

The Advertised Salary Premium for AI and Data Skills across UK HR Roles

A Quantitative Content Analysis
of Online Job Advertisements

Dianzi Wang

UCL School of Management

MSc People Analytics and Human-Centric Management

Abstract

The application of artificial intelligence is transforming the requirements for human resources work, while existing salary evidence often combines different occupations or groups several technical skills within broad categories. This study examines how AI and data skill requirements are associated with advertised salaries in UK HR job advertisements (RQ1) and whether these associations differ by seniority or HR functional area (RQ2).

The study employed quantitative content analysis and cross-sectional regression analysis. The sample contained 2,229 full-time, salary-disclosing HR advertisements collected from Reed.co.uk on 14 July 2026. AI-related requirements and Strict Advanced Data requirements were examined as separate measures.

After considering job vacancy characteristics, AI-related requirements were associated with 16.34% higher advertised fixed or basic salaries. The Strict Advanced Data estimate was 1.97% and was not statistically significant. Neither association varied systematically by seniority or HR function, and the main AI result remained positive across the robustness checks.

These findings provide focused vacancy-level evidence that advertised salary associations within HR can differ across specific technical requirement measures. Broadly stating that "both artificial intelligence and data skills have a wage premium" may obscure differences among various skill requirements. Since the study used a single-platform, cross-sectional sample with only 34 AI-positive advertisements, the generalizability of its findings may be limited.

Contents

Click a section to navigate.

Abstract	i
Chapter 1 Introduction	1
Chapter 2 Literature review	4
Chapter 3 Methodology	12
Chapter 4 Results	19
Chapter 5 Discussion	26
Chapter 6 Conclusion	33
References	35

Chapter 1: Introduction

1.1 Background and research context

Artificial intelligence has become increasingly widespread in British businesses. Among firms with at least 10 employees, the reported use of one or more AI technologies increased from approximately 12% in late 2023 to 35% in June 2026. However, adoption varied by firm size and industry. Only 10% of businesses using AI described their use as extensive (Office for National Statistics [ONS], 2026). AI alone does not create economic value. To use it effectively, firms must also invest in new processes, business models and human capital (Brynjolfsson et al., 2021). Evidence from U.S. companies further links investment in AI to growth and product innovation (Babina et al., 2024). As AI diffuses, firms face an organisational capability problem. They must identify the knowledge and expertise needed to implement these technologies and adjust their workforce accordingly.

Recent evidence shows that this adjustment can influence workforce composition as well as wider labour-market outcomes. Drawing on employee resumes and job-posting data, Babina et al. (2025) find that U.S. firms investing in AI tend to develop more highly educated workforces, with greater specialisation in STEM and IT skills. A similar pattern appears in 16 European countries. Albanesi et al. (2025) find that employment shares increased in occupations with greater exposure to AI, particularly those employing higher proportions of younger and skilled workers. However, they find little evidence that AI exposure is similarly related to relative wages across occupations. Changes in employment and workforce composition can therefore occur without a uniform wage pattern. In response, organisations may adapt their recruitment, employee development, role design, and career structures.

Current UK business data indicate that firms most commonly develop AI-related skills by training or retraining existing employees. The use of external recruitment rises with firm size, while a lack of expertise remains an important barrier to AI adoption (ONS, 2026). These organisational responses have direct relevance to HR. Drawing on discussions with 37 Fortune 500 chief human resource officers and their technology advisers, Nyberg et al. (2025) identify challenges associated with generative AI across recruitment, learning and development, performance management, job and work design, talent management, and compensation and benefits. These developments give HR two connected roles. As an organisational function, it helps firms adapt their wider workforce

to technological change. Concurrently, HR practitioners face changing employer expectations about the capabilities they should possess. This raises a labour-market question: are these changing expectations reflected in the salaries advertised for HR roles?

1.2 Research gap and rationale

Vacancy-level studies using broad U.S. and UK labour-market samples have found positive associations between AI-related requirements and advertised salaries (Alekseeva et al., 2021; Bone et al., 2025). Evidence from the UK energy sector further shows that estimated salary differences vary according to the type of digital capability required (Garcia-Lazaro et al., 2025). However, these studies cover multiple occupations or industries, making it impossible to determine whether the salary associations related to artificial intelligence exist within specific human resources domains.

A second limitation concerns how technical capability is defined. Existing classifications sometimes combine machine learning with data analysis, programming and broader digital skills, although these capabilities may support different forms of work and carry different labour-market value (Stephany & Teutloff, 2024). Such categories cannot reveal whether employers price AI-specific capability differently from advanced-data expertise. Within HR, it remains unclear whether AI-related and advanced-data requirements are associated with advertised salary after relevant job characteristics are taken into account, or whether those associations vary by seniority or HR functional area. These issues define the specific empirical gap that this study aims to address.

1.3 Research aim, questions and scope

The aim of this study is to examine whether candidate-facing AI-related requirements and advanced-data requirements are associated with advertised salaries in the sampled UK HR vacancies, and whether these associations vary across HR roles. The term AI/data skills is used in the research questions as an umbrella for two related but distinct constructs. In the empirical analysis, AI-related capability requirements are measured separately from Strict Advanced Data requirements, and both are included in the same models.

The study combines quantitative content analysis with cross-sectional regression analysis. It uses 2,229 full-time, salary-disclosing HR advertisements obtained through Reed.co.uk's official Jobseeker API in a snapshot collected on 14 July 2026. The analysis

examines advertised fixed or basic salary and accounts for observed vacancy characteristics. It estimates conditional associations within this Reed sample.

Accordingly, the study examines the following two research questions:

RQ1. *How are AI and data skill requirements associated with advertised salary in UK HR job advertisements?*

RQ2. *Do these associations differ by job seniority and HR functional area?*

By comparing two specific requirement measures within the same professional field, this study provides focused evidence on their associations with advertised salaries in HR vacancies. Exploratory analyses consider whether these average associations vary across seniority levels and HR functions.

1.4 Structure of the dissertation

Following this introduction, Chapter 2 reviews the literature on AI-related change in HR work and the pricing of technical capabilities before developing the hypotheses. Chapter 3 outlines the research design, data sources, sampling procedures, variable measurement, and analytical strategies. Chapter 4 presents the descriptive statistics, main regression results, heterogeneity analyses and robustness checks. Chapter 5 compares the research findings with existing studies, discusses their practical implications, limitations, and directions for future research. Chapter 6 concludes the main findings and outlines the study's contribution briefly.

Chapter 2: Literature review

2.1 AI and the reconfiguration of HR work

AI can change the tasks performed within HR roles, then influencing the skill requirements set by employers. Some routine and codable tasks can be automated. At the same time, AI can also support other activities while creating new responsibilities, such as interpreting outputs and overseeing automation. The initial unit of change is therefore the task, not the occupation. Acemoglu et al. (2022) provide labour-market evidence to support this argument. Using establishment-level vacancy data from the United States, they find that establishments whose existing tasks were more compatible with AI subsequently posted more AI-related vacancies. These establishments also stopped listing some previously requested skills and added new skill requirements. This evidence illustrates how changes in particular tasks can alter employer skill demand without eliminating the occupation itself. Hence, in HR, an existing role may remain while the skills required to perform it change.

Evidence from HR illustrates how AI affects different activities. Based on evidence from empirical studies on AI–HRM, Dima et al. (2024) categorized the impacts into five forms: task automation, more effective use of HR data, enhanced human capabilities, a reshaped work context, and changes to the social and relational dimensions of work. Vrontis et al. (2022) pointed out these changes in recruitment, training, decision-making, and job performance. Among these activities, intelligent technologies can perform selected tasks, provide information or feedback, and assist human judgement. However, existing knowledge is unevenly distributed across HR activities. Pan and Froese (2023) show that research has focused most heavily on recruitment, while other functions and organisational outcomes have received less systematic attention. Different tasks are reallocated in different activities, meaning that the resulting capability requirements depend on where and how AI is used.

The reallocation of HR tasks changes the capabilities practitioners are expected to bring to their roles. This shift is visible beyond HR. In their analysis of 18,947 AI-related vacancies on LinkedIn, Nakash and Peretz (2025) find that technical requirements, including Python and machine learning, frequently appear alongside communication and leadership skills. Although their study covers a range of occupations, it shows that AI-related roles are commonly associated with hybrid skill profiles. The AI–HRM literature

provides a more structured account of this combination. Drawing on the dynamic capabilities view, Deepa et al. (2024) organise the capabilities discussed in this literature into managerial cognition, human capital, and social capital. The associated competencies are diverse. This framework includes technical expertise, problem-solving, ethical decision-making, leadership, work redesign, and human–machine collaboration.

Overall, existing research explains how AI-induced task change can generate demand for technical, analytical, and human-centred capabilities. It still remains unclear whether these capabilities are associated with different levels of advertised pay. Section 2.2 therefore examines why particular skills may receive different levels of pay.

2.2 The sources of salary premiums for particular skills: tasks, scarcity and signalling

Section 2.1 established that AI can alter the capabilities required in HR work. The presence of a skill in a vacancy, however, does not establish that employers will reward it with higher pay. Skill-biased technological change (SBTC) offers a useful starting point. When technology raises employers' relative demand for skilled labour faster than the supply of that labour grows, the relative wages of skilled workers may increase (Acemoglu & Autor, 2011). This framework generally compares broad worker groups, such as college-educated and non-college workers. It therefore provides limited insight into how technological change affects demand for a particular capability listed in a vacancy.

The task-based approach provides a finer account. It treats tasks as activities that generate output and skills as the capabilities workers use to perform them (Acemoglu & Autor, 2011). Autor et al. (2003) show that computerisation can substitute for labour in routine tasks governed by explicit rules, while complementing labour in non-routine problem-solving and complex communication. Acemoglu and Restrepo (2019) make a related distinction between displacement, which occurs when technology takes over tasks previously performed by workers, and reinstatement, which occurs when new tasks expand the role of labour. If technology assumes responsibility for a task, demand for the skills used to perform it may decline. Where technology complements workers or creates new tasks, demand for the relevant capabilities may increase. These task-based mechanisms help explain shifts in skill demand, but they do not establish whether such demand leads to higher pay.

Whether higher demand becomes a salary premium also depends on relative supply. If many available workers already possess a capability, employers may meet their needs

without offering substantially higher pay. A premium becomes more likely when demand is high relative to the number of workers able to supply the capability. Skills are also rarely used in isolation. Using completed projects from a United States online freelance platform, Stephany and Teutloff (2024) find that skill premia are associated with relative demand and supply and with the number, diversity and value of complementary skills. This is skill-to-skill complementarity, rather than the relationship between technology and tasks described above. The same capability may therefore have different market value when combined with different skill portfolios. AI-specific capabilities, advanced data and analytics skills, and broad digital skills may support different tasks, differ in relative scarcity and form different combinations with other capabilities. They should therefore be examined separately rather than treated as a single category with a common premium. Their relative value remains an empirical question.

Employers must set recruitment requirements and pay before observing how an applicant will perform in the role. Spence (1973) describes hiring as a decision under uncertainty. Employers use observable attributes, including education and experience, as signals of likely productivity when determining wage offers. These signals are presented by the candidate. An advertised skill requirement indicates in advance which capability the employer expects applicants to demonstrate through such evidence. A vacancy advertisement records this requirement together with the remuneration offered for the role. The requirement shows which capability the employer seeks, while the salary shows how the vacancy is priced before hiring. Because salary reflects the vacancy as a whole, a simple comparison may also capture differences in seniority, HR function, location and other role characteristics. The analysis therefore focuses on the conditional association between a stated skill requirement and offered pay after these characteristics are taken into account. Advertised salary differs from realised earnings. It remains directly relevant to the present study because it records employer valuation at the point of recruitment. For the present study, the relevant signal is a candidate-facing AI-related requirement, which may concern practical capability or broader understanding, openness, willingness to learn, or readiness to engage with AI. References to AI tools used or supplied by the employer do not, by themselves, constitute such a requirement.

2.3 Evidence on AI and data skill premia from job postings

Section 2.2 explained why a skill may attract additional pay. The next empirical question is whether employers offer higher salaries when they require AI or data skills.

Online job advertisements make this relationship observable because when disclosed they record the capabilities, as well as the salary packages offered at the time of posting. Deming and Kahn (2018) demonstrate the value of vacancy-level information: skill requirements vary substantially within narrowly defined occupations. Since only 13 per cent of their advertisements disclose wages, they link average skill requirements to external occupation-by-location wage data rather than analyse skills and pay within the same advertisement. Salary-disclosing advertisements instead allow both to be observed within a vacancy, providing a basis for estimating advertised skill premia.

Studies comparing skill requirements and salaries within individual job advertisements generally find that AI skill requirements are associated with higher advertised salaries, although the estimated premia vary. In United States vacancies, Alekseeva et al. (2021) report an 11% premium after controlling for firm and occupation-by-local-labour-market differences. Adding standardised job-title fixed effects reduces the estimate to 5.1%. In the United Kingdom, Bone et al. (2025) estimate an AI skill premium of 23% after controlling for education, experience, year, region, industry and broad occupation. Garcia-Lazaro et al. (2025) further show that the association differs across digital capabilities. Within comparable occupations in the United Kingdom energy sector, AI and advanced digital skill requirements are associated with salaries around 9% higher, while intermediate digital skills correspond to a difference of 2.5%. These studies consistently identify positive salary associations, but not a single premium that applies across skills and labour-market settings.

One reason for these different estimates is the level at which vacancies are compared. Economy-wide studies combine salary differences between occupations with differences among vacancies within them. The reduction in Alekseeva et al.'s (2021) estimate after job-title controls shows that the premium changes when comparison is restricted to more similar roles. Deming and Kahn (2018) also find substantial variation in skill requirements within narrowly defined occupations, indicating that occupational averages can conceal relevant vacancy-level differences. Bone et al. (2025) provide a national benchmark by combining vacancies across the United Kingdom economy, whereas Garcia-Lazaro et al. (2025) provide a sector-level benchmark using energy-industry vacancies. These studies establish national and sectoral evidence, but leave the relationship within HR unanswered. Analysing HR vacancies separately can therefore test whether AI and data skill requirements are associated with additional advertised pay within one organisational function.

An HR-specific comparison also depends on how these skill requirements are defined. Existing studies often group AI with related digital technologies. Pan and Froese (2023) find that only 21% of the AI–HRM studies in their review provide a clear definition of AI; they also identify broad usages that include digitalisation, automation and the Internet of Things within the AI category. Similar conceptual overlap appears in labour-market skill classifications. Stephany and Teutloff’s (2024) AI skill category combines machine learning and deep learning with data analysis, cloud computing and general programming languages. A premium estimated from such a combined category cannot identify which capability is associated with higher pay. Classification thresholds also affect the measured relationship: Bone et al. (2025) find a larger AI premium when vacancies require multiple rather than single AI skills. Exact keyword matching creates a further problem because it can overlook the context in which a term is used and produce misclassification (Khaouja et al., 2021). This evidence supports distinguishing AI-specific capabilities, advanced data capabilities and broad digital capabilities, while using context to identify what the advertisement actually requires. Their relative premia remain an empirical question. The next section explains why HR is a theoretically informative setting in which to make this comparison.

2.4 Why HR is a theoretically informative setting for AI and data skill pricing

The evidence reviewed in Section 2.3 does not establish whether economy-wide AI and data skill premia also apply within HR. HR differs from many technology-intensive settings because these capabilities are used in employment decisions shaped by data limitations, interpersonal judgement, and institutional accountability. Such conditions may influence whether employers request the skills and how much they are willing to pay for them. This section first explains why the within-HR premium remains uncertain. It then considers how the premium may vary by seniority and HR function.

Employee performance can be difficult to define and measure, while events such as dismissals may be too rare to support reliable models. Decisions about hiring, promotion, and dismissal must also meet fairness, legal, and explainability requirements. Employee acceptance of these technologies may further affect organisational outcomes (Tambe et al., 2019). In addition, HR involves substantial interpersonal work. This is why human judgement remains important in many HR tasks involving AI. Deming (2017) shows that social interaction is difficult to automate because it relies on tacit abilities to interpret and

respond to others. Effective AI use in HR consequently requires technical competence to be combined with contextual judgement, communication, and accountability. These conditions may increase the value of employees who possess this combination of capabilities, but they may also limit the range of tasks for which AI is adopted. The expected salary premium is therefore uncertain.

The uneven adoption of digital technologies in HR adds to this uncertainty. In a survey of German firms, Chugunova and Danilov (2023) find that basic digital communication was widespread, while more than half of the organisations used neither data-driven nor AI-based HR tools. General digitalisation does not necessarily indicate the use of advanced analytical or AI capabilities. Marler and Boudreau (2017) identify analytical capability, managerial support, and HR information systems as recurring conditions for effective HR analytics. They also find that shortages of analytically skilled HR practitioners can impede adoption. Similarly, McCartney et al. (2021) show that HR analyst work combines technical knowledge and data analysis with consulting and communication skills, alongside knowledge of HR and business. Advanced technologies are therefore used selectively, while the roles associated with them require a combination of technical and human capabilities. These patterns may affect the frequency of AI and data skill requirements differently from the salaries attached to them. The two outcomes must therefore be examined separately.

After the frequency of a skill requirement is separated from its pricing, seniority becomes one possible source of variation in the premium. AI and data skills may complement different responsibilities at different organisational levels. In senior roles, their value may come from combining technical understanding with strategic judgement, implementation, and governance. Alekseeva et al. (2026) find that firms with greater AI adoption advertise more managerial vacancies and shift managerial requirements away from routine administration towards interpersonal and growth-oriented capabilities. Junior roles may follow a different mechanism. An AI requirement at this level could represent a new technical threshold, and AI-related productivity gains can also be substantial among less experienced workers. For example, Brynjolfsson et al. (2025) find that less experienced and lower-skilled customer-support workers achieve the largest productivity and learning gains. A senior-level premium may reflect a scarce combination of technical and strategic capabilities. In junior roles, it may arise from a higher technical threshold or a larger productivity gain. Both mechanisms are plausible, so the theory does not predict whether the advertised premium should be higher at senior or junior levels.

Variation may also arise across HR functions because their task structures and existing skill requirements differ. Chugunova and Danilov (2023) show that the use of advanced digital tools varies across HR processes. McCartney et al. (2021) similarly find that HR analyst roles require a different capability profile from more general HR roles. These differences can affect skill pricing in competing ways. In data-intensive functions, stronger task complementarity may increase the value of AI and data capabilities. However, when data skills already form part of the expected skill set, their additional value as a signal, and the associated premium, may be smaller. AI-specific capability may be less common in human-centred functions and could attract a higher premium when it complements contextual judgement. Rarity alone is insufficient. Employers may have little reason to pay more for a capability that has limited relevance to the function's core tasks. HR function may consequently affect both the frequency of AI and data skill requirements and the premium associated with them. These are distinct outcomes and should be examined separately. Taken together, the seniority and function mechanisms explain why AI and data skill premia may vary within HR. The next section integrates these arguments and presents the study's research questions and hypotheses.

2.5 Synthesis, research gap and research questions

Overall, previous research indicates that AI can reallocate tasks within HR and alter the capabilities expected of practitioners. Employers may offer additional pay for capabilities needed to perform particular tasks when demand for those capabilities is high relative to their supply. Studies comparing skill requirements and salaries within individual job advertisements generally report positive AI skill premiums, although the estimates vary. Evidence from national and sector-level analyses doesn't establish whether a similar association exists within HR. This gap is important because AI adoption in HR remains selective, while its use requires interpersonal judgement, governance, and accountability. The first aim of this study is therefore to examine whether AI and data skill requirements are associated with higher advertised salaries in UK HR vacancies.

This question also requires clearer construct boundaries. The term "AI and data skills" refers collectively to two related but analytically distinct requirement measures. Candidate-facing AI-related requirements may cover practical AI capability or broader understanding and readiness, whereas Strict Advanced Data requirements capture advanced data and analytics expertise. Existing classifications that combine AI, data analytics, programming, and broader digital skills cannot identify which requirement

measure is associated with any observed premium. The study therefore codes the two measures separately.

Any salary premium associated with AI and data skill requirements may vary across seniority levels. These capabilities can support different tasks and responsibilities in junior and senior roles. At senior levels, technical understanding may be valuable when joined with strategic judgement, implementation and governance. At junior levels, an AI/data requirement may instead mark a visible technical threshold or productivity increment. Since these mechanisms support opposing patterns, the literature does not justify a directional prediction. Seniority is therefore examined as an exploratory source of variation in the premium.

HR function may produce variation for similar reasons. Data-intensive functions may derive stronger task complementarity from AI/data capabilities, but the premium may be smaller when these skills are already part of the expected role profile. Human-centred functions may value a rarer combination of technical capability and contextual judgement, although rarity will not generate a premium where the capability has little relevance to core tasks. The balance among these mechanisms cannot be determined in advance. HR function is therefore examined as a second exploratory source of variation.

RQ1. *How are AI and data skill requirements associated with advertised salary in UK HR job advertisements?*

H1a. *Candidate-facing AI-related requirements are positively associated with fixed or basic advertised salary after controlling for observed job characteristics.*

H1b. *Strict Advanced Data requirements are positively associated with fixed or basic advertised salary after controlling for observed job characteristics.*

RQ2. *Do these associations differ by job seniority and HR functional area?*

H2a (exploratory). *The associations identified in RQ1 vary by job seniority.*

H2b (exploratory). *The associations identified in RQ1 vary by HR functional area.*

RQ1 provides the primary test of the salary associations. H2a and H2b are non-directional, exploratory tests under RQ2. Candidate-facing AI-related requirements and Strict Advanced Data requirements are treated as distinct constructs and entered jointly in the main models. Alternative coding boundaries, salary treatments, sample definitions and inference procedures are assessed through robustness checks. Descriptive profiles of AI-positive advertisements are treated as supplementary analyses.

Chapter 3: Methodology

3.1 Research design

To address the research questions, this study combined quantitative content analysis of online job advertisements with cross-sectional regression analysis to estimate conditional associations between the focal requirements and advertised salaries. The unit of analysis was a single HR vacancy advertisement retrieved under the study's sampling strategy.

Content analysis coded whether each advertisement included a candidate-facing AI-related requirement, covering either practical AI skills or explicit understanding, openness, willingness to learn, or readiness to engage with AI. It also coded whether each advertisement included a Strict Advanced Data requirement. For brevity, the first measure is referred to as the AI-related requirement throughout the remainder of the dissertation. The dependent variable was the natural logarithm of annualised and winsorised advertised fixed or basic salary. Regression models then assessed how the two focal requirements were associated with advertised salary after accounting for observed job characteristics.

3.2 Data source and sampling frame

This study collected search data via Reed.co.uk's official Jobseeker API on July 14, 2026. Reed was selected because it's a large general UK job board covering HR vacancies across functions, seniority levels and regions. For advertisements that disclosed pay, the API provided minimum and maximum salary fields, allowing advertised pay to be measured directly without relying on predicted salary estimates. The structured data and unique advertisement identifier, *jobId*, also enabled reliable within-platform deduplication.

Single-platform designs have precedent in UK vacancy research. Arthur (2021) analysed approximately 4.7 million advertisements obtained through the Reed API, while Ridgway et al. (2025) analysed 158 HR advertisements collected from People Management Jobs, the official job site of the Chartered Institute of Personnel and Development (CIPD). A multi-platform design was not used because job boards lack shared vacancy identifiers. Cross-platform duplicates, particularly agency reposts, could therefore not be identified reliably.

Because the API did not expose the Human Resources category displayed on Reed's website, the sampling frame was constructed using 12 separate free-text searches.

Together, these searches covered generalist HR roles, specialist functions, business partnering, and HR technology and analytics. Results were retrieved until each search was exhausted or the API's per-search 10,000-result pagination ceiling was reached; this ceiling affected two of the 12 searches. All retrievable records were then pooled and deduplicated using Reed's unique *jobId*. The complete search terms, query-level retrieval counts, title-matching rules, eligibility criteria, and intermediate sample counts are reported in Appendix A.

The main stages of sample construction are summarised below.

Table 3.1
Construction of the sampling frame

Stage	Processing step	N
Deduplicated search pool	Merged results from the twelve searches and removed duplicate <i>jobId</i> records	23,253
Records with disclosed salary fields	Retained advertisements with salary information; predicted salaries were not used	19,533
Initial HR title-based frame	Applied HR title inclusion and exclusion rules	2,940
Final eligible HR vacancy frame	Excluded fee-charging training programmes that were not vacancies, zero or non-GBP salary records, and ineligible or non-HR roles	2,539

Full advertisement descriptions were retrieved through the per-advertisement endpoint and archived with the search responses. All subsequent data preparation and analysis used these frozen records. The final frame therefore represents salary-disclosing Reed vacancies captured by the specified search strategy. Section 3.3 describes the additional salary and employment-basis restrictions used to form the primary analytical sample.

3.3 Data preparation and salary construction

The salary measure was the fixed or basic pay advertised for each vacancy (hereafter, fixed pay). On-target earnings (OTE), earning potential, and commission-inclusive amounts were not treated as fixed pay because they contained variable components that might not be realised. The API salary fields for the 2,539 advertisements in the final frame were therefore checked against the archived full descriptions.

Salary fields were corrected only when the full description provided unambiguous fixed-pay bounds or a clear pay period. Advertisements without a fixed component that could be separated from OTE, commission, earning potential or total earnings remained in

the vacancy frame but were excluded from salary analysis. These checks left 2,360 advertisements with recoverable fixed pay. The search terms, contextual rules, correction decisions and intermediate counts used in the salary audit are reported in Appendix B.

Employment basis was obtained from the archived detail records. The 131 part-time advertisements were excluded from the primary salary models because the records did not consistently distinguish actual part-time pay from full-time-equivalent salary. No full-time-equivalent conversion was applied. This produced a primary analytical sample of 2,229 full-time advertisements, which was used consistently in Models 1–3.

For salary ranges, the midpoint of the retained fixed-pay bounds was used; where only one value was reported, that value was retained. Annual salaries required no conversion. Hourly salaries were multiplied by 37.5 hours per week and 52 weeks per year, giving an annualisation factor of 1,950. Daily salaries were multiplied by five days per week and 52 weeks per year, giving a factor of 260.

After the salary audit, part-time exclusion and annualisation, fixed pay in the primary sample had a skewness of 3.151. To reduce the influence of extreme values, salaries were winsorised at the 1st and 99th percentiles. Values below £25,000 were set to £25,000, and values above £131,170 were set to £131,170. This reduced skewness to 2.097 without removing observations. The winsorised measure was then log-transformed, producing a skewness of 1.023. The dependent variable in all regression analyses was therefore the natural logarithm of winsorised annualised advertised fixed pay.

3.4 Measurement and validation of AI and advanced-data requirements

AI-related and advanced-data requirements were coded from the archived full advertisement descriptions. The text was cleaned before matching, with word-boundary and case-sensitive rules applied where required. Dictionary matches were used only to identify advertisements for full-text contextual review; a match did not automatically produce a positive classification. The complete dictionaries, preprocessing steps, matching rules and contextual coding criteria are reported in Appendix C.

The AI retrieval dictionary covered general AI concepts and named AI systems. Following a coverage-validation check, named systems, including ChatGPT, Claude, Gemini and Copilot, together with their spelling variants, were added to the dictionary. The revised dictionary was then reapplied to the frozen initial HR title-based frame.

Applying the dictionary to the 2,940 advertisements in the initial frame identified 72 unique candidates. Each candidate was assigned to one of four mutually exclusive

categories after full-text review. Hard indicated an explicit requirement for AI capability, experience, proficiency, use or implementation responsibility. Soft indicated an explicit expectation of openness, interest, willingness to learn or readiness to engage with AI. Tool-provided indicated that the employer provided or used AI tools without requiring AI capability or readiness from the candidate. None covered AI references unrelated to candidate requirements.

Hard and Soft were combined to construct the binary candidate-facing AI-related requirement variable, *ai_true*. Although Soft advertisements did not set a specific technical proficiency threshold, they still expressed an AI-related readiness requirement. Tool-provided and None were coded as non-AI. Full-text review classified the 72 candidates as 32 Hard, seven Soft, six Tool-provided and 27 None. After the eligibility restrictions described in Sections 3.2 and 3.3, the final frame contained 38 AI-positive advertisements, comprising 31 Hard and seven Soft. The primary analytical sample contained 34 AI-positive advertisements, comprising 29 Hard and five Soft.

The consistency of the contextual coding was assessed through a test–retest exercise. After five days, the same researcher independently recoded 20 AI candidate advertisements without consulting the original classifications. This assessed intra-rater rather than inter-rater reliability. Cohen’s κ was 1.000 for both the four-category and binary classifications, indicating perfect agreement within the recoded subset.

Advanced-data requirements were measured separately from AI requirements. The Strict Advanced Data retrieval dictionary covered HR and workforce analytics, data analysis and data science, data visualisation, statistical analysis, and named tools including Power BI, Tableau and SQL. An advertisement was classified as positive only when the matched term described a candidate skill or qualification, required experience, or a substantive analytical responsibility of the advertised role. Matches confined to search-engine-optimisation lists, descriptions of the employer’s business, the responsibilities of other employees, or software names without a candidate-facing analytical context were excluded.

Within the final frame, the Strict dictionary identified 99 candidates for contextual review. This produced 98 Strict Advanced Data-positive advertisements, of which 89 remained in the primary analytical sample. The Strict definition was used in the formal models, while broader combined AI/data definitions were reserved for the supplementary construct-boundary checks reported in Appendix D. Six advertisements in the primary

sample met both the AI and Strict Advanced Data definitions. The two requirements were entered as separate binary predictors in the joint models described in Section 3.6.

3.5 Job characteristics and control variables

The control variables were constructed from job titles and structured Reed fields using rule-based classifications. Table 3.2 summarises their construction and the reference categories used in the regression models. The complete title dictionaries, matching priorities, and London place-name rules are reported in Appendix E.

Table 3.2
Construction of job characteristics and control variables

Variable	Construction	Reference category
Seniority	Coded from jobTitle using an ordered title hierarchy as Entry, Mid, Senior, Director+ or Unclassified. One deliberately mixed-level title was retained as Unclassified.	Entry
HR function	Coded from jobTitle using specific combined-title decisions followed by ordered specialist rules: Analytics/HRIS, Employee Relations, Reward/Payroll, L&D and Recruitment/TA. Unmatched eligible titles formed the residual HRBP/Generalist category.	HRBP/Generalist
Greater London	Coded from exact audited London and Greater-London place labels in locationName. All other location labels were coded as non-London.	Non-London
Agency-recruiter role	Title-based indicator for recruitment consultant, recruiter, resourcer and talent-acquisition-consultant roles. Explicit internal or in-house recruiter titles were coded as non-agency roles.	No
Advertisement posted by a recruitment agency	Rule-expanded indicator combining recruitment-intermediary employer-name matches with high-precision agency boilerplate in the full description; the employer-level expansion was then applied consistently across that employerName group.	No
Salary basis	Recorded whether pay was originally quoted per annum, per hour or per day. The original basis remained as a control after hourly and daily pay had been annualised.	Per annum
Posting period	Derived from datePosted as Earlier (before 1 June 2026), June 2026 (1–30 June) or July 2026 (1–14 July, the collection date).	June 2026

Note. Rules reproduce the variables in the final primary sample ($N = 2,229$). Seniority and HR function entered Models 2 and 3; the remaining variables entered Model 3. Full dictionaries, matching priorities and category counts are reported in Appendix E.

Seniority and HR function were also used as the dimensions of heterogeneity examined under RQ2. The interaction specifications and the grouped HR-function measure are described in Section 3.6.

3.6 Analytical strategy

Main models and statistical inference

The analyses used the primary sample of 2,229 full-time advertisements. The dependent variable was the natural logarithm of winsorised annualised advertised fixed pay. The two focal predictors, *ai_true* and *data_skill_strict*, were entered jointly in every model.

Three nested ordinary least squares (OLS) models were estimated. Model 1 contained the two focal predictors. Model 2 added seniority and the six-category HR function measure. Model 3 added London, agency recruiter, agency posting, salary basis, and posting period. Model 3 was the primary specification for H1a and H1b. The nested sequence showed how the focal estimates changed as observed job characteristics were added.

Because the two requirements were entered together, each coefficient represented its association with advertised salary after accounting for the other requirement and the controls included in that model. The estimates therefore represented conditional associations rather than causal effects.

HC1 heteroscedasticity-robust standard errors were used for Models 1–3 and the primary H2 analyses. A Wald contrast based on the HC1 covariance matrix was used in Model 3 to test whether the coefficients for the two focal requirements differed. All tests were two-sided. Statistical significance was assessed at $p < .05$, and 95% confidence intervals were reported. Coefficients and confidence intervals for the focal predictors were converted from log points into estimated percentage salary differences using $100[\exp(b) - 1]$.

Exploratory heterogeneity analyses

H2a and H2b were examined as exploratory extensions of Model 3. For H2a, $AI \times$ seniority and Strict Advanced Data $\times$ seniority interaction terms were added. One advertisement coded as Unclassified was excluded because it could not be assigned to a single seniority category. This produced an analytical sample of $N = 2,228$.

For H2b, each focal requirement was interacted with a grouped HR function measure. The original six categories were combined into three broader groups because some contained very few AI-positive advertisements. The resulting groups were Core/Generalist and Recruitment (HRBP/Generalist and Recruitment/TA), Data-Intensive (Reward/Payroll and Analytics/HRIS), and Human-Centred (L&D and Employee Relations). This grouped measure replaced the original HR function variable, while seniority and the remaining Model 3 controls were retained.

Each heterogeneity analysis used 3 joint Wald tests. These assessed the AI interaction terms, the Strict Advanced Data interaction terms, and both sets of interactions together. Adjusted estimates and 95% confidence intervals were calculated for each seniority or function group. Heterogeneity was evaluated using the joint tests, with individual point estimates interpreted descriptively.

Robustness checks and diagnostics

Both focal requirements and the applicable Model 3 controls were retained in every robustness specification. For the inference check, Model 3 and the H2a and H2b interaction models were re-estimated with standard errors clustered across 536 employer-name groups derived from the *employerName* field. Outcome sensitivity was assessed in three ways: without winsorisation, with winsorisation at the 2.5th and 97.5th percentiles, and by trimming observations outside the 1st and 99th percentiles.

Sample checks used the preliminary post-OTE salary sample, its full-time subset, and the final fixed-pay sample with part-time advertisements retained. Further specifications retained one record from each group of advertisements with identical standardised descriptions, restricted the sample to salaries originally reported per annum, or excluded recruitment-agency roles. Coding checks recoded the Soft-AI and Copilot boundary cases as non-AI, applied a hard-only AI definition, and restored one Strict Advanced Data boundary case as positive. Leave-one-positive-out checks removed each of the 34 AI-positive advertisements and each of the 89 Strict Advanced Data-positive advertisements in turn. Winsorisation thresholds were recalculated whenever sample composition changed.

A Pearson correlation matrix and variance inflation factors were used to assess multicollinearity. Full specifications, sample sizes, and results are reported in Appendix F. Analyses were conducted in Python 3.13.9 using *statsmodels* 0.14.5.

3.7 Ethical considerations

This study collected vacancy data through Reed's official Jobseeker API and followed its access and rate-limit requirements. Each job advertisement served as the unit of analysis. No personal data relating to jobseekers or applicants were collected. Recruiter names or contact details occasionally appeared in the advertisements, but they were not coded as research variables. The study reported only aggregate statistical results.

The frozen data were stored securely on a local device, with API credentials kept separately from the code and reproduction materials. Any identifiable advertisement excerpts quoted in the dissertation were appropriately de-identified. Public sharing of the complete raw advertisement data would be subject to UCL requirements and Reed's platform terms.

Chapter 4: Results

4.1 Descriptive statistics and sample characteristics

The primary analysis included 2,229 full-time HR vacancy advertisements. Table 4.1 reports annualised fixed or basic salary before winsorisation, after winsorisation, and after the natural-log transformation used in the regression models.

Table 4.1

Continuous-variable descriptive statistics

Salary measure	Unit	N	Mean	SD	Minimum	Maximum	Skewness
Annualised advertised fixed/basic salary before winsorisation	GBP	2,229	£44,062.22	£21,718.06	£14,560	£234,000	3.151
Annualised advertised fixed/basic salary winsorised at 1st and 99th percentiles	GBP	2,229	£43,695.41	£19,411.07	£25,000	£131,170	2.097
Annualised advertised fixed/basic salary, winsorised and natural-log transformed	log GBP	2,229	10.611	0.361	10.127	11.784	1.023

Note. $N = 2,229$. Salary figures refer to annualised advertised fixed or basic salaries. The regression outcome was salary winsorised at the 1st and 99th percentiles and then transformed using the natural logarithm. *SD* = standard deviation.

Annualised salary was strongly right-skewed before winsorisation (skewness = 3.151). Winsorisation reduced skewness to 2.097, and the natural-log transformation reduced it further to 1.023. The models therefore used the log-transformed measure; Appendix Figure F1 shows the corresponding distributions.

Table 4.2

Frequencies and percentages for categorical variables in the primary analytical sample

Variable	Category	n	%
AI-related requirement	No	2,195	98.5
AI-related requirement	Yes	34	1.5
Strict Advanced Data requirement	No	2,140	96.0
Strict Advanced Data requirement	Yes	89	4.0
Seniority	Entry	742	33.3
Seniority	Mid	717	32.2

Variable	Category	n	%
Seniority	Senior	671	30.1
Seniority	Director+	98	4.4
Seniority	Unclassified	1	<0.1
HR function	HRBP/Generalist	967	43.4
HR function	Recruitment/TA	547	24.5
HR function	Reward/Payroll	540	24.2
HR function	L&D	70	3.1
HR function	Employee Relations	64	2.9
HR function	Analytics/HRIS	41	1.8
Greater London location	No	1,728	77.5
Greater London location	Yes	501	22.5
Recruitment-agency role	No	1,894	85.0
Recruitment-agency role	Yes	335	15.0
Advertisement posted by recruitment agency	No	920	41.3
Advertisement posted by recruitment agency	Yes	1,309	58.7
Salary basis	Per annum	1,845	82.8
Salary basis	Per hour	285	12.8
Salary basis	Per day	99	4.4
Posting period	June 2026	1,227	55.0
Posting period	July 2026	922	41.4
Posting period	Earlier	80	3.6

AI-related requirements occurred in 34 advertisements (1.5%) and Strict Advanced Data requirements in 89 (4.0%); six advertisements met both definitions. Entry, Mid and Senior roles each accounted for approximately one third of the sample, and HRBP/Generalist was the largest functional group. Appendix Figure F2 profiles the 38 AI-positive advertisements in the final HR frame.

4.2 Main joint OLS results

Table 4.3 presents three nested OLS models for the 2,229-advertisement sample. Model 1 included the two focal requirements, Model 2 added seniority and HR function, and Model 3 added the remaining controls and served as the primary specification.

Table 4.3

Joint OLS models predicting the natural logarithm of winsorised annualised advertised fixed or basic salary

Variable	Model 1	Model 2	Model 3
AI-related requirement	0.313*** (0.072)	0.184*** (0.051)	0.151*** (0.046)
Strict Advanced Data requirement	0.148*** (0.042)	0.025 (0.031)	0.020 (0.025)
Seniority: Mid	—	0.249*** (0.012)	0.241*** (0.010)
Seniority: Senior	—	0.499*** (0.014)	0.462*** (0.012)
Seniority: Director+	—	0.863*** (0.030)	0.786*** (0.029)
Seniority: Unclassified	—	-0.036*** (0.010)	0.051*** (0.011)
HR function: Analytics/HRIS	—	0.152* (0.063)	0.104** (0.039)
HR function: Employee Relations	—	0.105** (0.033)	0.083** (0.030)
HR function: L&D	—	-0.035 (0.032)	-0.030 (0.028)
HR function: Recruitment/TA	—	-0.148*** (0.013)	-0.089*** (0.017)
HR function: Reward/Payroll	—	0.007 (0.015)	-0.025* (0.012)
Greater London location	—	—	0.180*** (0.012)
Recruitment-agency role	—	—	-0.104*** (0.018)
Advertisement posted by recruitment agency	—	—	0.004 (0.009)
Salary basis: per day	—	—	0.575*** (0.037)
Salary basis: per hour	—	—	0.022 (0.013)
Posting period: July 2026	—	—	0.006 (0.009)
Posting period: Earlier	—	—	0.012 (0.021)
Constant	10.601*** (0.008)	10.370*** (0.009)	10.324*** (0.012)
N	2229	2229	2229
R ²	0.019	0.528	0.701
Adjusted R ²	0.018	0.526	0.698
HC1 robust F-statistic	17.69	695.74	627.77


Note. Entries are unstandardised OLS coefficients with HC1 heteroscedasticity-robust standard errors in parentheses. Reference categories are Entry, HRBP/Generalist, non-London, not an agency-recruiter role, not posted by a recruitment agency, per annum salary and June 2026. Dashes indicate exclusion from a model. The Unclassified category contains one advertisement. All overall robust F-tests had $p < .001$. * $p < .05$, ** $p < .01$, *** $p < .001$.

In Model 3, the candidate-facing AI-related requirement coefficient was positive and significant ($b = 0.151$, HC1 $SE = 0.046$, $p < .001$), corresponding to an estimated advertised salary difference of 16.34% (95% CI [6.36%, 27.25%]). The Strict Advanced Data coefficient was $b = 0.020$ (HC1 $SE = 0.025$, $p = .443$), corresponding to an estimated difference of 1.97% (95% CI [-2.99%, 7.19%]); this adjusted association was not statistically significant. A Wald contrast indicated that the two adjusted coefficients differed significantly, $b_{\text{diff}} = 0.132$, $SE = 0.053$, 95% CI [0.028, 0.235], Wald $\chi^2(1) = 6.21$, $p = .013$.

Across Models 1–3, the AI estimate decreased from 36.79% to 20.22% and 16.34% but remained significant. The Strict Advanced Data estimate decreased from 15.90% to 2.54% after seniority and HR function were introduced and remained close to zero in Model 3. Figure 4.1 shows these paths.

Figure 4.1

Estimated salary differences associated with AI-related capability and Strict Advanced Data requirements across joint Models 1–3



Note. Points are percentage salary differences; error bars are 95% confidence intervals based on HC1 standard errors. Both requirements are included jointly in every model.

Adjusted R^2 increased from .018 in Model 1 to .526 in Model 2 and .698 in Model 3. The fully adjusted results supported H1a but not H1b. Variance inflation factors ranged from 1.004 to 2.465, providing no indication of problematic multicollinearity. The full

correlation matrix and VIF results are reported in Appendix Tables F1 and F2, while Appendix Figure F3 visualises the correlations.

4.3 Heterogeneity by seniority

H2a examined whether the adjusted associations varied across seniority categories. Table 4.4 reports the formal interaction tests and seniority-specific estimates.

Table 4.4

Seniority heterogeneity in the adjusted salary differences associated with AI-related requirements and Strict Advanced Data requirements

Panel A. Joint interaction tests						
Joint test		Wald χ^2		df	p	
AI-related requirement $\times$ seniority interactions		2.910		3	.406	
Strict Advanced Data $\times$ seniority interactions		0.945		3	.814	
All AI-related requirement and Strict Advanced Data requirement interactions with seniority		4.150		6	.656	
Panel B. Adjusted estimates by seniority						
Requirement	Seniority	Group N	Positive n	Estimated difference	95% CI	p
AI-related requirement	Entry	742	5	20.6%	[6.7%, 36.3%]	.003
	Mid	717	11	27.7%	[5.1%, 55.2%]	.014
	Senior	671	11	7.1%	[−6.3%, 22.5%]	.315
	Director+	98	7	11.1%	[−6.1%, 31.6%]	.221
Strict Advanced Data requirement	Entry	742	18	−1.4%	[−9.2%, 7.1%]	.737
	Mid	717	35	3.1%	[−4.4%, 11.3%]	.425
	Senior	671	29	1.4%	[−7.7%, 11.3%]	.777
	Director+	98	7	6.9%	[−12.0%, 29.8%]	.500

Note. Panel A reports joint Wald tests; Panel B reports adjusted percentage differences. The H2a sample excludes one Unclassified advertisement ($N = 2,228$). Positive n is the number meeting each requirement. H2a is assessed using the separate requirement-specific interaction tests; the combined row is a supplementary omnibus test.

Neither the AI–seniority interactions (Wald $\chi^2(3) = 2.910$, $p = .406$) nor the Strict Advanced Data–seniority interactions (Wald $\chi^2(3) = 0.945$, $p = .814$) were jointly significant. The supplementary combined test was also non-significant ($p = .656$).

Re-estimation with standard errors clustered across 536 employer-name groups produced the same conclusion. Neither the AI–seniority interactions, Wald $\chi^2(3) = 3.483$, $p = .323$, nor the Strict Advanced Data–seniority interactions, Wald $\chi^2(3) = 1.084$, $p = .781$, were jointly significant. The full employer-clustered results are reported in Panel A of Appendix Table F3.

Descriptive results show that AI estimates were higher for Entry (20.6%) and Mid (27.7%) than for Senior (7.1%) and Director+ (11.1%) advertisements. The significant Entry and Mid estimates did not demonstrate between-group differences because the interaction test was non-significant. Strict Advanced Data estimates were small and non-significant in every category. H2a was not supported; Appendix Figure F4 presents the estimates graphically.

4.4 Heterogeneity by HR function

H2b examined whether the adjusted associations varied across the three HR function groups. Table 4.5 reports the interaction tests and group-specific estimates.

Table 4.5

Salary differences associated with AI-related requirements and Strict Advanced Data requirements by HR function group

Panel A. Joint interaction tests						
Joint test		Wald χ^2		df	p	
AI-related requirement $\times$ HR-function-group interactions		1.776		2	.411	
Strict Advanced Data requirement $\times$ HR-function-group interactions		0.541		2	.763	
AI-related requirement and Strict Advanced Data requirement interactions (combined)		2.478		4	.649	
Panel B. Adjusted estimates by HR function group						
Requirement	HR function group	Group N	Positive n	Estimated difference	95% CI	p
AI-related requirement	Core/generalist & recruitment	1514	28	12.4%	[1.8%, 24.2%]	.021
	Data-intensive	581	3	26.6%	[1.1%, 58.5%]	.040
	Human-centred	134	3	34.5%	[−2.3%, 85.2%]	.069
Strict Advanced Data requirement	Core/generalist & recruitment	1514	47	3.7%	[−2.8%, 10.6%]	.271
	Data-intensive	581	39	7.6%	[−0.3%, 16.2%]	.058
	Human-centred	134	3	4.3%	[−19.8%, 35.6%]	.753

Note. Panel A reports joint Wald tests; Panel B reports adjusted percentage differences. Function groups are Core/generalist & recruitment, Data-intensive and Human-centred. $N = 2,229$. Positive n is the number meeting each requirement. H2b is assessed using the separate requirement-specific interaction tests; the combined row is supplementary.

Neither the AI–function-group interactions (Wald $\chi^2(2) = 1.776$, $p = .411$) nor the Strict Advanced Data–function-group interactions (Wald $\chi^2(2) = 0.541$, $p = .763$) were jointly significant. The supplementary combined test was also non-significant ($p = .649$).

Employer-clustered inference likewise produced the same conclusion. Neither the AI–function-group interactions, Wald $\chi^2(2) = 1.619$, $p = .445$, nor the Strict Advanced

Data–function–group interactions, Wald $\chi^2(2) = 0.691$, $p = .708$, were jointly significant. The full employer-clustered results are reported in Panel B of Appendix Table F3.

AI estimates were 12.4%, 26.6% and 34.5% for Core/generalist & recruitment, Data-intensive and Human-centred advertisements, respectively. Individually significant estimates in the first two groups did not establish between-group differences. Strict Advanced Data estimates were small and non-significant across all groups. H2b was not supported; Appendix Figure F5 presents the estimates graphically.

4.5 Robustness checks

Robustness checks varied inference, salary treatment, sample definition, duplicate handling and coding boundaries. Full Joint Model 3 robustness results are reported in Appendix Table F4; the point-estimate checks in Panels A–C are visualised in Appendix Figure F6.

The AI estimate remained positive and significant throughout. Alternative salary treatments produced estimates of 15.3%–16.9%, while alternative samples and coding checks produced 14.1%–17.7% (all $p \leq .004$). Sequential deletion of each AI-positive advertisement produced 14.2%–18.1% (all $p < .003$).

Strict Advanced Data remained non-significant across the alternative specifications: estimates ranged from –0.1% to 4.0% in sample and coding checks and from 1.3% to 2.6% in leave-one-positive-out models. No robustness check changed the substantive conclusions.

4.6 Hypothesis and research-question summary

Table 4.6 summarises the four hypothesis decisions.

Table 4.6
Summary of hypothesis tests

Hypothesis	Test basis	Key result	Conclusion
H1a. Positive association for AI-related requirements	Joint Model 3	16.34%, 95% CI [6.36%, 27.25%], $p < .001$	Supported
H1b. Positive association for Strict Advanced Data requirements	Joint Model 3	1.97%, 95% CI [–2.99%, 7.19%], $p = .443$	Not supported
H2a. Variation by seniority	Requirement-specific joint interaction tests	AI-related requirement: Wald $\chi^2(3) = 2.910$, $p = .406$; Strict Advanced Data requirement: Wald $\chi^2(3) = 0.945$, $p = .814$	Not supported
H2b. Variation by HR function group	Requirement-specific joint interaction tests	AI-related requirement: Wald $\chi^2(2) = 1.776$, $p = .411$; Strict Advanced Data requirement: Wald $\chi^2(2) = 0.541$, $p = .763$	Not supported

Note. ‘Not supported’ indicates insufficient statistical evidence; it does not establish that an association or group difference is exactly zero.

For RQ1, AI-related requirements were associated with significantly higher adjusted advertised salaries, whereas Strict Advanced Data requirements were not. For RQ2, the interaction analyses provided insufficient evidence that either association varied systematically by seniority or HR function. Chapter 5 discusses the implications of these findings.

Chapter 5: Discussion

5.1 Overview of the key findings

The central finding was that the two requirement measures had different advertised salary associations in UK HR vacancies. Advertisements with an AI-related requirement carried an adjusted advertised salary premium of 16.34%, which remained positive and statistically significant across the robustness checks. Strict Advanced Data requirements, by contrast, carried no statistically significant premium, with an adjusted estimate of 1.97%. H1a was therefore supported, whereas H1b was not. There was also no evidence of a clear gradient across HR roles. The interaction tests did not show that either premium varied systematically by seniority or HR function, so H2a and H2b were not supported. Although the group-specific AI estimates differed descriptively, these differences did not establish between-group variation. Overall, the findings show a clear difference in the pricing of the requirement measures. No systematic variation was identified across the seniority or HR function groups examined. The following sections explain this selective pricing pattern and its implications.

5.2 Selective Pricing of AI-Related Requirements

The findings show selective salary associations across the two requirement measures examined. In Model 1, AI-related requirements were associated with a salary difference of 36.79%, while Strict Advanced Data requirements were associated with a difference of 15.90%. Adding seniority and HR function reduced these estimates to 20.22% and 2.54% respectively. In the fully adjusted model, the AI estimate remained statistically significant at 16.34%, whereas the Strict Advanced Data estimate was small and statistically non-significant at 1.97%.

The sharp reduction in the Strict Advanced Data coefficient after seniority and HR function were introduced suggests that much of its initial association reflected the types of HR roles in which these requirements appeared. The AI estimate also declined, indicating that observed job characteristics accounted for part of its initial association. Nevertheless, a substantial positive association remained after these characteristics were taken into account.

The persistence of the AI premium is consistent with positive vacancy-level estimates reported across broader US and UK labour-market samples (Alekseeva et al., 2021; Bone et al., 2025). Garcia-Lazaro et al. (2025) found that different types of digital skills were associated with different salary premia within the same occupations. The present study identifies a similar distinction within HR vacancies: AI-related requirements retained a significant salary association, whereas Strict Advanced Data requirements did not. Because both requirements were included in the same model, the AI estimate was obtained after accounting for the overlap with Strict Advanced Data requirements. The resulting premium therefore represents a salary association specific to AI-related requirements. The AI premium may reflect the combination of technical knowledge, professional judgement and governance responsibility required to use AI in HR. Based on interviews with 42 recruitment experts, Cruz (2024) found that recruiters used professional judgement to interpret, question and sometimes work around AI recommendations. Their responses were also shaped by organisational oversight and trust in the technology. Soleimani et al. (2025) reached a related conclusion from interviews with 39 HR professionals and AI developers. Reducing bias in AI recruitment required iterative collaboration over model requirements, data management and model retraining. HR professionals therefore needed technical and data literacy alongside an understanding of organisational context and people-related concerns. Together, these studies suggest that AI requirements signal a broader combination of technical, interpretive and governance responsibilities, which may help explain the observed salary premium.

The weak independent salary association for Strict Advanced Data requirements may reflect how advanced data work contributes to the wider HR analytics process. Drawing on interviews with 17 HR analytics experts, Wirges and Neyer (2023) found that analytics specialists, HR business partners and other departments needed to work together. Technical analysis had to address a business problem and be translated into an HR decision. Gerber et al. (2024), using a survey of 133 Swiss companies and 12 expert interviews, similarly found that the benefits of HR analytics depended on converting analytical findings into organisational action and business-relevant outcomes. Together, these studies

show that advanced data capability creates value within a wider process involving problem definition, cross-functional coordination and action. Responsibility for this process is likely to vary across seniority levels and HR functions. Strict Advanced Data requirements may therefore be concentrated in higher-paid levels and functions. Once these job characteristics were included, the estimated salary difference fell from 15.90% to 2.54%.

Together, AI-related requirements may involve a broader set of responsibilities, including interpreting AI outputs, overseeing the use of AI and integrating it into HR decisions. Strict Advanced Data requirements may cover only one technical stage of the analytical process, with subsequent interpretation and decision-making assigned to other roles. This difference in responsibility may help explain the different salary associations of the two requirements. Section 5.3 examines whether these associations varied systematically across seniority levels and HR functions.

5.3 No clear gradient across HR roles

The interaction tests found no systematic variation in either salary association by seniority or HR function. The joint tests by seniority were nonsignificant for AI-related requirements ($p = .406$) and Strict Advanced Data requirements ($p = .814$). The corresponding tests by HR function were also non-significant ($p = .411$ and $p = .763$, respectively). H2a and H2b were therefore not supported. Formal comparisons did not distinguish the groups, despite variation in their point estimates.

The AI point estimates by seniority peaked at Mid, fell sharply at Senior and rose again at Director+, forming no ordered hierarchy. The Entry and Mid estimates were individually significant, but the interaction test did not establish that they exceeded the estimates for the other levels. Across HR functions, the Data-intensive and Human-centred estimates were higher than the Core group estimate. Each of these two groups contained only three AI-positive advertisements, leaving the apparent differences imprecise. Strict Advanced Data estimates remained small and non-significant across both sets of groups.

A possible explanation for this irregular pattern is that seniority and HR function do not fully capture differences in job content. Atalay et al. (2024) used vacancy text to show that tasks and technologies vary substantially within occupations and contribute to wage differences within them. HR-specific evidence supports this explanation across both dimensions examined here. Espegren (2025) found that HR analytics work crossed organisational levels and involved HR, IT, finance and line management; none of the three public administrations studied had a formal operational analytics team. Kashive and

Khanna (2023) found that technical, business and communication capabilities were required across entry, middle–senior and top-level HR analytics roles. Together, these studies provide a task-based interpretation of the irregular subgroup estimates.

Another possibility is that the binary AI measure obscured differences in the scope of AI responsibility. Every advertisement with an AI-related requirement was coded as AI-positive, giving the same binary value to a role requiring basic tool use and one involving AI implementation or governance. Jobs within the same seniority or function group could therefore carry very different levels of AI responsibility. These unmeasured differences, together with the three AI-positive advertisements in each of the Data-intensive and Human-centred groups, may have made systematic group differences difficult to detect. The results direct attention to the specific AI tasks and responsibilities described in each advertisement. Section 5.4 considers how these differences may be reflected in advertised salaries.

5.4 Interpreting the AI Premium: Adoption Context and Employer Composition

What, then, does the 16.34% AI premium capture? The estimate remained after observed vacancy characteristics were controlled, yet it followed no stable seniority or function gradient. Seniority and HR function provide only part of the explanation. Two possible explanations are considered: the wider context of AI adoption and the types of employers posting AI-related requirements.

The first explanation concerns the diffusion of AI use and related hiring requirements. Candidate-facing AI-related requirements were identified in 34 of the 2,229 advertisements in the primary analytical sample (1.5%), indicating that such requirements were uncommon in the sampled vacancies. McElheran et al. (2024) found that early AI use was rare and concentrated among a limited set of U.S. firms. More recent UK evidence shows that AI adoption among businesses has expanded but remains uneven and often limited in depth (ONS, 2026). Under such conditions, potential applications and expected returns may remain uncertain (Ameye et al., 2023). Real options theory suggests that early investments in technology and workforce capability can preserve opportunities to expand or redirect future applications (Bhattacharya & Wright, 2005; Fichman, 2004). Employers may therefore be willing to offer higher salaries to recruit candidates who can help develop the AI capacity needed for future applications. The value of this capacity may extend

across seniority levels and HR functions, which could help explain the absence of a stable gradient.

The second explanation concerns employer composition. Employers posting AI-related requirements may already be more innovative, productive, or higher-paying. Babina et al. (2024) associate firm-level AI investment with growth and product innovation. Cirera and Martins-Neto (2023) similarly found that innovative firms paid higher wages before beginning innovation and that wages rose after innovation started. Because the present models do not account for employer-specific wage levels, part of the AI premium may reflect these employers' general pay policies. However, innovation-related wage gains may differ between existing employees and new recruits. Kline et al. (2019) found that the approval of valuable patents increased the earnings of existing employees but had no detectable effect on the starting wages of new employees. Because the present study measures salaries advertised to prospective employees, the same incumbent rent-sharing mechanism may not fully explain the observed AI premium. Employer composition may nevertheless matter if AI-related requirements are disproportionately posted by firms that already maintain higher pay levels.

The broader implication is that vacancy-level AI premiums may change as technologies mature and the range of employers recruiting for them expands. Kalyani et al. (2025) found that technology-related hiring spread across a wider range of locations and skill levels as technologies matured. The 16.34% estimate offers a point of reference for understanding the advertised salary association of AI-related requirements in UK HR vacancies. Employer-linked longitudinal vacancy data could show whether the association persists, narrows, or moves across roles over time, while separating changes in the salary association from changes in employer composition.

5.5 Implications for HR practice

AI requirements may refer to very different types of work, and the analysis did not identify a stable salary gradient across seniority levels or HR function groups. This makes it important to define AI responsibilities at the individual job level. Hiring managers and recruiters should state whether the post-holder will use or evaluate AI outputs, implement or validate a system, or make and oversee AI-supported employment decisions. Advertisements should also specify expected outputs, frequency of use, decision authority and accountability for errors. These details can guide the choice between internal capability development and specialist recruitment. Selection methods can then match the

work, using practical tool exercises for users, implementation cases for specialists, and judgement or governance scenarios for decision owners.

AI-supported HR work can integrate technical expertise, professional judgement, and governance capabilities, with responsibilities shared across multiple functions. Workforce planning and learning and development teams should therefore create separate pathways for users, implementers and decision owners. Users need AI literacy and practice in questioning outputs; implementers need skills in data quality, validation and system integration; decision owners need training in fairness, human review and escalation. For each AI-enabled process, organisations should name the owner at each stage and document handovers among HR, data or IT, and legal or compliance teams (Cruz, 2024; Soleimani et al., 2025; van den Broek et al., 2021).

The Strict Advanced Data result draws attention to the process through which technical analysis is converted into organisational value. Analysis must be connected to a defined business problem, communicated to decision-makers and followed by action. People analytics teams, HR business partners and line managers should identify who defines the question, approves the data, conducts and checks the analysis, communicates the findings, makes the decision and reviews the outcome. Analyst job descriptions should specify the decisions and stakeholders supported and connect each listed tool to its intended use. Analysts should be assessed on problem framing and communication alongside technical competence; HR business partners and line managers should be able to interrogate the evidence and act on it (Gerber et al., 2024; Wirges & Neyer, 2023)

The estimated 16.34% AI premium provides an aggregate market signal. The appropriate pay adjustment for an individual role depends on the work involved, and the market benchmark may change as AI adoption spreads. Reward specialists should therefore assess technical depth, implementation scope, decision consequences and governance accountability when setting pay. External comparisons should be based on roles with similar tasks and responsibilities in organisations at a comparable stage of AI adoption. Employers can review these benchmarks using time-to-fill, offer-acceptance rates, accepted salaries and retention data to determine whether recruitment conditions continue to support a premium.

5.6 Limitations and future research

Firstly, the cross-sectional design cannot determine whether AI requirements are associated with higher advertised salaries because the work is valued more highly or because employers requesting AI already pay more. Employer-clustered standard errors addressed dependence among advertisements from the same named employer but did not control for differences in employer wage levels. Information on industry, firm size, productivity, digital maturity and detailed location was also unavailable. Future research could use employer-linked longitudinal data to compare salaries before and after AI requirements appear within the same organisation.

The second limitation concerns statistical power for RQ2. The primary sample contained only 34 AI-positive advertisements, while the Data-intensive and Human-centred groups contained three each. Combining six HR functions into three made estimation possible but could not produce precise estimates from these small cells. The interaction tests therefore show that no systematic gradient was detected; they do not demonstrate that the salary associations were equal across groups. Future studies should increase the number of AI-positive advertisements within each subgroup and define minimum cell sizes and power requirements before testing heterogeneity.

Thirdly, the binary measures captured whether a capability requirement appeared but not its level or scope. The AI measure combined Hard and Soft requirements without distinguishing tool use, implementation, evaluation or governance. The Strict Advanced Data measure similarly did not capture the intensity of the analytical work required. The test–retest exercise assessed the same researcher’s consistency only among dictionary-matched advertisements; it did not assess inter-rater agreement or requirements missed by the dictionaries. Future research should independently double-code samples of matched and unmatched advertisements and construct graded measures of task intensity and responsibility. Employer interviews could then assess whether these measures correspond to the work expected in practice.

Finally, the sample was a single-date stock of vacancies that were live on Reed, rather than advertisements newly posted on that date. This improved comparability and avoided repeatedly observing the same live advertisements. Collecting several adjacent daily snapshots would capture many of the same records, but a single snapshot cannot represent annual patterns and may overrepresent advertisements that remain open for longer. The sample also excluded vacancies on other platforms, those without recoverable

fixed or basic salaries, and part-time or variable-pay roles. Two searches reached the API limit of 10,000 records, which may have further restricted coverage. Future research should collect data over a longer period, distinguish new, continuing and reposted vacancies, and apply cross-platform duplicate detection. Linking advertisements to accepted offers, payroll records and organisational adoption data would also show whether the observed associations extend to realised pay and capability use.

Chapter 6: Conclusion

Using 2,229 full-time UK HR vacancy advertisements, this study examined whether AI-related requirements and Strict Advanced Data requirements were associated with advertised fixed or basic salaries, and whether these associations varied by seniority or HR function. In the fully adjusted model, AI-related requirements carried a 16.34% salary premium, and the result remained positive across the robustness checks. The 1.97% estimate for Strict Advanced Data was statistically nonsignificant. Joint interaction tests found no detectable variation across seniority or HR function for either requirement. The evidence therefore identifies a selective AI premium. It does not support a common premium across the two requirement measures examined or a clear salary gradient across HR roles.

Overall, the results show that broad references to a technical-skill premium can conceal differences between specific capabilities. The AI association remained after Strict Advanced Data and the observed vacancy characteristics were taken into account. In comparison, the initial Strict Advanced Data association fell substantially after seniority and HR function were introduced. This pattern may be consistent with AI-related requirements functioning as a stronger vacancy signal, while Strict Advanced Data requirements may be more closely embedded in established HR role structures. The observed AI association may reflect both the responsibilities associated with these requirements and the pay practices of employers posting such vacancies. At a more specific level, the study provides focused evidence on job-level skill pricing within HR by showing that two requirements can have different conditional associations with advertised salary.

For HR practice, AI requirements should be tied to clearly specified responsibilities and valued within the whole role. Capability development and pay benchmarks should be reviewed over time because the observed premium may change as employer demand develops. The significance extends beyond pay: AI-supported HR systems shape

recruitment and employee decisions, making clear accountability and effective human oversight essential. The evidence comes from one platform and one cross-section, includes 34 AI-related advertisements, uses binary requirement measures and observes advertised fixed or basic pay. Employer-linked longitudinal evidence and richer task measures are needed to establish whether the premium persists and how far it reflects employer composition. The findings provide a field-specific point of comparison for future research tracking how the pay, responsibilities and distribution of AI-related requirements across HR roles change over time.

References

- Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of labor economics* (Vol. 4B, pp. 1043–1171). Elsevier. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293–S340. <https://doi.org/10.1086/718327>
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Albanesi, S., Dias da Silva, A., Jimeno, J. F., Lamo, A., & Wabitsch, A. (2025). New technologies and jobs in Europe. *Economic Policy*, 40(121), 71–139. <https://doi.org/10.1093/epolic/eiae058>
- Alekseeva, L., Azar, J., Giné, M., & Samila, S. (2026). Artificial intelligence adoption and the demand for managerial expertise. *Strategic Management Journal*. Advance online publication. <https://doi.org/10.1002/smj.70099>
- Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for AI skills in the labor market. *Labour Economics*, 71, Article 102002. <https://doi.org/10.1016/j.labeco.2021.102002>
- Ameye, N., Bughin, J., & van Zeebroeck, N. (2023). How uncertainty shapes herding in the corporate use of artificial intelligence technology. *Technovation*, 127, Article 102846. <https://doi.org/10.1016/j.technovation.2023.102846>

- Arthur, R. (2021). Studying the UK job market during the COVID-19 crisis with online job ads. *PLOS ONE*, 16(5), Article e0251431.
<https://doi.org/10.1371/journal.pone.0251431>
- Atalay, E., Sotelo, S., & Tannenbaum, D. (2024). The geography of job tasks. *Journal of Labor Economics*, 42(4), 979–1008. <https://doi.org/10.1086/725360>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, Article 103745.
<https://doi.org/10.1016/j.jfineco.2023.103745>
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2025). Firm investments in artificial intelligence technologies and changes in workforce composition. In S. Basu, L. Eldridge, J. Haltiwanger, & E. Strassner (Eds.), *Technology, productivity, and economic growth* (pp. 75–117). University of Chicago Press.
<https://www.nber.org/books-and-chapters/technology-productivity-and-economic-growth/firm-investments-artificial-intelligence-technologies-and-changes-workforce-composition>
- Bhattacharya, M., & Wright, P. M. (2005). Managing human assets in an uncertain world: Applying real options theory to HRM. *The International Journal of Human Resource Management*, 16(6), 929–948.
<https://doi.org/10.1080/09585190500120574>
- Bone, M., González Ehlinger, E., & Stephany, F. (2025). Skills or degree? The rise of skill-based hiring for AI and green jobs. *Technological Forecasting and Social Change*, 214, Article 124042. <https://doi.org/10.1016/j.techfore.2025.124042>

- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372. <https://doi.org/10.1257/mac.20180386>
- Chugunova, M., & Danilov, A. (2023). Use of digital technologies for HR management in Germany: Survey evidence. *CESifo Economic Studies*, 69(2), 69–90. <https://doi.org/10.1093/cesifo/ifad005>
- Cirera, X., & Martins-Neto, A. S. (2023). Do innovative firms pay higher wages? Micro-level evidence from Brazil. *Research Policy*, 52(1), Article 104645. <https://doi.org/10.1016/j.respol.2022.104645>
- Cruz, I. F. (2024). Expert-AI pairings: Expert interventions in AI-powered decisions. *Information and Organization*, 34(4), Article 100527. <https://doi.org/10.1016/j.infoandorg.2024.100527>
- Deepa, R., Sekar, S., Malik, A., Kumar, J., & Attri, R. (2024). Impact of AI-focussed technologies on social and technical competencies for HR managers—A systematic review and research agenda. *Technological Forecasting and Social Change*, 202, Article 123301. <https://doi.org/10.1016/j.techfore.2024.123301>
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4), 1593–1640. <https://doi.org/10.1093/qje/qjx022>
- Deming, D., & Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1), S337–S369. <https://doi.org/10.1086/694106>

- Dima, J., Gilbert, M.-H., Dextras-Gauthier, J., & Giraud, L. (2024). The effects of artificial intelligence on human resource activities and the roles of the human resource triad: Opportunities and challenges. *Frontiers in Psychology*, 15, Article 1360401. <https://doi.org/10.3389/fpsyg.2024.1360401>
- Espegren, Y. (2025). Reasons for HR analytics adoption in public sector organisations: Evidence from Swedish public administrations. *Personnel Review*, 54(7), 1621–1642. <https://doi.org/10.1108/PR-03-2024-0219>
- Fichman, R. G. (2004). Real options and IT platform adoption: Implications for theory and practice. *Information Systems Research*, 15(2), 132–154. <https://doi.org/10.1287/isre.1040.0021>
- Garcia-Lazaro, A., Mendez-Astudillo, J., Lattanzio, S., Larkin, C., & Newnes, L. (2025). The digital skill premium: Evidence from job vacancy data. *Economics Letters*, 250, Article 112294. <https://doi.org/10.1016/j.econlet.2025.112294>
- Gerber, M., Krause, A., Probst, J., & Heimann, M. (2024). HR analytics between ambition and reality: Current state and recommendations for the contribution of work and organizational psychology. *Gruppe. Interaktion. Organisation. Zeitschrift für Angewandte Organisationspsychologie (GIO)*, 55(2), 225–236. <https://doi.org/10.1007/s11612-024-00743-7>
- Kalyani, A., Bloom, N., Carvalho, M., Hassan, T., Lerner, J., & Tahoun, A. (2025). The diffusion of new technologies. *The Quarterly Journal of Economics*, 140(2), 1299–1365. <https://doi.org/10.1093/qje/qjaf002>
- Kashive, N., & Khanna, V. T. (2023). Emerging HR analytics role in a crisis: An analysis of LinkedIn data. *Competitiveness Review: An International Business Journal*, 33(6), 1179–1204. <https://doi.org/10.1108/CR-03-2022-0029>

- Khaouja, I., Kassou, I., & Ghogho, M. (2021). A survey on skill identification from online job ads. *IEEE Access*, 9, 118134–118153.
<https://doi.org/10.1109/ACCESS.2021.3106120>
- Kline, P., Petkova, N., Williams, H., & Zidar, O. (2019). Who profits from patents? Rent-sharing at innovative firms. *The Quarterly Journal of Economics*, 134(3), 1343–1404. <https://doi.org/10.1093/qje/qjz011>
- Marler, J. H., & Boudreau, J. W. (2017). An evidence-based review of HR analytics. *The International Journal of Human Resource Management*, 28(1), 3–26.
<https://doi.org/10.1080/09585192.2016.1244699>
- McCartney, S., Murphy, C., & McCarthy, J. (2021). 21st century HR: A competency model for the emerging role of HR analysts. *Personnel Review*, 50(6), 1495–1513.
<https://doi.org/10.1108/PR-12-2019-0670>
- McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L., & Zolas, N. (2024). AI adoption in America: Who, what, and where. *Journal of Economics & Management Strategy*, 33(2), 375–415. <https://doi.org/10.1111/jems.12576>
- Nakash, M., & Peretz, O. (2025). Decoding the AI job market: Mapping skills and classifying careers through online job ads. *Employee Relations: The International Journal*. Advance online publication. <https://doi.org/10.1108/ER-07-2025-0566>
- Nyberg, A. J., Schleicher, D. J., Bell, B. S., Boon, C., Cappelli, P., Collings, D. G., Dalle Molle, J. E., Feuerriegel, S., Gerhart, B., Jeong, Y., Korsgaard, M. A., Minbaeva, D., Ployhart, R. E., Tambe, P., Weller, I., Wright, P. M., & Yakubovich, V. (2025). A brave new world of human resources research: Navigating perils and identifying grand challenges of the GenAI revolution. *Journal of Management*, 51(6), 2677–2718. <https://doi.org/10.1177/01492063251325188>

Office for National Statistics. (2026, July 20). *Artificial intelligence in UK businesses: 2023 to 2026*.

<https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/articles/artificialintelligenceinukbusinesses/2023to2026>

Pan, Y., & Froese, F. J. (2023). An interdisciplinary review of AI and HRM: Challenges and future directions. *Human Resource Management Review*, 33(1), Article 100924. <https://doi.org/10.1016/j.hrmr.2022.100924>

Ridgway, M., Oldridge, L., & Mavin, S. (2025). ‘Leading’ by example? Gendered language in human resource job adverts. *Human Resource Management Journal*, 35(1), 1–24. <https://doi.org/10.1111/1748-8583.12549>

Soleimani, M., Intezari, A., Arrowsmith, J., Pauleen, D. J., & Taskin, N. (2025). Reducing AI bias in recruitment and selection: An integrative grounded approach. *The International Journal of Human Resource Management*, 36(14), 2480–2515. <https://doi.org/10.1080/09585192.2025.2480617>

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374. <https://doi.org/10.2307/1882010>

Stephany, F., & Teutloff, O. (2024). What is the price of a skill? The value of complementarity. *Research Policy*, 53(1), Article 104898. <https://doi.org/10.1016/j.respol.2023.104898>

Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. *California Management Review*, 61(4), 15–42. <https://doi.org/10.1177/0008125619867910>

van den Broek, E., Sergeeva, A., & Huysman, M. (2021). When the machine meets the expert: An ethnography of developing AI for hiring. *MIS Quarterly*, 45(3), 1557–1580. <https://doi.org/10.25300/MISQ/2021/16559>

- Vrontis, D., Christofi, M., Pereira, V., Tarba, S., Makrides, A., & Trichina, E. (2022). Artificial intelligence, robotics, advanced technologies and human resource management: A systematic review. *The International Journal of Human Resource Management*, 33(6), 1237–1266. <https://doi.org/10.1080/09585192.2020.1871398>
- Wirges, F., & Neyer, A.-K. (2023). Towards a process-oriented understanding of HR analytics: Implementation and application. *Review of Managerial Science*, 17(6), 2077–2108. <https://doi.org/10.1007/s11846-022-00574-0>