Journal of Engineering Innovation and Global Impact (JEIGI)
ISSN: 0000-0000
Volume 1, Issue 1 | PP 15-28 | https://jeigionline.com/ |

Submitted : 18-05-2026 | Revised : 11-06-2026 | Accepted : 23-06-2026 | Published : 30-06-2026

Rethinking Job Displacement in the Age of Automation: A Conceptual Model of Workforce Transition and Reskilling

Dr Arpit Jain

Department of Computer Science and Engineering,
Koneru Lakshmaiah Educational Foundation,Vaddeswaram, Vijayawada,
Andhra Pradesh, 522502, India
Emaill: arpitjain@kluniversity.in

Abstract

The debate about automation and jobs has been centred on the question of how many jobs are susceptible to being replaced by machines, and has relied on the task-based estimates of job exposure to machines that range from a factor of ten times between studies. In this paper, the question is challenged as being under-specified. Exposure is an attribute of tasks, whereas displacement is a consequence for workers and both are linked by a transition process that is not modelled by exposure metrics. The fates of workers in similarly exposed jobs vary widely based on where nearby jobs are located, whether their employer redeploys or lets them go, and the institutional and regional infrastructure that surrounds them—workers reskilling is considered a post hoc, undifferentiated supply-side solution. This paper presents the Adaptive Transition Capacity (ATC) model, a multi-level model which reframes displacement as a product of the task exposure on the one hand and the capacity to transition on the other, both at the individual, organisational and institutional-spatial level, in a bounded latency window between onset of exposure and separation. The study is guided by an approach of theory-synthesis that combines the literature on task-based automation, human capital and the relevance of skills, displacement economics and evaluation of active labour market policies. The model results in a typology redeployment/reconversion/stranded proximity/displacement cascade and seven testable propositions, as well as an operationalisation protocol that specifies constructs, indicative measures, and placeholder estimands. The contribution is theoretical - no estimates of the contribution are reported. Employer implications, labour market institutions and design and timing of reskilling interventions are discussed.

Keywords: automation; job displacement; workforce transition; reskilling; skill adjacency; labour market institutions

1. Introduction

There is perhaps no more frequently asked empirical question or answer that is more inconsistent than how much work machines will take. Two years later, a large share of jobs were estimated to be at high risk of computerisation (Frey & Osborne, 2017), while within another two years, the task-level estimates were an order of magnitude smaller (Arntz et al., 2016), and the international estimates were an order of magnitude in between (Nedelkoska & Quintini, 2018). The divide is not so much a data

problem. It is a definitional one because the studies focus on the technical substitutability of the task bundles that they measure and then assume that the measure is a prediction of what will happen to the people who are doing them.

It is not. There is a difference between being technically feasible and being adopted, between being adopted and being separated, between being separated and being displaced (in a welfare-relevant way), that being a separated worker is not the same as being displaced, even if they are separated and they move to comparable work. But what lies between the automatability of a task and the fate of a worker is a process of anticipation, retraining, internal redeployment, job search, geographic movement, and institutional support that can't be measured by exposure metrics, nor by construction be measured. The literature has therefore built up considerable knowledge regarding the supply of automation, and a significantly less amount of knowledge regarding the capacity for transition.

The implications are more than just academic. The framing is the inheritance of policy. While reskilling is presented as the main remedy to automation, it is generally used as an undifferentiated and separated strategy: the displaced worker is trained once they have lost their job, is provided with a training that is often not related to the configuration of the local labour markets and reskilling is supposed to translate into a job through the individual effort of the worker. There is evidence of evaluation of active labour market programmes that this is near to the worst configuration in terms of training effects, which are small in the short term, larger in the medium term, and highly variable by programme type, timing, and context (Card et al., 2018). The displacement literature over the past several decades has well documented that the impact of separation from a long-term job is manifested by long-term losses in earnings, which are not likely to be offset by a short training course. (Jacobson et al., 1993) Because there is no vocabulary in the framing, the field knows that the standard remedy is weak, and has continued to prescribe it.

In this paper, the problem is re-formulated. It itemizes three questions. RQ1: What conceptual framework links task-level exposure to automation to worker-level displacement results? RQ2: What capacities, at what levels, are responsible for the ability of an exposed worker to make a successful transition, a poor transition, or no transition? RQ3: The timing of intervention in relation to the onset of exposure impacts the effectiveness of reskilling.

There are three contributions. First, the paper introduces the Adaptive Transition Capacity (ATC) model that defines displacement as a function of exposure and multi-level transition capacity, instead of exposure alone. Second, it characterizes a four-trajectory typology based on the “displaced or not” versus the “displaced to what” binary, by considering skill adjacency and the capacity to move between skills as a basis for determining skill trajectories.Second, it proposes a 4-typology of skills based on the interaction of skill adjacency and skill trajectory capacity, as a way to define skill trajectories, not on the binary of “displaced or not” but on the new binary of “displaced to what.” Third, it brings the latency window as a temporal entity, as a concept that has more significance than it does as a detail of implementation. There are no estimates or tests of hypotheses in the paper.

2. Literature Review

2.1 The Task-Based Framework and Its Exposure Estimates

The task approach is the intellectual basis of the modern debate, which breaks occupations down to tasks and asks what's machineable. Initial models differentiated routine from non-routine, as well as cognitive from manual tasks, forecasting that codifiable routine tasks would be replaced and complementary tasks would be improved (Autor et al., 2003). A comprehensive part of the occupational polarisation observed in the advanced economies was captured by the framework (Autor & Dorn, 2013; Goos et al., 2014). Its formal extension sees automation as a substitution of labour for capital, and a countervailing reinstatement effect as new tasks appear, suggesting that net employment is a question of the economy, not technology (Acemoglu & Restrepo, 2019). Empirical studies on robot adoption reveal negative local employment and wage effects (Acemoglu & Restrepo, 2020), and

machine learning-specific studies reveal a diffuse impact on wage distribution and few occupations that are fully automatable (Brynjolfsson et al., 2018). The different estimates of exposure result mostly due to the unit of analysis (whole occupations (Frey & Osborne, 2017) versus task bundles within occupations (Arntz et al., 2016; Nedelkoska & Quintini, 2018)).

2.2 What Exposure Estimates Do Not Model

Three omissions matter. First, exposure and adoption are different: exposure happens when the technology is technically feasible, but adoption requires complementarities with investment, and relative prices (Bessen, 2019; Brynjolfsson & McAfee, 2014). Secondly, task reconfiguration in ongoing jobs: Most automation does not eliminate jobs, but reconfigures tasks (Autor, 2015). Third, a focus in this paper is the transition process itself. Exposure is calculated in an occupation but outcomes are received by the worker, and the mapping of the two is not explicitly stated as if there were no friction between the two (Autor et al., 2022).

2.3 Human Capital, Skill Adjacency, and the Structure of Movement

Human capital theory differentiates between general investment and firm-specific investment, and suggests that the more specific the investment, the higher the cost of stopping. (Becker, 1964). Task-based extensions of this kind reveal that human capital can be transferred as long as there is a sufficient overlap between tasks, meaning that a shift between tasks with similar profiles does not reduce (or even eliminate) earnings, while a shift to a task very different from the previous one does. (Gathmann & Schönberg, 2010) Relatedness research extends the learning: skills are relationally structured, and the movement takes place along the paths of relatedness in an economy's skill space (Neffke & Henning, 2013). The implication is on-point and largely underutilized: the question that needs to be asked is not if a worker's occupation is exposed, but where the unexposed adjacent work is, and how far away it is, in terms of skills.

2.4 Displacement, Scarring, and the Limits of Remedial Reskilling

Displacement economics has concluded that the earnings effects of separations from long-term jobs are large and enduring, with the impact being strongest for individuals who have human capital that is not easily transferable from job to job (Jacobson et al., 1993). The small or negative short-term effects of classroom training, and the modest positive long-term effects, are predicted from programme evaluation (Card et al., 2018); job-search assistance is more effective and employer-related training is more effective. Together, both literatures suggest that remedial reskilling is used at the moment when there is lowest capacity once income has ceased, specific capital has been written off, and search urgency has made the long conversion options impossible: but this is where the policy default is.

2.5 Organisational and Institutional Mediation

The employer is basically not involved in the automation discussion (which is odd as they determine if an exposed worker is redeployed or released). Adjustment is likely to take place within the firm's boundary when the firm has internal labour markets, when resources can be easily reconfigured within the firm under technological change, and when the firm has a long investment horizon (Teece, 2007). Downstream are all interconnected with institutional and spatial context conditions that vary widely across the country and widely known to create variation in programme effects: credential portability; the thickness and diversification of the local labour market; the design of active labour market policy (ALMPs); and others.

2.6 Synthesis and Research Gap

Table 1 consolidates the streams. Each is internally strong; none connects exposure to outcome through a specified transition process, and none treats reskilling as a construct with structure timing, adjacency, and sponsorship rather than as a homogeneous treatment.

Table 1. *Synthesis of literature streams and identified gaps*

Stream	Core contribution	Principal limitation	Representative sources
Task-based automation	Exposure decomposition; polarisation; reinstatement effect	Exposure equated with outcome; transition unmodelled	Autor et al. (2003); Acemoglu & Restrepo (2019)
Exposure estimation	Quantified risk shares at occupation and task level	Estimates diverge by unit of analysis; no worker-level process	Frey & Osborne (2017); Arntz et al. (2016)
Human capital and adjacency	Portability depends on task overlap; skill space structures movement	Not integrated with automation exposure	Becker (1964); Gathmann & Schönberg (2010); Neffke & Henning (2013)
Displacement and ALMP	Persistent scarring; heterogeneous, timing-sensitive training effects	Timing treated as implementation, not theory	Jacobson et al. (1993); Card et al. (2018)
Firm and institution	Redeployment capability; institutional variation	Employer largely absent from automation debate	Teece (2007); Autor et al. (2022)

3. Research Methodology

3.1 Research Design

It is a conceptual paper using a method of theory synthesis – the process of merging previously disjoint literatures into a model provided by none of them alone (Jaakkola, 2020). The design is suitable when a field's constructs are well theorised, but the relationships between constructs are still unclear, as is the case in the present field: exposure, adjacency, scarring, programme effects are each well theorised, but none of the relationships between them are in a framework. Adequacy is judged on the basis of the presence of constructs (what), specification of relationships (how), theoretical justification (why) and the limitation of boundary conditions (when, where, who) (Whetten, 1989), and requirement that a conceptual contribution does identifiable work here (differentiation and delineation rather than summary) (MacInnis, 2011).

3.2 Procedure

The development was divided into five stages (Table 2). Theoretical saturation of constructs was the criterion for purposive selection of sources, not exhaustive. The construction of harmonisation was not straightforward, as the word “skill” is used in many different ways in the policy literature, in task research to refer to the skills required for a task, and in the relatedness literature to describe positions in a relation; the model is used in its relational sense, and this is clearly stated. It was assumed that each level sets constraints on the level below it, and cannot be limited to it, and propositions were drawn at each interface.

Journal of Engineering Innovation and Global Impact (JEIGI)
ISSN: 0000-0000
Volume 1, Issue 1 | PP 15-28 | https://jeigionline.com/ |

Table 2. *Stages of model development*

Stage	Activity	Output
1	Problem identification	Mis-specification critique; RQ1–RQ3
2	Literature synthesis across five streams	Gap analysis (Table 1)
3	Construct extraction and harmonisation	Definitions and levels (Table 3)
4	Model synthesis	ATC model (Figure 1); trajectory typology (Figure 2); latency window (Figure 3)
5	Derivation of testable content	Propositions P1–P7 (Table 4); operationalisation protocol (Table 5)

3.3 Operationalisation and Evaluation Strategy

The model is to be empirically tested, and this is outlined in Table 5. Three designs are suggested: (i) linked employer-employee panel analysis, where exposure is measured at the task level and transition capacity is constructed using the skill-space distance and the redeployment experience of the firm; (ii) quasi-experimental evaluation exploiting the variation in timing of the interventions compared to the onset of the exposure, which is the natural test of the latency window; (iii) comparative institutional analysis across regions with different institutional thickness and programme design. Placeholders are used for all estimands.

3.4 Scope and Boundary Conditions

The model covers workers in paid employment in economies with formal labour market institutions and is restricted to displacement caused by automation of tasks, but it is expected that the transition machinery will become more widely applicable. It does not simulate overall employment, macro re-engagement with new tasks, or politics of adoption. No data has been collected and no estimates are made.

4. Results and Findings: The ATC Model

4.1 The Core Reframing

The model's foundational claim is a re-specification. Conventional treatments imply that displacement is a function of exposure. The ATC model proposes instead that the transition outcome is a function of exposure conditioned by transition capacity, itself a multi-level construct, within a bounded latency window:

$$\textit{Transition outcome} = f(\textit{task exposure} \times \textit{transition capacity [individual, organisational, institutional–spatial]} \mid \textit{latency window})$$

Two features follow immediately. Exposure becomes a necessary but insufficient condition it initiates the process rather than determining it which dissolves the paradox of divergent exposure estimates by making them estimates of an input rather than of an outcome. And capacity becomes the variable of interest, which relocates the policy question from how many jobs are at risk to which capacities are missing, at which level, for whom.

Journal of Engineering Innovation and Global Impact (JEIGI)
ISSN: 0000-0000
Volume 1, Issue 1 | PP 15-28 | https://jeigionline.com/ |

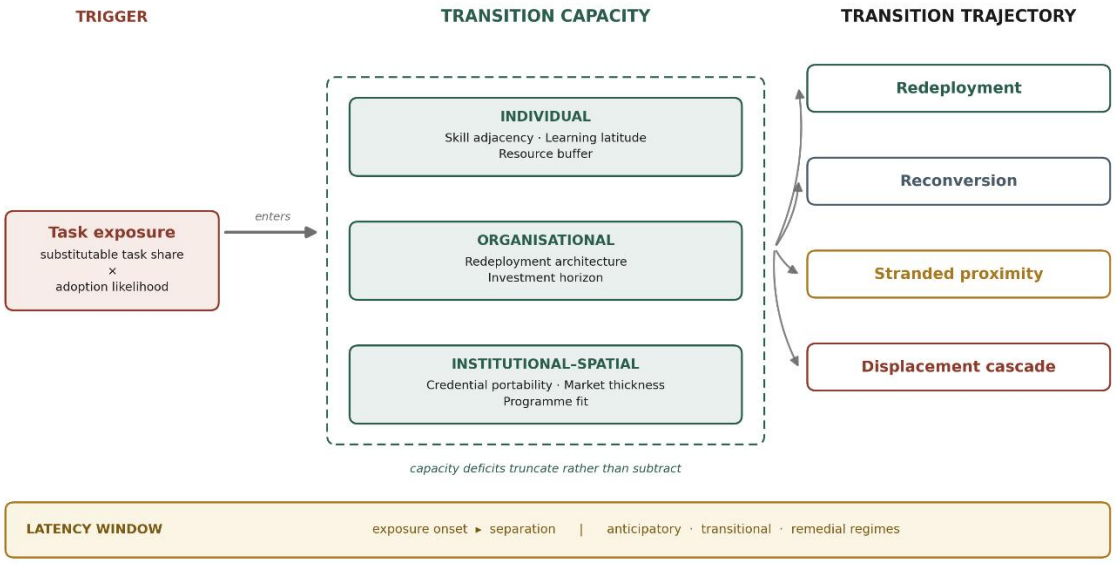

Figure 1. *The Adaptive Transition Capacity (ATC) model: task exposure conditioned by transition capacity at three levels, within a bounded latency window, yielding four transition trajectories.*

4.2 Levels and Constructs

Table 3. *Levels, constructs, and indicative measures in the ATC model*

Level	Construct	Definition	Indicative measure
Trigger	Task exposure	Share of a job's task bundle technically substitutable, weighted by adoption likelihood	Task-level exposure index × sectoral adoption rate
Individual	Skill adjacency	Distance in skill space between the worker's task profile and the nearest viable unexposed destination	Task-overlap distance to feasible destination occupations
Individual	Learning latitude	Capacity to acquire distant skills: prior learning history, formal attainment, time availability	Training participation history; hours available for learning
Individual	Resource buffer	Ability to sustain a non-earning or low-earning conversion period	Months of consumption financeable without earnings
Organisational	Redeployment architecture	Existence of internal mobility paths and the machinery to route workers along them	Internal transfer rate of exposed workers; skills inventory presence

Level	Construct	Definition	Indicative measure
Organisational	Investment horizon	Employer's willingness to fund adjustment it may not capture	Employer-funded training expenditure per exposed worker
Institutional–spatial	Credential portability	Degree to which acquired competence is recognised across employers	Presence of portable, employer-recognised credentials
Institutional–spatial	Market thickness	Density and diversity of employers demanding adjacent skills within commuting range	Count and industrial diversity of adjacent-skill vacancies
Institutional–spatial	Programme fit	Alignment of available reskilling provision with locally demanded adjacent skills	Overlap between programme content and adjacent-vacancy skill demand
Temporal	Latency window	Interval between credible exposure onset and separation	Months from adoption announcement to separation

The levels are not cumulative. Capacity at a higher level replaces deficits at lower levels, and a deficit at any level culls the trajectories that could grow from lower levels, and a thick market with portable credentials can support a worker whose employer is unwilling to redeploy the worker, while a deficit at any level truncates the potential trajectories that could grow from the levels below it, which is why single level interventions are disappointing.

4.3 The Trajectory Typology

Crossing skill adjacency with transition capacity yields four qualitatively distinct trajectories (Figure 2), replacing the displaced/not-displaced binary.

Journal of Engineering Innovation and Global Impact (JEIGI)
ISSN: 0000-0000
Volume 1, Issue 1 | PP 15-28 | https://jeigionline.com/ |

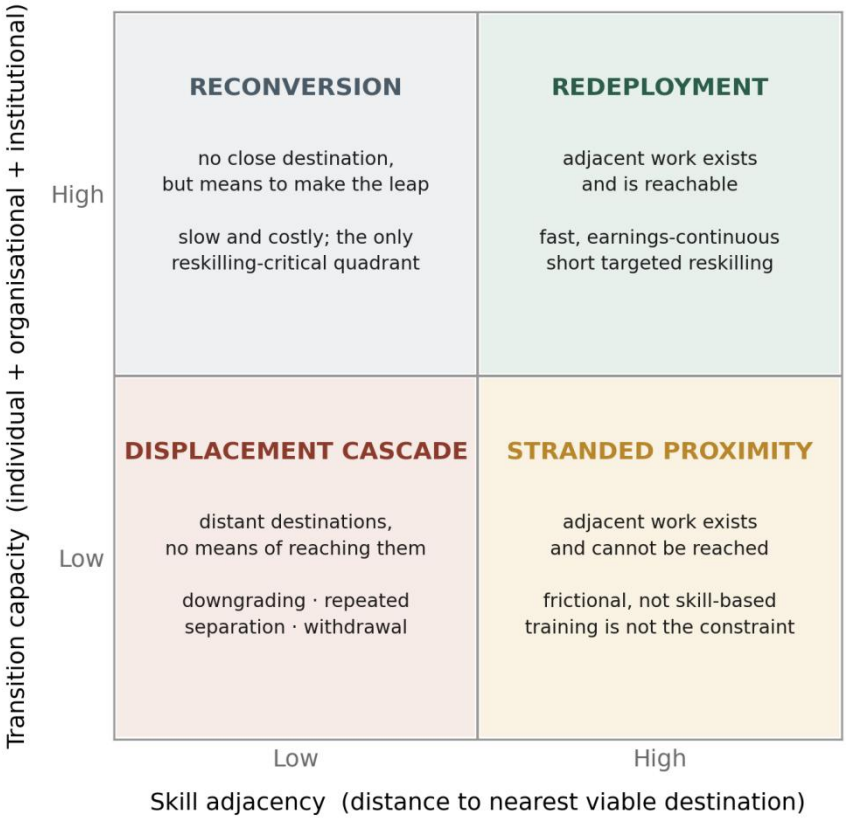

Figure 2. Transition trajectory typology. Quadrant positions are conceptual, not estimated.

Figure 2. *Transition trajectory typology: skill adjacency crossed with transition capacity.* Quadrant positions are conceptual, not estimated.

Redeployment (high adjacency, high capacity). There is adjacent work and the worker can access it, usually within the firm or occupation. Adjustment is rapid, continuity of earnings is high, and reskilling is short and focused. It's the assumption that is made in rosy reports, and it's a minority of cases.

Higher capacity, lower adjacency – reconversion. There is no close destination, but via resources, time and institutional support is the opportunity to make a true occupational step. Slow and costly to adjust; can be positive. It is the only quadrant in which there is a need for significant reskilling and where that reskilling is possible, and it is also the quadrant that is least capable of being served by remedial and post-separation programmes, since separation leads to the depletion of the resource buffer that is needed for reconversion.

Stranded proximity (high adjacency, low capacity). The case most under-theorised and the most common, as this paper suggests: appropriate neighbouring work is available nearby, but the worker cannot get to that job because qualifications are not transferable, or because the employer won't relocate, or because the conversion will take a few months and won't be paid for, or because the job is thirty kilometres away and would be a difficult commute. This failure is not a skill failure but a frictional failure and is not discernible by any framework that looks at outcomes as a function of exposure. But it's an uncomfortable policy implication for the reskilling consensus: training is not the constraint, and training is not going to solve the constraint.

Displacement cascade (low adjacent, low capacity). Long distance and no transportation. The observed trajectories include a downward trend into lower income work, multiple separations, and withdrawal, which are documented by displacement economics (Jacobson et al., 1993). As the conventional notion

Journal of Engineering Innovation and Global Impact (JEIGI)
ISSN: 0000-0000
Volume 1, Issue 1 | PP 15-28 | https://jeigionline.com/ |

of reskilling is being rendered almost irrelevant by the lack of destination and means, it also seems to be irrelevant when it comes to reskilling as a process.When it comes to reskilling as a process, it appears to be equally irrelevant, as the conventional notion of reskilling is already close to being irrelevant.

4.4 The Latency Window and the Timing of Intervention

A surprise exposure is rare. There is an interval between the exposure becomes credible and the point of separation, which is the latency window, and is preceded by announcement and procurement phases for adoption. The window is the temporal construct of the model, and it has the most important practical implication of the paper: the value that this reskilling intervention is expected to bring varies depending on its location in the window (Figure 3).

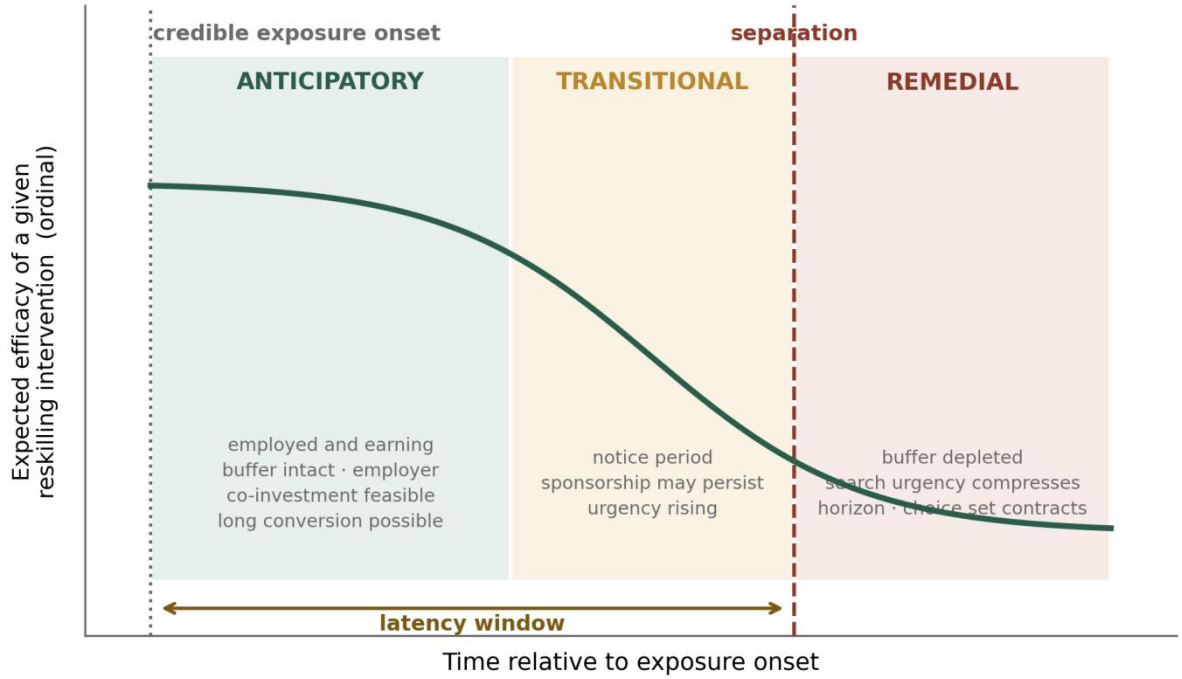

Figure 3. The latency window and intervention regimes. Schematic and ordinal: the curve represents theorised expected efficacy, not estimated values.

Figure 3. *The latency window and intervention regimes.* Schematic and ordinal: the curve represents theorised expected efficacy, not estimated values.

The following three regimes are distinguished. Anticipatory intervention takes place during the worker's employment and income: Resource buffer is still intact, the employer can co-invest, and long conversions are possible. Around separation, there's a mix of urgency and availability of notice periods and employer sponsorship, with transitional intervention. Remedial intervention happens after separation, when the buffer is depleted, search urgency tightens up horizons, and the set of feasible choices narrows down to what can be reached quickly exactly the time when classroom training works the least well (Card et al., 2018). The model therefore suggests that much of the difference in reskilling effectiveness between programmes is actually due to programme timing and that, if this is the case, the literature on evaluating reskilling programmes has been comparing the wrong things.

4.5 Propositions

Table 4. *Propositions derived from the ATC model*

ID	Proposition
P1	Task exposure predicts transition outcome only weakly once transition capacity is specified; the exposure–outcome relation is fully moderated at the three capacity levels.
P2	Skill adjacency dominates exposure intensity in predicting earnings continuity after separation.
P3	Organisational redeployment architecture substitutes for individual capacity: among workers of equivalent adjacency and individual capacity, those in firms with internal mobility paths exhibit better trajectories.
P4	Institutional–spatial capacity conditions the returns to individual reskilling: identical training yields divergent outcomes across thick and thin local labour markets.
P5	The expected efficacy of a given reskilling intervention declines monotonically across the anticipatory, transitional, and remedial regimes.
P6	Stranded proximity is frictional rather than skill-based: in this quadrant, training investment yields negligible improvement relative to credential portability or mobility support.
P7	Capacity deficits are truncating rather than additive: a deficit at any level constrains the trajectory set reachable from capacity at other levels.

4.6 Operationalisation Protocol

Table 5. *Constructs, measurement approaches, and placeholder estimands*

Construct	Measurement approach	Estimand (placeholder)
Exposure–outcome moderation (P1)	Panel interaction of exposure index with capacity indices	βint = [to be estimated]
Adjacency effect (P2)	Task-overlap distance regressed on post-separation earnings ratio	γadj = [to be estimated]
Redeployment substitution (P3)	Matched comparison across firms differing in internal mobility	$\Delta redeploy$ = [to be estimated]
Spatial moderation (P4)	Cross-region comparison of training returns by market thickness	$\theta thick$ = [to be estimated]
Timing gradient (P5)	Quasi-experimental variation in intervention timing within the window	$\lambda 1 > \lambda 2 > \lambda 3$ = [to be estimated]
Frictional stranding (P6)	Decomposition of non-transition into skill vs. friction components	$\varphi fric$ = [to be estimated]

5. Discussion

5.1 Theoretical Implications

The model will simulate three moves. It takes exposure out of the equation and thus makes the often-conflicting exposure literature consistent by re-classifying what they measure as an input rather than an outcome of the labour market. It shifts the focus from explaining the weight to explaining capacity, bringing the relational structure of skill space into a debate that has mostly considered skills as levels on a ladder (Gathmann & Schönberg, 2010; Neffke & Henning, 2013). And it is not a timing detail, it's a moderator in your design with a predicted sign, making it a falsifiable statement: the latency window is not a design parameter but a moderator with a predicted sign.

5.2 Practical Implications

The model changes the paradigm for the redeployment as an asset for the employer rather than a cost. When P2 is dominant (i.e., adjacency) and P3 is the substitute for individual capacity (i.e., redeployment), the firm that has an inventory of skills and internal mobility paths turns an exposure event into a redeployment event instead of a severance event and retains the specific capital that separation writes off. From the viewpoint of workers, the pertinent questions are not whether their job is going to be taken over by automation, where is the closest place to which they could be redirected and, what would it cost to get there. A stranded proximity quadrant for programme designers points to the mismatches for which more training is insufficient and for which credential recognition, support for credential mobility and employer brokerage may be possible.

5.3 Policy Implications

The model's policy thesis is that the reskilling consensus is well motivated, and is ill-timed. Intervention is focused in areas of the model's low separation after no sponsor and no local adjacency structure. The idea is that the latent period is predicted to be dominated by anticipatory intervention activities, which would mean the mechanisms of the institution barely exist: disclosure obligations in the process of adoption, portable training entitlements, and co-investment instruments, such that employer-funded conversion makes sense even if the employer might not keep the worker. All are disputed on the grounds of cost, incidence, and feasibility, and the model does not address these issues, but states why timing is the focus of the argument.

5.4 Limitations

Conceptual, untested model, propositions are argued, not demonstrated. For skill adjacency, a task level data set is needed and credible distance metrics are required; for latency window, an observable adoption timing is needed, and firms are not obligated to disclose this timing. The synthesis is deliberate and there may be selection bias. The model is restrictive because it excludes informal labour markets, self-employment and platform work, and the institutional features rely on formal credentialing systems which many jobs in the world do not have. Finally, the model assumes labour demand is exogenous, given adoption, which implies nothing about whether new tasks emerge – the focal issue in the reinstatement literature (Acemoglu & Restrepo, 2019) and a real constraint on any statement about aggregate adjustment.

5.5 Future Research

The propositions are followed by priority tests. P1 and P2 are measurable now in jurisdictions with linked employer–employee data in which task-level measures are matched with separations. P5 is the most valuable test and the most challenging, because of the need to vary the timing of interventions relative to the onset of exposures, which is plausibly available from firm-level adoption announcements and programme records. P6 needs to be broken down into skill and friction parts, to determine the extent to which the reskilling consensus is misallocating. In addition to proposition

testing, the model opens for inclusion of informal and platform labour, includes the reinstatement literature to endogenise destinations, and tests the possibility of learning transition capacity through repeated exposure events.

6. Conclusion

The automation debate has been honing an estimate of the wrong amount for 20 years. Task exposure doesn't necessarily tell us what happens to the people substituted, since it doesn't include a transition process. It has contributed to the Adaptive Transition Capacity model of displacement as a product of exposure and multi-level transition capacity within a bounded latency window, resulting in four qualitatively different trajectories of re-deployment, reconversion, stranded proximity and displacement cascade. The core arguments of the model are: Adjacency is more important than exposure intensity, Capacities deficits do not diminish exposure, and the moment of reskilling versus separation may account for more of the efficacy measured than the content of reskilling. All claims come with a disclaimer. The uncomfortable message for policy is that the dominant way to respond to automation is done at the time when those who face it are least equipped to do it; and rethinking displacement requires a rethinking of when we do it, rather than how.

Declarations

Conflict of Interest
The authors declare no conflicts of interest related to this study.

Funding Statement
This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Ethical Approval
This study did not involve any human or animal subjects and, therefore, did not require ethical approval.

References

Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3–30. https://doi.org/10.1257/jep.33.2.3

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. Journal of Political Economy, 128(6), 2188–2244. https://doi.org/10.1086/705716

Arntz, M., Gregory, T., & Zierahn, U. (2016). The risk of automation for jobs in OECD countries: A comparative analysis (OECD Social, Employment and Migration Working Papers No. 189). OECD Publishing. https://doi.org/10.1787/5jlz9h56dvq7-en

Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. Journal of Economic Perspectives, 29(3), 3–30. https://doi.org/10.1257/jep.29.3.3

Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. American Economic Review, 103(5), 1553–1597. https://doi.org/10.1257/aer.103.5.1553

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. The Quarterly Journal of Economics, 118(4), 1279–1333. https://doi.org/10.1162/003355303322552801

Autor, D., Mindell, D. A., & Reynolds, E. B. (2022). The work of the future: Building better jobs in an age of intelligent machines. MIT Press.

Becker, G. S. (1964). Human capital: A theoretical and empirical analysis, with special reference to education. National Bureau of Economic Research.

Bessen, J. (2019). Automation and jobs: When technology boosts employment. Economic Policy, 34(100), 589–626. https://doi.org/10.1093/epolic/eiaa001

Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. W. W. Norton.

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? AEA Papers and Proceedings, 108, 43–47. https://doi.org/10.1257/pandp.20181019

Card, D., Kluve, J., & Weber, A. (2018). What works? A meta analysis of recent active labor market program evaluations. Journal of the European Economic Association, 16(3), 894–931. https://doi.org/10.1093/jeea/jvx028

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change, 114, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

Gathmann, C., & Schönberg, U. (2010). How general is human capital? A task-based approach. Journal of Labor Economics, 28(1), 1–49. https://doi.org/10.1086/649786

Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. American Economic Review, 104(8), 2509–2526. https://doi.org/10.1257/aer.104.8.2509

Jaakkola, E. (2020). Designing conceptual articles: Four approaches. AMS Review, 10(1–2), 18–26. https://doi.org/10.1007/s13162-020-00161-0

Jacobson, L. S., LaLonde, R. J., & Sullivan, D. G. (1993). Earnings losses of displaced workers. American Economic Review, 83(4), 685–709.

MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. Journal of Marketing, 75(4), 136–154. https://doi.org/10.1509/jmkg.75.4.136

Nedelkoska, L., & Quintini, G. (2018). Automation, skills use and training (OECD Social, Employment and Migration Working Papers No. 202). OECD Publishing. https://doi.org/10.1787/2e2f4eea-en

Neffke, F., & Henning, M. (2013). Skill relatedness and firm diversification. Strategic Management Journal, 34(3), 297–316. https://doi.org/10.1002/smj.2014

Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. Strategic Management Journal, 28(13), 1319–1350. https://doi.org/10.1002/smj.640

Whetten, D. A. (1989). What constitutes a theoretical contribution? Academy of Management Review, 14(4), 490–495. https://doi.org/10.5465/amr.1989.4308371