

Who Pays for the Journey?

Designing the transition to an AI-adopting economy, from the inside of a small company

Dr Jussi Kajala – CEO and co-founder, 3DBear Oy, Helsinki. Doctor of Science in Technology (Aalto University), with a doctoral thesis on optimisation by singular value decomposition, the linear algebra that underpins modern machine learning; M.Sci. in Physics, University of Cambridge. Before founding 3DBear he was Senior Adviser at Tekes, the Finnish innovation funding agency, where he led the national 3D-printing programme, ran a ten-person analytics team, and was named Tekes Employee of the Year in 2015, its youngest ever recipient. 3DBear's Simuna platform is used by around half of Finland's vocational schools. He is active in the Finnish startup community's policy work and has been on the boards of Mobie Oy and Seppo.io. He grew up in Juankoski, a North Savonian town built around a board mill that passed through successive owners before closing; the municipality lost its independence to Kuopio in 2017.

Glyn J. Eggar – Actuary and mathematician, Linlithgow, Scotland. He read mathematics at the University of Cambridge and has spent his career as an actuary, latterly through his own consultancy. He now works across several concurrent engagements rather than for a single employer, which makes the subject of this paper also his own situation. An actuary prices risks that have not happened yet, and that discipline shapes three arguments this paper leans on hardest: that what matters is the rate at which people leave a job market rather than the size of the payment they get when they do; that a working life should be assessed over its whole span rather than year by year; and that the question worth asking is who carries the risk, not who caused it. He lives in a country which even now still remembers the impact of the coal mines closing in the 1980s and what that did to the towns built around them.

Mitchell Weisburgh – Entrepreneur, investor and author, United States. Founded Personal Computer Learning Centers of America in 1981 and grew it to more than 130 employees teaching computing to working adults; co-founded Academic Business Advisors in 2005, advising organisations building and selling into education, and later co-founded the non-profits Games4Ed and Edchat Interactive. Author of the *MindShifting* series and a forthcoming book on complexity, and he writes on inequality and institutional capture. Four decades of building and financing education businesses inform his reading of who captures the returns from a new technology.

July 2026

Table of Contents

Abstract

Part 1 – The situation

- 1.1 The transition is the problem, not the end state
- 1.2 What the economists now agree on
- 1.3 What is actually measured, and what is not
- 1.4 Where the gains are going
- 1.5 Three countries, three different exposures

Part 2 – Our analysis

- 2.1 The half-time question
- 2.2 Why goodwill cannot deliver it
- 2.3 Specialisation is efficient and fragile
- 2.4 What decays during unemployment is not skill
- 2.5 AI removed the rungs that made experts
- 2.6 Nobody is entitled to the whole surplus
- 2.7 Money and a placement are not a life
- 2.8 Proposed solutions
- 2.9 What is missing from all of it

Part 3 – Proposals for acting in the age of AI

- For workers
- For employers
- For policymakers

Part 4 – The recommendations in brief

Closing

References

Abstract

The argument about whether artificial intelligence will reshape work is effectively over. In July 2026 more than 250 economists and AI researchers, including several who had spent years disputing displacement forecasts, signed a statement calling for urgent preparation [1]. The discussion about how it will happen has barely started, and the sharpest unanswered question inside it is who pays for the adjustment. This paper offers three things: the case that the transition, not the end state, is where the harm lands; the reason firms will not share the gains voluntarily, which is competitive rather than moral; and a set of concrete proposals for how to act in the age of AI, addressed to workers, employers and policymakers.

We also benchmark against existing solutions which could work in the AI era, for example Sweden's transition councils have run since 1974 [2], because a proposal that misdescribes what already exists is not worth acting on. Finally, when we make proposals, every proposal is specified concretely enough that a reader can say what it would cost, and who would carry that cost.

Part 1 – The situation

1.1 The transition is the problem, not the end state

Almost all public argument about AI and work is argument about the end state: how many jobs go, and whether we arrive at broad abundance or widespread hardship. End-state speculation is cheap, because nobody can be shown wrong for a decade. The journey towards it is where the mistakes get made and where people end up with a bad deal. The debate is not new – the labour economics of technological transition goes back through deindustrialisation debates that long precede AI – but treating it as the primary question rather than a hedge on a forecast is rare in current discourse.

A society where AI and robotics do half the work across the economy, and people work half as much for the same living standard, is one acceptable end state among several. A society where the adjustment is pushed onto individuals over fifteen years while the gains turn into equity immediately is not, even if it arrives at the same place.

The environmental costs of AI – energy, water, land and emissions, and who pays for them – are a real externality and a separate argument. This paper is about the employment transition, and touches them only where they intersect with it.

The British pit closures earn their place as the analogy because the pits were going to close on any plausible energy-policy path; the catastrophe was the absence of a transition mechanism, not the closure. Coal mining employed around 200,000 people in Britain when the 1984 strike began, concentrated in towns that had no other industry, and the closures that followed removed that employment within a decade. Forty years

on, Sheffield Hallam's *State of the Coalfields 2024* finds those communities still showing median hourly earnings 6–7% below the national average [3]. Beatty and Fothergill document how the damage stayed invisible: over 90,000 displaced miners were diverted to incapacity benefits, so claimant unemployment in the coalfields was no higher than when the pits were open [4]. This was an example of an unmanaged transition, not proof that all transitions must fail.

1.2 What the economists now agree on

On 13 July 2026 Stanford's Digital Economy Lab published a four-sentence statement, "We Must Act Now", organised by Erik Brynjolfsson, Ajay Agrawal, Anton Korinek and Tom Cunningham and signed by more than 250 economists and AI researchers including sixteen Nobel laureates [1]. It states the transformation "could ... be larger than the Industrial Revolution, but unfolding over a vastly shorter time frame," warning that policy systems built for past shocks may not adapt in time, and taking no position on the magnitude of job loss. Daron Acemoglu and Simon Johnson signed despite having spent years pushing back on displacement forecasts, naming by omission the gap this paper is written into: there is no mechanism for who pays for the adjustment.

When a firm automates part of a job, someone has to fund whatever comes next for the person who did it – the months or years without income, the retraining, the search for different work, the lower wage in the new role. Four parties could carry that cost, and it is worth naming them, because "the firm" is a construct that hides which of them actually pays. The worker and their household pay first, out of savings, out of a lower future salary, and out of the years it takes to recover. Society pays next, through unemployment benefit, through the public services a displaced household draws on, and through the income tax and social contributions that person no longer remits. The state is not straightforwardly a loser here, and it is worth being precise about why. A firm that automates books a higher profit, and that profit is taxed; if the saving is passed to shareholders as dividends, those are taxed again in the shareholders' hands. The state therefore recovers part of what it loses on the employment side. Whether it comes out ahead or behind is an empirical question with a clear form: the state loses when the corporation and dividend tax it gains is smaller than the benefit payments it must now make plus the income tax and social contributions it no longer collects. Two things make the second quantity the larger one in most cases. Labour is taxed more heavily than capital in every country discussed in this paper – an employed worker generates income tax and employer and employee social contributions on the whole wage, where the same money routed through profit attracts corporation tax on the margin and dividend tax only when distributed. And the receipts land in different places: corporation tax is payable wherever the firm books its profit, which for a multinational need not be the country whose benefit system is now supporting the displaced worker. The state is therefore the payer of last resort without being the party that made the decision, and it is a partial rather than a total loser – but partial is enough for the argument, because it is still carrying a cost it did not choose. The shareholders and the executives who chose the automation keep the saving net of the tax on it; they are the change agents, and

they are the only one of the four groups whose position improves. Every proposal in this paper is an answer to the same question: out of which of those pockets, on what trigger, and how much.

1.3 What is actually measured, and what is not

Three camps organise the current debate, and Carnegie Endowment's April 2026 taxonomy by Teddy Tawil is the most useful map of them [5]. The **Alarmed** expect large, fast displacement; they are led by Dario Amodei, whose May 2025 warning that AI could eliminate half of entry-level white-collar jobs within five years became the reference point everyone else argues with. The **Patient** expect the same technology to arrive far more slowly than the forecasts imply, because integrating it into real organisations takes years. The **Excited** expect it to raise the value of ordinary work rather than destroy it; their strongest case is David Autor's, that AI used as decision support can extend the reach of non-elite expertise and so reverse forty years of computer-driven wage polarisation, conditional on the investments and regulation he calls for [6]. During 2026 the Alarmed and the Patient partly converged on the urgency of policy while still disagreeing about magnitude and speed. The live question has moved from whether to how fast, how large, and who pays.

The productivity effect is not yet visible at scale. A survey of some 6,000 executives across four advanced economies, run by the Federal Reserve Bank of Atlanta, one of the twelve regional banks of the US central bank, with other central banks, found 89% reporting no measurable AI effect on labour productivity over three years, and forecasting only a 1.4% gain over the next three [7].

The employment effect is real but narrow. Brynjolfsson, Chandar and Chen's analysis of payroll data on 4.6 million workers finds a roughly 16% relative employment decline for 22–25-year-olds in the most exposed occupations, arriving through reduced hiring rather than through firing, with no comparable decline for experienced workers [8]. Entry-level postings across knowledge fields fell around 35% between January 2023 and late 2025 [9], yet the occupations themselves grew: since ChatGPT's launch the US added roughly three million white-collar jobs, with 7% more software developers and 21% more paralegals than in 2022. The bottom of the ladder is thinning while the occupations are not shrinking. The measured evidence is concentrated in white-collar work because that is where the payroll data is; the mechanism is not.

The systems are also far from uniformly capable. Dell'Acqua and colleagues' pre-registered experiment with 758 BCG consultants found GPT-4-assisted consultants completing 12.2% more tasks 25.1% faster on work the model handled well, and 19 percentage points *less* likely to reach a correct solution on work just outside that range [10]. The boundary between the two is invisible from inside a session, which is why the paper refers to it as a jagged frontier. METR's 2025 randomised study of experienced open-source developers measured them 19% *slower* with AI tools while they estimated they had been 20% faster [11]; its 2026 redesign still returns negative, though weak, evidence [12].

And the price of the technology is not settled. The frontier labs are loss-making and are funding capability growth from capital rather than from revenue, so today's inference prices are not a stable basis on which to book a permanent cost saving. If prices rise toward cost recovery, the saving that any obligation would attach to shrinks with them. That cuts both ways for the argument here: a duty sized to the saving is self-limiting by construction, because a firm that saved little owes little.

The institutional forecasts are built around these findings. The IMF's April 2025 baseline of 40% of global jobs exposed, and 60% in advanced economies, is carried into a January 2026 follow-up that shifts emphasis toward creation as well as displacement [13]. The OECD's 2026 *Employment Outlook* is the best empirical anchor: adjustment happens through unemployment spells and through openings for newcomers, not through retraining incumbents, and displaced workers carry lasting scars [14]. The ILO and NASK's joint exposure index is the most conservative estimate, putting only 3.3% of global employment in the highest-exposure category and finding, across 135 countries, transformation rather than replacement [15].

At EU level three regulatory tracks run at once. The AI Act's high-risk employment provisions moved from 2 August 2026 to 2 December 2027 under the Digital AI Omnibus agreed on 7 May 2026 [16]. The Platform Work Directive restricts algorithmic management and creates a rebuttable presumption of employment, and its 2 December 2026 transposition deadline is running badly behind, with only four member states holding an in-force rule and eighteen not started [17]. The Quality Jobs Act covers algorithmic management and AI at work, and its consultation closes on 28 September 2026 [18].

1.4 Where the gains are going

Three groups capture the surplus: those who own the technology, large firms deploying it at scale, and their shareholders. In the United States the top 1% share of household net worth rose from roughly 23% in 1989 to 30.9% in 2024, reaching a record 31.7% in the third quarter of 2025 [19]. In the euro area the top 10% share rose from 54.0% in 2009 to 57.2% in 2024 [20]. Income has moved the same way: the Congressional Budget Office's January 2026 series shows the top 1%'s share of pre-tax income roughly doubling, from 9% to 18%, between 1979 and 2022, while the lowest quintile's fell from 5% to 4% [21].

The loose version of this argument – that everyone outside the top decile is losing ground – is not true of the recent period, and the case for acting does not need it. The bottom 50% wealth share was about 2.5% in early 2025, above its post-2008 trough, and real hourly wages at the 10th percentile were back above their 2021 level by mid-2024, an unexpected compression documented by Autor, Dube and McGrew and driven by a tight low-wage labour market and minimum-wage increases rather than by AI [22]. What is happening is narrower and harder to argue with: the very top is pulling away at an accelerating rate.

“Shareholders” needs splitting. Shareholders of *AI-supplying* firms are clearly winning. Shareholders in general are not obviously winning at all, given that 89% of firms report no measurable productivity effect. The US labour share of income fell to 53.8% in the third quarter of 2025, its lowest since the series began in 1947, but this is not yet an AI story: the New York Fed’s June 2026 analysis attributes the post-COVID decline to margin and market-power dynamics without naming AI [23]. Treat it as a forward risk, not a measured fact.

1.5 Three countries, three different exposures

The authors work in Finland, Scotland and the United States, and the three make a useful comparison, because each has a different amount of protective machinery in place and each therefore fails in a different way.

Finland has the most protection and the least visible effects so far. A 2026 study by ETLA, the research institute of the Finnish business community, finds AI exposure has *not* displaced young workers: employment among young people in high-exposure occupations tracks almost identically with low-exposure occupations, AI users report higher work engagement and no elevated fear of job loss, and around 40% of employed Finns already use generative AI at work [24]. ETLA attributes the buffering to the Nordic labour-market model and strong dismissal protection. For Finland, this is the window in which to build the mechanisms, because you cannot build them once the displacement has arrived. A second ETLA finding cuts against complacency: AI’s Finnish impact so far lands most on the elite and knowledge workers, and the less-educated have more to lose than to gain [25].

Scotland shows the problem of a system that would catch people if it had the money. The support exists in principle and is delivered largely through local councils, which run the retraining, the employability services and the social care that a displaced household draws on. Scottish councils face a collective structural gap of roughly £529 million for 2026–27, with Audit Scotland warning that services will have to be stopped or redesigned [26]. The 2026–27 settlement rose 2.9% in real terms, so the accurate claim is structural rather than an annual cut: demand is growing faster than funding. Any proposal that routes displaced workers towards council-provided support is proposing it to bodies currently being told to contract.

Britain also illustrates a subtler failure, in how the benefits system treats part-time and irregular work. Universal Credit withdraws support as earnings rise, at a taper of 55%: a claimant keeps 45p of each extra pound. With income tax and National Insurance on top, effective marginal deduction rates exceed 70% for around 9% of Universal Credit-eligible workers [27]. Consider a single parent who takes a weekend job paying £200 a month. The taper leaves her about £90; tax and National Insurance cut that to roughly £70, and she weighs whether that is worth the childcare and the lost Saturdays. Often it is not, and declining is rational: the design produced that choice, not her character. This is a real disincentive, and it is also an improvement on what came before – under the pre-2013 legacy system, 26% faced deduction rates above 70%, and


Income Support’s 100% withdrawal was a genuine cliff edge [28]. The trap has been reduced, not solved, and it matters here because everything this paper says about holding more than one source of income depends on part-time work being worth taking.

The United States has notice and almost nothing else. The WARN Act requires 60 days’ written notice for plant closings and mass redundancies at firms of 100 or more employees, and that is the whole of it: there is no proactive enforcement, the remedy is a private civil suit for back pay capped at the 60-day violation period, and no transition is funded at any point [29]. Its reach is narrower than it sounds. It is triggered by closures and large-scale events at a single site, so it does nothing for job losses that arrive as a steady trickle across an organisation, which is the shape AI displacement is currently taking. Underneath it, most American employees are employed at will and can be dismissed at any time for any reason that is not specifically unlawful. The 60-day notice is therefore a narrow exception to a default of no protection, rather than a floor beneath it. State legislatures have begun to move: New York’s proposed SHIELD Act would require a 90-day transition period during which affected workers must be offered continued employment, equivalent wages, or employer-funded retraining, and California’s proposed SB 951 would require 90 days’ advance notice where AI or automation causes 25 or more job losses at employers with more than 100 staff [30]. Both are notice-plus-offer duties. Neither scales what the employer owes to what the employer saved, which is the gap this paper is about.

Read together, the three tell one story. Finland has machinery and a closing window. Scotland has machinery without funding, and a benefits taper that penalises the flexible work a transition requires. The United States has a notice period and a private lawsuit. None of the three has an instrument that connects what a firm gains from automating a role to what it owes the person who held it.

Three countries, three ways of being unready

Where the authors work, and what each system does when a role is automated



Three countries, three different exposures

Figure 1. Finland, Scotland and the United States: different protection, the same gap.

Part 2 – Our analysis

2.1 The half-time question

Suppose AI does half your job. In theory everyone could work half the hours at the same pay and the firm's profit would be unchanged. In practice firms lay people off to raise profit, and frame that as a duty to shareholders. So what does a company owe employees for the transition from a low-AI to a high-AI working life – a **choice** rather than a technological inevitability? As an accounting claim the half-time question fails four ways, and it is better to concede them here than leave them for a critic.

It assumes fixed output. If AI halves the labour needed per unit of output, a profit-maximising firm often produces more rather than employing fewer. A translation agency that halves its cost per word need not let half its translators go: it can halve its prices, win twice the volume, and keep everyone busy. Cheaper work can mean more work – the Jevons paradox, named for the nineteenth-century observation that more efficient steam engines increased rather than reduced coal consumption. Amodeli invoked it in May 2026 for exactly this reason. **It assumes AI is free:** integration, verification and the wage premium for people who can supervise the systems absorb much of the surplus before any saving reaches the bottom line. **It assumes the augmented worker is the same worker.** Brynjolfsson, Li and Raymond, studying a rollout to 5,179 customer-support agents, found productivity up 14% overall but 34% for novices and roughly zero for experienced agents [31]. The half-time firm needs a smaller, differently skilled group, a change that headcount arithmetic conceals. **The premise is empirically premature:** nobody is saving half their labour cost at scale.

What survives is the question itself: if a firm captures a saving by automating a role, something is owed to the person whose role it was, and it should track the size of the saving.

2.2 Why goodwill cannot deliver it


One of our main points is that a firm cannot do this on its own, however much it wants to.

A firm that keeps everyone at full pay for half the hours is fine in a monopoly and dead in a competitive market, because a rival passes the saving into the price. If a behaviour is competitively suicidal, appeals to responsibility will not produce it. Only a rule binding all competitors simultaneously will.

Consider two parcel-sorting centres in the same city. One automates its sorting, cuts everyone's hours by a fifth, maintains full pay, and holds its prices where they were. The other automates, lets a fifth of its staff go, and passes the saving into lower prices.

Within a quarter the first centre is losing contracts. Its owner would like to keep the jobs, but keeping the jobs means going out of business.

The strongest objection to intervening is that shareholders bore the risk and should therefore keep the return. Under current rules a firm's productivity gains are supposed to go to maximising returns, not to keeping staff the firm no longer strictly needs; a manager who does otherwise is failing at the job as it is currently defined. But that is a fact about the rules, and rules can be written differently. A duty binding every firm in a market is priced into that market at once, as a cost of doing business, with its incidence shared between shareholders, consumers and workers as the market determines, which is how every other mandatory employment cost already works, from the pension contribution to holiday pay.



Why the decent employer cannot act alone

Figure 2. Two firms in the same market, and why goodwill cannot produce the transition.

Corporate social responsibility fails for the same reason. Alexander's study of apparel-sector living-wage pledges found "seriously lacking" evidence that they raised wages, with firms outsourcing the commitment while leaving intact the practices that produced low pay [32]. Voluntarism can work: Ingka Group reskilled about 8,500 call-centre staff into remote interior-design advisers, and its 2026 redundancies fell elsewhere, which is evidence that reskilling can pay for itself commercially [33]. But it is not reliable. IBM's chief executive promised more graduate hiring on 5 November 2025, then cut thousands of roles in AI-linked restructuring days later [34]. Europe moved from voluntary CSR to mandatory due diligence because voluntarism did not work, and in February 2026 the Omnibus amendment cut the sustainability reporting regime's scope by roughly 90% [35]. The clearest test is in training. Germany's dual apprenticeship

system, the world's most admired, is voluntary for most firms, and the number of firms offering apprenticeships has fallen [36]. Denmark's scheme obliges every employer to contribute whether they train or not, and has been evaluated and found to have positive employment effects [37]. Voluntary provision of a public good erodes under competitive pressure even where the reputational reward is real.

So what could be done instead? The lesson of the British coalfields is not that closures should have been prevented. It is that leaving the adjustment entirely to the market produced forty years of measurable damage in places where the market simply had nothing else to offer. Pure market rules do not work here, for a reason that is structural rather than sentimental: no individual firm can carry the cost of an orderly transition and survive competition, no individual worker can insure against a shock that hits their whole occupation at once, and no local labour market can absorb a displacement it did not cause. What is missing is a set of mechanisms that move workers into new roles quickly and at somebody's expense other than their own – which means deciding, in advance and in law, whose expense that is.

2.3 Specialisation is efficient and fragile

Nobody should have only one job, true before AI: for example, a specialist blue-collar paper-machine operator might have no transferable skills, so when the mill closes there is nothing to move to. This follows Lauri Järvillehto, the Aalto University professor of practice and entrepreneur, who frames societal resilience as the capacity to absorb shocks: the “never normal” era asks for continuous learning and, often, several sources of income.

Jacobson, LaLonde and Sullivan found earnings losses averaging around 25% of expected earnings five years after displacement, with little recovery after year three, and the largest losses among long-tenured workers in declining industries [38] – the mill-town profile, and the case for holding more than one source of income. But specialist jobs pay the highest salaries, and specialised training pays more and retains workers better precisely *because* the skills are non-transferable.

Specialisation is efficient and fragile; the question is who bears the fragility. Today the individual bears it alone. Our position is to keep specialisation and socialise its downside, which is better than “everyone should hustle” – advice that asks workers to accept lower lifetime income as the price of insurance they should not have to self-fund.

There is no single population of people holding more than one job. Lalé finds the typical multiple jobholder well-educated, holding a full-time primary job and taking a second briefly for extra income [39]; Indeed's Hiring Lab found in June 2026 that active job seekers with multiple jobs were concentrated in hourly and gig roles, with 79% describing their income as tight or inadequate [40]. A senior data scientist who consults two days a month and a care worker who drives for a delivery app to cover the rent both appear in the statistics as multiple jobholders, but the first is exercising an option and

the second is absorbing a shortfall. Telling her to build a portfolio describes her situation back to her as though it were a strategy.

2.4 What decays during unemployment is not skill

Everyone assumes that what decays during unemployment is skill. It does not. Cohen, Johnston and Lindner found no skill decline during unemployment [41], which forces other explanations for why a longer spell out of work is harder to reverse. These include lost match-specific value, the fit between a particular worker and a particular employer; signalling, where time out of work is read by recruiters as information about the candidate; and network decay, the erosion of the relationships through which most work is actually found.

A translator made redundant in June finds that by October the agencies that used to ring her have stopped ringing, though her languages are exactly as current as they were. What has gone is the web of people who thought of her first, and the quiet assumption that she was in work. Holding one smaller engagement alongside the main one – editing a trade journal, teaching an evening class – keeps a thread of that web alive.

That redirects the remedy. If skills rotted, training would be the answer, which is what almost every transition instrument supplies. If networks decay and signals sour, training is beside the point. The OECD's finding fits: adjustment happens through unemployment spells and openings for newcomers rather than through retraining incumbents [14], which is what you would expect if the binding constraint is connection rather than capability. A retraining plan that leaves someone alone with a curriculum for nine months addresses what was not decaying and leaves untouched what was.


2.5 AI removed the rungs that made experts

The traditional path to expertise runs through the bottom rungs: the junior does easy work badly, learns from doing it, and becomes senior. AI automated the easy work, so the load-bearing rung went first.

A junior paralegal used to spend her first year reading discovery documents, summarising depositions and checking citations: clumsy, slow work that taught her how a case is actually built. A language model now does those tasks in seconds, and the firm gets better output sooner and cheaper. What nobody buys any more is the year she spent learning the grain of the material, and she will need it on the day her job is to decide whether a summary she did not write can be trusted.

The rungs that made experts are the ones AI took first

How a junior became a senior, and what is left of the route



The missing rungs

Figure 3. The junior tasks that produced expert judgment are the ones automated first.

The evidence is unusually specific: a roughly 16% relative employment decline for 22–25-year-olds in the most exposed occupations, arriving through hiring rather than firing [8]; entry-level postings down around 35% in three years [9]; and the occupations themselves still growing.

If AI compresses the gap between novice and expert, as the 34%-for-novices result suggests, a firm gets usable output from a novice much sooner and has less reason to fund the years of graduated responsibility that used to convert novices into experts. The firm captures the compression; nobody captures the formation. This is a textbook example of a positive externality – a benefit that accrues to everyone and that no individual actor has a reason to pay for. And the value of what is no longer being formed is rising, not falling: AI-assisted consultants were 19 points worse outside the frontier, so the scarce skill is knowing which side of it you stand on. You need to know when to use these systems and when not to, and above all when to look outside the session for knowledge you do not have. That judgment was built on the rungs that are now being removed.

2.6 Nobody is entitled to the whole surplus

Every proposal to redistribute the gains from AI runs into the same objection: they built it, so they own it.

The return to a breakthrough goes to the idea *plus* an enormous amount of luck. Somebody else would have had the same idea a little later; what separates the firm that captured the market from the firm that arrived eighteen months afterwards is timing and market structure, not merit or effort. Winner-take-most markets in a general-purpose technology produce two or three winners and a long tail of near-identical losers, and the gap between first and second is set by the shape of the market rather than by the size of the contribution. A return calibrated to luck is a legitimate object of policy in a way that a return calibrated to contribution is not.

That is the moral premise most redistributive proposals skip, and skipping it is why they are so easily dismissed. A levy on the surplus is a claim on something the winner did not entirely earn. Without that reason, it is confiscation with a spreadsheet attached.

2.7 Money and a placement are not a life

For a lot of people, what matters is what they are achieving in a career and a journey. Take that away and you detract from society rather than improve it, even if it is more economical for owners. This cuts against naive leisure-utopianism. Marie Jahoda's research on unemployment in the 1930s identified five things a job supplies beyond the wage: a structure to the day, contact with people outside the family, a shared purpose larger than oneself, social status and identity, and regular activity [42]. A 2022 meta-analysis of longitudinal evidence finds a robust negative relationship between unemployment and wellbeing, around $d = -0.57$ for mental health, worsening with duration and – importantly – *not* moderated by the generosity of benefits [43]. A better safety net does not neutralise the cost of not working.

Every transition instrument we surveyed – Sweden's transition councils, the Dutch statutory payment, the German Sozialplan, the French redundancy plan, training levies – is a money-and-placement mechanism. None of them addresses structure, purpose or status. Any "AI frees up your time" proposal that supplies income but not the other four functions predicts harm. That gap is the whole case for funding community work as something people are actually employed to do, which is an expensive way of delivering what a simple transfer delivers more cheaply. Money is not sufficient, and the evidence for that is a century old.

2.8 Proposed solutions

Alternatives are now being proposed, and several countries have been running them for years. Before adding to the pile it is worth being accurate about what already exists, because most of the machinery a transition needs has been built somewhere already, and the gap is narrower and more specific than the volume of commentary suggests.

Sweden has run an employer-funded transition obligation since 1974. TRR, founded that year by the Confederation of Swedish Enterprise and the white-collar bargaining cartel PTK, is the closest existing thing to what this paper proposes: it is employer-funded, and provides guidance, placement, self-employment support and an income top-up to 80% of former salary, across 30,000 to 35,000 affiliated companies [2].

The levy is not flat. Confederation members pay 0.4% of payroll in 2026, down from 0.55%; firms affiliated through a *hångavtal*, a supplementary agreement for non-member employers, pay 0.95%; and state support of up to 0.15% is available [44]. Sweden did not stop in 1974: the 2022 reform, comprising an overhaul of the Employment Protection Act, a main agreement between the social partners, and a statutory transition study allowance paying up to 80% of prior income for up to 44 weeks of study, is the largest change since 1982 [45]. TRR reports that around nine in ten people leaving its programmes move on to something else, though that figure counts a new job, self-employment **or entry into study**, so it is broader than a re-employment rate.

Germany and the Netherlands compel a plan or a payment. Germany's Sozialplan process legally requires a plan jointly authored by the employer and the works council, compellable through an arbitration board whose award carries the force of a works agreement [46]. France's *plan de sauvegarde de l'emploi* can be reached either by a validated collective agreement or by a unilateral employer document approved by the labour administration [47]. The Dutch statutory transition payment is capped at €102,000 in 2026 [48].

What all of these have in common is the **metric**. Every regime links the obligation to tenure, to salary, or to a flat payroll levy. None links it to the saving. An employer saving €1 million a year and an employer saving €10,000 pay the same rate. That is the dial this paper proposes to change, and it is the one genuine gap in an otherwise well-built set of institutions.

Others are converging on the same territory from different directions. **A tax on AI usage has been worked out in detail:** Irwin, Wu and Barez's "Token Taxes" proposes collecting a surcharge where tokens are consumed rather than where the model is hosted, so that countries which buy AI services also collect revenue from them [49]. That matters for a country like Finland, where spending on US model providers may currently be taxed nowhere in Finland at all. **The largest supplier has proposed something along these lines, and stopped short of the mechanism.** OpenAI's April 2026 policy document argues that AI will shift economic activity from labour income toward corporate profits and capital gains, eroding the tax base that funds social insurance, and that governments should respond by taxing capital more heavily and exploring "new approaches such as taxes related to automated labor" [50]. It goes further than a tax proposal: it concedes that the gains may concentrate in a small number of firms, OpenAI among them, and that workers may see productivity rise without seeing any of the benefit. It proposes efficiency dividends, converting freed capacity into shorter hours or better benefits, and portable accounts for pensions and training that follow the worker rather than the job. Both are close to what this paper argues for. But the tax is fiscal, raising general revenue to keep social insurance solvent, and the efficiency dividend is an incentive offered to employers rather than a duty binding them. Neither connects the saving a particular employer books to what that employer owes the particular person whose job produced it. The most interested party in the industry has now written thirteen pages on the transition without proposing that

connection, which is a reasonable measure of how far the gap is from being closed. OpenAI is not disinterested here – a levy on automation is compatible with selling the technology that drives it – and the same question can be asked of Anthropic’s retraining commitments.

Taxing income over a lifetime rather than a year is an old idea with a recent revival: assessing cumulative lifetime earnings rather than discrete years removes the penalty that annual assessment imposes on people with spiky incomes [51]. Its extension to companies, allocating development costs across a product’s whole lifetime and across all the countries it sells into, is already the subject of an international framework in the OECD’s work on base erosion [52].

Conditional forms of basic income tie income to contribution. Daniel Susskind’s conditional basic income pays people for socially useful work rather than paying unconditionally [53]. It is state-funded and universal; this paper’s version of the idea is narrower, firm-specific and employer-funded, offered as an alternative to a redundancy. **The idea of a working life built from several concurrent engagements is Charles Handy’s**, from *The Empty Raincoat* in 1994 [54], though Handy wrote about knowledge workers choosing variety rather than the involuntary population this paper is most concerned with.

And the combination of a levy with portable benefits is already on the table. IPPR proposed in May 2026 a legal duty to disclose and consult on AI deployment, a worker-support levy funding retraining, and benefits that follow the worker rather than the job [55]. That is close to what follows here.

2.9 What is missing from all of it

The consensus now agrees the transition is the risk, and it still has almost no mechanism for who pays.

The mechanisms that do exist sit in three disconnected places. Labour law answers the question “what does an employer owe a displaced worker” with concrete, litigated numbers, but none of it was designed with AI in mind and all of it is keyed to tenure or payroll rather than to any saving. AI policy, mostly written in 2025 and 2026, is normatively engaged and mechanically thin: it argues that employers should do more, gestures at gain-sharing, and converges on no formula. Fiscal instruments raise revenue without naming a destination for it. Nobody combines the three.

Finland illustrates the shape of the gap. AI policy activity there is real – there are party working groups, parliamentary questions, a joint statement from the main labour-market organisations, a prime-ministerial review of AI in public services, and a national AI advisor appointed in April 2026. But it runs on a competitiveness track and a public-sector-efficiency track. None of it engages shared transition responsibility, work built from several income sources, or redistribution. The gap is one of framing rather than of attention.

There is a live opportunity to change that framing. The European Commission's Quality Jobs Act consultation closes on 28 September 2026 [18], and it is the natural home for a rule that sizes an employer's transition duty to the saving the employer books. That consultation runs under Article 154 of the Treaty on the Functioning of the European Union, which makes it a bilateral dialogue between the Commission and recognised EU-level social partners: the first phase consulted 34 organisations, and it is not an open portal. We cannot take part in it. But we have something worth saying from a ringside seat – as people who employ, who have paid the costs described here, and who work in three countries with three different sets of protections – and that is why we wrote this paper and offer the proposals that follow.

The end state is not the problem; the journey is. AI will change how work is done, and whether it ends in abundance or in hardship, the fifteen years in between are where people lose houses, careers and health. The damage is not evenly spread: it falls on the young, on those entering occupations whose bottom rungs have been automated away, and on specialists whose skills do not transfer. The gains, meanwhile, capitalise immediately and accrue to a small number of owners. No individual firm can fix this, because a firm that carries the cost of an orderly transition loses to a firm that does not; which rules out every voluntary solution. Nor can individuals insure themselves against a shock that hits an entire occupation at once. Money alone does not fix it either, because a job supplies structure, contact, purpose and status alongside the wage, and a transfer supplies only the wage. Meanwhile the institutions built for the last transition – notice periods, severance keyed to tenure, retraining schemes aimed at skills that were not the thing decaying – are aimed slightly to one side of the actual problem. And the one instrument nobody has built is the one that connects what a firm gains from automating a role to what it owes the person who held it.

That is why the proposals below take the form they do. They are addressed to three audiences, because the responsibility is shared: to workers, who can do a few concrete things that materially reduce their exposure; to employers, who capture the saving and are therefore the natural payer; and to policymakers, who alone can make a duty universal and so make it survivable for the firms carrying it. They rest on our best understanding of the evidence and on what is already done elsewhere – in Sweden, Denmark, Germany, Austria and Finland – rather than on invention. Where something is untested we say so.

Part 3 – Proposals for acting in the age of AI

These are proposals for consideration rather than settled conclusions. We have written them at the level of specificity where they can be argued with – what each would cost, who would pay it, and what would show it was working – because a vague proposal cannot be tested and a reader cannot usefully disagree with it. That is not the same as claiming they are right. Complex situations have no known solutions worked out in

advance, so you set a direction – the greatest good for the greatest number – and try things: prefer small reversible experiments over one grand instrument, build in measurement, and expect to be wrong about some of it.

Two caveats first. **Who is in the room:** three men – a CEO, a contractor, an investor – proposing a transition design for workers, with no worker or union voice among the authors. Worth admitting that when we first talked this through, unions appeared in our own reasoning mainly as a lever for reaching a large number of workers at once, which is precisely the instrumental posture the caveat is about. We frame the collective-agreement proposal as co-design *with* labour rather than a channel *through* it, and a reader is entitled to check whether the rest of the paper honours that.

Asymmetry of exposure: the most exposed roles are disproportionately entry-level and held by recent graduates and by women, so a mechanism blind to distribution will not reach the people hit first. Global exposure is large too – Philippine business-process outsourcing, Indian IT services, freelance work sold from lower-income countries into rich markets – and a tax on AI usage collected where it is consumed funds the rich country's transition while the displacement lands in Manila. We suggest how we could change that picture, and refuse to ignore it.

For workers

Find out what AI does to your own occupation, task by task

Do not “learn AI.” Find which of the tasks your job is made of are affected, and how. This is the only actionable version of the advice, because the boundary of what these systems do well is jagged and invisible from inside: the consultants in the BCG experiment could not tell which side of it they were on, and were confidently wrong on the far side. An occupation is a bundle of tasks – some the systems now do better than you, some they do worse while appearing to do better, some they cannot touch at all – and the mix differs by role, so it cannot be learned in the abstract.

Write down the tasks in your week at the granularity of a timesheet, try current systems on each, and record whether the output was better, indistinguishable, or confidently wrong. The confidently wrong column is the valuable one: your judgment is now worth more there, and it disappears if you only use the tools on work you cannot check. Where a task is done better by the system, develop supervision and verification rather than competing; where it sits outside the frontier, that is where your bargaining power lives. For anyone early in a career this is urgent, because the rungs that converted novices into experts are being removed from exactly those tasks, so the formation that used to arrive passively must now be sought deliberately.

Most people cannot do this alone, and should not have to. Working out which parts of a job should be handed to a machine is a genuine analytical exercise, and firms are frequently no better at it than their employees – they automate what is easy to automate

rather than what is sensible to automate. Where neither the firm nor the worker can do this unaided, they should be able to buy in help, and it should be paid for either by the employer or out of innovation funding, which is what innovation funding is for. In Finland around 40% of employed people already use generative AI at work [24], so for most this is auditing what has already happened rather than gaining new access.

Treat dependence on a single employer as a priced risk

Hold transferable skills and, where you can, more than one income relationship. You are trading a little average income for a large reduction in variance, and against a 25% five-year earnings loss after displacement that is usually a good trade. It is not universal advice: for someone already holding two jobs to make rent, this describes a problem rather than a solution.

But do not take on a second income source without checking whether the surrounding systems recognise it, because in most countries they do not, and the gap only becomes visible when it is too late to do anything about it. Pension accrual is the first: OECD figures put self-employed workers at 78% of an equivalent employee's pension on average, and in the UK only about 20% of self-employed people earning over £10,000 contribute to a private pension at all, against over 80% of employees, because automatic enrolment does not reach them [56]. Mortgage underwriting is the second: it relies on payslip verification, so someone with two smaller income streams can be assessed as a worse credit risk than someone with one larger stream of the same total. Unemployment insurance is the third, and in Finland it is the sharpest: you cannot currently accrue earnings-related unemployment insurance as an employee and as an entrepreneur at the same time, so multi-form workers must guess their status in advance and the resilient choice can leave a person with less insurance than the one who did not make it [57].

The United States adds a fourth, and it bites hardest of all. Employer health insurance and retirement contributions are commonly tied to a weekly-hours threshold in the region of 25 to 30 hours, so two half-time jobs can leave a worker below the line at both and insured at neither. Self-employment compounds it: an employee and their employer each pay half of the payroll tax funding Social Security and Medicare, while someone working for themselves pays both halves, roughly 15.3% against 7.65%. Splitting a working week across two engagements can therefore cost a person their health cover and raise the tax on the same total income. Where the answers are no, a second income source transfers risk to you without compensating you for it.

Build and keep networks, because that is what actually produces options

Of the three things that decay during a long spell out of work – the fit with a particular employer, the signal your CV sends, and your network – the network is the one you can maintain in advance. Keep live relationships outside your employer: former colleagues, people in adjacent firms, the professional community in your field. Do it while employed, when it costs nothing and you have something to offer in return. Its function is not

sociability but visibility, since this is the mechanism through which suitable openings surface at all, which is why breadth across firms matters more than depth within one.

For employers

Scale what you owe to what you saved, and offer a choice

Where a firm above a size threshold reduces or removes a role because a task has been automated, the employer chooses. **Either** redeploy the person into an equivalent role they accept, and keep the entire saving. **Or** commit a defined share of the first-year saving to a jointly authored retraining and placement plan for that person. The obligation attaches to the saving rather than to the termination, and it scales with firm size and profitability.

A mid-sized insurer automates most of its claims assessment, removes twelve roles, and books a first-year saving of €610,000. If the required share were half, the firm owes €305,000 towards retraining and placing those twelve people, or a little over €25,000 each. A small accountancy practice that automates one bookkeeper's tasks and saves €34,000 owes €17,000 for one person on the same rule. The per-person figures differ because the savings do, which is the point: what an employer owes tracks what that employer gained, and the two vary by an order of magnitude across firms. Neither is a fine; each is a fraction of money the employer has gained, which is why the duty is affordable by design. And in both cases the alternative is open: find the person another role they accept, and keep everything.

Sizing the duty to the saving, not to the sacking


What every existing system measures, and what we propose measuring instead

What existing systems key the obligation to



An employer saving €1 million a year and one saving €10,000 pay the same rate.

What we propose keying it to: the saving the employer actually booked



And the retraining need not fall on the employer alone



The new job may be at the same firm or at another one. Either counts.

Sizing the duty to the saving

Figure 4. What existing systems measure, and what we propose measuring instead.

One sceptical case deserves an answer we do not fully have. Nordic and Dutch systems deliberately do not care *why* someone is made redundant, on the reasoning that carving out an AI-caused category invites years of litigation about causation. Our answer is that the trigger should be the *saving* rather than the *cause*: a cause is a contested narrative, whereas a saving at least leaves an accounting trace. But a saving is itself gameable, through transfer pricing, through timing, or through reclassifying roles so that no position is identifiably the one that went. A firm automates three of five tasks in its customer-service roles, renames the surviving hybrid jobs “client relationship managers”, and books the saving against an old pure-agent role that is no longer on the organisation chart: new title, no identifiable displaced person, and the obligation dissolves into a reclassification exercise. This moves the argument from causation to valuation. That is the better argument to be having, since valuation has audit machinery behind it and causation does not, but the design owes that admission.

Fund the transition before the exit

Transition support should begin before separation and be delivered as a programme rather than paid as a lump sum. Nokia’s Bridge programme, established in 2011 as phone R&D centres and factories closed across thirteen countries, eventually covered

around 18,000 employees, and remains the most useful case of employer-funded managed transition at scale, including funding for employees to start companies [58].

It should be used carefully, because the headline figures do not mean what they are usually taken to mean. The commonly repeated numbers – 60%, over 1,000 new companies, 85% satisfaction – are company-reported, drawn from Nokia-commissioned and Harvard Business School material, and are not independently audited. The 60% figure means knowing one's next step on the day of exit, which is much broader than re-employment, and no independent source confirms the satisfaction figure. The one independent study, an Aalto dissertation surveying eight sites, reports improved wellbeing only in qualitative terms [59]; a companion Hitotsubashi study found only around 10% took the entrepreneurship track [60].

Two things follow. The first is that Bridge worked in a specific way: what it delivered was orientation, contact and a plan while people were still employed, rather than a high placement rate afterwards. That is worth having, and it is what the pre-separation timing buys. The second is that Nokia was cash-rich and paid for the whole thing itself, which is why it cannot be the model. If managed transition only happens when a profitable company chooses to fund it out of goodwill, it will not happen at the moments it is most needed. The lesson to take from Bridge is to start early and staff it properly; the lesson not to take is that employers will do this unprompted.

Write three contract clauses, and know which one decides everything

The institutional route in Finland is the sectoral collective agreement, where terms negotiated centrally then apply across a sector including to non-organised employers, so changing the template changes conditions for many people at once. One qualification: bargaining has decentralised since 1 January 2025, permitting local deals even in sectors covered by a generally applicable agreement [61], so the lever still exists but now runs through local channels too. Our position is central floor-setting plus local flexibility in how the floors are met. None of this requires weakening dismissal protection, and we do not propose it.

First, narrow non-compete clauses to genuine trade-secret protection. A non-compete that protects a commercial interest which no longer exists once the employee has left is suppressing a competitor's labour supply rather than protecting a secret, and it is the single largest obstacle to holding more than one income relationship. Scope it to what a trade secret actually requires and it survives; leave it as a general restraint and it damages those least able to litigate.

Second, flip the default on side work from permission-required to permitted-unless-conflicted. A permission regime makes every side engagement a negotiation in which the employee has the weaker hand. A permitted-unless-conflicted default instead requires the employer to identify an actual conflict, which is what the clause is legitimately for. Same protection, opposite outcome.

Third, and this decides what the instrument is: write into the contract that the employee decides where in an elastic hours band they sit. A band of 60–100% is a flexibility instrument if the employee holds that decision and a precarity instrument if the employer does, and nothing else distinguishes them. **If the employer holds the decision, an elastic band is a zero-hours contract with better branding.**

Fund and regulate senior supervision hours, because that is the scarce input

Denmark already operates most of the machinery. Its Employers' Reimbursement Fund obliges every employer, public or private, to contribute whether or not they train anyone, and reimburses training firms for a large share of an apprentice's salary while the apprentice is at school. Layered on top, a second scheme sets each employer a trainee-number target, pays those who exceed it and charges those who fall short – a 2016 tripartite response to a shortage of training places, with evaluations finding significant positive employment effects [37]. A bonus and a penalty on training capacity, funded by all employers including those who train nobody, is Danish law.

Point that instrument at the occupations where the rungs are disappearing fastest – paralegals, junior developers, analysts, copywriters – and specify the regulated unit as **senior supervision hours** rather than headcount, because headcount is what firms can satisfy by reclassifying roles they already have. The binding constraint on producing an expert is not the number of juniors hired but the amount of experienced attention available to them, and that is the input which is now scarce.

Two cautions. **A ratio without funding becomes a cap.** The US froze the number of funded medical residency slots in 1997, and it stayed frozen until 1,200 positions were authorised in 2021 and 2023, which limited the supply of doctors for a generation [62]. Denmark pairs its target with money, and this must too. And **supervision hours are harder to audit than headcount**, which is a real weakness in the design; the more administrable pairing may be headcount targets plus a supervision requirement.

The public sector should demonstrate this, not just fund it

Every fiscal proposal here treats the state as the funder of transitions. But the state is also one of the largest employers of exactly the administrative roles most exposed: a tax office with complete data can automate assessment almost entirely, and the support functions around it are then not needed. So the public sector is where a transition model can be mandated and demonstrated rather than merely paid for. Finland's prime-ministerial review of AI in public services runs to 31 October 2026 with a brief of efficiency, and a pilot costs one sentence in that mandate.

For policymakers

Make the transition duty universal, at the level where firms compete

A duty that binds one firm kills that firm; a duty that binds all of them becomes a cost of business. So the duty has to be set at the level where the competition actually happens, which for most sectors in Europe means the European Union rather than any single member state. Set it nationally and firms move the activity, or simply lose to competitors who did not have to pay it.

There is an existing instrument to attach it to. The Collective Redundancies Directive already requires an employer making 20 or more redundancies in 90 days to consult on “accompanying social measures aimed at redeploying or retraining workers” [63]. That is the right obligation with the wrong sizing rule: it says the employer must discuss retraining, but says nothing about how much the employer should spend on it, which in practice means the answer is often nothing. Amending it to size the obligation to the saving the automation produced would turn a consultation duty into a funded one, and it would do so through machinery that already exists, that firms already comply with, and that already operates at the level where firms compete. The Quality Jobs Act consultation closing on 28 September 2026 is the live opening for exactly this.


Implement a tax on AI usage, and earmark it to a named purpose

A token tax collects a surcharge at the point where AI is used rather than where the model is hosted, with compute providers acting as intermediaries since vendors already meter usage by token [49]. The design matters because inference can run wherever tax is lowest, whereas a charge at the point of use cannot relocate. That is a rare property for a digital tax base and it is worth using while the market is still forming.

The revenue should be earmarked, to the transition duty above and to the community work below, so that the levy funds the adjustment that the technology causes. Hypothecation is normally poor tax design and we propose it deliberately: an unearmarked levy gets spent on whatever is politically pressing, and transition is what nobody funds. A named destination also makes the levy defensible to the firms paying it, ourselves included.

How a token tax would work

A small charge on AI use, collected where the use happens, spent on the transition it causes



How a token tax would work

Figure 5. A charge on AI use, collected where the use happens and earmarked to the transition it causes.

A threshold for small firms is a condition of our support, not a detail. 3DBear is a heavy consumer of tokens and would pay this levy, so we have an interest to declare in the shape of it. A flat rate hits hardest the firms least able to absorb it, and would fall on small adopters at the point where adoption is doing the most good. Exempt them.

One concession is unresolved. Collecting at the point of consumption funds the consuming country's transition, so a rich-country levy funds rich-country adjustment while the displacement lands in supplier economies. We have no fix for that.

Fund community work through a statutory pooled fund, and name the work

Channel released capacity into paid community work through a statutory, pooled fund rather than through any firm's CSR budget. Money alone is insufficient, because a cheque supplies income and none of the other four things a job provides – the structure of the day, the contact with others, the shared purpose and the status.

Councils cannot currently fund work that visibly needs doing and requires little training. Name it: streets and parks; refuse collection; daily support for older and disabled people, meaning shopping, company and help getting to appointments; libraries and youth facilities; local food markets and community events; small home repairs; and nature management, path clearing, habitat work and invasive-species removal of the

kind conservation bodies currently run on volunteers. These are real jobs left undone for lack of a budget line, not make-work.

The evidence that this works is Austrian. The Marienthal job guarantee pilot, which ran from October 2020 to March 2024, offered guaranteed jobs to people who had been unemployed for nine months or more, using pairwise matched randomisation and a pre-registered synthetic control. It produced a large drop in long-term unemployment, no negative spillovers into existing employment, and gains in financial situation, health, self-efficacy and social inclusion [64]. Austria has committed €50 million to a wider rollout [65].

Five conditions separate this from workfare, and they belong in the enabling legislation rather than in implementation guidance. The work is genuinely voluntary. It is paid at the market rate, not as a top-up to a benefit. It is chosen by the worker, not assigned. It does not displace existing public employment. And declining carries no sanction.

Put them in the statute, because each is the thing an implementing authority under budget pressure erodes first. A scheme that begins as an offer becomes a condition of benefit within a few years unless the law forbids it, and at that point it stops delivering what makes it work and starts delivering harm. Funding must be pooled and statutory for the same reason: schemes financed at discretion get cut in the downturns when they are most needed.

Fix pensions, mortgages and unemployment insurance before promoting flexible work

Portable pension accrual, mortgage underwriting that recognises multi-source income, and unemployment insurance that runs continuously across employed and self-employed status must be in place *before* elastic contracts and multi-source working are promoted as policy. The sequence is the recommendation. Without it, promoting the norm reaches the involuntary worker as an instruction to bear more risk for less protection. Finland's combination-insurance legislation is the model for the third of these, and the closest existing proposal for benefits that follow the worker rather than the job is IPPR's [55].

Adopt a basic income, and decide the withdrawal rate first

We propose a basic income: a regular payment made to everyone as of right, replacing the means-tested patchwork that currently withdraws support as earnings rise. The case for it here is specific rather than utopian. If work is going to arrive in irregular quantities – three days a week this year, a contract next year, a gap after that – then a system which assumes a full-time job and treats everything else as an exception will misfire constantly, and it will misfire hardest on the people with the least slack to absorb it. A floor that does not need to be claimed, reassessed or justified is administratively simpler than what it replaces, and it makes irregular work rational to take rather than something to be weighed against losing your benefit.

Whatever the floor is called, the withdrawal rate decides whether it works. This is the point that survives whether or not you accept the basic income argument. If support is withdrawn steeply as earnings rise, a part-time job is not worth taking, and every other proposal in this paper that depends on flexible working collapses. Universal Credit's 55% taper, with the single parent keeping about £70 of her £200, is the illustration: the arithmetic, not her character, produced the decision. Denmark is the cautionary case in the other direction – its flexicurity model paired easy dismissal with strong security, and the security has eroded since the 1980s while the flexibility survived every reform [66]. Specify the taper first, at a rate slow enough that a job at 40% of full-time hours is rational, and then decide what to call the floor.

The limit of the evidence deserves stating. Finland's 2017–2018 Kela experiment paid €560 a month to 2,000 unemployed people for two years. Employment effects were small and *positive*, with recipients averaging six more days employed than the control group's 78, and wellbeing effects were unambiguously positive: higher life satisfaction, less mental strain, less loneliness, better cognitive function [67]. Calling that result "mixed", as it is often summarised, understates the wellbeing findings considerably. What no pilot has tested is a floor against simultaneous, correlated displacement: every one of them ran during a broadly functioning labour market, and the mechanism observed – recipients using the security to find *better* work – presupposes that better work exists to be found. That is the scenario this paper worries about, and it is untested.

Design for robustness across timelines, not for a forecast

Prefer instruments that are cheap to reverse and that work at both 25-year and 5-year speeds, with measurement built in. At 25 years, normal churn and retirement absorb most of the adjustment and urgent instruments are the wrong answer. At 5 years, nothing proposed here is fast enough. Nobody knows which it is, and the honest response to that is to give every instrument a sunset clause requiring evidence to renew, rather than tuning the whole design to a forecast.

Tax cumulative lifetime income, not discrete years

People with irregular income are systematically overtaxed by annual assessment relative to salaried workers with identical lifetime earnings, because a good year is taxed at a high marginal rate while a bad year wastes the allowances. Assess cumulative lifetime earnings instead: a redundancy year then generates a rebate, spiky contract income is treated fairly, and a year spent training is recognised rather than penalised. An individual cannot elect into a regime that does not exist, so this is a demand on tax policy rather than advice to a worker. Finland's five-year loss carry-forward is a partial version already in law.

Part 4 – The recommendations in brief

For workers

1. Audit your job task by task: which parts systems do better, which worse while convincing, which they cannot touch. Get outside help if you need it, paid for by your employer or from innovation funding.
2. Price dependence on one employer as a risk, not a loyalty.
3. Keep live professional relationships outside work – the network decays when work stops, not the skill.
4. Aim for more than one income source; check whether pension, mortgage and unemployment insurance recognise the second before you rely on it.

For employers

5. When you automate a role, either redeploy the person into a job they accept and keep the whole saving, or spend a defined share of that saving retraining them.
6. Start transition support before exit, as a programme, not a payment.
7. Cut non-competes to real trade secrets, permit side work unless it conflicts, let employees choose where in an elastic hours band they sit.
8. Fund and regulate senior supervision hours – AI removed the rungs that made experts, and senior attention is now the scarce input.
9. Public sector: run the model on your own staff first, and publish the results.

For policymakers

10. Legislate the transition duty at the level firms compete, sized to the saving rather than to tenure, through the Collective Redundancies Directive.
11. Implement a tax on AI usage, earmark the revenue to transition and community work, and exempt small adopters.
12. Fund named community work – streets and parks, refuse, daily support for older and disabled people, libraries and youth facilities, markets and events, small repairs, nature management – from a statutory pooled fund.
13. Write the five voluntariness conditions into legislation: voluntary, paid at market rate, chosen not assigned, no displacement of public jobs, no sanction for declining.
14. Fix pension accrual, mortgage underwriting and unemployment insurance for multi-source income before promoting flexible work.
15. Adopt a basic income, and decide the withdrawal rate before the floor – the taper decides whether part-time work pays.
16. Tax cumulative lifetime income, not discrete years, so that redundancy and training years stop being punished.
17. Give every instrument a sunset clause requiring evidence to renew: nobody knows whether this runs five years or twenty-five.

Closing

The argument about whether the AI job transformation will happen is over. The argument about how it will happen has barely started.

What we are confident of is this. The harm lands in the transition rather than in the end state. A firm acting alone cannot fund a decent transition and survive competition, so the rule has to bind everyone or it binds no one. What an employer owes should be sized to what that employer gained, because tenure and payroll have no connection to the saving and the saving is the thing that was actually captured. And the ladder that produced expert judgment is being dismantled from the bottom at exactly the moment that judgment becomes more valuable, which is a problem no firm has an incentive to solve alone.

From those four things the proposals follow. Give employers a choice between redeploying a person and funding their move to something else, and make that duty universal so that the decent employer is not punished for being decent. Fund the transition before the exit rather than after it. Pay for the supervision that turns novices into experts, since no single firm captures the return on it. Tax AI usage where it is consumed and spend the revenue on the adjustment it causes. Pay people to do community work that visibly needs doing, with the conditions that keep it voluntary written into the law rather than into guidance. Put a floor under everyone and set the withdrawal rate before arguing about the name. And fix the pension, mortgage and insurance treatment of irregular income before telling anybody that irregular income is the future.

None of this is exotic. Sweden has run the employer-funded transition obligation since 1974. Denmark funds training capacity through a levy on every employer including those who train nobody. Austria has evaluated a job guarantee and committed money to extending it. Finland has run a basic income experiment and is legislating on combination insurance. What does not exist anywhere is the connection between the size of the gain and the size of the obligation, and that is the one genuinely new thing proposed here.

It is worth being clear about what success would look like, because the pessimistic version of this argument is much easier to write. A society that gets this right does not simply avoid a disaster. It ends up with people working somewhat less for the same living standard; with the years around a career change funded rather than survived; with the option of building a working life from more than one thing, taken up because it is attractive rather than because the alternative collapsed; with community work that needed doing for decades finally done, by people paid properly to do it; and with expertise still being formed, because somebody decided to pay for the formation once the market stopped doing it by accident. That is a better world than the one we have

now, and it is reachable from here. The technology does not decide which version we get. The rules do, and the rules are still being written.

We argue for a stronger employer obligation from inside a small company that has paid to let someone go and felt the cost. That is not a contradiction we are embarrassed by; it is the reason to trust the argument. The duty should attach to the profitable firm that automates a role and books the saving, not to the small firm shedding a role it cannot afford. Current law does not make that distinction. It should.

References

1. “We Must Act Now,” Stanford Digital Economy Lab, 13 July 2026.
<http://wemustactnow.ai/> ·
<https://digitaleconomy.stanford.edu/news/wemustactnow/>
2. TUAC, “The Swedish Job Security Councils: social partners’ led transitions” (TRR).
<https://tuac.org/news/the-swedish-job-security-councils-a-case-study-on-social-partners-led-transitions/> The 0.3% levy figure circulating in secondary sources of this vintage is out of date; see reference 44.
3. Sheffield Hallam CRESR, “The State of the Coalfields 2024,” 19 April 2024.
<https://www.shu.ac.uk/centre-regional-economic-social-research/publications/state-of-the-coalfields-2024>
4. Beatty & Fothergill, “The impact on welfare and public finances of job loss in industrial Britain.”
[https://shura.shu.ac.uk/15994/8/Beatty%20-%20impact%20on%20welfare%20and%20public%20finances%20\(VoR\).pdf](https://shura.shu.ac.uk/15994/8/Beatty%20-%20impact%20on%20welfare%20and%20public%20finances%20(VoR).pdf)
5. Teddy Tawil, “The AI Labor Debate: Three Views on the Future of Work,” Carnegie Endowment, 23 April 2026.
<https://carnegieendowment.org/research/2026/04/the-ai-labor-debate-three-views-on-the-future-of-work>
6. David Autor, “Applying AI to Rebuild Middle Class Jobs,” NBER w32140.
<https://www.nber.org/papers/w32140>
7. Federal Reserve Bank of Atlanta, “Artificial Intelligence, Productivity, and the Workforce: Evidence from Corporate Executives,” March 2026.
<https://www.atlantafed.org/-/media/Project/Atlanta/FRBA/Documents/research/publication/working-paper/2026/03/25/04-artificial-intelligence-productivity-and-the-workforce-evidence-from-corporate-executives.pdf>
8. Brynjolfsson, Chandar & Chen, “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence,” Stanford Digital Economy Lab, November 2025, updated February 2026.
<https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>

9. Washington Monthly, “How AI Broke the Entry-Level Job,” 29 May 2026.
<https://washingtonmonthly.com/2026/05/29/ai-entry-level-jobs-college-graduates/>
10. Dell’Acqua et al., “Navigating the Jagged Technological Frontier,” SSRN 4573321; *Organization Science*, March 2026.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321
11. METR, “Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity,” July 2025.
<https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>
12. METR, “We are Changing our Developer Productivity Experiment Design,” 24 February 2026. The redesign followed a selection problem; updated estimates remain negative, but METR calls them weak evidence rather than a reversal.
<https://metr.org/blog/2026-02-24-uplift-update/>
13. IMF, “AI Adoption and Inequality,” April 2025.
<https://www.imf.org/en/publications/wp/issues/2025/04/04/ai-adoption-and-inequality-565729> · IMF, “Bridging Skill Gaps for the Future: New Jobs Creation in the AI Age,” Staff Discussion Note SDN/2026/001, 9 January 2026 (the follow-up cited in 1.3).
<https://www.imf.org/en/publications/staff-discussion-notes/issues/2026/01/09/bridging-skill-gaps-for-the-future-new-jobs-creation-in-the-ai-age-572136>
14. OECD, *Employment Outlook 2026*, July 2026.
https://www.oecd.org/en/publications/2026/07/oecd-employment-outlook-2026_a41e8b9f.html
15. ILO/NASK, “Generative AI and Jobs: A Refined Global Index of Occupational Exposure,” WP 140, May 2025.
https://www.ilo.org/sites/default/files/2025-05/WP140_web.pdf
16. artificialintelligenceact.eu, implementation timeline (high-risk employment provisions to December 2027).
<https://artificialintelligenceact.eu/implementation-timeline/>
17. Iuslaboris, Platform Work Directive transposition tracker.
<https://iuslaboris.com/insights/eu-platform-work-directive-which-countries-have-implemented/>
18. European Commission, second-phase consultation on the Quality Jobs Act, opened 20 July 2026, closes 28 September 2026.
https://employment-social-affairs.ec.europa.eu/news/commission-opens-second-phase-consultation-quality-jobs-act-2026-07-20_en See also the first-phase consultation listing the 34 EU-level social partner organisations consulted under Article 154 TFEU:
https://employment-social-affairs.ec.europa.eu/policies-and-activities/eu-employment-policies/social-dialogue/consultations-social-partners_en
19. Federal Reserve, Distributional Financial Accounts.
<https://www.federalreserve.gov/releases/efa/efa-distributional-financial-accounts.htm> · FRED series WFRBST01134 and WFRBSB50215.
<https://fred.stlouisfed.org/series/WFRBST01134>

20. Positive Money Europe, on ECB Distributional Wealth Accounts (euro-area top-decile share).
<https://positivemoney.org/eu/update/wealth-inequality-in-the-eurozone-in-8-charts/>
21. CBO, “The Distribution of Household Income, 2022,” January 2026.
<https://www.cbo.gov/system/files/2026-01/61911-Household-Income-2022.pdf>
22. EPI, “Strong wage growth for low-wage workers bucks the historic trend” (the Autor, Dube and McGrew compression).
<https://www.epi.org/publication/strong-wage-growth-for-low-wage-workers-bucks-the-historic-trend/> · Primary source: Autor, Dube & McGrew, “The Unexpected Compression: Competition at Work in the Low Wage Labor Market,” NBER w31010. <https://www.nber.org/papers/w31010>
23. Federal Reserve Bank of New York, “The Post-COVID Decline in the Labor Share,” June 2026.
<https://libertystreeteconomics.newyorkfed.org/2026/06/the-post-covid-decline-in-the-labor-share/>
24. ETLA, “Kauhukuva ei käynyt toteen”: AI has not displaced young workers in Finland.
<https://www.etla.fi/ajankohtaista/uutiset-ja-tiedotteet/etla-selvitti-kauhukuva-ei-kaynyt-toteen-tekoaly-ei-ole-syrjayttanyt-nuoria-suomen-tyomarkkinoilla/>
25. ETLA, AI’s Finnish impact concentrated among the elite and knowledge workers.
<https://www.etla.fi/ajankohtaista/uutiset-ja-tiedotteet/tekoalyn-tulo-suomeen-osuudeniten-eliittiin-ja-tietotylaisiin-vahemman-koulutetuilla-on-enemman-voitettavaa/>
26. Audit Scotland, councils’ budget gap.
<https://audit.scot/news/councils-face-half-billion-budget-gap> · Scottish Parliament research briefing SB 2604, 16 January 2026: the 2026–27 settlement rose 2.9% in real terms, so the accurate claim is a structural gap rather than an annual cut.
<https://www.parliament.scot/chamber-and-committees/research-prepared-for-parliament/research-briefings/2026/1/16/sb-2604>
27. Commons Library, Universal Credit taper rate.
<https://commonslibrary.parliament.uk/reducing-the-universal-credit-taper-rate-and-the-effect-on-incomes/>
28. IFS, Universal Credit briefing note 116, source for the legacy-system marginal deduction rates. https://ifs.org.uk/sites/default/files/output_url_files/bn116.pdf
29. US Department of Labor, “Plant Closings and Layoffs” (WARN Act).
<https://www.dol.gov/general/topic/termination/plantclosings> · Legal Aid at Work, WARN Act factsheet.
<https://legalaidatwork.org/factsheet/the-w-a-r-n-act-mass-layoffs-or-businessplant-closings/>
30. New York State Senate, S8589B (SHIELD Act).
<https://www.nysenate.gov/legislation/bills/2025/S8589/amendment/B> California SB 951 would require 90 days’ notice where AI or automation causes 25 or more job losses at employers with more than 100 staff.

31. Brynjolfsson, Li & Raymond, "Generative AI at Work," NBER w31161; *QJE* 2025.
<https://www.nber.org/papers/w31161>
32. Alexander, "The Ineffectiveness of CSR: Understanding Garment Company Commitments to Living Wages in Global Supply Chains," *RIPE*, 2021.
<https://www.tandfonline.com/doi/full/10.1080/13563467.2021.1926954>
33. Employer Branding News, "IKEA's AI Reskilling Case, Reassessed After 2026 Layoffs."
<https://employerbranding.news/resources/what-ikeas-ai-reskilling-story-looks-like-now-it-has-made-layoffs-too/>
34. Fortune, on IBM's graduate-hiring pledge and subsequent cuts, 5 November 2025.
<https://fortune.com/2025/11/05/ibm-ceo-arvind-krishna-promise-hire-more-gen-z-college-graduates-but-thousands-laid-off-ai-restructuring/>
35. Council of the EU, sustainability reporting and due diligence simplification, 24 February 2026.
<https://www.consilium.europa.eu/en/press/press-releases/2026/02/24/council-signs-off-simplification-of-sustainability-reporting-and-due-diligence-requirements-to-boost-eu-competitiveness/>
36. BIBB, German dual vocational training and trends in employer training participation. <https://www.bibb.de/de/4355.php>
37. Denmark, Employers' Reimbursement Fund (AUB) and the Læreplads-AUB trainee-target scheme, from the 2016 tripartite agreement.
<https://uvm.dk/ministeriet/love-regler-og-trepartsaftaler/oversigt-over-trepartsaftaler/trepartsaftale-om-tilstraekkelig-og-kvalificeret-arbejdskraft-og-laerepladser/laerepladsafhaengigt-aub-bidrag/>
38. Jacobson, LaLonde & Sullivan, "Earnings Losses of Displaced Workers," Upjohn WP 92-11. https://research.upjohn.org/up_workingpapers/11/
39. Upjohn Institute, on Lalé's multiple-jobholding research.
<https://www.upjohn.org/research-highlights/new-research-debunks-popular-notions-about-multiple-jobholding> · Lalé, "Multiple Jobholding: Knowing the Facts to Draw Proper Policy Conclusions," Upjohn policy brief 11.
https://research.upjohn.org/up_policybriefs/11/
40. Indeed Hiring Lab, "Active Job Seekers Are Already Working Multiple Jobs," 9 June 2026.
<https://www.hiringlab.org/2026/06/09/active-job-seekers-are-already-working-multiple-jobs/>
41. Cohen, Johnston & Lindner, "Skill Depreciation during Unemployment," NBER w31120. <https://www.nber.org/papers/w31120>
42. Marie Jahoda, latent deprivation model of employment. See also *Marienthal: The Sociology of an Unemployed Community* (Jahoda, Lazarsfeld & Zeisel, 1933).
43. "The relationship between unemployment and wellbeing: an updated meta-analysis of longitudinal evidence," *European Journal of Work and*

Organizational Psychology, 2022.

<https://www.tandfonline.com/doi/full/10.1080/1359432X.2022.2106855>

44. Fora and TRR on the 2026 *omställningsavgift*: 0.4% of payroll for Confederation-affiliated companies (reduced from 0.55%), 0.95% via *hängavtal*, with up to 0.15% available as state transition support.
<https://www.fora.se/prislista> · <https://www.trr.se>
45. Sweden's 2022 transition reform: the Employment Protection Act overhaul, the *Trygghetsöverenskommelsen* main agreement, and the statutory *omställningsstudiestöd*.
<https://www.csn.se/bidrag-och-lan/omstallningsstudiestod-for-dig-mitt-i-arbetslivet/vad-ar-omstallningsstudiestod.html> · <https://www.regeringen.se/regeringens-politik/trygghet-och-omstallning/ett-nytt-offentligt-studiestod--omstallningsstudiestod/>
46. German Sozialplan (§§112–112a BetrVG).
<https://globallawexperts.com/social-plan-germany-2026-works-council/>
47. French *plan de sauvegarde de l'emploi*. Joint authorship by validated collective agreement is one lawful route; a unilateral document homologated by the DREETS is the other.
<https://www.service-public.gouv.fr/particuliers/vosdroits/F2811?lang=en>
48. Dutch transition payment 2026.
<https://dutch-law.com/employment-law/transition-payment.html>
49. Irwin, L., Wu, T-Y., & Barez, F., “Token Taxes: mitigating AGI’s economic risks,” arXiv:2603.04555 v2, 11 June 2026, Oxford Martin School.
<https://arxiv.org/html/2603.04555v2> · <https://www.oxfordmartin.ox.ac.uk/publications/token-taxes-mitigating-agis-economic-risks> See also Windfall Trust Policy Atlas, “Token Taxes,” 22 July 2026.
<https://windfalltrust.org/policy-atlas/token-taxes>
50. OpenAI, “Industrial Policy for the Intelligence Age: Ideas to Keep People First,” April 2026, p. 6.
<https://openai.com/index/industrial-policy-for-the-intelligence-age/>
51. Vickrey, “Averaging of Income for Income-Tax Purposes,” *Journal of Political Economy* 47(3), 1939, pp. 379–397; cumulative lifetime income averaging revived by Batchelder and by Gordon and Wen.
<https://www.journals.uchicago.edu/doi/10.1086/255390>
52. OECD, base erosion and profit shifting framework (BEPS, Pillar One).
<https://www.oecd.org/tax/beps/>
53. Daniel Susskind, research page (conditional basic income; Digitalist Papers, 2025). <https://www.danielsusskind.com/research>
54. Charles Handy, *The Empty Raincoat* (Hutchinson, London, 1994; published in the United States as *The Age of Paradox*), on the portfolio career. Book, no online edition.
55. IPPR, “Strike while AI is hot: rebuilding worker power for the age of AI,” May 2026. <https://www.ippr.org/articles/strike-while-ai-is-hot-worker-power>

56. OECD, *Pensions at a Glance 2025*, relative pensions of the self-employed.
https://www.oecd.org/en/publications/2025/11/pensions-at-a-glance-2025_76510fe4/full-report/theoretical-relative-pensions-of-the-self-employed_f764691a.html · IFS, “Private pensions for the self-employed: challenges and options for reform.”
<https://ifs.org.uk/publications/private-pensions-self-employed-challenges-and-options-reform>
57. STM (Finland), *yhdistelmävakuutus* project STM047:00/2025.
<https://stm.fi/en/project?tunnus=STM047:00/2025> · PAM response.
<https://www.pam.fi/artikkelit/pam-yleistuessa-korjattavaa-yhdistelmavakuutus-ei-ratkaise-olennaisia-asioita-ja-oikeus-opiskeluun-hyva/>
58. HBR/HBS, “Nokia’s Bridge Program: Outcome and Results (B).”
Company-reported figures, not independently audited; “60%” means knowing one’s next step on the day of exit rather than re-employment.
<https://store.hbr.org/product/nokia-s-bridge-program-outcome-and-results-b/315003>
59. Aalto University, “The Bridge Program: Participant Perspectives,” the genuinely independent study. <https://aaltodoc.aalto.fi/handle/123456789/16082>
60. Kang et al., Hitotsubashi, on the Nokia entrepreneurship track.
<https://pubs.iir.hit-u.ac.jp/admin/en/pdfs/file/2056>
61. TEM (Finland), increasing local collective bargaining.
<https://tem.fi/en/increasing-local-collective-bargaining>
62. US Government Accountability Office on Medicare graduate medical education funding and the cap on funded residency positions.
<https://www.gao.gov/products/gao-21-391> · GAO-26-107686 on the initial distributions of the new Medicare-funded residency positions.
<https://www.gao.gov/products/gao-26-107686>
63. EUR-Lex, Collective Redundancies Directive 98/59/EC.
<https://eur-lex.europa.eu/eli/dir/1998/59/oj/eng>
64. Kasy, Lehner et al., MAGMA (Marienthal) evaluation, IZA DP 16088.
<https://docs.iza.org/dp16088.pdf>
65. INET Oxford, Austrian government €50m rollout commitment.
<https://www.inet.ox.ac.uk/news/austrian-government-commits-50-million-to-rollout-job-guarantee-scheme-in-budget>
66. CEPR VoxEU, “The Danish flexicurity model in the Great Recession.”
<https://cepr.org/voxeu/columns/danish-flexicurity-model-great-recession>
67. VATT, results of the basic income experiment.
<https://vatt.fi/en/-/results-of-the-basic-income-experiment-small-employment-effects-better-perceived-economic-security-and-mental-wellbeing>