


Automation anxiety

New AI technologies and employment outcomes in Malaysia



CERT

POLICY PAPER

Centre for Responsible Technology

The Centre for Responsible Technology (CERT) is an autonomous research centre established by the Institute of Strategic & International Studies (ISIS) Malaysia with the foundational support of the Malaysian Communications and Multimedia Commission. CERT's mission is to provide government and key stakeholders with the knowledge and insights needed to shape technological development guided by the principle of responsibility, ensuring that such development prioritises public interest, human rights, accountability and effective regulation.

CERT's research mandate is intersectional, focusing on the relationship between technology and four primary domains: society, sustainability, safety and security. Through research, dialogue and stakeholder engagement, CERT contributes to informed policy discussions on emerging technologies and their societal implications, working at the intersection of academia, industry and government.

[POLICY PAPER]

Automation anxiety

New AI technologies and employment outcomes in Malaysia

By

Calvin Cheng

Foreword by

Datuk Prof Dr Mohd Faiz Abdullah

Illustrated by

Ashikin Hussin

[AUTHOR]

Calvin Cheng is the director of the Economics, Trade and Regional Integration division at the Institute of Strategic & International Studies (ISIS) Malaysia. His research covers various topics in applied economics, centring around labour markets, wellbeing and the design of social transfer programmes. Some of his current projects examine how emerging technologies, such as generative AI, shape the future of work, wellbeing and inequality. He also leads ISIS Malaysia's economics research across its work areas and coordinates cross-team projects and external collaborations.

[ACKNOWLEDGEMENTS]

This paper benefited from the support of the PERKESO team, who generously provided data access, as well as helpful comments and assistance. The author is grateful to Allyson Wong for her assistance with the literature review and note-taking.

Foreword

Few questions animate the zeitgeist quite so completely as this one: what does the rise of artificial intelligence mean for humans who work? It is a question that has spread from academic journals to dinner tables with a speed that mirrors the technology under scrutiny. Indeed, the ubiquity of this question itself speaks to a growing collective anxiety, however inchoate, that the AI epoch may truly be different from the transformations of technologies past.

And yet, for all the commotion, the discourse on artificial intelligence and jobs has, until recently, been conducted largely on the terrain of ideology. On one hill stand the techno-evangelists, who extol the virtues of mega-tech corporations and “thinking machines” alike as humanity’s salvation, a *deus ex machina* for our myriad economic and social woes. On the opposing hill stand the catastrophists, who see in every new large language model an ever-larger threat of the obsolescence of human cognition. While both camps marshal their arguments with certitude, neither has been especially encumbered by empirical evidence.

Perhaps this is because the weight of evidence itself has been wanting. In its place, we have a deluge of anecdotes, consultancy reports, and statements from the AI companies themselves. While these sources are certainly not without value, they cannot by themselves constitute the evidential foundations upon which public policy must rest. For that, we require analysis and inquiry examining the material conditions of this change.

This is difficult work, owing to the novelty of generative AI technologies and the remarkable pace at which they evolve in both form and function. It is with this in mind that the present paper represents a significant first step towards an answer. Drawing on high-frequency administrative data from PERKESO and building on ISIS Malaysia’s sustained programme of future of work research, this study offers one of the first serious empirical investigations of the association between AI adoption and employment outcomes in the Malaysian context.

While the findings here are far from a final verdict — they are, if anything, a beginning — perhaps it is precisely what we need for the time being: enough to begin bridging the gap between the techno-optimists and the catastrophists, between the ideological and the material, and between theory and practice.

Datuk Prof Dr Mohd Faiz Abdullah

Executive Chairman

Institute of Strategic & International Studies (ISIS) Malaysia

Executive summary

- Loss-of-employment (LoE) claims in Malaysia have risen roughly three-fold since 2022, coinciding with the mainstream diffusion of generative AI tools. This rise is driven disproportionately by employer-initiated separations.
- Occupations with higher AI exposure account for the bulk of this increase. High-AI-exposed occupations — in tech and professional services, such as multimedia design — show markedly elevated LoE rates, while the rates for non-high-AI-exposure occupations have risen only modestly.
- Regression estimates suggest that jobs with the highest AI exposure experience around 69% more LoE than the sample average. Critically, jobs where AI is used predominantly for automation (task replacement) face about 1.5 times the average LoE rate, while jobs where AI is primarily augmentative are associated with lower separation rates. The AI gradient appears to be strengthening over time.
- Early-to-mid-career workers, tertiary-educated employees and women in AI-exposed occupations are more highly affected. These stratified patterns are consistent with prior theoretical works, as well as emerging global evidence on which socio-demographics are most vulnerable to AI-driven displacement.
- Although this work's results are associative, they provide early empirical evidence that generative AI tools are beginning to reshape Malaysia's labour market. While the effects are not yet broad-based, the strengthening trend could warrant proactive policy attention.
- In response, policymakers could double down on "human edge" skills and institutions, embed AI literacy into education without displacing foundational skills, and strengthen labour standards and fiscal incentives so that firms adopt AI in labour-complementary ways, while protecting workers least able to bear adjustment costs.
- Policy may need to move beyond "AI readiness" towards active management of realised labour market risks. This may mean updating PERKESO systems to monitor AI linked separations in highAI-exposure sectors in near real time, as well as providing workers with tools to identify skill adjacent transition pathways and rebuilding on the job experience pipelines for entry level roles, to preserve opportunities for youths to access mid- and senior-level knowledge-economy positions.

1. Introduction and literature review

Perhaps the first “generative AI” tool was ELIZA. Created in 1966, ELIZA could already simulate conversation through pattern matching and scripted responses in ways that foreshadow the rise of the modern chatbot.¹ Yet, for most people, having conversations with a machine would not become commonplace until the launch of ChatGPT in November 2022. ChatGPT thrust a new generation of AI technologies into mainstream public consciousness, and its diffusion was extraordinarily rapid. It quickly became the fastest-growing consumer app in history,² with back-of-the-envelope estimates suggesting that generative AI has spread orders of magnitude faster than past waves of technology, such as mobile and cloud computing.³ Indeed, few other technologies have managed to embed themselves so thoroughly in the textures of daily life that it is now increasingly difficult to go a single day without encountering — knowingly or otherwise — text, images, video or music generated by these new technologies.⁴

Recent labour market impacts of generative AI

The rapid diffusion of these technologies has also reignited deep-seated anxieties about job displacement and, more broadly, about the future of human labour. While its aggregate impacts have so far been modest,^{5,6} a growing body of empirical, theoretical and anecdotal evidence has indicated that AI indeed has a causal effect on employment, and that the weight of evidence suggests that these effects may tilt towards the negative.

Emerging empirical studies suggest that negative impacts are concentrated among younger and less-experienced workers, particularly in AI-exposed industries. Brynjolfsson et al. (2025), using an event study approach, estimated that young workers in occupations with the most AI exposure experienced a 16% relative employment decline after the launch of ChatGPT, while senior-level roles remained stable or grew.⁷ Hosseini and Lichtinger (2025),⁸ Klein Teeselink (2025)⁹ and Liu et al. (2025)¹⁰ demonstrated a similar pattern, finding entry-level and junior hiring in AI-adopting firms declining since 2022 and intensifying over time. While studies exist that do not detect aggregate effects on employment,^{11,12} on balance, the available evidence is consistent with modest, concentrated displacement effects beginning to emerge first in the most exposed occupations — potentially extending past waves of routine-biased technological change to new frontiers.

Whether new AI technologies will ultimately be a net “replacing” or “enabling” technology remains an open question.¹³ Where new tasks are emerging due to AI use, the picture is uneven. Humlum and Vestergaard (2026) identified new task creation in limited areas, such as content review, ethics and compliance, and AI integration, though the scale is limited. Marguerit (2025) found that augmentation AI is creating new work primarily among high-skilled workers — in turn raising their wages but not expanding aggregate employment — even as automation unambiguously reduces employment and wages.¹⁴ As to be discussed in Section 2.1, “augmentation” and “automation” are not mutually exclusive: augmentation may still lead to displacement in aggregate if productivity gains do not translate into expanded demand for labour.

Studies in the Malaysian context

While global studies have emerged in a global context, few focus on developing country contexts, and still fewer focus on Malaysia. In 2025, an Institute of Strategic & International Studies (ISIS) Malaysia–World Bank joint study constructed a generative AI exposure index for all 20,000-plus Malaysian job tasks at the Malaysian Standard Classification of Occupations (MASCO) 6-digit level. Linking this index to Labour Force Survey microdata, the study found that 28% of Malaysia's labour force are highly exposed to generative AI, and that AI exposure is higher for women, younger, and tertiary-educated workers.¹⁵ Similarly, TalentCorp (the Ministry of Human Resources' think tank) in its 2024 report, drawing on sector-level surveys, similarly found that about a fifth of formal-sector roles could be affected within the next three to five years.¹⁶

These studies are forward-looking by nature, in that they aim to estimate what could be at risk, not what has already been affected. Empirical evidence of realised AI impacts on the Malaysian labour market remains scant, likely due to data limitations and the novelty of this technology. This policy paper aims to provide early evidence that fills this empirical gap. Drawing from the Social Security Organisation's (PERKESO) individual-level loss-of-employment data in 2020–2025, I document the rise in loss-of-employment rates in Malaysia since 2022 and investigate if AI technologies played a role. The data, methodology and key limitations are described in Section 2. Section 3 presents the empirical analysis and main findings. Section 4 discusses the results and outlines their implications for policy design.

2. Data and methods

2.1 Theoretical framework

The paper adopts a task-based framework, in which technology affects labour demand by affecting the tasks that comprise an occupation.^{17,18} In the case of AI, at the task level, a task is “automated” when AI can perform it without human involvement, directly substituting for labour. On the other hand, a task is “augmented” when AI raises a worker's productivity on that task, but the human remains in the loop. AI can also have no effect on tasks — primarily for physical-manual tasks or those requiring physical embodiment.

These task-level effects aggregate to the macro-level as the net change in an occupation's optimal labour requirements through four channels, illustrated in Fig. 1. These can have four effects: displacement effect, efficiency effect, productivity effect and reinstatement effect.


The displacement effect operates as a result of task-level *automation*: when AI takes over tasks previously requiring human labour, firms need fewer workers to produce the same output. On the other hand, the efficiency effect channel is a result of task-level *augmentation*. When AI makes workers more productive on a task, fewer worker-hours are needed to produce a given quantity of output. Counterintuitively, both channels reduce the amount of human labour required to produce a

given level of output, just through different mechanisms: one removes humans from tasks entirely, while the other makes the human workers who remain more productive than necessary at previous employment levels.

These demand-reducing channels can be partially offset by two further effects (Fig. 1). The first of these is the productivity effect: as AI reduces costs, prices may fall and output demand may rise depending on how responsive demand is to price changes — in turn increasing demand for the labour still required in production. The second is the reinstatement channel, which operates over longer time frames. As some tasks are automated, new human-specific tasks could emerge at the frontier of what machines cannot yet do, as occurred historically with electrification and computerisation.¹⁹

Subsequently, in this framework, employers can adjust across four margins in response to changing optimal labour needs. Employers can reduce current workers, reduce new hiring or cut wages and/or hours. This paper focuses only on the first margin of adjustment: the effect of employers reducing current workers or separations (labelled “B” in Fig. 1) — specifically, the loss of employment among formal-sector incumbent workers as recorded in PERKESO’s Employment Insurance System (EIS). This means AI-driven labour demand reductions operating through other margins, such as reduced hiring or wage/hour cuts. The issue is discussed further in the limitations section (Section 2.3).

Fig. 1. Stylised theoretical framework of impact of technology on labour market outcomes



Source: Author’s illustrations, expanded from Autor et al. (2003) and Acemoglu and Restrepo (2018, 2019)

2.2 Data and hypotheses

The paper draws on three main data sources. The first is monthly administrative data provided by PERKESO, comprising individual-level loss of employment cases via the EIS system under Act 800. The data spans January 2020 through December 2025 and are disaggregated by key sociodemographic characteristics. While extensive in coverage, the dataset has several important limitations, which are outlined in the limitations section (Section 2.3) and again in the discussion section (Section 4). For the main analysis in this paper, I restrict the PERKESO sample to loss-of-employment cases that occurred due to employer-initiated separations (excluding cases that occurred due to voluntary resignation and business closure). This strengthens the correspondence to the separation margin of employer adjustment (margin “A” in Diagram 1), as these likely represent idiosyncratic and/or orthogonal processes to occupation-targeted demand changes from technology. Nonetheless, the full-sample reports are presented in Technical Appendix B for robustness.

The second source is an AI exposure index developed by Cheng, Chong, Dornan and Jasmin (2025), constructed to measure occupation-level exposure to generative AI. As per the strand of literature on ex-ante exposure indices, such as Gmyrek et al. (2023),²⁰ the index assigns automatability scores for all 20,000-plus tasks at the MASCO 6-digit level and then aggregates them to the occupation level. In this paper, these scores were used to identify the AI exposure level of the occupations identified in the PERKESO data — i.e., the degree to which the task bundle of an occupation overlaps with what generative AI can do. This enabled a gauge of whether employment losses were concentrated in AI-exposed occupations.

The third source is the Anthropic Economic Index (AEI).²¹ The AEI maps actual Claude usage to the task descriptions in the US Occupational Information Network (O*NET), enabling the classification of interactions as either “automation-style” or “augmentation-style”. This is then aggregated to produce occupation-level scores, representing the share of job tasks falling into each pattern. The AEI is used in this paper to complement the AI exposure index, treating the AEI’s *automation* score as the estimated fraction of an occupation’s tasks currently performed directly by AI tools, and the *augmentation* score as the corresponding share of AI use involving interactive, human-in-the-loop contexts, such as validation, iterative editing and learning. AEI scores are defined at the O*NET–SOC level, and I map these to MASCO occupations using an SOC–MASCO crosswalk by taking the unweighted mean across matched SOC codes in cases where a single MASCO occupation matched multiple O*NET categories.²² Crucially, by distinguishing automation-style from augmentation-style AI use, the AEI adds an independent second measure that maps more directly onto the automation/augmentation distinction in the framework (task-level Channels 1 and 2, respectively, in Fig 1).

Overall, these data sources serve to empirically investigate the theoretical framework in Section 2.1 and produces two hypotheses:

- H1** If AI is driving employer-initiated job losses, separation rates should be strongly positively associated with AI exposure.
- H2** Higher AEI-augmentation occupations should be a weaker or negative predictor of near-term separation rates relative to high AEI-automation occupations.

As suggested by the theoretical framework in Section 2.1, automation leads to direct task substitution through the displacement channel. Meanwhile, augmentation also reduces optimal labour demand but through the efficiency channel (Diagram 1), and as such is expected to do so primarily via reduced new hiring rather than increasing current separations over the short study window.

2.3 Limitations and caveats

Several limitations and caveats arise from the nature of the data sources. Firstly, the EIS system covers only formal private-sector employees, which could exclude the self-employed, public servants, gig workers and some in contract categories.²³ In addition, eligibility conditions and benefit caps for the Job Search Allowance may lead to under-reporting, particularly among higher-paid or younger workers. This introduces a selection process of unobserved loss-of-employment cases. To the extent that under-reporting is concentrated among groups more exposed to AI (as is likely the case for younger and higher-wage workers),²⁴ the estimates in this paper are likely biased downwards and thus may be interpreted as lower bounds.

Second, detailed occupational information (at the MASCO 4-digit level) is only available from 2023 onwards. This restricts my ability to link loss-of-employment observations to the AI exposure index to the post-2023 period, which also coincides with the rollout of generative AI tools, such as ChatGPT (launched in November 2022). Likewise, the absence of firm and individual identifiers prevents the use of standard causal inference strategies, such as difference-in-differences with firm-time fixed effects common in this literature.

Third, the EIS data captures only employment separations and does not allow the observation of other margins of labour market adjustment. As previously discussed in the theoretical framework section (Section 2.1), I only observe impacts that arise from formal-sector employment losses. If AI's effect on labour demand instead primarily comes from a reduction in new hiring of workers rather than creating separations of current workers, the data could show no rise in loss-of-employment cases even if there were an AI-driven contraction in labour demand. This is important because much of the existing evidence on AI's employment effects indeed appears to operate through slower hiring rather than an outright increase in separations.²⁵ Consequently, this could also cause the paper's analysis to understate the overall impact of AI on labour demand and therefore strengthen the interpretation of the estimates as lower bound of the true extent

of negative employment impact associated with AI tools.

2.4 Empirical strategy

Given the data limitations outlined, the paper's analysis is descriptive and does not seek to identify causal effects of AI on employment. Instead, I aim to document early evidence on whether AI exposure is associated with elevated employment losses in Malaysia's formal sector, and whether this association has strengthened since 2022. The PERKESO sample was restricted to loss-of-employment cases that occurred due to employer-initiated separations (e.g., downsizing, dismissals), excluding cases that occur due to voluntary resignation and business closures. Nonetheless, full-sample estimates are provided in Technical Appendix A.

My analysis proceeds in two parts. The first is descriptive. The paper documents trends in loss-of-employment rates and their composition across AI exposure groups from 2020 to 2025, covering the pre-ChatGPT baseline, the COVID-19 disruption period and the post-2022 period of rapid generative AI diffusion. Specifically, I use the Cheng et al. (2025) exposure index to define these groups, so that divergence in loss-of-employment trends across high-and low-exposure occupations constitutes a first descriptive check on the displacement channel.

The second part estimates the empirical association between AI exposure and loss-of-employment rates, via constructed panel of 4-digit MASCO occupation-month cells over 2023–2025. My baseline specification regresses the monthly loss-of-employment rates per 1,000 workers on the occupation-level AI exposure index, conditional on month-MASCO-1-digit fixed effects. This directly estimates the AI exposure gradient in the separation margin outlined in Section 2, with the coefficient of interest capturing the within-month, within-occupation-group average difference in loss-of-employment rates across 4-digit occupations with differing levels of AI exposure. A full set of alternative specifications are reported in Technical Appendix A. Additionally, I also estimate the AI gradient dynamically across different months, as well as across different demographic subgroups — by gender, education, and career stage — to characterise which workers in AI-exposed occupations bear the adjustment burden.

Finally, to further explore mechanisms by which AI affects employment, I complement the Cheng et al. (2025) task-based index with occupational-level measures derived from the AEI. This tests the second hypothesis from Section 2.2 that the displacement channel, rather than the augmentation channel, is the operative mechanism by which AI exposure links to employment losses in the short-term.

3. Analysis and results

3.1 Why are loss-of-employment rates rising?

The starting point of this paper is to document that the loss-of-employment rate has risen markedly since 2022 (Fig. 2). Since the start of data collection in 2020, loss-of-employment rates spiked during the first phases of the COVID-19 crisis and then flattened significantly in the post-pandemic phase in 2021 to 2022. However, since late 2022, loss-of-employment rates have risen consistently month-on-month and, as of December 2025, stands at approximately 8,000 monthly claims. This is three times higher than that in 2022 and happens to coincide with the launch of generative AI tools into the mainstream.

Alongside the rise in volumes, there has been a structural shift in the composition of separations: the post-2022 increase has been driven predominantly by employer-initiated separations (restructuring and redundancy) rather than by business closures or voluntary resignations. This is notable: even at the height of the COVID-19 pandemic, the composition was not dominated to this degree by employer-initiated events (Fig. 3).

Fig. 2. Number of LoE cases has steadily risen 3-fold since 2022





Fig. 3. Restructuring as share of LoE cases has risen



Source: ISIS Malaysia analysis of PERKESO data

Notes:

Employer-initiated separations exclude LoE cases due to “business closure”, “voluntary resignation” and “others”.
LoE = loss of employment as defined by PERKESO’s EIS system.


What could explain this steady upward trend? Could the rapid diffusion of generative AI tools play a role? In the theoretical framework in Section 2.1, this could possibly operate through the displacement and/or efficiency effects: firms initiating separations to reduce their labour inputs. Empirical and anecdotal evidence also suggests that this is plausible: the global literature reviewed in Section 1 suggests that this is indeed starting to occur in some economies.²⁶

To investigate this empirically, I use the AI exposure index to classify all 4-digit MASCO occupations (available from 2023 onwards) by automation score, and decompose loss-of-employment claims into those from “high-AI-exposed” occupations (the top quartile of AI exposure) and “non-high-AI-exposed” occupations (the bottom three quartiles). As shown in Fig. 4, the loss-of-employment rates for high-AI-exposed and non-high-AI-exposed occupations have diverged markedly since 2023: high-AI-exposed occupations show elevated and rising employment losses, while the losses for non-high-AI-exposed occupations have risen only modestly. This could suggest that the recent rise in loss-of-employment cases is not a broad-based phenomenon but, rather, disproportionately concentrated among occupations most exposed to generative AI.

Examining the top occupations and industries with the highest employment losses provides further corroborating detail. High-AI-exposed occupations are heavily over-represented in the top 10 industries by loss-of-employment claims (Fig. 5), and the leading industries are information and communication technology (ICT) and professional services (Fig. 6). These are among the highest-AI-exposed industries in the AI exposure index measure, and this pattern tracks closely with anecdotal observations of AI deployment in technology (e.g., usage of AI tools to write code and apps), multimedia and marketing (e.g., AI-generated copy, visual content, graphic design) and professional services (e.g., document drafting, data analysis).

However, as Fig. 6 shows, manufacturing — a sector not expected to be uniquely exposed to generative AI tools — also appears amongst the industries with top employment losses. This raises the concern that employment losses may be driven by factors unrelated to AI adoption. Certainly, the period under analysis (2022–2025) was one of exceptional turbulence, marked by major geopolitical ruptures and aggressive reshaping of tariff regimes. These factors risk confounding the signal in aggregate employment trends. As such, I turn to a more detailed analysis to attempt to overcome the issue in the following section.

Fig. 4. Rise in LoE cases after 2022 appears to be driven by high-AI-exposure occupations



Source: ISIS Malaysia analysis of PERKESO data and Cheng et al. (2025)

Notes:
Classification of high and low AI exposures used AI exposure index developed in Cheng et al. (2025). “High AI” group represents top quartiles of AI exposure, while “Non-high AI” represents bottom three quartiles of AI exposure. LoE = loss of employment as defined by PERKESO’s EIS system.

Fig. 5. High-AI-exposure jobs are over-represented in top LoE rankings





Fig. 6. ICT and professional services record highest LoE cases



Source: ISIS Malaysia analysis of PERKESO data and Cheng et al. (2025)

Notes:
Classification of high and low AI exposures used AI exposure index developed in Cheng et al. (2025). “High AI” group represents top two quartiles of AI exposure, while “Low AI” represents bottom two quartiles of AI exposure. Data for top 4 industries calculated via mean LoE cases per 1,000 workers. LoE = loss of employment as defined by PERKESO’s EIS system.

3.2 Does AI exposure predict loss-of-employment rates?

Baseline regressions

To more carefully answer the question of whether new AI technologies have driven employment losses, I estimate a series of regressions of loss-of-employment rates on AI exposure by constructing a 4-digit MASCO occupation by month panel. The baseline specification includes month-by-MASCO-1-digit fixed effects. I also explore alternative specifications, weighting schemes and clustering approaches, as presented in Technical Appendix A. The fixed-effects specification controls for time-invariant differences across occupations and for common month-to-month shocks within broader occupational groups. While residual confounding cannot be fully ruled out, I also explore alternative explanations, as presented in Section 3.4. Overall, I cautiously interpret the results as evidence consistent with the role of AI exposure in contributing to differences in loss-of-employment rates.

The baseline estimates, scaled by the sample-average loss-of-employment rate (1.15 per 1,000 workers), suggest that occupations with the highest AI exposure recorded approximately 69% more loss-of-employment cases than the sample average after accounting for occupational group differences and common macro shocks (Fig. 7). This is an economically and practically significant magnitude, suggesting that an occupation at maximal AI exposure is estimated to have two-thirds more job separations per worker than an average in the same major occupational group in the same month. This pattern suggests that AI exposure is shaping how labour market risks are distributed across occupations, rather than merely reflecting pre-existing differences between jobs or broader macroeconomic trends.

One important qualification is that this effect is sensitive to weighting. Weighting the baseline specification by occupational employment shares (such that larger-AI-exposed occupations contribute more to the estimated effect) reduces the estimated effect of AI exposure on employment loss (Fig. 6, second column). Specifically, the coefficient attenuates substantially to approximately 3% more loss-of-employment cases versus the average, as well as turning statistically insignificant (the right-hand-side bar in Fig. 7). I take this divergence as evidence that AI-related displacement is currently driven by smaller pockets of AI-exposed occupations rather than being broadly distributed across the workforce. The unweighted estimates capture whether AI displacement is occurring, while the weighted estimates capture how widespread this effect is for the average worker. This finding is consistent with both the global evidence base and certain theoretical frameworks, which predict that displacement effects initially concentrate in the most technologically substitutable task bundles before diffusing more broadly.²⁷

Automation vs. augmentation

Next, I also estimate the model jointly with job-level automation and augmentation scores from the AEI. This captures the average share of each occupation's AI-exposed tasks that are coded by

Anthropic as being automation-type workflows versus augmentative, human-in-the-loop workflows (see the data section above). In this specification, I estimate that occupations in the top decile of the AEI automation score exhibit roughly 1.5 times the sample-average loss-of-employment rate (higher than just AI exposure alone), whereas those in the top decile of the augmentation score exhibit about 1.3 times lower employment losses than the average.

Interpreted through the theoretical framework in Section 2.1, this result is directly consistent with the displacement channel. When a large share of an occupation's tasks is being performed by AI without human involvement (i.e., "automation" type), firms face reduced optimal labour demand for that occupation and adjust, in part, through reducing current workers. In addition, the framework predicts that augmentation also reduces optimal labour demand, only this time through the efficiency channel. As such, I interpret the relatively lower employment-loss rate as evidence that the two channels trigger different firm-level adjustment responses. Automation-style AI substitutes for specific tasks unambiguously, while augmentation-style AI raises worker productivity gradually, which firms may absorb through reducing future hiring and natural attrition — a margin that the data do not capture. This interpretation is consistent with recent evidence from the US suggesting that AI's primary effect on young workers operates through reduced labour market entry rather than elevated exits.²⁸

Dynamic effects

Lastly, I estimate the dynamic effect of AI exposure on loss-of-employment cases by re-estimating the baseline regression separately for each month in the sample (Fig. 9). The resulting sequence of coefficients on the AI exposure index shows an upward trend: while the impact of exposure on loss-of-employment is modest and imprecisely estimated in early 2023, it roughly doubles by 2025 and remains elevated thereafter. Although it should be cautioned that the monthly estimates remain somewhat imprecisely estimated, this pattern is nonetheless consistent with findings in other studies, such as Liu et al. (2025), which also find that AI's displacement effects grow stronger over time.²⁹

Fig. 7. Highest-AI-exposed jobs have about 69% more LoE than average, after accounting for occupation and monthly macro shifts




Fig. 8. The most “automatable” jobs have 1.5× more LoE than average, while the most “augmentable” have 1.3× less





Fig. 9. Impact of AI exposure on LoE cases appears to be rising over time



Source: ISIS Malaysia estimates using PERKESO data and Anthropic Economic Index

Notes:
 *Indicates statistical significance at 1% level based on wild bootstrap p-values. Baseline specification estimated with time-occupation group fixed effects clustered by month. See Technical Appendix A for further details on regression specifications and standard errors. “Automated” and “augmentable” scores are from AEI. LoE = loss of employment as defined by PERKESO’s EIS system.


3.3 Who are most affected?

Like many past waves of technological change, AI exposure is a deeply unequal phenomenon, and its employment impacts are likely to be no different. To explore this, baseline regressions were run on stratified samples separated by age, gender and educational attainment, and the result broadly suggests that loss-of-employment risks are not shared equally across demographic groups. Fig. 10 presents the estimated regression coefficients by career stage, education and gender. In terms of career stage, early-career workers (aged 25 to 34) in AI-exposed occupations experience approximately 25% higher employment loss risk than the sample average, while mid-career workers (aged 35 to 44) experience about 19% higher risk. Senior workers (aged 55 and above) show no statistically significant effect, and pre-career youth (aged 15 to 24) exhibit a non-statistically significant negative association between AI exposure and employment loss, likely reflecting low levels of EIS coverage and/or selection effects on the EIS in this age group.

In terms of education, the effects are particularly strong for degree-holders: AI-exposed jobs held by tertiary-educated workers show approximately a 78% higher loss-of-employment rate, compared with no significant effect for non-degree workers. This mirrors the nature of the AI exposure index, which disproportionately classifies knowledge-intensive, cognitive occupations (e.g., software development, professional services, ICT, and creative industries) as highly exposed. These are jobs typically held by degree-level workers.³⁰ In terms of gender, women workers in high-AI-exposed occupations are associated with an approximately 62% higher loss-of-employment risk than the sample average, while insignificant results were recorded for men. This gender asymmetry likely reflects the type of occupational sorting long-described in the literature:³¹ women in the formal sector are over-represented in service roles, which are both highly AI-exposed and highly automatable.

Fig. 10. Early-career workers, degree-holders and women may face higher AI-related job loss risk

Estimated regression coefficients vs. sample-average LoE rate (stratified samples, highlighted for statistical significance)



Source: ISIS Malaysia estimates using PERKESO data and Anthropic Economic Index

Notes:

*Data points highlighted for statistical significance at 5% level. Estimated separately with stratified samples with time-occupation group fixed effects clustered by month. Age groups: senior = 55+, late-career = 45–54, mid-career = 35–44, early-career = 25–34, pre-career = 15–24.

3.4 Alternative explanations

The results so far provide some early indications that new AI technologies may have been a driver of employment loss over the past few years. Nonetheless, I attempt to rule out three main alternative explanations for these findings.

First, amongst the most consequential external trade shocks since 2022 have been the US “Liberation Day” tariffs announced in April 2025, which saw the imposition then-24% tariffs affecting Malaysian exports to the US. As a highly trade-dependent economy, this could have driven some of the rise in loss of employment in Malaysia since then, particularly in manufacturing and industrial occupations. This could bias the AI coefficient downwards. That is, if tariff shocks disproportionately affect non-AI-exposed manufacturing or industrial occupations, this could raise loss-of-employment rates in lower-AI-exposed jobs and reduce the estimated AI effect. To assess this, the baseline was re-estimated on a subsample that excluded the period from post-April to December 2025. Here, my estimates remain broadly similar to the full-period results, suggesting that these results are not likely to be materially affected by tariff shocks (Technical Appendix B).

Second, global technology sector cycles, including layoff waves in multinational tech firms from 2023 to 2025, could have contributed to elevated employment losses in ICT and professional services independently of AI adoption. This could bias the AI coefficient upward, since the ICT sector is both high-AI-exposed and subject to tech-cycle volatility. This was investigated by re-estimating with the ICT sector (MSIC 1-digit) excluded. The direction and magnitude of these resulting estimates remain consistent, implying that these results were not purely driven by the tech sector (Technical Appendix B). Similarly, to rule out other industry-specific cyclical effects that are orthogonal to AI, I re-run the baseline regression with industry-by-month fixed effects. The results remain directionally consistent, if attenuated, as this specification absorbed a substantial share of variation that may be correlated with both AI exposure and employment dynamics within industries.

Third, to rule out the possibility that these results are driven by a transient post-COVID labour market adjustment rather than AI-induced displacement, I re-run the model on a sample restricted to 2024–2025. Malaysia’s wage subsidy, employment retention and other pandemic-era policies could have suppressed retrenchments and, as these programmes expired in 2023, it is plausible that this would have produced a spike in recorded employment losses. To the extent that AI-exposed industries were also disproportionate Wage Subsidy Programme (WSP) recipients (as could be the case in finance, business services, and administrative support), this clearing wave could inflate the estimated relationship between AI exposure and employment loss in samples during that period. Confining the sample to 2024 onwards yields figures that remain directionally consistent with the baseline, suggesting that the results were not due to artefactual spikes from policy changes.

4. Discussion

4.1 Summary of findings

First, I document a nearly three-fold rise in loss-of-employment cases since 2022, coinciding with the mainstream diffusion of generative AI tools. Using occupation-level AI exposure measures from Cheng et al. (2025), I decompose loss of employment cases into high-AI-exposed and non-high-AI-exposed groups and find that the post-2022 rise in loss of employment is disproportionately driven by high-AI-exposure occupations, while non-high-AI occupations show only a modest increase (Fig. 3). Top industries with the highest loss of employment cases are ICT and professional services, in line with anecdotal and ex-ante evidence of sectors most susceptible to generative AI exposure.

Second, regression estimates that controlled for occupational group differences and common macro shocks indicated that AI exposure is associated with approximately 69% higher employment losses than the sample average — a result cautiously interpreted as early evidence that generative AI is beginning to affect labour markets in Malaysia (Fig. 7). The AI gradient appears to be growing over time (Fig. 9), though it currently remains concentrated in pockets of most-exposed occupations rather than being broad-based: employment-weighted estimates are close to zero and statistically insignificant (Fig. 7), suggesting the average Malaysian worker may yet be materially affected.

Third, this effect is significantly amplified for occupations with higher AEI automation scores relative to augmentation scores. Jobs in the top decile of the automation distribution experience approximately 1.5 times the average employment loss rate, while the most augmentable occupations are associated with below-average employment losses (Fig. 8). I infer this result as automation-driven displacement being more immediately converted into employment losses, while augmentation-driven displacement may instead manifest through reduced new hiring instead of active separations, if not offset by productivity or reinstatement effects (see Section 2.1).

Fourth, I find that the burden of adjustment to the diffusion of AI technologies is uneven. Early-to-mid-career workers, tertiary-educated employees and women in AI-exposed occupations all record statistically significant increases in loss-of-employment cases versus the sample average (Fig. 10). These stratified patterns are consistent with prior theoretical work, as well as emerging global evidence on which socio-demographics are most vulnerable to AI-driven displacement.

4.2 Policy and economic implications

Overall, the findings in this paper lend some weight to predictions in the literature that suggest generative AI possibly extending displacement risks beyond routine manual and clerical work into a wider set of cognitive and professional occupations.³² As in earlier waves of technological change, the effects of generative AI are unlikely to be evenly distributed. Workers whose tasks are more easily

codified or reorganised, younger workers with less developed human-edge skills, and less statutorily protected workers (e.g., contract workers) in AI-exposed industries are more exposed to employment risks, while those in roles where AI complements human judgement may see weaker displacement effects or at least adjustment through slower hiring rather than outright separations.

For developing countries, such as Malaysia, the implications may be even more consequential. If AI's augmentation potential is greatest in advanced economies, while its automation potential is highest in developing and middle-income economies, as outlined in Cheng and Chong (forthcoming),³³ then Malaysia may face a less favourable balance of benefits and risks than many advanced economies, not least because its social protection and skills institutions are less mature and less able to absorb the adjustment costs of a rapid technological shock.

Of course, these negative near-term effects on workers could, over the longer term, be offset by the countervailing productivity and reinstatement channels that the theoretical framework identifies (see Section 2.1). Whether or not this will occur remains an open empirical question. Nonetheless, a few considerations currently temper my optimism about the distributional and economic implications of this possibility. To start, the empirical evidence on both offsetting channels is, to date, relatively modest. On the productivity channel, as outlined in Acemoglu (2024), AI's aggregate productivity effects are likely to be limited in scope because the share of tasks where AI generates meaningful cost reductions remains narrow, and early macroeconomic data have yet to register the productivity gains that would be needed to substantially expand labour demand³⁴ — a redux of the “productivity paradox” that have puzzled many in the decades since the advent of the computer age.³⁵

Likewise, on reinstatement and new task creation, Acemoglu and Restrepo (2019) documented that the rate of new task creation has slowed markedly relative to the pace of automation in the pre-AI era, suggesting the structural conditions for rapid reinstatement could be weakening even before generative AI arrived.³⁶ More directly, Humlum and Vestergaard (2026) identified new task creation in the wake of AI adoption but found it concentrated in various areas, such as content review, ethics and compliance, and AI integration and oversight, potentially at a scale that falls well short of offsetting displacement in occupations most affected.³⁷

Finally, even successful aggregate adjustment and reinstatement do not guarantee good outcomes for workers or for economies writ large. Braxton and Taska (2023) showed that workers displaced by past waves of technological change who transition to new occupations typically do so at a permanent wage penalty, with the displaced workers rarely recovering the earnings trajectory they would have followed absent displacement.³⁸ This could mean that even under the most optimistic aggregate scenario, in which productivity and reinstatement effects eventually expand total labour demand, workers who lose jobs in the near term are structurally unlikely to be the primary beneficiaries of those longer-run gains.

5. Policy recommendations

In this section, I outline some brief policy recommendations motivated by the findings. As discussed prior, the analysis in this paper remains descriptive associations, and I do not interpret these estimates as causal. Nonetheless, the findings establish a pattern of occupational and demographic concentration in AI-driven separations that is sufficiently clear to motivate policy attention.

First, it is worth restating some of the recommendations set out in the ISIS Malaysia–World Bank report, which remain relevant in relation to AI-driven employment impacts.³⁹ On skills and education, the report recommended mainstreaming AI literacy into national school curricula in a way that reinforces rather than replaces foundational numeracy, language and critical thinking. In other words, a renewed focus on building human-edge competencies — skills that, for now, humans retain a comparative advantage in performing — which the report identifies as commanding rising wage premiums as AI adoption progresses. At the tertiary level, the report recommends building a national stackable micro-credential framework anchored to Malaysian Qualifications Agency guidelines so that workers can top up specific competencies without abandoning the wider intellectual and social benefits of higher education.

On labour market institutions, the ISIS Malaysia–World Bank report also suggested increasing job quality across the board through strengthening labour market institutions, standards and enforcement, while using fiscal tools to rebalance the relative cost of labour versus capital investment to make it cheaper for employers to invest in workers' skills. Likewise, stronger incentives for firms to adopt AI in labour-complementary rather than labour-substituting ways could begin to tilt these incentives in a more constructive direction. Taken together, these recommendations are designed to maximise the benefits of AI diffusion, while ensuring that adjustment costs are not borne disproportionately by workers least equipped to absorb them.

Second, the new empirical results above point to several additional priorities that go beyond general preparedness. This may mean that the policy imperative needs to shift from just encouraging AI adoption and upskilling to a more pragmatic and active management of realised labour market risks. Current national AI plans and governance frameworks have focused primarily on accelerating AI adoption and capturing productivity gains.⁴⁰ The findings here suggest these institutions may need to be updated to include explicit, adequately resourced provisions for managing adjustment costs among displaced workers, particularly in high-AI-exposure sectors, such as ICT, professional services and advertising.

Likewise, the concentration of rising loss-of-employment rates in automation-oriented, high-AI-exposure occupations also strengthens the case for targeted early-warning systems. PERKESO already has granular, real-time administrative data on separations by occupation. Combining this specifically with measures of AI exposure as per the index in this paper could allow policymakers

to detect emerging displacement pressure earlier and track whether the high-AI-exposure quartile divergence documented above continues to widen.

The flipside of early-warning systems for policymakers is an accessible intelligence tool for workers most directly affected. Displacement pressure is most actionable at the individual level, when workers can identify not only that their occupation is at risk but where feasible transition pathways lead and what skills gaps stand between their current role and more insulated alternatives. In this vein, open-access tools, such as Chong's (2026) pilot skills-based occupation network explorer (see cert.isis.org.my/skill-pathways-tool), could allow workers, jobseekers and career advisors to query the AI exposure of specific occupations, map skill adjacencies across the occupational network and identify transition pathways away from highly automated roles towards occupations with lower displacement risk and sufficient skills overlap to make retraining feasible.

Third, policymakers could also invest in rebuilding on-the-job experience pipelines that entry-level workers in AI-exposed occupations are losing. The concentration of AI-related loss of employment among early-career workers points to a structural risk in Malaysia's youth human capital pipeline. Tertiary graduates are being trained for occupations that are among the most AI-exposed (such as in software development, advertising, financial analysis, and policy and planning) and consequently may increasingly risk being displaced before they can accumulate the experience and tacit knowledge needed to transition into senior, less-automated roles that sit upstream in the same occupational ladders.

If the entry-level rungs of knowledge-economy careers are progressively hollowed out by AI, the pipeline feeding experienced mid- and senior-level professionals in these fields will thin over time, with long-run consequences for productivity and organisational capability that are not yet visible (nor perhaps even knowable). Addressing this could entail a combination of firm-level incentives to preserve graduate development programmes in high-AI-exposure sectors — potentially integrated with the Progressive Wage Policy — and/or structured apprenticeship and mentoring mechanisms that accelerate tacit knowledge transfer between experienced and entry-level workers.

To conclude, these findings suggest that Malaysia is unlikely to face a mass technology-driven jobs apocalypse for now, but it would be equally misguided to assume that the long-term equilibrium will be benign by default. Rather, the impacts of new AI technologies on the future of work will be far-reaching and wide-ranging, experienced as a series of cumulative shocks concentrated in specific occupations, sectors and demographic groups. Over time, these quantitative changes may well culminate in qualitative transformations in the future of human labour — though whether these shifts ultimately prove beneficial or harmful for human flourishing will yet depend on the policy, political and institutional choices made in the years to come.

Endnotes

- ¹ He, R., Cao, J., & Tan, T. (2025). Generative artificial intelligence: A historical perspective. *National Science Review*, 12(5), nwaf050. <https://doi.org/10.1093/nsr/nwaf050>
- ² Berg, A. & Ho, A. (2025). After the ChatGPT moment: Measuring AI's adoption. Epoch.ai. <https://epoch.ai/gradient-updates/after-the-chatgpt-moment-measuring-ais-adoption>
- ³ Cheng, C., Chong, H., Dornan, M., & Jasmin, A. F. (2025). Novel AI Technologies and the Future of Work in Malaysia [Policy brief]. Institute of Strategic & International Studies (ISIS) Malaysia. <https://www.isis.org.my/wp-content/uploads/2025/08/Novel-AI-technologies-PB.pdf>
- ⁴ Gomez Sarmiento, I. (2026, May 2). AI music is flooding streaming platforms. But listeners like it less and less. NPR. <https://www.npr.org/2026/05/02/nx-s1-5804489/music-listeners-dislike-ai-music-study>
- ⁵ Hampole, M., Papanikolaou, D., Schmidt, L. D. W., & Seegmiller, B. (2025). Artificial Intelligence and the Labor Market (NBER Working Paper No. 33509). National Bureau of Economic Research (NBER). <https://doi.org/10.3386/w33509>
- ⁶ Acemoglu, D. (2024). The Simple Macroeconomics of AI (NBER Working Paper No. 32487). National Bureau of Economic Research. <https://doi.org/10.3386/w32487>
- ⁷ Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Stanford Digital Economy Lab. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>
- ⁸ Hosseini Maasoum, S. M. & Lichtinger, G. (2025). Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5425555
- ⁹ Klein Teeselink, B. (2025). Generative AI and Labor Market Outcomes: Evidence from the United Kingdom. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5516798
- ¹⁰ Liu, Y., Wang, H., & Yu, S. (2025). Labor Demand in the Age of Generative AI: Early Evidence from the U.S. Job Posting Data (Policy Research Working Paper No. 11263). World Bank. <https://doi.org/10.1596/1813-9450-11263>
- ¹¹ Humlum, A. & Vestergaard, E. (2026). Still Waters, Rapid Currents: Early Labor Market Transformation Under Generative AI (NBER Working Paper No. 33777). National Bureau of Economic Research. <https://doi.org/10.3386/w33777>
- ¹² Johnston, A. C. & Makridis, C. (2025). AI, Output, and Employment. SSRN. <https://ssrn.com/abstract=5375017>
- ¹³ Frey, C. B. (2018). *The Technology Trap: Capital, Labor, and Power in the Age of Automation*. Princeton University Press.
- ¹⁴ Marguerit, D. (2025). Augmenting or automating labor? The effect of AI development on new work, employment, and wages. arXiv, arXiv:2503.19159. <https://arxiv.org/abs/2503.19159>
- ¹⁵ Cheng et al. (2025).
- ¹⁶ Talent Corporation Malaysia Berhad. (2024). Impact Study of Artificial Intelligence, Digital, and Green Economy on the Malaysian Workforce: Volume 2 — Sector: Information and Communications Technology. <https://www.talentcorp.com.my/images/uploads/publication/199/Impact-Study-of-Artificial-Intelligence-Digital-and-Green-Economy-on-the-Malaysian-Workforce-Volume-2-Sector-Information-and-Communications-Technology-1747879475.pdf>
- ¹⁷ Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355503322552801>

- ¹⁸ Acemoglu, D. & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542. <https://doi.org/10.1257/aer.20160696>
- ¹⁹ Frey (2018).
- ²⁰ Gmyrek, P., Berg, J., & Bescond, D. (2023). Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality (ILO Working Paper No. 96). International Labour Organization. <https://doi.org/10.54394/FHEM8239>
- ²¹ Massenkoff, M., Lyubich, E., McCrory, P., Appel, R., & Heller, R. (2026, March 24). Anthropic Economic Index Report: Learning Curves. <https://www.anthropic.com/research/economic-index-march-2026-report>
- ²² Ibid.
- ²³ Social Security Organisation (n.d.). Employment Insurance (Lindung Kerjaya). <https://www.perkeso.gov.my/en/our-services/protection/employment-insurance.html>
- ²⁴ Cheng et al. (2025).
- ²⁵ Brynjolfsson et al. (2025).
- ²⁶ Thomas, D. (2026, May 9). Companies name AI as top reason for job cuts for second straight month: Analysis. Yahoo! News. <https://news.yahoo.com/companies-name-ai-top-reason-134642138.html>
- ²⁷ Gmyrek, P., Viollaz, M., & Winkler, H. (2026). Disruption Without Dividend? How the Digital Divide and Task Differences Split GenAI's Global Impact (Policy Research Working Paper No. 11328). World Bank. <https://doi.org/10.1596/1813-9450-11328>
- ²⁸ Brynjolfsson et al. (2025).
- ²⁹ Liu et al. (2025).
- ³⁰ Cheng et al. (2025).
- ³¹ Ngai, L. R. & Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4), 1–44. <https://doi.org/10.1257/mac.20150253>
- ³² Brollo et al. (2024). Broadening the Gains from Generative AI: The Role of Fiscal Policies (Staff Discussion Note No. SDN/2024/002). International Monetary Fund. <https://doi.org/10.5089/9798400277177.006>
- ³³ Cheng, C. & Chong, H. (forthcoming). AI and the Future of Work in ASEAN and East Asia. Economic Research Institute for ASEAN and East Asia (ERIA).
- ³⁴ Acemoglu (2024).
- ³⁵ Polák, P. (2017). The productivity paradox: A meta-analysis. *Information Economics and Policy*, 38, 38–54. <https://www.sciencedirect.com/science/article/abs/pii/S0167624516301524>
- ³⁶ Acemoglu, D. & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- ³⁷ Humlum & Vestergaard (2026).
- ³⁸ Braxton, J. C. & Taska, B. (2023). Technological change and the consequences of job loss. *American Economic Review*, 113(2), 279–316. <https://doi.org/10.1257/aer.20210182>

³⁹ Cheng et al. (2025).

⁴⁰ Chong, H. (2026). Skill Pathways for Technology-Driven Occupation Transitions: The Case for an Occupation Network Explorer [Research note]. Centre for Responsible Technology.

Technical Appendix

A Full regression tables

A1 Main regressions

Table A1

	Full Sample				
	M1 Baseline	M2 + Salary	M3 Emp.wtd	M4 Crossed	M5 4D clust
AI exposure index	0.809*** (0.149)	0.666*** (0.150)	0.018 (0.058)	0.813*** (0.152)	0.809 (0.774)
lg (Mean salary)		0.000*** (0.000)			
N	6520	6520	6520	6504	6520
Within R ²	0.075	0.082	0.097	0.085	0.075
Fixed effects	month + 1D	month + 1D	month + 1D	month + 1D	month + 1D
p (WCB)	0.000***	0.002**	0.344	n.a.	n.a.
Cluster	Month	Month	Month	Month	4-digit occupation

Table A1r

	Employer-initiated separations only				
	M1 Baseline	M2 + Salary	M3 Emp.wtd	M4 Crossed	M5 4D clust
AI exposure index	0.792*** (0.104)	0.720*** (0.105)	0.033 (0.042)	0.798*** (0.107)	0.792 (0.499)
lg (Mean salary)		0.000*** (0.000)			
N	5416	5416	5416	5395	5416
Within R ²	0.110	0.114	0.109	0.123	0.110
Fixed effects	month + 1D	month + 1D	month + 1D	month + 1D	month + 1D
p (WCB)	0.000***	0.000***	0.111	n.a.	n.a.
Cluster	Month	Month	Month	Month	4-digit occupation

Notes:

1) *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$

2) Standard errors in parentheses: HC1 throughout; M5/M5r used 4-digit MASCO occupation-clustered SEs.
Outcome: monthly loss-of-employment (LoE) events per 1,000 LFS 2021 workers. Sample: 4-digit MASCO occupations, January 2023–December 2025.

3) p (WCB): wild cluster bootstrap p -value (Roodman et al. (2019)), clustered by month ($G = 36$), Rademacher weights, null imposed, $B = 999$. n.a. = M4/M4r crossed FE not supported by WCB; M5/M5r used 4D-clustered SEs (WCB not required)

All models are estimated on a panel of 4-digit MASCO occupation-month cells over January 2023–December 2025. The dependent variable is monthly loss-of-employment events per 1,000 workers, where the employment denominator is drawn from the LFS 2021 count (not time-varying). Table A1 presents estimates on the full LoE sample. Table A1r replicates Table A1 on employer-initiated separations only (restructuring and redundancy). Tables A2 and A3 report stratified and alternative index estimates, respectively.

Main specifications

I estimated five specifications. M1 (Baseline) is the preferred specification used throughout the main text: it regresses LoE per 1,000 on the AI exposure index with month and one-digit occupation-group (MASCO 1D) fixed effects. M2 (+Salary) adds log mean monthly salary as a covariate, testing whether the AI gradient is driven by an omitted wage channel, though the study did not include this in the baseline specification due to caution of overfitting, since salary is contemporaneously determined with the treatment instead of being a pre-treatment covariate.

M3 (Emp.wtd) re-weights observations by the 2021 LFS employment share, so that larger occupations contribute proportionally more to the estimated effect. This tests whether the gradient holds across the employment-weighted distribution and answers the question of: “what is the estimated effect for the average worker in Malaysia?”.

M4 replaces the additive month + 1D structure with a fully interacted month $\times$ 1D fixed-effects specification, absorbing all occupation-group-specific time trends. Identification is then solely from cross-occupation AI variation within each occupation group and month. M5 (4D clust) retains the M1 structure but clusters standard errors at the 4-digit MASCO level rather than by month.

Interpreting the coefficient of interest

The AI exposure index is time-invariant at the 4-digit MASCO occupation level by construction, as it is built from task-level characteristics of occupations that do not change over the sample window. This means four-digit occupation fixed effects cannot be included — they would be perfectly collinear with the regressor of interest. Month fixed effects absorb all common macro shocks. Consequently, the identifying variation is therefore entirely cross-sectional, i.e., within a given one-digit occupation group and a given month, do four-digit occupations with higher AI exposure record higher LoE rates? I interpret the resulting coefficient as a conditional AI gradient rather than as a causal effect.

A one-unit change in the AI exposure index corresponds to the theoretical maximum: a shift from a fully non-exposed occupation (index = 0) to a fully exposed one (index = 1). This range is large relative to the empirical support of the index. Throughout the main text, coefficients are expressed as a percentage of the full-sample mean monthly LoE rate (1.15 per 1,000 workers), which provides

an economically interpretable benchmark. The baseline estimate (M1) indicates that a one-unit increase in the AI exposure index is associated with an additional 0.809 LoE events per 1,000 workers per month in the full-sample specification (SE = 0.149; WCB p = 0.000). This corresponds to approximately 70% of the sample-mean monthly LoE rate of 1.15 per 1,000 workers. For employer-initiated separations, the gradient is somewhat sharper in statistical precision (M1r: β = 0.792, SE = 0.104, WCB p = 0.000), reflecting lower residual variance in the restricted outcome variable, and the point estimate is similar in magnitude.

Adding log mean salary as a control (M2/M2r) attenuates the AI coefficient modestly (from 0.809 to 0.666 for the full sample) but it remains highly significant (WCB p = 0.002), indicating that the gradient is not primarily an artefact of the salary-AI exposure correlation. The employment-weighted specification (M3/M3r) produces a coefficient close to zero (β = 0.018, SE = 0.058; p = 0.344), consistent with AI displacement being concentrated in a subset of occupations rather than broadly distributed across the employment-weighted occupational distribution. The crossed fixed-effects specification (M4/M4r) delivers a coefficient essentially identical to the baseline (β = 0.813, SE = 0.152), suggesting that occupation-group-specific time trends are not responsible for the estimated gradient. Finally, clustering at the 4-digit level (M5/M5r) substantially widens the standard error (SE = 0.774 for M5; SE = 0.499 for M5r) without changing the point estimate, reflecting the small-cluster problem at this level of aggregation (181 four-digit occupations). The coefficient remains directionally consistent, though not significant, at conventional thresholds.

A note on standard errors

The panel has 36 monthly clusters (January 2023–December 2025). With a relatively small number of clusters, conventional cluster-robust standard errors can be biased upwards. As such, I report wild cluster bootstrap p -values (p (WCB)) using the `fwildclusterboot` R package (Fischer and Roodman, 2021), with B = 999 bootstrap replications and Rademacher weights, imposing the null hypothesis. Most significance stars in the main text are based on WCB p -values. Notably, WCB is not available for M4 (crossed FEs are not supported by the bootstrap procedure) or M5 (which already uses a more conservative cluster structure at the 4-digit level, making WCB redundant). These are noted as “n.a.” in the tables.

A2 Stratified subsample regressions

Table A2

Subsample	β	SE	Sig.	p (HC1)	% of avg. LoE
Female	0.708	0.086	***	0.000	61.7%
Male	-0.129	0.124		0.301	-11.2%
Degree (Batch+)	0.891	0.104	***	0.000	77.6%
Non-degree	-0.039	0.113		0.731	-3.4%
Pre-career	-0.110	0.067	*	0.099	-9.6%
Early-career	0.281	0.085	***	0.001	24.5%
Mid-career	0.216	0.064	***	0.001	18.9%
Late-career	0.072	0.051		0.157	6.3%
Senior	0.010	0.025		0.680	0.9%

Notes:

1) *: $p < 0.1$, **: $p < 0.01$, ***: $p < 0.001$

2) Baseline specification as per Table A1 estimated separately within each demographic category.

Employment denominator: full LFS 2021 (not stratum-specific); % of avg. LoE: coefficient as a percentage of the full-sample mean monthly LoE rate (1.15 per 1,000); HC1 SEs throughout.

Table A2 reports the baseline specification (M1: $\text{loe_per1k} \sim \text{AI_exposure_index} \mid \text{month_id} + \text{MASCO_code1D}$) estimated separately within each of nine demographic strata, constructed by filtering the joint PERKESO–LFS panel to the relevant subpopulation prior to building the occupation-by-month panel. The employment denominator throughout the full LFS 2021 count (not stratum-specific), and so coefficients are interpretable on a common scale and can be compared directly against the full-sample benchmark.

Gender: The AI gradient is positive and highly significant for women ($\beta = 0.708$, $\text{SE} = 0.086$; $p < 0.001$), corresponding to approximately 62% of the average LoE rate. For men, the coefficient is negative, small, and not significant ($\beta = -0.129$, $\text{SE} = 0.124$; $p = 0.301$). This marked gender asymmetry points to a pattern of occupational sorting: women in Malaysia’s formal sector are over-represented in administrative, clerical and professional support roles, which are both highly AI-exposed and highly automatable, while men are more broadly distributed across a range of exposure levels, diluting the gradient. The finding has direct implications for gender-targeted skills and social protection policy.

Education: Among degree-holders (bachelor’s degree and above), the AI gradient is strongly positive and significant ($\beta = 0.891$, $\text{SE} = 0.104$; $p < 0.001$), representing approximately 78% of the

sample-average LoE rate. Among non-degree workers, the coefficient is indistinguishable from zero ($\beta = -0.039$, $SE = 0.113$; $p = 0.731$). This contrast is consistent with the structure of the AI exposure index, which assigns high exposure scores primarily to knowledge-intensive, cognitive occupations, i.e., ICT, professional services and creative industries, which are predominantly staffed by tertiary-educated workers. Non-degree occupations in the Malaysian formal sector tend to be in manufacturing, construction, and trade — all sectors where AI exposure is lower and where the current displacement channel is less active.

Age group: Early-career workers (25–34) show a positive and significant AI gradient ($\beta = 0.281$, $SE = 0.085$; $p = 0.001$), equivalent to approximately 24% of the mean LoE rate. Mid-career workers (35–44) also show a positive and significant effect ($\beta = 0.216$, $SE = 0.064$; $p = 0.001$), equivalent to about 19%. Late-career workers (45–54) show a positive but statistically insignificant coefficient ($\beta = 0.072$, $SE = 0.051$; $p = 0.157$). Senior workers (55+) show no meaningful effect ($\beta = 0.010$, $SE = 0.025$; $p = 0.680$). Pre-career youth (15–24) exhibit a borderline-negative coefficient ($\beta = -0.110$, $SE = 0.067$; $p = 0.099$, marginal at the 10% level), which likely reflects the low LFS-based employment denominator and thin EIS coverage for this age group rather than genuine insulation from AI displacement. The age gradient is broadly consistent with global evidence that AI displacement is currently most acute among early-to-mid-career workers who hold the routine cognitive and entry-level tasks most susceptible to automation.

A3 Anthropic Economic Index (AEI) scores

Table A3

	Baseline (AEI)	+AI exposure
AI exposure index		0.950*** (0.159)
Automation score (AEI)	1.751*** (0.302)	1.341*** (0.309)
Augmentation score (AEI)	-1.469*** (0.250)	-1.513*** (0.249)
N	6184	6184
Within R ²	0.075	0.081
Fixed effects	month + 1D	month + 1D
p (WCB) — auto.	0.000***	0.000***
p (WCB) — augment.	0.000***	0.000***

Notes:

1) ***: $p < 0.001$

2) M_auto: automation and augmentation scores as regressors (employer-initiated LoE sample); M_auto2: added Cheng &Chong's (forthcoming) AI exposure index. Scores from the Anthropic Economic Index (March 2025 release): automation = directive + feedback_loop; augmentation = task_iteration + validation + learning. Standard errors in parentheses: HC1 throughout; M5/M5r used 4-digit MASCO occupation-clustered SEs. Outcome: monthly loss-of-employment (LoE) events per 1,000 LFS 2021 workers. Sample: 4-digit MASCO occupations, January 2023–December 2025. 3) p (WCB): wild cluster bootstrap p-value (Roodman et al.(2019)), clustered by month (G = 36), Rademacher weights, null imposed, B = 999. n.a. = M4/M4r crossed FE not supported by WCB; M5/M5r used 4D-clustered SEs (WCB not required)

Table A3 supplements the Cheng et al. (2025) task-based AI exposure index with occupation-level automation and augmentation scores from the Anthropic Economic Index (AEI) by Massenkoff et al. (2026). AEI scores are derived from empirical Claude usage patterns matched to O*NET task descriptions: the automation score captures the share of an occupation’s tasks currently performed in directive, high-autonomy AI workflows, while the augmentation score captures the share performed in human-supervised, iterative workflows such as validation, editing, and learning. Both scores are defined at the O*NET–SOC level and mapped to MASCO four-digit occupations via a SOC–MASCO crosswalk; where multiple SOC codes match a single MASCO code, I take the unweighted mean.

The baseline AEI specification includes only automation and augmentation scores as regressors, with month and 1D fixed effects retained. The automation score enters positively and significantly ($\beta = 1.751$, $SE = 0.302$; WCB $p < 0.001$): occupations where a higher share of tasks is currently performed through automation-style AI workflows record markedly elevated LoE. The augmentation score enters negatively and significantly ($\beta = -1.469$, $SE = 0.250$; WCB $p < 0.001$): occupations where AI is used primarily in collaborative, human-supervised contexts are associated with below-average employment losses. This decomposition is consistent with the theoretical channels outlined in Section 2.1 and the hypotheses in Section 2.2, i.e., that automation operates through the displacement channel, while augmentation does not.

The extended specification adds the Cheng et al. (2025) AI exposure index jointly. The automation and augmentation coefficients remain essentially unchanged in magnitude and significance (automation: $\beta = 1.341$, $SE = 0.309$; augmentation: $\beta = -1.513$, $SE = 0.249$), while the coefficient on AI exposure index is positive and significant on its own ($\beta = 0.950$, $SE = 0.159$). The stability of the AEI coefficients upon controlling for AI exposure suggests that the automation/augmentation decomposition captures distinct dimensions of AI-related task substitutability that are only partially correlated with the task-based exposure index.

B Robustness and alternative explanations


In Section 3.4, three principal alternative explanations for the observed AI-loss-of-employment gradient were discussed: (1) the US *Liberation Day* tariff shock of April 2025, (2) the global technology-sector cycle of 2023–2025 and (3) industry-specific business cycles and demand fluctuations. I assess robustness against each of these threats, plus a fourth check on the sensitivity of results to the 2023 deferred-claims accumulation period. The table below presents robustness specifications R0–R4 alongside the full-sample baseline.

R0 is the full sample baseline for reference. R1 excludes the post-April 2025 subsample, specifically, it retains only January 2023–March 2025, to test whether the results are driven by the tariff shock announced on 2 April 2025 (which imposed a then-24% tariff on Malaysian exports). R2 excludes MSIC Section J (Information and Communication) to test whether the gradient is an artefact of the

tech-sector layoff cycle that simultaneously elevated AI exposure scores and sector-specific employment losses in ICT. R3 replaces month fixed effects with industry-by-month fixed effects, absorbing all industry-specific time-varying shocks. Identification is then from within-industry, within-1D-group, cross-occupation AI variation. R4 restricts the sample to January 2024 onwards, removing the 2023 deferred-claims accumulation period in which early-period backlogs of COVID-era cases from the expiration of pandemic-era employment retention and support policies may have created artificial employment loss spikes.

Spec	Details	Key question
R0	Full-sample baseline (=M1 above)	Reference point
R1	Excluded April-December 2025	Are results driven by the Trump “Liberation Day” tariff shock (24% tariff on Malaysian exports)?
R2	Excluded ICT sector	Are results an artefact of the global tech cycle that simultaneously raises AI exposure scores and sector specific employment losses?
R3	Industry x month fixed effects	Do industry-specific business cycles, commodity-price shocks, or export-demand fluctuations explain the AI gradient?
R4	2024–2025 sub-period only	Did the expiration of COVID-era policies, such as employment retention and wage subsidy programmes in 2023, bias/ influence the findings?

Coefficient as % of sample mean LoE rate: 95% CI (HC1 SEs)



Results: The main finding is relatively robust across all four checks. Excluding the tariff-shock period (R1) attenuates the coefficient somewhat, from 0.809 to 0.603 for the full sample, and from 0.792 to 0.600 for employer-initiated separations, but both remain highly significant (WCB not reported for robustness specifications, instead HC1 p-values are reported). This confirms that the April 2025 tariff shock does not drive the main result, which is consistent with the directional prediction that tariff shocks primarily affect non-AI-exposed manufacturing occupations, which would compress the estimated AI gradient. Excluding the ICT sector (R2) produces coefficients essentially identical to the baseline ($\beta = 0.853$ for full sample; $\beta = 0.851$ for employer-initiated), indicating that tech-sector cyclicalities is not the primary source of the gradient. The industry-by-month fixed-effects specification (R3) attenuates the coefficient more substantially — to 0.395 (SE = 0.175) for the full sample and 0.460 (SE = 0.123) for employer-initiated separations — consistent with a degree of upward bias in the baseline from industry-specific time-varying confounders, but the coefficient remains positive and statistically significant. Finally, restricting to 2024–2025 (R4) increases the coefficient ($\beta = 1.158$ for full sample, $\beta = 0.945$ for employer-initiated), suggesting that the AI gradient has strengthened in more recent years, in line with the dynamic coefficient evidence in Figure 8 of the main text.

	<i>Employer-initiated separations only</i>				
	Baseline	Excl.2025 Q2+	Excl.ICT	IndxMonth	2024-2025
AI exposure index	0.809*** (0.149)	0.603*** (0.166)	0.853*** (0.161)	0.395* (0.175)	1.158*** (0.170)
N	6520	4729	5707	6472	4423
Within R ²	0.075	0.066	0.058	0.151	0.087
Fixed effects	month + 1D	month + 1D	month + 1D	indxmonth + 1D	month + 1D
	<i>Employer-initiated separations only</i>				
AI exposure index	0.792*** (0.104)	0.600*** (0.110)	0.851*** (0.110)	0.460*** (0.123)	0.945*** (0.133)
N	5416	3902	4670	5366	3701
Within R ²	0.110	0.101	0.089	0.165	0.113
Fixed effects	month + 1D	month + 1D	month + 1D	indxmonth + 1D	month + 1D

Notes:

1) *: $p < 0.05$, **: $p < 0.01$, ***: $p < 0.001$

2) Standard errors in parentheses: HC1 throughout; M5/M5r used 4-digit MASCO occupation-clustered SEs. Outcome: monthly loss-of-employment (LoE) events per 1,000 LFS 2021 workers. Sample: 4-digit MASCO occupations, January 2023–December 2025.

3) R0 = Full 2023–2025 baseline (replicated M1/M1r); R1 = excluded April–December 2025 (US Liberation Day tariff shock); R2 = excluded Information & Communication sector (global tech-sector cycle); R3 = industry×month FE (absorbed industry-specific business cycles); R4 = January 2024 onwards (removed 2023 deferred-claims accumulation). HC1 SEs throughout. No WCB for robustness specifications.

