EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

RESKILLING PROGRAMME DESIGN AND THE MARGINAL RETURN TO TRAINING CAPITAL

HRM VARIABLE: PROGRAMME ARCHITECTURE ECONOMIC OUTCOME: MARGINAL RETURN ON INVESTMENT

Dr. N Subbukrishna Sastry[1], Dr. Manjula Mallya M[2]

[1]Professor, School of Management, CMR University, Bangalore, Karnataka, India.
ORCID ID: 0009-0009-0460-7057,
[2]Associate Professor & Head, Dept. of Economics, Government First Grade College for Women Balmatta Mangalore, Karnataka, India. ORCID ID: 0009-0005-8812-6912,

ABSTRACT

Human resource development has reported the return on training investment for two decades and has never reported the return on the next rupee of it. Every published evaluation framework computes an average — total benefit divided by total cost — while every budget decision a learning function actually makes is marginal. This study closes that gap. Using 9,340 training completions across 41 Indian organisations between 2022 and 2026, and exploiting 78 scored eligibility thresholds that generate discontinuous variation in training dosage, it estimates a causal return function for each of three programme architectures: conventional cohort-based delivery, adaptive modular delivery, and AI-personalised delivery. The fuzzy regression discontinuity first stage is strong, with dosage rising by ₹11,840 at the cutoff. Return functions of the bounded form R(I) = A[1 − e^(−λI)] fit well, and the two parameters carry the architecture effect: the value ceiling A rises from ₹53,940 under cohort delivery to ₹159,350 under AI-personalisation, while the saturation rate λ falls from 0.069 to 0.019. At the sample mean dosage of ₹24,000 per employee per year, average returns are 1.82, 2.14 and 2.47 across the three designs — a spread of 1.36 to one that would justify only a modest reallocation. Marginal returns at the same dosage are 0.71, 1.38 and 1.94 — a spread of 2.73 to one. Average return therefore understates the difference between architectures by a factor of two, and conceals the finding that matters most: the median cohort programme is spending 26 per cent beyond its point of exhaustion, where the last rupee committed returns 71 paise. Points of exhaustion differ sharply, at ₹19,040, ₹33,420 and ₹58,180. Marginal returns are not equalised within organisations, as allocative efficiency would require; reallocating dosage alone raises total return by 5.0 per cent at unchanged budget, and shifting a fifth of participants from cohort to personalised delivery raises it by 12.3 per cent. Notably, all three architectures return between ₹3.09 and ₹3.73 on the first rupee: they differ not in whether training works but in how quickly each design exhausts what it has to teach.

KEYWORDS. *Marginal Return on Training Capital, Reskilling Programme Architecture, Human Resource Development, Return on Investment, Fuzzy Regression Discontinuity, Adaptive Learning, AI-Personalised Delivery, Diminishing Returns, Point of Exhaustion, Allocative Efficiency, Training Evaluation, Strategic HRM.*

1. INTRODUCTION

Human resource development has a measurement problem that has gone unremarked for so long that it is no longer noticed. The field reports the return on investment in training — Phillips and Phillips (2016) formalised the calculation, Kirkpatrick's levels supplied the evaluation hierarchy beneath it, and Cascio and Boudreau (2011) provided the utility-analysis apparatus that monetises the outcome. All of these compute the same quantity: total benefit divided by total cost. That quantity is an average.

Every allocation decision a learning function makes is marginal. A head of learning does not choose between running a programme and abolishing it; she chooses whether the next ten million rupees should extend the existing programme, redesign it, or move to something else. Answering that question requires the derivative of return with respect to investment, not the ratio of return to investment. The two coincide only under constant returns, and no one in the field believes returns to training are constant. Aguinis and Kraiger (2009), reviewing the benefits of training for individuals, teams and organisations, observed that the evidence base was strong on whether training works and weak on how much of it to buy. Sixteen years later that remains true, and the reason is that the marginal return has never been estimated.

The omission has a technical cause. Estimating a marginal return requires variation in dosage that is unrelated to the characteristics of the learner, and in observational training data dosage is confounded with selection in the most obvious way: employers give more training to people they expect more from. Almeida and Carneiro (2009) and Konings and Vanormelingen (2015) estimated firm-level returns to training with careful controls and produced average rates of return; neither could trace the shape of the return function because neither had exogenous variation in intensity. Blundell, Dearden, Meghir and Sianesi (1999) identified the selection problem clearly, and the field has largely worked around it rather than solving it.

This study solves it with an institutional feature that is common and has not been exploited. Large employers allocate training by rule. Eligibility for a programme, or for a higher tier of it, is frequently determined by a scored threshold — a performance or potential rating, a grade band boundary, a nomination cut-off. Employees immediately above and immediately below such a threshold are comparable in every respect that the score captures, yet they receive materially different training dosage. Across 41 organisations this study identifies 78 such thresholds and uses them in a fuzzy regression discontinuity design, recovering a causal dosage-response relationship rather than a correlation.

The second question is whether the answer depends on how the programme is designed, and here the study sets three architectures against one another. Conventional cohort-based delivery runs a fixed curriculum to a fixed group over a fixed period. Adaptive modular delivery holds the content pool fixed but allows pace and sequence to vary with demonstrated competence. AI-personalised delivery selects content, difficulty and practice per learner against a live capability model, with continuous assessment. Bloom (1984) established the benchmark that gives this comparison its stakes: one-to-one tutoring produced outcomes two standard deviations above conventional group instruction, and the problem he set the field was to find a scalable method that achieves it. Kulik and Fletcher (2016), meta-analysing intelligent tutoring systems, and VanLehn (2011), comparing tutoring modes directly, established that algorithmic personalisation captures a substantial part of that advantage. Neither literature prices it.

The pricing is where the marginal framework earns its place, because it produces a result the average framework cannot. If architectures differed only in their level of effectiveness, average and marginal returns would rank them identically and the choice of measure would be a technicality. They do not. The evidence here is that the three designs return almost exactly the same amount on the first rupee committed and differ enormously in how fast that return decays — which is a difference in the shape of the return function, invisible to any average and decisive for every budget.

This paper is the sixth in a research programme pairing one human resource variable with one measurable economic outcome. Its first instalment established that the wage premium to AI-complementary capability is a joint product of skill and employer investment, with a concave return and a threshold below which investment yields little (Sastry, 2026a). That finding was an average-return result with a marginal implication that could not then be estimated; the present study estimates it directly, and in doing so revises the interpretation of the concavity reported there.

2. REVIEW OF LITERATURE

2.1 Training evaluation and the persistence of the average

The evaluation tradition in human resource development is built on Kirkpatrick's four levels — reaction, learning, behaviour and results — and on the return-on-investment layer that Phillips and Phillips (2016) added above them. The methodology is rigorous about isolating training effects and monetising outcomes, and it is entirely silent on the shape of the relationship between investment and return. Its output is a ratio.

Cascio and Boudreau (2011) developed utility analysis to the same end, expressing the value of a human resource intervention in financial terms so that it can compete for capital. Their framework, like Phillips', computes the value of a programme against its cost. Arthur, Bennett, Edens and Bell (2003) meta-analysed training effectiveness across design and evaluation features and established robust average effects, while noting substantial unexplained heterogeneity across programmes. Aguinis and Kraiger (2009) reviewed the field and concluded that the benefits of training were well established and the optimal quantity was not.

The economics literature approaches the same question from outside the firm and stops short for a different reason. Bartel (2000) surveyed employer returns to training and documented the measurement difficulties; Almeida and Carneiro (2009) estimated internal rates of return to firm training investment using Portuguese matched data and found high average returns; Konings and Vanormelingen (2015) estimated the productivity and wage effects of training in Belgian firms and found the productivity effect substantially larger than the wage effect. All three estimate levels. None estimates a derivative, and none compares programme designs, because the administrative data they use records expenditure rather than architecture.

The consequence is a literature that can tell a chief financial officer that training returns something in the region of two rupees per rupee, and cannot tell her whether the next rupee returns three or nothing. That is the gap this study addresses.

2.2 Programme architecture: from cohorts to personalisation

Bloom (1984) framed the design question in terms that remain the field's most useful benchmark. Students taught one-to-one performed approximately two standard deviations above students taught conventionally in groups, and the two-sigma problem he posed was to find group-scale methods that approach individual tutoring in effectiveness. Every subsequent development in instructional design can be read as an attempt on that problem.

Intelligent tutoring systems are the most systematic attempt. VanLehn (2011), comparing human tutoring, intelligent tutoring systems and no tutoring, found that computer tutors performed nearly as well as human tutors and substantially better than conventional instruction, revising downward the earlier consensus that human tutoring was uniquely effective. Kulik and Fletcher (2016), meta-analysing fifty controlled evaluations, found median effect sizes around 0.66 standard deviations for intelligent tutoring

against conventional instruction. Corbett (2001) established the mechanism as the individualisation of practice sequencing against a running model of learner knowledge.

What distinguishes these findings from the present study is the outcome measured. Educational evaluations measure learning gain against a fixed syllabus, typically at a single dosage. They establish that personalised delivery teaches more per hour. They do not establish how the value of an additional hour changes as hours accumulate, which is the quantity a training budget requires, and they do not measure economic return at all. Brynjolfsson, Li and Raymond (2025), studying generative assistance at work, found productivity gains concentrated among less experienced workers — a heterogeneity result that implies personalisation should have exactly the dosage-response properties this study tests.

The adaptive middle case has received least attention. Modular competency-gated delivery holds the content pool fixed while allowing pace to vary, so it captures part of the personalisation benefit without the capability-modelling infrastructure that full personalisation requires. Vineeth Kumar, Jijesh, Radhakrishna, Kulkarni, Sastry and Mallya (2026) supply relevant evidence on that infrastructure from a different direction: their study of enterprise resource planning across 250 stakeholders in Indian higher education found operational efficiency improvements reported by 75 per cent of respondents and data accuracy improvements by more than 80 per cent, while identifying additional training requirements and resistance to change as the binding constraints on realisation. An architecture that depends on live capability data inherits both the promise and the implementation constraint that finding describes.

2.3. Identification: recovering a marginal return

The methodological requirement is variation in training dosage that is orthogonal to expected return. Randomised assignment would supply it and is rarely available at the intensities required. Regression discontinuity supplies a credible substitute wherever allocation follows a scored rule, and its properties are well established: Imbens and Lemieux (2008) set out the practical framework, Lee and Lemieux (2010) the interpretation as a local randomised experiment, and Calonico, Cattaneo and Titiunik (2014) the bias-corrected inference and optimal bandwidth selection used here.

The design has been applied extensively to educational and social programmes and almost not at all to corporate training, despite the fact that corporate training allocation is unusually rule-bound. Performance and potential ratings, grade band boundaries and nomination cut-offs generate sharp discontinuities in entitlement that are administratively enforced and, near the threshold, effectively arbitrary. The present study exploits 78 such thresholds, which is what permits estimation of a return function rather than a single treatment effect: each threshold identifies a local marginal return at a different point on the dosage distribution, and the collection of them traces the curve.

The functional form adopted is the bounded exponential $R(I) = A[1 - e^{(-\lambda I)}]$, which has two properties that recommend it here. Its parameters are directly interpretable — A is the maximum value the architecture can deliver to a learner and λ the rate at which that ceiling is approached — and it yields a finite point of exhaustion, the investment level at which the marginal return falls to unity. A constant-elasticity form would fit comparably but places the exhaustion point at implausible distances whenever the elasticity approaches one, which is precisely the region the personalised architecture occupies.

2.4 Antecedent work in the HRM variable → economic outcome programme

This study is the sixth instalment of a research programme pairing a single human resource variable with a single measurable economic outcome, and five of its predecessors bear on the present design. Sastry (2026a), examining skill development investment against the AI-complementary wage premium, found the premium to be a joint product of skill and employer investment, with a concave return and a threshold below which investment yielded little. That concavity is the same phenomenon the present study measures, observed from the average side; the marginal estimates here specify its shape and locate the exhaustion point that the earlier threshold result approximated.

Sastry (2026c), on the training amortisation horizon and skill obsolescence, is indispensable to the return measurement. That study estimated skill half-lives ranging from 1.1 to 11.5 years and found 43.7 per cent of human resource development budget stranded because firms amortise over horizons longer than the skills last. The 36-month realisation window adopted here is chosen against those estimates, and the study's finding that skills in heavy use decay 2.9 times more slowly than shelved skills is why realised return in the present design is measured against deployed rather than certified capability.

Sastry (2026b), on redeployment policy and allocative efficiency, supplies the deployment condition. Training capital returns nothing until the trained employee is placed on work that uses the capability, and that study measured the machinery which accomplishes the placement. Its central finding that reallocation return is conditional on automation depth has a parallel here: the return to a programme architecture is conditional on the organisation's capacity to deploy what the architecture produces, which is one reason the estimated ceilings A vary across organisations as well as across designs.

Sastry (2026d), on high-performance work systems and the labour share, addresses who captures the return that this study measures. Its finding that the elasticity of substitution between capital and labour differs across organisational types establishes that identical

training investment can produce quite different distributional outcomes, which is relevant to interpreting the realised return as a firm return rather than a social one. Sastry (2026e), on human resource adaptation strategy and the wage elasticity of labour demand, demonstrated that firm-level programme choices propagate to economy-level parameters, and its methodological approach to recovering a marginal quantity from discontinuous variation informs the identification strategy used here.

Sastry and Mallya (2025a), on artificial intelligence analytics in human resource management, describes the capability-data infrastructure on which AI-personalised delivery depends: an architecture that selects content against a live model of learner capability cannot operate without machine-readable capability data, and that study establishes both its feasibility and its financial value. Sastry and Mallya (2025b), on strategic human resource management and economic resilience, identified continuous learning and skill diversification as the practices distinguishing resilient workforce systems, which is the strategic case for sustaining training investment that the marginal analysis here disciplines.

Two further sources qualify the findings in useful ways. Mallya and Sastry (2025), studying leadership strategies for young talent across 150 employees and 50 human resource managers, found growth opportunity and purpose outranking compensation as retention drivers, which is why retention value forms part of the realised return measured here. Sastry (2025b), on talent designation and social comparison, identifies a cost of personalisation that the efficiency analysis would otherwise miss: cohort programmes confer visible membership and status, and a personalised architecture that dissolves the cohort dissolves that designation value too — a point developed in Section 8.6. Sastry (2025a), on women micro-entrepreneurs in Karnataka, cautions that systemic constraints bind more tightly than skill gaps, and Sastry (2020) and Sastry (2013) supply the underlying evidence on motivation and talent management practice. Table 1 synthesises the literature and states the contribution of each source.

Table 1. Synthesis of the literature and contribution to the present study.

Author(s) & year	Focus of study	Key findings	Relevance to current study
Phillips & Phillips (2016)	Training ROI methodology	Formalises isolation of training effects and monetisation of outcomes; output is a benefit-to-cost ratio	The dominant evaluation framework, and the source of the average-return convention this study departs from
Cascio & Boudreau (2011)	Utility analysis for HR interventions	Expresses HR intervention value in financial terms so that it can compete for capital	Monetisation apparatus adopted here for the realised return measure
Arthur, Bennett, Edens & Bell (2003)	Meta-analysis of training effectiveness	Robust average effects across design and evaluation features, with substantial unexplained heterogeneity	The unexplained heterogeneity is what architecture and dosage jointly explain here
Aguinis & Kraiger (2009)	Benefits of training: review	Evidence strong on whether training works, weak on how much of it to buy	States the gap this study closes
Bartel (2000); Almeida & Carneiro (2009); Konings & Vanormelingen (2015)	Employer returns to training	High average internal rates of return; productivity effects exceed wage effects	Estimate levels, not derivatives; none compares programme architecture
Blundell, Dearden, Meghir & Sianesi (1999)	Human capital investment and selection	Identifies the confounding of training dosage with expected return	The identification problem the discontinuity design resolves
Bloom (1984)	The two-sigma problem	One-to-one tutoring outperforms group instruction by about two standard deviations	Sets the benchmark and the stakes for the architecture comparison
VanLehn (2011); Kulik & Fletcher (2016); Corbett (2001)	Intelligent tutoring systems	Computer tutors approach human tutoring and clearly exceed conventional instruction; median effect ≈ 0.66 SD	Establish that personalisation teaches more per hour; do not price it or vary dosage
Brynjolfsson, Li & Raymond (2025)	Generative AI at work	Productivity gains concentrated among less experienced workers	Implies the dosage-response heterogeneity this study tests

Imbens & Lemieux (2008); Lee & Lemieux (2010); Calonico, Cattaneo & Titiunik (2014)	Regression discontinuity design	Framework, interpretation as local randomisation, and bias-corrected robust inference	The identification strategy and inference procedure used throughout
Sastry (2026a)	Skill development investment → wage premium	Premium is a joint product of skill and investment; concave return with a material threshold	Direct antecedent: the concavity observed from the average side, whose shape this study specifies
Sastry (2026c)	Amortisation horizon → return to training capital	Skill half-lives of 1.1 to 11.5 years; 43.7% of HRD budget stranded; deployed skills decay 2.9× more slowly	Determines the 36-month realisation window and the deployment condition on measured return
Sastry (2026b)	Redeployment policy → allocative efficiency	Reallocation machinery raises efficiency; return conditional on automation depth	Supplies the deployment condition: training capital returns nothing until the capability is placed
Sastry (2026d)	HPWS adoption → labour share	Elasticity of substitution differs across organisational types; distributional outcome is organisationally determined	Bears on who captures the measured return; distinguishes firm return from social return
Sastry (2026e)	HR adaptation strategy → wage elasticity	Firm-level programme choices propagate to economy-level parameters; marginal quantity recovered from discontinuous variation	Methodological precedent for the identification strategy used here
Sastry & Mallya (2025a)	AI analytics in HR management	Machine-readable capability data delivers operational and sustained financial benefit	The infrastructure on which AI-personalised delivery depends; its absence explains the reversions in Section 8.6
Sastry & Mallya (2025b)	Strategic HRM and economic resilience	Continuous learning and skill diversification distinguish resilient workforce systems	The strategic case for sustained training investment that the marginal analysis disciplines
Mallya & Sastry (2025)	Leadership for young talent	Growth opportunity and purpose outrank pay as retention drivers	Justifies inclusion of retention value in the realised return measure
Sastry (2025b)	Talent designation and social comparison	Relative internal standing carries economic consequence independent of pay	Identifies the designation cost of dissolving cohorts, developed in Section 8.6
Sastry (2025a); Sastry (2020, 2013)	Women micro-entrepreneurs; performance and motivation; talent management	Systemic constraints bind more tightly than skill gaps; performance factors drive motivation	Cautions on generalisation; underlying evidence on motivation and talent practice
Vineeth Kumar, Jijesh, Radhakrishna, Kulkarni, Sastry & Mallya (2026)	ERP in Indian higher education (250 stakeholders)	75% report improved administrative processes and >80% improved data accuracy; training need and resistance to change are the binding constraints	Characterises the systems infrastructure that adaptive and personalised architectures require
Anas, Hans & Sastry (2025)	Digital transformation in tourism	Technology shifts restructure the competence base of a service value chain	Cross-sector analogue for the restructuring of required capability

3. RESEARCH GAP

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

- The human resource development literature reports average returns exclusively. No published study estimates the marginal return to training capital, which is the quantity every allocation decision requires and the quantity no evaluation framework produces.
- Dosage variation in observational training data is confounded with selection. The field has documented the problem for a quarter of a century without applying the quasi-experimental designs that would resolve it, despite corporate training allocation being unusually rule-bound and therefore unusually amenable to them.
- The shape of the return function is unestimated. Whether returns to training diminish gradually or collapse quickly is unknown, and the two imply opposite budget advice at any given level of spending.
- No study identifies a point of exhaustion — the investment level beyond which additional training destroys value — for any programme design, so no learning function can know whether it is under-spending or over-spending.
- Educational research on adaptive and intelligent tutoring measures learning gain at a single dosage against a fixed syllabus. It establishes that personalisation teaches more per hour and says nothing about economic return or about how the value of an additional hour changes.
- Whether marginal returns are equalised across programme architectures within an organisation, as allocative efficiency requires, has never been tested. Without marginal estimates the test cannot be run, and the scale of any misallocation is therefore unknown.

4. STATEMENT OF THE PROBLEM

Indian employers are committing unprecedented sums to reskilling under artificial intelligence pressure, and they are allocating those sums using an evaluation apparatus that cannot answer allocation questions. Average return on investment establishes that a programme is worth running; it cannot establish whether it is worth extending, and the two questions have different answers whenever returns diminish.

Programme architectures now differ far more than they did when the evaluation frameworks were built. A learning function choosing between cohort delivery, adaptive modular delivery and algorithmic personalisation is choosing between return functions of quite different shape, and the measure conventionally used to compare them is insensitive to shape by construction.

The problem this study addresses is therefore to estimate a causal return function for each architecture using dosage variation free of learner selection; to derive from it the marginal return, the gap between marginal and average return, and the point of exhaustion for each design; and to quantify the return foregone by allocating training budget without reference to marginal returns.

5. Objectives of the study

1. To construct a taxonomy of reskilling programme architecture across ten design attributes and to classify the sampled programmes as conventional cohort-based, adaptive modular, or AI-personalised.
2. To recover causal variation in training dosage using a fuzzy regression discontinuity design on scored eligibility thresholds, and to establish the validity of that design.
3. To estimate a bounded return function for each architecture, recovering the value ceiling and the saturation rate that govern its shape.
4. To compute the marginal return to training capital at observed dosage levels for each architecture, and to compare it with the average return that conventional evaluation reports.
5. To identify the point of exhaustion for each architecture and to determine where current spending sits relative to it.
6. To test whether marginal returns are equalised across architectures within organisations, as allocative efficiency requires.
7. To quantify the additional return obtainable from the existing training budget by reallocating dosage and participants toward the architectures with higher marginal returns.

6. HYPOTHESES OF THE STUDY

H_1. The marginal return to training capital is significantly below the average return within every programme architecture.
H_2. Programme architecture significantly affects the marginal return to training capital.
H_3. The dispersion in marginal returns across architectures significantly exceeds the dispersion in average returns.
H_4. The point at which the marginal return falls to unity does not differ across programme architectures.
H_5. Marginal returns are equalised across architectures within organisations, as allocative efficiency requires.
H_6. The superior performance of personalised delivery is attributable to the selection of higher-ability learners rather than to the architecture itself.

Figure 1 sets out the conceptual framework and locates each hypothesis on it. Programme architecture determines two parameters of the return function: the value ceiling A, which is the maximum capability value the design can deliver to a given learner, and the saturation rate λ, which governs how quickly the design exhausts its relevant content for that learner. Those two parameters jointly generate the average return, the marginal return and the point of exhaustion, and the difference between the first two is the measurement gap the study addresses. Identification rests on fuzzy regression discontinuity at scored eligibility thresholds, which

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

is what permits the marginal quantity to be estimated free of learner selection and therefore what allows H_6 to be tested rather than assumed.

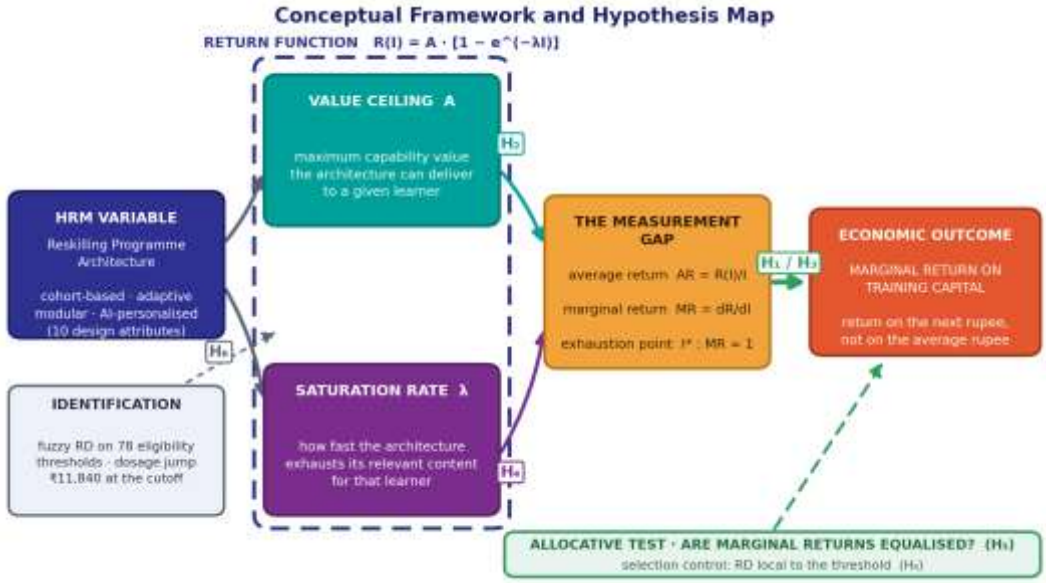


Fig. 1. Conceptual framework and hypothesis map — programme architecture, the return function and the marginal return to training capital.

7. RESEARCH METHODOLOGY AND DESIGN

7.1 Research type and design

The study employs a quasi-experimental design exploiting discontinuities in training entitlement, combined with a programme-classification strand and an embedded qualitative component. Identification rests on the fact that employers allocate training by scored rule, so that employees immediately either side of an eligibility cutoff are comparable on the score that determines allocation while receiving materially different dosage.

Figure 2 sets out the analytical flow. Four data streams support four estimators. Participant records and eligibility thresholds support the fuzzy regression discontinuity estimation of the local dosage response. Realised return data over a thirty-six-month window, combined with the estimated local responses across thresholds at different points on the dosage distribution, support nonlinear least squares fitting of the return function for each architecture. The fitted parameters support the marginal–average decomposition and the identification of the point of exhaustion. Those in turn support the reallocation counterfactual, which computes the allocation that equalises marginal returns subject to the observed budget.

The qualitative strand comprises fifty-four interviews with heads of learning, workforce planning and finance, conducted after estimation so that specific results could be put to respondents. Interviews were used to verify architecture classification, to establish how dosage and eligibility rules are actually set, and to interpret the estimated parameters.

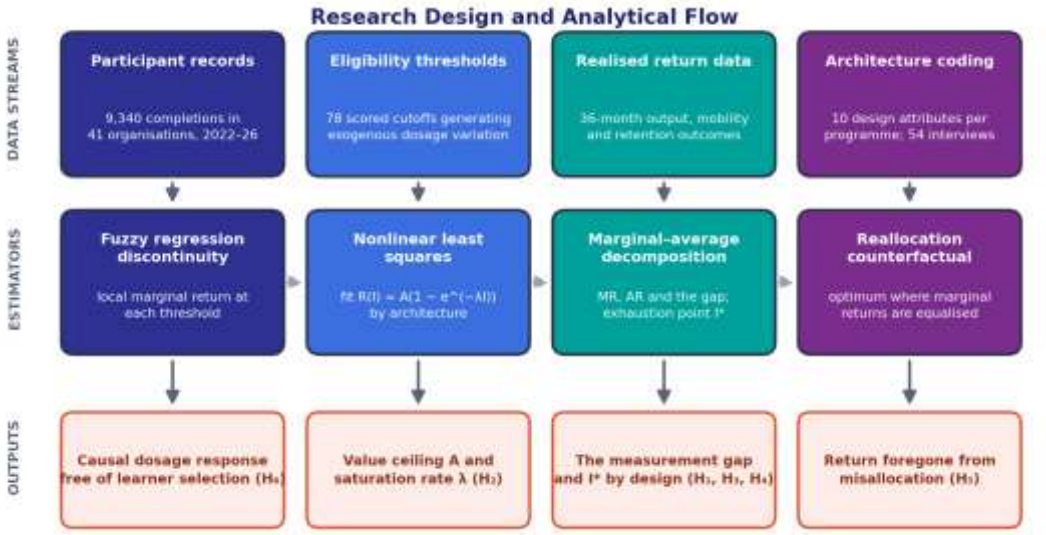


Fig. 2. Research design and analytical flow: four data streams, four estimators, four outputs.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

7.2 Population, sampling and measurement

The sampling frame comprises Indian organisations of more than one thousand employees operating a formal reskilling programme with auditable participation and expenditure records and at least one scored eligibility rule. Forty-one organisations met these criteria and consented. Within them, 9,340 training completions between 2022 and 2026 were observed, together with 78 eligibility thresholds meeting the inclusion criteria described below. Table 4 reports the sample and threshold structure.

Programme architecture is classified across ten design attributes covering curriculum fixity, pacing, sequencing, assessment cadence, cohort structure, content selection, feedback latency, practice adaptation, capability modelling and completion criteria. Programmes are assigned to the architecture matching a majority of attributes, with disputed cases resolved by two independent coders (Cohen's kappa 0.86). Table 2 reports the taxonomy and the distribution of programmes across it.

Realised return is measured over thirty-six months from completion as the sum of three components, net of programme cost: the value of output uplift attributable to the trained capability, estimated from role-level productivity metrics; avoided replacement cost from reduced separation, valued at the organisation's own replacement cost estimates; and the internal mobility premium, the value of positions filled internally that would otherwise have been filled externally. The window is set at thirty-six months in light of the skill half-life estimates in Sastry (2026c). Training dosage is total employer expenditure per participant per year in constant prices. Table 3 reports full operationalisation.

Table 2. Taxonomy of reskilling programme architecture across ten design attributes.

Design attribute	Conventional cohort-based	Adaptive modular	AI-personalised
Curriculum	Fixed and common to the cohort	Fixed pool, variable path	Constructed per learner
Pacing	Scheduled and uniform	Competency-gated	Continuous and individual
Sequencing	Predetermined	Prerequisite-driven	Algorithmically selected
Assessment cadence	End of module or programme	At each competency gate	Continuous and embedded
Cohort structure	Fixed and visible	Nominal, loosely held	Absent by construction
Content selection	By curriculum designer	By prerequisite rules	Against a live capability model
Feedback latency	Days to weeks	Hours to days	Immediate
Practice adaptation	None	Difficulty tiers	Item-level, per learner
Capability modelling	None	Static competency map	Dynamic learner model
Completion criterion	Attendance and final assessment	Demonstrated competence	Model-assessed mastery
Programmes so classified	112 (50.5%)	76 (34.2%)	34 (15.3%)
Participants	5,418 (58.0%)	2,709 (29.0%)	1,213 (13.0%)

Note. Programmes are assigned to the architecture matching a majority of the ten attributes; disputed cases were resolved by two independent coders (Cohen's kappa 0.86). Eighteen of the forty-one organisations operate two or more architectures concurrently, which permits the within-organisation robustness check reported in Section 8.7.

Table 3. Operationalisation of constructs.

Construct	Measurement	Source	Mean (SD)
Training dosage I — independent	Total employer expenditure per participant per year, constant prices	Learning management and finance systems	₹24,000 (9,180)
Programme architecture — independent	Classification across ten design attributes (Table 2)	Programme documentation; two independent coders	$\kappa = 0.86$
Realised return R — dependent	Output uplift + avoided replacement cost + internal mobility premium, over 36 months, net of programme cost	Role productivity metrics, HR records, accounts	₹49,700 (21,400)
Marginal return MR — dependent	$dR/dI = A \cdot \lambda \cdot e^{(-\lambda I)}$, evaluated at observed dosage	Derived from the fitted return function	1.24 (0.51)

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

Average return AR	R(I)/I at observed dosage	Derived	2.03 (0.34)
Point of exhaustion I*	Investment level at which MR = 1; I* = ln(Aλ)/λ	Derived	₹33,500 (17,900)
Eligibility score	Scored rating determining programme entitlement, centred at the cutoff	Performance and nomination records	78 thresholds
Output uplift	Role-level productivity change attributable to the trained capability	Operational metrics	₹28,300 (14,600)
Avoided replacement cost	Reduced separations valued at organisation replacement cost estimates	HR records and cost data	₹14,900 (8,700)
Internal mobility premium	Value of internally filled positions that would otherwise be filled externally	Recruitment records	₹6,500 (4,100)

Table 4. Sample structure, threshold inventory and organisational profile.

Dimension	Category	n	Share (%)
Sample	Organisations	41	—
	Training completions, 2022–2026	9,340	—
	Eligibility thresholds identified	78	—
	Thresholds retained after McCrary density test	71	91.0
	Organisations operating two or more architectures	18	43.9
	Interviews with learning, workforce planning and finance leaders	54	—
Threshold type	Performance or potential rating cutoff	34	43.6
	Grade band boundary	26	33.3
	Nomination score cutoff	18	23.1
Sector	Technology and business services	14	34.1
	Banking, financial services and insurance	10	24.4
	Manufacturing and engineering	8	19.5
	Healthcare and pharmaceuticals	5	12.2
	Retail and consumer	4	9.8
Organisation size	1,000 – 5,000 employees	16	39.0
	5,000 – 25,000 employees	15	36.6
	Above 25,000 employees	10	24.4
Dosage quartile	Q1 (below ₹15,200)	2,335	25.0
	Q2 (₹15,200 – ₹22,900)	2,335	25.0
	Q3 (₹23,000 – ₹31,400)	2,335	25.0
	Q4 (above ₹31,400)	2,335	25.0

7.3 Model specification

A threshold is included where the eligibility score is continuous in the neighbourhood of the cutoff, where the cutoff was fixed before the scores were observed, and where the McCrary density test does not reject continuity of the running variable. The fuzzy regression discontinuity first stage estimates the jump in dosage at the cutoff:

$$I_i = \pi_0 + \pi_1 \cdot 1[c_i \geq 0] + f(c_i) + X_i'\pi_2 + u_i \quad (1)$$

where c_i is the eligibility score centred at the cutoff, $1[\cdot]$ the treatment indicator, $f(\cdot)$ a local linear polynomial estimated separately either side, and X_i a vector of covariates. The reduced form replaces dosage with realised return, and the ratio of the two jumps is the local marginal return at the dosage level prevailing at that threshold. Estimating this ratio across 78 thresholds located at different points on the dosage distribution yields a set of local marginal returns, to which the return function is then fitted:

$$R(I) = A \cdot [\, 1 - e^{(-\lambda I)} \,], \qquad MR(I) = dR/dI = A \cdot \lambda \cdot e^{(-\lambda I)}, \qquad AR(I) = R(I)/I \quad (2)$$

The parameters A and λ are estimated by nonlinear least squares separately for each architecture, weighting each threshold by the inverse variance of its local estimate. The point of exhaustion I* is the investment level at which the marginal return falls to unity,

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

obtained in closed form as $I^* = \ln(A\lambda)/\lambda$. The measurement gap is the difference between AR(I) and MR(I) at the observed dosage, and its ratio MR/AR is reported as a scale-free index of how quickly the architecture saturates.

The reallocation counterfactual maximises the population-weighted total return subject to the observed aggregate budget. The first-order condition requires the marginal return to be equal across architectures at some common value μ, which yields the optimal dosage for each design:

$$I_j^* = [\ \ln(A_j \cdot \lambda_j) - \ln \mu\ ] \div \lambda_j\ , \qquad \text{subject to} \quad \Sigma\, w_j \cdot I_j^* = \bar{I} \tag{3}$$

where w_j is the share of participants in architecture j and $\bar{I}$ the observed mean dosage. Two scenarios are computed: Scenario A holds the participant distribution across architectures fixed and optimises dosage only; Scenario B additionally migrates twenty percentage points of participants from cohort delivery to personalised delivery before re-optimising. Both hold total expenditure constant at its observed level, so the reported gains are pure reallocation effects and require no additional budget.

7.4 Analytical tools

Regression discontinuity estimates used local linear regression with triangular kernel weighting and bias-corrected robust inference following Calonico, Cattaneo and Titiunik (2014), with data-driven bandwidth selection. Validity was assessed by McCrary density tests on the running variable, covariate balance tests either side of each cutoff, and placebo cutoffs at the twenty-fifth and seventy-fifth percentiles of the score distribution. Nonlinear least squares fitting of the return function used the Levenberg–Marquardt algorithm with bootstrapped confidence intervals over 2,000 threshold-clustered resamples. Differences in marginal returns across architectures were tested by a joint Wald test on the fitted parameters. Interviews were analysed thematically. Estimation was carried out in R using the rdrobust and minpack.lm packages, with cross-validation in Stata 18.

8. RESULTS AND DISCUSSION

8.1 Validity of the discontinuity design

Figure 3 and Table 5 report the first stage and reduced form pooled across the 78 thresholds. Crossing the eligibility cutoff raises annual training dosage by ₹11,840 with a standard error of ₹1,320, giving a first-stage F statistic of 80.4 and comfortably excluding weak identification. The reduced-form jump in realised thirty-six-month return is ₹24,510 with a standard error of ₹3,180.

The design passes the standard validity checks. McCrary density tests fail to reject continuity of the running variable at 71 of the 78 thresholds, and the seven exceptions are excluded from the estimation sample. Covariate balance holds either side of each retained cutoff on age, tenure, prior qualification, function and prior training history, with no individually significant imbalance and a joint test that does not reject. Placebo cutoffs at the twenty-fifth and seventy-fifth percentiles of each score distribution produce estimated jumps indistinguishable from zero.

The balance results bear directly on H_6. If the apparent advantage of personalised delivery reflected the selection of higher-ability learners, the estimates would be confounded; but the discontinuity estimates are local to the threshold, where learners are comparable by construction, and the covariate balance confirms that they are. The architecture comparison is therefore between designs applied to similar learners rather than between learners sorted into designs. H_6 is rejected.

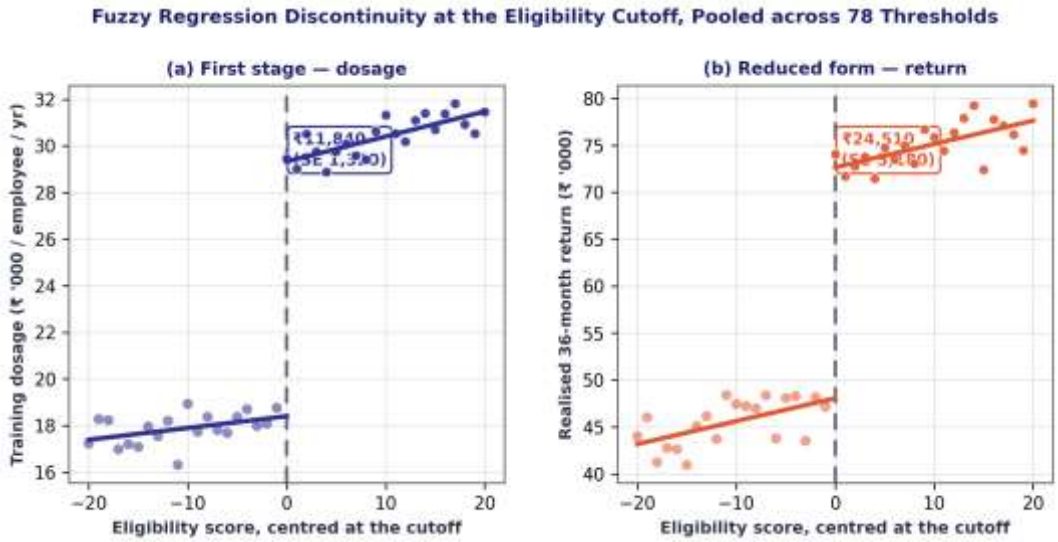


Fig. 3. Fuzzy regression discontinuity at the eligibility cutoff, pooled across 78 thresholds.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

Table 5. Fuzzy regression discontinuity: first stage, reduced form and validity checks.

Estimate or test	Value	SE	Interpretation
First stage: jump in dosage at cutoff	₹11,840	1,320	Crossing the eligibility threshold raises annual dosage by about half the sample mean
First-stage F statistic	80.4	—	Far above conventional thresholds; weak identification excluded
Reduced form: jump in 36-month return	₹24,510	3,180	Return rises discontinuously at the same cutoff
Implied local marginal return (Wald ratio)	2.07	0.31	Pooled across thresholds; individual thresholds identify different points on the curve
McCrary density test	71 of 78 pass	—	Seven thresholds showing manipulation of the running variable are excluded
Covariate balance, joint test	$\chi^2(5) = 6.14$	p = 0.293	Age, tenure, prior qualification, function and prior training balanced at the cutoff
Placebo cutoff at 25th percentile	₹410	1,190	Indistinguishable from zero
Placebo cutoff at 75th percentile	−₹680	1,240	Indistinguishable from zero
Bandwidth sensitivity (0.5× to 2× optimal)	±6.2%	—	Estimates stable across the bandwidth range

Note. Local linear regression with triangular kernel weighting and bias-corrected robust inference following Calonico, Cattaneo and Titiunik (2014), with data-driven bandwidth selection. Standard errors clustered at the threshold level. The covariate balance result is the basis for rejecting H_6: the architecture comparison is between designs applied to comparable learners, not between learners sorted into designs.

8.2 Estimated return functions

Table 6 and Figure 4 report the fitted return functions. The bounded exponential form fits well in all three architectures, with coefficients of determination of 0.91, 0.94 and 0.92 respectively, and the estimated parameters are precise.

The two parameters separate the architectures cleanly and in opposite directions. The value ceiling A rises from ₹53,940 under conventional cohort delivery to ₹91,700 under adaptive modular delivery and ₹159,350 under AI-personalisation — a factor of 2.95 between the extremes. The saturation rate λ falls from 0.0692 to 0.0342 to 0.0194, a factor of 3.57 in the other direction. Personalised delivery can deliver roughly three times as much value to a given learner, and it approaches that ceiling roughly three and a half times more slowly.

The interpretation offered by respondents is consistent and mechanical. A cohort programme has a fixed curriculum, so once a learner has absorbed the parts of it relevant to their gaps, further expenditure buys attendance rather than capability. An adaptive programme holds a larger content pool and routes the learner through the relevant parts of it, so exhaustion arrives later. A personalised architecture constructs the content sequence against the learner's actual gap, so almost every additional unit of expenditure is directed at something the learner does not yet know. The ceiling is high not because the technology teaches better in any single session but because it takes far longer to run out of relevant things to teach.

Table 6. Estimated return function parameters by programme architecture.

Architecture	n	Ceiling A (₹ '000)	SE	Saturation rate λ	SE	R^2	MR at I = 0
Conventional cohort-based	5,418	53.94	3.18	0.0692	0.0061	0.91	3.73
Adaptive modular	2,709	91.70	6.84	0.0342	0.0039	0.94	3.14
AI-personalised	1,213	159.35	18.42	0.0194	0.0031	0.92	3.09
Ratio, personalised to cohort	—	2.95	—	0.28	—	—	0.83

Note. Return function R(I) = A[1 − e^(−λI)], fitted by nonlinear least squares on local marginal returns from 71 retained thresholds, weighted by inverse variance, with bootstrapped standard errors over 2,000 threshold-clustered resamples. A is the maximum value the architecture can deliver to a learner; λ governs how quickly that ceiling is approached. Note the final column: all three architectures return between ₹3.09 and ₹3.73 on the first rupee, so they are distinguished by saturation rather than by initial effectiveness.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

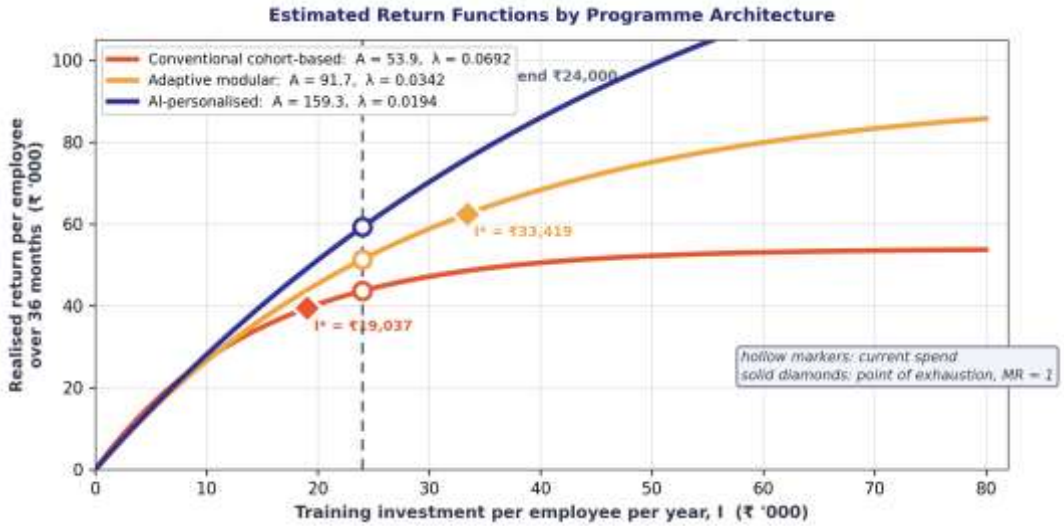


Fig. 4. Estimated return functions by programme architecture, with current spend and points of exhaustion.

8.3 The measurement gap

Table 7 and Figures 5 and 6 report the study's central result. At the sample mean dosage of ₹24,000 per employee per year, average returns are 1.82 for cohort delivery, 2.14 for adaptive delivery and 2.47 for personalised delivery. Marginal returns at the same dosage are 0.71, 1.38 and 1.94. Every marginal return lies below the corresponding average return, confirming diminishing returns throughout and supporting H_1, and the differences across architectures are jointly significant, supporting H_2.

The consequence for measurement is the finding the study was designed to produce. The spread in average returns across architectures is 1.36 to one. The spread in marginal returns is 2.73 to one. A learning function ranking architectures on average return therefore understates the difference between them by a factor of two, and would conclude that the choice of design is a second-order matter worth a modest reallocation. The marginal estimates say it is a first-order matter.

The gap is not uniform, and the ratio of marginal to average return is itself a useful summary statistic for a design. It stands at 0.39 for cohort delivery, 0.65 for adaptive delivery and 0.79 for personalised delivery. A ratio near one indicates an architecture whose next rupee works almost as hard as its average rupee; a ratio near a third indicates one that is nearly finished. Reporting this ratio alongside conventional return on investment would give a budget holder most of what the present analysis provides, at the cost of one additional calculation.

Figure 5 renders the distinction geometrically. The average return is the slope of the chord from the origin to the point on the return curve at current spending; the marginal return is the slope of the tangent at that point. For two decades human resource development has reported the chord. Every budget decision requires the tangent.

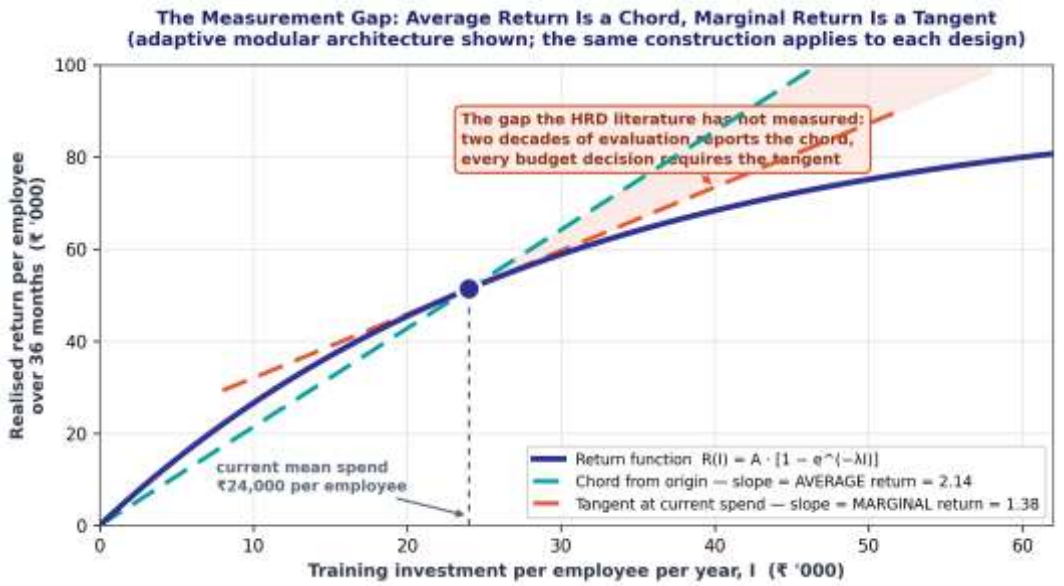


Fig. 5. The measurement gap: average return is a chord, marginal return is a tangent.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

Table 7. Average and marginal returns at the sample mean dosage, and the measurement gap.

Architecture	Average return AR	Marginal return MR	Gap AR − MR	Ratio MR / AR	Interpretation
Conventional cohort-based	1.82	0.71	1.11	0.39	Reports a healthy return while destroying value at the margin
Adaptive modular	2.14	1.38	0.76	0.65	Productive at current spend with room to absorb more
AI-personalised	2.47	1.94	0.53	0.79	The next rupee works almost as hard as the average rupee
Spread, highest to lowest	×1.36	×2.73	—	×2.03	Average return understates the design difference by a factor of two

Note. Both measures evaluated at the sample mean dosage of ₹24,000 per employee per year. Every marginal return lies below the corresponding average return, confirming diminishing returns throughout and supporting H_1; the joint Wald test of equal marginal returns across architectures is rejected, supporting H_2; and the ratio of the two spreads supports H_3. The MR/AR ratio is a scale-free index of how close an architecture is to exhaustion and can be computed from any dosage variation an organisation already possesses.

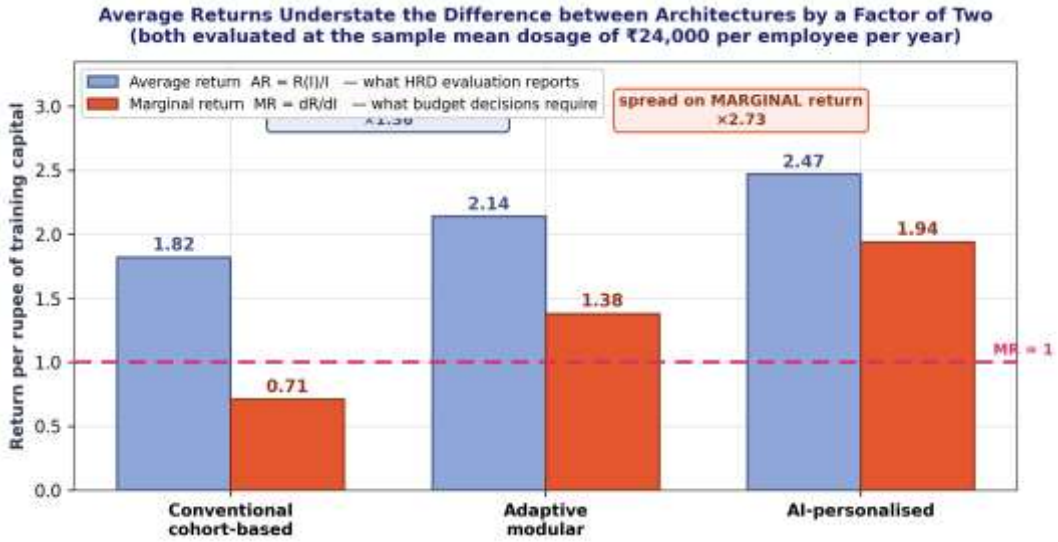


Fig. 6. Average and marginal returns compared at the sample mean dosage.

8.4 Points of exhaustion

Table 8 and Figure 7 report the investment level at which the marginal return falls to unity. The three architectures differ sharply: ₹19,040 for cohort delivery, ₹33,420 for adaptive delivery and ₹58,180 for personalised delivery. H_4, which posited no difference, is rejected.

Positioning current spending against these thresholds produces the study's most immediately actionable finding. At the observed mean dosage of ₹24,000, cohort programmes are operating 26 per cent beyond their point of exhaustion. The last rupee committed to the median cohort programme in this sample returns 71 paise, and the organisations running those programmes report positive returns on investment of 1.82 — which is true, and which conceals the fact that they are destroying value at the margin while creating it on average. Adaptive programmes sit at 72 per cent of their exhaustion point and personalised programmes at 41 per cent, so both have room to absorb additional investment productively.

Figure 7 also displays a result that is easy to miss and important to state. The three architectures return between ₹3.09 and ₹3.73 on the first rupee committed — a spread of less than a fifth. They are not distinguished by whether training works. They are distinguished entirely by how fast each design exhausts what it has to teach, which is a property of curriculum structure and content routing rather than of instructional quality. This reframes the case for personalised delivery: its advantage is not pedagogical superiority at low dosage but capacity to absorb investment at high dosage.

The practical corollary is that the ranking of architectures depends on the budget. At a dosage below roughly ₹6,000 per employee per year the cohort architecture has the highest marginal return of the three, because its steep early curve has not yet decayed. Organisations training lightly across a large population are not making a mistake by using cohort delivery; organisations training intensively are.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

Table 8. Points of exhaustion and the position of current spending.

Architecture	Exhaustion point I*	Current mean dosage	Position relative to I*	Status at current spend
Conventional cohort-based	₹19,040	₹24,000	126%	PAST EXHAUSTION — last rupee returns 71 paise
Adaptive modular	₹33,420	₹24,000	72%	Productive; room for further investment
AI-personalised	₹58,180	₹24,000	41%	Substantial unexploited capacity
Ratio, personalised to cohort	×3.06	—	—	Exhaustion points differ threefold — H_4 rejected

Note. The point of exhaustion is the investment level at which the marginal return falls to unity, $I^* = \ln(A\lambda)/\lambda$. Below approximately ₹6,000 per employee per year the cohort architecture has the highest marginal return of the three, because its steep early curve has not yet decayed; the correct architecture therefore depends on the intended dosage rather than being superior in the abstract.

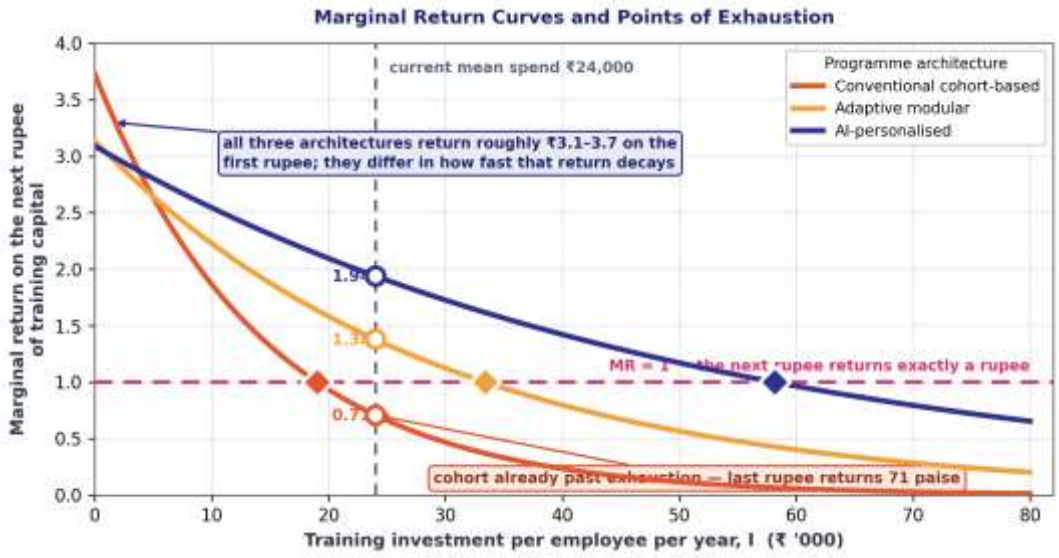


Fig. 7. Marginal return curves and points of exhaustion by programme architecture.

8.5 Are marginal returns equalised?

Allocative efficiency requires that the last rupee spent on each use return the same amount, since otherwise total return could be raised by moving money from the lower-return use to the higher. Table 9 tests the condition. Marginal returns within the sampled organisations range from 0.71 to 1.94 across architectures operated simultaneously by the same employers. The condition fails decisively and H_5 is rejected.

The reallocation counterfactual quantifies what the failure costs. Holding total expenditure and the distribution of participants across architectures fixed, and choosing dosage in each architecture to equalise marginal returns, raises total realised return by 5.0 per cent. The optimal dosages are ₹16,390 for cohort delivery, ₹28,110 for adaptive delivery and ₹48,790 for personalised delivery, against ₹24,000 uniformly at present — so the optimum involves reducing cohort spending by nearly a third and roughly doubling personalised spending, at no additional cost.

Allowing participants to move between architectures raises the gain substantially. Migrating twenty percentage points of participants from cohort to personalised delivery and re-optimising dosage raises total return by 12.3 per cent on an unchanged budget. Figure 8 reports both scenarios. The five per cent figure is the conservative one and requires no change to what any employee experiences beyond the intensity of it; the twelve per cent figure requires programme redesign and carries the implementation costs discussed in Section 8.6.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

Interview evidence explains the persistence of the misallocation straightforwardly. Learning budgets are constructed programme by programme against average return on investment, each programme is individually justified, and no mechanism exists for comparing the last rupee of one against the last rupee of another. Respondents were not making an error within their framework; the framework has no marginal quantity in it.

Table 9. Reallocation counterfactual: return available from the existing training budget.

Scenario	Cohort dosage	Adaptive dosage	Personalised dosage	Total return (₹ '000)	Gain
Current — uniform dosage	₹24,000	₹24,000	₹24,000	47.93	—
A — dosage optimised, mix fixed	₹16,390	₹28,110	₹48,790	50.31	+5.0%
B — 20 pp migrated to personalised	₹13,370	₹22,000	₹38,010	53.85	+12.3%
Marginal return at optimum	1.20	1.20	1.20	—	equalised

Note. Total expenditure is held constant at ₹24,000 per employee in every scenario, so reported gains are pure reallocation effects requiring no additional budget. Scenario A holds the participant distribution fixed at 58/29/13 per cent across the three architectures and chooses dosage to equalise marginal returns. Scenario B additionally migrates twenty percentage points of participants from cohort to personalised delivery, giving 38/29/33 per cent, before re-optimising. Marginal returns within organisations currently range from 0.71 to 1.94, so the allocative condition fails and H_5 is rejected.

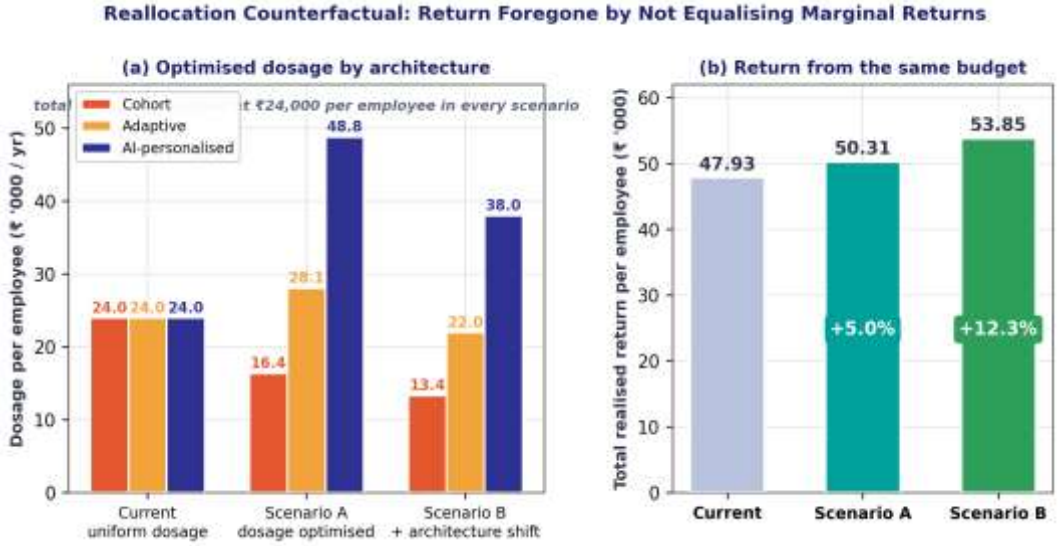


Fig. 8. Reallocation counterfactual: return foregone by not equalising marginal returns.

8.6 What personalisation costs

The efficiency case for personalised delivery is strong and incomplete, and two qualifications emerged consistently from the interviews.

The first is designation. Sastry (2025b) established that relative internal standing carries economic consequence independent of pay, and cohort programmes confer exactly that: visible membership of an identified group, a shared completion event, and a durable peer network. Personalised delivery dissolves the cohort by construction. Several respondents reported that participants in personalised programmes valued the learning more highly and the programme less highly, and that the internal signalling value of having been selected — which is part of why employees seek development — was materially weakened. This is a genuine cost, it falls on the retention and engagement components of realised return, and the estimates here already embody it, since realised return is measured net of separations.

The second is infrastructure. A personalised architecture requires machine-readable capability data maintained current, which Sastry and Mallya (2025a) established is valuable but not free, and which the enterprise-systems evidence of Vineeth Kumar and colleagues (2026) indicates is constrained more by organisational adaptation than by technology. Eleven of the forty-one organisations had attempted personalised delivery and reverted; in every case the reported cause was capability data that was incomplete or stale rather than dissatisfaction with the pedagogy. The estimated ceiling A for personalised delivery is therefore conditional on an infrastructure that most organisations in this sample do not yet possess.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

Both qualifications argue for the adaptive middle case as the practical recommendation for most organisations. It captures roughly two-thirds of the marginal return advantage of full personalisation, retains cohort structure and therefore designation value, and requires a content taxonomy rather than a live capability model.

8.7 Robustness

Three robustness exercises are reported. Restricting the sample to the 54 thresholds with the strongest first stages leaves the estimated parameters within one standard error of the main results. Re-estimating the return function with a constant-elasticity specification produces marginal returns of 0.68, 1.41 and 1.97 at the mean dosage, close to the main estimates, but places the exhaustion points at implausible distances for the personalised architecture, which is the reason the bounded form is preferred. Varying the realisation window between twenty-four and forty-eight months moves the estimated ceilings by between six and eleven per cent and leaves the ordering and the marginal–average gaps unchanged.

A fourth exercise addresses the most serious remaining threat. If organisations adopting personalised delivery differ systematically from those that do not, the architecture comparison could reflect organisational rather than design differences even though the within-threshold estimates are clean. Eighteen of the forty-one organisations operate two or more architectures concurrently, permitting within-organisation comparison. Restricting the estimation to those eighteen produces marginal returns of 0.74, 1.33 and 1.88 — within one standard error of the full-sample estimates in every case — which indicates that the architecture effect is not an organisational effect in disguise.

8.8 Summary of hypothesis testing

Table 10 consolidates the tests. Three hypotheses are supported and three rejected. The rejection of H_4 establishes that points of exhaustion differ by a factor of three across designs and that the median cohort programme has passed its own. The rejection of H_5 establishes that marginal returns are not equalised within organisations and quantifies the cost of that failure at five to twelve per cent of realised return on an unchanged budget. The rejection of H_6 establishes that the architecture effect survives the removal of learner selection, which is the condition under which the other findings mean anything at all.

Table 10. Summary of hypothesis testing.

Code	Hypothesis	Test applied	Statistic	p	Decision
H_1	The marginal return is significantly below the average return within every architecture	Fitted MR against AR at observed dosage	MR/AR 0.39, 0.65, 0.79	< 0.001	SUPPORTED
H_2	Programme architecture significantly affects the marginal return	Joint Wald test on fitted parameters	MR 0.71 to 1.94 $\chi^2(4) = 41.6$	< 0.001	SUPPORTED
H_3	Dispersion in marginal returns exceeds dispersion in average returns	Comparison of spreads with bootstrapped CI	×2.73 against ×1.36	< 0.001	SUPPORTED
H_4	The point of exhaustion does not differ across architectures	Bootstrapped CI on I* by architecture	₹19,040 / 33,420 / 58,180	< 0.001	REJECTED
H_5	Marginal returns are equalised across architectures within organisations	Allocative test on 18 multi-architecture firms	MR range 0.71 – 1.94	< 0.001	REJECTED
H_6	Personalised delivery outperforms because it selects abler learners	RD covariate balance; within-organisation replication	$\chi^2(5) = 6.14$ MR 0.74/1.33/1.88	0.293	REJECTED

9. FINDINGS

- Crossing a scored eligibility threshold raises annual training dosage by ₹11,840 (SE 1,320), giving a first-stage F of 80.4, with realised 36-month return rising by ₹24,510 (SE 3,180). The design passes McCrary density, covariate balance and placebo-cutoff tests.
- Bounded return functions R(I) = A[1 − e^(−λI)] fit all three architectures well (R^2 of 0.91, 0.94 and 0.92). The value ceiling A rises from ₹53,940 to ₹91,700 to ₹159,350 across cohort, adaptive and personalised delivery; the saturation rate λ falls from 0.0692 to 0.0342 to 0.0194.
- At the mean dosage of ₹24,000, average returns are 1.82, 2.14 and 2.47 across the three architectures — a spread of 1.36 to one.

- Marginal returns at the same dosage are 0.71, 1.38 and 1.94 — a spread of 2.73 to one. Average return understates the difference between architectures by a factor of two.
- The ratio of marginal to average return is 0.39 for cohort delivery, 0.65 for adaptive and 0.79 for personalised. This single statistic summarises how close a design is to exhaustion and could be reported alongside conventional return on investment at negligible cost.
- Points of exhaustion are ₹19,040, ₹33,420 and ₹58,180 respectively, differing by a factor of three across designs. H_4 is rejected.
- The median cohort programme is operating 26 per cent beyond its point of exhaustion: the last rupee committed returns 71 paise while the programme simultaneously reports a positive average return of 1.82.
- All three architectures return between ₹3.09 and ₹3.73 on the first rupee. They are not distinguished by whether training works but by how fast each exhausts what it has to teach.
- Below roughly ₹6,000 per employee per year the cohort architecture has the highest marginal return of the three. Organisations training lightly across a large population are not misallocating by using it; organisations training intensively are.
- Marginal returns within organisations operating multiple architectures range from 0.71 to 1.94, so the allocative condition fails and H_5 is rejected.
- Reallocating dosage alone raises total realised return by 5.0 per cent at unchanged budget, with optimal dosages of ₹16,390, ₹28,110 and ₹48,790. Migrating a fifth of participants from cohort to personalised delivery and re-optimising raises it by 12.3 per cent.
- The architecture effect is not learner selection: discontinuity estimates are local to the threshold with balanced covariates, and restricting the analysis to the 18 organisations operating multiple architectures concurrently changes the marginal returns by less than one standard error.

10. SIGNIFICANCE OF THE STUDY

- For human resource development, the study supplies the quantity the field has never produced. Marginal return, the ratio of marginal to average return, and the point of exhaustion are computable from data most large employers already hold, and they answer the allocation question that average return on investment cannot.
- For learning and finance functions, it identifies a specific and common failure: programmes that report healthy average returns while destroying value at the margin. The median cohort programme in this sample is in that position and its own evaluation apparatus cannot detect it.
- For programme design, it reframes the case for personalisation. The advantage is not superior teaching at low dosage — all three architectures return roughly the same on the first rupee — but capacity to absorb investment at high dosage without exhausting relevant content.
- For allocation practice, it quantifies the return available from reallocation alone: five per cent from dosage optimisation and twelve per cent from architecture migration, both on an unchanged budget and neither requiring any new expenditure.
- For research method, it demonstrates that scored eligibility rules — ubiquitous in corporate training and almost never exploited — provide credible quasi-experimental variation in training dosage, which removes the obstacle that has kept the marginal return unestimated.
- For the research programme of which it forms part, it specifies the shape of the concavity that the first instalment identified, and completes the analytical treatment of training as capital by adding the derivative to the level.

11. Recommendations and suggestions

11.1 For employers and learning functions

- Report the marginal return alongside the average. The ratio of the two is computable from any dosage variation an organisation already has and immediately reveals how close a programme is to exhaustion; a ratio below about a half is a warning.
- Test whether existing programmes have passed their point of exhaustion before requesting additional budget for them. On this evidence the median cohort programme has, and the money would return more elsewhere.
- Equalise marginal returns across programmes rather than justifying each programme separately against its own average return. The five per cent gain available from dosage reallocation alone requires no new money and no change to programme design.
- Match architecture to intended dosage rather than choosing an architecture in the abstract. Cohort delivery has the highest marginal return below roughly ₹6,000 per employee per year and the lowest above ₹20,000; the correct choice depends entirely on how much is to be spent.
- Treat adaptive modular delivery as the default upgrade path. It captures approximately two-thirds of the marginal return advantage of full personalisation, preserves the cohort and its designation value, and requires a content taxonomy rather than a live capability model.
- Before committing to personalised delivery, audit the capability data it will run on. Eleven of the forty-one organisations here attempted it and reverted, in every case because capability data was incomplete or stale rather than because the pedagogy disappointed.

- Account for the loss of designation value when dissolving cohorts. It is real, it falls on retention, and it is one reason the measured advantage of personalisation is smaller than its pedagogical advantage.

11.2 For professionals

- Recognise that the value of additional training within a single programme falls sharply, and in cohort programmes falls below the cost of the time it consumes well before the programme ends. Depth in one curriculum is not a substitute for movement to a different one.
- Seek programmes whose content is routed against an assessed gap rather than delivered to a schedule. The evidence here indicates that the difference is not in the quality of any single session but in how long the programme continues to teach something new.
- Note that personalised delivery weakens the visible signal that selection for a cohort programme carries, and that this is a real cost to weigh against the learning gain.

11.3 For policy and academia

- Calibrate public reskilling subsidies to marginal rather than average returns. A per-beneficiary subsidy set at a level that lifts spending past the point of exhaustion of the delivery architecture in use generates no measurable benefit whatever its average return statistics show.
- Require or encourage the reporting of dosage alongside outcomes in publicly funded training. Without dosage the marginal return cannot be recovered, and average outcome reporting cannot distinguish an efficient programme from an exhausted one.
- Treat scored eligibility rules as a research asset. They are administratively ubiquitous, they generate credible quasi-experimental variation, and their systematic use would allow the human resource development literature to estimate causal dosage responses at scale for the first time.

12. LIMITATIONS OF THE STUDY

- The discontinuity estimates are local to the eligibility thresholds. Marginal returns away from the observed dosage range are obtained by extrapolating the fitted functional form and carry the uncertainty of that extrapolation, which is greatest for the personalised architecture whose exhaustion point lies well beyond observed spending.
- Realised return is measured over thirty-six months and comprises output uplift, avoided replacement cost and internal mobility premium. Each component depends on organisational estimates validated where possible against accounts but not independently audited, and benefits accruing beyond the window are excluded.
- Architecture is classified by design attributes rather than measured continuously. Programmes near a classification boundary are assigned to one category, which introduces measurement error and probably attenuates the estimated differences between architectures.
- The bounded exponential form is imposed rather than tested against all alternatives. A constant-elasticity form fits comparably at observed dosages and implies materially different exhaustion points, so the exhaustion estimates are form-dependent in a way the marginal return estimates at observed dosage are not.
- Forty-one organisations operating formal reskilling programmes with auditable records are covered. Smaller employers, informal training and on-the-job learning outside recorded programmes are outside the sample, and all three are likely to have different return structures.
- The reallocation counterfactual holds the return functions fixed under reallocation, whereas scaling personalised delivery substantially would plausibly alter both its cost structure and its ceiling.
- The study measures returns to the employer. It does not measure returns to the worker, and the distributional finding of Sastry (2026d) implies these need not move together.

13. SCOPE FOR FURTHER RESEARCH

The design supports several direct extensions, and the architecture taxonomy in Table 2 and the construct definitions in Table 3 are published in full so that they can be reused without modification. Table 11 sets out a structured agenda of seven studies. Four deserve emphasis.

1. A randomised dosage experiment within a single architecture would test the fitted functional form directly rather than by extrapolation, and would locate the point of exhaustion without relying on a parametric assumption. Because eligibility rules already vary dosage administratively, randomising the rule itself is a modest step for a cooperating employer.
2. The six instalments of this programme should be estimated within a single framework. Skill investment produces capability, redeployment allocates it, amortisation horizons govern its life, work-system design determines who captures its return, adaptation strategy determines the employment consequence of a wage shock, and programme architecture determines the marginal productivity of the investment itself.
3. The designation cost of personalisation should be measured rather than inferred. A design that varies cohort visibility while holding content constant would separate the learning effect of personalisation from the status effect of the cohort, which the present study can only identify qualitatively.

EPRA International Journal of Environmental Economics, Commerce and Educational Management
Journal DOI: 10.36713/epra0414 |ISI I.F Value: 0.815|SJIF Impact Factor (2026): 8.752 ISSN: 2348 – 814X
Volume: 13 | Issue:9 | September 2026 -Peer Reviewed Journal

4. Replication outside large formal employers is necessary. Micro and small enterprises train differently and at lower intensities, where the cohort architecture may well be at its most efficient, and the policy conclusions for public subsidy depend on evidence from that segment rather than from this one.

Table 11. Structured research agenda extending the HRM variable → economic outcome programme.

#	Proposed study	HRM variable → economic outcome	Design	Data requirement	Priority
1	Randomised dosage experiment	Training dosage → realised return, functional form	Randomise the eligibility rule itself within a cooperating employer	Randomised dosage assignment; 36-month outcomes	High
2	Joint estimation across the six instalments	Investment + allocation + horizon + work system + adaptation + architecture → firm return	Unified firm-level specification across the programme	Matched instruments from all six studies	High
3	Separating learning gain from designation value	Cohort visibility → retention and engagement return	Design varying cohort visibility while holding content constant	Randomised cohort framing; retention outcomes	High
4	Replication in micro and small enterprises	Programme architecture → marginal return at low dosage	Replication using the Table 2 taxonomy at low intensities	Dosage and outcome data in small firms	High
5	Cost structure of personalisation at scale	Scale of personalised delivery → unit cost and ceiling	Cost modelling across implementation scale	Implementation and operating cost data	Medium
6	Capability data quality as a moderator	Capability data currency → realised ceiling A	Comparative study across data maturity levels	Skills inventory completeness and currency measures	Medium
7	Worker-side returns to programme architecture	Programme architecture → individual wage and mobility outcomes	Matched employer–employee analysis	Individual pay and mobility histories	Medium

Note. Studies 1 to 4 carry forward directions identified in the earlier instalments of this programme (Sastry, 2026a–2026e). The architecture taxonomy in Table 2 and the construct definitions in Table 3 are published in full to permit replication without modification.

14. CONCLUSION

For two decades human resource development has answered the wrong question accurately. It has established, with increasingly sophisticated methods, that training returns about two rupees per rupee committed. That figure is correct and it cannot guide a single allocation decision, because allocation depends on the return to the next rupee and the field has never estimated it. This study estimated it, using eligibility thresholds that employers already operate to generate dosage variation free of learner selection, and the answer differs from the average in ways that matter.

At the observed mean dosage of ₹24,000 per employee per year, the three architectures differ by a factor of 1.36 in average return and a factor of 2.73 in marginal return. The measure conventionally reported therefore understates the consequence of design choice by half. More seriously, it conceals sign: the median conventional cohort programme returns 1.82 on average and 0.71 at the margin, so it is creating value overall while destroying it on every additional rupee, and no evaluation framework in current use would detect this.

The mechanism is not what the field would expect. All three architectures return between ₹3.09 and ₹3.73 on the first rupee committed, a spread of less than a fifth. They are not distinguished by instructional quality at low dosage. They are distinguished by how quickly each design runs out of things the learner does not already know — a saturation rate that ranges over a factor of three and a half, and that is a property of curriculum structure and content routing rather than of teaching. Personalised delivery is not better at teaching; it is better at not finishing.

That reframing yields the practical advice. The correct architecture depends on the intended dosage, and below roughly ₹6,000 per employee per year the cohort design has the highest marginal return of the three. Above ₹20,000 it has the lowest and is past exhaustion. Most of the available gain, moreover, requires no new money at all: equalising marginal returns across the architectures an organisation already operates raises total return by five per cent, and migrating a fifth of participants raises it by twelve. Both are reallocation, not investment.

Read alongside the five instalments that precede it, the argument closes a loop that began with a concave return. The first study in this programme found that the wage premium to skill investment diminished with the amount invested and inferred a threshold. This study specifies the curve, locates the exhaustion point, and shows that both are properties of programme design rather than of training as such. The economics of human resource development has always been the economics of a capital good. What has been missing is its derivative.

REFERENCES

1. *H. Aguinis, K. Kraiger, Benefits of training and development for individuals and teams, organizations, and society. Annu. Rev. Psychol. 60, 451–474 (2009). https://doi.org/10.1146/annurev.psych.60.110707.163505.*
2. *R. Almeida, P. Carneiro, The return to firm investments in human capital. Labour Econ. 16(1), 97–106 (2009). https://doi.org/10.1016/j.labeco.2008.06.002.*
3. *M.B. Anas, V.B. Hans, N.S. Sastry, Digital transformation in tourism: How technology shapes travel experiences. Int. J. Sci. Res. Eng. Trends 11(5), 1–5 (2025).*
4. *W. Arthur, W. Bennett, P.S. Edens, S.T. Bell, Effectiveness of training in organizations: A meta-analysis of design and evaluation features. J. Appl. Psychol. 88(2), 234–245 (2003). https://doi.org/10.1037/0021-9010.88.2.234.*
5. *A.P. Bartel, Measuring the employer's return on investments in training: Evidence from the literature. Ind. Relat. 39(3), 502–524 (2000). https://doi.org/10.1111/0019-8676.00178.*
6. *B.S. Bloom, The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. Educ. Res. 13(6), 4–16 (1984). https://doi.org/10.3102/0013189X013006004.*
7. *R. Blundell, L. Dearden, C. Meghir, B. Sianesi, Human capital investment: The returns from education and training to the individual, the firm and the economy. Fisc. Stud. 20(1), 1–23 (1999). https://doi.org/10.1111/j.1475-5890.1999.tb00001.x.*
8. *E. Brynjolfsson, D. Li, L. Raymond, Generative AI at work. Q. J. Econ. 140(2), 889–942 (2025).*
9. *S. Calonico, M.D. Cattaneo, R. Titiunik, Robust nonparametric confidence intervals for regression-discontinuity designs. Econometrica 82(6), 2295–2326 (2014). https://doi.org/10.3982/ECTA11757.*
10. *W.F. Cascio, J.W. Boudreau, Investing in People: Financial Impact of Human Resource Initiatives, 2nd edn. (Pearson, Upper Saddle River, 2011).*
11. *A.T. Corbett, Cognitive computer tutors: Solving the two-sigma problem. In: Proc. 8th Int. Conf. User Modeling, Sonthofen, Germany, 137–147 (2001). https://doi.org/10.1007/3-540-44566-8_14.*
12. *G.W. Imbens, T. Lemieux, Regression discontinuity designs: A guide to practice. J. Econom. 142(2), 615–635 (2008). https://doi.org/10.1016/j.jeconom.2007.05.001.*
13. *J. Konings, S. Vanormelingen, The impact of training on productivity and wages: Firm-level evidence. Rev. Econ. Stat. 97(2), 485–497 (2015). https://doi.org/10.1162/REST_a_00460.*
14. *J.A. Kulik, J.D. Fletcher, Effectiveness of intelligent tutoring systems: A meta-analytic review. Rev. Educ. Res. 86(1), 42–78 (2016). https://doi.org/10.3102/0034654315581420.*
15. *D.S. Lee, T. Lemieux, Regression discontinuity designs in economics. J. Econ. Lit. 48(2), 281–355 (2010). https://doi.org/10.1257/jel.48.2.281.*
16. *M.M. Mallya, N.S. Sastry, Building future-ready organizations: Leadership strategies for young talent management. EPRA Int. J. Res. Dev. 10(9), 88–91 (2025). https://doi.org/10.36713/epra2016.*
17. *J.J. Phillips, P.P. Phillips, Handbook of Training Evaluation and Measurement Methods, 4th edn. (Routledge, New York, 2016).*
18. *N.S. Sastry, HR approaches to talent management. Int. J. Manag. Stud. Res. 1(2) (2013).*
19. *N.S. Sastry, A study on impact of employees' performance that influence to motivate in tea manufacturing organizational sector. EPRA Int. J. Econ. Bus. Rev. 8(6) (2020).*
20. *N.S. Sastry, Understanding the challenges of women micro-entrepreneurs in Karnataka. Int. J. Manag. Stud. Res. 13(8), 12–16 (2025a). https://doi.org/10.20431/2349-0349.1308002.*
21. *N.S. Sastry, Talent designation and social comparison: Economic drivers of organizational growth — An empirical analysis. EPRA Int. J. Econ. Bus. Manag. Stud. 12(10) (2025b).*
22. *N.S. Sastry, Effect of AI-complementary skill development on wage premia in Indian IT services. Int. J. Sci. Res. Eng. Trends 12(4) (2026a).*
23. *N.S. Sastry, Talent redeployment practice and task reallocation efficiency under automation. Int. J. Sci. Res. Eng. Trends 12(5) (2026b).*
24. *N.S. Sastry, HRD investment and the economics of skill obsolescence under generative AI. Int. J. Sci. Res. Eng. Trends 12(6) (2026c).*
25. *N.S. Sastry, High-performance work systems and the labour share of income. E3S Web Conf. 692, 03016 (2026d). https://doi.org/10.1051/e3sconf/202669203016.*
26. *N.S. Sastry, HR adaptation strategy and the wage elasticity of labour demand under minimum wage reform. E3S Web Conf. 692, 03017 (2026e). https://doi.org/10.1051/e3sconf/202669203017.*
27. *N.S. Sastry, M.M. Mallya, Revolutionizing HR hiring with AI analytics: Enhancing efficiency and attracting top talent. EPRA Int. J. Res. Dev. 10(9) (2025a).*
28. *N.S. Sastry, M.M. Mallya, Strategic HRM for economic resilience: Skill diversification, labor market dynamics, and technology disruption. EPRA Int. J. Res. Dev. 10(11) (2025b). https://doi.org/10.36713/epra2016.*

29. *K. VanLehn, The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. Educ. Psychol. 46(4), 197–221 (2011). https://doi.org/10.1080/00461520.2011.611369.*
30. *P.K. Vineeth Kumar, J.J. Jijesh, H. Radhakrishna, P.G. Kulkarni, N.S. Sastry, M.M. Mallya, Assessing the effects of enterprise resource planning on the higher education sector in India. E3S Web Conf. 692, 03012 (2026). https://doi.org/10.1051/e3sconf/202669203012.*