption: firms place AI first where errors are correctable and accountability is traceable.

To track the penetration of the technological frontier, Wave 3 added measures of six frontier applications reflecting developments since 2025. Among adopting firms, generative course content creation and AI interviewers conducting live conversational assessment spread fastest, the latter reaching 23.2% in the full sample — close to the level of established recruitment practices. Copilot tools embedded in office software and human–machine task allocation occupy the middle range. Agentic applications in which AI links the entire recruitment process or orchestrates workflows autonomously remain early, with full-sample penetration around 8%. Overall, frontier applications of the generative and augmenting kind are diffusing considerably faster than those involving autonomous decision-making: firms are willing to let AI produce material and participate in assessment, but are not yet comfortable letting it decide alone. This corroborates the dominance of augmentation motives discussed in Chapter 6.

Exhibit 4-4

Generative and agentic applications emerge; content generation and AI interviewers move fastest

Six frontier AI applications newly measured in 2026 (Wave 3)



Source: AIHRM Survey Wave 3 (n=526; adopters n=252).

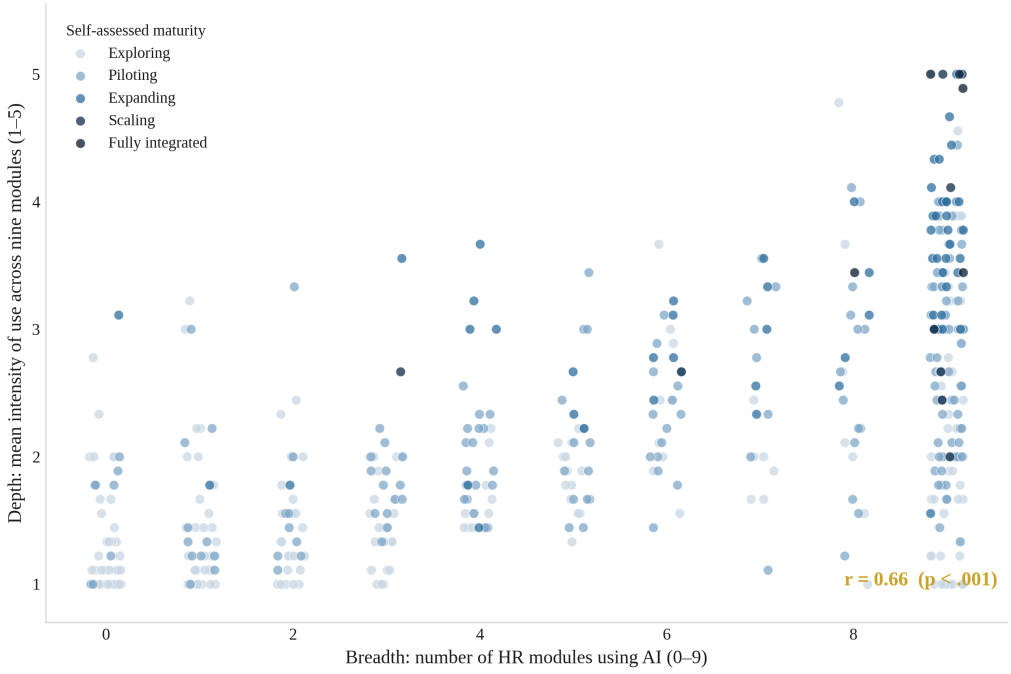
4.3 The relationship between breadth and depth

To characterise application at firm level, we measure breadth by the number of modules using AI and depth by the mean intensity rating across the nine modules. The two are significantly positively correlated, but the joint distribution shows most firms clustered in a region of relatively high breadth and middling-to-low depth: spreading across modules has run ahead of deepening within any one of them. Shading by self-assessed maturity shows that firms which have genuinely reached the expansion stage or beyond combine high breadth with high depth, whereas the large number of firms still piloting have touched several modules while remaining at a preliminary level within each.

Exhibit 4-5

Breadth outruns depth: firms spread across modules but stay shallow within them

Joint distribution of breadth and depth at firm level (Wave 3, n=526)



Source: AIHRM Survey Wave 3.

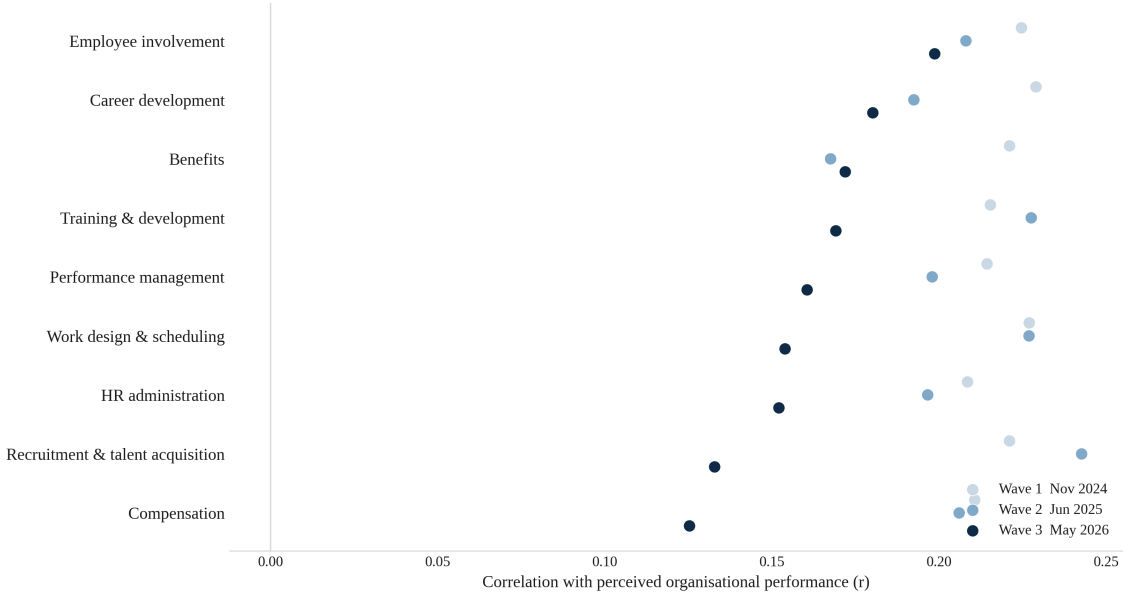
4.4 Module-level application and perceived performance

Correlating the number of applications in each module with firms' subjective performance ratings produces consistently positive associations across all three waves, with recruitment and training showing the most stable relationships, statistically significant in every wave. These results should not be read causally — better-performing firms may simply have more capacity to invest in AI. Still, the differences in strength across modules are informative: they echo the finding in Chapter 8 that recruitment and administration rank highest in perceived return on investment, suggesting that at this stage the value of AI is being realised chiefly in process-efficiency modules.

Exhibit 4-6

AI use correlates positively with perceived performance; recruitment and training most consistently

Correlation between module-level AI use and perceived organisational performance (large dots: p<.05)



Source: AIHRM Longitudinal Survey. Correlation does not imply causation; performance is self-rated relative to competitors.

4.5 Interpretation and implications

Module-level data yield a practical rule for sequencing investment: the value of AI is realised first where tasks are standardised, text-intensive and low in decision risk. It is no accident that recruitment, HR administration and work design lead simultaneously in penetration and in perceived return; what they share is clear process, high volume and correctable error. Compensation and performance, by contrast, bear directly on the distribution of rewards and on compliance risk, and firms' caution there is not conservatism but rational risk pricing. Executives can sequence their roadmaps accordingly: build capability and trust in the first category of modules and extend gradually into the second, rather than trying to prove the technology's value in the most sensitive areas first.

The distribution of breadth and depth exposes a common misallocation. Most firms sit in the region of high breadth and insufficient depth, and this usually arises from an understandable pressure: management wants to see activity in every module, so resources are spread across nine directions and each remains at the level of demonstration and trial. Firms that have genuinely reached the expansion stage show the opposite shape — breadth and depth together in a small number of modules. The concrete recommendation for HR leaders is to re-examine the current project list: if there are live projects in all nine modules and not one module in which the process has actually been rebuilt, the sensible move is to narrow the front, concentrate resources on one or two use cases with demonstrated return, and take those far enough to replace the original process before talking about replication.

The penetration pattern of frontier applications offers a third signal. Copilot-style tools embedded in office software are spreading fastest, while agentic applications that orchestrate processes autonomously remain early. This ordering follows a general regularity of technology adoption: tools that augment people are accepted more readily than tools that replace them, because the former leave accountability unchanged. Executives introducing agentic applications should expect a longer acceptance cycle and more extensive process redesign, and should treat the definition of accountability boundaries as part of the solution design rather than as a note appended after launch.

5

CHAPTER 5

Drivers of Adoption: Organisational Conditions and the HR Leader

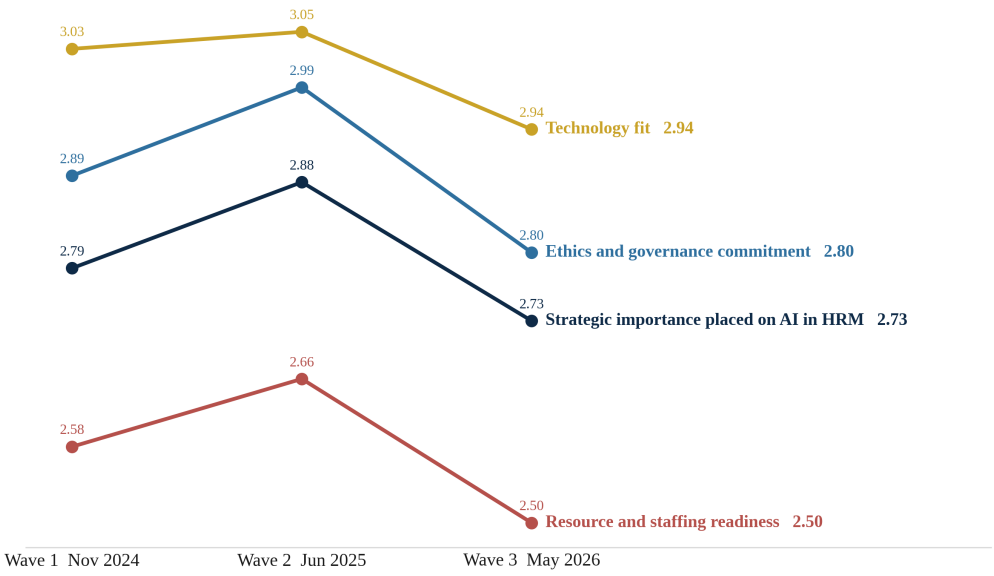
5.1 Three-year curves in organisational mindset and readiness

The four core indicators in the technology–organisation–environment framework trace a curve that rises and then falls. The strategic importance placed on AIHRM rose from 2.79 in Wave 1 to 2.88 in Wave 2 before falling back to 2.73 in Wave 3; attention to ethics and governance, technology fit and resource readiness follow the same shape, with resource readiness consistently the lowest of the four and standing at 2.50 in Wave 3. This collective pullback occurred against a background of continuously rising adoption, and the divergence deserves careful reading: as actual deployment proceeded, firms' assessments of their own preparedness shifted from the optimism of the conceptual phase to the caution of the implementation phase, and the field as a whole entered a more sober stage. Changes in Wave 3 sample composition may account for part of the decline, but checks against the matched panel show the same direction.

Exhibit 5-1

Enthusiasm rose then cooled: all readiness measures pulled back in Wave 3

Mean scores on the technology–organisation–environment dimensions (1–5 scale)



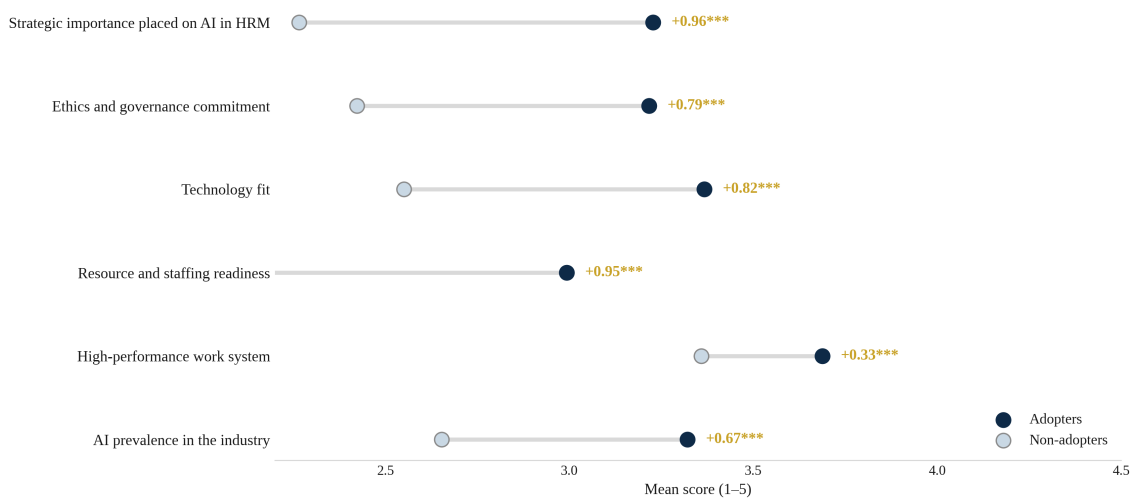
Source: AIHRM Longitudinal Survey.

Comparing adopters with non-adopters locates the crux of the divergence more precisely. Wave 3 data show significant differences on all six organisational indicators, with the two largest gaps in the strategic importance placed on AIHRM and in resource and staffing readiness; differences in technology fit and ethical attention are comparatively small. What separates adopters from observers, in other words, is less a matter of awareness and attitude than of actual resource commitment and accumulated staff capability. This finding corroborates the latent profile analysis in Chapter 6 and forms an important basis for the recommendations in this report.

Exhibit 5-2

The widest gap between adopters and non-adopters is resource readiness, not attitude

Differences in mean scores, Wave 3 (* p<.05, ** p<.01, *** p<.001)



Source: AIHRM Survey Wave 3.

5.2 High-performance work systems and adoption: a relationship that depends on the frame

The relationship between the high-performance work system and adoption cannot be summarised as simply positive or negative, because the two statistical framings point in opposite directions. In terms of direct correlation, all three waves are positive: firms scoring higher on the high-performance work system also adopt at somewhat higher rates, a fact reflected in the mean advantage of adopters on that dimension in Exhibit 5-2. This is not hard to understand, since firms with mature systems tend to be larger and better resourced, and size and resources are themselves positively associated with adoption. However, the team's regression analyses in the first two waves showed that once firm size, strategic importance and readiness are controlled, the coefficient on the high-performance work system turns significantly negative. The divergence between the two framings means the negative relationship is a conditional effect: among otherwise comparable firms, the more mature the HR system, the lower the marginal propensity to adopt. One explanation is path dependence — established processes and vested capabilities within a mature system reduce the relative return to change. In the integrated Wave 3 model this conditional effect remains negative in direction but is no longer significant (see Exhibit 5-8), suggesting that as tools become more available and peer demonstration accumulates, the dampening influence of a mature system is weakening.

5.3 Characteristics of the HR leader: age and professional background

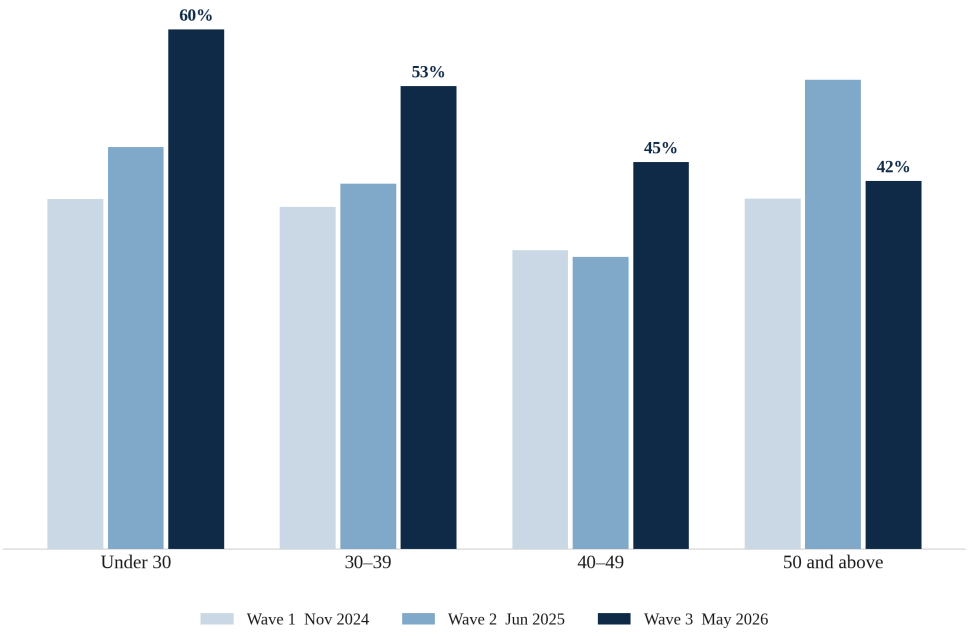
There is a stable association between the individual characteristics of the HR leader and the firm's adoption rate. Grouped by age, the older the HR leader, the lower the adoption rate of the firm, and this gradient runs in the same direction in all three waves. Grouped by professional background, firms whose HR leader also

holds a cross-functional background in engineering, finance or strategy consistently adopt at higher rates than firms whose HR leader holds an HR-only background, with the gap holding at roughly ten percentage points.

Exhibit 5-4

A robust age gradient: the older the HR leader, the lower the firm's adoption rate

AI adoption rate by age group of the HR leader

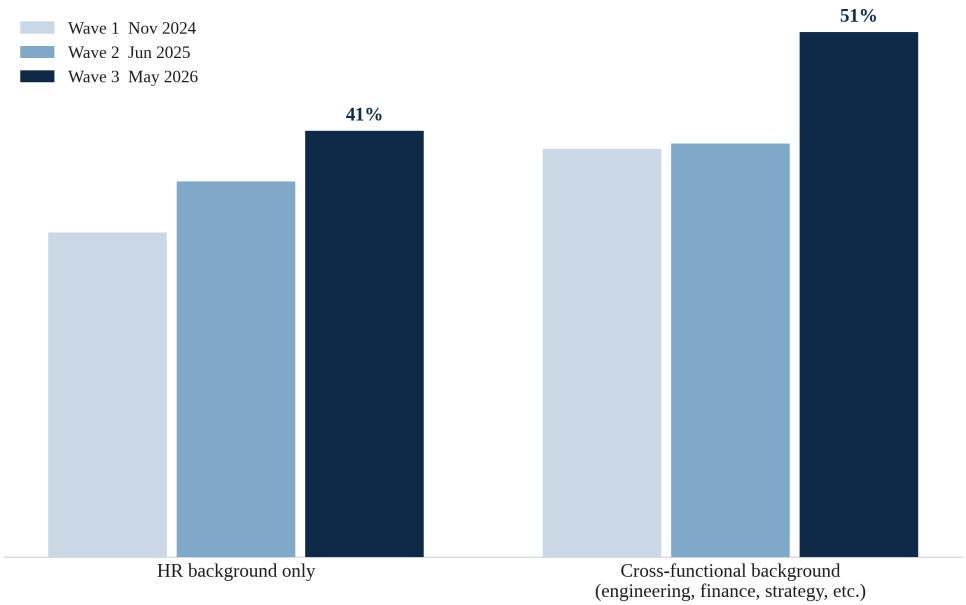


Source: AIHRM Longitudinal Survey. Wave 2 collected demographics only from new participants; returning respondents' data carried across matched waves.

Exhibit 5-5

Firms led by cross-functional HR leaders stay roughly ten points ahead on adoption

AI adoption rate by professional background of the HR leader



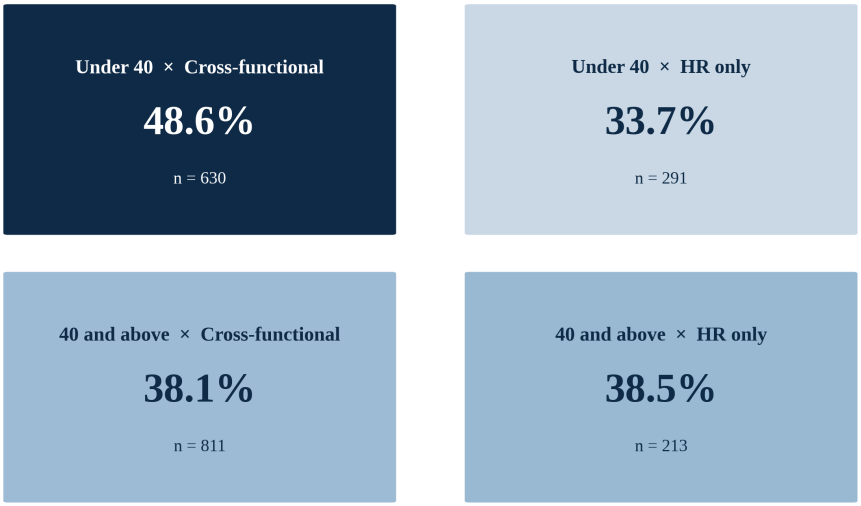
Source: AIHRM Longitudinal Survey. Cross-functional = holds any background outside HR (engineering, finance, legal, marketing, strategy or operations).

Crossing age with background produces a more textured picture. Comparing four groups in the pooled three-wave sample, firms whose HR leader is under forty and holds a cross-functional background adopt at 48.6%, the highest of the four; firms whose leader is under forty with an HR-only background adopt at 33.7%, the lowest of the four; the two groups aged forty and above fall in between and close to one another. Logistic regression confirms that the main effects of age and background and their interaction are all statistically significant. The interaction means that the adoption advantage conferred by a cross-functional background is concentrated among younger leaders and largely disappears among older ones. One reading is that leaders who are both young and cross-functional combine appetite for change with cross-domain knowledge, making them the most active force behind the AI agenda inside the organisation, while the limitations of a single-track professional path show up most clearly among the younger managers for whom cross-domain judgement matters most.

Exhibit 5-6

Young, cross-functional HR leaders drive adoption hardest — a 15-point spread

AI adoption rate by HR leader age × professional background (three waves pooled)



Source: AIHRM Longitudinal Survey, pooled sample. Logistic regression: age, background and their interaction are all significant (p<.01).

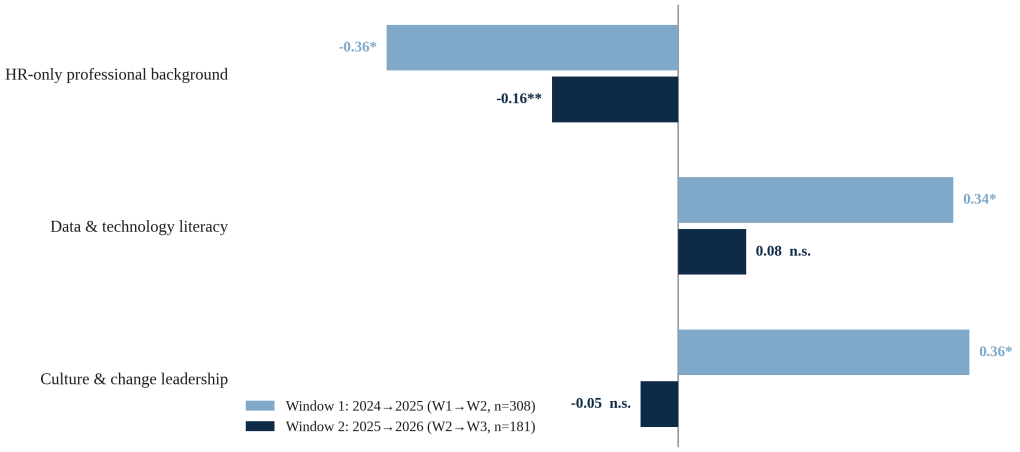
5.4 Lagged tests: replication across two time windows

Cross-sectional associations cannot rule out reverse causation or omitted variables, so we exploit the longitudinal design to run lagged tests. The first window uses firms matched across Waves 1 and 2, examining how leader characteristics predict Wave 2 adoption after controlling for Wave 1 adoption. The team's earlier analysis found a coefficient of −0.36 for an HR-only background, 0.34 for the leader's data and technology literacy, and 0.36 for culture and change leadership, all significant. This report constructs a second, independent window from firms matched across Waves 2 and 3. Controlling for Wave 2 adoption, firm size and the high-performance work system, the standardised coefficient of an HR-only background on Wave 3 adoption is −0.16, significant at the .001 level and fully consistent with the first window. The coefficient on data and technology literacy is positive in direction but not significant, and culture and change leadership is close to zero. Taken together, the dampening effect of professional background is the most robust of the leader characteristics, replicating across two years and two different sets of firms; the predictive power of the competency dimensions awaits confirmation from a longer observation window and a larger matched sample.

Exhibit 5-7

The HR-only penalty replicates across two independent time windows

Standardised coefficients predicting next-wave adoption, controlling for prior adoption



Source: AIHRM matched panel. Window 1 from the team's earlier analysis; Window 2 new to this report, controlling for firm size and HPWS. † p<.10, * p<.05, ** p<.01, *** p<.001.

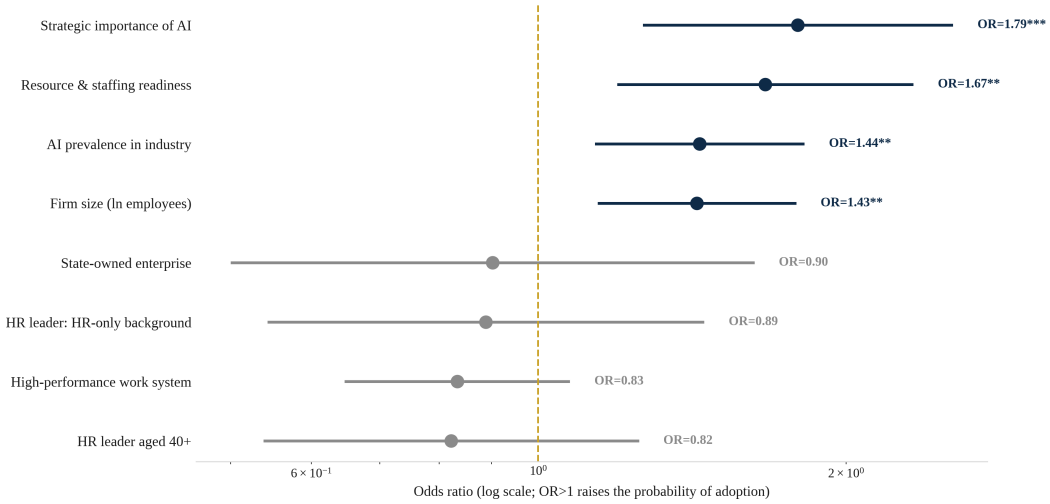
5.5 The integrated model: direct effects and transmission paths

Entering organisational factors and leader characteristics together in a Wave 3 logistic regression shows their relative strength. The strategic importance placed on AIHRM and resource readiness carry the largest odds ratios: a one standard deviation increase raises the odds of adoption by roughly 79% and 67% respectively. Firm size and the prevalence of AI within the industry follow, both significant. The age and professional background of the HR leader are no longer significant once organisational variables enter. Combined with the robust background effect in the lagged tests above, an interpretation consistent with the data is that leader characteristics act on adoption largely through organisational states: a young, cross-functional leader raises the organisation's strategic commitment to AI and its resource investment, and thereby drives adoption indirectly, with the organisational variables in the cross-sectional model absorbing that indirect path. Rigorous testing of this mechanism requires a longer panel, but the present evidence already supports one managerial judgement: changing the organisation's state of readiness is the practical lever through which an HR leader influences the course of adoption.

Exhibit 5-8

Strategic importance, readiness, size and industry conditions drive adoption directly

Logistic regression of AI adoption (Wave 3, n=515; continuous predictors standardised)



Source: AIHRM Survey Wave 3. Grey = not significant. HR-leader characteristics lose significance once organisational variables enter, suggesting their effect runs through organisational readiness; see Exhibit 5-7 for the lagged direct effect.

5.6 Interpretation and implications

The central finding of this chapter is that the gap between adopters and observers lies mainly in resource readiness rather than in awareness or attitude. This means HR leaders seeking support from senior management need to shift the emphasis of their argument. In most firms, executives do not need to be persuaded again that AI matters; disagreement about its importance is far smaller than disagreement about resource commitment. The more effective proposal is not another recital of the trend but a clear statement of the specific shortfalls in readiness — the availability of data, the degree to which systems can be connected, the technical understanding within the team, and the dedicated headcount available — with the cost and time required to close each one. With mean readiness at just 2.50 in Wave 3, the lowest of the four indicators, this is not an issue for particular firms but a shared weakness across the field: whoever closes the gap first gains a substantive head start.

The weakening of the conditional negative effect of the high-performance work system in Wave 3 matters particularly for firms with mature systems. Regressions in the first two waves suggested that among firms comparable in size and resources, the more mature the HR system, the lower the marginal propensity to advance AI, with path dependence a reasonable explanation; by Wave 3 that effect is no longer significant. For such firms this is a signal that they can re-enter: the very standardisation of their existing processes is turning from an obstacle into an asset, because standardised processes and accumulated data are precisely the preconditions for applying AI. Their earlier hesitation had its logic; continued hesitation has less.

Findings about the individual characteristics of HR leaders need to be used carefully and honestly. The data show that firms led by younger, cross-functional HR leaders adopt at significantly higher rates, and that firms whose leader holds an HR-only background show lower gains in adoption across two independent

time windows. The correct use of this conclusion is not to turn it into an age label in staffing decisions: age is not itself a capability, and background differences largely disappear within the group aged forty and above. The more reasonable reading is that these individual characteristics mark the availability of cross-domain knowledge and of experience in driving change — both of which can be acquired through appointment and development. The recommendation to firms is therefore twofold. In appointment, weigh data literacy and cross-functional experience when selecting HR leaders. In development, provide incumbent leaders with systematic training in technical understanding and data analysis, and bring members with technical backgrounds into the HR team so that the knowledge structure of the team compensates for the limits of any individual's background. For HR leaders themselves the data carry a direct message: technical and data capability is no longer a bonus qualification but part of the basic equipment of the role.

That leader characteristics lose significance once organisational variables are controlled should not be read as evidence that the leader does not matter, but as clarification of the path through which the leader matters. Leaders do not determine adoption directly through personal preference; they shape it by raising the organisation's strategic commitment to AI, securing resources and pushing preparatory work forward. That is precisely an actionable path: changing the organisation's state of readiness is the most realistic lever an HR leader holds.

6

CHAPTER 6

Motives and Firm Profiles: Divergence in Readiness, Not in Motivation

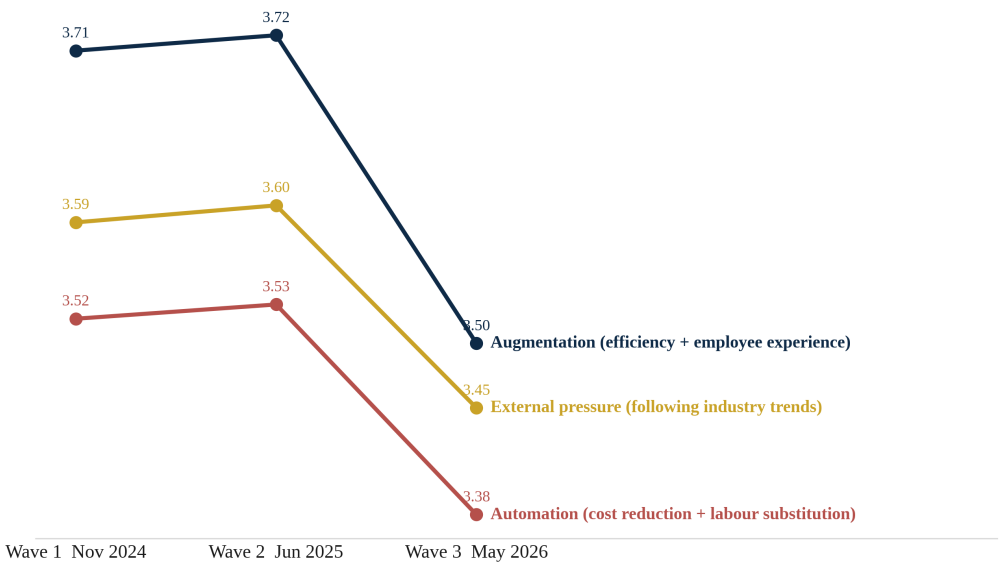
6.1 Three motives cool in parallel

Firms' motives for adopting AI fall into three categories: an augmentation logic centred on raising efficiency and improving the employee experience, a substitution logic centred on replacing labour and compressing cost, and external pressure to follow industry trends. Across the three waves the means of all three fall together while their ordering holds, with augmentation consistently first. This general cooling of attributed motives is part of the same phenomenon as the pullback in readiness observed in Chapter 5: the high expectations of the conceptual phase are giving way to the sober assessment of the implementation phase, and firms' answers to why they use AI have become more restrained and more practical. Substitution motives show no relative rise, indicating that — at least at the level of what HR leaders report — cost substitution has not become the dominant motive. The tension between this and the contraction in headcount observed in Chapter 8 is taken up in the concluding chapter.

Exhibit 6-1

All three motives cool in parallel; augmentation logic remains dominant

Mean attributed motives for adopting AI in HRM (1–5 scale)



Source: AIHRM Longitudinal Survey.

6.2 Identifying four firm profiles

To move beyond single-variable description, we conducted a latent profile analysis on the 2,238 pooled firm observations, using as inputs the four state indicators — strategic importance, ethical attention, technology fit and resource readiness — together with the three motive indicators, estimated with a Gaussian mixture model. Balancing model fit against interpretability, a four-profile solution was adopted. The first profile scores high on every state indicator, accounts for 26% of the sample and adopts at 70%; we

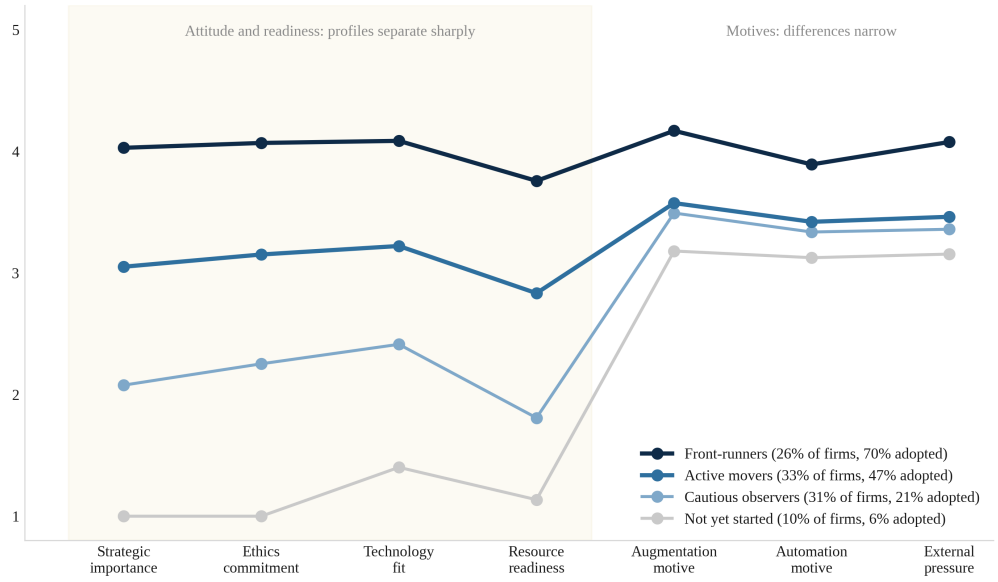
term these the front-runners. The second scores moderately high across indicators, accounts for 33% and adopts at 47%: the active movers. The third scores moderately low, accounts for 31% and adopts at 21%: the cautious observers. The fourth sits near the bottom of the scale on state indicators, accounts for 10% and adopts at just 6%: those not yet started.

The discriminating structure of the profiles reveals a crucial fact: differences among the four groups on the motive dimensions are far smaller than their differences on the state dimensions. Cautious observers endorse the augmentation logic as well, and their motive scores are not far from those of the front-runners; what genuinely separates them is resource readiness and strategic commitment. This means Chinese firms have converged close to consensus on whether AI should be embraced, and that divergence occurs in whether the corresponding resources and capabilities can be assembled. The bottleneck in adoption is a problem of execution, not of awareness.

Exhibit 6-2

Four latent profiles: firms diverge on readiness, not on motivation

Profile means on seven dimensions (Gaussian mixture model, three waves pooled, n=2,238)

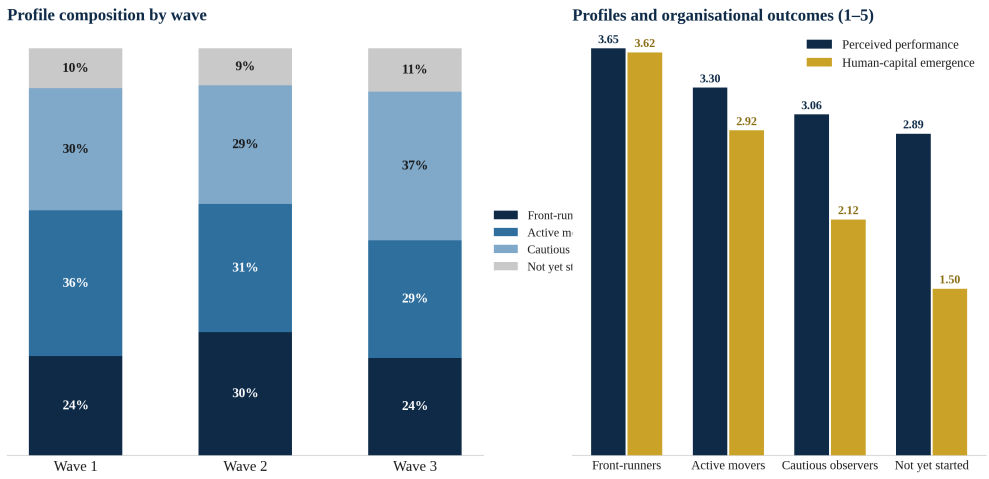


Source: AIHRM Longitudinal Survey. Model estimated on standardised variables; original scale means shown.

The profile structure is highly stable across waves, with the share of each type moving by no more than three percentage points, indicating that this classification captures relatively durable positional differences among firms. The gradient on the outcome dimensions, by contrast, is steep: from those not yet started to the front-runners, both subjective performance and human-capital emergence rise step by step, with the front-runners significantly ahead on both. The profiles can therefore serve as a compact diagnostic framework: a firm can locate itself against the four sets of characteristics and direct its catch-up effort at the weakest state indicators.

Exhibit 6-3

Profile structure holds steady across waves, but the outcome gradient is steep



Source: AIHRM Longitudinal Survey.

6.3 Interpretation and implications

The discriminating structure of the four profiles yields a counter-intuitive but important conclusion: Chinese firms have converged close to consensus on why AI should be used, and cautious observers endorse the augmentation logic almost as strongly as front-runners do. Divergence occurs in whether resources and capabilities can be assembled. It follows that internal debate about AI, if still conducted at the level of principle, is largely solving a problem that no longer exists. The agenda for management meetings should move from whether to act to what is missing, who will close it, and by when.

The profile framework can be used directly as a self-diagnostic. An HR leader can locate the firm against the dimensional scores of the four types. If all state indicators are moderately low while motive scores are not, the firm is a typical cautious observer: the shortfall lies in preparation rather than willingness, and the remedy is to concentrate resources on closing gaps in data and staffing rather than continuing to conduct feasibility studies. If state indicators are moderately high, the firm is an active mover, and its risk is stalling at the pilot stage; the remedy is to select a use case, take it deep, and put a quantifiable evaluation mechanism in place. If the firm is already a front-runner, the real challenge shifts to governance and organisational fairness — precisely the subjects of Chapters 7 and 8.

The high stability of the profile structure over three years is itself a reason for managerial vigilance. It indicates that a firm's relative position in this distribution is sticky: laggards do not catch up automatically as the field as a whole advances, and changing position requires deliberate, structural investment. Against this, the outcome gradient is steep, with performance and emergence rising step by step from the not-yet-started group to the front-runners. Stickiness of position combined with a steep outcome gradient points to a practical judgement: the window for catching up exists, but it does not open by itself.

7

CHAPTER 7

Challenges and Governance: Application Has Outrun the Rules

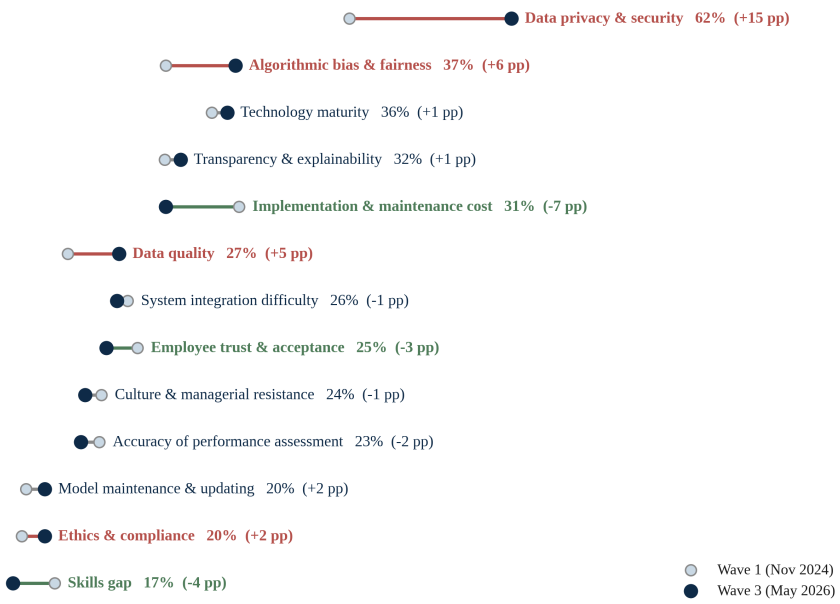
7.1 How implementation challenges have shifted

Tracking implementation challenges shows a structural shift in the list of difficulties firms face. Mentions of data privacy and security rose from 47.3% in Wave 1 to 61.8% in Wave 3, an increase of nearly fifteen percentage points and the largest of any challenge. Mentions of algorithmic bias and fairness rose from 30.8% to 37.1%, and data quality from 22.0% to 26.6%. Over the same period, mentions of implementation and maintenance cost fell from 37.4% to 30.8%, and mentions of the skills gap also declined. This pattern of rises and falls describes a clear transition between phases: as tool prices fall and experience accumulates, entry barriers of cost and skill are easing, while endogenous risks that grow as application deepens — privacy, bias and data quality — are becoming firms' principal concerns.

Exhibit 7-1

Data privacy jumps to the top; cost and skills concerns recede

Share of firms citing each challenge: Wave 1 vs Wave 3 (red = up >2 pp, green = down >2 pp)



Source: AIHRM Survey Waves 1 and 3 (multiple response). pp = percentage points.

7.2 The lag in governance

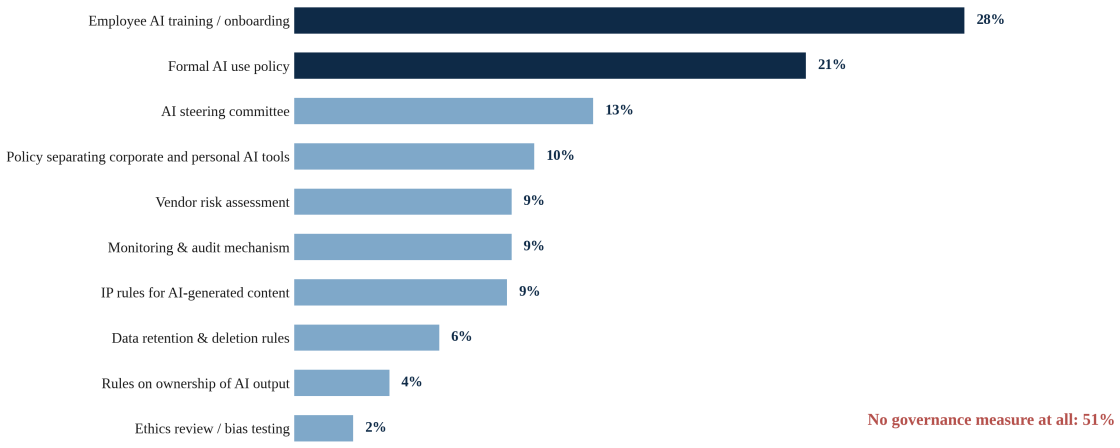
In sharp contrast to rising risk awareness stands a widespread absence of governance. Wave 3 listed ten common AI governance measures; 51% of firms had established none of them. Even employee AI training, the most widely implemented measure, is in place at only 28.1% of firms; formal use policies cover 21.5%; 12.5% have established a dedicated steering committee or decision body; and ethics review and bias testing

cover less than 3%. Setting the chapter's two findings side by side reveals an obvious gap: firms know clearly where the risks lie, yet most have not translated that awareness into institutional arrangements. Given that 86% of firms remain in the exploration or pilot stage, there is still a window for building governance — but the window will close quickly as application scales.

Exhibit 7-2

A governance vacuum: over half of firms have no AI governance measure in place

Share of firms with each governance measure established (Wave 3, n=526, multiple response)



Source: AIHRM Survey Wave 3.

7.3 Interpretation and implications

The shift in AIHRM implementation challenges provides executives with an updated risk register. Falling mentions of cost and skills indicate that the barrier to entry is genuinely coming down, while rising concern about data privacy, algorithmic bias and data quality identifies risks that share a common property: they grow as application deepens, and they are endogenous to use. The managerial implication is that the risk assessment for an AI project cannot simply carry over the list drawn up at approval. What firms worry about during pilots is whether they can afford the tool and whether people can use it; what determines success once deployment scales is whether the data can be used, whether results can be explained, and who is accountable. HR leaders should initiate a deliberate reassessment of the risk register at the point where a project moves from pilot to rollout.

The gap revealed by the governance data is more urgent still. More than six in ten firms name data privacy as their leading challenge, yet 51% have no governance measure of any kind, the most widely adopted measure reaches only 28.1%, and ethics review and bias testing sit below 3%. This fissure between risk awareness and institutional arrangement is the clearest managerial shortfall observed in this study. It bears emphasis that this is not an issue that can be deferred: the cost of establishing governance during the pilot phase is far lower than after scaling, and once AI is embedded in the formal processes of recruitment

screening or performance appraisal, retrofitting review mechanisms means overturning established processes and confronting compliance consequences that may already have occurred.

The concrete recommendation for HR leaders is to treat governance as a precondition for scaling rather than a parallel option, beginning with the measures that cost least and generate least controversy. First, a clear use policy defining which situations permit AI use and which decisions must retain human review. Second, a definition of data boundaries specifying which employee data may enter models, how long it is retained and how it is destroyed. Third, training for users, currently the most widely adopted measure and the easiest to advance. Fourth, bias-testing mechanisms for high-risk uses such as recruitment and performance. With 86% of firms still exploring or piloting, this window is precisely the best moment to establish a governance framework, and the cheapest — but it will close quickly as application scales.

8

CHAPTER 8

Organisational Consequences: Returns, Gaps and Restructuring

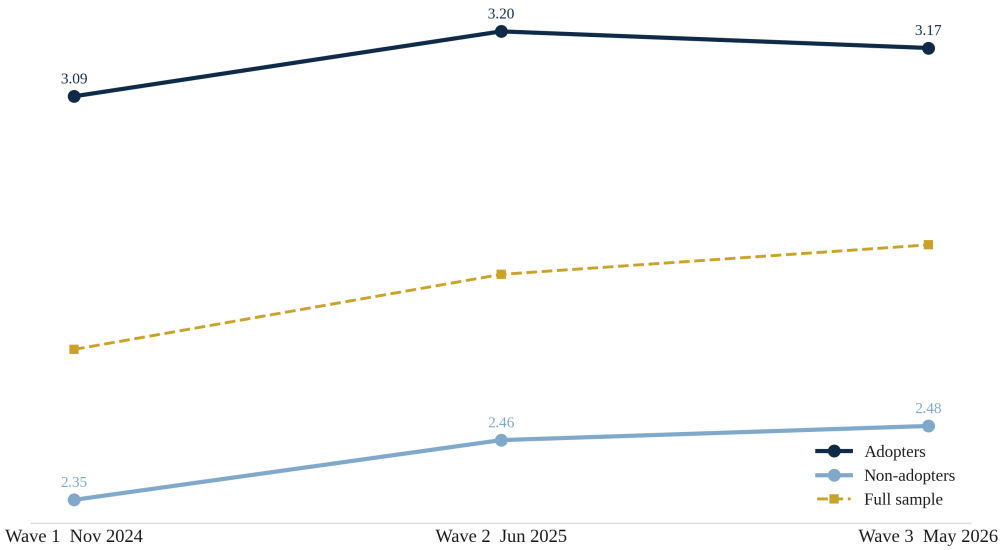
8.1 Human-capital emergence continues to strengthen

Against a background in which most attitudinal indicators pulled back, the emergence scale — which measures the extent to which generative AI converts individual knowledge and skill into collective team capability — is one of the few indicators to rise wave on wave, with the full-sample mean climbing from 2.62 in Wave 1 to 2.81 in Wave 3. Adopters score significantly higher than non-adopters in all three waves, and the gap between them has not narrowed. The steady rise in emergence sends a signal different from the sentiment indicators: within organisations actually using AI, gains in individual capability are settling into capability at the team level — precisely the step through which, in human capital theory, technological investment becomes organisational capability.

Exhibit 8-1

Human-capital emergence strengthens wave on wave — a rare consistently rising outcome

Eight-item scale: generative AI converting individual knowledge into team capability (1–5)



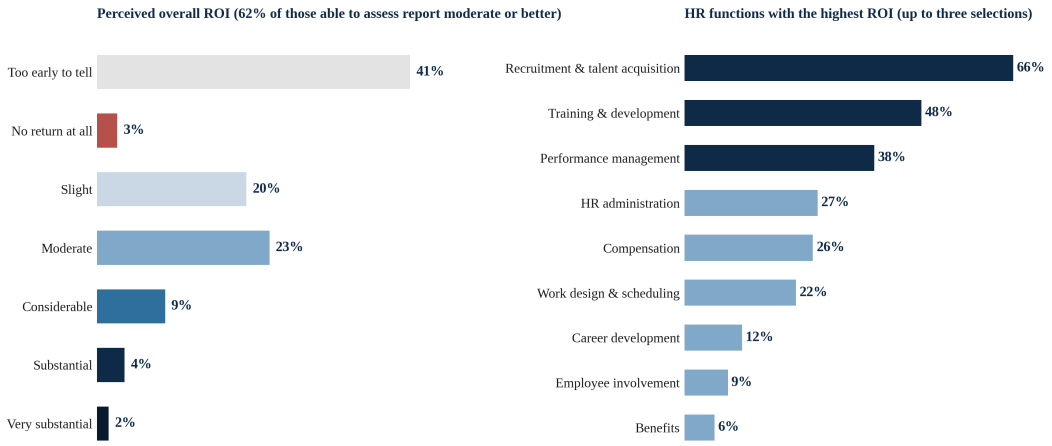
Source: AIHRM Longitudinal Survey.

8.2 Return on investment: early realisation within an observation period

Direct measurement of return on investment in Wave 3 shows that assessment remains in an observation period: 41.1% of firms consider it too early to draw conclusions. Among firms able to make an assessment, roughly six in ten perceive a moderate or better return, and very few report no return at all. By function, recruitment and talent acquisition and HR administration are named by the most firms as the areas of highest return, overlapping closely with the leading tier of module penetration in Chapter 4. The path to realising returns is thus visible: AI first produces perceptible efficiency gains where processes are dense and tasks standardised, while in judgement-intensive areas such as performance and compensation the validation of value still requires time.

Exhibit 8-2

Returns enter an observation period: four in ten say it is too early to judge



Source: AIHRM Survey Wave 3 (n=526).

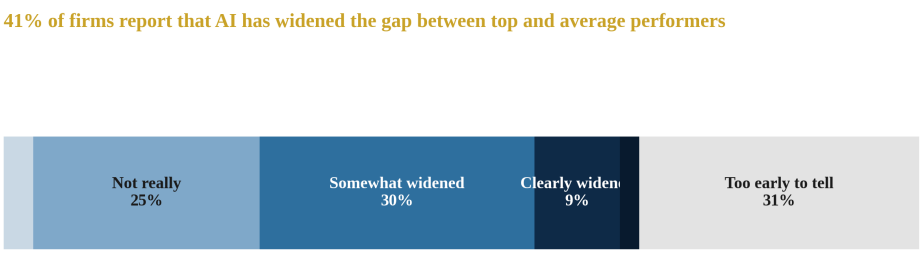
8.3 The widening performance gap

The benefits of AI are not evenly distributed across employees. Wave 3 asked whether firms had observed AI widening the gap between high and average performers; 41.4% answered affirmatively to some degree, with about one in ten reporting that the gap had widened clearly or substantially. This observation runs in the same direction as recent scholarly discussion of generative tools amplifying returns to ability: employees adept at working with the tools gain greater output leverage. For executives this is at once a new signal for identifying high-potential talent and a new problem of fairness management. Allowing the gap to widen unchecked risks eroding both collaboration and the perceived internal fairness of the reward system.

Exhibit 8-3

AI is amplifying the star-employee premium: nearly half see a widening gap

"Has AI use widened the gap between your top and average performers?" (Wave 3, n=526)



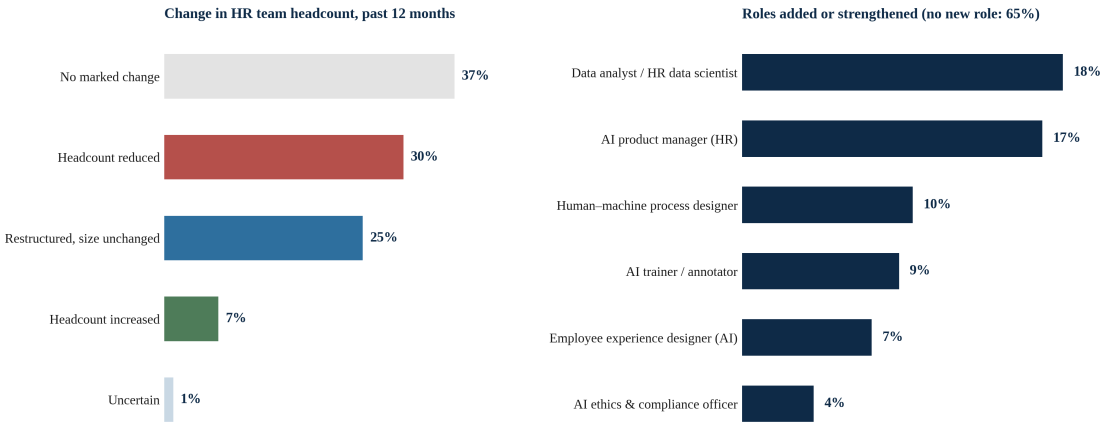
Source: AIHRM Survey Wave 3.

8.4 The HR function restructures itself

AI is changing not only the objects of human resource management but the function itself. Over the past twelve months, 30.2% of firms reduced HR team headcount overall, 25.1% restructured while holding headcount constant, and only 6.8% increased it. Running alongside contraction is renewal of roles: data analysts are the role most frequently added or strengthened, and new positions such as AI product manager and human–machine process designer are beginning to appear. Comparing identical items across Waves 2 and 3 shows across-the-board increases in the share of firms eliminating existing positions, redefining job duties, and launching reskilling and redeployment programmes: structural adjustment is accelerating. The parallel movement of contraction and renewal indicates that what the HR function is undergoing is not simple downsizing but a substitution in the composition of skills.

Exhibit 8-4

Contraction and renewal run together: teams shrink while data roles are added

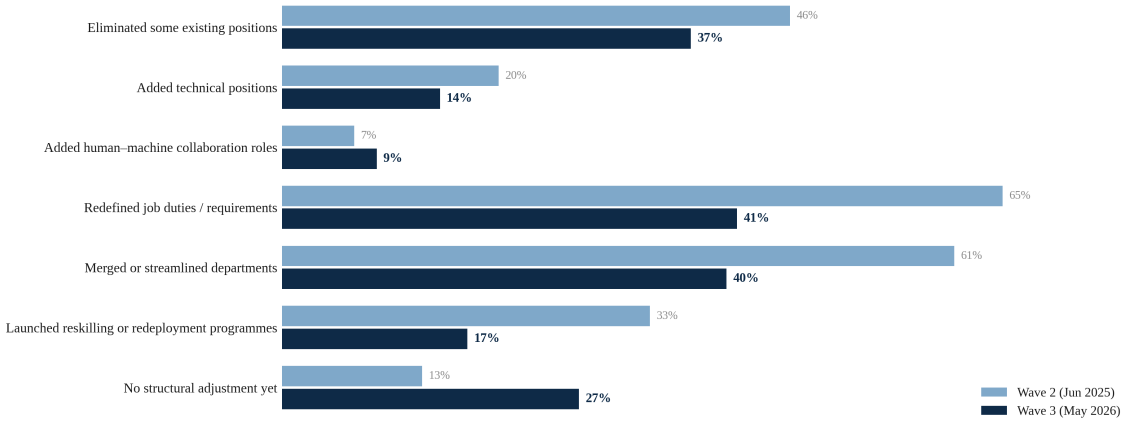


Source: AIHRM Survey Wave 3 (n=526).

Exhibit 8-5

Structural adjustment accelerates: redefined roles and reskilling lead the way

Share of firms making each HR structural adjustment in the past six months



Source: AIHRM Survey Waves 2 and 3 (multiple response).

8.5 Depth of application and organisational outcomes

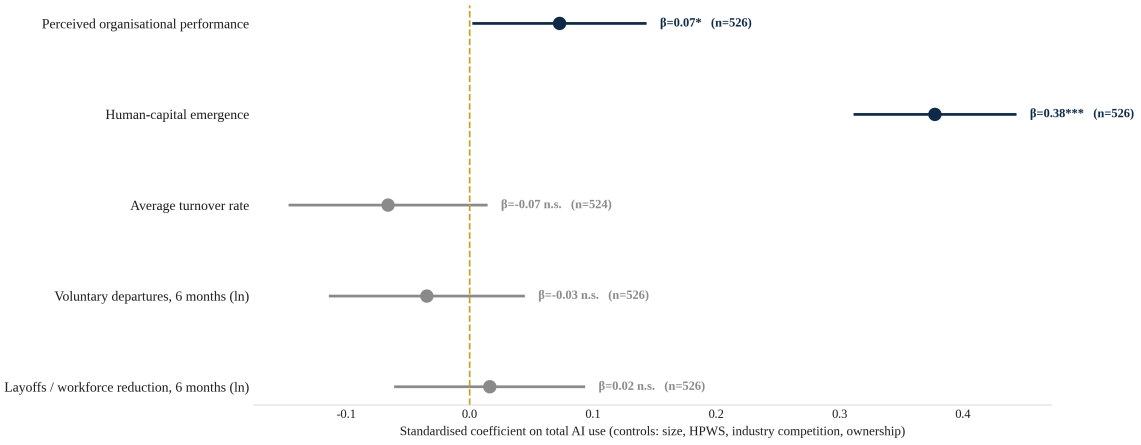
Regressions using total AI applications as the explanatory variable, controlling for firm size, the high-performance work system, industry competition and ownership, show that firms with deeper application report higher human-capital emergence and higher subjective performance, both coefficients significant. At the same time, total application is also significantly positively associated with employee turnover, voluntary departures and workforce reductions. The coexistence of positive outcomes and higher workforce churn is consistent with the team's earlier conclusion from two waves of matched data: the effect of AI in

reducing headcount appears relatively quickly, while its effect on performance requires longer to verify. Cross-sectional data cannot establish the direction of causation — accelerated churn may be a consequence of application, or may be intertwined with concurrent organisational adjustment — but the simultaneous rise of both categories of indicator is enough to remind executives that implementing AI is also a workforce transition requiring active management.

Exhibit 8-6

Deeper AI use goes with stronger capability emergence — and with higher workforce churn

Association between total AI use and organisational outcomes (Wave 3 cross-section, 95% CI)



Source: AIHRM Survey Wave 3. Cross-sectional associations do not establish causation.

8.6 Interpretation and implications

The data in this chapter show that the benefits and the costs of AI arrive together. On one side, emergence strengthens wave on wave, roughly six in ten firms able to assess returns perceive a moderate or better return, and the value of efficiency-oriented modules is being realised. On the other, firms with deeper application show higher workforce churn, three in ten have contracted their HR teams, and four in ten observe a widening performance gap among employees. Executives who prepare to capture the former without preparing to manage the latter will find the organisational costs of transition appearing as attrition and damaged morale.

The state of return assessment points to a second managerial lesson. That 41.1% of firms consider it too early to judge is itself a signal: most firms did not establish measurement standards when they committed to AI. Rather than asking after the fact what the return was, HR leaders would do better to fix baselines and definitions at approval — time and cost per hire, staff hours consumed by administrative work, first-contact resolution rates for employee queries. The concentration of perceived returns in recruitment and HR administration is instructive here: metrics in these two areas are the easiest to quantify and therefore the easiest to perceive as returns. Choosing measurable use cases first is itself a strategy for making the transition sustainable.

The widening performance gap is this chapter's most direct challenge to management. That 41.4% of firms observe AI amplifying the gap between high and average performers is both a new signal for identifying

potential and a new source of pressure on internal fairness. If the gap stems from unequal access to tools, the remedy is broader provision; if it stems from differences in the ability to use them, the remedy is training; if it stems from differences in how amenable particular tasks are to AI, then the evaluation system needs correction so that a technological dividend is not misread as individual capability. The risk of allowing the gap to widen unchecked is that it erodes both the willingness to collaborate and the perceived internal fairness of the reward system.

The restructuring of the HR function requires leaders to confront a question about their own position. Contraction in headcount and renewal of roles are occurring together, with data analysts the most common addition, which indicates that the function is undergoing not simple downsizing but a substitution in the composition of skills. The recommendation for incumbent teams is to begin a skills audit early: identify which work will be absorbed by tools, which capabilities need strengthening, and what transition paths exist for current members, and plan reskilling and redeployment as part of the deployment programme alongside technology investment rather than as an afterthought once reductions have been decided. The HR function occupies a particular position in this transition — it is both the designer of the change and its object — and that dual role means the whole company will watch how it handles the task.

9

CHAPTER 9

Findings from the Three-Wave Panel

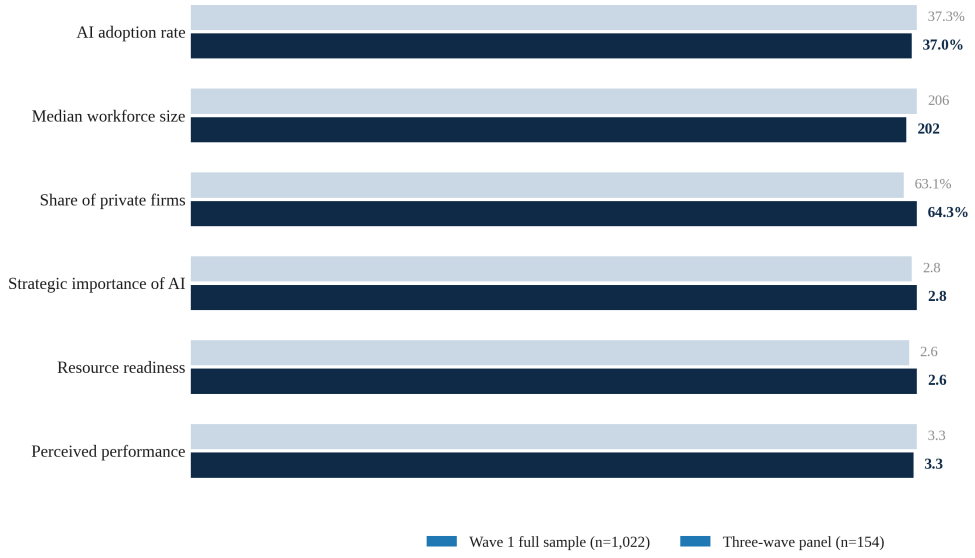
9.1 Representativeness of the panel

The 154 firms that took part in all three waves are the most valuable asset of this study: three consecutive observations of the same firms allow change in the environment to be separated from change within firms. Using a panel for inference presupposes representativeness, so we compared panel firms against the full Wave 1 sample on baseline characteristics item by item. The two are closely aligned on adoption rate, workforce size, ownership composition, the strategic importance placed on AIHRM, resource readiness and performance, with no statistically significant differences detected. The panel therefore shows no marked self-selection bias, and its conclusions have a basis for extrapolation to the wider population.

Exhibit 9-1

The 154 continuously tracked firms closely resemble the full sample

Baseline characteristics, panel vs full Wave 1 sample (bars normalised to the larger value in each pair)



Source: AIHRM Longitudinal Survey. No statistically significant differences were found.

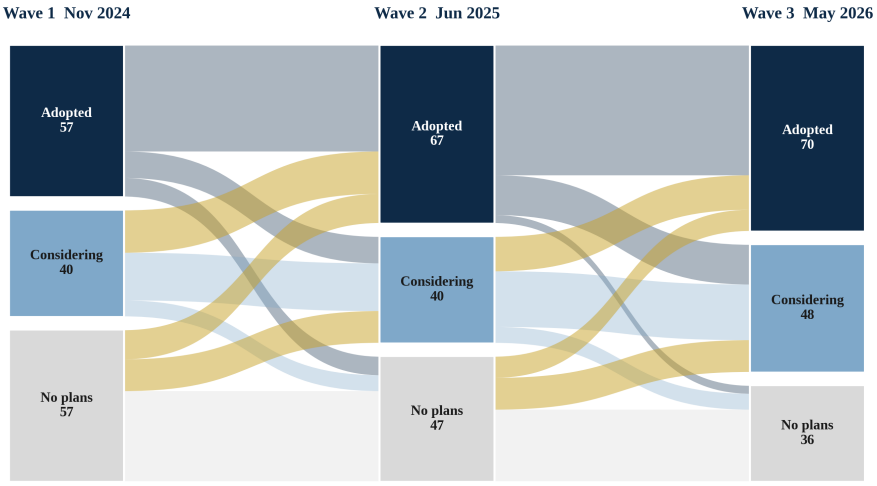
9.2 Wave-to-wave migration in adoption status

Classifying panel firms as adopted, considering, or no plans, we traced their migration wave by wave. Adopters rose from 57 firms in Wave 1 to 67 in Wave 2 and 70 in Wave 3, while firms with no plans fell from 57 to 36. Across the two intervals there were 48 transitions into adoption from the considering or no-plans states, against 35 movements in the opposite direction. Migration shows a clear net inflow, but exit is not rare: roughly two exits accompany every three entries, indicating that early adoption at some firms was exploratory and that they retrenched after testing it in practice. This detail corrects any impression of diffusion as a smooth one-way process: beneath the steady rise of the aggregate curve lies active trial and error at firm level.

Exhibit 9-2

Three years of migration among 154 firms: movement runs overwhelmingly toward adoption

Wave-to-wave transitions in adoption status (Sankey diagram)



Gold flows = moves toward adoption. Across the two intervals there were 48 entries into adoption and 35 exits — net movement is clearly inward.

Source: AIHRM three-wave panel (n=154).

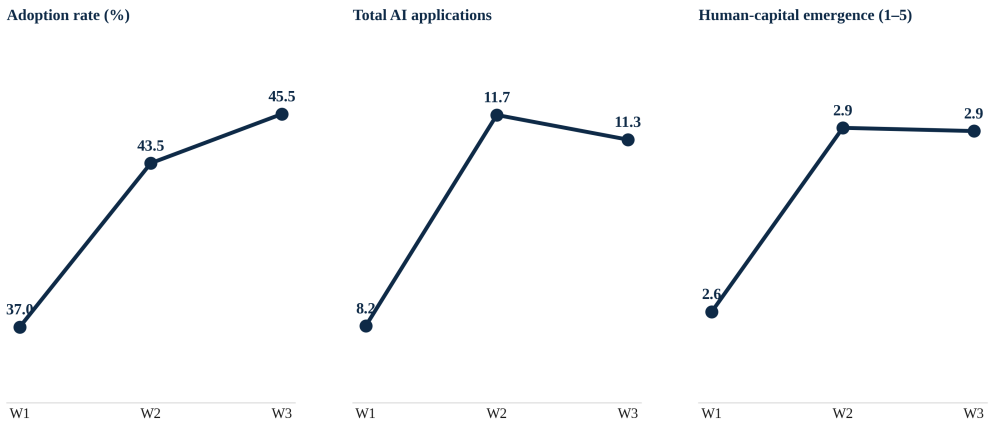
9.3 Trajectories of core indicators among the same firms

Within-firm paired tests provide the most rigorous confirmation of the overall trend. Among the 154 firms, the total number of AI applications rose from an average of 8.2 in Wave 1 to 11.3 in Wave 3, and mean emergence rose from 2.57 to 2.91, both increases significant at the .001 level; the adoption rate rose from 37% to 45%, an increase that is marginally significant given the size of the panel. The plateauing of total applications at a high level between Waves 2 and 3 also deserves attention: it corroborates the pullback in readiness observed in Chapter 5, suggesting that the fastest phase of expansion may have passed and that firms are shifting from adding applications to absorbing and deepening those they have.

Exhibit 9-3

The same firms, three years on: adoption, application count and emergence all rise

Within-firm change among the three-wave panel (n=154)



Source: AIHRM three-wave panel. Paired tests: total AI use and emergence rise significantly (p<.001); the adoption increase is marginal (p=.069) given panel size.

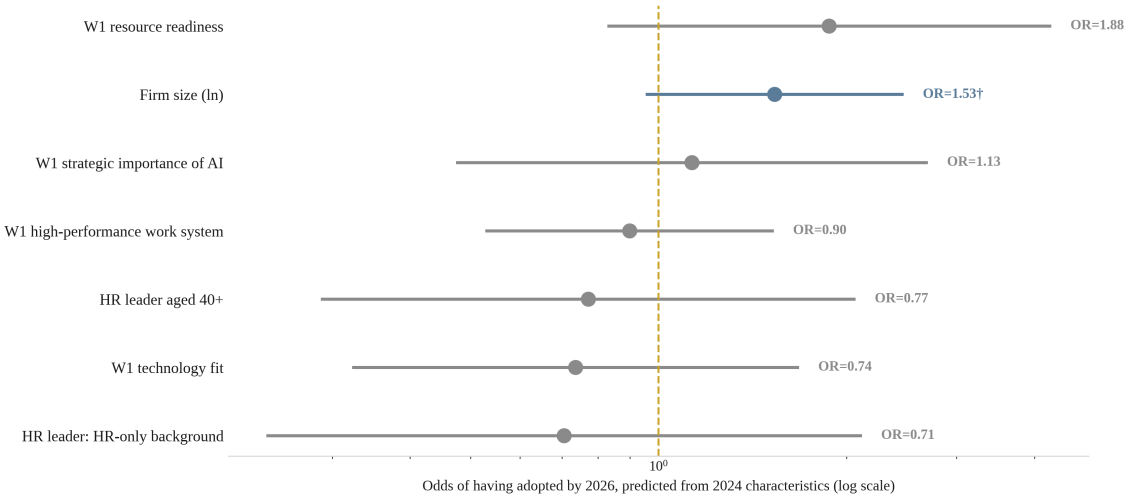
9.4 From watching to doing: which firms converted

The panel design allows a forward-looking question: among firms that had not adopted in 2024, which characteristics foreshadowed implementation two years later? A logistic regression on the 93 panel firms that had not adopted in Wave 1 shows a two-year conversion rate of 32.3%. Wave 1 resource readiness and firm size show directional advantages — an odds ratio of 1.88 per standard deviation of readiness and 1.53 for size — but given the sample size these fall respectively outside and at the margin of statistical significance, and no other variable shows predictive power. The prudent reading is that early resource preparation may constitute a precondition for later conversion, but the test based on 93 firms has limited power, and confirmation awaits a fourth wave with an expanded panel.

Exhibit 9-4

From watching to doing: early readiness and size point the way, though not decisively

Conversion among panel firms that had not adopted in Wave 1 (n=93; 32% converted)



Source: AIHRM three-wave panel. † p<.10. Readiness OR=1.88 (p=.13); size OR=1.53 (p=.08).

9.5 Interpretation and implications

The value of panel data is that it removes the confound of sample turnover and lets executives see genuine change within the same set of firms. Three conclusions carry direct relevance. First, the broad direction of transition is unmistakable: migration among the three adoption states shows a clear net inflow, and the no-plans camp contracted from 57 firms to 36. Second, the transition is not smooth: 48 entries were accompanied by 35 exits, roughly two exits for every three entries, indicating that early adoption at a substantial share of firms was exploratory and that they narrowed their front after testing it. This detail is both a caution and a reassurance for executives: withdrawing from a pilot that does not hold up is not failure but a normal feature of diffusion. The real failure is to neither deepen nor withdraw in the absence of an evaluation mechanism, leaving projects suspended and consuming resources.

Third, total applications among the same firms plateaued at a high level between Waves 2 and 3 rather than continuing to expand rapidly. Combined with the pullback in readiness observed in Chapter 5, a reasonable judgement is that the dividends of the spreading phase have largely been harvested and that competition in the next phase will occur in absorption and deepening. For firms that have completed initial deployment, the marginal return to adding further applications is falling; taking existing applications deep enough to replace the original process is now the more valuable direction of investment.

Although the conversion analysis does not reach statistical significance given the sample size, its direction is still worth noting: firms with higher resource readiness and larger size in 2024 were more likely to have adopted two years later, while the strategic importance they placed on AI at the time and their technology fit showed no predictive power. Should this direction be confirmed in later waves, the managerial implication would be clear — what determines whether a firm moves from watching to doing is whether it

was quietly laying foundations two years earlier, not how enthusiastically it spoke at the time. For firms still watching, this points to a concrete action: even if application is not yet under way, preparatory work on data governance, system connectivity and staff capability should begin now, because these are the real preconditions for later conversion.

10

CHAPTER 10

Conclusions, Implications and the Road Ahead

10.1 What three years of observation establish

The three waves support the following conclusions. AIHRM adoption among Chinese firms has moved past the early novelty stage: an adoption rate approaching half, the continuing contraction of the wait-and-see camp, and genuine conversion within the matched panel together indicate that the entry of AI into human resource management is now an irreversible trend. Yet the distribution of maturity shows this is a transformation that has only just entered deep water — fewer than 3% of firms have reached scale or deep integration, and breadth of application runs well ahead of depth. On drivers, disagreement at the level of awareness and motive is now small, resource readiness has become the variable that separates firms, and the professional background of the HR leader shows a robust effect on gains in adoption across two independent time windows, most likely transmitted through the organisation's state of readiness. On consequences, returns in efficiency-oriented modules are beginning to materialise and human-capital emergence continues to strengthen, but a widening performance gap, accelerating workforce churn and the contraction and restructuring of the HR function itself are occurring alongside them: benefits and costs arrive together. On institutions, governance lags markedly behind application — risk awareness has clearly risen, yet more than half of firms still have no governance arrangement of any kind.

10.2 Recommendations for executives

First, shift the centre of gravity from spreading to deepening. The data show breadth near its ceiling while depth remains generally insufficient, and returns concentrated in a few modules. Rather than distributing effort evenly across nine modules, firms would do better to select use cases with demonstrated returns — recruitment, administration — and drive a single application to the depth of process redesign, establishing a replicable internal benchmark before expanding. Second, move governance ahead of scaling. The cost of establishing use policies, data boundaries, ethics review and training is far lower during the pilot phase than the cost of retrofitting after scaling; data privacy has become the leading challenge while more than half of firms have no governance arrangement at all, and this gap should be closed before application scales. Third, treat leadership appointment as part of the transformation plan. The data show that firms led by younger, cross-functional HR leaders progress faster, and that firms whose leader holds a single professional background show lower gains in adoption across both time windows. Firms should weigh data literacy and cross-functional experience when appointing HR leaders, and provide incumbent teams with systematic training in technical and data capability. Fourth, manage the human side of the transition actively. The co-occurrence of deeper application with workforce churn and headcount adjustment indicates that reskilling and redeployment should not be remedial measures taken after reductions, but components of the deployment plan scheduled alongside technology investment. Performance dispersion created by tool use should be addressed through broader training and adjustments to evaluation, so that it does not harden into a new form of internal inequality.

10.3 The technological frontier: from assistive tools to digital labour

Before offering any judgement about the end state, it is worth being clear about where the technology itself is heading. Over the past three years, the form AI takes inside organisations has moved through three steps. It entered first as point tools performing well-defined tasks such as screening, matching and calculation. It then embedded itself as a copilot within everyday work software, collaborating with users in real time to generate content and process information. The current frontier is the agent: a system that can autonomously link multiple steps toward a goal, invoke a range of tools and make decisions dynamically. Wave 3 data provide a cross-sectional gauge of this ladder. Résumé screening, belonging to the first step, is in routine use at 62.4% of firms. At the second step, generative course content creation and copilot tools embedded in office software reach 28.9% and 11.3% respectively. At the third, autonomous end-to-end recruiting agents and workflow-orchestration agents both sit at around 8%. The coexistence of all three forms with steadily declining penetration indicates that Chinese firms are in the midst of a large-scale migration from the first step to the second, with only scattered probing of the third.

International frontier practice suggests what the upper end of the ladder may look like. Some leading firms have handed sourcing, initial screening, communication and interview scheduling to agents executing end to end, with people retaining only final review and involvement in critical roles. Others have begun creating employee-like management records for agents, bringing them into headcount and permission systems, so that digital labour ceases to be a metaphor and becomes an object of management. In the underlying logic of talent management, systems organised around skills rather than jobs, together with internal talent marketplaces, are spreading, with firms continuously parsing employees' skill composition to make deployment, development and succession decisions. Organisational structure has produced its own landmark: the American biotechnology firm Moderna merged its human resources and information technology functions into a single people and digital technology organisation, headed by its former chief human resources officer. The logic, as industry commentary put it, is that HR understands work and technology understands tools, and when work and tools become deeply entangled the two must be designed at the same table.

The lag between technological capability and organisational application is the key variable in judging the pace of what follows. A report published by Anthropic in early 2026 noted that AI already possesses the capability for most occupational tasks, while actual application trails the capability frontier significantly. This report's maturity data echo that observation: adoption approaches half while fewer than 3% have achieved scale or deep integration, and the distance between capability and application constitutes the principal potential energy behind the changes of the next several years. A separate study by our team, based on recruitment records from Chinese listed companies, points in the same direction: average AI exposure across all advertised positions peaked in 2019 and has declined since. That decline is not a sign of weakening impact — quite the reverse. It means hiring demand for highly exposed positions is contracting, and that substitution effects have begun to show up in the structure of the labour market.

10.4 Projecting the end state: possible forms of AIHRM

Extending the line from three waves of data and current frontier practice, we offer five projections about the medium- to long-term form of AIHRM. These are extrapolations from evidence rather than predictions, and the speed at which they are realised will depend on the joint evolution of technology, cost and institutions.

First, the transactional layer will approach complete absorption by systems. Chapter 4 shows that the leading practices in all nine modules today are almost entirely screening, monitoring and process automation — tasks with clear boundaries, stable rules and correctable errors, precisely the territory automation exhausts first. Extending that line, the great majority of transactional work — attendance and leave, payroll calculation, policy enquiries, records management, onboarding and offboarding — will be absorbed wholesale by the platform layer within the next several years. The daily interface between employees and the HR system will shift from a person to a system, and human involvement will contract to exception handling and appeals. The disappearance of the transactional layer does not mean the disappearance of the HR function; it means a fundamental rewriting of its cost structure and its staffing.

Second, recruitment will be the first module to be rebuilt completely. Résumé screening is already routine, interview applications performing live conversational assessment are climbing quickly at 23.2% penetration, and autonomous end-to-end recruiting agents already appear in 8% of firms. The full presence of all three technological forms within a single module makes recruitment the most likely to complete the journey from tool to agent first: for standard roles, sourcing, screening, communication and assessment will be orchestrated end to end by agents, and the human role will contract to setting the standards, adjudicating exceptions and giving final approval for critical positions. At that point the core competitive capability in recruitment shifts from execution efficiency to two human capabilities: the judgement to set the right standards for the agent, and the relational skill to build trust with high-value candidates.

Third, the object of management will extend from people to hybrid human–machine teams. As agents take on ever more complete units of work, the organisation effectively acquires a new class of labour whose deployment, authorisation, evaluation and retirement all require rules. A saying has circulated in the industry that the HR department will cease to exist, but the original formulation ran in precisely the opposite direction: information technology departments will need to manage agents the way HR departments manage employees. That formulation points to a more likely trajectory in which the methodology of human resource management spills over to become the template for governing digital labour, with job descriptions, permission boundaries, performance standards and behavioural norms — institutional instruments built for people — rewritten for the management of machines. Designing the division of labour in human–machine collaboration, and setting behavioural boundaries for agents, may become the most value-adding new task for the HR function in the next phase.

Fourth, the basic unit of talent management will move from the job to skills and tasks. Our team's analysis of position data from S&P 500 firms shows that a large number of job titles have remained unchanged while their skill cores have been replaced, so that two different kinds of people now coexist under the same title. Once job labels cease to be reliable, a skills-based talent ledger becomes the basis for deployment and

development decisions, and internal talent marketplaces displace static job hierarchies as the main channel of talent mobility. At the same time, the widening performance gap documented in this report means that the performance distribution itself is being reshaped by technology, with those adept at using the tools gaining amplified output leverage. Systems for identifying and evaluating talent will need to be rewritten accordingly, so as to distinguish the dividend of the tool from genuine individual capability.

Fifth, the path to the end state will not be a straight line. The panel data in this report show roughly two exits for every three entries into adoption: trial, error and reversal are the norm of diffusion rather than the exception. International experience gives equal reason for tempered optimism. The Swedish fintech firm Klarna replaced several hundred customer-service staff with AI and subsequently rehired human employees after service quality declined; and statistics on the more than one million layoffs announced in the United States in 2025 show that fewer than five per cent were explicitly attributed by the firms themselves to AI, with a great many other motives mixed into workforce adjustments made in technology's name. Economics has long held that the net employment effect of technology depends on firms' choices between augmentation and substitution, and that this choice is ultimately determined by institutions and management rather than by technology itself. The direction of the end state is clear; how organisations arrive there remains in their own hands.

10.5 HR professionals in the age of AI: the market's verdict and the rebuilding of capability

Once AI has reshaped human resource management, who will do the work of HR? This was the question respondents raised most often during the study. Rather than rely on opinion, we let the labour market vote. Our team built a database of more than 710,000 HR job advertisements spanning 2014 to 2025; what firms actually wrote in their job descriptions constitutes the most honest record of how this occupation is changing. The analysis shows that over twelve years the hard-skill requirements of HR positions shifted in three directions — toward data, toward performance and toward strategy. What the market expects of practitioners is moving from familiarity with system operation and policy detail toward the ability to analyse data and understand business strategy. On soft skills, communication and collaboration hold steady at a high level as baseline requirements, while genuine growth concentrates in two categories: taking responsibility and influencing others. The steepest decline is in skills related to standardised content delivery; in training and development, market demand for roughly seventy per cent of such skill items is contracting — the territory where generative AI substitutes most completely.

More striking still is the divergence within the occupation. Comparing the same HR positions grouped by salary reveals that requirements in the high-salary group are rising across the board, with data analysis, strategic understanding and team leadership all gaining weight, while requirements in the low-salary group have essentially stayed where they were. The same job title is heading toward two different fates: the divergence occurs not between occupations but within one. This pattern is structurally identical to the widening employee performance gap documented in Chapter 8 — technology amplifies differences in the returns to capability. Summing up twelve years of the market's voting, the conclusion can be stated plainly: what is being squeezed is not the HR occupation itself but the portion of it built on executing standardised

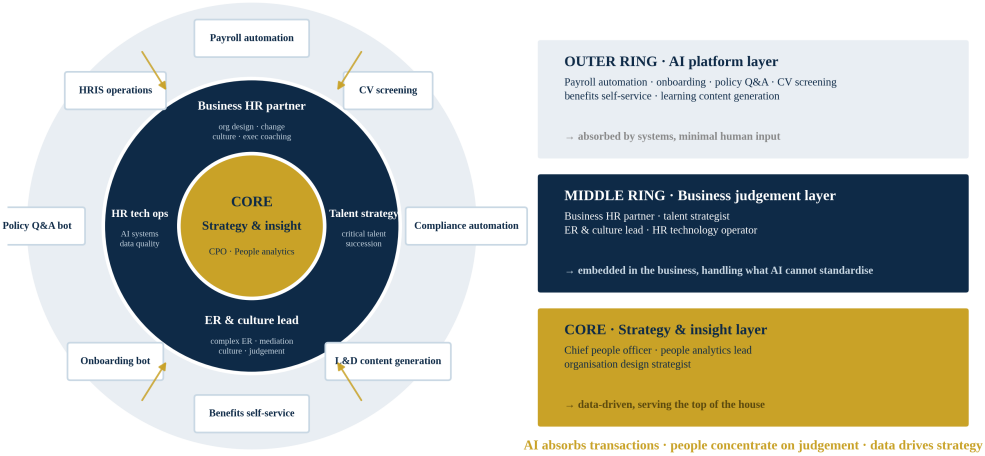
processes. The market value of practitioners who combine business understanding, data fluency and command of AI continues to rise.

Turning to organisational form, there is a structural question to answer. For the past twenty years, the organisational design of HR functions worldwide has been modelled almost universally on the three-pillar model proposed by Dave Ulrich — the division among centres of expertise, business partners and shared service centres. The arrival of AI is shaking the foundation of that architecture, because the standardised delivery carried by the shared service centre is precisely what systems absorb first and most completely. It may be time for the three pillars to evolve into another form. Our team summarises the structure of the future HR function as a three-ring model. The outer ring is the AI platform layer, where payroll automation, intelligent onboarding, policy enquiry, initial recruitment screening, benefits self-service and training content generation are absorbed by systems with minimal human input. The middle ring is the business judgement layer, composed of business partners, talent strategy, employee relations and culture, and HR technology operations — deeply embedded in the business and handling what AI cannot standardise. The core is the strategy and insight layer, composed of the chief people officer, the people analytics lead and the organisation design strategist, data-driven and serving the top of the house directly. The division of labour across the three rings can be captured in a sentence: AI absorbs transactions, people concentrate on judgement, data drives strategy. It should be said that the boundaries between the rings are not fixed. As AI capability continues to advance, some judgement work handled by people in the middle ring today may be absorbed into the outer ring tomorrow; the dividing lines will keep contracting toward the centre, and the territory on which people concentrate will migrate toward higher-order judgement and strategy. The three-ring model therefore describes not a static distribution of positions but a dynamic process of contraction. Staffing ratios across the rings will vary by firm size and industry, but the direction — value converging inward, the machine boundary advancing inward — is the same.

Exhibit 10-1

From three pillars to three rings: the boundaries keep contracting inward

Gold arrows: as AI capability advances, work migrates from the middle ring outward into the platform layer



Source: Research team. A directional projection based on current evidence; staffing ratios across rings will vary by firm size and industry.

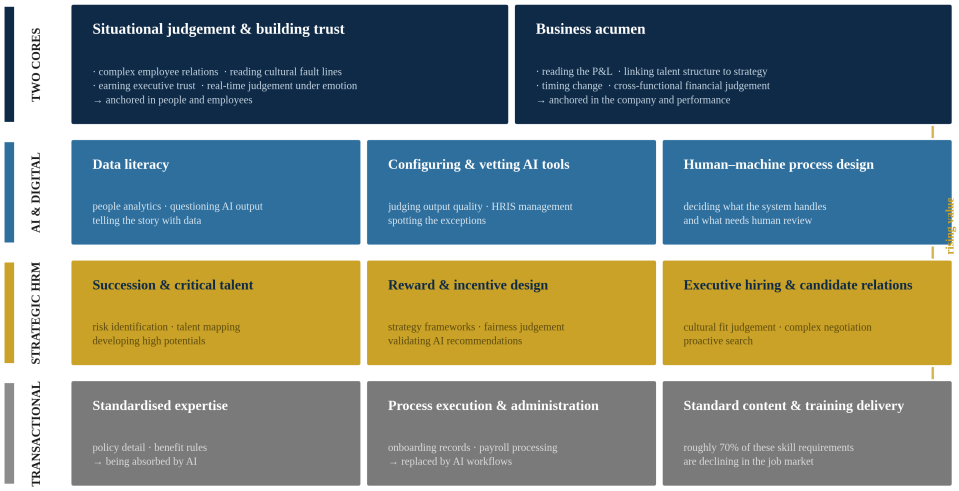
At the individual level, the capability map of the future practitioner takes the form of a hierarchy built from the bottom up. At the base lies transactional HR capability: standardised professional knowledge, process execution and administration, and the production of standard content and training delivery. Work at this level has clear boundaries and stable rules; it is being ceded to systems, and market demand for it continues to fall. Above it sits strategic HR capability: succession planning and critical talent management, reward and incentive design requiring judgements of fairness, executive hiring and candidate relationships, and organisation design and change implementation. These belong to the higher reaches of the traditional functional division of HR; they depend on experience and professional depth, and machines cannot take them over wholesale. On top of these two traditional layers, future practitioners must add a new one: AI and digital capability, comprising the data literacy to explain problems with data and to remain sceptical of algorithmic conclusions, the ability to judge the quality of AI output and identify exceptions, the process design skill to define the division of labour between people and machines, and an awareness of ethics and compliance covering privacy boundaries and algorithmic accountability. At the apex, and most critical of all, sit two core capabilities. The first is situational judgement and the building of interpersonal trust — reading emotion, forming judgements and earning trust in complex employee relations and in situations calling for executive confidence — anchored in a commitment to people and to employees. The second is business acumen — reading the profit and loss statement, connecting talent structure to business strategy, and anticipating the right moment for change — anchored in a commitment to the company and to performance. What these two core capabilities share is that they are situational, involve trust and resist standardisation, placing them precisely in the region hardest for machines to replicate; and they correspond, respectively, to the dual mission of HR work in serving people and serving the organisation. For practitioners, the practical step is to audit their own capability structure against this map: the more their

work concentrates in the transactional base, the more urgent their transition; and any gap in AI and digital capability or in the two capabilities at the apex should be the focus of remedial effort over the next two to three years.

Exhibit 10-2

The capability stack: transactional skills give way, two human cores rise to the top

Based on the team's analysis of 710,000+ HR job advertisements (2014–2025)



Source: Research team, HR job-advertisement database.

The Moderna case provides a concrete footnote to this chapter: when human resources and digital technology were merged into a single function, the person chosen to lead it was the former chief human resources officer rather than the technology head. The judgement behind that appointment is worth remembering. When firms begin to answer seriously which work goes to people and which to machines, the centre of that question is the organisation, the work and the people — precisely the domain in which the HR profession has accumulated the deepest expertise. Technology does not become benign on its own. Whether this transformation is defined by the logic of efficiency alone, or guided by professionals who understand the technology while retaining warmth toward people, depends on who occupies the driver's seat. Twelve years of the market's voting show that the ticket to that seat is a composite capability structure; and another fact recurring throughout three waves of our data is that the bottleneck in most firms' transformation is precisely the scarcity of HR professionals who understand both the business and the technology. For practitioners, that is at once a demand and an opportunity.

10.6 Limitations and future plans

The conclusions of this report are subject to several conditions. Sample composition differs across waves; although all key conclusions have been checked against the matched panel, the overall statistics should be

understood as a picture of participating firms rather than a nationally representative estimate in the strict sense. The principal variables are self-reported by HR leaders and may be affected by common method bias. Causal interpretation of cross-sectional associations remains cautious, and while the lagged tests provide stronger evidence, their sample sizes are still limited. The forward-looking material in sections 10.3 to 10.5 consists of extrapolations from current evidence and does not constitute a prediction of any specific timetable. Our team plans to conduct a fourth wave in 2027, with priority on expanding the panel of firms tracked across three or more waves, matching in objective operating data, and measuring governance development, agentic applications and employee-level usage behaviour in finer detail, so as to continue documenting this organisational transformation among Chinese firms.

Thanks

Acknowledgements

This study formally began in August 2024. In the preparatory phase, with the assistance of the Human Resources Association (HRA) under Beijing Foreign Enterprise Human Resources Service Co., Ltd. (FESCO) — and with special thanks to Ms Zhang Shuyan — our team conducted in-depth interviews with nearly ten firms already applying AI in human resource management. Those interviews allowed us to understand real application scenarios, difficulties and reservations before designing the questionnaire. Building on the interviews and a review of the literature, the team completed the questionnaire design and, with the support of the China HRD Club (with special thanks to Mr Cai Yuanqi), the HRA and the Peking University MBA Human Resources Association, launched the online survey in October 2024. The first phase concluded at the end of December 2024 and produced the preliminary findings of this report. Recruitment for the two subsequent waves relied principally on the same channels. Without these organisations' investment in firm outreach, channel access and sustained mobilisation, a three-wave study spanning nearly two years could not have been carried out, and we extend our sincere thanks.

From its inception the study has been funded throughout by the Guanghua Thought Leadership programme at the Guanghua School of Management, Peking University. The programme was established to encourage faculty to undertake exploratory work close to enterprise practice, laying foundations for future research without publication as the sole objective. It was precisely this orientation that gave our team the confidence to follow a rapidly changing subject at an unhurried pace: three waves of data collection, two years of tracking, and a large body of interviews and data that could not be converted immediately into papers were all completed under that support. The work also produced an unplanned dividend: during the study our team developed a course for MBA students, Business Applications of Artificial Intelligence, bringing the real practices of firms back into the classroom. We hope that this report and that course have honoured the founding intent of the Guanghua Thought Leadership fund.

We thank the two doctoral students on the team. Li Hangtao carried out the design of the online questionnaire in the preparatory phase and a great deal of the analysis across three waves; much painstaking technical labour lies behind the numbers in this report. Hou Wenxin managed the fielding of all three waves, follow-up contact with participating firms, and the logging of gift dispatch, along with a great many small and concrete administrative tasks — it was this interlocking execution work that kept a two-year longitudinal study running smoothly. It is worth noting that this is precisely the kind of work AI cannot yet fully replace: the professional judgement behind a questionnaire, and the patience and sense of responsibility required to stay in contact with several hundred firms, can still only be supplied by people. Without their efforts this study could not have been carried out.

We are especially grateful to every HR leader who took the time to respond across the three waves. The questionnaire was not an easy one: it was long, wide-ranging, and many items required respondents to retrieve specific internal information before answering, so that completing it once took considerable time and patience. That you answered carefully amid demanding schedules, and did so again when the second and third invitations arrived, is the most important foundation of this report. The 154 firms that participated in all three waves have allowed us to observe genuine movement within the same set of organisations over two years — evidence of a kind that is not common in management research. The time you gave has been

converted into public knowledge that peers can reference and those who follow can build on: a contribution to management scholarship and management practice alike.

As the three waves progressed, interim findings were presented at a number of academic and industry venues, including the Digital and Intelligent Organisation Symposium hosted by the School of Management at Zhejiang University in 2024; the 2025 Biennial Conference of the International Association for Chinese Management Research (IACMR), held alongside the tenth anniversary of the journal Management Insights; two consecutive years of the Super A Conference organised by the HRA under FESCO; the 2026 annual forum of the HR Excellence Center; and the launch meeting of the New-Quality Organisation study conducted with Ruihua Zhice. The questions, challenges and suggestions received at each of these occasions entered, to varying degrees, into subsequent revisions of the questionnaire and into the analytical framework of this report. We thank the organisers who provided these opportunities and the colleagues in academia and industry who offered feedback.

During the study we also received many letters. A number of correspondents expressed keen interest in the subject while also confessing to a certain anxiety: knowing that AI is changing human resource management, yet finding that their own organisations could not yet move, constrained by data foundations, resources or organisational awareness. We would offer this response: there is no need to be anxious. The data from three waves show that firms which have not adopted remain the majority, that the overwhelming majority of those that have adopted are still exploring or piloting, and that fewer than three per cent have reached scaled application. The gap today is far narrower than industry narratives suggest. As the technology matures and the cost of use continues to fall, availability will cease to be the privilege of a few firms, and broad diffusion can reasonably be expected. What we would most like to see at that point is HR professionals freed from heavy transactional work to do things of greater value and greater strategic significance — serving employees better and expressing more fully the humane care that HR work ought to embody. That is both the direction technological progress should take and the reason this study continues.

A final note. The writing of this report was substantially supported by Claude Fable 5; AI participated deeply in the production process, from data preparation and chart production to the revision of the manuscript, with all content reviewed carefully and adjusted repeatedly by our team before being finalised. We state this both as a matter of transparency in the use of AI tools — we believe that honest disclosure of how artificial intelligence participates in knowledge production should become a basic norm of scholarly work in this era — and because, for a report whose subject is artificial intelligence, this was also a gain in efficiency and a first-hand attempt by our team to embrace the technology we study.

The views and conclusions in this report are the responsibility of the research team; we welcome comment and correction on any shortcomings in the text.