



## RECRUITMENT PROCESS DESIGN AND SEARCH-MATCHING EFFICIENCY HRM VARIABLE: SOURCING AND SELECTION DESIGN ® ECONOMIC OUTCOME: MATCHING EFFICIENCY

**Costs present vacancy duration and mismatch against projected efficiency gains from AI-mediated matching, applying labour search theory directly to recruitment practice rather than treating the two literatures separately.**

**Dr. N. Subbu Krishna Sastry<sup>1</sup>, Dr. Manjula Mallya M<sup>2</sup>**

<sup>1</sup>*School of Management, CMR University, Bengaluru, Karnataka, India*

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

### ABSTRACT

Recruitment practice and labour search theory address the same phenomenon from opposite ends and rarely speak to each other. Human resource management treats sourcing and selection as a process to be optimised on time-to-fill and cost-per-hire; labour economics treats matching efficiency as an unexplained residual in an aggregate matching function. This paper joins the two by writing recruitment process design directly into the matching function, and then costing what that design buys. A recruitment process design index built from sourcing breadth, screening precision, process velocity and offer conversion enters a constant-returns Cobb–Douglas matching function as a multiplicative efficiency shifter. Vacancy duration, mismatch loss and process resource cost are placed on a common per-hire footing and calibrated to Indian IT services for FY2026, using NASSCOM sector aggregates, PLFS labour force estimates and published Indian time-to-fill benchmarks. Four design regimes, from traditional recruitment to fully AI-mediated matching with a human-in-the-loop final interview, are evaluated at point estimates and through a 50,000-draw Monte Carlo simulation. A well-executed AI-mediated design raises matching efficiency by 44.8 per cent, cuts expected time-to-fill from 45 to 31.1 days and lowers total search-and-matching cost per hire by 20.7 per cent. The unconditional expectation across implementation quality is far weaker: mean net gain falls to ₹14,552 per hire and 28.3 per cent of simulated implementations destroy value outright. Velocity gains are near-certain but economically small; the payoff turns on screening precision and on the cost of the algorithmic stack, including bias-audit compliance. Because replacement hiring accounts for 85.1 per cent of gross hires in the sector, roughly a third of the private gain is reallocation between employers rather than new output, implying a private–social wedge of 33.0 per cent that firm-level business cases systematically ignore.

**KEYWORDS** — Recruitment Process Design; Search And Matching; Matching Efficiency; Vacancy Duration; Mismatch Cost; AI-Mediated Hiring; Beveridge Curve; Indian IT Services; Monte Carlo Simulation.

### I. INTRODUCTION

Two mature literatures describe the same event — a firm and a worker finding each other — and they almost never cite one another. On one side, the human resource management literature on recruitment studies sourcing channels, applicant attraction, screening instruments and selection validity, and evaluates them on operational metrics: applications per requisition, time-to-fill, cost-per-hire, offer acceptance. On the other, the labour economics literature descended from Diamond, Mortensen and Pissarides represents the entire process as a matching function, in which a stock of job seekers and a stock of vacancies produce a flow of hires at a rate governed by a scalar efficiency parameter [15], [20]. That parameter is estimated as a residual, interpreted as everything not captured by the two stocks, and left largely unexplained [19].

The separation is costly in both directions. Search theory cannot say why matching efficiency differs across firms, sectors or decades, because it contains no description of what recruiters actually do. Recruitment practice cannot say what its process improvements are worth in economic terms, because time-to-fill is an operational statistic with no price attached to it. A firm that reduces time-to-fill from forty-five days to thirty-one has produced a number; whether it has produced value depends on what an unfilled position costs, on whether the faster process selected as well as the slower one, and on what the acceleration itself cost to buy.

The question has become urgent rather than academic. Artificial intelligence has entered recruitment at the point in the process where search theory locates the friction — the screening of a large applicant pool against an imperfectly specified vacancy. Industry surveys for 2026 report near-universal adoption of some AI tool in hiring, time-to-hire reductions of 40 per cent for AI screening combined with a human final interview, and reductions as high as 75 per cent in vendor-reported cases [10], [12]. Those are claims about matching efficiency, made in the language of process metrics, without the economic accounting that would make them assessable.

This paper supplies that accounting. It writes recruitment process design into the matching function as an explicit, decomposable efficiency shifter, so that the sourcing and selection choices a firm actually makes map onto the parameter that search theory leaves as a residual. It then costs the three channels through which that design affects economic outcomes — the



duration of the vacancy, the quality of the resulting match, and the resources consumed by the process itself — on a single per-hire footing. Finally, it evaluates AI-mediated matching against that accounting, at point estimates and under uncertainty, and separates the return that accrues to the hiring firm from the return that accrues to the economy.

### 1.1 Research Questions

The study addresses four questions. First, how can sourcing and selection design be represented inside a matching function without abandoning the constant-returns structure that makes that function tractable? Second, what is the annual economic cost of search-and-matching friction in Indian IT services, decomposed into vacancy duration, mismatch and process resource use? Third, how much of that cost does AI-mediated matching recover, and with what dispersion across plausible implementations? Fourth, how much of the recovered value is new output rather than reallocation of workers between employers?

### 1.2 Contribution

The contribution is fourfold. Theoretically, the paper gives the matching efficiency parameter a firm-level micro-foundation in recruitment process design, expressed as a weighted geometric index of four measurable sub-dimensions. Methodologically, it places vacancy duration and mismatch — usually measured in incommensurable units of days and turnover percentages — on a common monetary footing that permits a break-even test. Empirically, it calibrates the model to a large, well-documented labour market and reports both a point estimate and a full distribution, rather than the single optimistic scenario characteristic of practitioner business cases. Analytically, it identifies a private–social wedge in AI recruitment gains that follows directly from the replacement-heavy composition of hiring in the sector, and which no firm-level business case has reason to notice.

## II. REVIEW OF LITERATURE

### 2.1 The Aggregate View: Matching Functions and the Efficiency Residual

The canonical search-and-matching framework treats hiring as the output of a production function whose inputs are unemployed job seekers and posted vacancies [15], [20]. Under constant returns, the vacancy filling rate depends only on market tightness, and the Beveridge curve traces the locus of vacancy and unemployment pairs consistent with steady-state flows. Petrongolo and Pissarides, surveying three decades of estimates, place the elasticity with respect to job seekers between 0.5 and 0.7 and observe that the scale parameter — matching efficiency — absorbs everything the specification omits: search intensity, screening technology, geographic and skill dispersion, and the institutions of recruitment [19]. Shimer showed how much rests on that residual, since the volatility of vacancies and unemployment cannot be reproduced without it [32].

Recent work has begun to interrogate the residual rather than accept it. Gottfries reinterprets the post-2010 outward drift of the Beveridge curve as a change in flow composition rather than a deterioration in matching technology [9]. The Federal Reserve Bank of St. Louis pushes the point further, separating vacancies posted to draw workers out of unemployment from poaching vacancies that shuffle already-employed workers between firms; the aggregate vacancy rate rose from roughly 2 per cent before 2010 to roughly 8 per cent in the 2020–2025 period, while vacancies aimed at the unemployed stayed below 2 per cent throughout, so that virtually the entire surge — and hence the apparent collapse in matching efficiency — was poaching [8]. If that decomposition is right, matching efficiency is not falling; the composition of what is being matched is changing, and the distinction matters greatly for what an improvement in recruitment technology is actually worth.

Sastry and Mallya M. approach the same territory from the practice side, arguing that skill diversification is what allows an organisation to remain resilient as labour market conditions and technology shift beneath it [29]. Read against the matching function, that argument is a statement about the composition of the job-seeker stock rather than its size: a workforce whose skills are diversified presents a wider set of feasible assignments to any given vacancy, which is precisely the mechanism by which the efficiency parameter rises without any change in the number of seekers or postings. Their treatment of technology disruption as a labour-market-wide phenomenon rather than a firm-level shock also anticipates the aggregation problem this paper takes up in Section V: a recruitment technology that improves one firm's position in the queue does not necessarily improve the queue.

### 2.2 Recruitment as a Firm-Level Search Technology

The HRM literature has long treated recruitment as a designed system rather than an administrative function. Barber's framework separates applicant generation, applicant maintenance and job choice influence, each with distinct levers [2]; Breaugh's review consolidates evidence that source, message and timing all shape both the size and the composition of the applicant pool [5]. Rynes, Bretz and Gerhart demonstrated that recruitment experiences influence job choice independently of the job itself, which in matching terms means the process affects the probability that an identified match is consummated [22].

What this literature has not done is convert those levers into the language of search. Sastry's study of departmental structure and task achievement offers a way in [23]. Its finding — that how work is partitioned across departments conditions whether tasks are actually completed — has a direct implication for recruitment that the study itself does not draw out. A vacancy is a description of a bundle of tasks; if the underlying structure allocates tasks ambiguously across units, then the requisition inherits that ambiguity and the vacancy becomes a noisy signal of what is being sought. Search theory would call this a degradation of



the meeting technology that occurs before any recruiter acts, and it implies that a share of what is usually charged to sourcing failure is in fact organisational design failure upstream of it.

Entry-level pipelines are a second locus where design and search meet. Sastry's analysis of leadership strategies for young talent management makes the case that future-readiness depends on how deliberately organisations build and hold early-career pipelines rather than on how aggressively they compete for experienced hires [25]. In the framework used here, a maintained campus and early-career pipeline raises sourcing breadth persistently rather than episodically, and does so in the segment where job seekers are most numerous and least attached — exactly the segment where, following the St. Louis Fed decomposition, a filled vacancy is more likely to represent new employment than a transfer between employers.

### **2.3 Selection Design, Screening Precision and Match Quality**

Selection research supplies the one parameter that search theory most conspicuously lacks: how accurately a screen predicts realised productivity. Schmidt and Hunter's validity generalisation places structured procedures and work samples well above unstructured interviews and credential screens [31], while Spence's signalling model explains why candidates invest in observable markers whose informational content erodes as they become universal [33]. Both point to the same variable — the correlation between what the screen measures and what the job requires — which this paper labels screening precision and treats as the decisive design parameter.

Screening precision is, however, only observable after the fact. Sastry and Selvarasu's study of performance appraisal and employee engagement is useful here for a reason unrelated to its own framing [30]. If appraisal outcomes and engagement move together, then the appraisal record becomes the practical instrument through which a firm learns, *ex post*, whether its *ex ante* screen was accurate. That closes a loop the recruitment literature usually leaves open: without a credible post-hire read on realised fit, screening precision cannot be estimated at all, and a firm adopting a new screening technology has no way to know whether it has bought precision or merely speed. The measurement architecture for validating a screen is thus an appraisal-system property, not a recruitment-system property.

Offer conversion is the fourth design lever, and Sastry's analysis of talent designation and social comparison speaks to it directly [26]. Designations function as signals whose value to the recipient depends on the comparison set rather than on the title itself; the paper treats this comparative dimension as an economic driver of organisational growth rather than a matter of status. Applied to recruitment, it implies that offer acceptance is not a fixed conversion rate to be optimised in isolation but a function of how the offered position sits relative to the candidate's reference group — which is why identical compensation packages convert at different rates across firms, and why a design that improves speed and precision can still fail at the final step.

### **2.4 The Economic Cost of Vacancy Duration and Mismatch**

Hiring costs have been measured with reasonable precision in the European literature. Blatter, Muehlemann and Schenker estimate the cost of hiring a skilled worker at several months of wages once search, adaptation and productivity ramp are included [3], and Muehlemann and Strupler Leiser show these costs rise convexly with labour market tightness [16]. Guvenen and co-authors quantify the other channel, showing that multidimensional skill mismatch carries large and persistent earnings and productivity penalties [11]. The two channels are usually studied apart, which obscures the trade-off between them: a process can always be accelerated by screening less, and the saving in vacancy days is then paid for in mismatch.

The mechanism by which mismatch becomes a cost is worth stating carefully, and Sastry's work on workforce burnout supplies it [24]. Burnout is treated there as a leadership and workload phenomenon requiring organisational rather than individual remedy. Placed in the present framework, it identifies the pathway through which a mismatched hire converts into a monetary loss: misfit between the person and the demands of the role does not present as a hiring failure but as sustained strain, declining discretionary effort, and eventual separation. The loss per mismatched hire used in this paper is therefore not a single event cost but the discounted sum of that trajectory, which is why estimates that count only replacement expenditure understate it.

Sourcing breadth interacts with both channels, and Sastry's study of women employees in Bangalore IT startups makes the interaction visible [27]. Where the effective sourcing pool is narrowed — by retention conditions, by role design, or by the informal channels through which startups recruit — the firm searches a smaller distribution for the same requirement. In matching terms this raises expected vacancy duration and simultaneously worsens the expected quality of the best available candidate, because both are order statistics of the pool that was actually searched. The two costs this paper separates for analytical purposes therefore share a common upstream determinant, which is one reason interventions that widen the pool tend to outperform interventions that accelerate the process.

### **2.5 AI-Mediated Matching: Precision, Gaming and Compliance**



Online job boards were the first technology claimed to transform matching; Autor argued they would reduce search costs and thicken markets [1], though Kuhn and Mansour found the employment effects more modest than anticipated [13]. Algorithmic screening makes a stronger claim, because it operates on the screening stage rather than the meeting stage. Sastry and Mallya M. set out the case in detail, arguing that AI analytics in hiring both compresses the process and improves the firm's ability to attract and identify strong candidates [28]. That paper is the direct antecedent of the present one, and the relationship between them is deliberate: it establishes the mechanism and the operational case, and this paper subjects that case to an economic test by asking what the compressed process is worth net of what it costs to run and net of what it does to selection accuracy.

Three findings complicate the optimistic reading. First, the precision gain is contested. Cowgill finds algorithmic résumé screening can outperform human screeners, particularly on non-traditional candidates [6], while Raghavan and co-authors document how widely audited practice diverges from vendor claims [21]. Second, candidates adapt: 2026 survey evidence reports that 74 per cent of job seekers use AI in their search and 41 per cent admit to using prompt injection or comparable techniques to defeat automated screening, while 91 per cent of recruiters report encountering suspected candidate deception [10]. A screen that is gamed loses precision in exactly the dimension it was adopted to improve. Third, the technology carries a compliance overhead: New York City's Local Law 144 requires annual independent bias audits of automated employment decision tools [18], and the EU Artificial Intelligence Act classifies recruitment and selection systems as high-risk under Annex III, with conformity, documentation and human oversight obligations attached [7]. These are recurring per-hire costs, not one-time implementation costs, and the model below treats them as such.

Adoption evidence should be read with care. The 2026 figures most widely cited — 99 per cent of hiring managers reporting some AI use, 40 per cent time-to-hire reduction for AI screening with human final interview, 25 per cent improvement in first-year retention, 21 per cent better role alignment — originate in vendor and platform surveys of self-selected adopters [10], [12]. They are used in this paper as calibration anchors and as an upper bound against which model-implied estimates can be compared, not as unbiased population parameters.

## **2.6 The Indian IT Services Context**

India's technology sector provides an unusually well-instrumented setting. NASSCOM projects industry revenue of approximately US\$315 billion in FY2026 on 6.1 per cent growth, with headcount near 5.95 million, net additions of roughly 135,000 and AI-attributable revenue of US\$10–12 billion [17]. The gap between revenue growth and headcount growth is itself the phenomenon under study: output is expanding while net hiring has flattened, which means the majority of hiring activity is replacement rather than expansion. On the supply side, the Periodic Labour Force Survey for calendar 2025 reports all-India unemployment of 3.1 per cent, urban unemployment of 4.8 per cent and urban youth unemployment of 13.6 per cent, with a labour force participation rate of 59.3 per cent [14]. Published Indian benchmarks place time-to-fill for mid-level software and data roles at 35 to 60 days, rising to 60 to 90 days for data science, machine learning and cybersecurity roles, with offer acceptance for IT roles between 68 and 75 per cent and interview-to-offer ratios of 5:1 to 7:1 against a 3:1 target [34].

## **2.7 Research Gap**

The review identifies a specific and tractable gap. Search theory possesses a rigorous aggregate framework with an unexplained efficiency parameter. Recruitment research possesses a detailed account of the design levers that plausibly constitute that parameter but no method for pricing them. Neither literature has costed vacancy duration and mismatch against one another on a common footing, and neither has evaluated AI-mediated matching in a way that distinguishes a well-executed implementation from the expected outcome across implementations, or private from social return. This paper closes that gap by construction rather than by estimation, building the bridge as an explicit model and calibrating it to observable sector data.

# **III. RESEARCH METHODOLOGY**

## **3.1 Research Design**

The study uses an analytical model calibrated to secondary data and evaluated by Monte Carlo simulation. No primary data are collected and no parameters are estimated econometrically; the model is built from established structural relationships and calibrated so that its baseline reproduces observed sector benchmarks. This design is appropriate because the counterfactual of interest — the economy-wide effect of a recruitment technology that is being adopted contemporaneously and unevenly — is not observable in historical data, and because firm-level recruitment cost data of the granularity required are not publicly available in India. The trade-off is explicit: the model buys internal consistency and transparency of assumptions at the cost of external validation, and Section VII treats this as the study's principal limitation.




**Fig. 1.** Analytical framework. Recruitment process design is decomposed into four measurable sub-indices which enter the matching function as a multiplicative efficiency shifter. The shifter propagates to three cost channels, whose sum is the total search-and-matching cost per hire. Private and social returns diverge because a share of matches is reallocation between employers rather than new employment.

### 3.2 The Augmented Matching Function

Let  $U$  denote the stock of job seekers in the relevant skill segment and  $V$  the stock of posted vacancies. Hires are produced by a constant-returns Cobb–Douglas matching function

$$M = A \cdot U^\alpha \cdot V^{1-\alpha} \quad (1)$$

where  $\alpha$  is the elasticity with respect to job seekers, set at 0.60, within the 0.5–0.7 interval established in the literature [19]. Constant returns imply that the vacancy filling rate depends only on market tightness  $\theta = V/U$ :

$$q(\theta) = M/V = A \cdot \theta^{-\alpha} \quad (2)$$

and expected vacancy duration is the reciprocal of the filling rate,

$$D = 1/q(\theta) = \theta^\alpha / A \quad (3)$$

The departure from the standard treatment is to decompose  $A$  rather than estimate it. Let  $A = A_0 \cdot \Omega$ , where  $A_0$  is the matching efficiency delivered by a traditional recruitment process and  $\Omega$  is a recruitment process design index formed as a weighted geometric mean of four sub-indices, each normalised to unity under the traditional design:

$$\Omega = S^{\theta_S} \cdot P^{\theta_P} \cdot R^{\theta_R} \cdot E^{\theta_E} \quad \Sigma \theta_K = 1 \quad (4)$$

Here  $S$  is sourcing breadth (the effective size and diversity of the pool actually searched),  $P$  is screening precision (the correlation between the screen and realised job performance),  $R$  is process velocity (the inverse of internal cycle time, net of external market conditions) and  $E$  is offer conversion. The unit-sum restriction on the elasticities preserves the interpretation of  $\Omega$  as a pure efficiency shifter: doubling every dimension of design quality doubles matching efficiency. The elasticities are set at  $\theta_S = 0.25$ ,  $\theta_P = 0.20$ ,  $\theta_R = 0.35$  and  $\theta_E = 0.20$ , weighting velocity most heavily because it operates on every requisition while precision operates only at the screening stage. The endogeneity that complicates econometric estimation of the matching function [4] does not arise here, since  $\Omega$  is imposed rather than fitted.




**Fig. 2.** Steady-state Beveridge loci implied by the calibrated model. Raising matching efficiency shifts the locus toward the origin: the same flow of matches is sustained with fewer vacancies open at any given job-seeker rate. At the PLFS 2025 urban unemployment rate of 4.8 per cent, the D3 design sustains steady state at a 2.4 per cent vacancy rate against 6.0 per cent under D0.

### 3.3 Match Quality and the Gaming Penalty

Screening precision governs the probability that a completed hire is a mismatch. Let  $\pi_{m0}$  be the mismatch probability under the traditional design. A design with precision multiplier  $P$  yields

$$\pi_m = \pi_{m0} \cdot P^{-\lambda} \cdot (1 + \gamma \cdot g) \quad (5)$$

where  $\lambda$  is the elasticity of mismatch with respect to precision, set at 1.10, and the final term is the gaming penalty. The variable  $g$  measures the intensity of algorithmic mediation and  $\gamma$  the rate at which candidates and intermediaries adapt to it. This term is the model's formal representation of the survey finding that a substantial minority of candidates now deliberately manipulate automated screens [10]: an algorithmic screen is not a fixed instrument but one whose informational content decays as it is optimised against. Its inclusion is what makes the mismatch channel capable of deteriorating even as speed improves.

### 3.4 The Cost Function

The three channels are placed on a common per-hire footing. The vacancy-duration cost is

$$C_v = D \cdot c_v \cdot \phi \quad (6)$$

where  $c_v$  is the daily contribution forgone while the position is unfilled and  $\phi$  is a realisation factor capturing the share of that contribution that the existing team cannot absorb. The mismatch cost is

$$C_m = \pi_m \cdot L_m \quad (7)$$

with  $L_m$  the discounted loss per mismatched hire, comprising lost productivity, the separation event and the replacement cycle. The process cost is the sum of recruiter time, external agency and portal fees, assessment expenditure, algorithmic licensing and the recurring cost of bias audit and conformity documentation:

$$C_p = c_{rec} + c_{agy} + c_{asm} + c_{lic} + c_{aud} \quad (8)$$

and total search-and-matching cost per hire is

$$C = C_v + C_m + C_p \quad (9)$$

Two features of this specification deserve comment. First, treating  $c_{lic}$  and  $c_{aud}$  as recurring per-hire costs rather than amortised capital expenditure follows the regulatory structure: Local Law 144 requires an annual independent bias audit [18] and the EU Artificial Intelligence Act imposes continuing conformity and human-oversight obligations on high-risk employment systems [7], so the compliance burden scales with continued use rather than with initial deployment. Second, the specification permits  $C$  to rise: a design that raises  $c_{lic}$  and  $c_{aud}$  faster than it lowers  $C_v$  and  $C_m$  destroys value, and Section IV shows this outcome is not rare.

### 3.5 Design Regimes

Four recruitment designs are compared. D0 is the traditional process: job-board and agency sourcing, résumé screening by recruiters, unstructured or semi-structured interviews. D1 adds AI-assisted sourcing, résumé parsing and candidate matching at the top of the funnel, leaving human judgement at every decision point. D2 adds AI-assisted structured screening and automated assessment scoring. D3 is fully AI-mediated matching with a human-in-the-loop final interview, the configuration for which the strongest empirical claims are made [10], [12]. Sub-index values and cost parameters for each regime are set out in Table 2.



### 3.6 Calibration

The model is calibrated to Indian IT and business services for FY2026. Sector revenue of US\$315 billion on headcount of 5.95 million gives annual revenue per employee of approximately ₹46.6 lakh at ₹88 per US dollar [17]. Rather than work with that sector average, which is inflated by onsite and subcontracted staff, the representative unit is a mid-level offshore engineering role with a fully loaded annual cost of ₹14 lakh against attributable annual revenue of ₹24 lakh, giving a daily contribution of ₹2,740 while the seat is filled. Gross hiring is derived as voluntary attrition of 13 per cent applied to sector headcount plus net additions of 135,000, giving 908,500 gross hires, of which replacement hiring accounts for 85.1 per cent. Baseline time-to-fill of 45 days is the midpoint of the published Indian benchmark of 35 to 60 days for mid-level software and data roles [34]. Market tightness in the skill segment is set at  $\theta = 0.75$ , consistent with urban unemployment of 4.8 per cent alongside persistent reported difficulty in filling specialist roles [14]. Full calibration is given in Table 1.

**Table 1. Model parameters, baseline values and sources.**

| Parameter                                | Value                  | Basis                                                               |
|------------------------------------------|------------------------|---------------------------------------------------------------------|
| $\alpha$ matching elasticity             | 0.60                   | Survey range 0.5–0.7 [19]                                           |
| $\theta$ market tightness (V/U)          | 0.75                   | PLFS urban unemployment 4.8% [14]                                   |
| $D_0$ baseline time-to-fill              | 45 days                | India mid-level IT benchmark 35–60 d [34]                           |
| $c_v$ daily contribution forgone         | ₹2,740                 | Role revenue ₹24L less cost ₹14L, $\div 365$                        |
| $\phi$ realisation factor                | 0.70                   | Partial absorption by existing team                                 |
| $\pi_{m0}$ baseline mismatch probability | 0.14                   | Early-attrition and misfit incidence                                |
| $L_m$ loss per mismatched hire           | ₹7,00,000              | $0.50 \times$ annual role cost [3], [11]                            |
| $\lambda$ mismatch–precision elasticity  | 1.10                   | Selection validity literature [31]                                  |
| $\gamma$ gaming penalty coefficient      | 0.35                   | Candidate adaptation to AI screens [10]                             |
| $\theta_S, \theta_P, \theta_R, \theta_E$ | 0.25, 0.20, 0.35, 0.20 | Imposed; unit sum                                                   |
| $C_p$ process cost, $D_0$                | ₹65,000                | Recruiter ₹28k + agency ₹30k + assessment ₹7k                       |
| Gross hires FY2026                       | 908,500                | $13\% \text{ attrition} \times 5.95\text{m} + 135\text{k net}$ [17] |
| Replacement share of hires               | 85.1%                  | Derived from [17]                                                   |
| Exchange rate                            | ₹88 / US\$             | FY2026 working assumption                                           |

**Table 2. Recruitment design regimes: sub-index values, gaming intensity and process cost components (₹ per hire).**

| Regime              | S    | P    | R    | E    | g    | AI cost | Total $C_p$ |
|---------------------|------|------|------|------|------|---------|-------------|
| D0 Traditional      | 1.00 | 1.00 | 1.00 | 1.00 | 0.00 | 0       | 65,000      |
| D1 AI sourcing      | 1.35 | 1.02 | 1.15 | 1.00 | 0.05 | 7,500   | 62,400      |
| D2 + AI screening   | 1.45 | 1.18 | 1.35 | 1.05 | 0.12 | 15,000  | 58,860      |
| D3 Full AI-mediated | 1.55 | 1.30 | 1.70 | 1.12 | 0.18 | 25,500  | 59,940      |

### 3.7 Simulation Protocol

Point estimates describe a well-executed reference implementation in which every design parameter takes its central value. Because real implementations vary in quality, vendor capability and organisational absorption, the model is also evaluated across 50,000 Monte Carlo draws. Triangular distributions are used for parameters with a defensible central case and bounded support; uniform distributions where no central case is defensible; and a Dirichlet distribution for the four design elasticities so that the unit-sum restriction holds in every draw. Critically, the D3 precision multiplier is drawn from a triangular distribution on [0.92, 1.48] with mode 1.28, allowing precision to deteriorate rather than improve. This admits the possibility, supported by the survey evidence on candidate gaming and recruiter-reported deception [10], that an algorithmic screen performs worse than the human screen it replaced. Offer conversion is similarly allowed to fall below unity, reflecting reported erosion of candidate trust where AI use is disclosed. Without these downside branches the simulation would merely restate the point estimate with noise.

### 3.8 Private and Social Returns

Let  $\sigma$  denote the share of gross hires that represent net employment creation, and  $(1 - \sigma)$  the replacement share. Following the vacancy-composition argument [8], a filled replacement vacancy transfers a worker between employers; the hiring firm gains,

the losing firm incurs its own vacancy, and the net addition to output is limited to the improvement in assignment quality. The social gain is therefore

$$G_{social} = \Delta C_v [\sigma + (1 - \sigma) \psi] + \Delta C_m + \Delta C_p \quad (10)$$

where  $\psi$  is the fraction of the vacancy-duration saving on replacement hires that represents genuine output rather than transfer, set conservatively at 0.25. The mismatch and process channels enter in full, because better assignment and lower resource use are social gains regardless of whether the worker came from unemployment or from a competitor. Equation (10) is the paper's analytical bridge between the firm-level business case and the aggregate matching efficiency question, and Section IV.7 shows it drives a wedge of a third between the two.

## IV. RESULTS AND ANALYSIS

### 4.1 Matching Efficiency and Vacancy Duration

Table 3 reports the point-estimate outcome of each design regime. The recruitment process design index rises from 1.0000 under the traditional design to 1.1364 under AI-assisted sourcing, 1.2723 when AI-assisted structured screening is added, and 1.4483 under full AI-mediated matching. Since  $\Omega$  multiplies matching efficiency directly, the fully mediated design delivers a 44.8 per cent increase in  $\Omega$ , raising the calibrated efficiency parameter from 0.01870 to 0.02708. Expected time-to-fill, which is inversely proportional to  $\Omega$  at fixed tightness, falls from 45.0 days to 39.6, 35.4 and 31.1 days respectively — a reduction of 30.9 per cent at D3.

Figure 2 translates this into the aggregate representation. At the PLFS 2025 urban unemployment rate of 4.8 per cent, the steady-state vacancy rate consistent with the D0 matching technology is 6.0 per cent; under D3 it is 2.4 per cent. The amplification arises because the steady-state locus depends on efficiency raised to the power  $1/(1 - \alpha)$ , so a 44.8 per cent efficiency gain produces a far larger inward movement of the Beveridge locus than of the filling rate itself. This is the sense in which recruitment process design is a macroeconomic variable: identical labour supply and identical labour demand generate materially different observed vacancy rates depending on the screening technology in use.

**Table 3. Point-estimate outcomes by design regime. Matching efficiency  $\Omega$  rises from 0.01870 at D0 to 0.02708 at D3.**

| Costs in ₹ per hire. |          |      |         |         |          |         |
|----------------------|----------|------|---------|---------|----------|---------|
| Regime               | $\Omega$ | Days | $\pi_m$ | Vacancy | Mismatch | Total   |
| D0 Traditional       | 1.0000   | 45.0 | 14.00%  | 86,301  | 98,000   | 249,301 |
| D1 AI sourcing       | 1.1364   | 39.6 | 13.94%  | 75,940  | 97,566   | 235,906 |
| D2 + AI screening    | 1.2723   | 35.4 | 12.16%  | 67,834  | 85,118   | 211,812 |
| D3 Full AI-mediated  | 1.4483   | 31.1 | 11.15%  | 59,586  | 78,059   | 197,585 |




**Fig. 3.** Speed and quality do not move together. Expected time-to-fill falls 30.9 per cent between D0 and D3 while the matching efficiency multiplier rises 44.8 per cent (upper panel), but the mismatch probability falls only from 14.00 to 11.15 per cent (lower panel), and barely moves at D1 because AI-assisted sourcing adds almost no screening precision.

### 4.2 Match Quality

The lower panel of Figure 3 records a result that the operational literature systematically understates. Mismatch probability falls from 14.00 per cent under D0 to only 13.94 per cent under D1, because AI-assisted sourcing widens the pool without improving the screen — and the small gain in precision is almost exactly offset by the gaming penalty the technology introduces. Meaningful improvement begins only at D2, where screening itself becomes algorithmically structured, and reaches 11.15 per

cent at D3: a relative reduction of 20.4 per cent. Speed improves by half as much again as quality, and the divergence widens with the intensity of mediation, since the gaming penalty scales with  $g$  while the precision gain does not.

#### 4.3 Where the Money Is

Figure 4 decomposes total cost per hire. Under the traditional design, the ₹249,301 per-hire friction divides into ₹86,301 of vacancy-duration cost, ₹98,000 of mismatch cost and ₹65,000 of process cost. The largest single component is mismatch, not speed — a point obscured whenever recruitment performance is managed on time-to-fill alone. Full AI mediation cuts the total to ₹197,585, a reduction of 20.7 per cent, of which ₹26,715 comes from the vacancy channel, ₹19,941 from the quality channel and only ₹5,060 from the process channel.

The process channel behaves non-monotonically and this is the model's first counter-intuitive result. Process cost falls from ₹65,000 to ₹62,400 and then to ₹58,860 as automation displaces recruiter and agency expenditure, but rises again to ₹59,940 at D3, because licensing and bias-audit compliance grow faster than the remaining recruiter time they displace. The savings that recruitment automation is usually sold on — fewer recruiter hours, lower agency fees — are the smallest of the three channels and are largely consumed by the cost of the algorithmic stack. What pays for AI-mediated matching is the shortened vacancy and the better match, not the leaner recruiting function.



**Fig. 4.** Cost decomposition per hire by design regime (₹ thousands). Mismatch is the largest component under every regime. The process-cost band contracts and then expands, because licensing and bias-audit compliance at D3 outgrow the recruiter and agency expenditure they displace.

**Table 4.** Sector-level aggregation, FY2026. Applied to 908,500 gross hires; US\$ converted at ₹88.

| Regime              | Friction (₹ crore) | Saving (US\$ mn) | % of sector revenue    |
|---------------------|--------------------|------------------|------------------------|
| D0 Traditional      | 22,649             | —                | 0.82% (total friction) |
| D1 AI sourcing      | 21,432             | 138              | 0.044%                 |
| D2 + AI screening   | 19,243             | 387              | 0.123%                 |
| D3 Full AI-mediated | 17,951             | 534              | 0.169%                 |

Aggregated to the sector, total search-and-matching friction under the traditional design is ₹22,649 crore, or approximately US\$2.57 billion — 0.82 per cent of FY2026 sector revenue. Full AI mediation recovers ₹4,698 crore, about US\$534 million or 0.169 per cent of revenue. The recovery is real but modest relative to the rhetoric surrounding recruitment automation, and it is bounded above by the fact that four-fifths of the friction is not addressable by faster screening at all.

#### 4.4 The Distribution of Outcomes

The point estimates describe a competent implementation. Table 5 and Figure 5 report what happens across 50,000 draws spanning the plausible range of implementation quality, market conditions and algorithmic cost. Mean net gain falls to ₹14,552 per hire with a standard deviation of ₹26,068, and the median to ₹15,677 — roughly 30 per cent of the reference-implementation figure of ₹51,716. The probability that D3 outperforms D0 is 71.7 per cent, which means 28.3 per cent of plausible implementations destroy value.

**Table 5. Monte Carlo results, 50,000 draws. Net gain per hire, D3 versus D0.**

| Statistic               | Value (₹) | Interpretation              |
|-------------------------|-----------|-----------------------------|
| Mean                    | 14,552    | Unconditional expectation   |
| Median                  | 15,677    | Typical implementation      |
| Standard deviation      | 26,068    | Dispersion exceeds the mean |
| 5th percentile          | -30,254   | Adverse outcome             |
| 25th percentile         | -2,766    | Lower quartile is negative  |
| 75th percentile         | 33,033    |                             |
| 95th percentile         | 55,295    | Near the point estimate     |
| P(net gain > 0)         | 71.7%     | 28.3% destroy value         |
| P(mismatch worsens)     | 24.6%     | Quality can deteriorate     |
| P(time-to-fill worsens) | 0.0%      | Speed gain is certain       |




**Fig. 5.** Distribution of net gain per hire across 50,000 simulated implementations. The left tail is not a modelling artefact: it is generated by draws in which screening precision fails to improve while licensing and audit costs are high.

The asymmetry between the two operational outcomes is the central empirical finding. Time-to-fill improves in every single draw — the probability that D3 is slower than D0 is zero — while mismatch worsens in 24.6 per cent of draws. Acceleration is a near-deterministic consequence of automating a manual funnel; accuracy is not. Conditioning makes the point sharper. In the 8 per cent of draws where the precision multiplier stays below 1.05, mean net gain is -₹13,181 and only 31.2 per cent of implementations create value. In the 29 per cent of draws where precision reaches 1.30 or above, mean net gain is ₹28,744 and 88.1 per cent create value. The same technology, the same cost, opposite economics — separated entirely by whether the screen actually got better at predicting performance.

At sector level, the mean expected gain is ₹1,322 crore, about US\$150 million or 0.048 per cent of revenue, with a 5th-to-95th percentile range from -₹2,749 crore to +₹5,024 crore. The unconditional expectation is roughly a quarter of the reference-implementation estimate, and the interval includes substantial aggregate value destruction.

#### 4.5 What the Result Depends On

Figure 6 ranks the fourteen parameters by their influence on the reference-case net gain of ₹51,716 per hire. The cost of the algorithmic stack dominates, moving net gain by -₹22,500 to +₹10,500 across its plausible range. The three parameters that follow — loss per mismatched hire, daily contribution forgone and baseline time-to-fill — are environmental: they scale the size of the prize but are not levers a firm pulls. Among the levers that are chosen, screening precision is decisive at -₹12,749 to +₹10,122, and it ranks fourth overall. The dimensions on which recruitment automation is generally marketed rank near the bottom of the fourteen: process velocity moves the result by at most ₹3,411, sourcing breadth by ₹2,094 and offer conversion by ₹817. Only two of the four most influential parameters are within the firm's control, and they are the price of the algorithmic stack and the accuracy of the screen it runs.




**Fig. 6.** One-at-a-time sensitivity around the reference case, showing the nine most influential of fourteen parameters. Algorithmic cost dominates and screening precision is the decisive design lever; process velocity, sourcing breadth and offer conversion — the dimensions on which recruitment automation is usually marketed — rank tenth, thirteenth and fourteenth and are omitted from the chart.

#### 4.6 The Break-Even Frontier

The lower panel of Figure 7 inverts the question: given a precision gain, what is the most a firm can spend on the algorithmic stack before the investment stops paying? At a precision multiplier of 1.30, break-even expenditure is ₹77,236 per hire; at 1.10 it is ₹59,465; at 1.00 — no precision gain at all, so that the entire return comes from speed and process saving — it is ₹47,911; and at 0.95, where the screen is worse than the one it replaced, ₹41,217. The frontier is comfortably above the modal observed spend of ₹28,000 per hire, which is why the median implementation still creates value. It is not, however, above the upper end of the observed range: at ₹95,000 per hire, no precision multiplier in the plausible interval breaks even. The loss region in Figure 5 is populated precisely by such draws.

#### 4.7 Private Gain and Social Gain

The upper panel of Figure 7 decomposes the reference-case gain of ₹51,716 per hire into its three channels and then applies equation (10). Replacement hiring accounts for 85.1 per cent of gross hires in the sector, so most vacancy-duration saving is a faster transfer of an already-employed worker between employers rather than a faster movement from non-employment into work. Applying  $\sigma = 0.149$  and  $\psi = 0.25$ , the vacancy channel's social value falls from ₹26,715 to ₹9,656. The mismatch and process channels carry over in full, giving a social gain of ₹34,657 against a private gain of ₹51,716 — a ratio of 67.0 per cent and a wedge of 33.0 per cent. At sector level, the ₹4,698 crore of private gain corresponds to about ₹3,149 crore of social gain, US\$358 million against US\$534 million.




**Fig. 7.** Upper: decomposition of the reference-case gain and the deduction for reallocation between employers. Lower: break-even algorithmic expenditure as a function of the screening-precision multiplier, against the observed range of per-hire AI spend.



**Table 6. Private versus social gain, D3 relative to D0 (₹ per hire).**

| Channel          | Private       | Social        | Retained     |
|------------------|---------------|---------------|--------------|
| Vacancy duration | 26,715        | 9,656         | 36.1%        |
| Match quality    | 19,941        | 19,941        | 100%         |
| Process cost     | 5,060         | 5,060         | 100%         |
| <b>Total</b>     | <b>51,716</b> | <b>34,657</b> | <b>67.0%</b> |

## V. DISCUSSION

### 5.1 Speed Is Cheap; Precision Is Dear

The most robust finding is also the least convenient for the recruitment technology market. Automating a hiring funnel reliably makes it faster — the simulation produces no draw in which D3 is slower than D0 — and reliably makes it faster by an amount that is worth relatively little. Vacancy days are cheap at the margin because the existing team absorbs part of the load, and because the modal vacancy in this sector is a replacement whose contribution was being produced somewhere in the economy the day before. Precision, by contrast, is expensive to obtain, uncertain to obtain, and worth a great deal when obtained, because mismatch is the largest cost component under every regime tested. Recruitment functions managed on time-to-fill are therefore optimising the dimension that automation improves most reliably and that matters least, which is a plausible explanation for the gap between reported operational improvement and unreported economic improvement.

### 5.2 Model-Implied Against Vendor-Reported Gains

The model was calibrated independently of the headline claims and can be checked against them. It implies a 30.9 per cent reduction in time-to-fill for a well-executed AI-mediated design with human final interview, against 40 per cent reported for that configuration and 75 per cent reported in the most favourable vendor cases [10], [12]. It implies a 20.4 per cent reduction in mismatch, against a reported 25 per cent improvement in first-year retention and 21 per cent better role alignment. The quality figures agree closely; the speed figure is materially more conservative. That pattern is consistent with selection in the reporting: firms that adopt early and report results are those whose baseline processes had the most slack, and process cycle time is the metric most exposed to that slack. The agreement on quality, which is harder to manipulate and slower to observe, provides modest external support for the mismatch specification.

### 5.3 The Composition Problem

The vacancy-composition literature has so far been read as a diagnostic caution for monetary policy [8], [9]. The present results suggest it has a direct application to HRM economics. If 85 per cent of hiring in a sector is replacement, then a technology that helps every firm fill vacancies faster mostly helps firms take workers from each other faster. Each firm's business case is correct on its own terms and the aggregate case is a third weaker, because a third of the measured gain is the private value of winning a race whose social value is limited to the improvement in assignment. The wedge is not a modelling artefact of the chosen  $\psi$ : at  $\psi = 0$  it widens to 44.0 per cent and at  $\psi = 0.5$  it narrows to 22.0 per cent, so a substantial wedge survives across the whole range.

This has a corollary for how AI recruitment adoption should be expected to diffuse. Because the private return exceeds the social return, adoption will be faster than efficiency alone would warrant, and firms that do not adopt will lose ground even where adoption creates little aggregate value — the structure of an arms race rather than of a productivity improvement. The channel through which genuine aggregate gains arise is the mismatch channel, which is exactly the channel with the weakest and most contested evidence.

### 5.4 Compliance as a Structural Cost

Treating bias audit and conformity documentation as recurring per-hire costs rather than amortised implementation costs materially changes the conclusion, and the regulatory design justifies it. Annual independent bias audits under Local Law 144 [18] and the continuing obligations attaching to high-risk employment systems under the EU Artificial Intelligence Act [7] scale with continued operation. In the model this is what makes process cost turn upward at D3 and what makes algorithmic cost the single most influential parameter in Figure 6. Indian IT services firms serving US and EU clients face these obligations through their delivery relationships even where domestic law imposes none, so the cost is not avoidable by jurisdiction.

### 5.5 Theoretical Implication

The paper's theoretical claim is narrow and testable: matching efficiency is not a residual but a construct, and recruitment process design is its firm-level content. Equation (4) makes A a function of four measurable quantities, each of which a firm chooses. If the decomposition is right, then cross-firm variation in matching efficiency should be predictable from recruitment design audits, and the sectoral drift in aggregate matching efficiency should be traceable to the diffusion of screening technologies. Both are estimable propositions given firm-level requisition data, which is the natural next step and is set out in Section VII.



## VI. IMPLICATIONS

### 6.1 For Firms

Four implications follow directly. First, evaluate recruitment technology on realised match quality, not on time-to-fill; the metric that improves most reliably is the one worth least. Second, instrument the screen before automating it — screening precision cannot be managed without a post-hire read on realised fit, which means the appraisal architecture is a prerequisite for the recruitment investment rather than an unrelated system [30]. Third, treat licensing and audit expenditure as a recurring per-hire cost and test it against the break-even frontier in Figure 7 rather than against an implementation budget; above roughly ₹77,000 per hire the investment requires a precision multiplier of 1.30 or better, which fewer than a third of simulated implementations achieve. Fourth, expect the gaming penalty to grow: a screen that is optimised against by candidates loses precision over time, so the business case should be tested at a lower precision multiplier in later years than in the first.

### 6.2 For Policy

For policy the implications concern the gap between private and social return. Public support for recruitment technology adoption — through skilling missions, digital public infrastructure for labour markets, or procurement preference — is best directed at the channels where private and social returns coincide, namely assignment quality and the movement of workers from non-employment into work, rather than at speed of replacement hiring. A public employment exchange or labour market information system that widens the pool searched, particularly for first-time and returning entrants, addresses the  $\sigma$  term in equation (10) directly, which no firm-level investment does. Regulators designing audit regimes should note that per-hire compliance cost is the single most influential parameter in the model: an audit regime that is expensive per hire rather than per system will suppress adoption of exactly the precision-improving configurations that generate the social gains the regime exists to protect.

## VII. LIMITATIONS AND FUTURE RESEARCH

The model is calibrated, not estimated. The design elasticities in equation (4) are imposed on judgement rather than fitted, and while the Dirichlet draws in the simulation test sensitivity to that choice, they cannot establish it. The mismatch specification in equation (5) is functional-form dependent, and  $\lambda$  is not separately identified from  $\pi_{m0}$  in the data used. Adoption statistics are drawn from vendor and platform surveys of self-selected respondents [10], [12] and are treated as anchors and upper bounds rather than population parameters. The representative-role calibration abstracts from heterogeneity across role families, where time-to-fill ranges from 15 days for entry-level positions to 120 days for architects [34]; the friction cost is likely understated for specialist roles and overstated for high-volume entry-level hiring. Finally, the private–social decomposition assumes the losing firm in a poaching transaction incurs an equivalent vacancy, which holds in a steady state but not during rapid sectoral contraction.

Four extensions follow. First, direct estimation of  $\Omega$  from firm-level requisition data, using variation in the timing of applicant-tracking-system upgrades as the source of identification. Second, measurement of the gaming penalty by tracking screen validity over time within a single firm as candidate adaptation proceeds — the parameter is currently the least evidenced in the model and among the most consequential. Third, extension to a two-sided model in which candidates also search with algorithmic assistance, which the 74 per cent candidate AI adoption figure suggests is now the empirically relevant case. Fourth, replication in a labour market with a lower replacement share of gross hiring, where the private–social wedge should be correspondingly smaller; the model predicts the wedge scales approximately with  $(1 - \sigma)$ , which is a falsifiable cross-country prediction.

## VIII. CONCLUSION

This paper set out to apply labour search theory directly to recruitment practice rather than treating the two literatures separately, and to cost present vacancy duration and mismatch against the projected efficiency gains from AI-mediated matching. Writing sourcing and selection design into the matching function as a decomposable efficiency shifter turns matching efficiency from a residual into a construct with measurable firm-level content, and permits vacancy duration and mismatch — ordinarily reported in incommensurable units — to be compared on a common monetary footing.

Calibrated to Indian IT services for FY2026, the model puts total search-and-matching friction at approximately US\$2.57 billion, or 0.82 per cent of sector revenue. A well-executed AI-mediated design raises matching efficiency by 44.8 per cent, shortens expected time-to-fill from 45.0 to 31.1 days, lowers mismatch from 14.00 to 11.15 per cent and recovers 20.7 per cent of the friction, worth about US\$534 million at sector level. Across 50,000 simulated implementations, however, the expected gain is roughly a quarter of that, and 28.3 per cent of implementations destroy value. Time-to-fill improves in every draw; mismatch worsens in a quarter of them.

Three conclusions follow. The economics of AI-mediated matching rest on screening precision and on the cost of the algorithmic stack, not on the speed and headcount savings on which the technology is generally sold. The dispersion of outcomes is wide enough that adoption is a decision under genuine uncertainty rather than a settled optimisation, and the difference between the good and bad cases is an implementation property, not a technology property. And because replacement hiring dominates gross hiring in this sector, a third of the private return is reallocation rather than new output — a wedge that every



correct firm-level business case will nonetheless miss, and that belongs in any assessment of what recruitment technology is doing to aggregate matching efficiency.

## REFERENCES

1. Autor, D. H. (2001). Wiring the labor market. *Journal of Economic Perspectives*, 15(1), 25–40.
2. Barber, A. E. (1998). *Recruiting Employees: Individual and Organizational Perspectives*. Thousand Oaks, CA: Sage Publications.
3. Blatter, M., Muehlemann, S., & Schenker, S. (2012). The costs of hiring skilled workers. *European Economic Review*, 56(1), 20–35.
4. Borowczyk-Martins, D., Jolivet, G., & Postel-Vinay, F. (2013). Accounting for endogeneity in matching function estimation. *Review of Economic Dynamics*, 16(3), 440–451.
5. Breaugh, J. A. (2013). Employee recruitment. *Annual Review of Psychology*, 64, 389–416.
6. Cowgill, B. (2020). *Bias and Productivity in Humans and Algorithms: Theory and Evidence from Résumé Screening*. Working Paper, Columbia Business School, New York.
7. European Parliament and Council of the European Union (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Annex III: high-risk AI systems in employment, workers management and access to self-employment. *Official Journal of the European Union*.
8. Federal Reserve Bank of St. Louis (2025). *The Beveridge curve's predictive power: why job vacancy types matter for monetary policy*. On the Economy, November 2025.
9. Gottfries, A. (2025). The Beveridge curve, matching, and labour market flows: a reinterpretation. *Oxford Bulletin of Economics and Statistics*, early view.
10. Greenhouse (2026). *The State of Hiring: Candidate and Recruiter Survey Findings 2026*. Reported in *SelectSoftwareReviews*, AI recruiting statistics 2026.
11. Guvenen, F., Kuruscu, B., Tanaka, S., & Wiczer, D. (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, 12(1), 210–244.
12. Insight Global (2026). *Hiring Manager AI Adoption Survey 2026*. Reported in *SelectSoftwareReviews*, AI recruiting statistics 2026.
13. Kuhn, P., & Mansour, H. (2014). Is internet job search still ineffective? *The Economic Journal*, 124(581), 1213–1233.
14. Ministry of Statistics and Programme Implementation, Government of India (2026). *Periodic Labour Force Survey (PLFS) Annual Report, January–December 2025*. New Delhi: National Statistics Office.
15. Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3), 397–415.
16. Muehlemann, S., & Strupler Leiser, M. (2018). Hiring costs and labor market tightness. *Labour Economics*, 52, 122–131.
17. NASSCOM (2026). *Technology Sector in India: Strategic Review 2026*. New Delhi: National Association of Software and Service Companies.
18. New York City Department of Consumer and Worker Protection (2023). *Local Law 144 of 2021: Automated Employment Decision Tools – Bias Audit and Notice Requirements*. New York.
19. Petrongolo, B., & Pissarides, C. A. (2001). Looking into the black box: a survey of the matching function. *Journal of Economic Literature*, 39(2), 390–431.
20. Pissarides, C. A. (2000). *Equilibrium Unemployment Theory* (2nd ed.). Cambridge, MA: MIT Press.
21. Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT '20)*, 469–481.
22. Rynes, S. L., Bretz, R. D., & Gerhart, B. (1991). The importance of recruitment in job choice: a different way of looking. *Personnel Psychology*, 44(3), 487–521.
23. Sastry, N. Subbu Krishna (2024). An impact of departmental structure on task achievement. *EPRA International Journal of Environmental Economics, Commerce and Educational Management*, 11(9).
24. Sastry, N. Subbu Krishna (2024). Revitalizing leadership to combat workforce burnout. *EPRA International Journal of Environmental Economics, Commerce and Educational Management*, 11(9).
25. Sastry, N. Subbu Krishna (2025). Building future-ready organizations: leadership strategies for young talent management. *EPRA International Journal of Research and Development*, 10(9), 88–91.
26. Sastry, N. Subbu Krishna (2025). 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. Subbu Krishna (2025). Women employee perspective strategies on Bangalore IT startups. *EPRA International Journal of Research and Development*, 10(7).
28. Sastry, N. Subbu Krishna, & Malliya M., Manjula (2025). Revolutionizing HR hiring with AI analytics: enhancing efficiency and attracting top talent. *EPRA International Journal of Research and Development*, 10(9), 125–131.
29. Sastry, N. Subbu Krishna, & Malliya M., Manjula (2025). Strategic HRM for economic resilience: skill diversification, labor market dynamics and technology disruption. *EPRA International Journal of Research and Development*, 10(11), 408–414.
30. Sastry, N. Subbu Krishna, & Selvarasu, A. (2014). A study of impact on performance appraisal on employee's engagement. *International Journal of Managerial Studies and Research*, 2(11), 10–22.
31. Schmidt, F. L., & Hunter, J. E. (1998). The validity and utility of selection methods in personnel psychology. *Psychological Bulletin*, 124(2), 262–274.
32. Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1), 25–49.
33. Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374.
34. Voltech HR Services (2026). *How Long Does It Take to Fill a White-Collar Role in India? Benchmarks for 2026*. Chennai.