



EFFECT OF AI-COMPLEMENTARY SKILL DEVELOPMENT ON WAGE PREMIA IN INDIAN IT SERVICES

HRM VARIABLE: SKILL DEVELOPMENT INVESTMENT ECONOMIC OUTCOME: WAGE PREMIUM

Dr. N Subbukrishna Sastry¹, Dr. Manjula Mallya M²

¹Professor, School of Management, CMR University, Bangalore, Karnataka, India.

ORCID ID: 0009-0009-0460-7057,

²Associate Professor & Head, Dept. of Economics, Government First Grade College for Women Balmatta Mangalore, Karnataka, India. ORCID ID: 0009-0005-8812-6912,

ABSTRACT

Indian IT services stands at the intersection of two forces: a USD 300 billion industry employing close to 5.8 million people, and a generative-AI transition that is simultaneously automating routine delivery work and creating acute scarcity in AI-complementary capability. This study examines a single, tightly specified HRM–economics linkage — the effect of employer skill development investment (SDI) on the wage premium earned by AI-complementary professionals relative to matched non-AI peers. Using a mixed-methods design covering 480 IT professionals and 60 learning-and-development managers across Bengaluru, Hyderabad, Pune, Chennai and Delhi-NCR, the study estimates an augmented Mincerian earnings function in which AI-complementary certification, skill development investment, and their interaction are the focal regressors. The model explains 71.2 per cent of the variance in log annual compensation. AI-complementary skill holding carries a base premium of 20.6 per cent at mean investment, rising to 36.2 per cent for employees in the highest investment quartile — evidence of a significant positive interaction ($\beta = 0.061$, $p = 0.004$) between employer investment and individual skill. The dose–response relationship is concave rather than linear, with a threshold near ₹18,000 per employee per year below which premium capture is weak. Decomposition attributes 36 per cent of the mean 31.1 per cent premium to scarcity rent, 32 per cent to genuine productivity complementarity, 17 per cent to signalling and 15 per cent to task-portfolio reallocation — implying that roughly a third of the premium is transitional and will erode as supply catches up. Employer-type effects are pronounced: global capability centres invest 2.9 times what mid-tier Indian firms invest and capture a 16.9 percentage-point higher premium. A residual gender gap of 4.7 per cent persists after controlling for skill, investment and experience. The study argues that skill development investment operates as a produced factor of wage formation rather than a welfare expense, and offers a threshold-based investment framework for HR leaders, a first-year return schedule for finance functions, and calibration guidance for national reskilling missions.

KEYWORDS — AI-Complementary Skills, Skill Development Investment, Wage Premium, Mincerian Earnings Function, Human Capital Theory, Indian IT Services, Strategic HRM, Reskilling, Labour Economics, Generative AI, Global Capability Centres, Skills-Based Pay, Task-Biased Technological Change.

I. INTRODUCTION

The Indian technology industry crossed a defining threshold in the financial year 2026. Aggregate revenue approached the USD 300 billion mark, having reached an estimated USD 282.6 billion in FY2025 on growth of 5.1 per cent, with exports of USD 224.4 billion and a domestic market of USD 58.2 billion (NASSCOM, 2025). Yet the employment story diverged sharply from the revenue story. Total headcount stood at approximately 5.8 million, but net addition for the year was only about 126,000 — a fraction of the intake the industry absorbed during the pandemic-era expansion. For the first time in three decades, the industry is growing its revenue substantially faster than its workforce. The linear relationship between billable heads and billable revenue, which underwrote the Indian IT services business model since the 1990s, has been broken.

The proximate cause is artificial intelligence. More than ninety per cent of the top twenty Indian services companies now embed AI, cloud, data and generative-AI capability into their delivery stack, and sixty-three per cent of technology chief executives project AI investment exceeding ten per cent of total technology spending in



FY2026 (NASSCOM, 2025). The consequence for work is not uniform displacement but a recomposition of tasks within existing roles. Code generation, first-pass testing, documentation, ticket triage and routine data transformation are increasingly machine-executed; specification, orchestration, evaluation, governance and domain translation are not. Autor, Levy and Murnane (2003) anticipated precisely this pattern when they argued that technology substitutes for routine cognitive tasks while complementing non-routine ones, and Acemoglu and Restrepo (2019, 2022) formalised the resulting distinction between the displacement effect, which lowers the labour share, and the reinstatement effect, which raises it where new labour-intensive tasks emerge.

The labour market has repriced accordingly, and the magnitude of the repricing is the striking fact. PwC's Global AI Jobs Barometer, drawing on close to a billion job advertisements across six continents, reported a 56 per cent wage premium in 2025 for workers holding AI skills relative to same-occupation colleagues without them — more than double the 25 per cent premium recorded a year earlier (PwC, 2025). Indian evidence points in the same direction with different magnitudes. An analysis of 950 technology postings across five metropolitan markets found AI roles commanding 20 per cent more than traditional IT roles at entry level, widening to 38 per cent beyond ten years of experience (1 Finance Research, 2026). EY's Future of Pay 2026 survey placed the skill-base premium for AI, generative AI, machine learning and cybersecurity capability at 30 to 40 per cent, against an all-India average increment of 9.1 per cent, and reported that 45 to 50 per cent of Indian organisations are actively migrating from role-based to skills-based pay architectures (EY, 2026).

For the human resource function, these figures pose a question of investment rather than of observation. If AI-complementary capability commands a substantial premium, the operative issue is who produces that capability and who captures its return. Becker's (1964) classical answer distinguishes general human capital, whose returns are portable and therefore financed by the worker, from firm-specific human capital, whose returns are captured jointly and therefore financed by the firm. AI-complementary skill sits awkwardly between these poles. A generative-AI engineering certification is portable, which argues for worker financing; but its productive value depends on deployment against the employer's client portfolio, delivery methodology and domain data, which argues for employer financing. Lazear's (2009) skill-weights approach resolves the tension conceptually — firm specificity arises from the firm's particular combination of general skills rather than from any single skill — but leaves the empirical question of who invests, how much, and with what wage consequence entirely open in the Indian context.

This study addresses that question directly. It isolates one HRM variable — skill development investment, measured as employer expenditure, released learning hours and certification attainment per employee — and one economic outcome — the wage premium, measured as the log-wage differential against a matched non-AI peer. It asks whether the premium is a property of the skill alone or a joint product of the skill and the investment that produced it; whether the investment–premium relationship is linear or threshold-bounded; how much of the observed premium reflects durable productivity complementarity as against transitory scarcity rent; and how far employer type, skill cluster and gender moderate premium capture.

The paper proceeds as follows. Section II reviews the theoretical and empirical literature on human capital, task-biased technological change and AI wage premia, and situates the study within an existing programme of research linking specific HRM variables to measurable economic outcomes. Section III identifies the research gap and Section IV states the problem. Sections V and VI set out objectives and hypotheses. Section VII details the methodology, the operationalisation of constructs and the estimating equation. Section VIII presents results and discussion across seven analytical strands. Sections IX to XIV report findings, significance, recommendations, limitations, scope for further work and conclusions.

II. REVIEW OF LITERATURE

2.1 Theoretical Foundations

The analytical foundation of this study rests on four theoretical traditions. The first is human capital theory. Schultz (1961) and Becker (1964) established that expenditure on education and training is an investment yielding a measurable rate of return rather than a consumption item, and that the financing of that investment follows from the portability of the skill produced. Mincer (1974) converted the insight into an estimable earnings function in which log earnings are a linear function of schooling and a concave function of experience — the specification that remains, five decades later, the workhorse of applied wage analysis, and whose properties and limitations were comprehensively assessed by Heckman, Lochner and Todd (2006).



The second tradition concerns the task content of technological change. Autor, Levy and Murnane (2003) demonstrated that computerisation substitutes for routine tasks, whether cognitive or manual, while complementing non-routine problem-solving and complex communication. Autor (2015) extended the argument to explain why aggregate employment persists despite automation: complementarity, elastic final demand and the inelastic supply of human labour to non-automated tasks all counteract displacement. Acemoglu and Restrepo (2019) formalised the dual mechanism, showing that automation exerts a displacement effect that lowers labour demand while the creation of new tasks exerts a reinstatement effect that raises it, and Acemoglu and Restrepo (2022) attributed a substantial share of the post-1980 rise in United States wage inequality to task displacement concentrated among routine-intensive occupations.

The third tradition is signalling. Spence (1973) showed that in the presence of asymmetric information, a costly and observable credential can raise wages even where its direct productive contribution is limited, because it separates high-ability from low-ability workers. In fast-moving technology labour markets where employers cannot readily verify capability at the point of hire, vendor and platform certifications plausibly carry a signalling value distinct from their productive value — a distinction this study attempts to quantify rather than assume away. The fourth tradition is the skill-obsolescence literature. Deming and Noray (2020) showed that the wage returns to technology-intensive STEM occupations decay rapidly as the skill content of those occupations changes, so that early-career premia erode unless continuously renewed. Goldin and Katz (2008) framed the broader dynamic as a race between education and technology, in which the wage structure is determined by the relative pace of skill supply and skill-biased demand. Together these strands imply that an AI wage premium observed at a point in time is an equilibrium outcome of a supply–demand race, not a permanent property of the skill — a proposition central to the decomposition attempted in Section VIII.

2.2 Empirical Evidence on Artificial Intelligence Wage Premia

Direct empirical estimation of AI wage premia is recent. Alekseeva, Azar, Giné, Samila and Taska (2021), analysing the near-universe of United States online job postings, found that AI skill requirements carry a substantial wage premium that survives the inclusion of firm and occupation fixed effects, implying that the premium reflects genuine within-firm skill scarcity rather than the sorting of AI skills into high-paying firms. Felten, Raj and Seamans (2021) constructed occupational exposure measures showing that AI capability maps unevenly onto occupations, with the highest exposure in analytically intensive white-collar work — precisely the composition of Indian IT services delivery.

On the productivity side, Brynjolfsson, Li and Raymond (2025), in a staged deployment of a generative-AI assistant to more than five thousand customer-support agents, recorded an average productivity gain of about fourteen per cent, concentrated overwhelmingly among novice and low-skilled workers, with negligible or slightly negative effects for the most experienced. The finding is significant for premium analysis because it suggests that generative AI compresses rather than widens within-occupation productivity dispersion at the task level, which would imply that observed wage premia arise less from productivity differentials among AI users and more from the scarcity of those able to design, govern and integrate such systems. Bessen (2019) offered the complementary demand-side argument that automation raises employment where product demand is elastic and reduces it where demand is saturated — a distinction that separates AI-enabled new service lines from AI-automated legacy maintenance work within the same Indian IT firm.

PwC's (2025) Global AI Jobs Barometer provides the largest-scale premium estimate available, comparing workers in the same occupation who differ only in whether they hold AI skills, and reporting a 56 per cent premium in 2025 against 25 per cent in 2024. The doubling within a single year is itself informative: a premium that moves that fast is unlikely to reflect a stable productivity differential and is more consistent with a rapidly widening supply–demand gap.

2.3 The Indian Evidence Base

Indian evidence combines industry-body assessment with market-pricing analysis. The Deloitte and NASSCOM (2024) study projected India's AI talent demand rising from 600,000–650,000 professionals to more than 1.25 million by 2027, against an AI market expanding at 25 to 35 per cent annually. NASSCOM's subsequent analysis of the AI talent inflection point (NASSCOM & Deloitte, 2025) recorded that 43 per cent of the Indian workforce has used AI tools within their organisation and that roughly 60 per cent of workers — and 71 per cent of Generation Z workers — recognise AI capability as career-enhancing. The India Skills Report 2026 placed overall



employability at 56.35 per cent, an improvement that nonetheless leaves a substantial share of the annual graduate cohort outside the employable pool for AI-adjacent work (Wheebox, 2026).

On pricing, the 1 Finance Research (2026) analysis of 950 technology postings across Mumbai, Pune, Delhi-NCR, Bengaluru and Hyderabad produced the experience-graded premium schedule that this study adopts as its external benchmark: 20 per cent at zero to three years, 33 per cent at three to five years, 29 per cent at five to ten years and 38 per cent beyond ten years. The same analysis found that global capability centres pay 12 to 20 per cent above traditional Indian IT services firms for comparable roles, widening to 30 to 50 per cent for AI and data specialists, and that 64 per cent of newly created capability-centre positions specify AI, analytics or automation requirements. Demand for such talent is reported growing at approximately 25 per cent annually against supply growth near 15 per cent — an arithmetic that generates a widening gap and, by implication, a rising scarcity component in the premium.

EY's (2026) Future of Pay survey adds the compensation-architecture dimension. Against a projected all-India increment of 9.1 per cent for 2026, with capability centres leading at 10.4 per cent, the survey placed skill premia for AI, machine learning, cybersecurity and cloud capability at 30 to 40 per cent, and found 45 to 50 per cent of organisations transitioning to skills-based pay frameworks. Attrition in the hi-tech and IT segment stood at 20.5 per cent against 14.1 per cent in capability centres, and top performers were receiving 120 to 150 per cent of target payout against 60 to 80 per cent for average performers. The evidence collectively indicates that Indian technology compensation is being restructured around demonstrable capability rather than tenure or grade.

2.4 Antecedent Work Linking HRM Variables to Economic Outcomes

The present study extends a continuing research programme that pairs a specific human resource variable with a measurable economic outcome. Sastry and Mallya (2025b) examined strategic human resource management as a determinant of economic resilience, comparing developed and emerging economies and finding that advanced economies deploy AI-driven HR analytics and structured reskilling while emerging markets emphasise digital literacy and inclusive upskilling; by correlating HR practice with productivity and wage dynamics, that study established the analytical template — an HR practice as cause, an economic aggregate as effect — on which the present work builds, and identified continuous learning and digital skill development as the practices most strongly associated with workforce sustainability under AI pressure.

Sastry and Mallya (2025a) analysed AI analytics in recruitment, showing that machine learning and natural language processing applied to candidate screening deliver both immediate operational gains and sustained financial benefit through reduced turnover cost and faster acquisition of scarce talent, while raising the fairness question that any algorithmic capability assessment must answer. That study is directly relevant here because the same analytics infrastructure that prices capability at the point of hire also determines which internal capabilities an employer chooses to fund. Mallya and Sastry (2025) examined leadership strategy for young talent management across 150 young employees and 50 HR managers, finding that growth opportunity and purpose ranked above compensation as retention drivers for Millennial and Generation Z professionals — a finding that bears directly on the non-wage return to skill development investment measured in Section VIII.

Sastry (2025b) investigated talent designation and social comparison as economic drivers of organisational growth, demonstrating empirically that the relative positioning of employees against visible internal reference groups shapes effort and retention independently of absolute pay — a mechanism that helps explain why an AI wage premium, once visible within a delivery organisation, alters the behaviour of non-recipients as well as recipients. Sastry (2025a), studying 150 women micro-entrepreneurs across five Karnataka districts using the Garrett ranking technique, concluded that socio-cultural and systemic constraints bind more tightly than skill or knowledge gaps — a caution against treating training expenditure as a sufficient condition for economic advancement, and one that informs this study's treatment of the residual gender gap.

Earlier work in the same programme established the underlying performance mechanisms. Sastry (2020) examined the influence of performance factors on motivation in the tea manufacturing sector, and Sastry (2013) set out approaches to talent management within the human resource function. Anas, Hans and Sastry (2025) demonstrated, in the tourism setting, how a technology shift restructures an entire service value chain and the competences required to operate it — the same structural logic that this study applies to software services. Table 1 synthesises the literature and states the specific contribution each source makes to the present design.



Table 1 · Synthesis of the Literature and Contribution to the Present Study

Author(s) & Year	Focus of Study	Key Findings	Relevance to Current Study
Becker (1964); Schultz (1961)	Human capital theory	Training is investment yielding a measurable return; financing follows the portability of the skill produced	Supplies the general-versus-specific framing for who should fund AI-complementary skilling
Mincer (1974); Heckman et al. (2006)	Earnings function	Log earnings are linear in schooling and concave in experience; the specification remains the standard for wage estimation	Provides the estimating equation this study augments with the HRM variable
Autor, Levy & Murnane (2003)	Task content of technical change	Technology substitutes for routine tasks and complements non-routine problem-solving and communication	Establishes why AI-complementary and AI-substitutable skills should price differently (H ₂)
Acemoglu & Restrepo (2019, 2022)	Displacement and reinstatement	Automation lowers labour demand through displacement while new tasks raise it; task displacement drives wage inequality	Grounds the premium as a task-composition rather than a headcount phenomenon
Spence (1973)	Job market signalling	A costly observable credential raises wages under information asymmetry independently of productive contribution	Justifies isolating a distinct signalling component in the premium decomposition
Deming & Noray (2020)	Skill obsolescence in STEM	Returns to technology-intensive occupations decay rapidly as skill content shifts	Predicts erosion of the premium absent continuous renewal; informs Section 8.4
Alekseeva et al. (2021)	AI skill demand in postings	AI skill requirements carry a wage premium robust to firm and occupation fixed effects	Confirms the premium is within-firm scarcity rather than firm sorting
Brynjolfsson, Li & Raymond (2025)	Generative AI at work	Roughly 14 per cent productivity gain, concentrated among novices; negligible for the most experienced	Suggests the premium reflects scarcity of designers and integrators more than user productivity
PwC (2025)	Global AI Jobs Barometer	56 per cent wage premium for AI skills in 2025, up from 25 per cent in 2024, across a billion job advertisements	Global benchmark; the speed of change signals a scarcity-driven rather than stable premium
Deloitte & NASSCOM (2024, 2025)	India AI talent pool	Demand rising from 600–650 thousand to over 1.25 million by 2027; market growing 25–35 per cent annually	Supplies the supply-expansion path against which premium erosion is projected
EY (2026)	Future of Pay, India	9.1 per cent average increment; 30–40 per cent AI skill premium; 45–50 per cent of firms moving to skills-based pay	Confirms the compensation-architecture shift the study analyses
1 Finance Research (2026)	India AI vs IT pay bands	AI roles pay 20–38 per cent more across experience bands; GCCs pay 30–50 per cent more for AI and data specialists	External benchmark for the descriptive premium schedule in Table 4
Sastry & Mallya (2025b)	Strategic HRM and economic resilience	Continuous learning and digital skill development enhance workforce sustainability; HR practice correlates with productivity and wage dynamics	Direct antecedent; establishes the HRM-variable-to-economic-outcome template used here
Sastry & Mallya (2025a)	AI analytics in HR hiring	AI screening delivers operational gain and sustained financial benefit through lower turnover cost and faster acquisition	Links capability pricing at hire to internal investment decisions
Mallya & Sastry (2025)	Leadership for young talent	Growth opportunity and purpose outrank pay as retention drivers for Millennial and Gen Z professionals	Underpins the non-wage retention return quantified in Table 8
Sastry (2025b)	Talent designation and social comparison	Relative internal positioning shapes effort and retention independently of absolute pay	Explains spillover effects of a visible premium on non-recipients



Sastry (2025a)	Women micro-entrepreneurs, Karnataka	Socio-cultural and systemic constraints bind more tightly than skill or knowledge gaps	Frames interpretation of the residual gender differential (H_6)
Anas, Hans & Sastry (2025)	Digital transformation in tourism	Technology restructures the service value chain and the competences required to operate it	Cross-sector analogue for capability restructuring under a technology shift
Sastry (2020); Sastry (2013)	Performance, motivation and talent management	Performance factors drive motivation; structured talent management approaches within the HR function	Establishes the performance mechanisms underlying the investment–outcome link

III. RESEARCH GAP

Five gaps emerge from the review, and together they define the space this study occupies.

- The large-scale premium estimates that dominate the literature are drawn from United States or global posting data; Indian estimates exist but are commercial market scans rather than controlled estimations, and none isolates the Indian IT services delivery workforce specifically.
- The literature treats AI skill as a binary attribute of the worker. It does not treat that skill as the output of a measurable employer investment, and therefore cannot say whether the premium is a property of the skill or a joint product of the skill and the investment that produced it. The interaction term is absent from existing specifications.
- No study decomposes the premium. Reporting that AI skills pay 56 per cent more is analytically incomplete unless one can say how much of that figure is durable productivity complementarity and how much is transitory scarcity rent, because the two have opposite implications for how long an investment payback period may safely be assumed.
- The functional form of the investment–premium relationship is untested. Practitioner guidance implicitly assumes linearity, which would imply that any increment in training budget yields proportionate return; if the true relationship is concave with a threshold, sub-threshold investment is largely wasted and the policy implication reverses.
- The employer's return is unquantified. The literature documents what AI-skilled employees earn but not what the firms that trained them recover, leaving the business case for skill development investment resting on assertion rather than measurement.

IV. STATEMENT OF THE PROBLEM

Indian IT services firms are being asked to fund a large-scale reskilling transition at precisely the moment their headcount-linked revenue model is weakening. Learning-and-development budgets are therefore under scrutiny, and the case for expanding them rests on an unproven premise: that employer investment in AI-complementary capability raises the value of the workforce by more than it costs.

At the same time, employees observe a visible and widening wage gap between AI-complementary and traditional delivery roles, but face no clear guidance on what quantum or type of skill acquisition converts into realised compensation. Certification proliferation has outpaced evidence on which certifications pay.

The problem this study addresses is therefore twofold: to establish whether, and by how much, employer skill development investment causally augments the AI wage premium; and to determine the shape of that relationship, so that both firms and individuals can identify the investment level at which returns become material.

V. OBJECTIVES OF THE STUDY

1. To estimate the wage premium accruing to AI-complementary skill holding within Indian IT services, net of experience, education, employer type and location.
2. To test whether employer skill development investment independently raises compensation and whether it interacts with AI-complementary skill to amplify the premium.
3. To determine the functional form of the investment–premium relationship and to identify any investment threshold below which premium capture is negligible.
4. To decompose the observed premium into scarcity rent, productivity complementarity, signalling and task-portfolio reallocation components.
5. To examine heterogeneity in premium capture across AI skill clusters, employer types, experience bands and gender.
6. To quantify the employer's first-year return on skill development investment and to derive an evidence-based investment framework for HR and finance leaders.



VI. HYPOTHESES OF THE STUDY

H₁: Skill development investment exerts a positive and statistically significant effect on individual compensation in Indian IT services.

H₂: AI-complementary skills command a significantly higher wage premium than AI-substitutable routine digital skills.

H₃: Skill development investment positively moderates the AI-skill–wage relationship, such that the premium rises with the level of employer investment.

H₄: Mean realised wage premium differs significantly across employer types (global capability centre, Tier-1 Indian IT, mid-tier Indian IT, product firm).

H₅: The relationship between skill development investment and wage premium is linear across the full investment range.

H₆: Gender does not moderate realised premium capture once skill, investment and experience are controlled.

VII. RESEARCH METHODOLOGY AND DESIGN

7.1 Research Type and Design

The study employs a mixed-methods, cross-sectional explanatory design. The quantitative component estimates an augmented Mincerian earnings function on individual-level compensation and training data. The qualitative component comprises semi-structured interviews with learning-and-development leaders, used to interpret the estimated coefficients and to establish the mechanisms through which investment translates into premium. The design is explanatory rather than exploratory because the direction of the hypothesised relationships is theoretically established; the contribution lies in magnitude, functional form and decomposition rather than in discovery.

7.2 Population, Sampling and Instrument

The target population comprises professionals employed in Indian IT services delivery, engineering research and development, business process management and global capability centres across the five principal technology labour markets. A stratified purposive sample of 480 professionals was drawn, stratified by city, employer type and experience band, together with 60 learning-and-development and compensation managers. Stratum sizes were set proportional to published city-level and firm-type employment shares, with deliberate over-sampling of the ten-plus-year band, where premium dispersion is widest and precision therefore most costly to obtain.

The instrument comprised a structured questionnaire capturing compensation, role, skill portfolio, certification history, employer-funded and self-funded learning expenditure, released learning hours and AI deployment intensity, with attitudinal items measured on a five-point Likert scale. Compensation was collected in banded form and converted to band midpoints to reduce non-response. Internal consistency was satisfactory (Cronbach's alpha = 0.87 across attitudinal scales). Sampling adequacy for the factor structure was confirmed (Kaiser–Meyer–Olkin = 0.842; Bartlett's test of sphericity, $p < 0.001$). Multicollinearity diagnostics returned a maximum variance inflation factor of 2.31, well within the conventional threshold of 5. Tables 2 and 3 report construct operationalisation and the realised sample profile.

Table 2 · Operationalisation of Constructs

Construct	Indicator / Proxy	Measurement	Data Source
Skill Development Investment (SDI) — independent	Employer expenditure per employee per year; released learning hours; employer-sponsored certifications completed	₹ per employee per annum (continuous, log-transformed and mean-centred); hours per annum; count	Employee questionnaire, validated against L&D manager estimates
AI-Complementary Skill (AISkill) — independent	Verified capability in generative AI, MLOps, data engineering for AI, agent design or AI governance	Binary (1 = verified capability, 0 = otherwise), corroborated by certification or project evidence	Questionnaire, certification records
Deployment Intensity — mediator	Share of billable hours worked on AI-enabled client engagements	Percentage of annual billable hours (continuous)	Questionnaire, project allocation records
Wage Premium — dependent	Log-wage differential against a matched non-AI peer in the same role, experience band and firm tier	$\exp(\beta_1 + \beta_3 \cdot \ln SDI) - 1$, expressed as a percentage	Derived from banded compensation data
Compensation (W)	Annual cost-to-company	₹ lakh per annum, collected in bands and converted to midpoints; natural log used in estimation	Questionnaire



Experience (Exp)	Total years of relevant professional experience	Years, entered in linear and quadratic form	Questionnaire
Employer Type (Firm)	Global capability centre; Tier-1 Indian IT; mid-tier Indian IT; product firm or AI start-up	Categorical, dummy-coded with mid-tier Indian IT as reference	Questionnaire
Controls	Post-graduate qualification; metropolitan labour market; gender	Binary indicators	Questionnaire

Table 3 · Realised Sample Profile (N = 480 professionals; 60 L&D and compensation managers)

Stratification Dimension	Category	n	Share (%)
Metropolitan market	Bengaluru	150	31.3
	Hyderabad	95	19.8
	Pune	80	16.7
	Delhi-NCR	80	16.7
	Chennai	75	15.6
Employer type	Tier-1 Indian IT services	165	34.4
	Global capability centre	145	30.2
	Mid-tier Indian IT services	110	22.9
	Product firm / AI start-up	60	12.5
Experience band	0–3 years	120	25.0
	3–5 years	130	27.1
	5–10 years	140	29.2
	10+ years	90	18.8
Skill classification	AI-complementary capability	268	55.8
	AI-substitutable / routine digital	212	44.2
Gender	Male	302	62.9
	Female	178	37.1

7.3 Model Specification

The estimating equation augments the standard Mincerian form with the focal HRM variable, its interaction with AI-complementary skill holding, and controls for deployment, employer type, location and gender:

$$\ln W_i = \beta_0 + \beta_1 AISkill_i + \beta_2 \ln SDI_i + \beta_3 (AISkill \times \ln SDI)_i + \beta_4 Exp_i + \beta_5 Exp_i^2 + \beta_6 PG_i + \beta_7 Deploy_i + \beta_8 Firm_i + \beta_9 City_i + \beta_{10} Gender_i + \varepsilon_i \quad \dots(1)$$

where W_i is annual cost-to-company for individual i , $AISkill_i$ is a binary indicator of verified AI-complementary capability, $\ln SDI_i$ is the natural logarithm of annual employer skill development investment per employee (mean-centred so that β_1 is interpretable as the premium at average investment), Exp_i is years of experience entered in quadratic form to capture concavity, $Deploy_i$ is the share of billable hours on AI-enabled engagements, and the remaining terms are controls. The wage premium is recovered as $\exp(\beta_1 + \beta_3 \cdot \ln SDI) - 1$, evaluated at the relevant investment level, which makes the premium a function of investment rather than a constant.

Figure 1 sets out the conceptual framework the equation operationalises. Skill development investment acts on wage outcomes through two mediating channels — the acquisition of AI-complementary capability and the intensity with which that capability is deployed against billable work — which together drive four economic transmission mechanisms: realisation-rate improvement, task-portfolio upgrading, scarcity rent and external market signalling. Employer type, experience band, role family, metropolitan labour market and gender enter as moderators.



Figure 1 · Conceptual Framework — Skill Development Investment → Wage Premium




Figure 1 · Conceptual framework linking skill development investment to wage premium through capability acquisition and deployment intensity.

7.4 Analytical Tools

Ordinary least squares with heteroskedasticity-robust standard errors was used for the earnings function. One-way analysis of variance with Scheffé post-hoc comparison tested differences across investment quartiles, and Welch's analysis of variance, which does not assume equal variances, tested differences across employer types. Independent-samples t-tests compared AI-complementary against AI-substitutable skill holders. Premium decomposition followed an Oaxaca-type additive attribution, allocating the log-wage gap across scarcity, productivity, signalling and task-mix components using deployment data, vacancy-to-candidate ratios and interview evidence. Thematic analysis was applied to the interview transcripts. Analysis was conducted in SPSS version 29 with cross-validation in R.

VIII. RESULTS AND DISCUSSION

8.1 Descriptive Wage Differentials

Table 4 and Figure 2 report median compensation for AI-complementary and matched traditional roles by experience band. The gradient is monotonic in absolute terms and non-monotonic in proportional terms. Entry-level professionals with AI-complementary capability earned a median ₹6.2 lakh against ₹5.1 lakh for matched peers, a premium of 21.6 per cent. The premium widened to 34.8 per cent at three to five years, moderated to 30.6 per cent at five to ten years, and reached its maximum of 39.4 per cent beyond ten years, where the absolute gap of ₹10.4 lakh exceeds the entire entry-level salary of a non-AI peer. The sample-weighted mean premium is 31.1 per cent.

The dip at five to ten years is consistent across sub-samples and merits comment. It corresponds to the band in which Indian IT services firms concentrate their people-management ladder: professionals at this stage are typically moving into delivery-lead and project-manager roles where compensation is set by span of control rather than technical scarcity. Those who remain on the technical track capture the full premium; those who transition into management do not, and the band average is pulled down accordingly. The pattern echoes Deming and Noray's (2020) finding that technical premia decay when workers move away from the frontier skill content of their occupation, and carries a direct managerial implication addressed in Section XI.

All four band-level differences are significant at the one per cent level, with t-statistics ranging from 4.82 to 9.16. The overall difference between AI-complementary and AI-substitutable skill holders is large: mean premia of 31.1 per cent and 8.9 per cent respectively, $t(478) = 12.63$, $p < 0.001$, Cohen's $d = 1.15$, indicating an effect of substantial practical as well as statistical significance and providing initial support for H2.



Table 4 · Median Compensation and Realised Wage Premium by Experience Band

Experience Band	n	Traditional IT median CTC (₹ lakh)	AI-complementary median CTC (₹ lakh)	Absolute gap (₹ lakh)	Premium (%)	t-statistic
0–3 years	120	5.1	6.2	1.1	21.6	4.82***
3–5 years	130	9.2	12.4	3.2	34.8	7.94***
5–10 years	140	17.3	22.6	5.3	30.6	6.51***
10+ years	90	26.4	36.8	10.4	39.4	9.16***
All bands (weighted)	480	13.8	18.1	4.3	31.1	12.63***

Note: *** significant at $p < 0.001$. Benchmark comparison: 1 Finance Research (2026) reports 20, 33, 29 and 38 per cent for the corresponding bands.

Figure 2 · AI-Complementary vs. Traditional IT Compensation by Experience Band (survey estimates, N = 480; green callout = realised wage premium over matched peer)




Figure 2 · Median compensation and realised premium by experience band (N = 480).

8.2 Mincerian Estimation

Table 5 reports the augmented earnings function. The model explains 71.2 per cent of variance in log compensation (adjusted $R^2 = 0.703$; $F = 118.4$, $p < 0.001$), which is high for individual wage data and reflects the relative homogeneity of the sampled occupations.

Three coefficients carry the argument. AI-complementary skill holding is associated with a coefficient of 0.187 ($p < 0.001$), implying a premium of 20.6 per cent at mean investment. Skill development investment enters positively and independently ($\beta = 0.094$, $p < 0.001$), confirming H_1 : employer investment raises compensation even for employees without AI-complementary certification, consistent with a general capability effect. Decisively for this study, the interaction term is positive and significant ($\beta = 0.061$, $p = 0.004$), confirming H_3 . The premium is therefore not a fixed property of the skill. Evaluated at one standard deviation above mean investment it rises to 28.1 per cent, and at two standard deviations to 36.2 per cent — a spread of roughly sixteen percentage points attributable to how much the employer invested in producing and sustaining the capability rather than to the capability itself.



AI deployment intensity is strongly significant ($\beta = 0.112$, $p < 0.001$), which is theoretically important: it indicates that holding an AI skill without applying it to billable work yields substantially less than holding and deploying it. Certification without deployment is worth roughly half of certification with deployment at the sample mean. Experience enters with the expected positive linear and negative quadratic terms, confirming Mincerian concavity with an estimated turning point near 24.4 years. Employment in a global capability centre carries an independent premium of 9.0 per cent ($\beta = 0.086$, $p = 0.001$), while the Tier-1 Indian IT indicator and the metropolitan location indicator are not significant at conventional levels — the latter suggesting that within the five metropolitan markets studied, location has ceased to be an independent wage determinant. The gender coefficient is negative and significant ($\beta = -0.048$, $p = 0.030$), implying a residual differential of 4.7 per cent unexplained by skill, investment, experience, education or employer, and requiring rejection of H_6 .

Table 5 · Augmented Mincerian Earnings Function — OLS Estimates (dependent variable: ln annual CTC)

Predictor	β	Robust SE	t	p	Interpretation
Constant	12.418	0.094	132.11	< 0.001	—
AI-complementary skill (1 = yes)	0.187	0.026	7.19	< 0.001	20.6% premium at mean investment
ln Skill Development Investment	0.094	0.019	4.95	< 0.001	Independent investment effect — supports H_1
AI-skill $\times$ ln SDI (interaction)	0.061	0.021	2.90	0.004	Investment amplifies the premium — supports H_3
Experience (years)	0.078	0.008	9.75	< 0.001	Standard Mincerian return to experience
Experience ²	-0.0016	0.0004	-4.00	< 0.001	Concavity; turning point $\approx$ 24.4 years
Post-graduate qualification	0.069	0.024	2.88	0.004	7.1% return to postgraduate education
AI deployment intensity (% billable)	0.112	0.027	4.15	< 0.001	Deployment matters nearly as much as certification
Global capability centre employer	0.086	0.025	3.44	0.001	9.0% independent employer-type premium
Tier-1 Indian IT employer	0.041	0.023	1.78	0.076	Not significant at 5%
Metropolitan location (BLR / HYD)	0.033	0.021	1.57	0.117	Location no longer an independent determinant
Gender (female = 1)	-0.048	0.022	-2.18	0.030	4.7% residual differential — H_6 rejected

Note: $N = 480$. $R^2 = 0.712$; adjusted $R^2 = 0.703$; $F(11, 468) = 118.4$, $p < 0.001$. Maximum VIF = 2.31. Heteroskedasticity-robust standard errors. ln SDI is mean-centred, so β_1 is the premium at average investment.

8.3 Dose-Response and the Investment Threshold

Table 6 and Figure 3 examine the shape of the investment–premium relationship across quartiles of employer skill development investment. Mean realised premium rises from 9.4 per cent in the lowest quartile, where annual investment is below ₹8,000 per employee and released learning time averages fourteen hours, to 38.9 per cent in the highest quartile, where investment exceeds ₹35,000 and released time averages ninety-six hours. Analysis of variance confirms that the quartile means differ significantly, $F(3, 476) = 64.8$, $p < 0.001$, with an eta-squared of 0.29 indicating that investment quartile alone accounts for close to thirty per cent of premium variance. Scheffé post-hoc comparison separates all four quartiles at the five per cent level.

The relationship is, however, concave rather than linear. A logarithmic specification fits materially better than a linear one, and the marginal premium per additional rupee of investment declines across the range: the second quartile buys 10.2 percentage points over the first, the third buys 10.5 over the second, and the fourth buys only 8.8 over the third, despite the fourth quartile representing the largest absolute increment in spending. H_5 , which posited linearity, is therefore rejected.

Two practical inferences follow. First, there is a threshold near ₹18,000 per employee per year — corresponding to approximately forty released learning hours — below which premium capture is weak and unreliable. Sub-threshold investment does not produce a proportionally smaller premium; it produces very little premium at all, because the capability it generates does not reach the depth at which clients and internal staffing systems will price it. Learning-and-development budgets set below this level are, on this evidence, close to a dead loss. Second, because returns are concave, an employer facing a fixed budget maximises aggregate premium capture by lifting

a moderate number of employees decisively past the threshold rather than by distributing a token allocation across the entire workforce. Interview evidence supports this reading: L&D leaders in high-investment firms consistently described a deliberate concentration of spend on identified cohorts, while those in low-investment firms described uniform per-capita allocation and reported difficulty demonstrating any return.

Table 6 · Wage Premium by Quartile of Employer Skill Development Investment

Quartile	Annual SDI range (₹ per employee)	Mean released learning hours	n	Mean premium (%)	SD	Scheffé grouping
Q1 (lowest)	Below ₹8,000	14	118	9.4	6.8	a
Q2	₹8,000 – ₹18,000	32	121	19.6	8.1	b
Q3	₹18,001 – ₹35,000	58	124	30.1	9.4	c
Q4 (highest)	Above ₹35,000	96	117	38.9	10.7	d
All	Mean ₹22,100	50	480	24.6	13.9	—

Note: ANOVA $F(3, 476) = 64.8, p < 0.001, \eta^2 = 0.29$. Quartiles sharing no common letter differ at $p < 0.05$ (Scheffé). Marginal gains of 10.2, 10.5 and 8.8 percentage points across successive quartiles indicate concavity; a logarithmic specification outperforms the linear form.

Figure 3 · Dose-Response Between Skill Development Investment and Wage Premium (SDI quartiles; ANOVA $F(3,476) = 64.8, p < 0.001, \eta^2 = 0.29$)




Figure 3 · Concave dose-response between employer investment and realised premium, with the ₹18,000 threshold.

8.4 Decomposition of the Premium

Figure 4 decomposes the mean 31.1 percentage-point premium into four components. Scarcity rent, estimated from vacancy-to-candidate ratios by skill cluster, accounts for 11.2 percentage points or 36 per cent of the total. Productivity complementarity, estimated from realisation-rate and delivery-metric differentials reported by employers, accounts for 9.8 points or 32 per cent. Signalling and certification effects account for 5.3 points or 17 per cent, and task-portfolio reallocation — the movement of AI-skilled staff onto intrinsically higher-value work — accounts for 4.8 points or 15 per cent.

The decomposition carries the study's most consequential implication. Only about a third of the current premium rests on measurable productivity complementarity, which is the component most likely to persist. The scarcity component, more than a third of the total, is by construction transitional: it exists because demand for AI-complementary talent is growing at roughly 25 per cent annually against supply growth near 15 per cent, and it will compress as the Deloitte and NASSCOM (2024) projection of a 1.25 million-strong AI talent pool by 2027 is realised. The signalling component is similarly fragile, since certification scarcity erodes fastest of all as credential supply expands — the copilot-fluency cluster in Figure 5, with a supply index of 62 and a premium of only 16.8 per cent, offers a preview of where the scarcer clusters are heading.



For employees this means that a premium captured today should not be projected forward as an annuity. For employers it means that retention strategies priced against today's premium will over-pay within two to three years, while skill development investment justified solely on avoided scarcity-rent payments will disappoint. The durable case for investment rests on the productivity and task-mix components, which together account for 47 per cent of the current premium and are the components an employer can actually influence through deployment design.

Figure 4 · Decomposition of the AI-Complementary Skill Wage Premium




Figure 4 · Four-way decomposition of the mean 31.1 per cent premium into scarcity, productivity, signalling and task-mix components.

Figure 5 · Wage Premium by AI Skill Cluster
(supply index = certified candidates per 100 open roles; lower index = scarcer skill)




Figure 5 · Premium by AI skill cluster against supply index (certified candidates per 100 open roles).



8.5 Heterogeneity Across Skill Clusters

Figure 5 disaggregates the premium by AI skill cluster and pairs each with a supply index expressing certified candidates per hundred open roles. The ordering is close to an inverse function of supply. Generative AI and large language model engineering, the scarcest cluster with a supply index of 18, commands the highest premium at 46.3 per cent. MLOps and AI platform engineering follows at 39.7 per cent, data engineering for AI at 33.5 per cent, AI product and agent design at 31.2 per cent, and domain–AI hybrid capability in banking, financial services and healthcare at 28.6 per cent.

Two clusters depart from the scarcity ordering and are therefore diagnostic. AI governance, risk and assurance is scarce, with a supply index of 21, yet commands only 24.1 per cent — below several better-supplied clusters. Interview evidence attributes this to demand rather than supply: governance capability is currently priced as a compliance cost rather than a revenue enabler, and Indian delivery organisations have not yet built client-facing propositions around it. As the European Union AI Act and comparable regimes move into enforcement, this cluster is the most probable site of premium expansion, and it is presently under-invested relative to that prospect. Conversely, AI-assisted software development — general copilot fluency — has a supply index of 62 and the lowest premium at 16.8 per cent, confirming that a capability which becomes universal ceases to be a differentiator. Copilot fluency is rapidly becoming a condition of employment rather than a source of premium, which argues for funding it as baseline enablement rather than as premium-generating investment.

8.6 Employer-Type Effects

Figure 6 sets skill development investment against realised premium by employer type. Global capability centres invest the most, at ₹41,200 per employee annually, and capture the highest premium at 38.4 per cent. Tier-1 Indian IT services firms invest ₹27,600 and capture 29.7 per cent. Mid-tier Indian IT firms invest ₹14,300 — below the identified threshold — and capture 21.5 per cent. Welch's analysis of variance confirms significant differences across types, $F(3, 183.6) = 27.9$, $p < 0.001$, supporting H4. The gap between capability centres and mid-tier firms is 16.9 percentage points on a 2.9-fold difference in investment.

Product firms and AI start-ups break the pattern in an instructive way, capturing 34.9 per cent on investment of only ₹22,800. The excess return is not produced by training expenditure. Interviews indicate three sources: deployment intensity approaching total, since essentially all billable work is AI-enabled; selection at hire, since these firms recruit capability rather than build it; and equity-linked compensation that raises measured cost-to-company. The lesson for services firms is that the deployment channel in Figure 1 is at least as powerful as the acquisition channel, and considerably cheaper to operate. A firm that trains without restructuring its staffing and engagement model to place trained staff on AI-enabled work will systematically under-realise on its investment — which is precisely the failure mode described by mid-tier respondents.



Figure 6 · Employer-Type Variation in Skill Investment and Premium Capture
(Welch ANOVA $F(3,183.6) = 27.9, p < 0.001$)




Figure 6 · Skill development investment and realised premium by employer type.

Table 8 · Employer First-Year Return on Skill Development Investment, by Investment Cohort

Investment cohort	Mean SDI (₹ / employee)	Billing-rate uplift (%)	Attrition change (percentage points)	Replacement cost avoided (₹ / employee)	Net first-year return
Q4 (above ₹35,000)	42,600	+18.4	-5.2	1,86,000	4.4×
Q3 (₹18,001–35,000)	26,300	+12.1	-3.4	1,21,000	3.9×
Q2 (₹8,000–18,000)	13,100	+6.8	-1.1	39,000	2.4×
Q1 (below ₹8,000)	6,400	+2.3	+0.8	-28,000	0.6×

Note: Return computed as realised billing-rate uplift plus avoided replacement cost, net of investment, over the twelve months following the intervention. A negative replacement-cost figure denotes cost incurred rather than avoided. Across the upper three cohorts the retention channel contributes 38–44 per cent of total return.

8.7 The Employer's Return

Table 8 addresses the question the literature leaves open: what the investing firm recovers. Cohorts are ranked by investment quartile, with return computed as the sum of realised billing-rate uplift and avoided replacement cost, net of investment, over the first year following the training intervention.

The pattern mirrors the concavity of the employee-side relationship, but with a sharper break. The lowest-investment cohort records a net first-year return of 0.6 times — a loss — because a billing-rate uplift of 2.3 per cent is insufficient to recover even a modest per-capita outlay, and because attrition in this cohort rose by 0.8 percentage points, partial training having improved external marketability without producing internal advancement. This is the most consequential single result for practice: token training is worse than no training, since it funds a competitor's hiring pipeline. The second cohort reaches 2.4 times, the third 3.9 times, and the highest-investment cohort 4.4 times, driven by an 18.4 per cent billing-rate uplift and a 5.2 percentage-point attrition reduction worth approximately ₹1.86 lakh per employee in avoided replacement cost.



The retention channel is worth isolating. Across the upper three cohorts, avoided replacement cost contributes between 38 and 44 per cent of total return — comparable in magnitude to the billing-rate channel. This is consistent with Mallya and Sastry (2025), who found growth opportunity ranking above compensation as a retention driver for younger professionals, and it means that the business case for skill development investment cannot be evaluated on revenue uplift alone. A finance function that scores L&D proposals only against billing rates will systematically under-fund them by close to half.

8.8 Summary of Hypothesis Testing

Table 7 consolidates the hypothesis tests. Four hypotheses are supported and two are rejected. The rejections are substantively more interesting than the confirmations. The rejection of H_5 establishes that the investment–premium relationship is concave with a material threshold, which invalidates the proportional-return assumption embedded in most L&D budgeting practice. The rejection of H_6 establishes a residual gender differential of 4.7 per cent that survives controls for skill, investment, experience, education, employer type and location, and which therefore cannot be attributed to differential human capital. The finding echoes Sastry (2025a), where socio-cultural and systemic constraints were found to bind more tightly on women's economic outcomes than skill or knowledge gaps, and indicates that skill development investment alone will not close the differential.

Table 7 · Summary of Hypothesis Testing

Code	Hypothesis	Test Applied	Statistic	p	Decision
H ₁	Skill development investment exerts a positive and significant effect on compensation	OLS coefficient	$\beta = 0.094$ $t = 4.95$	< 0.001	SUPPORTED
H ₂	AI-complementary skills command a higher premium than AI-substitutable routine digital skills	Independent-samples t-test	$t(478) = 12.63$ $d = 1.15$	< 0.001	SUPPORTED
H ₃	Skill development investment positively moderates the AI-skill–wage relationship	Interaction term	$\beta = 0.061$ $t = 2.90$	0.004	SUPPORTED
H ₄	Mean realised premium differs significantly across employer types	Welch ANOVA	$F(3, 183.6) = 27.9$	< 0.001	SUPPORTED
H ₅	The investment–premium relationship is linear across the full investment range	Model comparison; ANOVA with Scheffé	Log form superior; marginal gains $10.2 \rightarrow 8.8$	< 0.001	REJECTED
H ₆	Gender does not moderate realised premium capture after full controls	OLS coefficient	$\beta = -0.048$ $t = -2.18$	0.030	REJECTED

IX. FINDINGS

- AI-complementary skill holding carries a mean wage premium of 31.1 per cent within Indian IT services, ranging from 21.6 per cent at entry level to 39.4 per cent beyond ten years of experience.
- Employer skill development investment raises compensation independently ($\beta = 0.094$, $p < 0.001$) and amplifies the AI-skill premium through a significant positive interaction ($\beta = 0.061$, $p = 0.004$). The premium is a joint product of skill and investment, not a property of skill alone.
- The premium ranges from 20.6 per cent at mean investment to 36.2 per cent at two standard deviations above it — a sixteen-point spread attributable to employer investment rather than to individual capability.
- The investment–premium relationship is concave, not linear, with a threshold near ₹18,000 per employee per year (approximately forty released learning hours) below which premium capture is weak.
- Investment quartile alone explains 29 per cent of premium variance ($\eta^2 = 0.29$), a larger share than experience or education.
- Only 32 per cent of the premium reflects measurable productivity complementarity; 36 per cent is transitional scarcity rent and 17 per cent is signalling. Roughly half the current premium should be expected to erode as supply expands toward the projected 1.25 million AI talent pool by 2027.
- Deployment intensity is independently significant ($\beta = 0.112$, $p < 0.001$). Certified but undeployed capability realises approximately half the premium of certified and deployed capability.
- Premium varies inversely with supply across skill clusters, from 46.3 per cent for generative AI engineering (supply index 18) to 16.8 per cent for general copilot fluency (supply index 62). AI governance is the notable exception — scarce yet under-priced, and the most probable site of future premium expansion.



- Global capability centres invest 2.9 times what mid-tier Indian firms invest and capture a 16.9 percentage-point higher premium; product firms achieve high premia on moderate investment through deployment intensity and selection rather than training volume.
- Employer first-year return rises from 0.6 times in the lowest investment cohort — a net loss — to 4.4 times in the highest. Sub-threshold training raises attrition by 0.8 percentage points, making token investment worse than none.
- Avoided replacement cost contributes 38 to 44 per cent of total employer return, comparable to the billing-rate channel, so appraisal on revenue uplift alone understates the business case by nearly half.
- A residual gender differential of 4.7 per cent persists after full controls, indicating that skill development investment alone will not equalise premium capture.

X. SIGNIFICANCE OF THE STUDY

- For human resource practice, the study reframes skill development investment as a produced factor of wage formation rather than a welfare or engagement expense, and supplies the threshold, functional form and return schedule needed to defend an L&D budget in financial terms.
- For finance and strategy functions, it provides a first-year return schedule by investment cohort and identifies the retention channel as roughly two-fifths of total return, correcting a systematic appraisal bias against learning investment.
- For employees and career entrants, it identifies which skill clusters currently pay, which are converging toward commodity status, and — through the deployment coefficient — establishes that securing AI-enabled billable work matters nearly as much as securing the certification itself.
- For policy, the decomposition indicates how much of the present premium is transitional, which bears directly on the calibration of national reskilling missions: interventions justified against today's premium will over-state their benefit if the scarcity component is treated as permanent.
- For the academic literature, the study contributes the first India-specific estimate of an interaction between employer training investment and AI skill in a Mincerian framework, together with a four-way premium decomposition that separates rent from productivity.

XI. RECOMMENDATIONS AND SUGGESTIONS

11.1 For Employers

- Set skill development investment above the ₹18,000 per employee threshold for targeted cohorts rather than distributing a uniform sub-threshold allocation across the workforce. Concavity of returns makes concentration the budget-maximising strategy under a fixed envelope.
- Pair every training intervention with a staffing commitment. Deployment intensity carries a coefficient nearly as large as certification itself; capability that is funded but not placed on AI-enabled billable work realises roughly half its potential return.
- Do not fund token training. The lowest investment cohort records a net first-year loss and an increase in attrition, because partial capability improves external marketability without producing internal advancement.
- Rebuild the technical career ladder at the five-to-ten-year stage, where the premium dips as professionals migrate to people-management roles. A parallel technical track with compensation parity would retain frontier capability that is currently being converted into supervisory capacity.
- Fund AI governance, risk and assurance capability now, ahead of the premium. It is presently scarce and under-priced; as enforcement regimes mature it is the cluster most likely to reprice upward, and acquiring it at today's cost is an option worth holding.
- Treat general copilot fluency as baseline enablement funded from operating cost, not as premium-generating investment. At a supply index of 62 it has ceased to differentiate.
- Include avoided replacement cost in the L&D business case. Excluding it understates return by roughly two-fifths and biases capital allocation against learning investment.

11.2 For Professionals

- Prioritise employer-funded programmes carrying a deployment commitment over self-funded certification without one; the interaction and deployment coefficients both indicate that where and how a skill is applied determines much of its realised value.



- Recognise that roughly half the present premium is transitional. Skills acquired at today's scarcity should be treated as a window for advancement rather than as a durable earnings floor, and renewed continuously in line with the obsolescence dynamics documented by Deming and Noray (2020).
- Where a management transition is available at the five-to-ten-year stage, weigh it against the technical premium foregone, which this study measures at approximately nine percentage points relative to the adjacent bands.

11.3 For Policy and Academia

- Calibrate public reskilling subsidies to the concave return curve: per-beneficiary allocations below the threshold will generate little measurable wage effect, however wide their coverage.
- Address the residual gender differential through pathways other than training volume. The gap survives full controls for skill and investment and is therefore unlikely to respond to skilling interventions alone.
- Extend national skill accounting to capture deployment as well as certification, since certification counts alone materially over-state realised capability value.

XII. LIMITATIONS OF THE STUDY

- The design is cross-sectional; the interaction between investment and skill is identified from between-individual variation, and causal interpretation, while theoretically grounded, cannot be established with the strength a panel or quasi-experimental design would permit.
- Compensation was self-reported in banded form and converted to midpoints, which introduces measurement error and may compress dispersion at the tails.
- The sample covers five metropolitan markets and therefore does not represent tier-2 delivery locations, where investment levels and premium structures are likely to differ.
- Skill development investment was reported by employees and validated against L&D manager estimates where available, but firm-level accounting data were not accessible for the full sample.
- The premium decomposition relies on an attribution procedure combining vacancy ratios, employer-reported realisation data and interview evidence, and the component boundaries — particularly between signalling and productivity — are estimates rather than direct measurements.
- AI skill definitions are evolving rapidly, and cluster boundaries valid in 2026 may not remain comparable over time, which constrains longitudinal replication.

XIII. SCOPE FOR FURTHER RESEARCH

- A panel design tracking the same individuals across two or more compensation cycles would permit identification of the causal effect of investment on premium and would allow direct measurement of premium decay as supply expands.
- Firm-level analysis linking audited L&D expenditure to realisation rates and margin would test the employer-return schedule against accounting rather than survey data.
- Extension to tier-2 delivery centres and to the business process management segment would establish whether the threshold identified here is a national constant or a metropolitan-market artefact.
- A dedicated study of the residual gender differential, using matched pairs and promotion-history data, would separate access effects from negotiation and allocation effects.
- Replication after the projected expansion of the AI talent pool toward 2027 would provide a direct test of this study's central prediction that the scarcity component of the premium is transitional.

XIV. CONCLUSION

This study set out to determine whether employer skill development investment affects the wage premium accruing to AI-complementary capability in Indian IT services, and if so with what magnitude and shape. The evidence indicates that it does, substantially, and that the relationship is more structured than current practice assumes. The premium is not a property of the skill; it is a joint product of the skill and the investment that produced and deployed it, varying from 20.6 to 36.2 per cent according to how much the employer invested. The relationship is concave with a threshold near ₹18,000 per employee per year, below which investment produces little premium and, in the lowest cohort, a net loss to the firm through increased attrition.

The decomposition qualifies the optimism that surrounds AI wage premia. Only about a third of the current 31.1 per cent premium rests on measurable productivity complementarity. More than a third is scarcity rent that will compress as India's AI talent pool expands toward the projected 1.25 million by 2027, and a further sixth is



signalling value that erodes faster still. Both employers pricing retention and policymakers appraising reskilling missions against today's premium are, on this evidence, working from a number that will not hold.

What follows for the human resource function is a reframing rather than an incremental adjustment. Skill development investment is not a welfare expenditure to be defended on engagement grounds, nor a compliance obligation to be discharged at minimum cost. It is a produced factor of wage formation with a measurable, non-linear and firm-appropriable return — 4.4 times in the first year at adequate scale, negative at token scale. The distinction between those two outcomes is not the presence of a training programme but whether the investment crosses the threshold and whether the resulting capability is placed on work that prices it. Firms that grasp this will convert the AI transition into margin; those that fund training uniformly and deploy it inattentively will fund their competitors' hiring pipelines instead.

Finally, the persistence of a 4.7 per cent gender differential after full controls is a reminder that skill investment operates within a social structure it does not by itself reform. Consistent with earlier findings on the primacy of socio-cultural over skill constraints in Karnataka's entrepreneurial economy, the evidence here suggests that equalising investment will not, on its own, equalise return. That is a separate problem, and it requires separate instruments.

REFERENCES

- [1] Acemoglu, D., & Restrepo, P. (2019). *Automation and new tasks: How technology displaces and reinstates labor*. *Journal of Economic Perspectives*, 33(2), 3–30.
- [2] Acemoglu, D., & Restrepo, P. (2022). *Tasks, automation, and the rise in U.S. wage inequality*. *Econometrica*, 90(5), 1973–2016.
- [3] Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). *The demand for AI skills in the labor market*. *Labour Economics*, 71, 102002.
- [4] Anas, M. B., Hans, V. B., & Sastry, N. S. (2025). *Digital transformation in tourism: How technology shapes travel experiences*. *International Journal of Scientific Research & Engineering Trends*, 11(5), 1–5.
- [5] Autor, D. H. (2015). *Why are there still so many jobs? The history and future of workplace automation*. *Journal of Economic Perspectives*, 29(3), 3–30.
- [6] Autor, D. H., Levy, F., & Murnane, R. J. (2003). *The skill content of recent technological change: An empirical exploration*. *Quarterly Journal of Economics*, 118(4), 1279–1333.
- [7] Becker, G. S. (1964). *Human capital: A theoretical and empirical analysis, with special reference to education*. National Bureau of Economic Research / Columbia University Press.
- [8] Bessen, J. (2019). *AI and jobs: The role of demand* (NBER Working Paper No. 24235). National Bureau of Economic Research.
- [9] Brynjolfsson, E., Li, D., & Raymond, L. (2025). *Generative AI at work*. *Quarterly Journal of Economics*, 140(2), 889–942.
- [10] Deloitte & NASSCOM. (2024). *Bridging the AI talent gap to boost India's tech and economic impact*. Deloitte India.
- [11] Deming, D. J., & Noray, K. (2020). *Earnings dynamics, changing job skills, and STEM careers*. *Quarterly Journal of Economics*, 135(4), 1965–2005.
- [12] Ernst & Young. (2026). *Future of Pay 2026: India Inc. compensation and workforce trends*. EY India.
- [13] Felten, E., Raj, M., & Seamans, R. (2021). *Occupational, industry and geographic exposure to artificial intelligence*. *Strategic Management Journal*, 42(12), 2195–2217.
- [14] Goldin, C., & Katz, L. F. (2008). *The race between education and technology*. Harvard University Press.
- [15] Heckman, J. J., Lochner, L. J., & Todd, P. E. (2006). *Earnings functions, rates of return and treatment effects: The Mincer equation and beyond*. In *Handbook of the Economics of Education* (Vol. 1, pp. 307–458). Elsevier.
- [16] Lazear, E. P. (2009). *Firm-specific human capital: A skill-weights approach*. *Journal of Political Economy*, 117(5), 914–940.
- [17] Mallya, M. M., & Sastry, N. S. (2025). *Building future-ready organizations: Leadership strategies for young talent management*. *EPRA International Journal of Research and Development*, 10(9), 88–91.
- [18] Mincer, J. (1974). *Schooling, experience, and earnings*. National Bureau of Economic Research / Columbia University Press.
- [19] NASSCOM. (2025). *Strategic Review 2025: The technology industry in India*. National Association of Software and Service Companies.
- [20] NASSCOM & Deloitte. (2025). *India's AI talent inflection point: From skill gaps to competitive advantage*. NASSCOM Insights.
- [21] 1 Finance Research. (2026). *AI versus traditional IT compensation in India: An analysis of 950 technology job postings*. 1 Finance.
- [22] PwC. (2025). *Global AI Jobs Barometer*. PricewaterhouseCoopers.
- [23] Sastry, N. S. (2013). *HR approaches to talent management*. *International Journal of Managerial Studies and Research*, 1(2).



- [24] Sastry, N. S. (2020). A study on impact of employees' performance that influence to motivate in tea manufacturing organizational sector. *EPRA International Journal of Economic and Business Review*, 8(6).
- [25] Sastry, N. S. (2025a). Understanding the challenges of women micro-entrepreneurs in Karnataka. *International Journal of Managerial Studies and Research*, 13(8), 12–16. <https://doi.org/10.20431/2349-0349.1308002>
- [26] Sastry, N. S. (2025b). Talent designation and social comparison: Economic drivers of organizational growth – An empirical analysis. *EPRA International Journal of Economics, Business and Management Studies*, 12(10).
- [27] Sastry, N. S., & Mallya, M. M. (2025a). Revolutionizing HR hiring with AI analytics: Enhancing efficiency and attracting top talent. *EPRA International Journal of Research and Development*, 10(9).
- [28] Sastry, N. S., & Mallya, M. M. (2025b). Strategic HRM for economic resilience: Skill diversification, labor market dynamics, and technology disruption. *EPRA International Journal of Research and Development*, 10(11). <https://doi.org/10.36713/epra2016>
- [29] Schultz, T. W. (1961). Investment in human capital. *American Economic Review*, 51(1), 1–17.
- [30] Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.
- [31] Vermeulen, B., Kesselhut, J., Pyka, A., & Saviotti, P. P. (2018). The impact of automation on employment: Just the usual structural change? *Sustainability*, 10(5), 1661.
- [32] Wheebox. (2026). India Skills Report 2026. Wheebox, Confederation of Indian Industry and All India Council for Technical Education.