



HRD INVESTMENT AND THE ECONOMICS OF SKILL OBSOLESCENCE UNDER GENERATIVE AI

HRM Variable: Training Amortisation Horizon → Economic Outcome: Return to Training Capital

Dr. N Subbukrishna Sastry¹, Dr. Manjula Mallya M²

¹Professor, School of Management, CMR University, Bangalore, Karnataka, India. ORCID ID: 0009-0009-0460-7057

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

ABSTRACT

Firms treat training expenditure as an investment recovered over a horizon, but they choose that horizon administratively while the asset it buys depreciates at a rate set by the technology frontier. This study measures the distance between the two. The training amortisation horizon (A) is the period over which a firm expects to recover a training outlay, observed through four verifiable practices: the certification currency window in the staffing system, the curriculum refresh cycle, the management-accounting payback horizon and the skills-inventory decay rule. The skill half-life (H) is the observed time until half the economic value of the trained capability has decayed, estimated from exponential fits to skill-value trajectories. Their ratio, $\Gamma = A/H$, is the study's central quantity, and the resulting stranded share of expenditure is directly interpretable as misallocated HRD budget. Evidence comes from a retrospective longitudinal cohort of 1,240 training completions in 38 organisations, tracked across four waves over 36 months, with 55 learning and finance leader interviews. Half-lives range from 1.1 years for generative-AI prompt and application patterns to 11.5 years for leadership capability, while the horizon firms apply ranges only from 4.2 to 5.3 years. Half-life varies 10.4 times more across skill clusters than the horizon applied to it, and 84.0 per cent of HRD spend sits on the over-amortised side of the crossover at $H \approx 4.6$ years. The spend-weighted horizon is 4.65 years against a spend-weighted half-life of 3.12 years, a gap of 1.53 years and a ratio of 1.49. Discounting the decaying benefit flow at 12 per cent against the recovery multiple firms themselves assume ($2.84\times$) yields a realised multiple of $1.60\times$ and a stranded share of 43.7 per cent — ₹43.70 of every ₹100 committed. Realised return falls by 0.284 for each unit increase in Γ . Half-life is not invariant to use: skills deployed on more than sixty per cent of billable hours decay 2.9 times more slowly than those below fifteen per cent, so obsolescence is partly an allocation outcome rather than a purely technological one. The eight organisations that differentiate horizon by cluster strand ₹27.90 against ₹48.60 for those applying a uniform horizon. The study supplies a cluster-differentiated amortisation schedule and argues that skills with half-lives beyond eight years should be reclassified from amortising training capital to capability endowment.

KEYWORDS — Training Amortisation Horizon, Skill Half-Life, Skill Obsolescence, Return to Training Capital, Human Resource Development, Generative AI, Human Capital Depreciation, Stranded Budget, Exponential Decay, Weibull Accelerated Failure Time, Capital Budgeting for HRD, Indian IT Services, Strategic HRM.

I. INTRODUCTION

Every organisation that funds training makes an implicit capital-budgeting decision, and most make it without noticing. Expenditure is committed today against a benefit expected to accrue over some future period. The length of that period is rarely stated as such, but it is embedded in observable practice: in how long a certification remains valid in the staffing system, in how often a curriculum is rebuilt, in the payback assumed in a learning business case, and in the rule that flags a recorded skill as stale. Together these constitute a training amortisation horizon — the period over which the firm behaves as though its training outlay is being recovered.

The asset that outlay buys does not consult the horizon. Human capital depreciates at a rate set by the pace at which the underlying technology, tooling and task content change, and under generative AI that pace has accelerated sharply for exactly the skills into which most technology-sector training expenditure now flows. Deming and Noray (2020) established that returns to technology-intensive occupations decay rapidly as skill content shifts, and that the decay is fastest at the frontier — precisely where firms are spending most. A framework-specific engineering skill acquired in one release cycle may be substantially obsolete within two, while the firm's accounting and workforce systems continue to treat it as current for four or five years.

The distance between these two quantities is the subject of this study. Let A denote the amortisation horizon a firm applies and H the observed half-life of the skill it has funded. Their ratio, $\Gamma = A/H$, expresses how many half-lives of economic value the firm expects to recover over. Where Γ exceeds one substantially, the organisation is recognising benefit from an asset that has already largely decayed; the shortfall between planned and realised recovery is not an accounting abstraction but a quantity of money, and it is directly interpretable as misallocated human resource development budget. Where Γ falls below one, the opposite error is present: durable capability is under-credited, and the firm under-invests in skills whose value it fails to recognise.

The problem is sharpened by three features of the present moment in Indian IT services. First, training expenditure is large and rising: the industry is funding reskilling at scale against a workforce approaching 5.8 million while net revenue per employee is under pressure. Second, the composition of that expenditure has shifted decisively toward the fastest-decaying skill clusters, since



those are where the visible market premium sits. Third, and least examined, the accounting treatment has not moved at all. Indian Accounting Standard 38 requires training expenditure to be expensed rather than capitalised, on the reasoning that the firm does not control the asset; but management accounting nonetheless reasons in payback terms, and it is the management horizon rather than the statutory treatment that governs how much is spent and on what. That horizon has remained essentially constant while the half-life of what it funds has collapsed.

There is a substantial literature on human capital depreciation and a substantial literature on training evaluation, and they do not meet. The depreciation literature, from Rosen (1975) and Mincer and Ofek (1982) through to contemporary work on skill obsolescence, estimates decay rates but treats the firm's investment horizon as exogenous. The training-evaluation literature, following Kirkpatrick's levels and the utility-analysis tradition of Cascio and Boudreau (2011), estimates returns but assumes the benefit stream is stationary over the appraisal period. Neither asks the question that follows immediately from putting them together: what happens to the measured return when the appraisal horizon and the depreciation rate are estimated from the same data and compared?

This study asks it. It is the third instalment of a research programme pairing one human resource variable with one measurable economic outcome. The first examined skill development investment against the wage premium accruing to AI-complementary capability, finding the premium to be a joint product of skill and investment with a concave return (Sastry, 2026a). The second examined redeployment policy against allocative efficiency, finding that the return to reallocation machinery is conditional on automation depth and that policy components behave as complements (Sastry, 2026b). The present study completes the sequence by asking how long the capability produced in the first and allocated in the second actually lasts — and it finds that the answer materially revises the returns reported in both. Section VIII establishes, further, that decay is itself partly an allocation outcome: skills in heavy billable use decay at roughly a third the rate of skills that are trained and shelved, which connects this study's dependent variable directly to the deployment coefficient of the first and the redeployment machinery of the second.

II. REVIEW OF LITERATURE

2.1 Human Capital Depreciation and Skill Obsolescence

The theoretical starting point is that human capital, like physical capital, depreciates. Becker (1964) treated training as investment but said comparatively little about the rate at which its returns decay. Rosen (1975) and Mincer and Ofek (1982) supplied the missing element, showing that skills atrophy during periods of non-use and that a substantial part of the observed wage penalty to labour-force interruption reflects depreciation rather than signalling. Their central insight for present purposes is that depreciation is not purely a function of elapsed time; it depends on whether the skill is exercised.

The obsolescence literature distinguishes wear-and-tear depreciation, in which unused skills atrophy, from technological or economic obsolescence, in which skills remain intact but the demand for them disappears. De Grip and Van Loo (2002) set out this taxonomy and argued that technological obsolescence dominates in fast-moving sectors, where the worker has not forgotten anything but the market has moved. Neuman and Weiss (1995) estimated that depreciation rates in high-technology industries exceed those in low-technology industries by a substantial margin, and that the difference widens with the age of the vintage of training. Deming and Noray (2020) provided the most directly relevant contemporary estimate, showing that the wage premium to technology-intensive STEM occupations erodes rapidly as the skill content of those occupations changes, with new entrants continually displacing incumbents at the frontier.

Generative AI plausibly accelerates all of this. Brynjolfsson, Li and Raymond (2025) documented that generative assistance compresses productivity dispersion between novices and experts, which mechanically shortens the period over which any given increment of trained expertise remains distinguishing. Acemoglu and Restrepo (2019) supplied the task-level account of why: where a technology assumes task content, the human capital specific to that task loses value regardless of the worker's proficiency. The implication is that the half-life of a skill is a property of the task it addresses rather than of the person holding it, which is what makes cluster-level estimation the appropriate unit of analysis.

2.2 Training as Capital: Appraisal, Amortisation and Horizon

The treatment of training as capital has a long practical tradition and a weak measurement base. Cascio and Boudreau (2011) developed the utility-analysis framework that remains the standard for monetising human resource interventions, and Phillips and Phillips (2016) extended Kirkpatrick's evaluation levels to an explicit return-on-investment calculation. Both frameworks compute a return over an appraisal period, and both leave the choice of that period to the analyst. In practice it is chosen by convention, and the convention has proven remarkably stable across a period in which the underlying decay rates have not been.

The accounting treatment reinforces the disconnect. Under Indian Accounting Standard 38, and equivalently under IAS 38, expenditure on staff training is expensed as incurred because the entity cannot demonstrate control over the expected future economic benefits — employees may leave. Lev and Zarowin (1999) argued that the systematic expensing of intangible investment progressively degrades the informativeness of financial statements, and Lev and Radhakrishnan (2005) developed organisation-capital measures precisely to recover what conventional accounts omit. The consequence for the present study is specific: because training is expensed, no external discipline forces a firm to state its recovery horizon or to test it against evidence. The horizon survives unexamined inside management accounting, where this study locates and measures it.



Where a horizon is disciplined, it is disciplined contractually rather than economically. Indian IT services firms commonly attach service agreements or training bonds of twelve to twenty-four months to funded programmes. That period reflects a legal judgement about enforceability and a retention calculation, not an estimate of how long the capability will remain valuable, and Section 8.2 shows it to be markedly shorter than the horizon the same firms apply in their accounting and staffing systems — an internal inconsistency that appears to have gone unnoticed because the two numbers live in different functions.

2.3 Measuring Decay: Method Precedents

Estimating a half-life requires a decay model. The exponential specification $V(t) = V_0 e^{(-\lambda t)}$, with half-life $H = \ln 2 / \lambda$, is the standard form in depreciation analysis and has the property that the decay rate is constant in proportional terms — each year removes the same fraction of what remains. Arrazola and de Hevia (2004) applied exponential decay to human capital using panel wage data, and the approach has since become conventional where repeated observation of skill value is available.

Where the question is time to an obsolescence event rather than the shape of a continuous decline, survival methods are appropriate. The Weibull accelerated failure time model, unlike the semi-parametric proportional-hazards approach, yields directly interpretable time ratios — the multiplicative effect of a covariate on expected survival time — and permits the shape parameter to indicate whether the hazard of obsolescence is rising or falling with age. A shape parameter above one indicates accelerating obsolescence, which is the theoretically expected pattern where a technology frontier is moving. This study uses both: exponential fits for cluster half-lives, and a Weibull accelerated failure time model to identify what lengthens or shortens the survival of a trained skill.

For the financial translation, the relevant precedent is standard capital budgeting. Where a benefit flow $B(t) = B_0 e^{(-\lambda t)}$ is discounted at rate r over a horizon A , the present value is $B_0 [1 - e^{-(\lambda+r)A}] / (\lambda + r)$, whereas a firm assuming a stationary flow computes $B_0 [1 - e^{(-rA)}] / r$. The ratio of the two gives the realised share of planned recovery, and its complement is the stranded share. This is elementary, which is part of the point: the calculation is not difficult, and the finding of this study is that it is not being done.

2.4 Antecedent Work in the HRM Variable → Economic Outcome Programme

This study is the third instalment of a research programme pairing a single human resource variable with a single measurable economic outcome, and it depends closely on the two that precede it. Sastry (2026a) examined skill development investment against the AI-complementary wage premium and established that the premium is a joint product of the skill and the employer investment that produced it, with a concave return and a threshold below which investment yields little. That study reported a substantial deployment coefficient — capability applied to billable work realised roughly twice the premium of capability held but unapplied — without being able to explain it. The present study supplies the explanation: deployment does not merely price the skill higher, it slows the skill's decay, so deployed capability is both more valuable per period and valuable for longer.

Sastry (2026b) examined redeployment policy against allocative efficiency and found the return to reallocation machinery conditional on automation intensity, with policy components functioning as complements above a critical mass. That study measured how quickly displaced capability is rematched to productive work. Read together with the present findings, its reassignment intervals acquire an economic price: at the observed decay rates, an additional seventy-one days on the bench — the difference between the lowest and highest redeployment-policy quartiles reported there — destroys a measurable fraction of the skill value the firm has already paid to build, and destroys most of it in the fastest-decaying clusters. The three studies are therefore a single argument in three parts: capability is produced by investment, realised by allocation, and consumed by time.

Sastry and Mallya (2025b), on strategic human resource management and economic resilience, identified continuous learning and skill diversification as the practices most strongly associated with workforce sustainability under technological disruption. The present study qualifies that finding usefully: continuous learning is not merely desirable but structurally required, because a portfolio with a spend-weighted half-life of 3.12 years cannot be maintained by episodic investment. Sastry and Mallya (2025a), on artificial intelligence analytics in hiring, demonstrated the financial value of machine-readable capability data; the skills-inventory decay rule examined here is a direct application of that infrastructure, and its mis-specification is one of the four components of the horizon this study measures.

Mallya and Sastry (2025), studying leadership strategies for young talent, found growth opportunity and purpose to outrank compensation as retention drivers. That result bears on the interpretation of the retention rationale for training bonds: if development opportunity is itself a retention instrument, then a bond period set to protect against attrition may be solving a problem that the training expenditure has already reduced. Sastry (2025b), on talent designation and social comparison, established that relative internal standing carries economic consequence independent of pay, which helps explain why employees pursue visibly current certifications in the fastest-decaying clusters even where the durable return is lower — the signalling value decays more slowly than the productive value, and the individual optimises on the former while the firm pays for the latter.

Sastry (2025a), examining women micro-entrepreneurs in Karnataka, concluded that socio-cultural and systemic constraints bind more tightly on economic outcomes than skill or knowledge gaps; the corresponding caution here is that a training portfolio optimised purely on decay rates would still not deliver equitable returns. Sastry (2020) examined performance factors influencing motivation in the tea manufacturing sector, evidence that informs the treatment of refresher training as a behavioural as well as a technical intervention, and Sastry (2013) set out approaches to talent management within the human resource function. Anas, Hans



and Sastry (2025) demonstrated in tourism how a technology shift restructures the competence base of an entire service value chain, the cross-sector analogue of the obsolescence process modelled here. Table 1 synthesises the literature and states each source's specific contribution to the present design.

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

Author(s) & Year	Focus of Study	Key Findings	Relevance to Current Study
Becker (1964)	Human capital as investment	Training is investment yielding a measurable return; financing follows the portability of the skill	Establishes training as capital, but is silent on the rate at which its return decays
Rosen (1975); Mincer & Ofek (1982)	Human capital depreciation	Skills atrophy during non-use; a substantial part of the interruption wage penalty is depreciation, not signalling	Establishes that depreciation depends on use — the theoretical basis for the H_0 test
De Grip & Van Loo (2002)	Taxonomy of skill obsolescence	Distinguishes wear-and-tear atrophy from technological obsolescence, where the skill survives but demand does not	Frames the interpretation of the deployment gradient in Section 8.6
Neuman & Weiss (1995)	Depreciation by industry	Depreciation rates in high-technology industries substantially exceed those in low-technology industries	Supports cluster-level rather than portfolio-level estimation of half-life
Deming & Noray (2020)	Skill obsolescence in STEM careers	Returns to technology-intensive occupations decay rapidly as skill content shifts, fastest at the frontier	Closest contemporary precedent; motivates the study's central question
Acemoglu & Restrepo (2019)	Task displacement and reinstatement	Where technology assumes task content, human capital specific to that task loses value regardless of proficiency	Explains why half-life is a property of the task, not of the individual
Brynjolfsson, Li & Raymond (2025)	Generative AI at work	Generative assistance compresses productivity dispersion between novices and experts	Implies shortening of the period over which trained expertise remains distinguishing
Cascio & Boudreau (2011); Phillips & Phillips (2016)	Utility analysis and training ROI	Frameworks monetise HR interventions over an appraisal period chosen by the analyst	Identifies the unexamined horizon assumption this study measures
Lev & Zarowin (1999); Lev & Radhakrishnan (2005)	Intangibles and organisation capital	Systematic expensing of intangible investment degrades the informativeness of financial statements	Explains why no external discipline forces a firm to state or test its recovery horizon
Arrazola & de Hevia (2004)	Exponential decay of human capital	Applies the exponential specification to human capital using panel wage data	Method precedent for the decay fits in Table 2
Sastry (2026a)	Skill development investment → wage premium	The AI wage premium is a joint product of skill and employer investment; deployment doubles realised premium	First instalment; this study explains its unexplained deployment coefficient through the decay channel
Sastry (2026b)	Redeployment policy → allocative efficiency	Reallocation return is conditional on automation depth; policy components are complements above critical mass	Second instalment; this study prices its reassignment intervals as decay of training capital
Sastry & Mallya (2025b)	Strategic HRM and economic resilience	Continuous learning and skill diversification are the practices most associated with workforce sustainability	The 3.12-year portfolio half-life makes continuous learning structurally required, not merely desirable
Sastry & Mallya (2025a)	AI analytics in HR hiring	Machine-readable capability data delivers operational and sustained financial benefit	The skills-inventory decay rule is an application of this infrastructure and one of the four horizon components
Mallya & Sastry (2025)	Leadership for young talent	Growth opportunity and purpose outrank pay as retention drivers	Questions the retention rationale for training bonds set independently of the recovery horizon
Sastry (2025b)	Talent designation and social comparison	Relative internal standing carries economic consequence independent of absolute pay	Explains why individuals pursue fast-decaying visible certifications the firm pays for
Sastry (2025a)	Women micro-entrepreneurs, Karnataka	Socio-cultural and systemic constraints bind more tightly than skill or knowledge gaps	Cautions that a decay-optimised portfolio still would not deliver equitable returns
Anas, Hans & Sastry (2025); Sastry (2020, 2013)	Digital transformation in tourism; performance and motivation; talent management	Technology shifts restructure the competence base of a service value chain; performance factors drive motivation	Cross-sector analogue of obsolescence; informs treatment of refresher training as behavioural intervention



III. RESEARCH GAP

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

- The depreciation literature estimates decay rates and the training-evaluation literature estimates returns, but no study estimates both from the same data and compares them. The amortisation horizon is treated as exogenous in the first tradition and as an analyst's choice in the second.
- Skill half-lives are asserted far more often than they are measured. Practitioner claims about the shrinking half-life of technical skills circulate widely without an underlying longitudinal estimate, and none exists for the Indian IT services skill portfolio.
- The firm's amortisation horizon has never been operationalised as an observable, verifiable quantity. Because training is expensed under IAS 38 and its Indian equivalent, the horizon exists only implicitly in management practice and has therefore not been measured.
- No study quantifies the stranded fraction of training expenditure. Return-on-investment frameworks report a return; they do not report how much of the planned return fails to arrive because the asset decayed inside the appraisal period.
- Whether decay rates respond to deployment has not been tested in a contemporary technology setting. The wear-and-tear tradition implies that use preserves skill, but the technological-obsolescence tradition implies that use is irrelevant because the market, not the worker, has changed. The two make opposite predictions and the question is empirical.
- No cluster-differentiated amortisation guidance exists. Firms apply a single horizon across a portfolio whose components differ by an order of magnitude in durability, and no published schedule offers an alternative.

IV. STATEMENT OF THE PROBLEM

Indian IT services organisations are committing substantial and rising human resource development budgets under generative-AI pressure, and they are committing them predominantly to the skill clusters that decay fastest. The horizon over which they expect to recover that expenditure is set by administrative convention, is nearly uniform across clusters, and has not been revised as decay rates have accelerated.

The consequence is a measurable but unmeasured loss. Where the horizon exceeds the half-life, planned recovery is not realised, and the difference is money already spent against value that no longer exists. Because training is expensed rather than capitalised, no financial statement discloses this and no audit tests it.

The problem this study addresses is therefore to measure the two quantities on comparable terms — the amortisation horizon the firm applies and the half-life of the skill it funds — to express their relationship as a single interpretable ratio, to convert the mismatch into a stranded share of budget, and to establish whether that share is reducible by changes in horizon policy and deployment practice.

V. OBJECTIVES OF THE STUDY

1. To operationalise the training amortisation horizon as a verifiable composite of four observable organisational practices, and to measure it across a sample of Indian IT services organisations.
2. To estimate the economic half-life of trained skills by cluster, using exponential decay fits to longitudinal skill-value trajectories.
3. To compute the amortisation gap, expressed both as a difference in years and as the ratio $\Gamma = A / H$, and to test whether the horizon varies less across clusters than the half-life does.
4. To convert the gap into a stranded share of human resource development expenditure and to express it as rupees lost per hundred rupees committed.
5. To estimate the effect of the amortisation gap on the realised return to training capital, controlling for deployment, spend intensity, refresh cadence and firm characteristics.
6. To identify, through a Weibull accelerated failure time model, which characteristics of a trained skill lengthen or shorten its survival, and to test whether half-life is invariant to deployment intensity.
7. To derive a cluster-differentiated amortisation schedule specifying the horizon, the amortisation method and the refresh trigger appropriate to each skill cluster.

VI. HYPOTHESES OF THE STUDY

H₁: The training amortisation horizon applied by firms exceeds the observed half-life of the skills funded ($\Gamma > 1$) for the majority of human resource development expenditure.

H₂: The amortisation horizon varies significantly less across skill clusters than the observed half-life does — horizon rigidity.

H₃: The realised return to training capital declines significantly as the amortisation gap Γ widens.

H₄: Skill clusters exposed to generative AI exhibit significantly shorter half-lives than non-exposed clusters of comparable technical depth.

H₅: Organisations that differentiate the amortisation horizon by skill cluster achieve a significantly higher realised return to training capital than those applying a uniform horizon.

H₆: Skill half-life is invariant to deployment intensity — trained capability decays at the same rate whether or not it is applied to billable work.

VII. RESEARCH METHODOLOGY AND DESIGN

7.1 Research Type and Longitudinal Structure

The study employs a retrospective longitudinal cohort design with an embedded qualitative strand. A cohort of training completions is identified at a common baseline and the economic value of the skill acquired is measured at four waves — completion, twelve, twenty-four and thirty-six months. This structure permits direct estimation of a decay function for each skill



cluster rather than inference of decay from cross-sectional differences, which is the principal methodological advance over the existing obsolescence literature in this setting.

Four analytical strands follow. Exponential decay models estimate cluster half-lives. A Weibull accelerated failure time model identifies covariates that extend or shorten skill survival. A discounted capital-budgeting model converts half-lives and horizons into realised recovery multiples and stranded shares. A regression of realised return on the amortisation gap tests the study's central relationship with controls. Interviews with learning and finance leaders establish how horizons are set in practice and are used to interpret rather than to generate the estimates.

7.2 Population, Sampling and Measurement

The sampling frame comprises Indian IT services delivery organisations, engineering research and development centres and global capability centres that maintain auditable learning management records permitting retrospective cohort construction. Thirty-eight organisations met this criterion and consented. Within them, 1,240 training completions were sampled, stratified across nine skill clusters in proportion to observed expenditure share, and fifty-five learning, workforce-planning and finance leaders were interviewed.

The training amortisation horizon A is measured as the mean of four verifiable practices, each obtained from documentation rather than perception: the certification currency window applied in the staffing and skills system, the curriculum refresh cycle for the programme, the recovery or payback horizon used in the learning business case, and the interval after which a recorded skill is flagged stale in the skills inventory. The training bond or service agreement period is recorded separately as a contractual rather than an economic horizon, and is reported alongside for contrast. Table 3 gives the observed values.

Skill value $V(t)$ is a composite index anchored at 100 at completion, combining four standardised indicators: the share of the employee's billable task volume still applying the skill, internal staffing-system demand for the skill, the external market wage premium attaching to it, and manager-rated relevance of the proficiency to current work. Composite reliability is 0.86 and average variance extracted 0.61, satisfying convergent validity. Inter-wave measurement used identical instruments and identical rater panels where possible. Table 4 reports full construct operationalisation and the realised sample profile.

Table 3 · The Training Amortisation Horizon — Component Practices and Observed Values

#	Component Practice	Operational Definition (read from documentation, not perception)	Mean (years)	SD
1	Certification currency window	Period for which a completed certification is treated as valid by the staffing and skills allocation system	4.9	1.7
2	Curriculum refresh cycle	Interval between substantive rebuilds of a training programme's content	3.8	1.4
3	Business-case recovery horizon	Payback or recovery period assumed in the management-accounting case supporting the learning investment	5.2	1.9
4	Skills-inventory decay rule	Interval after which a recorded skill is automatically flagged stale in the capability inventory	5.1	2.1
—	COMPOSITE AMORTISATION HORIZON A	Mean of the four component practices, spend-weighted across clusters	4.65	1.6
—	Training bond / service agreement	Contractual period during which the employee is bound following a funded programme — recorded for contrast, not included in A	1.9	0.6

Note: Observed range of the composite horizon across the 38 organisations: 1.5 to 9.0 years. The contractual horizon is less than half the composite accounting horizon applied by the same organisations, and no organisation in the sample had reconciled the two figures.

Table 4 · Construct Operationalisation and Realised Sample Profile

Construct / Dimension	Measurement or Category	Value / n	Statistic
Amortisation horizon A — independent	Mean of four documented practices (Table 3), in years	Mean 4.65	SD 1.6; range 1.5–9.0
Skill half-life H — independent	$\ln 2 / \lambda$ from nonlinear least squares fit to $V(t)$	Mean 3.12	Spend-weighted
Amortisation gap Γ — central quantity	Ratio A / H ; gap in years reported as $A - H$	$\Gamma = 1.49$	Gap = 1.53 yrs
Return to training capital — dependent	Realised recovery multiple on training outlay, discounted at $r = 12\%$	Mean $1.60\times$	Planned $2.84\times$
Stranded share S — dependent	$1 - PV(\text{actual}) / PV(\text{planned})$	43.7%	₹43.70 per ₹100
Skill value $V(t)$	Composite of billable task application, internal demand, market premium, manager-rated relevance; anchored 100 at completion	4 waves	CR = 0.86; AVE = 0.61
Deployment intensity — moderator	Share of billable hours applying the trained skill	Mean 38.4%	SD 21.7
Horizon differentiation — moderator	Organisation applies cluster-specific rather than uniform horizon	8 of 38	21.1%



Discount rate r	Weighted average cost of capital reported by sampled organisations	12.0%	Sensitivity 8–16%
Sample — training completions	Completions with auditable baseline and four measurement waves	1,240	Level 1
Sample — organisations	Firms maintaining auditable learning management records	38	Level 2
Sample — interviews	Learning, workforce-planning and finance leaders	55	Qualitative strand
Employer type	Tier-1 Indian IT services	13 firms	34.2%
	Global capability centre	12 firms	31.6%
	Mid-tier Indian IT services	9 firms	23.7%
	Engineering R&D / product	4 firms	10.5%
Completion profile — experience	Below 5 yrs / 5–10 yrs / above 10 yrs	462 / 468 / 310	37.3 / 37.7 / 25.0%
Completion profile — funding	Employer-funded / co-funded / self-funded	874 / 219 / 147	70.5 / 17.7 / 11.9%

7.3 Model Specification

Decay is estimated cluster by cluster by nonlinear least squares on the exponential form, from which the half-life follows directly:

$$V(t) = V_0 \cdot e^{(-\lambda t)} \quad H = \ln 2 / \lambda \quad \Gamma = A / H \quad \dots(1)$$

The financial translation compares the present value of the actual decaying benefit flow against the present value the firm's own business case assumes, both taken over the firm's horizon A and discounted at r:

$$\begin{aligned} PV(actual) &= B_0 \cdot [1 - e^{-(\lambda+r)A}] / (\lambda + r) & PV(planned) &= B_0 \cdot [1 - e^{(-rA)}] / r \\ Stranded \text{ share } S &= 1 - PV(actual) / PV(planned) & Realised \text{ multiple} &= M(planned) \cdot (1 - S) \end{aligned} \quad \dots(2)$$

where r is set at 12 per cent, consistent with the weighted average cost of capital reported by the sampled organisations, and M(planned) is the recovery multiple assumed in their own learning business cases, which averages 2.84 times the outlay. The stranded share S is the study's operative measure of misallocated human resource development budget: it is the fraction of planned recovery that does not arrive because the asset decayed inside the horizon over which the firm was recognising it.

The central relationship is then estimated as a regression of the realised return to training capital on the amortisation gap and controls, with standard errors clustered at the organisation level to account for the nesting of completions within firms. Survival is modelled separately using a Weibull accelerated failure time specification in which the event is the skill-value index falling below fifty, with time ratios reported so that coefficients express the multiplicative effect of each covariate on expected skill survival time.

Figure 1 sets out the conceptual framework. The amortisation horizon is an administrative property of the firm; the half-life is an economic property of the skill; the gap between them determines what fraction of the training capital is recovered. Deployment intensity and horizon differentiation enter as the two moderators the firm can actually control.

Figure 1 · Conceptual Framework — Amortisation Horizon, Skill Half-Life and Return to Training Capital




Figure 1 · Conceptual framework — the amortisation horizon is a property of the firm, the half-life a property of the skill, and their ratio determines the recovered share of training capital.

7.4 Analytical Tools



Decay parameters were estimated by nonlinear least squares with bootstrapped confidence intervals on the half-life. Levene's test compared the dispersion of horizons against the dispersion of half-lives across clusters, and coefficients of variation were computed for the rigidity test. Weibull accelerated failure time models were fitted by maximum likelihood, with model adequacy assessed by Cox–Snell residuals. The return regression used ordinary least squares with organisation-clustered robust standard errors. Group comparisons between differentiating and non-differentiating organisations used Welch's t-test. Sensitivity of the stranded-share estimate to the discount rate was tested across the range 8 to 16 per cent and is reported in Section XII. Interview transcripts were analysed thematically. Estimation was carried out in R using the survival, minpack.lm and sandwich packages, with cross-validation in SPSS version 29.

VIII. RESULTS AND DISCUSSION

8.1 Estimated Skill Half-Lives

Table 2 and Figure 3 report the decay estimates. Fits are good across all nine clusters, with coefficients of determination between 0.89 and 0.96, confirming that the exponential specification describes the observed trajectories adequately.

The range is the first substantive finding. Generative AI prompt and application patterns decay at $\lambda = 0.630$ per year, giving a half-life of 1.1 years: half the economic value of such training has gone within thirteen months of completion. Framework-specific AI engineering follows at 1.4 years and cloud platform and tooling certification at 2.2 years. At the other end, foundational software and computer science engineering has a half-life of 9.2 years and leadership and client management 11.5 years. The ratio between the most and least durable cluster is 10.5 to one.

The pattern is not simply technical versus non-technical. AI governance, risk and assurance is a technical cluster with a half-life of 4.3 years, roughly four times that of generative-AI application patterns, because its content is regulatory and conceptual rather than tool-specific. Domain knowledge in banking, financial services and healthcare reaches 6.8 years. What separates the durable from the perishable is not subject matter but the degree of coupling to a specific tool, framework or vendor implementation — a distinction the Weibull results in Section 8.5 confirm directly.

Comparing generative-AI-exposed clusters against non-exposed clusters of comparable technical depth gives a mean half-life of 2.16 against 5.91 years, $t(1238) = 21.4$, $p < 0.001$, supporting H_4 . The exposure effect survives in the multivariate survival model reported later, where generative-AI exposure carries a time ratio of 0.594 — a forty-one per cent reduction in expected skill survival.

The spend-weighted mean half-life across the observed portfolio is 3.12 years. That figure is itself a finding about composition rather than about technology: it is low because 84.0 per cent of expenditure flows to clusters below the crossover, not because all skills decay quickly.

Table 2 · Estimated Skill Half-Lives by Cluster (exponential decay fits, 1,240 completions over 36 months)

Skill Cluster	Share of HRD spend (%)	Decay rate λ (per year)	Half-life H (years)	95% CI on H	R ² of fit	GenAI exposed
Generative AI prompt & application patterns	21.3	0.630	1.1	0.9 – 1.3	0.94	Yes
Framework-specific AI engineering	16.8	0.495	1.4	1.2 – 1.7	0.92	Yes
Cloud platform & tooling certification	22.6	0.315	2.2	1.9 – 2.5	0.96	Yes
Data engineering for AI	12.4	0.267	2.6	2.2 – 3.0	0.93	Yes
MLOps & AI platform architecture	8.1	0.224	3.1	2.6 – 3.7	0.95	Yes
AI governance, risk & assurance	2.8	0.161	4.3	3.4 – 5.4	0.89	Yes
Domain knowledge (BFSI / health / telecom)	4.9	0.102	6.8	5.5 – 8.4	0.91	No
Foundational software & CS engineering	6.4	0.075	9.2	7.3 – 11.6	0.91	No
Leadership & client management	4.7	0.060	11.5	8.9 – 14.8	0.90	No
Spend-weighted portfolio	100.0	0.222	3.12	—	—	—

Note: Half-life $H = \ln 2 / \lambda$. Confidence intervals bootstrapped over 2,000 resamples. GenAI-exposed clusters average 2.16 years against 5.91 for non-exposed clusters of comparable technical depth: $t(1238) = 21.4$, $p < 0.001$. Half-lives above roughly seven years are extrapolated beyond the 36-month observation window and should be read as indicative.




Figure 3 · Estimated exponential decay curves for four representative skill clusters, with observed measurement waves.

8.2 The Amortisation Horizon and Its Rigidity

Table 3 reports the horizon. The four component practices give closely similar values — a certification currency window of 4.9 years, a curriculum refresh cycle of 3.8 years, a business-case recovery horizon of 5.2 years and a skills-inventory decay rule of 5.1 years — producing a composite mean horizon of 4.65 years on a spend-weighted basis, with a standard deviation across organisations of 1.6 years.

The contractual horizon is a striking outlier. Training bonds and service agreements average 1.9 years, less than half the accounting horizon applied by the same organisations. Firms therefore protect their investment contractually for a period that implies they expect recovery within two years, while simultaneously recognising benefit in their management accounts for nearly five. Interviews confirmed that no organisation in the sample had reconciled the two figures, and in several cases the leaders responsible for each were unaware of the other. The bond period is in fact much closer to the observed half-life of the technology clusters into which most spend flows, which suggests that legal intuition about the durability of the investment has been more accurate than the accounting.

Figure 2 presents the study's central visual result. The horizon line is nearly flat across clusters, ranging only from 4.2 to 5.3 years, while the half-life line rises from 1.1 to 11.5. The coefficient of variation of the horizon across clusters is 0.075 against 0.784 for the half-life — half-life varies 10.4 times more than the horizon applied to it. Levene's test confirms the difference in dispersion ($p < 0.001$), supporting H_2 .

The two lines cross at a half-life of approximately 4.6 years. To the left of the crossover the firm amortises over a horizon longer than the asset lasts, and budget is stranded; to the right it amortises over a horizon shorter than the asset lasts, and durable capability is under-credited. Eighty-four per cent of expenditure sits to the left. The remaining sixteen per cent is the mirror-image error and deserves note: because the firm does not recognise that foundational engineering, domain knowledge and leadership capability retain value for eight to twelve years, it appraises investment in them against a five-year horizon and systematically under-funds them. Both errors follow from the same cause, which is a single horizon applied to a heterogeneous portfolio.

Figure 2 · Horizon Rigidity — a Near-Flat Amortisation Horizon Applied to Skills Whose Half-Lives Differ Tenfold




Figure 2 · Horizon rigidity: a near-flat amortisation horizon applied across clusters whose half-lives differ by a factor of ten.

8.3 The Amortisation Gap and the Stranded Budget

Table 5 computes the gap and its financial consequence cluster by cluster. The spend-weighted horizon of 4.65 years against a spend-weighted half-life of 3.12 gives a gap of 1.53 years and a ratio Γ of 1.49. H_1 is supported: for 84.0 per cent of expenditure the horizon exceeds the half-life, and for the largest single cluster it exceeds it fourfold.

Applying the discounted specification at $r = 12$ per cent produces a spend-weighted stranded share of 43.7 per cent. Against the recovery multiple of 2.84 times that firms assume in their own business cases, the realised multiple is 1.60 times. Expressed in the terms most useful to a budget holder: of every ₹100 committed to human resource development in this sample, ₹43.70 funds capability whose economic value has decayed before the firm's own recovery horizon closes.

The dispersion across clusters is wider than the aggregate suggests and is where the actionable content lies. Generative AI prompt and application patterns strand 62.4 per cent of committed budget and return 1.07 times — barely recovering the outlay. Framework-specific engineering strands 54.4 per cent and cloud certification 45.3 per cent. These three clusters together absorb 60.7 per cent of total expenditure. At the durable end, leadership development strands 13.0 per cent and returns 2.47 times, and foundational engineering 15.1 per cent and 2.41 times. The organisation is losing most where it spends most, which is the signature of a portfolio allocated on visible market premium rather than on durability.

Figure 4 shows the mechanism at cluster level, taking cloud certification as the illustration. The firm recognises benefit in equal annual instalments while the actual flow decays exponentially, so recognition runs ahead of realisation from the first year and the cumulative gap widens until the horizon closes. Figure 7 then decomposes the ₹43.70 for the portfolio as a whole. Value decayed before the recovery horizon closes accounts for ₹31.20, horizon over-extension beyond the economically justified period for ₹8.10, and the straight-line versus declining-balance shape error for ₹4.40. The decomposition matters because the three components have different remedies. The first is addressed by changing what is bought — shifting composition toward more durable clusters, or buying perishable capability closer to the point of use. The second is addressed by shortening the horizon. The third is addressed by changing only the amortisation method, at no cash cost whatever, and is therefore the cheapest ₹4.40 any learning function will ever recover.

Table 5 · The Amortisation Gap and the Stranded Share of HRD Budget, by Skill Cluster

Skill Cluster	Spend share (%)	H (yrs)	A (yrs)	Gap A – H	$\Gamma = A / H$	Stranded share (%)	₹ stranded per ₹100	Realised multiple
Generative AI prompt & application	21.3	1.1	4.4	+3.3	4.00	62.4	62.40	1.07×
Framework-specific AI engineering	16.8	1.4	4.2	+2.8	3.00	54.4	54.40	1.29×
Cloud platform & tooling certification	22.6	2.2	4.9	+2.7	2.23	45.3	45.30	1.55×
Data engineering for AI	12.4	2.6	4.6	+2.0	1.77	39.2	39.20	1.73×
MLOps & AI platform architecture	8.1	3.1	4.8	+1.7	1.55	35.6	35.60	1.83×
AI governance, risk & assurance	2.8	4.3	4.5	+0.2	1.05	26.6	26.60	2.09×
Domain knowledge (BFSI / health)	4.9	6.8	5.1	–1.7	0.75	20.0	20.00	2.27×
Foundational software & CS engineering	6.4	9.2	5.0	–4.2	0.54	15.1	15.10	2.41×
Leadership & client management	4.7	11.5	5.3	–6.2	0.46	13.0	13.00	2.47×
SPEND-WEIGHTED PORTFOLIO	100.0	3.12	4.65	+1.53	1.49	43.7	43.70	1.60×

Note: Stranded share $S = 1 - PV(actual)/PV(planned)$, with $PV(actual) = B_0[1 - e^{-(\lambda+r)A}]/(\lambda+r)$ and $PV(planned) = B_0[1 - e^{(-r)A}]/r$, at $r = 12\%$. Realised multiple $= 2.84 \times (1 - S)$, where $2.84\times$ is the recovery multiple assumed in the organisations' own business cases. Clusters with a positive gap are over-amortised; those with a negative gap are under-amortised and the durable capability is correspondingly under-funded.

Figure 4 · Straight-Line Amortisation Against an Exponentially Decaying Benefit




Figure 4 · Straight-line amortisation against an exponentially decaying benefit flow, and the resulting cumulative recovery gap.

Figure 7 · The Misallocated HRD Budget




Figure 7 · Decomposition of the stranded HRD budget, and the effect of horizon differentiation.

8.4 Return to Training Capital and the Effect of the Gap

Table 7 and Figure 5 report the central relationship. The model explains 58.6 per cent of variance in the realised recovery multiple, $F(8, 1231) = 217.8, p < 0.001$, with standard errors clustered at organisation level.

The amortisation gap carries a coefficient of -0.284 ($p < 0.001$), supporting H_3 : each unit increase in Γ reduces the realised recovery multiple by 0.284. Moving from a well-matched horizon at $\Gamma = 1$ to the generative-AI cluster's $\Gamma = 4.00$ costs approximately 0.85 on the multiple, which is close to a third of the planned return. Horizon differentiation by cluster carries a positive coefficient of 0.371 ($p < 0.001$), supporting H_5 , and the group comparison is equally clear: the eight organisations that differentiate strand ₹27.90 per ₹100 and realise 2.18 times, against ₹48.60 and 1.44 times for the thirty that apply a uniform horizon, $t(36) = 4.21, p < 0.001$.

Three controls carry substantive content. Deployment intensity is positive and significant ($\beta = 0.0094, p < 0.001$), which links this study to the first instalment of the programme and is developed in the next section. Curriculum refresh frequency is positive ($\beta = 0.163, p = 0.002$), indicating that organisations refreshing content more often recover more, consistent with refresh acting as partial re-investment that resets the decay clock. Generative-AI exposure of the cluster carries a negative coefficient ($\beta = -0.196, p = 0.004$) even after the gap is controlled, indicating that exposure damages returns through channels additional to the shortened half-life — most plausibly through the greater volatility of demand for those skills. Firm size and employer type are not significant. Figure 5 displays the relationship with marker area proportional to expenditure share. The visual point is uncomfortable: the largest markers cluster at the high- Γ , low-return end of the fitted line. The portfolio is weighted toward precisely the region where the return is worst.

Table 7 · Realised Return to Training Capital Regressed on the Amortisation Gap

Predictor	β	Clustered SE	t	p	Interpretation
Constant	2.412	0.086	28.05	< 0.001	Recovery multiple at centred predictors
Amortisation gap $\Gamma = A / H$	-0.284	0.031	-9.16	< 0.001	Each unit of Γ costs 0.284 on the multiple — supports H_3
Horizon differentiated by cluster (1 = yes)	0.371	0.078	4.76	< 0.001	Differentiation raises realised return — supports H_5
Deployment intensity (% billable hours)	0.0094	0.0021	4.48	< 0.001	Links to the deployment channel of Section 8.6
Curriculum refresh frequency (per year)	0.163	0.052	3.13	0.002	Refresh acts as partial re-investment, resetting the decay clock
Training spend per employee (ln ₹)	0.118	0.044	2.68	0.008	Scale of investment matters, but far less than its horizon
Generative-AI exposed cluster (1 = yes)	-0.196	0.067	-2.93	0.004	Exposure damages return beyond the shortened half-life
Firm size (ln headcount)	0.037	0.039	0.95	0.343	Not significant
Global capability centre employer (1 = yes)	0.081	0.062	1.31	0.192	Not significant

Note: $N = 1,240$ completions in 38 organisations. $R^2 = 0.586$; adjusted $R^2 = 0.583$; $F(8, 1231) = 217.8$, $p < 0.001$. Standard errors clustered at organisation level. Dependent variable is the realised recovery multiple on training capital. Group comparison of differentiating against non-differentiating organisations: $2.18\times$ vs $1.44\times$, $t(36) = 4.21$, $p < 0.001$.

Figure 5 · Realised Return to Training Capital Falls as the Amortisation Gap Widens
(marker area = share of HRD spend; colour = stranded share of that cluster's budget)




Figure 5 · Realised return to training capital against the amortisation gap Γ , with marker area proportional to expenditure share.

8.5 What Extends the Life of a Trained Skill

Table 6 reports the Weibull accelerated failure time model. The shape parameter is 1.34, significantly greater than one, indicating that the hazard of obsolescence rises with the age of the training — skills do not merely decay, they decay at an accelerating rate, which is the expected signature of a moving technology frontier.

Two characteristics dominate. Foundational or conceptual content, as against tool-specific content, carries a time ratio of 2.836 ($p < 0.001$): conceptually framed training survives 2.8 times as long as training built around a specific tool. Vendor-neutral certification survives 1.41 times as long as vendor-specific certification. These are the two largest modifiable factors in the model, and neither concerns how much is spent — both concern how the content is framed. A firm can extend the effective life of its training portfolio substantially without increasing its budget, simply by shifting the balance between conceptual and tool-specific framing.

Refresher training within twelve months of the original completion extends survival by 52.3 per cent, confirming that decay is partially reversible at modest cost and giving an empirical basis for a refresh cadence tied to half-life rather than to the annual planning cycle. Application of the skill across two or more client domains extends survival by 28.8 per cent, consistent with the matching logic that breadth of application both reinforces the capability and diversifies the demand supporting it. Generative-AI exposure shortens survival to 0.594 of the baseline. Employee experience above ten years and the cohort size of the training programme are not significant, which is worth stating explicitly: skill durability is a property of the skill and its use, not of the seniority of the person holding it or the scale on which it was delivered.

Table 6 · Weibull Accelerated Failure Time Model of Skill Survival (event: skill-value index falls below 50)

Covariate	Time ratio	95% CI	z	p	Effect on expected skill survival
Foundational / conceptual content (vs tool-specific)	2.836	2.41 – 3.34	12.44	< 0.001	Survives 2.8 $\times$ as long; largest modifiable factor in the model
Refresher training within 12 months	1.523	1.33 – 1.74	6.15	< 0.001	Extends survival 52.3%; decay is partially reversible
Vendor-neutral certification (vs vendor-specific)	1.412	1.24 – 1.61	5.20	< 0.001	Survives 41% longer at comparable cost
Skill applied across two or more client domains	1.288	1.13 – 1.47	3.79	< 0.001	Breadth of application reinforces and diversifies demand
Deployment intensity (per 10 percentage points)	1.147	1.108 – 1.187	7.79	< 0.001	Use preserves capability — H_0 rejected
Employee experience above 10 years	1.094	0.96 – 1.25	1.32	0.187	Not significant — durability is not a function of seniority



Programme cohort size (per 100)	0.981	0.94 – 1.02	–0.92	0.358	Not significant — scale of delivery does not affect survival
Generative-AI exposed cluster	0.594	0.52 – 0.68	–7.41	< 0.001	41% reduction in expected survival — supports H ₄

Note: 1,240 completions in 38 organisations. Weibull shape parameter $p = 1.34$ (95% CI 1.26–1.43), significantly above 1, indicating that the hazard of obsolescence accelerates with the age of the training. Log-likelihood = $-1,842.6$; $LR \chi^2(8) = 486.3$, $p < 0.001$. Model adequacy assessed by Cox–Snell residuals. Time ratios above 1 indicate longer expected survival.

8.6 Decay Is Partly an Allocation Outcome

Table 8 and Figure 6 test H₆, and reject it decisively. Half-life rises monotonically with deployment intensity: 1.6 years where the skill is applied to less than fifteen per cent of billable hours, 2.4 years at fifteen to thirty-five per cent, 3.4 years at thirty-six to sixty per cent, and 4.6 years above sixty per cent. Skills in heavy use decay 2.9 times more slowly than skills that are trained and shelved. The relationship is well described by $\lambda = 0.487 - 0.0051 \times \text{deployment percentage}$ ($t = -11.3$, $p < 0.001$), and the group differences are significant, $F(3, 1236) = 96.4$, $p < 0.001$.

This is the finding with the widest implications, and it revises the interpretation of both preceding studies in this programme. Deployment does not merely price capability higher, as the first instalment established; it preserves capability, so a deployed skill is more valuable per period and valuable for more periods. The two effects compound. It also gives the reassignment intervals measured in the second instalment an economic price: at the fastest observed decay rates, capability sitting undeployed loses value at roughly 0.63 per year in log terms, so the seventy-one-day difference between the weakest and strongest redeployment-policy quartiles reported there is not merely a utilisation loss but a permanent write-down of training capital already committed.

It also complicates the obsolescence literature's central distinction. The technological-obsolescence account holds that the worker has not changed but the market has, in which case use should be irrelevant to decay. The wear-and-tear account holds that unused skills atrophy. The pattern here is consistent with wear-and-tear operating alongside technological obsolescence rather than in place of it: the deployment gradient is steep, but even the most heavily deployed quartile records a half-life of only 4.6 years, well short of the durable clusters. Use slows decay considerably; it does not stop it.

For practice the implication is direct and cheap. A training investment that is not followed within months by deployment on billable work is not merely under-realised, it is actively depreciating faster than an otherwise identical investment that is deployed. Training calendars set independently of staffing plans therefore destroy value in a way that neither function currently measures, because the loss appears in neither the learning budget nor the utilisation report.

Table 8 · Skill Half-Life by Deployment Intensity Quartile (test of H₆)

Deployment quartile	Share of billable hours applying the skill	n	Decay rate λ (per year)	Half-life H (years)	Ratio to Q1	Mean skill value at 24 months
Q1 (lowest)	Below 15%	308	0.433	1.6	1.00	35.4
Q2	15 – 35%	312	0.289	2.4	1.50	50.1
Q3	36 – 60%	311	0.204	3.4	2.13	64.7
Q4 (highest)	Above 60%	309	0.151	4.6	2.88	74.0
All completions	Mean 38.4%	1,240	0.222	3.12	1.95	56.1

Note: ANOVA on estimated half-life across quartiles: $F(3, 1236) = 96.4$, $p < 0.001$. Continuous specification: $\lambda = 0.487 - 0.0051 \times \text{deployment percentage}$, $t = -11.3$, $p < 0.001$, $R^2 = 0.34$. H₆ — that half-life is invariant to deployment intensity — is rejected. Note that even the most heavily deployed quartile records a half-life of 4.6 years, well below the durable clusters: use slows decay considerably but does not stop it.




Figure 6 · Skill half-life and decay rate by deployment intensity quartile — the rejection of H_0 .

8.7 A Cluster-Differentiated Amortisation Schedule

Table 10 converts the findings into the schedule the study was designed to produce. For each cluster it specifies a recommended horizon, an amortisation method and a refresh trigger.

Three principles govern it. First, the horizon should be set as a multiple of the observed half-life rather than by convention, with 1.5 half-lives serving as a practical rule that captures the majority of discounted benefit while remaining short enough to be credible. Second, amortisation should be declining-balance at the estimated decay rate λ rather than straight-line, since the benefit flow is exponential; this change alone recovers ₹4.40 per ₹100 at no cash cost. Third, and least conventional, skills with half-lives beyond eight years should be removed from the amortising category altogether and reclassified as capability endowment. Foundational engineering and leadership capability persist longer than any credible planning horizon, and treating them as depreciating assets on a five-year schedule causes the systematic under-investment visible on the right-hand side of Figure 2. They are better managed as a stock to be maintained than as an outlay to be recovered.

The schedule also specifies refresh triggers, which the evidence suggests matter more than refresh calendars. For the fastest clusters the appropriate trigger is a version or model change in the underlying technology rather than the passage of a fixed interval, because that is what actually destroys the value. Given the 52.3 per cent survival extension from refresher training within twelve months, a trigger-based cadence in the perishable clusters is likely to be both cheaper and more effective than the annual cycle currently in use.

Table 10 · Recommended Cluster-Differentiated Amortisation Schedule

Skill Cluster	Observed H (yrs)	Current A (yrs)	Recommended horizon A*	Recommended method	Refresh trigger
Generative AI prompt & application	1.1	4.4	1.5 yrs	Expense in year incurred; treat as consumable	Model or platform version change
Framework-specific AI engineering	1.4	4.2	2.0 yrs	Declining balance at $\lambda = 0.495$	Major framework release
Cloud platform & tooling certification	2.2	4.9	3.5 yrs	Declining balance at $\lambda = 0.315$	Vendor certification refresh cycle
Data engineering for AI	2.6	4.6	4.0 yrs	Declining balance at $\lambda = 0.267$	Architecture or stack migration
MLOps & AI platform architecture	3.1	4.8	4.5 yrs	Declining balance at $\lambda = 0.224$	Platform generation change
AI governance, risk & assurance	4.3	4.5	6.5 yrs	Declining balance at $\lambda = 0.161$	Regulatory instrument in force
Domain knowledge (BFSI / health)	6.8	5.1	8.0 yrs	Declining balance at $\lambda = 0.102$	Sector regulation or product cycle
Foundational software & CS engineering	9.2	5.0	Not amortised	Reclassify as capability endowment	Maintain as stock; no expiry
Leadership & client management	11.5	5.3	Not amortised	Reclassify as capability endowment	Maintain as stock; no expiry



Note: Recommended horizons follow the rule $A^* \approx 1.5 H$, capped at eight years; beyond that the schedule recommends reclassification rather than a longer horizon, because capability persisting longer than any credible planning horizon is better managed as a stock to be maintained than as an outlay to be recovered. Adopting declining-balance in place of straight-line amortisation recovers approximately ₹4.40 per ₹100 of expenditure at no cash cost.

8.8 Summary of Hypothesis Testing

Table 9 consolidates the tests. Five hypotheses are supported and one is rejected. The rejection of H_6 is the substantively richest result, because it converts skill obsolescence from a purely exogenous technological process into one that human resource and staffing decisions partly govern. If half-life were invariant to use, the only available response to shortening half-lives would be to spend less or spend later; because it is not, the firm can extend the life of what it has already bought by changing how quickly and how intensively trained capability is put to work — which is precisely the machinery examined in the second instalment of this programme.

Table 9 · Summary of Hypothesis Testing

Code	Hypothesis	Test Applied	Statistic	p	Decision
H ₁	The amortisation horizon exceeds the skill half-life ($\Gamma > 1$) for the majority of HRD expenditure	Spend-weighted comparison of A against H	$\Gamma = 1.49$ 84.0% of spend	< 0.001	SUPPORTED
H ₂	The horizon varies significantly less across clusters than the half-life does — horizon rigidity	Levene's test; comparison of coefficients of variation	CV 0.075 vs 0.784 ratio 10.4×	< 0.001	SUPPORTED
H ₃	Realised return to training capital declines as the amortisation gap widens	OLS with organisation-clustered SE	$\beta = -0.284$ $t = -9.16$	< 0.001	SUPPORTED
H ₄	Generative-AI-exposed clusters exhibit shorter half-lives than non-exposed clusters	Independent-samples t-test; Weibull AFT	2.16 vs 5.91 yrs TR = 0.594	< 0.001	SUPPORTED
H ₅	Organisations differentiating the horizon by cluster realise a higher return	Welch's t-test; OLS coefficient	2.18× vs 1.44× $t(36) = 4.21$	< 0.001	SUPPORTED
H ₆	Skill half-life is invariant to deployment intensity	ANOVA across quartiles; continuous decay model	$F(3,1236) = 96.4$ $t = -11.3$	< 0.001	REJECTED

IX. FINDINGS

- Skill half-lives in the Indian IT services training portfolio range from 1.1 years for generative-AI prompt and application patterns to 11.5 years for leadership and client management, a ratio of 10.5 to one, with exponential decay fits of R^2 between 0.89 and 0.96.
- The amortisation horizon firms apply ranges only from 4.2 to 5.3 years. Half-life varies 10.4 times more across clusters than the horizon applied to it (CV 0.784 against 0.075) — the study's core finding of horizon rigidity.
- The spend-weighted horizon is 4.65 years against a spend-weighted half-life of 3.12 years: a gap of 1.53 years and a ratio $\Gamma = 1.49$. The two lines cross at a half-life of approximately 4.6 years, and 84.0 per cent of expenditure sits on the over-amortised side.
- The spend-weighted stranded share is 43.7 per cent. Of every ₹100 committed to human resource development, ₹43.70 funds capability whose value decays before the firm's own recovery horizon closes; the realised recovery multiple is 1.60× against the 2.84× assumed in business cases.
- Stranding is concentrated where spending is heaviest. The three fastest-decaying clusters absorb 60.7 per cent of expenditure and strand between 45.3 and 62.4 per cent of it; generative-AI application training returns 1.07× — barely the outlay.
- The stranded amount decomposes into ₹31.20 of decay before recovery, ₹8.10 of horizon over-extension and ₹4.40 of straight-line shape error. The last is recoverable by changing the amortisation method alone, at no cash cost.
- Realised return falls by 0.284 for each unit increase in Γ ($p < 0.001$). Moving from a matched horizon to the generative-AI cluster's $\Gamma = 4.00$ costs roughly a third of the planned return.
- Organisations that differentiate the horizon by cluster — 8 of 38 — strand ₹27.90 per ₹100 and realise 2.18×, against ₹48.60 and 1.44× for those applying a uniform horizon ($t(36) = 4.21$, $p < 0.001$).
- Half-life is not invariant to use. Skills deployed on more than sixty per cent of billable hours decay 2.9 times more slowly than those below fifteen per cent; $\lambda = 0.487 - 0.0051 \times \text{deployment per cent}$ ($t = -11.3$, $p < 0.001$). H_6 is rejected and obsolescence is established as partly an allocation outcome.
- Conceptually framed training survives 2.84 times as long as tool-specific training, and vendor-neutral certification 1.41 times as long as vendor-specific. Neither depends on how much is spent — both depend on how content is framed.
- Refresher training within twelve months extends skill survival by 52.3 per cent, and application across two or more client domains by 28.8 per cent. The Weibull shape parameter of 1.34 indicates that obsolescence accelerates with the age of the training.
- The contractual horizon — training bonds averaging 1.9 years — is less than half the accounting horizon applied by the same organisations, and is markedly closer to the observed half-life of the clusters receiving most expenditure. No organisation in the sample had reconciled the two figures.



X. SIGNIFICANCE OF THE STUDY

- For human resource development practice, the study replaces an administrative convention with a measured quantity, and supplies both the instrument for measuring a firm's own horizon and the cluster-differentiated schedule to replace it.
- For finance functions, it converts an unmeasured loss into a budget line: ₹43.70 stranded per ₹100 committed, decomposed into three components with distinct and differently priced remedies, one of which costs nothing.
- For accounting and disclosure, it demonstrates a consequence of expensing training under IAS 38 that has not previously been quantified — because no recovery horizon must be stated, none is tested, and a materially wrong one survives indefinitely inside management accounting.
- For staffing and workforce planning, it establishes that training calendars set independently of deployment plans destroy value measurably, since undeployed capability decays roughly three times faster than deployed capability.
- For the academic literature, it is the first study to estimate skill decay rates and firm amortisation horizons from the same data and compare them, and the first to express their mismatch as a stranded share of expenditure.
- For the research programme of which it forms part, it completes the sequence of production, allocation and consumption of capability, and shows that the returns reported in the first two instalments require downward revision once decay is priced.

XI. RECOMMENDATIONS AND SUGGESTIONS

11.1 For Employers

- Measure your own amortisation horizon before revising it. The four practices in Table 3 are documented inside every organisation and can be read off in a week; most firms do not know what horizon they are applying because it has never been stated in one place.
- Differentiate the horizon by skill cluster. This single change separates the organisations stranding ₹27.90 per ₹100 from those stranding ₹48.60, and it requires no additional expenditure.
- Switch from straight-line to declining-balance amortisation at the estimated decay rate. It recovers ₹4.40 per ₹100 at no cash cost and is the least contested change available.
- Reclassify skills with half-lives beyond eight years — foundational engineering, domain knowledge, leadership — as capability endowment rather than amortising training capital. Appraising them against a five-year horizon systematically under-funds the most durable assets in the portfolio.
- Pair every training commitment with a deployment commitment inside six months. Undeployed capability decays roughly three times faster, and the loss appears in neither the learning budget nor the utilisation report.
- Shift the balance of content from tool-specific to conceptual framing and from vendor-specific to vendor-neutral certification. These are the two largest modifiable determinants of skill survival and neither requires a larger budget.
- For the fastest-decaying clusters, buy capability closer to the point of use and trigger refresh on technology version change rather than on the annual calendar. Given a 1.1-year half-life, training delivered a year ahead of deployment has lost half its value before it is applied.
- Reconcile the training bond period with the accounting horizon. That the legal function assumes two years while the finance function assumes five is not a defensible position for either, and the evidence here suggests the legal estimate is the better one.

11.2 For Professionals

- Weigh durability alongside premium when choosing what to learn. The clusters paying the highest visible premium are also those decaying fastest, and a certification with a 1.1-year half-life must be renewed continuously to hold its value.
- Prefer conceptual and vendor-neutral capability where a choice exists. It survives roughly three times as long at comparable cost in effort.
- Secure application of a new skill on billable work quickly. Capability applied to more than sixty per cent of billable hours retains value nearly three times as long as capability that is certified and shelved.

11.3 For Policy and Academia

- Calibrate national reskilling programmes to observed half-lives. A programme certifying capability with a one-year half-life and reporting outcomes on a three-year cycle will report benefits that no longer exist at the point of reporting.
- Consider whether the expensing treatment of training under IAS 38 is serving its purpose. The standard's reasoning is that the firm does not control the asset; the consequence observed here is that no horizon is ever stated or tested, and a materially wrong one persists.
- Extend workforce skill accounting to record currency as well as attainment. A certification count that does not decay overstates the national skill stock by an amount this study suggests is substantial.

XII. LIMITATIONS OF THE STUDY

- The cohort is retrospective. Although learning management records fixed the baseline objectively and measurement waves used identical instruments, recall error may affect the earlier waves of the skill-value composite.
- The skill-value index is a constructed composite combining task application, internal demand, market premium and manager rating. It has satisfactory reliability but no external benchmark, and different weightings of its four components would shift the estimated half-lives somewhat.
- The stranded-share estimate depends on the discount rate. Sensitivity testing across 8 to 16 per cent moves the spend-weighted figure between 39.8 and 46.9 per cent, so the conclusion is robust in direction and approximate in magnitude.



- The planned recovery multiple of 2.84 times is taken from the organisations' own business cases and is therefore only as reliable as those cases. Where firms overstate planned benefit, the stranded share computed against it is correspondingly overstated.
- Thirty-six months of observation is short relative to the durable clusters. Half-lives above roughly seven years are extrapolated from the fitted curve rather than observed directly, and should be treated as indicative.
- The exponential form assumes a constant proportional decay rate. The Weibull shape parameter of 1.34 indicates accelerating hazard, which suggests the exponential slightly understates late-period decay; a Gompertz or Weibull decay specification would fit marginally better at the cost of interpretability.
- Deployment intensity and half-life are jointly observed, and reverse causation cannot be excluded — staffing systems may preferentially deploy skills that are independently more durable.
- The sample is confined to organisations maintaining auditable learning records, which biases toward larger and better-instrumented firms and probably understates the horizon rigidity present in the wider industry.

XIII. SCOPE FOR FURTHER RESEARCH

The design supports several direct extensions, and the instruments in Tables 3 and 4 are published in full so that they can be reused. Table 11 sets out a structured agenda of seven studies. Four points deserve emphasis.

1. A prospective panel with randomised or staggered variation in the timing of deployment relative to training completion would resolve the reverse-causation concern around the deployment–decay relationship, which is this study's most consequential and least well identified finding.
2. The three instalments of this programme should be estimated jointly in a single framework in which investment produces capability, allocation realises it and decay consumes it. Each has so far been estimated with the others held implicit, and the present results indicate that the returns reported in the first two require downward revision once decay is priced.
3. A decomposition of decay into wear-and-tear and technological components would settle a question this study can only bound. The deployment gradient establishes that both operate; separating them requires observing skill value for the same capability under differing rates of technological change, which a multi-country or multi-vintage design could supply.
4. Replication in sectors with slower technology cycles would test whether horizon rigidity is a general administrative property of organisations or a specific failure of technology-sector learning functions. The prior is that the horizon is equally rigid everywhere and simply less costly where half-lives are longer, which would make the finding general and its consequence sector-specific.

Table 11 · Structured Research Agenda Extending the HRM Variable → Economic Outcome Programme

#	Proposed Study	HRM Variable → Economic Outcome	Design	Data Requirement	Priority
1	Causal test of the deployment–decay relationship	Deployment timing → skill half-life	Prospective panel with staggered or randomised deployment timing after completion	Completion dates, deployment dates, repeated skill-value measurement	High
2	Joint estimation of production, allocation and decay	Skill investment + redeployment policy + amortisation horizon → output per employee-hour	Unified specification across all three instalments of this programme	Matched training, redeployment and decay data at unit level	High
3	Decomposition of decay into its two sources	Skill use → wear-and-tear versus technological obsolescence	Multi-vintage or multi-country comparison holding use constant	Skill value under differing rates of technological change	High
4	Horizon rigidity outside technology sectors	Amortisation horizon → return to training capital in slow-cycle industries	Replication using the Table 3 instrument in manufacturing, banking and healthcare	Documented horizon practices and skill-value trajectories	Medium
5	Optimal refresh cadence	Refresher timing → extension of skill survival	Field experiment varying refresher interval against half-life	Randomised refresher timing with survival tracking	Medium
6	Disclosure consequences of expensing training	Accounting treatment → quality of HRD investment decisions	Comparative study of firms with and without internal capitalisation policies	Management-accounting horizon documentation and outcome data	Medium
7	Portfolio optimisation under decay	Composition of training spend → realised portfolio return	Simulation and optimisation using estimated cluster half-lives	Cluster-level spend, half-life and premium data	Medium



Note: Studies 1 and 2 carry forward directions identified in the first and second instalments of this programme (Sastry, 2026a, 2026b). The horizon instrument in Table 3 and the skill-value composite in Table 4 are published in full to permit replication without modification.

XIV. CONCLUSION

Firms choose the period over which they expect to recover a training investment, and the technology frontier chooses how long the investment lasts. This study measured both on comparable terms and found that they have come apart. The horizon Indian IT services organisations apply is nearly flat across a portfolio whose components differ tenfold in durability: it varies by a factor of 1.26 across skill clusters while half-life varies by a factor of 10.5. Eighty-four per cent of human resource development expenditure is committed on the wrong side of the crossover.

The consequence is measurable and large. At a spend-weighted horizon of 4.65 years against a half-life of 3.12, and discounting the decaying benefit flow at the firms' own cost of capital, 43.7 per cent of planned recovery does not arrive. Against the 2.84-times multiple these organisations assume in their own business cases, the realised figure is 1.60 times. That is ₹43.70 of every ₹100 committed, and it is concentrated in the clusters receiving the most money — the fastest-decaying three absorb sixty per cent of the budget and strand roughly half of it.

Three of the remedies cost nothing. Differentiating the horizon by cluster separates organisations stranding ₹27.90 from those stranding ₹48.60. Switching from straight-line to declining-balance amortisation recovers ₹4.40 per ₹100 by changing an arithmetic convention. Shifting content toward conceptual framing and vendor-neutral certification extends survival by factors of 2.84 and 1.41 respectively without altering the budget. None of these requires spending more; all require knowing the half-life of what is being bought, which no organisation in this sample did.

The most consequential finding is that decay is not wholly exogenous. Skills applied to more than sixty per cent of billable hours decay 2.9 times more slowly than skills that are certified and shelved. Obsolescence is therefore partly an allocation outcome, which means the training calendar and the staffing plan are not two independent processes that happen to concern the same people — they jointly determine how long the firm's training capital survives. A learning function that schedules delivery without reference to deployment is not merely failing to realise value; it is accelerating the destruction of value it has already paid for.

Read alongside the two studies that precede it, the argument completes. Capability is produced by investment, realised by allocation and consumed by time, and a firm that measures only the first has no basis for believing its own return figures. The Indian IT industry is currently funding the largest reskilling effort in its history against a horizon assumption that its own contract lawyers, setting two-year bonds, appear to have understood better than its finance functions. Correcting that assumption is not an accounting refinement. On the evidence here it is worth close to forty-four per cent of the budget.

REFERENCES

- [1] Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- [2] Anas, M. B., Hans, V. B., & Sastry, N. S. (2025). Digital transformation in tourism: How technology shapes travel experiences. *International Journal of Scientific Research & Engineering Trends*, 11(5), 1–5.
- [3] Arrazola, M., & de Hevia, J. (2004). More on the estimation of the human capital depreciation rate. *Applied Economics Letters*, 11(3), 145–148.
- [4] Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.
- [5] Becker, G. S. (1964). *Human capital: A theoretical and empirical analysis, with special reference to education*. National Bureau of Economic Research / Columbia University Press.
- [6] Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942.
- [7] Cascio, W. F., & Boudreau, J. W. (2011). *Investing in people: Financial impact of human resource initiatives* (2nd ed.). Pearson.
- [8] De Grip, A., & Van Loo, J. (2002). The economics of skills obsolescence: A review. *Research in Labor Economics*, 21, 1–26.
- [9] Deloitte & NASSCOM. (2024). *Bridging the AI talent gap to boost India's tech and economic impact*. Deloitte India.
- [10] Deming, D. J., & Noray, K. (2020). Earnings dynamics, changing job skills, and STEM careers. *Quarterly Journal of Economics*, 135(4), 1965–2005.
- [11] Ernst & Young. (2026). *Future of Pay 2026: India Inc. compensation and workforce trends*. EY India.
- [12] Institute of Chartered Accountants of India. (2023). *Indian Accounting Standard (Ind AS) 38: Intangible Assets*. ICAI.
- [13] Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of Political Economy*, 117(5), 914–940.
- [14] Lev, B., & Radhakrishnan, S. (2005). The valuation of organization capital. In *Measuring Capital in the New Economy* (pp. 73–110). University of Chicago Press.
- [15] Lev, B., & Zarowin, P. (1999). The boundaries of financial reporting and how to extend them. *Journal of Accounting Research*, 37(2), 353–385.
- [16] Mallya, M. M., & Sastry, N. S. (2025). Building future-ready organizations: Leadership strategies for young talent management. *EPRA International Journal of Research and Development*, 10(9), 88–91.
- [17] Mincer, J., & Ofek, H. (1982). Interrupted work careers: Depreciation and restoration of human capital. *Journal of Human Resources*, 17(1), 3–24.
- [18] NASSCOM. (2025). *Strategic Review 2025: The technology industry in India*. National Association of Software and Service Companies.
- [19] Neuman, S., & Weiss, A. (1995). On the effects of schooling vintage on experience-earnings profiles: Theory and evidence. *European Economic Review*, 39(5), 943–955.
- [20] Phillips, J. J., & Phillips, P. P. (2016). *Handbook of training evaluation and measurement methods* (4th ed.). Routledge.
- [21] Rosen, S. (1975). Measuring the obsolescence of knowledge. In F. T. Juster (Ed.), *Education, Income, and Human Behavior* (pp. 199–232). NBER / McGraw-Hill.
- [22] Sastry, N. S. (2013). HR approaches to talent management. *International Journal of Managerial Studies and Research*, 1(2).



- [23] Sastry, N. S. (2020). A study on impact of employees' performance that influence to motivate in tea manufacturing organizational sector. *EPRA International Journal of Economic and Business Review*, 8(6).
- [24] Sastry, N. S. (2025a). Understanding the challenges of women micro-entrepreneurs in Karnataka. *International Journal of Managerial Studies and Research*, 13(8), 12–16. <https://doi.org/10.20431/2349-0349.1308002>
- [25] Sastry, N. S. (2025b). Talent designation and social comparison: Economic drivers of organizational growth – An empirical analysis. *EPRA International Journal of Economics, Business and Management Studies*, 12(10).
- [26] Sastry, N. S. (2026a). Effect of AI-complementary skill development on wage premia in Indian IT services. *International Journal of Scientific Research & Engineering Trends*, 12(4).
- [27] Sastry, N. S. (2026b). Talent redeployment practice and task reallocation efficiency under automation. *International Journal of Scientific Research & Engineering Trends*, 12(5).
- [28] Sastry, N. S., & Mallya, M. M. (2025a). Revolutionizing HR hiring with AI analytics: Enhancing efficiency and attracting top talent. *EPRA International Journal of Research and Development*, 10(9).
- [29] Sastry, N. S., & Mallya, M. M. (2025b). Strategic HRM for economic resilience: Skill diversification, labor market dynamics, and technology disruption. *EPRA International Journal of Research and Development*, 10(11). <https://doi.org/10.36713/epra2016>
- [30] Schultz, T. W. (1961). Investment in human capital. *American Economic Review*, 51(1), 1–17.
- [31] ServiceNow. (2025). India's leap to AI maturity begins with skills: Agentic AI and the future of Indian work. ServiceNow Research.
- [32] Syverson, C. (2011). What determines productivity? *Journal of Economic Literature*, 49(2), 326–365.
- [33] Wheebox. (2026). India Skills Report 2026. Wheebox, Confederation of Indian Industry and All India Council for Technical Education.