

Redesigning Work to Preserve Capability

A Practice Brief on Workplace Learning in Technology-Enabled Job Redesign

Abstract

Technology-enabled work redesign can improve efficiency while unintentionally removing the work through which employees build judgement and procedural fluency. This article examines that risk through Singapore's Model AI Governance Framework for Agentic AI, recent evidence on AI-exposed work, and an anonymised Singapore SME consulting case involving automation of manual operations and finance data-entry activities. The case shows why technology adoption should be coupled with task analysis, role redesign and employee transition support. Building on the case and evidence, the brief introduces the CARE Model for capability-preserving work redesign: identify capability-critical tasks, make explicit allocation decisions between people and technology, retain or replace developmental exposure, and evaluate retained capability. The model positions workplace learning as part of automation and job-redesign architecture rather than a training response added after work has been removed.

Keywords: *workplace learning; job redesign; organisational capability; agentic AI; automation*

1. Practice Context and Problem

Workforce transformation is often framed as a skills problem requiring a training response. The World Economic Forum (2025) reports that 39% of workers' existing skills are expected to be transformed or become outdated by 2030, while 63% of employers identify skills gaps as a major barrier to transformation. Reskilling and upskilling are necessary responses, but they do not by themselves preserve the conditions through which capability is formed.

Routine work is not developmentally neutral. Repeated performance, guided participation, observation and encounters with exceptions can build tacit understanding that formal instruction alone does not reproduce. Workplace-learning research treats participation in work as a central source of learning: workplaces afford or restrict

opportunities to engage in progressively demanding activity (Billett, 2001), and professional knowledge develops through combinations of participation, feedback, consultation and reflection (Eraut, 2004). Lave and Wenger's (1991) account of legitimate peripheral participation similarly explains how newcomers move towards fuller participation through access to practice. When automation absorbs entry-level or routine tasks, it may therefore remove an informal learning pathway even if the task appears low-value when assessed only for throughput.

The practical problem is not whether organisations should automate routine work. It is whether the redesign process makes the capability consequences visible early enough to act on them. Training supplies explanations, instruction and guided practice. Work exposure supplies repeated contact with real variation, consequences and exceptions. Once a workflow has been automated end to end, learning and development functions may no longer control whether employees encounter the situations through which judgement and procedural fluency develop. Capability preservation must therefore enter solution design, function allocation and job redesign before the original exposure disappears.

2. Regulatory and Conceptual Reframing

Singapore's Model AI Governance Framework for Agentic AI, version 1.5, reframes this issue as both a capability and continuity risk. The framework notes that when agents assume entry-level work that has served as training, organisations may face skill degradation and reduced ability to execute critical processes when systems fail or are unavailable. It recommends identifying core job capabilities and providing sufficient training and work exposure to retain foundational skills (Infocomm Media Development Authority [IMDA], 2026). The important contribution is the explicit recognition that work exposure, not training alone, belongs within governance.

This concern has a longer history. Bainbridge (1983) described an irony of automation: the human operator may be removed from routine control yet still be expected to intervene effectively in rare and difficult circumstances. Parasuraman and Manzey (2010) later distinguished automation-related complacency and automation bias, showing why monitoring and approval cannot be assumed to preserve active scrutiny.

Research on out-of-the-loop performance also indicates that reduced active control can weaken situation awareness and the ability to recover when automation fails (Endsley & Kiris, 1995). This literature focuses on system safety, monitoring, and intervention; it establishes the risk but does not provide an SME-oriented process for tracing the learning value of tasks through job redesign.

Work-design research adds a complementary point. Digital technologies do not determine job quality or human contribution on their own; managerial and design choices shape autonomy, task variety, feedback, skill use and learning opportunities (Parker & Grote, 2022). Beane's (2019) study of robotic surgery makes the learning mechanism particularly visible: where the approved work arrangement reduced trainees' access to hands-on participation, some sought informal or hidden routes to acquire skill. The wider lesson is that a technically successful implementation can still degrade the pathway by which novices become capable practitioners.

3. Evidence and Limits of Transfer

Two recent studies illustrate why the form of human participation matters. Brynjolfsson, Li and Raymond (2025), studying 5,172 customer-support agents, found that a generative AI assistant improved productivity, with larger gains among less-experienced workers, and reported evidence consistent with learning for some workers. In that configuration, the employee remained an active performer: the system supplied guidance while the employee interpreted the customer's problem and completed the task. The result should not be generalised to workflows in which AI performs the analysis, and the employee mainly approves the output. Assistance can increase learning when it scaffolds performance; substitution can reduce learning when it removes performance.

Brynjolfsson, Chandar and Chen (2026) report a widening employment gap for workers aged 22 to 25 in highly AI-exposed occupations in United States payroll data, with adjustment occurring substantially through reduced hiring. The study is descriptive, country-specific and concerned with labour-market entry rather than the capability of incumbent employees. It does not establish that Singapore SMEs are reducing entry-level hiring, nor that automation has caused skill loss. Its relevance here is narrower: it

identifies a plausible pipeline risk that organisations should investigate locally. If fewer newcomers enter roles and the remaining roles contain less foundational work, the supply of experienced future employees may weaken even when current incumbents remain productive.

4 Approach and Methodology

4.1 Inquiry design and organisational context

This article uses a practice-based reflective case design. It draws on a completed job-redesign engagement with an anonymised Singapore commodity-trading SME that employed fewer than 20 local employees at the time of application. The company was integrating technology into operations and finance processes while redesigning the Shipping Operations Executive role. The project spanned 24 weeks from April 2024, focusing on work redesign and transforming selected finance operations. The original engagement aimed at organisational improvement rather than research, and CARE was not specified or tested prospectively. The case is therefore used to generate and examine a practice proposition, not to demonstrate effectiveness or causal impact.

4.2 Author role and participants

The author served as a project consultant and coach and had direct access to the diagnosis, redesign and transition artefacts. The approved scope identified ten local employees as affected by the redesign. The final delivery record documents two separate one-hour coaching conversations: one with an impacted operations employee and one with the reporting officer or direct supervisor. The reports also refer more broadly to stakeholder discussions used to validate the redesigned role, but do not provide a complete attendance count for those discussions. This practitioner position provided detailed knowledge of the intervention but also creates a risk of confirmation bias because the same person contributed to the work and interpreted it for publication. The unit of analysis in this brief is the role and its task architecture, not the performance of named individuals.

4.3 Data sources and analysis

The retained project corpus comprised the job-task analysis, records of operational pain points and redesign opportunities, the redesigned role and career pathway, the Job

Role and Skills Transition Map, group-coaching content, and the implementation plan. These artefacts were reviewed in two stages. First, they were used to reconstruct the original sequence of planning, intervention and observed project outputs. Second, task categories were re-examined through four questions: what capability was developed through performance; how work was or could be allocated between employee and technology; what exposure should be retained or recreated; and what evidence would show that capability remained available. This second stage produced the retrospective CARE analysis in Table 3.

The inquiry can be mapped cautiously to a plan-act-observe-reflect cycle. Planning consisted of diagnosing operational challenges and analysing the role. Actions included technology-enabled process redesign, revised role expectations, transition mapping, and coaching. Observation was limited to project artefacts and stakeholder engagement during delivery; there was no structured post-implementation measurement of capability or performance. Reflection involved comparing the intervention with automation, workplace-learning and work-design research and identifying what the original design had not measured. Because observation was limited and CARE emerged retrospectively, this may not be fully presented as a complete action-research cycle.

4.4 Ethics, permission and confidentiality

The organisation is not assigned a pseudonym, and participants are not named. The source coaching logs contained personal identifiers, but these were excluded from the analysis and manuscript. Commercially sensitive process material is reported only at task-category level, and no individual performance data or attributable quotations are used. Written client permission for anonymised publication was obtained and retained before submission.

5. Practice intervention in a Singapore SME

The SME traded commodities across regional markets and relied on Google Sheets to track deals, contracts, shipments, payments and incentives. Manual entry supported reviews, checks, approvals and payment release, but the Phase 1 diagnosis identified time burden, error exposure, inconsistent processes, limited analytical capacity, data-quality concerns and constraints on scalability. Seven process areas were selected for automation or

system support: sales-contract generation, operations tracking-sheet entry, finance tracking workflows, automated shipment reminders, costing-sheet modification, packing-list preparation and certificate-of-analysis preparation.

The redesign covered the Shipping Operations Executive role and two affected local employees. The job-task analysis repositioned the role from a concentration on data entry and coordination towards four broader work functions: resolving trade disputes; facilitating product deployment; managing trade documentation and reconciliations; and supporting technology implementation and continuous improvement. The transition architecture then connected the redesigned role to a career pathway, skills inventory, Job Role and Skills Transition Map, two coaching conversations, an OJT plan and an implementation schedule. The available project record supports a before-and-after account at task-category level rather than a time-and-motion study. Table 1 reports documented task changes and capability implications. The records do not quantify the percentage of time previously spent on data entry or the hours saved after implementation; therefore, planned productivity benefits are not presented as measured outcomes.

Table 1. Before and after task architecture of the Shipping Operations Executive role

Task area	Before redesign	Technology-enabled change	Redesigned human contribution
Commercial and operations records	Manual Google Sheets entry for deals, contracts and shipment tracking.	Sales-contract generation and operations tracking selected for automation.	Check source data, reconcile records and investigate erroneous trade transactions.
Finance-related workflows	Manual updating of payment, incentive, finance-tracking and costing records.	Finance tracking and costing-sheet steps selected for system support.	Collaborate with sales and finance, validate documentation and resolve mismatches.
Shipment documentation	Manual reminders and preparation of packing lists and certificates of analysis.	Automated reminders and document-generation support proposed.	Monitor shipment requirements, coordinate providers and manage non-conformance or delays.

Task area	Before redesign	Technology-enabled change	Redesigned human contribution
Role contribution	Routine processing and coordination limited capacity for improvement work.	Reduced routine processing created scope for an enlarged operations-and-technology contribution.	Support technology implementation, analyse operations indicators and plan continuous improvement.

The project also recorded stakeholder-reviewed changes in job design. On a ten-point scale, assessed criticality increased from 3 to 7, impact on business outcomes from 3 to 8, and value-add from 4 to 8. Work autonomy increased from 3 to 7, decision-making responsibility from 3 to 7, and work complexity from 4 to 8. These ratings show that the redesigned role was intended to be larger and more valuable; they are design assessments, not evidence that productivity, capability or employee satisfaction subsequently improved. The Job Role and Skills Transition Map linked changes in work to changes in responsibility and capability. At a minimum, it made visible the shift from manual transaction processing to data validation, exception handling, process coordination, and higher-value problem solving. Group coaching helped employees understand the changed role, interpret new expectations, and plan the transition. These mechanisms were useful because they addressed adoption and role clarity rather than assuming that a system implementation would automatically produce a more capable workforce.

The engagement nevertheless left an important question unanswered. The intervention identified what employees would need to do in the redesigned role, but it did not test whether reduced transaction exposure weakened process understanding or manual fallback ability. Nor did it specify which original tasks were developmentally critical and which could be removed without consequence. That omission became the starting point for the CARE model.

5. The CARE Model for Capability Preserving Work Redesign

The CARE model is proposed from the case and literature; it was not applied during the original engagement. It adds a capability-retention lens to technology governance and job redesign through four linked actions: C, map capability-critical tasks; A, make deliberate allocation decisions; R, retain or replace developmental exposure; and E, evaluate retained capability. Its central question is: what must employees continue to know, experience and be able to do even when technology can perform the task?

Table 2. CARE model diagnostic questions

Stage	Design question	Required output
C Capability critical task mapping	Which tasks build judgement, process understanding, stakeholder awareness or exception-handling capability?	A task-capability map identifying what could weaken if exposure disappears.
A Allocation decisions	Which tasks should be automated, augmented or deliberately retained for human performance?	An explicit function-allocation decision with capability and resilience rationale.
R Retain or replace exposure	Where work is removed, what equivalent practice, rotation, simulation or supervised exception exposure is required?	A minimum exposure design by employee group and career stage.
E Evaluate retained capability.	Can employees explain, challenge and, where necessary, execute critical work when technology is unavailable?	Measures, thresholds, review cadence, owners and corrective actions.

6.1 What CARE adds

CARE does not replace function allocation, human-centred automation, skill-fade controls or the IMDA framework. Function allocation asks which functions should be assigned to humans or machines; human-centred automation seeks systems that support meaningful human control; and safety-critical skill-fade approaches use recurrent practice, checks and refresher training to preserve readiness. IMDA provides governance principles and explicitly identifies work exposure as a control. CARE adds a compact bridge between these ideas and ordinary job-redesign practice. It traces capability from the task that develops it, through the allocation decision, to the replacement exposure and evidence of retention. It also applies the logic before automation, not only after declining performance is detected.

The distinction between automation eligibility and automation desirability is central. A task may be technically feasible and economically attractive to automate but still warrant selective human performance because it develops diagnostic understanding or preserves recovery capability. This does not justify retaining routine work indiscriminately. It requires a reasoned decision about who needs exposure, to what type of case, how often and for what measurable purpose.

6.2 Retrospective application to the SME case

Table 3 applies CARE retrospectively to four task categories supported by the case record. It is a worked redesign exercise rather than evidence that the organisation adopted the actions shown. The analysis indicates what might have been decided differently if capability retention had been an explicit design objective.

Table 3. Retrospective CARE application to the SME case

Task and capability at risk	Allocation decision	Retained or replacement exposure	Evaluation evidence
Routine operations data entry: process sequence and data relationships.	Automate standard entries; retain human review of sampled transactions and process maps.	New hires complete a short supervised manual set before automation access; incumbents review a rotating sample monthly.	Explain data lineage and achieve at least 90% accuracy on a sampled manual case.
Finance-related entry: links between operational events and finance records.	Automate standard transfers; keep reconciliation and mismatch investigation human-led.	Use anonymised mismatch cases and supervised reconciliation practice.	Correctly identify and resolve at least four of five seeded discrepancies.
Status updating and coordination: end-to-end process awareness.	Automate routine status updates; retain ownership of delays, dependencies and stakeholder communication.	Rotate responsibility for a weekly exception review using real or simulated cases.	Accurate identification of material exceptions with documented escalation rationale.

Task and capability at risk	Allocation decision	Retained or replacement exposure	Evaluation evidence
Approval of system output: independent judgement and challenge.	Require independent assessment for higher-risk or unusual cases before revealing system recommendation.	Record reasoning first, then compare with the system and discuss differences with a supervisor.	Review agreement patterns with audit samples; investigate low override rates rather than treating them automatically as success.

The retrospective analysis would have changed the original engagement in three ways. It would have identified learning-generating tasks before they were reduced, specified minimum exposure by career stage, and established evidence of retained capability. It might also have differentiated incumbents from new hires: experienced employees may need periodic fallback and exception practice, whereas newcomers need a sequenced pathway that begins with foundational process exposure.

7. Operationalising Evaluation

Evaluation needs an owner, a threshold, and a review rhythm. A single metric can mislead. For example, a low override rate may reflect a highly accurate system, passive acceptance or both. Organisations should combine behavioural data with sampled case review and direct demonstration. Table 4 provides starting measures for an SME. Thresholds are deliberately presented as initial governance triggers rather than universal standards; adjust them to task risk, system accuracy, and the consequences of failure.

Table 4. Lightweight SME capability retention scorecard

Indicator	Starting trigger or threshold	Cadence	Owner	Action if below threshold
Manual fallback drill success	At least 90% of critical steps completed accurately within the agreed recovery time.	Six-monthly and after major system change.	Operations manager with process owner.	Targeted practice; repeat drill within one month; review workflow documentation.

Indicator	Starting trigger or threshold	Cadence	Owner	Action if below threshold
Exception-case accuracy	At least 80% initially, rising to 90% for experienced staff on a five-case sample.	Quarterly.	Line manager or subject-matter expert.	Coaching on error patterns; expand supervised exception exposure.
Independent judgement before AI reveal	Reasoning recorded for all sampled high-risk cases; material differences discussed.	Monthly sample.	Line manager.	Review decision criteria and whether AI is being shown too early.
Override rate plus audit sample	No universal target; investigate sustained near-zero override together with missed-error results.	Monthly trend; quarterly review.	Process owner and risk lead.	Sample approvals, test reviewer attention, and recalibrate escalation rules.
Time to competence for new hires	New hire meets defined independent-performance standard within the role-specific target period.	Each onboarding cohort.	Hiring manager with L&D or HR.	Review exposure sequence, coaching access and task progression.
Process explanation and data lineage	The employee can explain key hand-offs, controls, and failure points without system prompts.	At onboarding completion and annually.	Process owner.	Refresher walkthrough followed by reassessment.

7.1 Efficiency resilience trade-off and SME feasibility

Retaining exposure has a cost. Manual drills, sampled reviews and supervised cases use time that could otherwise be devoted to throughput. The relevant comparison, however, is not training time versus zero cost. It is the small recurring cost of readiness against the expected cost of recovery failure, undetected error, weak challenge or a longer new-hire ramp-up. Higher-risk processes justify more frequent practice; low-risk processes may require only sampled review and annual demonstration.

A resource-constrained SME can begin with a minimum viable version of CARE: identify no more than three critical capabilities for each redesigned role; select one representative task or exception for each; conduct a five-case quarterly check; run one

six-monthly fallback exercise; and review new-hire time to competence. The line manager can own the scorecard, with the process owner confirming standards. L&D or HR supports the exposure design but should not become the sole owner of operational capability.

8. Reflection and Limitations

The strongest feature of the SME engagement was its integration of task analysis, technology change, role redesign, coaching and transition planning. This created a clearer line of sight from process change to the future role than a stand-alone system implementation would have done. It also showed that job redesign can create capacity for higher-value coordination and exception handling when routine entry is reduced. What I would do differently is make capability preservation a design criterion at the start. The original engagement asked which tasks created inefficiency and how the role should change. A stronger diagnosis would also ask which tasks developed process understanding, which exceptions employees needed to experience, and what evidence would show that they could still recover when the system was unavailable. Baseline measures of time use, error patterns, exception exposure and manual fallback capability should be collected before implementation. Follow-up measures should then be reviewed after stabilisation, not inferred from system adoption.

The central tension is between efficiency and resilience. Clients understandably seek to reduce manual work, and a consultant can unintentionally reinforce that objective by treating all repetition as waste. Yet some repetition is practice, and some apparently routine tasks reveal the structure of the process. The design challenge is to preserve the learning function without preserving unnecessary volume. CARE makes that judgement explicit, but it does not eliminate the need for operational and managerial trade-offs.

The limitations are substantial. This is a single anonymised case interpreted by a practitioner who participated in the engagement. The underlying project was not designed as research; the retained record does not provide a quantified before-and-after task inventory, systematic participant feedback, capability outcomes or a prospective CARE cycle. The case involved system integration and data-entry

automation rather than an agentic workflow in which an employee mainly approves AI-generated outputs. The agentic-AI implications are therefore conceptual transfer, not empirical findings. The proposed measures and thresholds have not been validated in this organisation. Future work should apply CARE prospectively, collect baseline and follow-up data, include employee and manager perspectives, and compare the throughput cost of retained exposure with continuity, quality and time-to-competence outcomes.

9. Implications for Practice

For L&D, the implication is a shift in timing and scope. The function needs representation before workflows are automated, not only after deployment when users require system instruction. Its contribution is to identify learning-generating work, design replacement exposure and help define evidence that capability remains available. For operations and technology teams, each automation decision should include a capability-impact question alongside speed, cost, accuracy and control. For risk functions, continuity testing should assess whether people can recognise an error, challenge an output and perform critical fallback steps. For HR and workforce planners, entry-level redesign should include an explicit progression path and time-to-competence measure. Singapore's Workforce Development Grant Job Redesign Plus is relevant because it connects workforce consultancy, capability-building initiatives and workforce technology within a broader redesign approach (Skills and Workforce Development Agency, 2026). The immediate practice recommendation is modest: add a capability-impact review to the automation business case and assign ownership for the resulting exposure and evaluation controls. That step turns workplace learning from a remedial response into part of the work architecture.

10. Conclusion

Technology can increase throughput while reducing the work through which future capability is developed. The response is not to preserve low-value work for its own sake, but to identify what capability performance generates and decide deliberately how to retain it. The CARE model makes that sequence visible: map capability-critical tasks,

make allocation choices explicit, retain or replace developmental exposure, and evaluate whether capability remains available. In this brief, CARE is a proposition generated through reflection on one engagement and established literature. Its next test is prospective application with measurable outcomes. Until then, its practical value is as a disciplined question set that helps organisations consider workforce capability alongside automation, rather than after the learning pathway has disappeared.

References

- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775-779.
[https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
- Beane, M. (2019). Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly*, 64(1), 87-123.
<https://doi.org/10.1177/0001839217751692>
- Billett, S. (2001). Learning through work: Workplace affordances and individual engagement. *Journal of Workplace Learning*, 13(5), 209-214.
<https://doi.org/10.1108/EUM0000000005548>
- Brynjolfsson, E., Chandar, B., & Chen, R. (2026). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence (revised August 2026). Stanford Digital Economy Lab.
<https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889-942. <https://doi.org/10.1093/qje/qjae044>
- Endsley, M. R., & Kiris, E. O. (1995). The out-of-the-loop performance problem and level of control in automation. *Human Factors*, 37(2), 381-394.
<https://doi.org/10.1518/001872095779064555>

- Eraut, M. (2004). Informal learning in the workplace. *Studies in Continuing Education*, 26(2), 247-273. <https://doi.org/10.1080/158037042000225245>
- Infocomm Media Development Authority. (2026). Model AI governance framework for agentic AI (Version 1.5). Government of Singapore. <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf>
- Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381-410. <https://doi.org/10.1177/0018720810376055>
- Parker, S. K., & Grote, G. (2022). Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Applied Psychology*, 71(4), 1171-1204. <https://doi.org/10.1111/apps.12241>
- Skills and Workforce Development Agency. (2026). SkillsFuture Workforce Development Grant Job Redesign Plus. <https://www.swda.gov.sg/home/employers-industry-partners/workforce-skills-development/wdg-jr>
- World Economic Forum. (2025). *The future of jobs report 2025*. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>