

RESEARCH NOTE | 16 SEPTEMBER 2026

Will artificial intelligence be the jobs Armageddon so many people fear?

Clues are in the data and the relevant timelines

AUTHOR:

JLC Venter

✉ jlc@sun.ac.za

☎ +27 (21) 808 9784

Please refer to the glossary on the BER's **website** for explanations of technical terms.

Copyright & Disclaimer

*This publication is confidential and only for the use of the intended recipient.
Copyright for this publication is held by Stellenbosch University.*

Although reasonable professional skill, care and diligence are exercised to record and interpret all information correctly, Stellenbosch University, its division BER and the author(s)/editor do not accept any liability for any direct or indirect loss whatsoever that might result from unintentional inaccurate data and interpretations provided by the BER as well as any interpretations by third parties. Stellenbosch University further accepts no liability for the consequences of any decisions or actions taken by any third party on the basis of information provided in this publication. The views, conclusions or opinions contained in this publication are those of the BER and do not necessarily reflect those of Stellenbosch University.

Executive summary

Output holds, but headcount falls slowly at first.

Our collective response as a country matters.

Views on artificial intelligence (AI), automation and employment range from mass layoffs to a future in which technology raises productivity without reducing overall employment. The evidence so far supports neither extreme. This note reviews what is currently known and considers how the effects could play out in South Africa.

International evidence shows little sign of an aggregate labour-market shock from AI to date. Studies do, however, point to changes in the composition of employment, particularly among younger and entry-level workers in AI-exposed occupations, although the causal role of AI remains contested. If these effects are emerging, they may first appear through slower hiring, attrition and non-replacement rather than through large-scale retrenchments. That would make them difficult to detect in headline employment statistics.

South Africa's exposure to AI is likely lower than in advanced economies because a larger share of employment remains concentrated in manual and less-easily automated work. Lower aggregate exposure, however, does not necessarily imply lower vulnerability. The jobs most exposed are concentrated in the formal sector, including professional and entry-level occupations that matter disproportionately for skills formation, earnings and the tax base. In a labour market characterised by very high unemployment, long-term unemployment and weak absorption of young workers, even a gradual narrowing of entry-level opportunities could have important consequences.

The note draws five broad conclusions:

- **The nature of technological change matters as much as the pace of adoption.** Some technologies replace or commodify existing expertise, while others augment workers or create new tasks. The policy question is therefore not simply how much AI is adopted, but how it is used and what kinds of work it creates or replaces.
- **Exposure differs sharply across activities.** Tasks that are routine, structured or easily traded are generally more exposed to automation, while those requiring unpredictable physical work, interpersonal interaction or judgement in less structured settings are less so. Many of the least-exposed jobs are also relatively low-paid because the economics of replacing labour with technology remain unattractive.
- **The first labour-market effects may show up in hiring rather than firing.** South African evidence from the agricultural minimum-wage literature illustrates how employment can adjust through slower job creation rather than retrenchment. The underlying shock is different, so the magnitude should not be transferred to AI, but the mechanism is relevant.
- **The bigger long-term risk is not the technologies themselves, but failing to create new work.** Advanced economies have combined automation with the creation of new skilled occupations. South Africa's challenge is that it could import productivity-enhancing technologies without

generating enough new high-value work, particularly for younger workers entering the labour market.

- **AI could also have fiscal consequences.** South Africa relies heavily on a relatively small group of high-income taxpayers, many of whom work in professional occupations with relatively high AI exposure. AI may raise the productivity and earnings of some of these workers, while reducing demand or compressing earnings in others. The direction is uncertain, but the concentration of the tax base makes the issue worth monitoring.

A central implication is that South Africa should avoid unnecessarily raising the cost of employment, particularly where this strengthens firms' incentive to substitute technology for labour without generating corresponding productivity gains. At the same time, policy should focus on strengthening the entry points into formal employment through apprenticeships and training, attracting scarce skills, improving foundational education and expanding digital and ICT infrastructure.

There is no attempt to forecast a single employment effect from AI. The evidence is too limited, and the technology is evolving too quickly. The more useful question is where the first signs will appear. Entry-level hiring, occupational composition, vacancies and job advertisements, wages in exposed occupations, and firms' adoption and use of AI are likely to be more informative than headline unemployment alone.

A note before we start.

This is a fast-moving subject. Over the weekend before this note was finalised, some of the biggest names in AI themselves called for a slower pace of frontier-model development and stronger independent safety checks. Whether that actually happens remains to be seen, but it is a useful reminder of how quickly both the technology and the debate around it are changing. What follows is therefore our reading of the evidence as it stands now, not an attempt to pretend that the path ahead is settled.

Table of Contents

EXECUTIVE SUMMARY	3
TABLE OF CONTENTS	5
1. WHAT THE EVIDENCE SHOWS	6
1.1 The aggregate impact is flat for now	6
1.2 The compositional signal in the data	7
1.3 Classifying the technology: five categories, one of them pro-worker	8
1.4 Physical automation and rent dissipation	9
1.5 Why historical automation evidence only gets us so far	11
1.6 The attrition mechanism, measured in South Africa	11
2. HOW EXPOSED IS SOUTH AFRICA RIGHT NOW?	12
2.1 Exposure is lower here than in rich countries	12
2.2 The labour market cannot absorb a knock at the entry point	13
2.3 The informal sector is not a guaranteed buffer	13
2.4 Labour market institutions cut both ways	13
2.5 Which sectors cope and which do not	14
2.6 The labour share argument has changed	15
3. THE FISCAL CHANNEL	15
3.1 The tax base is narrow and sits where the exposure is	15
3.2 Two channels through which AI could affect the tax base	16
4. CONCLUDING REMARKS	16

1. What the evidence shows

AI has substantial potential to automate tasks and displace workers, but do the numbers show that happening yet? South African evidence is still sparse, making it difficult to isolate AI's effect from the many other factors shaping employment outcomes. Internationally, however, a substantial body of research exists, and the results are less alarming than the hype suggests.

1.1 THE AGGREGATE IMPACT IS FLAT FOR NOW

The most rigorous available test links large-scale adoption surveys to administrative payroll records in Denmark (Humlum and Vestergaard, 2026).¹ The study covers 25 000 workers across 7 000 workplaces in the 2024 survey round, across 11 occupations selected for high exposure to AI: accountants, customer support specialists, financial advisers, HR professionals, IT support specialists, journalists, legal professionals, marketing professionals, office clerks, software developers and teachers. Using a difference-in-differences design, which compares changes over time between similar groups, the authors find no measurable effect on pay or hours worked, for either individual workers or entire workplaces. Two years after ChatGPT's launch, they can rule out any impact larger than 2 per cent, in either direction. The ChatGPT of 2026 is significantly more sophisticated compared to that of 2024, so even this assertion is uncertain.

This result held even for daily ChatGPT users, early adopters, workplaces offering full support (encouragement, enterprise tools and training) and all 11 occupations studied, including those where pay is negotiated individually rather than set by fixed rules.

Please note that finding no effect does not mean ChatGPT had zero impact; it just means the data could not detect one. A true effect could still exist, but it may be too small, or too hard to isolate, given the methods and data used.

Denmark's suitability as a test case is not incidental. It has a flexible labour market, low hiring and firing costs and largely decentralised annual wage negotiations, similar to the United States.² That similarity helps explain why the result may travel reasonably well to the United States, but also why it may be less directly applicable to South Africa.

Adoption in Denmark is rapid and increasingly employer led. About 43 per cent of employers encourage chatbot use, and in workplaces combining encouragement, enterprise tools and training, 93 per cent of workers report having used chatbots and 19 per cent report saving more

¹ Humlum, A. & Vestergaard, E. (2025, revised March 2026). "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI." NBER Working Paper 33777. Previously circulated as "Large Language Models, Small Labor Market Effects." <https://www.nber.org/papers/w33777>

² Humlum & Vestergaard (2026), Section 1 and Appendix C.1. The authors argue Denmark's low hiring and firing costs and largely decentralised annual wage negotiation make it comparable to the United States, which is the basis of their claim that the result generalises to other advanced economies.

than an hour a day.³ Randomised controlled trials of the same occupations document large productivity gains from AI use.⁴ One might expect these productivity gains to show up over time in stronger output growth, lower labour costs, higher margins or changes in wages and employment.

So far, however, the Danish data shows little aggregate movement in pay or hours worked. The US Census Bureau's Business Trends and Outlook Survey polls thousands of firms on their use of AI to produce goods and services. The share of those reporting AI use grew from 3.7 per cent in late 2023 to about 10 per cent by the end of 2025.⁵

Among firms using AI, 44 per cent say it supplemented or enhanced existing work, 10 per cent say it performed a task an employee used to do, and 11 per cent say it introduced a task nobody had been doing before.⁶

Among firms using generative AI, the most common uses are writing or editing documents and email (85 per cent), searching for information (about half) and summarising documents (45 per cent), with coding a distant fourth (13 per cent). Most adopters (64 per cent) report changing nothing about the business to use AI, while 15 per cent trained existing staff, 15 per cent built new workflows and just over 1 per cent hired anyone with AI skills. Where AI has taken over some tasks, the degree of substitution is growing. The share of firms reporting that it took over a large number of tasks rose from 2.4 to 7.1 per cent between 2023 and 2025. Those saying it took over a moderate number of tasks rose from 13 to 22 per cent. This group remains small, however, at about a tenth of AI adopters, who themselves are about a fifth of all firms.

Headline numbers are not falling even though the underlying impacts are visible and widely agreed. The authors of the Danish study argue that employers continuously reorganise internal tasks around a more productive, AI-upskilled workforce rather than cut headcount. This mechanism may be less likely to play out in the same way in South Africa, where aggregate productivity growth has been much weaker, and firms face a very different labour market and operating environment.

1.2 THE COMPOSITIONAL SIGNAL IN THE DATA

A widely cited US finding is that employment among workers aged 22 to 25 in the most AI-exposed quintile fell by roughly 6 per cent between late 2022 and mid-2025, in ADP payroll records, while older workers in the same occupations recorded gains.⁷ PwC's 2026 Barometer, drawing on more than a billion job advertisements across 27 countries, puts the AI-skills wage premium at 62 per cent, up from 57 per cent, and finds that AI-exposed entry-level roles are

³ Humlum & Vestergaard (2026), Section 2.2 and Figure 2. The frequently quoted figure of 2.8 per cent average time savings comes from the earlier 2023 survey round reported in Humlum, A. & Vestergaard, E. (2025), "The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities Among Workers," PNAS 122(1), and should not be attributed to the revised paper.

⁴ Brynjolfsson, E., Li, D. & Raymond, L. (2025). "Generative AI at Work." Quarterly Journal of Economics; Noy, S. & Zhang, W. (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." Science 381(6654), 187–192.

⁵ US Census Bureau, Business Trends and Outlook Survey. <https://www.census.gov/data/experimental-data-products/business-trends-and-outlook-survey.html>

⁶ Marginal Revolution (August 2026). <https://marginalrevolution.com/marginalrevolution/2026/08/ai-and-employment-what-the-firms-say.html>

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.

seven times more likely to demand traditionally senior skills. Such roles have grown by 35 per cent since 2019, while other entry-level roles have fallen by 10 per cent.⁸

Together these findings are consistent with a process of ‘seniorisation’: firms relying less on entry-level workers while maintaining middle and senior skill sets.

The direct test cuts against this story. Humlum and Vestergaard replicate the early-career employment decline in exposed Danish occupations, then compare firms that adopted AI chatbots with those that did not. The decline is no larger at adopting firms, which undermines the claim that AI is the cause. Their estimates on the early-career employment share are precise enough to rule out effects larger than a third of a percentage point, implying that AI adoption accounts for, at most, roughly 0.14 percentage points of an aggregate decline of about 0.6 percentage points.⁹ They find no differential effect on total hours, wage bills, job creation or destruction, incumbent worker hours, or the composition of hires and separations at adopting workplaces.

The authors’ reading is that other forces may be driving the decline, including macroeconomic conditions or hiring freezes made in anticipation of AI rather than in response to actual adoption. Two 2026 papers press the same point from different directions: one arguing that AI-exposed occupations were already deteriorating before ChatGPT’s release, the other directly questioning the entry-level ladder story.¹⁰

1.3 CLASSIFYING THE TECHNOLOGY: FIVE CATEGORIES, ONE OF THEM PRO-WORKER

Acemoglu, Autor and Johnson (2026) supply a useful framework for tracing how technology affects workforces.¹¹ Building on Acemoglu, Kong and Restrepo (2025), they distinguish five economically distinct categories of technological change.¹² This framework supersedes the so-so technologies binary used in earlier versions of this note, which conflated the last three rows in Table 1.^{13,14}

⁸PwC (2026). Global AI Jobs Barometer. <https://www.pwc.com/gx/en/news-room/press-releases/2026/pwc-2026-ai-jobs-barometer.html>

⁹Humlum & Vestergaard (2026), Section 3.4 and Appendix D.2. Their arithmetic is as follows. Roughly 43 per cent of workplaces encourage chatbot use, and at the largest decline consistent with their confidence interval, adoption accounts for at most $0.43 \times 0.33 \approx 0.14$ percentage points of an aggregate early-career share decline of about 0.6 percentage points between October 2022 and December 2024.

¹⁰Frank, M. R., Javadian Sabet, A., Simon, L., Bana, S. H. & Yu, R. (2026). “AI-Exposed Jobs Deteriorated before ChatGPT.” arXiv:2601.02554; Iscenko, Z. & Curto Millet, F. (2026). “Looking for the Ladder: Is AI Impacting Entry-Level Jobs?” Economic Innovation Group, The American Worker Project.

¹¹Acemoglu, D., Autor, D. & Johnson, S. (2026). “Building Pro-Worker Artificial Intelligence.” The Hamilton Project, Brookings Institution, February 2026. <https://economics.mit.edu/sites/default/files/2026-03/Building%20Pro-Worker%20Artificial%20Intelligence.pdf>

¹²Acemoglu, D., Kong, F. & Restrepo, P. (2025). “Tasks at Work: Comparative Advantage, Technology and Labor Demand.” In Handbook of Labor Economics, vol. 6, eds. Dustmann & Lemieux. Elsevier. DOI: 10.1016/bs.heslab.2025.08.003

¹³“So-so” technologies refer to activities such as self-checkout services at supermarkets, where technology is deployed without any significant increase in productivity.

¹⁴Acemoglu, D. & Restrepo, P. (2019). “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” Journal of Economic Perspectives 33(2), 3–30. Source of the “so-so technologies” formulation used in earlier versions of this document and superseded here by the five-category taxonomy.

Table 1. Five categories of technological change and their likely South African impact

Category	Effect on the value of human expertise	Labour share	Pro-worker?	South African likely impact
Labour-augmenting	Expertise more useful, but higher output may cut price	Roughly neutral	Ambiguous	Cost per unit may imply less labour used
Capital-augmenting	Same ambiguity as above	Roughly neutral	Ambiguous	Uncertain
Automating	Existing expertise commodified or made obsolete	Falls	No	Negative
Expertise-levelling	Entrants gain; incumbents' scarcity value dissipates	Ambiguous	Ambiguous	Ambiguous and dependent on response. Increased skills available but may place downward pressure on wages
New task creating	Generates demand for novel expertise	Rises	Unambiguously yes	Positive if successfully applied

Notes: adapted from Acemoglu, Autor and Johnson (2026), building on Acemoglu, Kong and Restrepo (2025). The taxonomy supersedes the "so-so technologies" binary used in earlier versions of this note.

Three implications matter for what follows, and South Africa would do best to think proactively about how to shape the way AI is adopted. Even the term AI is often difficult to define, as it might indicate automation using AI, the adoption of AI in business processes using the classic LLMs or more engineering-led automation.

First, the policy question is not how much AI arrives but which kind. Automation and new-task creation can move the labour share in opposite directions, and both can result from AI. The balance between them will depend on how the technology is deployed.

Second, expertise-levelling is not pro-worker by definition. It enables less-expert workers to perform tasks that previously required expertise from another domain, benefiting entrants but diminishing the scarcity value of incumbents. In South Africa, with its persistent high-skills shortages, this could also be a positive, if lower-skilled workers use AI to reach higher productivity levels.

Third, the authors' definitional bar is demanding and perhaps narrower than the way 'pro-worker' is often used. A tool that lets anyone do anything without expertise would fully commodify labour and eliminate its scarcity value. Pro-worker AI is not AI that removes the need for skill. It is AI that makes skills more valuable, which requires it to be deployed with future tasks in mind rather than simply to replace tasks people currently do. South Africa's persistent skills shortage means AI could offer a significant opportunity to upskill rapidly, especially in the informal sector and the lower-skilled parts of the economy.

1.4 PHYSICAL AUTOMATION AND RENT DISSIPATION

So-called cognitive automation, where AI displaces tasks requiring judgement, tends to diffuse more gradually. Industrial robotics can move faster where tasks are routine and predictable.

Acemoglu and Restrepo find robust negative effects across US commuting zones where robots are deployed. One additional robot per thousand workers reduces the employment-to-

population ratio by 0.2 percentage points and wages by 0.42 per cent. This is equivalent to roughly 5.6 workers displaced per robot. Their 2026 co-authored work reaffirms this result.¹⁵

This is in the United States, where cutting-edge technology is constantly designed and deployed, so its relevance to South Africa depends on tradability. Where a task or product is tradable, a firm that deploys robots to cut costs abroad can reduce demand for the less-competitive version produced domestically, even if local firms themselves do not automate.

Over a one-to-three-year horizon, physical capital moves at the speed of capex cycles, safety certification and local regulation. Agility Robotics, a US firm, has verified that over 100 000 totes¹⁶ have been moved by its Digit humanoid robot at a single GXO logistics facility, a genuine first. But independent trackers put the total installed base at only 40 to 150 units, concentrated at that one site, which is far from a widespread phenomenon even in the US.¹⁷ Analysts expect limited production deployment in logistics and light manufacturing from around 2027, subject to many constraints. In the US, that would hardly qualify as an aggregate labour-market event over the next decade, let alone in South Africa.

The longer-term implications should not be dismissed once these frictions ease and deployment becomes more general, particularly given the scale of current investment in humanoid robotics.

Extending the timeframe into the 2030s, diffusion lags start to expire. Acemoglu and Restrepo find that between 50 and 70 per cent of changes in the US wage structure between 1980 and 2016 are accounted for by relative wage declines among worker groups specialised in routine tasks in rapidly automating industries.¹⁸ The pace of change from AI could plausibly be faster than that from digitisation and the microprocessor.

Firms target automation at high-rent tasks, meaning those where workers earn above their outside options because the work embodies firm-specific knowledge, and this targeting works to the detriment of worker earnings. More sharply, automation can be adopted unproductively purely to reduce worker bargaining power, raising profits at the expense of earnings with no gain in output or efficiency.¹⁹ Automating away from human labour carries a cost, so the lower the cost of labour per unit produced, the lower the incentive to automate, and the reverse also holds. This may matter particularly in South Africa, where a high-friction bargaining environment could strengthen the incentive to automate for bargaining reasons as well as productivity reasons

Legislation can protect these occupations over the short term, but that protection comes at the cost of efficiency and competitiveness. Certain job categories could conceivably be protected as human-reserved work in future.²⁰

¹⁵ Acemoglu, D. & Restrepo, P. (2020). "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128(6), 2188–2244. Reaffirmed in Acemoglu, Autor & Johnson (2026), endnote 10.

¹⁶ A tote is a standardised plastic storage bin or small load carrier that automated robots grip, stack, move and retrieve instead of handling loose items.

¹⁷ Agility Robotics (2025). "Digit Moves Over 100,000 Totes in Commercial Deployment." Installed-base estimates of roughly 40 to 150 units from the New Market Pitch Agility Digit Deployment Tracker (2026), an independent estimate rather than a company figure.

¹⁸ Acemoglu, D. & Restrepo, P. (2022). "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica* 90(5), 1973–2016. DOI: 10.3982/ECTA19815.

¹⁹ Acemoglu, D. & Restrepo, P. (2026). "Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity." *Quarterly Journal of Economics*, 30 January 2026, qjag006. DOI: 10.1093/qje/qjag006. Summarised in Acemoglu, Autor & Johnson (2026), §V.B.1 and endnote 9.

²⁰ SGates, B. (2026). "The turbulent AI era is here. The choices we make now are critical." Published 26 August 2026.

1.5 WHY HISTORICAL AUTOMATION EVIDENCE ONLY GETS US SO FAR

Most of the above evidence comes from studies of physical, single-purpose machinery, such as factory robots or specialised equipment, used in rich-country job markets between 1990 and 2016. Those studies produced elasticities showing how much automation typically affects jobs and wages.

AI does not behave like that older kind of capital. Physical machines are rivalrous, in that only one factory can use a given robot at a time, and they are built for a single job. AI is cheap to copy, can be used by unlimited people at once and can, optimists argue, be applied to almost any task. The old elasticities therefore do not transfer directly, and AI should not be assumed to affect jobs the way factory automation did.

This note, therefore, does not attempt to quantify AI's employment impact precisely. In the short term, adoption will be shaped by existing budgets, job structures, labour-market institutions and firms' ability to reorganise work. Over a longer horizon, those constraints may ease as technology diffuses and business models adjust. But uncertainty rises with the forecast horizon, so any longer-term conclusions should be treated cautiously.

1.6 THE ATTRITION MECHANISM, MEASURED IN SOUTH AFRICA

AI-specific evidence for the attrition mechanism is contested and far from clear-cut. The strongest support for the mechanism highlighted in this note is domestic and predates generative AI.

Agriculture is the one South African sector where the minimum-wage literature consistently finds large, negative employment effects, repeatedly flagged as an exception to the country's otherwise modest employment responses to minimum wages. Bhorat, Kanbur and Stanwix estimate a decline in the likelihood of employment of up to 13 per cent in the four years following the 2003 sectoral determination.²¹

The 2013 increase is the more instructive case. Piek, von Fintel and Kirsten, using anonymised SARS tax data for the 2010/11 to 2016/17 tax years, find that employment fell by approximately 14 percentage points, of which only 5 percentage points are attributable to job destruction, with the remainder attributable to slower job creation.²² They also find that only 37 per cent of the specified minimum wage increase was actually paid, which resolves the apparent paradox of large disemployment alongside large non-compliance. Employment adjusted through hiring (lack of new job creation), not through firing.

²¹ Bhorat, H., Kanbur, R. & Stanwix, B. (2014). "Estimating the Impact of Minimum Wages on Employment, Wages, and Non-Wage Benefits: The Case of Agriculture in South Africa." *American Journal of Agricultural Economics* 96(5), 1402–1419. The 13 per cent estimate appears in the earlier DPRU Working Paper 149 (2012) version.

²² Piek, M., von Fintel, D. & Kirsten, J. (2020). "Separating employment effects into job destruction and job creation." UNU-WIDER Working Paper 51/2020. DOI: 10.35188/UNU-WIDER/2020/808-5. Published as Piek, von Fintel & Kirsten (2023), "The impact of agricultural minimum wages on worker flows in South Africa," *South African Journal of Economics*, DOI: 10.1111/saje.12357.

In their final paper²³ they found that employment growth was temporarily affected by 54% for the average firm, but had recovered four years after the 2013 shock. The effect was mainly driven by slower entry with abrupt and relatively small effects on exit. There were no wage results four years later.

The rough initially observed 5-to-9 split between job destruction and slower job creation provides a useful illustration of how labour-market adjustment can occur through non-creation rather than retrenchment.²⁴ **A wage shock in agriculture is clearly not the same as a technology or productivity shock, and the underlying shock here is a wage increase, not a change in technology. The mechanism transfers. The magnitude should not.** The intention is to use this example to illustrate how AI's impact might transmit, not to quantify how much or when.

The practical implication concerns what to measure. If AI affects employment through this channel, the first signs may be less visible in retrenchment numbers and more apparent in slower hiring and changes in the composition of work.

2. How exposed is South Africa right now?

South Africa's exposure to AI is lower than that of rich countries, but the exposure is concentrated in particularly vulnerable parts of the labour market. The jobs most at risk are the formal ones young people use to enter the economy, and the ones the tax system depends on.

2.1 EXPOSURE IS LOWER HERE THAN IN RICH COUNTRIES

The IMF estimates that about 60 per cent of jobs in advanced economies are in occupations with high AI exposure, compared with about 40 per cent in emerging markets and about 26 per cent in low-income countries. The reason is straightforward as poorer countries have more manual and farm work. South Africa was included in the IMF sample.²⁵

The headline risk in South Africa is therefore smaller than in Europe or the United States, though that is not the same as being safe. The exposed jobs in South Africa are clerical and professional jobs in the formal sector. Those jobs pay most of the income tax, and they are the main route into the formal economy for many people.

²³ Piek, von Fintel & Kirsten, "The impact of agricultural minimum wages on worker flows in South Africa" South African Journal of Economics 10 September 2013.

²⁴ Some caveats apply. A 2024 review notes the identification strategy is not clearly specified, so the point estimate is indicative rather than settled.

²⁵ Cazzaniga, M. et al. (2024). "Gen-AI: Artificial Intelligence and the Future of Work." IMF Staff Discussion Note SDN/2024/001. The data sources include Labour Market Dynamics in South Africa.

2.2 THE LABOUR MARKET CANNOT ABSORB A KNOCK AT THE ENTRY POINT

Unemployment stood at 33.6 per cent in the second quarter of 2026, up 0.9 percentage points on the previous quarter. There were 8.5 million unemployed people, compared with 16.7 million in work. The number of unemployed youth aged 15 to 34 rose by 264 000 to 5.0 million.²⁶

The figure that should worry policymakers most is long-term unemployment. It has risen from 64.9 per cent of the unemployed in early 2016 to 77.4 per cent in early 2026. More than three in four unemployed people have been out of work for over a year.

Most international work on AI assumes displaced workers move into other jobs. That assumption does not hold in South Africa. The relevant question here is not only what happens to people who lose jobs, but whether young people get a first job at all.

2.3 THE INFORMAL SECTOR IS NOT A GUARANTEED BUFFER

Informal work accounts for between 20 and just over 30 per cent of total employment, depending on how it is measured. Informal trading is often assumed to absorb formal job losses, but this is unfortunately not always true. Research has found significant constraints hampering growth of the informal sector in South Africa - such as crime, zoning scheme limitations and lack of physical infrastructure.²⁷

While AI is unlikely to help informal traders overcome these obstacles, it could provide cheaper access to business advice, payment tools, or assistance with things like website or logo design. However, this depends on the user's ability to engage with AI tools. South Africa has extensive mobile coverage, but internet access does not imply meaningful use. Data costs and other constraints limit the extent to which South Africans have equal opportunity to adopt the suite of productivity benefits that AI can offer.²⁸

2.4 LABOUR MARKET INSTITUTIONS CUT BOTH WAYS

Danish and American results generalise to each other because both countries have flexible hiring and firing and decentralised wage setting. South Africa has neither. Bargaining council extensions, CCMA procedures, strong union presence and equity scorecards all make it expensive and complex to reduce staff. South Africa's labour-market institutions offer workers relatively strong protections, but these protections also affect firms' hiring and adjustment decisions

These protections may slow job losses, but they can also raise the cost of adjusting headcount and potentially the hurdle for hiring junior or entry-level workers. Where bargaining is difficult, firms have a stronger reason to automate to weaken workers' bargaining power rather than to raise output, which is the least-useful kind of automation.

²⁶ Statistics South Africa (2026). Quarterly Labour Force Survey, Q2:2026, released 11 August 2026. <https://www.statssa.gov.za/?p=19804>

²⁷ <https://www.gtac.gov.za/resource/low-informality-amidst-excessive-unemployment-south-africas-empirical-smoking-gun/>

²⁸ Gillwald & van der Spuy (2026). Measuring what matters: Digital inclusion and South Africa's growth and development agenda. BER Research note (forthcoming)

2.5 WHICH SECTORS COPE AND WHICH DO NOT

South Africans can sell professional work into hard-currency markets without a local AI industry, local data centres or state capacity. Global business services employment grew from 65 000 jobs in 2019 to about 150 000 in 2024, and export earnings rose from R18.9 billion to R53 billion.²⁹ The question is whether this trend can continue.

AI is likely to reduce the number of simple call-centre seats over time and increase the value of the harder ones, where someone must be accountable for the answer and exercise discretion rather than follow a script. The result is fewer seats, more value per seat and better-trained operators. That is bad for mass employment but good for the earnings of those who keep a job, and it widens inequality at home. That trade-off is not palatable, but it is worth stating clearly.

Three factors decide this.

1. How structured the workplace is. Machines handle controlled settings long before messy ones.
2. Whether the sector is traded. Cheap local labour only protects work that cannot be imported at a lower price.
3. Value per unit of output. High value pulls in capital even where the work is difficult.

The second point is the one most often missed. In a traded sector, when the decision to automate is taken abroad by competitor firms at foreign wages and cost structures, the world price will tend to fall to their cost base regardless of what South African employers pay their workers to produce the same product. Piek and von Fintel tested this by comparing agriculture with retail and found that small farms bore the damage while large farms were shielded.³⁰

On this framework, more exposed activities may include fruit grading and packing, low-complexity call-centre work, retail, junior audit and legal work, and light manufacturing. Less exposed activities are likely to include those requiring unpredictable actions, such as tourism, frontline care, domestic work and security. Mining and field agriculture probably come later, but the pressure is real. The cost of mechanical fruit harvesting still protects wages for now, but full mechanisation is feasible.

The uncomfortable part is that the safest sectors are also the poorest ones. What protects a job from a machine is usually that it was never worth deploying much capital to defend it in the first place.

²⁹BPESA. Refreshed National Value Proposition; GBS Sector Job Creation Report. <https://www.bpesa.org.za/>

³⁰Piek, M. & von Fintel, D. (2020). "Sectoral minimum wages in South Africa: disemployment by firm size and trade exposure." *Development Southern Africa* 37(3), 462–482.

2.6 THE LABOUR SHARE ARGUMENT HAS CHANGED

Popular narratives assume automation cuts labour's share of national income. Cross-country numbers do not support that. In 2020, labour's share of income was 60.5 per cent in North America and 58.0 per cent in Europe, against 51.8 per cent in Sub-Saharan Africa and 51.0 per cent in Latin America. The most automated economies have the highest labour shares.³¹

One important reason is the creation of new, expert work in advanced economies. More than six in ten American workers in 2018 were in job specialities that did not exist in 1940.³² South Africa, it could be argued, worries too much about the jobs and companies it has, and not enough about the jobs and companies it does not yet have.

The risk for South Africa, then, is not that automation arrives. It is that the country imports the automation and creates no new work. That is a harder problem to solve, and it requires genuine, longer-term strategic thinking.

3. The fiscal channel

In brief, AI could have fiscal implications because South Africa relies heavily on a relatively small group of high-income taxpayers, many of whom work in occupations with relatively high AI exposure. The country needs more people in employment.

3.1 THE TAX BASE IS NARROW AND HIGHLY EXPOSED TO AI AT THE TOP END

Personal income tax is expected to bring in about 36 per cent of total tax revenue in the 2026/27 framework, R844.8 billion out of an estimated R2.35 trillion. Personal and company tax together typically make up roughly 55 per cent of the total, which is high by OECD standards. Within personal tax, more than half comes from the 7.7 per cent of taxpayers earning above R1 million a year.³³

There is likely to be meaningful overlap between high-income taxpayers, workers in professional occupations with relatively high AI exposure, and those able to earn an AI-related skills premium.³⁴ The sustainability and growth of South Africa's tax base and future tax revenue, therefore, depend on the part of the labour market where AI's effects are potentially going to be the strongest. This means developments in AI-exposed professional occupations could matter disproportionately for the tax base.

³¹Acemoglu, D., Autor, D. & Johnson, S. (2026). "Building Pro-Worker Artificial Intelligence." The Hamilton Project, Brookings, February 2026, Figure 1 (Our World in Data, 2020).

³²Autor, D., Chin, C., Salomons, A. & Seegmiller, B. (2024). "New Frontiers: The Origins and Content of New Work, 1940–2018." Quarterly Journal of Economics.

³³National Treasury (2026). 2026 Budget Review, Chapter 4: Revenue Trends and Tax Proposals; Werksmans Attorneys (2026), "Budget Speech 2026/2027: Tax Overview," for the concentration of PIT in the above-R1m cohort.

³⁴PwC (2026), Global AI Jobs Barometer. See also fn. 10.

3.2 TWO CHANNELS THROUGH WHICH AI COULD AFFECT THE TAX BASE

The first channel operates at the top. If AI materially changes the earnings or location decisions of high-income professionals in exposed occupations, South Africa could lose tax revenue disproportionate to their numbers. There are also limits to how much additional revenue can realistically be raised from higher marginal personal income tax rates at the top end, because, in South Africa, higher rates can encourage restructuring and weaken the expected revenue gain. This leaves relatively little room to offset any erosion of the tax base by raising personal income taxes alone.

The second channel is less obvious and works from below. Danish payroll data shows that workers who adopt AI tools switch occupations more often, and those who switch move into better-paid roles, with earnings growing 12 percentage points faster than for non-switchers.³⁵ That is good news for the individual, but the same mechanism widens the pool of people capable of performing professional work, which could put downward pressure on wages. It is not clear that this holds in South Africa, given its relatively inflexible labour market. It most probably does not.

Ride-hailing offers a useful example. Once apps such as Uber removed the need for route knowledge and a government medallion, American taxi and chauffeur employment rose by roughly 60 per cent between 2000 and 2018, while the occupation's earnings fell from about 67 to 55 per cent of the national average.³⁶ More people entered the occupation, and those already there earned less, though the aggregate number employed rose significantly. The comparison is directional rather than a forecast, but it illustrates the mechanism. If professional work in South Africa follows a similar path, an indispensable part of the tax base could get poorer per head, even as more people join it. Equally, South Africa's skills deficit means AI could enable more employment as productivity rises. This channel needs monitoring rather than assuming that one or other direction will prevail.

4. Concluding remarks

Predicting the extent of employment losses from wide-scale AI adoption in South Africa is neither possible nor responsible at this stage. The data is insufficient and what exists does not provide clear results. It is likely, however, that the labour force will be materially affected over a longer horizon.

The Danish study shows that substantial adoption can produce real productivity and time-saving gains without any aggregate negative impact on employment or wages, but AI applications are changing rapidly, and this result should not be assumed to hold indefinitely.

³⁵Humlum, A. & Vestergaard, E. (2026). "Still Waters, Rapid Currents." NBER Working Paper 33777, Section 3.3 and Table 2.

³⁶Acemoglu, D., Autor, D. & Johnson, S. (2026). "Building Pro-Worker Artificial Intelligence." The Hamilton Project, Brookings, §II.B.5 and §III, including the ride-hailing case drawn from Current Population Survey calculations.

If AI affects employment through the channels discussed in this note, the first signs may be more visible in attrition, slower hiring and non-replacement than in retrenchments. A shrinking on-ramp for entry-level jobs is a genuine concern, because it implies a workforce of fewer, better-skilled workers over time – those who are able to harness AI's improved capabilities. The risk of greater automation will rise as the technology becomes cheaper and more capable. Many jobs in the formal and informal economy remain protected from AI for now because the cost of automating them exceeds the short-run savings from replacing workers. South Africa should be cautious about policy changes that unnecessarily raise the cost of employment, because these can strengthen the incentive to automate. The agricultural minimum-wage evidence in this note illustrates how increased wage costs can reduce employment, either directly or through slower job creation.

An immediate priority is therefore better measurement. Rather than trying to forecast a single employment number, South Africa should track where the first effects are emerging: entry-level hiring, occupational composition, vacancies and job advertisements, wages in exposed occupations, and firms' adoption and use of AI.

Key policy implications from the evidence reviewed are:

- Avoid unnecessarily raising the cost of employing people, especially at the lower end of the wage spectrum.
- Reform the apprenticeship system to boost the number of skilled, specialised young people and think carefully about creating more access points on the entry-level training ladder.
- Deliberately enable scarce skills to enter the country. This has dynamic value and also adds a fiscal buffer.
- Improve foundational education, extend digital access at all levels and ensure the ICT infrastructure is in place for South Africans to compete in the digital domain.