

MAS-ESS Essay Competition 2026

The Compression Hiding in Plain Sight

Singapore's AI Strategy needs a second axis

Dhruv Aggarwal

NUS High School of Mathematics and Science

Summary

Singapore's labour market in 2026 looks healthy on paper. Resident unemployment sits at 2.9%, real median wages rose 4.3% in 2025, vacancies still outnumber unemployed residents 1.58 to one (MOM, 2026). Yet PMET retrenchment has climbed to 10.1 per 1,000, above pre-recessionary norms, concentrated in financial services, professional services, and ICT, the verticals where Generative AI (GenAI) is most operational. DBS is restructuring its workforce. PM Wong, on Budget 2026, named "deep deep concerns" about the future of work. The standard response, reskill and grow, is necessary but calibrated for the wrong margin.

This essay argues that the distributional risk of GenAI is largely invisible to current policy. The three most cited field studies, Brynjolfsson, Li and Raymond on customer-service agents (*Quarterly Journal of Economics*, 2025), Noy and Zhang on writing tasks (*Science*, 2023), and Dell'Acqua on consultants (Harvard/BCG, 2023), find that AI compresses productivity within occupations: bottom performers gain heavily, top performers barely. The IMF (Khan, 2024) places 77% of Singaporean workers in the high-exposure band against a 60% advanced-economy average. Half face displacement risk; the other half face an erosion of the wage premium on median competence that no SkillsFuture course can address. Acemoglu (2025) and Korinek and Suh (2024) further show that even modest AI scenarios push the labour share down, while Singapore's already sits near 37%, against an OECD range of 55 to 60 percent.

Three instruments are proposed: a Complementarity Audit modelled on the Progressive Wage Model; an AI Capital Dividend channelling a small share of Temasek and GIC returns

into CPF; and an Intangibles-Adjusted Productivity Series to hold political nerve through the J-curve. Inclusive growth is not a side-effect of AI strategy. It is the condition that lets the strategy survive contact with politics.

Contents

1. The Paradox of an Inclusive Headline
 2. The Compression Pattern
 3. Why Reskilling Cannot Reach the Distributional Frontier
 4. The Capital-Share Corollary
 5. Three Proposals for a Second Axis
 6. Conclusion
- References

1. The Paradox of an Inclusive Headline

On paper, Singapore's labour market in 2026 looks reassuring. Resident unemployment is 2.9%, real median wages rose 4.3% in 2025, vacancies outnumber unemployed residents 1.58 to one (MOM, 2026). Household Gini after taxes and transfers fell to 0.379, a record low (DOS, 2026).

The same Manpower report contains a less comfortable set of figures. PMET retrenchment incidence rose from 8.6 to 10.1 per 1,000 residents in 2025, a level last seen in recessionary years rather than expansions. The increase clusters in financial services, ICT, and professional services, the white-collar verticals where GenAI has moved fastest from pilot to production. DBS announced a multi-market workforce restructuring affecting around 4,000 roles. The Singapore National Employers Federation's 2026 survey reports a majority of employers planning to freeze or slow headcount this year. PM Wong, presenting Budget 2026, allowed himself unusually personal phrasing: he had, he said, "deep deep concerns" about the workforce.

The official response, NAIS 2.0, the new National AI Council, an expanded AI Apprenticeship Programme, and Budget 2026's Champions of AI scheme, pursues two goals: grow the pie through faster AI adoption, and reskill the displaced into adjacent occupations. Both are sensible and well-funded, and neither is the subject of this essay's argument.

The argument is that a third risk has been smuggled in alongside the first two, and none of the existing instruments is shaped to catch it. GenAI does not only eliminate jobs (the IMF

estimates roughly half of Singapore's exposed workforce is at low-complementarity displacement risk; Khan, 2024). It also rearranges the distribution of productivity within occupations that survive, in ways politically benign in the short run and corrosive over a longer horizon. SkillsFuture, the workhorse of the existing distributional response, is not built for this margin.

2. The Compression Pattern

The empirical case rests on three field studies. Each looks at a different occupation, and each finds the same thing.

Study	Setting	Bottom-quartile gain	Top-quartile gain
Brynjolfsson, Li & Raymond (QJE, 2025)	5,172 customer-support agents	+34%	≈ 0
Noy & Zhang (Science, 2023)	Mid-skill writing tasks	Largest gains	Smallest gains
Dell'Acqua et al. (Harvard / BCG, 2023)	758 BCG consultants	+43%	+17%

Brynjolfsson, Li and Raymond (*Quarterly Journal of Economics*, 2025) study 5,172 customer-support agents at a Fortune 500 firm. An AI assistant raises problems resolved per hour by 14% on average, but the gain is 34% for new hires and statistically zero for the most experienced. Noy and Zhang (*Science*, 2023) run a writing-task experiment at MIT: time-to-completion drops 40% and quality rises four-tenths of a standard deviation, with the largest improvements among writers initially rated lowest. Dell'Acqua and co-authors (Harvard/BCG, 2023) study 758 BCG consultants: inside the AI's task envelope, bottom-half consultants improve by 43%, top-half by 17%.

A pattern of bottom performers catching up to the median sounds like good news, and in some sense it is. The second-order economics is, however, less comforting.

Take for example, the wage premium on median competence. If a bottom-quartile lawyer can now produce work indistinguishable from a former median lawyer, the median

lawyer's relative bargaining position weakens, even as her absolute output stays the same. The unit of inequality is no longer the gap between lawyer and paralegal. It is the much narrower gap between two lawyers in the same firm who use the same model differently.

Consider, too, where the surplus goes. Brynjolfsson and his co-authors document an 8.6 percentage-point fall in agent attrition once the AI is deployed. The natural reading is not that workers love their newly upskilled jobs (the survey evidence is mixed) but that firms find it cheaper to retain newly-productive lower-cost workers than to bid up wages for stars whose marginal product has, relatively, fallen. A compressing technology, in firms' hands, becomes a wage-suppressing one. The wage pass-through is, for now, predicted rather than identified, but elevated PMET retrenchment in a 2% unemployment market is the kind of early signal that makes the prediction harder to ignore.

Finally, consider how much of this is even visible to a Singaporean policymaker reading official statistics. Household Gini is at a record low. Real wages are rising. The metric that would actually capture compression, dispersion of marginal product within job titles, is not produced by SingStat or anyone else. The IMF's 2024 Selected Issues Paper places 77% of Singapore's workforce in the high-AI-exposure band, against a 60% advanced-economy average, and warns that inequality will rise "in the absence of appropriate policies" (Khan, 2024).

3. Why Reskilling Cannot Reach the Distributional Frontier

Singapore has built one of the world's more capable workforce-development systems. SkillsFuture credits, the AI Apprenticeship Programme, TechSkills Accelerator, Career Conversion Programmes, and Budget 2026's subsidised AI tools share a common architecture. They award discrete credentials supporting transitions between occupations or skill levels. A bookkeeper retrain as a data analyst. A plumber adds solar accreditation. The credential travels with the worker, the labour market re-prices her, policy moves on.

That architecture, sensible for inter-occupation shifts, is the wrong shape for the intra-occupation problem. The Brynjolfsson gap is not between a customer-service agent and a software engineer. It is between two agents in the same firm, sharing a job description, separated only by how skilfully each prompts the same model. Closing it does not require a new credential. It requires patient, in-work investment in what some term complementarity capital, the day-to-day practice of integrating an AI tool into an existing function.

Three problems get in the way. First, timing: SkillsFuture courses are taken before or between jobs, but complementarity capital can only be built during work. Second, granularity: the unit of compression is the task, not the course. A litigator and a transactional lawyer with identical job titles face entirely different AI-exposure profiles, and SkillsFuture has no instrument that sees the difference. Third, incentive: a firm capturing the surplus from compression has no private reason to redistribute it through wages. Only collective action or credibly co-funded public policy can change that.

Singapore has faced this kind of problem before. The Progressive Wage Model was designed for cleaning, security, and landscaping precisely because those sectors faced collective-action failure on training and wage-setting. Its design, sectoral coverage, mandatory wage ladders, productivity benchmarks, is a more useful template for AI complementarity than the SkillsFuture menu. The PMET workforce now needs its own analogue: a framework that operates inside firms, at task level, with a measurable instrument of distribution.

NAIS 2.0's tripling of the AI talent pool to 15,000 by 2028, the National AI Council's sectoral missions, the AI Park at one-north, all aim at raising the ceiling of Singapore's AI capability. None reduces the gradient of AI complementarity within firms. David Autor's well-known counter-thesis, that AI can rebuild middle-class work by lowering the skill threshold for expert tasks (NBER, 2024), is not wrong. But Autor's optimism is conditional. It assumes complementarity is distributed broadly inside firms rather than allowed to concentrate. Distribution does not happen by itself.

4. The Capital-Share Corollary

Even if Singapore solved the intra-occupation problem, a wider distributional question would remain. Acemoglu (*Economic Policy*, 2025) projects cumulative AI-driven TFP gains of at most 0.66% over a decade, modest by industry hype but meaningful in compounding terms. His insight is not about the level of TFP but its incidence: even when AI raises low-skill productivity, inequality can rise if the labour share falls, because the gains accrue disproportionately to capital. Korinek and Suh (NBER, 2024) reach a similar conclusion: if AI is even a partial gross substitute for labour, the labour share declines almost regardless of how much reskilling a country does.

For Singapore this is not a distant theoretical concern. Compensation of employees runs around 41% of GDP; the underlying labour share, on the standard adjustment, is closer to 37%, well below the OECD range of 55 to 60 percent. The Atlantic Council's 2026 *Freedom and Prosperity* analysis traces this to Singapore's extensive-growth model: FDI-led, capital-deep, with private consumption only 36% of GDP. AI does not invent this dynamic, it accelerates it. A reskilling-only response leaves the distribution of national income permanently more dependent on capital ownership, exactly as the rents from AI capital are scaling.

The implication is uncomfortable but must be stated. The most direct distributional instrument available to Singapore in the long run is not labour-market policy. It is ownership policy. Temasek and GIC have already built exposure to the leading US foundation-model labs across recent capital raises. The relevant question is not whether sovereign wealth participates in the AI dividend, but whether that participation is

institutionally connected to the household balance sheet. At present it is not. CPF balances are determined by contribution rates and an administered interest rate, insulated from the same AI capital stack that may, in time, erode the wage component of those contributions.

5. Three Proposals for a Second Axis

Hence, I propose three solutions. Each adds an instrument the existing strategy lacks, and each draws on Singapore's own institutional line-up of regulator-led frameworks, FEAT in 2018, Veritas in 2019, MindForge in 2023, that turned distributional and risk problems into governable instruments.

A *Complementarity Audit* would establish a sector-by-sector reporting framework, modelled on the Progressive Wage Model but adapted to PMET work. Firms above a size threshold would publish, annually, three pieces of information: which job functions are being AI-augmented; what share of the resulting productivity gain flows to wages versus margins; and what training investment per worker has been made. Government co-funding under Champions of AI would condition on participation. The instrument is informational, not coercive, in the same spirit as Veritas in finance. Over time, this kind of reporting creates the basis for a Progressive AI Model: sectoral wage ladders linked to AI-augmented productivity benchmarks.

An *AI Capital Dividend* would ring-fence a small portion of Temasek's AI-stack returns, perhaps 0.25% of net portfolio value annually, and channel it into CPF MediSave or a dedicated AI Adjustment Account. This is not a tax on capital. It is a direct transmission of capital-share gains to citizen balance sheets, an extension of the logic Singapore already applies through the Net Investment Returns Contribution. Norway's Government Pension Fund and Alaska's Permanent Fund operate analogous mechanisms. The novelty is the explicit AI ring-fence, which serves as a hedge against the labour-share dynamic Acemoglu and Korinek model and which does not require any forecast about AGI to justify. This is

not a universal basic income. It is a specifically scoped citizen claim against a specifically identifiable risk: the wage component of CPF being eroded by the same AI capital stack the SWFs are invested in.

An *Intangibles-Adjusted Productivity Series* is the cheapest of the three. MAS and the Department of Statistics would publish, alongside headline TFP, a parallel series adjusted for intangible capital, following Brynjolfsson, Rock and Syverson (2021), whose method has been shown to revise US TFP upward by as much as 15.9%. The reason is political-economic, not academic. The political life of an AI investment programme is dominated by the J-curve: heavy intangible spending depresses measured productivity for years before the payoff arrives. Without a parallel series, the next government will face pressure to retrench precisely when persistence is most valuable. A credible adjusted series buys the kind of patience capital the MAS Macroeconomic Review buys for inflation expectations.

6. Conclusion

None of the above is a critique of NAIS 2.0. Singapore's positioning, "deploy AI effectively, responsibly, and at speed", is the right one for a city-state in 2026, and the regulatory and adoption infrastructure built around it (PathFin.ai, Project MindForge, AI Verify, Pair Chat now used by most of the public service) is genuinely world-class. The more modest claim of this essay is that NAIS 2.0 optimises one axis well, aggregate productivity, and that distribution requires its own dedicated stack of instruments rather than being assumed as a downstream consequence of growth.

Singapore is not going to win the race to build the largest frontier model. That race is closed to a country of 5.9 million. The race that remains open, and that matters more for political consent, is over which country distributes the gains from AI most credibly. An economy that cannot show its workers a fair share of the AI dividend will lose patience with AI strategy long before the J-curve turns. Inclusive growth is not, in that sense, the soft side of an AI strategy. It is the part that decides whether the rest survives politics.

References

- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121), 13–58.
- Atlantic Council. (2026). *Freedom and Prosperity Atlas 2026*. Atlantic Council Freedom and Prosperity Centre.
- Autor, D. (2024). *Applying AI to rebuild middle-class jobs*. NBER Working Paper No. 32140.
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942.
- Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372.
- Dell'Acqua, F., et al. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality*. Harvard Business School Working Paper 24-013.
- Department of Statistics Singapore. (2026, February). *Key Household Income Trends 2025*.
- Infocomm Media Development Authority. (2025). *Singapore Digital Economy Report 2025*.
- Khan, S. (2024). *Impact of AI on Singapore's labor market*. IMF Selected Issues Paper SIP/2024/040.
- Korinek, A., & Suh, D. (2024). *Scenarios for the transition to AGI*. NBER Working Paper No. 32255.
- Ministry of Digital Development and Information. (2023). *National AI Strategy 2.0: AI for the public good, for Singapore and the world*.
- Ministry of Manpower. (2026, March). *Labour Market Report Fourth Quarter 2025*.
- Monetary Authority of Singapore. (2025). *Proposed Guidelines on AI Risk Management* (consultation paper).
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- Smart Nation Singapore. (2024). *Smart Nation 2.0*.
- Wong, L. (2026, February 12). Budget 2026: Harness AI as a strategic advantage. Ministry of Finance.