

Digital Transformation and Economic Inequality in the IT Workforce: Task Displacement, Firm Rents, and the Narrowing of Entry Pathways

AUTHOR:- Mrs.SHAILA ASHOK RANE

RESEARCH STUDENT

BHARATI VIDYAPEETH DEEMED UNIVERSITY, PUNE

Email Id: Shailarane780@gmail.Com

UNDER THE SUPERVISION OF

Dr.VINOD H. MANE (MA, M.Phil, G.D.C.& A., Ph.D.)

Abstract

Research on technology and inequality has generally treated information technology workers as beneficiaries of skill-biased technological change rather than as a population within which inequality is itself being restructured. This paper argues that the assumption no longer holds, and supports the argument with a documentary qualitative analysis of thirty-four publicly available sources combined with descriptive quantitative evidence. We develop a framework in which IT work is a bundle of separable tasks priced by heterogeneous employers, and derive four propositions: automation of routine implementation widens dispersion within rather than between IT occupations; a growing share of that dispersion is attributable to employer identity rather than worker skill; geographic convergence at senior levels coexists with divergence at junior levels; and entry-level compression converts a distributional problem into a cohort problem. Thematic analysis of the source corpus yields six themes, of which three (task unbundling, broken apprenticeship, and attribution ambiguity) are salient across every source category while two (rent asymmetry and contractual distance) appear almost exclusively in academic and industry sources respectively, indicating a systematic blind spot in how the phenomenon is discussed. Descriptive evidence shows generative AI adoption among developers rising to 84 percent while trust in output accuracy fell to 29 percent, alongside relative employment declines of 9 to 20 percent for early-career workers in exposed roles with no corresponding senior effect. We specify an empirical design capable of testing the propositions, state falsification conditions, and identify the attribution problem posed by the coincidence of AI diffusion with a monetary tightening cycle that produced over 700,000 tech layoffs between 2022 and 2025. The framework carries an uncomfortable implication: interventions raising individual skill will not address inequality generated by employer identity or by the accident of entry timing.

Keywords: Digital transformation; Economic inequality; IT workforce; Task displacement; Generative AI; Firm wage premia; Qualitative document analysis; Labour market entry

1. Introduction

For most of the period in which economists have studied skill-biased technological change, information technology workers occupied an odd position in the argument. They were the mechanism rather than the subject. Computerization replaced basic clerical and production jobs. It increased the need for analytical and interactive work that remained. The people who created the computers ended up on top of the new job distribution (Autor, Levy, & Murnane, 2003; Bresnahan, Brynjolfsson, & Hitt, 2002). Technology causes inequality in other jobs. The IT workforce benefited the most. That perspective has become challenging to maintain. Generative AI now handles a significant portion of the routine implementation tasks that have traditionally occupied the lower levels of the software career ladder. Product-

market rents in digital industries have concentrated in a small set of firms, and those firms pay differently from everyone else for nominally identical work. Offshore capability centres, which for two decades supplied cost arbitrage on peripheral tasks, have moved into core product engineering. Each development redistributes income inside the IT workforce. None is visible in analyses that treat software development as a single occupation with a single wage.

This matters more than it might appear, because within-occupation dispersion is where most of the action in wage inequality actually is. Kim and Sakamoto (2008) found that the majority of the increase in United States wage inequality between 1983 and 2002 occurred within detailed occupations rather than between them, and van der Velde (2020) reports that within-occupation dispersion accounts for roughly half of total wage inequality across European labour markets, with the widest internal dispersion in non-routine intensive occupations. Software development is about as non-routine intensive as occupations get. If we want to understand how digital transformation distributes its returns, the interesting variance sits inside the occupational category, not across it.

Three separate literatures speak to this and none speaks to the others. Work on generative AI and employment has established that early-career workers in exposed occupations have lost ground (Brynjolfsson, Chandar, & Chen, 2025; Hosseini & Lichtinger, 2025). Work on firm wage premia has established that employer identity explains a rising share of earnings variation (Card, Cardoso, Heining, & Kline, 2018; Song, Price, Guvenen, Bloom, & von Wachter, 2019). Work on superstar firms has established that concentration and labour share decline move together (Autor, Dorn, Katz, Patterson, & Van Reenen, 2020). Three identification strategies, three sets of findings, and no analysis of how the forces interact inside a single occupational labour market where all of them operate at once.

This paper makes four key contributions. First, it establishes a framework that considers IT work as a collection of separable tasks, priced by various employers. This approach enables the integration of automation and firm-rents channels as components of a single process rather than as competing explanations.

Third, it performs a thematic analysis of thirty-four publicly available sources, revealing a pattern that a quantitative synthesis would overlook: the mechanisms are discussed in separate, non-overlapping literatures, and this fragmentation contributes to the phenomenon remaining poorly measured. Fourth, it compiles the available descriptive evidence and outlines the design required for a proper test. This paper does not provide new causal estimates. The qualitative analysis is documentary rather than based on interviews, and the quantitative material is secondary.

2. Theoretical framework and propositions

The framework consists of three components: a task structure, an employer distribution, and a career pathway. Existing research typically addresses these elements separately. Task-based models of automation (Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2022) clarify which types of work are displaced but offer limited. Firm-based models of wage inequality (Card et al., 2018) explain why identical workers earn different amounts at various employers, but they assume that the content of tasks is fixed. Additionally, they do not address the sequencing problem: a worker's position largely depends on when they join the company. Figure 1 sets out how the components combine.




Figure 1. Conceptual framework linking four drivers of digital transformation to within-occupation earnings dispersion in the IT workforce. Solid arrows denote direct effects; curved arrows denote the two amplification pathways discussed in Section 2.5. Schematic; not derived from estimated parameters.

2.1 Task displacement inside a single occupation

Software development is not one activity. It bundles requirements elicitation, architectural decision-making, specification-to-code translation, testing, debugging, code review, deployment, incident response, and stakeholder negotiation. These tasks differ enormously in exposure to current models. Translating a well-specified requirement into working code in a common framework is close to fully automatable and in many organisations already largely automated. Deciding which requirement is worth specifying, arbitrating between conflicting stakeholder demands, or diagnosing a production failure whose symptoms match no documented pattern remains stubbornly manual.

Classical logic (Acemoglu & Restrepo, 2022) predicts that the wage effect depends on the elasticity of substitution between automated tasks and those that remain. When these tasks are complementary, automating the first increases the marginal product of the second. That is the optimistic reading, and the productivity experiments broadly support it: Peng, Kalliamvakou, Cihon, and Demirel (2023) found substantial completion-time reductions among developers using AI assistance, and Brynjolfsson, Li, and Raymond (2023) found large gains among customer support agents concentrated in the least experienced.

Complementarity is a characteristic of task pairs, not of entire occupations. Two developers with identical titles and nominal seniority may have significantly different task bundles. For instance, a developer whose work consists of 70 percent specification-to-code translation and another whose role involves 70 percent architectural arbitration. Automation raises the second one's wage and undercuts the first one's, and both effects vanish in the occupational mean.

There is a further complication: the optimistic reading tends to skip. Productivity gains accrue disproportionately to novices, which sounds egalitarian and is not. If a tool raises novice output toward expert output, it compresses the productivity distribution while leaving the wage distribution to be set by bargaining. Whether compression in output produces compression in pay depends on who captures the surplus, which is a question about employers rather than tasks. This is where the task framework alone stops being sufficient.

P1. *Generative AI diffusion widens wage dispersion within detailed IT occupations rather than between them, with the widening concentrated among workers whose task bundles are heavy in routine implementation and light in architectural or interpersonal content.*

2.2 Firm rents and the between-employer channel

A substantial body of evidence establishes that firms pay systematically different wages to observationally identical workers, that this employer component explains a growing share of earnings variance, and that its growth concentrates in the industries where product-market concentration has risen fastest (Autor et al., 2020; Card et al., 2018; Song et al., 2019). Digital industries are the paradigm case. Network effects and near-zero marginal costs of software reproduction generate rents; rents are shared, incompletely and unequally, with employees.

The consequence is that identical technical work is priced differently depending on where it sits in the industrial structure. An engineer maintaining an internal payments service at a regional insurer and an engineer maintaining an internal payments service at a large platform firm perform tasks a content classifier would code identically. Their compensation is not close, and the gap is not primarily explained by measured skill. Much of it is equity, which introduces a second issue: equity-heavy compensation ties a worker's income to firm-level rent capture, increasing exposure to the very concentration dynamics that generated the rents. Workers at the top of the IT distribution are not merely paid more; they are paid in an instrument whose value correlates with continued concentration.

This channel interacts with the first in a way that is easy to miss. If generative AI compresses the productivity gap between a strong and a weak implementer, the residual determinant of pay becomes even more heavily weighted toward employer identity. Automation of task-level differentiation does not eliminate wage dispersion. It relocates its source from the worker to the firm.

P2. *As task-level productivity differences compress, the share of within-occupation earnings variance attributable to employer fixed effects rises, and this shift is larger in labour markets with greater product-market concentration.*

2.3 Spatial reconfiguration and the compression of arbitrage

For roughly two decades the standard account of IT offshoring was labour cost arbitrage on separable, well-specified, peripheral tasks (Blinder, 2006). That arrangement is now dissolving in two directions at once.

At the senior end, offshore capability centres have moved into work the arbitrage model said would stay onshore. Industry sources place the number of such centres in India above two thousand as of mid-2026, employing well over two million professionals, with NASSCOM projecting continued growth through the decade. These organisations increasingly own product mandates rather than executing specifications written elsewhere. Where that happens, the wage gap for equivalent senior work narrows, because the firm is competing for a globally scarce input rather than purchasing a locally abundant one.

At the junior end, the opposite occurs. Generative AI attacks precisely the well-specified, separable, routine work that made cost arbitrage viable in the first place. The commodity tier faces substitution not by onshore labour but by capital. The result is a squeeze from both sides: senior offshore roles converge upward toward global rates while junior offshore roles face the same automation pressure as junior onshore roles, without the institutional cushions that sometimes attend the latter.

A related shift concerns contractual form rather than location. Weil's (2014) account of the fissured workplace describes the systematic relocation of work outside the boundaries of the firms that ultimately control it, and IT labour has proven unusually amenable to this. Contractor arrangements, vendor staffing, and platform freelancing all place technical workers outside the rent-sharing perimeter described in Section 2.2. A worker can perform core engineering work for a highly profitable firm and receive none of the firm wage premium, because the employment relationship runs through an intermediary. Contractual distance from the rent-generating entity may prove a stronger predictor of IT earnings than either skill or location, and it is almost entirely absent from the digital transformation literature.

P3. *Wage convergence between offshore and onshore IT labour is concentrated at senior levels and in wholly owned capability centres, while junior and vendor-mediated roles diverge; contractual distance from the rent-generating firm predicts earnings independently of task content and location.*

2.4 Entry pathways and the cohort problem

Career progression in software has historically run through a specific sequence. A junior developer is assigned bounded, well-specified implementation work; performs it under review; accumulates contextual knowledge about the system, the organisation, and the domain; and gradually takes on the ambiguous work that cannot be specified in advance. The routine tier is not merely the bottom of the wage distribution. It is the training apparatus.

Automating that tier eliminates the apparatus along with the jobs. This is what distinguishes the current episode from earlier waves of software automation, which removed toil without removing the entry point. Brynjolfsson, Chandar, and Chen (2025), using administrative payroll microdata covering millions of workers, document substantial relative employment declines for early-career workers in highly AI-exposed occupations with no corresponding effect on senior workers in the same firms. Hosseini and Lichtinger (2025) find that firms adopting generative AI reduced junior headcount relative to non-adopters while leaving senior headcount unaffected. Different data, different designs, compatible conclusions.

The distributional consequence is that inequality acquires a cohort dimension. If the entry rate falls for several years, the affected cohort carries a permanent deficit in accumulated firm-specific and domain-specific knowledge, and the scarring literature suggests such deficits persist long after conditions improve. Meanwhile, the existing senior tier faces reduced competitive pressure from below, raising its rents. Inequality between cohorts widens even if inequality within each cohort does not.

A supply-side response complicates this in a direction unfavourable to the framework. If entry becomes visibly harder, enrolment in computing programmes may fall, eventually tightening junior supply and possibly restoring or overshooting prior wage levels for whoever remains. The framework as stated has no dynamic supply response and would need one before its long-run predictions could be taken seriously.

P4. Reduced entry-level hiring in AI-exposed IT roles produces persistent cohort-level earnings deficits not eliminated by subsequent labour market recovery, and simultaneously raises the rents accruing to incumbent senior workers.

2.5 How the channels interact

The four propositions are not additive, and treating them as a list understates the problem. P4 amplifies P2, because rationing entry raises the value of incumbency at high-premium employers and thickens the sorting advantage of whoever got in first. P1 amplifies P2 through the compression mechanism already described. P3 interacts with P4 in a way that may partially offset the others, since a global reduction in entry-level positions concentrates the remaining ones wherever the marginal cost of a junior hire is lowest, which is not obviously the high-wage locations.

The practical implication is that any design estimating one channel while ignoring the others will attribute to that channel effects belonging to the others. This is not a general caution about omitted variables. It is a specific claim that these mechanisms are jointly determined and that single-channel estimates in this literature are biased in a known direction, namely toward overstating the channel under study.

3. Methods

The study combines a documentary qualitative analysis with a descriptive quantitative synthesis. Neither component involves primary data collection from human participants, and no interviews were conducted. The design is therefore best described as a mixed-method evidence synthesis rather than as mixed-method fieldwork, and its conclusions are correspondingly bounded.

3.1 Qualitative component: corpus and coding

Document analysis is an established qualitative method in which published texts are treated as data rather than as background (Bowen, 2009). It is appropriate here because the phenomenon of interest, the distribution of returns within a fast-moving occupational labour market, is being actively narrated by academics, employers, industry bodies, and workers themselves, and the differences between those narrations carry information.

A purposive corpus of thirty-four sources published between 2020 and 2026 was assembled, stratified across five categories to capture perspectives with different interests and evidentiary standards. Table 1 sets out the composition. Inclusion required that the source make a substantive claim about the distribution of work, pay, or opportunity within technology occupations; sources addressing technology's effect on other occupations were excluded.

Coding followed reflexive thematic analysis in the tradition of Braun and Clarke (2006), proceeding through familiarisation, initial code generation, theme construction, and review. Coding was inductive at the first pass and then checked against the framework in Section 2, which means the themes are not independent of the theory and should not be read as confirming it. Six themes were retained. Salience was recorded on a four-point ordinal scale (absent, low, moderate, high) reflecting both frequency and centrality within each source category.

Two limitations of the coding are worth stating at the outset. The corpus was assembled and coded by a single analyst, so no inter-rater reliability statistic is available. And the category boundaries are porous: several working papers were authored by researchers employed at firms with commercial interests in AI adoption, which places them ambiguously between the second and fourth categories.

Table 1. Composition of the qualitative corpus (n = 34)

Source category	n	Examples of source type	Interest position
Peer-reviewed articles	9	Journal articles on task content, firm wage premia, labour share	Disciplinary; incentives toward novelty and identification
Working papers and preprints	7	Payroll microdata studies, vacancy-based analyses of AI exposure	As above; several authored with firm affiliation
Practitioner surveys	4	Large-scale developer surveys, graduate outlook surveys	Self-selected respondents; platform-recruited samples
Industry and consultancy reports	8	Capability centre census data, compensation benchmarking, talent reports	Commercial interest in the trends reported
Trade and business press	6	Layoff trackers, sector coverage of hiring conditions	Attention incentives; event-driven coverage

Note: Purposive rather than systematic sampling. The corpus is not exhaustive and was not assembled through a pre-registered search protocol, which limits claims about thematic saturation.

3.2 Quantitative component

The quantitative material in Section 5 is descriptive and drawn entirely from secondary sources. Figures are reproduced as reported by the originating study rather than recomputed, and estimates from different studies are not pooled, because their designs, populations, and comparison groups differ enough that a pooled figure would be misleading. Where sources disagree, the disagreement is reported rather than resolved.

4. Qualitative findings

Six themes were constructed from the corpus. Figure 2 shows their salience across source categories, and Table 2 sets out the coding framework with representative evidence.




Figure 2. Theme salience across source categories. Cell values reflect the author's ordinal coding of frequency and centrality, not counts of occurrence. Note the non-overlap of T2 and T5: rent asymmetry is discussed almost exclusively in peer-reviewed work, contractual distance almost exclusively in industry sources.

Table 2. Thematic coding framework

Theme	Definition	Representative evidence	Relation to framework
T1 Task unbundling	Recognition that the job title has become a poor proxy for the work performed, and that exposure varies within roles	Distinction between AI-assisted implementation and unassisted judgement work recurs across academic and survey sources	Supports P1; supplies the mechanism
T2 Rent asymmetry	Attribution of pay differences to employer characteristics rather than worker characteristics	Firm wage premium and labour share literatures; largely absent from practitioner discourse	Directly P2
T3 Broken apprenticeship	Concern that automating junior work removes the mechanism by which seniors are produced	Strongly present in trade press and surveys; framed as a pipeline problem rather than a distributional one	Directly P4
T4 Verification burden	New unpriced work created by the need to check AI output	Survey findings on debugging time and accuracy concerns	Complicates P1; suggests task substitution is partial
T5 Contractual distance	Location of technical labour outside the employing firm's rent perimeter	Contractor share data; capability centre versus vendor distinctions	Directly P3

Theme	Definition	Representative evidence	Relation to framework
T6 Attribution ambiguity	Uncertainty about whether observed hiring changes reflect AI or macroeconomic conditions	Explicit in working papers; contested in trade press	Threat to identification of all four

Note: Themes constructed inductively then reviewed against the Section 2 framework, so the alignment in the final column is partly a product of the coding procedure and should not be treated as independent corroboration.

4.1 The fragmentation finding

The most informative result is not any single theme but the pattern of their distribution. Two themes are effectively siloed. Rent asymmetry (T2) appears with high salience in peer-reviewed sources and is absent from practitioner surveys entirely: developers discussing their own pay in survey free-text discuss skills, tools, and market conditions, and rarely the employer's position in the product market. Contractual distance (T5) shows the mirror pattern, appearing with high salience in industry and consultancy sources, which track contractor ratios closely because they sell into that market, and only moderately in academic work.

This matters for measurement rather than only for rhetoric. The two mechanisms most likely to determine who captures the returns from digital transformation are each visible to only one community, and neither community is positioned to measure the other's variable. Academics have the identification tools but not the contractor data; consultancies have the contractor data but no interest in a research design that could produce an unflattering result. The fragmentation is a plausible partial explanation for why the phenomenon remains poorly quantified after several years of intense attention.

A second observation concerns theme T4. Verification burden was among the most salient themes in practitioner sources and among the least salient in academic ones, and it complicates the automation story in a specific way. If a substantial share of AI-generated output requires checking, then the technology has not removed the junior task so much as converted it from production into review, which is a different and less codified skill that novices are worse positioned to supply. Whether this represents transitional friction or a durable feature is not determinable from the corpus.

4.2 Three organisational archetypes

Synthesising across the corpus produces three recurrent organisational forms in which IT work is performed, distinguished by their relationship to product-market rents rather than by their technical activity. These are analytic constructs assembled from published descriptions, not empirical cases, and should be read accordingly.

The platform firm generates rents directly, shares them through equity-heavy compensation, and has automated junior implementation aggressively while retaining senior judgement roles. The capability centre performs technically comparable work inside the same corporate boundary at lower cost, with an internal split between product-owning teams converging toward global rates and delivery teams facing direct substitution. The vendor or intermediated arrangement supplies technical labour to rent-generating firms without conferring access to the rents, and concentrates exactly the task profile most exposed to automation. Figure 3 positions these and related roles on the two dimensions the framework identifies as decisive.




Figure 3. Positioning of IT worker archetypes by proximity to the rent-generating entity and by the share of the task bundle resistant to automation. Positions are analytic judgements derived from the corpus, not measured coordinates. The framework predicts that inequality growth is steepest along the diagonal from lower left to upper right.

The positioning makes a prediction the framework would not otherwise generate explicitly. The platform firm junior engineer occupies an unstable position: high proximity to rents, low task protection. Whether that role converges toward the senior position through retention and training or toward the vendor position through substitution is the single most consequential open question for cohort inequality, and it is a question about firm policy rather than about technology.

5. Descriptive quantitative evidence

This section assembles the available figures. They are presented as motivation and as constraints on plausible interpretation, not as tests. Table 3 summarises the key facts with sources and reliability assessments.

5.1 Adoption without confidence

The Stack Overflow Developer Survey, fielded in mid-2025 with 49,009 respondents across more than 160 countries, reports that 84 percent of developers use or plan to use AI tools, up from 76 percent the previous year. Over the same period the share trusting the accuracy of AI output fell from 40 percent to 29 percent, and favourable sentiment fell from above 70 percent to 60 percent. Among respondents, 46 percent actively distrust output accuracy against 33 percent who trust it, and only 3 percent report high trust.




Figure 4. Divergence between adoption and confidence among developers, 2023-2025. Source: Stack Overflow Developer Survey 2025 (n = 49,009). Favourable sentiment for 2023 and 2024 is plotted at 70 percent, the lower bound of the reported figure. Sample is recruited through the platform's own channels and over-represents engaged users.

Two features of this pattern bear on the framework. First, the experience gradient runs opposite to what a simple adoption story predicts: early-career developers report the highest daily use while experienced developers report the lowest use and the highest distrust, with only 2.6 percent of experienced respondents highly trusting output against 20 percent highly distrusting it. The workers whose task bundles are most exposed are also the heaviest adopters, which is consistent with substitution of their own labour rather than augmentation of it.

Second, 66 percent report AI answers that are close but not correct, and 45 percent say debugging AI-generated code takes longer than writing it themselves. This is theme T4 in quantitative form, and it undermines any straightforward reading of the productivity experiments. A tool that produces near-correct output shifts work from production to verification, and verification is the harder skill.

5.2 The cohort divergence

Estimates of the early-career effect come from several designs and are summarised in Figure 5. Payroll microdata analyses report relative employment declines on the order of 13 percent for workers aged 22 to 25 in the most AI-exposed occupations since late 2022, with a steeper figure of roughly 20 percent for software developers specifically, while employment for older workers in the same firms grew. Firm-level comparisons find junior headcount reductions of about 9 percent among adopting firms relative to non-adopters, with senior headcount unaffected.




Figure 5. Relative employment changes by career stage in AI-exposed roles. Sources: Brynjolfsson, Chandar, and Chen (2025) for the age-based estimates; Hosseini and Lichtinger (2025) for the firm-adoption comparison. Estimates derive from different populations, comparison groups, and time windows and are not directly comparable; they are plotted together to show consistency of direction, not magnitude.

Convergent direction across independent designs is meaningful evidence. It is not, however, evidence about magnitude, and the temptation to average these figures should be resisted. A reader who takes away that early-career employment in exposed roles fell by somewhere between 9 and 20 percent relative to a comparison group has extracted about as much as the evidence supports.

5.3 The attribution problem

The central obstacle to interpreting any of this is that generative AI diffusion began within months of a monetary tightening cycle and a sharp correction in technology firm valuations. Figure 6 shows the layoff series against both events.




Figure 6. Tech sector layoffs by year, 2020-2025, in thousands. Source: Layoffs.fyi as reported by trade press. Figures for 2025 vary between roughly 90,000 and 130,000 across reporting outlets depending on coverage and cut-off date; the higher figure is plotted. The 2023 peak of approximately 263,000 predates meaningful production deployment of coding assistants at scale.

The sequencing is unhelpful for the automation hypothesis. The layoff peak occurs in 2023, before generative coding assistants had reached the deployment levels documented in Section 5.1, and the series declines through 2024 and 2025 even as adoption rises. That pattern is more consistent with a pandemic over-hiring correction amplified by interest rates than with technological displacement. What the layoff series cannot show, and what the microdata studies do show, is the composition of hiring rather than separations: a firm can hold headcount flat while shifting its intake from junior to senior, and that shift is where the framework locates the effect.

Both stories therefore remain live, and the honest summary is that the published literature has not yet separated them. Theme T6 registered this ambiguity explicitly in working papers and inconsistently elsewhere, with trade coverage frequently attributing to AI what could equally be attributed to the rate environment. Any credible design must address this directly, and Section 6.3 specifies how.

Table 3. Summary of key facts and figures

Indicator	Value	Source	Reliability
Developers using or planning to use AI tools, 2025	84% (from 76% in 2024)	Stack Overflow Developer Survey 2025, n = 49,009	Moderate; self-selected sample
Developers trusting accuracy of AI output, 2025	29% (from 40% in 2024)	As above	Moderate; as above
Developers reporting debugging AI code takes longer than writing it	45%	As above	Moderate; as above
Relative employment decline, age 22-25, most AI-exposed occupations	approx. 13%	Brynjolfsson, Chandar, and Chen (2025), payroll microdata	High; administrative data
Relative junior headcount decline at GenAI-adopting firms	approx. 9%	Hosseini and Lichtinger (2025)	Moderate; working paper, resume-based
Within-occupation share of total wage inequality, Europe	approx. 50%	van der Velde (2020), EU-SES	High; peer-reviewed
Tech layoffs, 2023 peak	approx. 263,000	Layoffs.fyi via trade press	Low to moderate; inconsistent coverage
Capability centres in India, mid-2026	2,100+ centres, 2.3m staff	Industry and NASSCOM-derived reporting	Low; commercially interested sources
Women's share of US technology occupations	approx. 25%	CompTIA State of the Tech Workforce	Moderate; definitional variation across sources

Note: Reliability assessments are the author's, based on data provenance, transparency of method, and the interest position of the originating organisation. Values are reported as published and are not recomputed or pooled.

5.4 Distributional composition

The channels described above do not act on a demographically neutral population. Industry workforce surveys, including CompTIA's, place women at roughly a quarter of technology occupations in the United States, with lower shares in several core engineering specialisms and higher shares in design-adjacent roles; sources vary between about 22 and 28

percent depending on how technology occupations are defined, which is a wide enough band to warrant caution about any specific figure. If routine implementation work is disproportionately performed at the entry level, and if entry-level access is itself unequally distributed, then automating that tier has consequences along dimensions a task-content analysis will not register. This paper does not model that interaction, which is a limitation rather than a judgement about its importance.

6. A proposed empirical design

Testing the framework requires linking three things rarely linked: worker earnings histories, employer identity, and task-level job content. This section specifies what such a design would look like and where it currently breaks down.

6.1 Constructs and measures

Table 4. Constructs, proposed measures, and candidate data sources

Construct	Symbol	Proposed measure	Candidate source
Within-role wage dispersion	D	Log 90-10 earnings ratio within detailed occupation, seniority band, and metropolitan area	Matched employer-employee administrative records; OEWS for cross-checks
Routine implementation share	R	Share of posted tasks classifiable as specification-to-code translation, test scaffolding, or boilerplate maintenance	Vacancy text coded against O*NET task statements
Firm wage premium	F	Employer fixed effect from a two-way fixed-effects earnings decomposition	Longitudinal employer-employee panel
Generative AI exposure	A	Task-weighted exposure index validated against observed tool adoption rather than assumed from task text	Published exposure indices; firm licence or repository telemetry
Entry-cohort share	E	Share of firm hires in a year with under two years of prior occupational experience	Employer-employee panel; professional network data as substitute
Offshore capability intensity	G	Share of engineering headcount in wholly owned offshore centres, distinguished from vendor contracting	Firm disclosures; industry census data
Contractual distance	C	Share of technical labour engaged as contractors, vendor staff, or freelancers rather than employees	Firm filings; contingent worker supplements

Note: Measures are proposed rather than implemented. Three of the seven are not currently measurable in public data, as discussed in Section 6.4.

6.2 Identification

For P1 and P4, a staggered difference-in-differences design exploiting variation in the timing of firm-level tool adoption, with later adopters as controls for earlier ones. The identifying assumption is that adoption timing is uncorrelated with unobserved determinants of the junior hiring rate, which is doubtful; firms adopting early are plausibly firms already inclined toward headcount reduction. Instrumenting adoption with vendor rollout schedules or regional licence availability offers a partial remedy.

For P2, a two-way fixed-effects earnings decomposition estimated separately in pre- and post-diffusion windows, with the change in the employer variance share as the object of interest. Limited mobility bias is the principal concern, and IT labour markets have high enough job mobility that it binds less tightly than in many settings.

For P3, comparison of within-firm wage gaps between onshore and offshore locations for matched roles, distinguishing wholly owned centres from vendor arrangements. This requires firm cooperation and will be feasible only for a non-random set of participants, which limits external validity in a direction that probably understates junior-end divergence, since firms willing to open their compensation data are unlikely to be the worst offenders.

6.3 Separating automation from the rate cycle

Given the sequencing shown in Figure 6, no design should be credited that lacks an explicit strategy here. Three are available. Within-firm comparison across occupations of differing exposure holds the firm's financing conditions fixed, so differential junior contraction in exposed relative to unexposed roles within the same firm cannot be attributed to the rate environment. Comparison of firms with differing capital structures separates the interest-sensitivity channel, since cash-rich incumbents faced a weaker rate shock than venture-dependent firms while facing the same technology. And composition rather than level should be the outcome: the ratio of junior to senior hiring is informative even in periods when both are falling.

6.4 Falsification and feasibility

Stating falsification conditions matters more than usual here, because a framework with four channels and interaction effects can accommodate a great deal of evidence if it is not disciplined in advance. If within-occupation dispersion narrowed following diffusion while between-occupation dispersion widened, P1 fails and the classical account is the better description. If the employer share of variance in IT occupations stayed flat while concentration rose, P2 fails. If offshore convergence appeared uniformly across seniority levels, P3 fails. If entry-cohort deficits close within three to four years of recovery, P4 fails in its strong form.

On feasibility, three of the seven constructs in Table 4 are not currently measurable in public data. Realised task allocation, firm-level tool adoption, and the contractor share of technical labour within specific firms each require either firm cooperation or new statistical collection. That is a substantive obstacle rather than a technical footnote. It means the questions this paper poses cannot at present be answered from outside the firms whose behaviour is at issue, and that most published estimates rely on proxies whose measurement error is plausibly correlated with the treatment of interest.

7. Implications

If the framework is broadly correct, several widely recommended responses are aimed at the wrong target.

Reskilling programmes address the task channel and nothing else. They assume that inequality among IT workers reflects differences in what workers can do, and where P2 holds, it substantially does not. A worker who acquires architectural skills while remaining at a low-premium employer, or outside the employment perimeter of the rent-generating firm, captures a fraction of the return an identical worker inside that perimeter captures. This does not make training useless. It means training is a partial instrument being asked to do a job that also requires attention to employment structure, contractual form, and labour market access.

For firms, the entry-level finding poses a problem individual firms cannot solve. If routine implementation is the training apparatus, each firm's rational decision to automate it degrades a pipeline all firms draw on. The senior engineers firms will want in 2032 are the junior engineers nobody is hiring in 2026. This has the structure of a common problem, and common problems are not usually resolved by participants noticing them. The verification burden documented in Section 5.1 offers one partial route: if checking AI output is now a substantial share of the work, structured review responsibility is a plausible replacement apprenticeship task, though it demands judgement that the old sequence built rather than assumed.

For educators, the implication is uncomfortable in a different way. Curricula optimised for producing competent implementers are optimised for the segment under most direct pressure. Shifting toward system design, requirement

elicitation, and diagnostic reasoning is the obvious response, but these are precisely the capabilities hardest to teach without the experience base that entry-level work used to supply.

For policymakers, the framework points toward labour market structure rather than skill supply. Portable benefits and classification rules affect who sits inside the rent-sharing perimeter. Pay transparency acts directly on the employer-premium channel by making firm-level differences legible. Support for early-career entry addresses a cohort problem markets have no mechanism to correct retrospectively. None of these instruments is new. The argument concerns which lever the evidence points toward, and it points away from the one current policy attention favours.

8. Conclusion

The paper reports no causal estimates. The qualitative component is documentary rather than interview-based, which means it captures how the phenomenon is discussed rather than how it is experienced; a study drawing on interviews with junior developers, hiring managers, and capability centre leads would test the archetypes in Section 4.2 far more directly. The corpus was assembled purposefully and coded by one analyst without inter-rater checks, so the theme salience recorded in Figure 2 should be read as a structured reading rather than as a measurement. The framework has no dynamic supply response. If computing enrolment falls in response to visible entry difficulty, junior supply tightens and the long-run predictions change, possibly substantially. Building an equilibrium version with endogenous entry is the most important extension.

The geographic analysis rests on industry sources with commercial interests in the trends they report. Consultancies selling capability centre services and staffing platforms selling contractor placement are not disinterested observers of capability centre growth or contractor share. Academic replication using firm filings and national statistics is needed before P3 should carry much weight. The distributional composition of the affected workforce is described but not modelled. Gender, ethnicity, disability, institutional pedigree, and immigration status all shape access to entry-level positions and to high-premium employers, and a framework about who loses from entry-level compression that does not model who was obtaining entry-level positions in the first place is incomplete in a way that matters.

Digital transformation is no longer a process that happens to other occupations while IT workers collect the returns. It is restructuring the IT workforce internally, through at least four mechanisms that current designs cannot separate, and the resulting inequality is generated as much by employer identity and entry timing as by anything a worker knows how to do. The qualitative analysis adds a specific reason for the measurement gap: the two mechanisms most likely to determine who captures the returns are each visible to only one community of observers, and neither is positioned to measure the other's variable. Closing that gap requires data access that mostly sits inside firms, and the research community has not yet made a serious collective case for it.

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.
- Acemoglu, D., & Restrepo, P. (2022). *Tasks, automation, and the rise in U.S. wage inequality*. *Econometrica*, 90(5), 1973–2016.
- 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.
- Autor, D. H., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). *The fall of the labor share and the rise of superstar firms*. *Quarterly Journal of Economics*, 135(2), 645–709.
- 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.
- Blinder, A. S. (2006). *Offshoring: The next industrial revolution?* *Foreign Affairs*, 85(2), 113–128.
- Bowen, G. A. (2009). *Document analysis as a qualitative research method*. *Qualitative Research Journal*, 9(2), 27–40.

- Braun, V., & Clarke, V. (2006). *Using thematic analysis in psychology*. *Qualitative Research in Psychology*, 3(2), 77–101.
- Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). *Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence*. *Quarterly Journal of Economics*, 117(1), 339–376.
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab working paper.
- Brynjolfsson, E., Li, D., & Raymond, L. (2023). *Generative AI at work*. NBER Working Paper No. 31161. National Bureau of Economic Research.
- Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). *Firms and labor market inequality: Evidence and some theory*. *Journal of Labor Economics*, 36(S1), S13–S70.
- CompTIA. (2025). *State of the tech workforce*. Computing Technology Industry Association.
- Deepak Ranjan Sahoo, P. D. K. (2017). *Emotional Intelligence: Managing Stress and Anxiety at Work Place*. *International Journal of Multifaceted and Multilingual Studies*, 4(10), 88–93.
- Deepak Ranjan Sahoo, V. D. S. (2019a). *A Study of Economic Literacy and sustainable development with context to India*. *COSMOS Bi-Annual Referred Journal*, 10(1), 47–53.
- Deepak Ranjan Sahoo, V. D. S. (2019b). *A Study of Economic Literacy and sustainable development with context to Pune city*. *Scholarly Research Journal for Interdisciplinary Studies*, 7(41), 382–389.
- Deepak Ranjan Sahoo, V. K. J. (2017). *DERIVATIVES: A device for Risk Management*. *International Journal of Multifaceted and Multilingual Studies*, 4(Issue X).
- Deming, D. J. (2017). *The growing importance of social skills in the labor market*. *Quarterly Journal of Economics*, 132(4), 1593–1640.
- Frey, C. B., & Osborne, M. A. (2017). *The future of employment: How susceptible are jobs to computerisation?* *Technological Forecasting and Social Change*, 114, 254–280.
- Goldin, C. (2014). *A grand gender convergence: Its last chapter*. *American Economic Review*, 104(4), 1091–1119.
- Hosseini, A., & Lichtinger, G. (2025). *Generative AI as seniority-biased technological change: Evidence from U.S. resume and job posting data*. Working paper.
- Kalleberg, A. L. (2011). *Good jobs, bad jobs: The rise of polarized and precarious employment systems in the United States, 1970s to 2000s*. Russell Sage Foundation.
- Katz, L. F., & Murphy, K. M. (1992). *Changes in relative wages, 1963–1987: Supply and demand factors*. *Quarterly Journal of Economics*, 107(1), 35–78.
- Kim, C., & Sakamoto, A. (2008). *The rise of intra-occupational wage inequality in the United States, 1983 to 2002*. *American Sociological Review*, 73(1), 129–157.
- Noy, S., & Zhang, W. (2023). *Experimental evidence on the productivity effects of generative artificial intelligence*. *Science*, 381(6654), 187–192.
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). *The impact of AI on developer productivity: Evidence from GitHub Copilot*. arXiv:2302.06590.
- Sahoo, D. R. (2017a). *Creative Accounting: A Critical Study*. *International Journal of Multifaceted and Multilingual Studies*, 4(IX).
- Sahoo, D. R. (2017b). *Emotional Intelligence: Managing Stress and Anxiety at work place*. *International Journal of Multifaceted and Multilingual Studies*, 4(IX).
- Sahoo, D. R. (2018). *Corporate Governance: Need of the Hour (Companies Act 2013)*. *International Journal of Management Studies (IJMS)*, 4(ii).

- Sahoo, D. R., & Deepak, S. V. (2024a). *Exploring the Ethical Perspectives of Sustainable Finance: A Research Study*. In *Advances in Business Strategy and Competitive Advantage* (pp. 119–142). IGI Global.
- Sahoo, D. R., & Deepak, S. V. (2024b). *Exploring the Nexus of Sustainable Finance: ESG Ethical Framework Practices*. In *Ethical Quandaries in Business Practices: Exploring Morality and Social Responsibility* (pp. 253–280). IGI Global.
- Sahoo, D. R., & Jagtap, V. K. (2017). *IND AS 109: Its impact on Financial Statements*. *International Research Journal of Multidisciplinary Studies*, 3(12), 1–5.
- Sahoo, V. D., & Sahoo, D. R. (2025a). *Promoting Inclusivity: Corporate Social Responsibility in Game-Based Learning and Professional Employment*. In *Game-Based Education Approaches to Inclusive Business Management* (pp. 101–126). IGI Global Scientific Publishing.
- Sahoo, V. D., & Sahoo, D. R. (2025b). *Prospective of Blockchain in Derivative Markets: An Empirical Review*. *Innovations in Blockchain-Powered Intelligence and Cognitive Internet of Things (CIoT)*, 229–252.
- Shende, M. A., & Sahoo, D. R. (2021). *Talent Management critical implications and Strategic approach reference to employer and employees prospective*. *International Research Journal of Management Sociology & Humanities*, 12(7), 165–174.
- Song, J., Price, D. J., Guvenen, F., Bloom, N., & von Wachter, T. (2019). *Firming up inequality*. *Quarterly Journal of Economics*, 134(1), 1–50.
- Stack Overflow. (2025). *2025 Developer Survey*. Stack Overflow.
- van der Velde, L. (2020). *Within occupation wage dispersion and the task content of jobs*. *Oxford Bulletin of Economics and Statistics*, 82(5), 1161–1197.
- Weil, D. (2014). *The fissured workplace: Why work became so bad for so many and what can be done to improve it*. Harvard University Press.