BUREAU for ECONOMIC RESEARCH

Impumelelo GROWTH LAB

RESEARCH NOTE | April 2026

Putting AI to work for South Africans

What South Africa's next AI policy should say

Johan Fourie

Stellenbosch UNIVERSITY IYUNIVESITHI UNIVERSITEIT

CORRESPONDING AUTHOR:

Johan Fourie
Department of Economics and LEAP, Stellenbosch University

johanf@sun.ac.za
+27 (21) 808 3590

ACKNOWLEDGEMENTS:

Prepared for the Bureau for Economic Research. The author thanks colleagues at the BER for guidance and feedback. The report draws on data from the BER's 2026Q1 outlook survey, Visa and Discovery Bank's *SpendTrend26* consumer survey, and the Anthropic Economic Index. Any errors are the author's own.

Executive Summary

The Department of Communications and Digital Technologies gazetted a Draft National AI Policy for public comment on 10 April 2026 and the Minister withdrew it on 26 April after fictitious references in the bibliography were confirmed (Department of Communications and Digital Technologies, 2026*b*). The withdrawal opens the door to a stronger foundation for the next draft. This note is offered as one such foundation, drawing throughout on the recent economics literature on AI, which was largely absent from the first draft and which remains the most reliable guide currently available on the questions a national AI policy must answer. The argument of this note is simple: a national AI policy will be judged by whether it equips and empowers South Africans to put the defining general-purpose technology of this era to work to build a more prosperous future, or stands between them and it.

The empirical position. Survey evidence suggests that AI use is already widespread among engaged, digitally connected South Africans. In a once-off Bureau for Economic Research survey in 2026Q1, 95 per cent of managers and 91 per cent of private individuals reported using AI in their working week, and adoption reaches well beyond the employed: more than 90 per cent of unemployed respondents and 71 per cent of retirees use AI in some form. In the Visa and Discovery Bank *SpendTrend26* survey, 40 per cent of credit-card-using consumers use AI weekly for purchase decisions and 43 per cent pay for an AI subscription. Eighty-two per cent of managers in the BER survey expect AI to add five per cent or more to their productivity over the next three years; 57 per cent of individuals expect the same. Recent US evidence values what users gain from these tools, over and above what they pay for them, at roughly twelve times what their providers earn.

Lessons from the withdrawal. Several instincts in the withdrawn draft should not be repeated in the next one. The draft proposed eight new institutional bodies before specifying the rules that would govern how each operated, anchored its economic case on an unsourced consultancy

projection (a USD 19.9 trillion global figure for 2030), and proposed an AI Insurance Superfund modelled on the Road Accident Fund without engaging with the RAF's accumulated deficit, backlog and governance difficulties. The verbs of the draft were *establish*, *regulate*, *monitor*, *audit*, *certify*, *enforce* and *mandate*; the verbs of growth, diffusion and reallocation appeared less often. An independent audit conducted for this report identified at least eight fabricated academic references in the draft's 67-entry bibliography, and the Minister withdrew the draft on the same grounds. The episode raises a direct question about the state's readiness to oversee, audit and certify the AI systems any draft would be asked to regulate.

What the literature says. Recent research in economics is largely consistent on the relevant points. Productivity gains come from the spread of AI and the practical investments around it (skills, organisational change, measurement), not from state deployment alone. Most of the welfare accrues to users, by roughly an order of magnitude over what providers earn. AI safety regulation is structurally vulnerable to capture by the firms it oversees; autonomous AI systems need clear rules on identity, registration and accountability more than new ethics boards. Concentration among AI providers is real, but the response is fair access on existing tools, not new licensing. When regulators attempt to use AI tools to monitor AI-using firms but lack the capacity to keep pace, monitoring expands without harm falling.

An adoption-first programme. (1) Roll out the digital infrastructure AI adoption depends on: cheap bandwidth, last-mile connectivity and exchange-control rules that do not penalise software-IP transactions. (2) Halve the institutional count and write each surviving body's rules of operation into the primary policy document. (3) Where a real market failure justifies state action, target it narrowly and evaluate it carefully, in the spirit of evidence-led industrial policy. (4) Build a national measurement capability in StatsSA, not a new monitoring centre. (5) Embed AI into every existing university and TVET programme rather than creating new AI-only qualifications, and let lecturers experiment.

Contents

Already underway

South Africa's draft National AI Policy was gazetted on 10 April 2026 and withdrawn on 26 April, after fictitious references were identified in its bibliography (Department of Communications and Digital Technologies, 2026*b*). The withdrawal creates an opportunity to write a new policy from a stronger foundation. This note is offered as one such foundation, anchored throughout in the recent economics literature on AI. That literature was largely absent from the first draft, and on the questions a national AI policy must answer, it remains the most reliable guide currently available.

By the time the first gazette was published, AI adoption was already well underway in parts of the South African economy.

In the Bureau for Economic Research's 2026Q1 outlook survey, 95 per cent of managers and professionals and 91 per cent of private individuals reported using AI in their working week. Forty-one per cent of managers and 30 per cent of individuals used it for more than five hours a week (Bureau for Economic Research, 2026). In the Visa and Discovery Bank *SpendTrend26* consumer survey of one thousand credit-card-using South Africans, 40 per cent used an AI tool weekly or more to decide what to buy, where to buy it, or whether to buy at all. Forty-three per cent paid for an AI subscription. Among subscribers, 67 per cent paid for ChatGPT, 35 per cent for Gemini and 27 per cent for Microsoft Copilot. AI subscription payment volume rose by 125 per cent in 2025, and the share of AI in total subscription payments doubled in a single year (Visa and Discovery Bank, 2026). In the Anthropic Economic Index for the week of 2026-02-05, South Africa registered 4,107 Claude.ai conversations, 0.41 per cent of global usage that week, with a workload mix dominated by software development, web development and tutoring (Anthropic, 2026).

AI adoption in South Africa is occurring at scale, in a middle-income economy, in the early years of a general-purpose technology of the kind that, like electricity or the internal combustion engine, eventually affects almost every sector, and the gains are already measurable. The policy question is not whether to permit AI adoption but whether the framework will help that

adoption broaden, deepen and become more productive, or whether it will impede it.

The draft, despite its intentions, addressed supervision before diffusion. Our argument rests on a single proposition that the rest of the report defends:

> *A national AI policy will be judged by whether it equips and empowers South Africans to put the defining general-purpose technology of this era to work to build a more prosperous future, or stands between them and it.*

What South Africans are doing with AI

Three sources, each useful and each with limits, allow an empirically grounded answer to three questions: how much South Africans use AI, who uses it and on what, and to what effect.

How much

The BER's 2026Q1 outlook survey included a once-off AI module conducted among registered users of the BER website. These respondents are individuals who, in their personal capacity, subscribe to access BER publications and were asked a small set of questions about their use of AI. The sample is 358 respondents, split roughly evenly between managers and professionals and private individuals. The survey is not nationally representative and reflects a self-selected group of BER readers. Respondents who use AI are likely to be over-represented, as they may have been more inclined to complete a survey on the topic. The figures should therefore be read as indicative of behaviour among this group, rather than as population-level adoption rates.

The natural way to discipline this is to compare the BER survey with two parallel datasets that follow different sampling logics: the Visa and

Discovery Bank *SpendTrend26* consumer survey, fielded among one thousand credit-card-using South Africans earning more than R100,000 a year, and the Anthropic Economic Index, which records every Claude.ai conversation associated with a South African geography. All three samples skew toward the banked, employed, upper-middle-income end of the distribution. While this limits their representativeness, the consistency of headline findings across these datasets suggests that AI use is already widespread within this segment of the population, and provides a useful, if partial, empirical picture.

Among managers, 5 per cent reported no AI use; 18 per cent up to one hour a week; 37 per cent one to five hours; and 41 per cent more than five hours a week. Among private individuals the figures were 9, 24, 37 and 30 per cent. About 73 per cent of all respondents used AI between one and five hours a week; roughly a third used it more intensively still. AI use has become routine for the majority of respondents (Bureau for Economic Research, 2026).

The Visa and Discovery Bank survey corroborates this from a different sample. Its sample of one thousand credit-card-using South Africans earning more than R100,000 a year is upper-middle-income and urban. Within that sample, 40 per cent use AI tools weekly or more for purchase decisions; 50 per cent among 18- to 30-year-olds; 42 per cent of men and 39 per cent of women; 19 per cent have not yet used AI for shopping. The behavioural consequences are visible in the spending data: 42 per cent of AI users found a cheaper substitute for a planned purchase in the past twelve months, 35 per cent switched brands or retailers, and 35 per cent avoided a purchase they judged risky (Visa and Discovery Bank, 2026).

The Anthropic Economic Index supplies the international benchmark. In the week of 2026-02-05 to 2026-02-12, South Africa accounted for 0.41 per cent of global Claude.ai conversations, or 4,107 conversations across the Free, Pro and Max tiers. The country sits below the high-volume comparators in absolute terms: the United States recorded 222,000 conversations in the same week, India 62,000, Brazil 26,000, Mexico 10,000, Nigeria 7,700 and Argentina 5,000 (Anthropic, 2026). South Africa is not at the frontier of AI use, but it is one of the measurable participants. On a per-capita basis, given a working-age population of about 40 million, South Africa sits in the same

broad band as several middle-income peers. Two cautions follow. First, both the BER and *SpendTrend26* samples skew toward the banked, employed, educated end of the income distribution, so broad-based national figures will be lower. Second, the AEI captures Anthropic's products only, which understates total AI use to the extent that South Africans rely on ChatGPT, Gemini, Copilot and other systems (Visa and Discovery Bank, 2026).

Who, and on what

The BER microdata shows AI use intensity among private individuals broken down by employment status. Figure 1 shows the distribution.

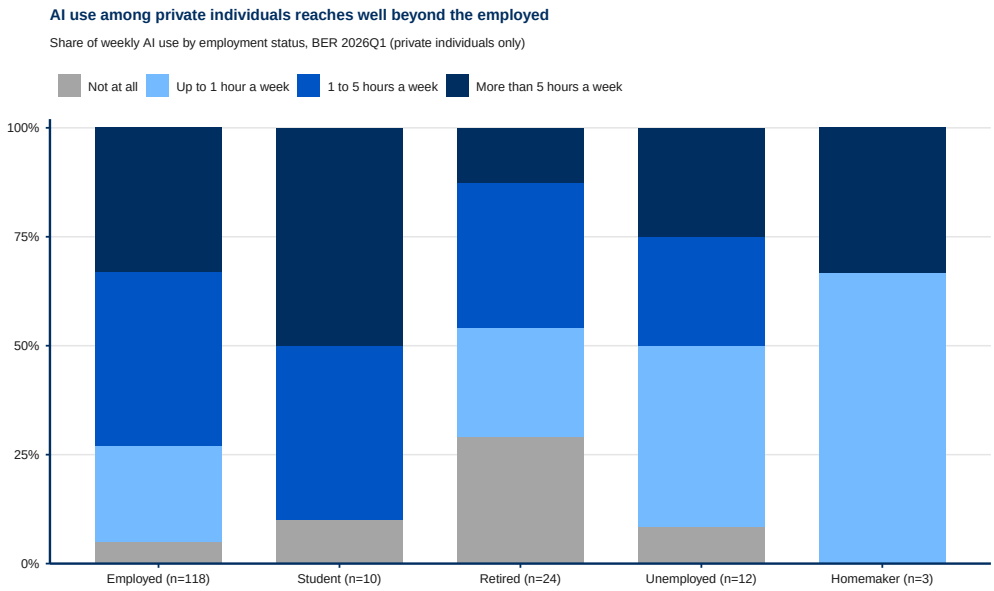

Figure 1: AI use among private individuals reaches well beyond the employed. Distribution of self-reported weekly AI use by employment status, BER 2026Q1 (private individuals only). Sample sizes are shown next to category labels; the homemaker category has only three respondents and should be read with caution. Source: Bureau for Economic Research (2026).

Among employed private individuals, 73 per cent use AI at least once a week and roughly a third use it more than five hours a week. Adoption does not stop at the employed. The student cell is small (n=10) but every student reports at least one hour of weekly AI use. Of the twelve unemployed

respondents, more than 90 per cent use AI in some form. Among retirees, 71 per cent do. AI has become part of how a wide cross-section of South Africans look for work, manage their lives, study and stay informed.

The Anthropic Economic Index also classifies each conversation by the form of human-AI exchange: *directive* (a single instruction yielding an output), *task iteration* (back-and-forth refinement), *learning* (the user is being taught), *feedback loop* (the AI is being corrected), *validation* (a check on user-generated work) and a small *none* residual. Figure 2 compares South Africa's mix to the usage-weighted global mean.

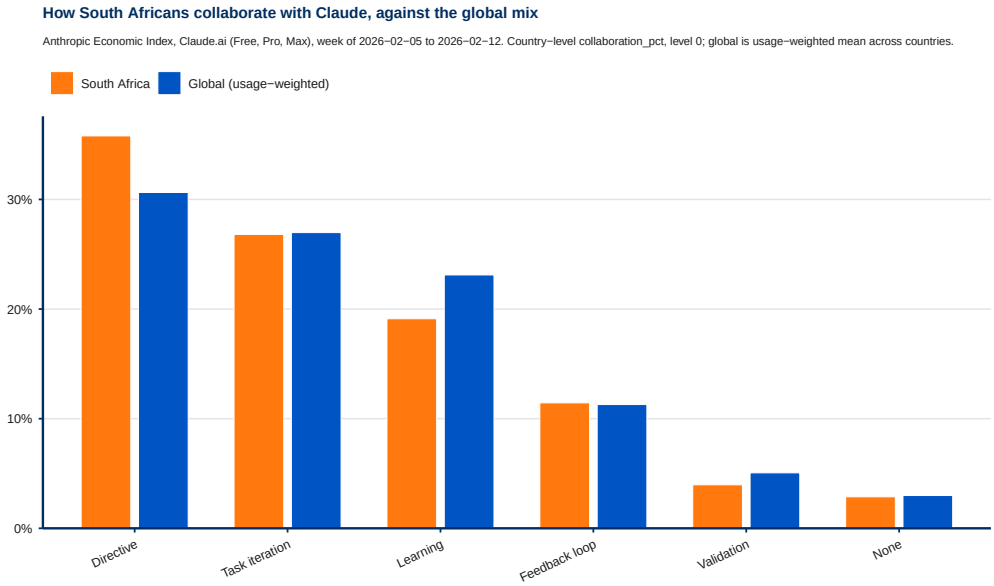

Figure 2: How South Africans collaborate with Claude, against the global mix. Country-level shares of conversation type; global is the usage-weighted mean across countries. Source: Anthropic (2026).

South Africans are slightly more likely than the global average to use Claude by giving it a single instruction and taking the answer (36 per cent versus 31 per cent), and slightly less likely to use it as a tutor (19 per cent versus 23 per cent). The other modes (back-and-forth refinement, correcting the AI, and using it to check user-generated work) track the global mix closely. South Africans, in short, use AI more as a tool that does work for them than as one that teaches them. The pattern is consistent with the productivity story the BER and *SpendTrend26* data tell, and with what we would expect of a country

whose AI users are mostly working professionals and adult consumers rather than school children.

The top ten task categories (from the standardised occupational task list maintained by the US Department of Labor, which Anthropic uses to classify what users are asking AI to do) in South Africa's Claude.ai usage are dominated by software work: modifying existing software (4.0 per cent of conversations), writing new programs (3.2 per cent), building and maintaining web sites (1.9 per cent), troubleshooting hardware and networks (1.6 per cent), and routine system administration (1.4 per cent). Four of the top ten are programming or IT support tasks. Two more are tutoring and coursework support: helping students with material, and reviewing class material with students. One is copy editing. AI in South Africa is not yet doing health, agriculture or public administration at volume in this dataset; it is helping software developers, IT support staff, tutors and writers.

To what effect

Reported productivity is the third question, and the answer matters most for policy. Figure 3 compares the share of BER respondents who selected each productivity-impact category for the past three years and for the next three years.

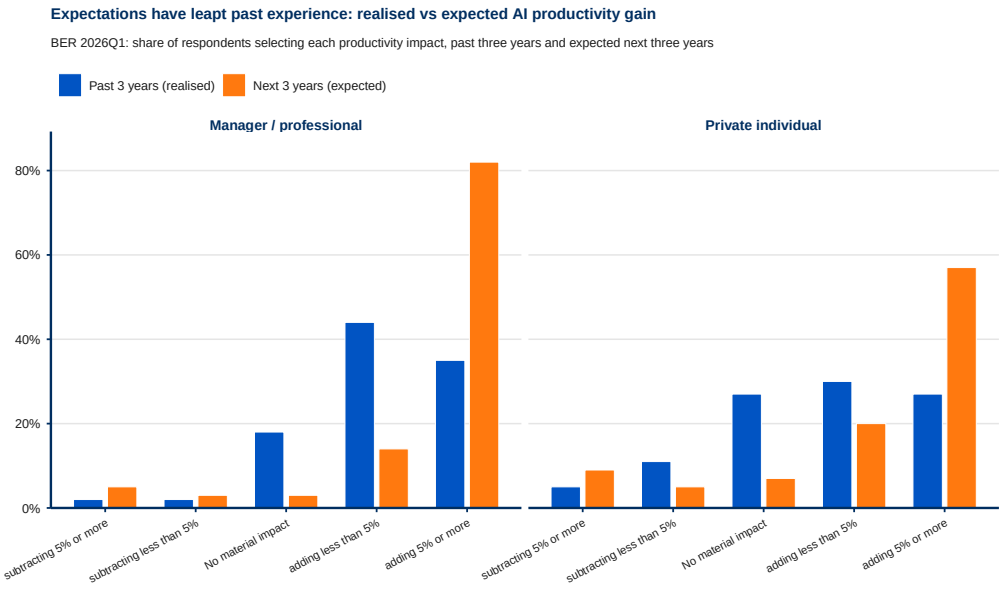

Figure 3: Expectations have leapt past experience. BER 2026Q1: share of respondents selecting each productivity impact, past three years and expected next three years, by capacity. Source: Bureau for Economic Research (2026).

Among managers and professionals, 35 per cent reported that AI had added five per cent or more to their productivity over the past three years. Eighty-two per cent expect AI to add five per cent or more over the next three. Among private individuals the corresponding figures are 27 and 57 per cent. Expectations are about twice as large as realised gains. Negative impacts are small at both horizons: four per cent of managers and 16 per cent of individuals report any net negative effect on past productivity; the corresponding figures for the next three years are eight and 14 per cent.

The 82 per cent figure is a self-report by people who already use AI of what they expect from their own work, rather than a forecast of total factor productivity, and that expectation is what shapes investment in retraining, in new tools, in workflow redesign, and ultimately in adoption depth. Expectations of this magnitude are the proximate cause of the diffusion described above. The realised-versus-expected gap is, in turn, consistent with the recent literature on AI in productive work. The international evidence has the same shape. Babina (2026) reviews evidence from senior executives

in the United States and across several advanced economies, all of which finds modest realised productivity gains from AI in 2025 alongside substantially larger expected gains over the medium term, and notes that productivity effects in firm-level data follow a J-curve, with short-run declines giving way to longer-run gains as firms reorganise around the technology. Agrawal et al. (2026) document a "jagged frontier" in scientific work in which AI's near-term effect is augmentation rather than replacement, with gains uneven across tasks and access to the best models constrained by processing power and skill. They report that the median foundation model used by frontier scientists in 2024 was 26 times smaller than the median model released in the same year. South African respondents report the same pattern: substantial but uneven realised gains, and large expected gains as access deepens and workflows reorganise around the tool.

SpendTrend26 adds a spending-side signal that is, in some ways, more striking than the work-side numbers. AI is the only subscription category whose share of South African card-payment volume doubled in 2025, and AI subscription payment volume itself grew by 125 per cent over the year, even as several entertainment subscription categories contracted. Within the AI category, however, individual products are replaced quickly: when AI subscribers pause a subscription, only 15 per cent later resume the same one. That is the lowest renewal rate of any subscription category, against 48 per cent for food and groceries and 35 per cent for online media. The combined picture is a category that is growing rapidly in aggregate while churning rapidly within. South Africans who cancel an AI subscription are not leaving AI; they are switching to a different AI product. Spending on AI is therefore competitive in a way that, say, streaming is not: providers retain customers month by month on product quality, not on switching costs (Visa and Discovery Bank, 2026).

Recent US evidence puts a dollar figure on what such use is worth. Brynjolfsson et al. (2026), surveying representative samples of US adults in July 2025 and March 2026 and asking what they would need to be paid to give up AI for a month, find that the average US generative-AI user would require $124.50 in compensation, up from $98 eight months earlier. The median valuation rose from $3.40 to $11.40 over the same period. Combined with the rise in the user base from 98 to 115 million US adults, the implied

total value to consumers (over and above what they pay for the tools) rose from $116 billion to $172 billion in under a year. That figure exceeds the estimated revenue these tools generate for their providers by an order of magnitude, in line with the long-run pattern that innovators capture only a small share of the total economic gain from a major new technology. The South African totals will be smaller, but the same accounting holds: most of the value of AI accrues to the people using it, not to the firms supplying it.

Taken across the four sources, the descriptive picture is one of broad and deepening adoption; intensity that places parts of the South African economy in the global middle band of AI use; work concentrated for now in software, education and writing; a directive-heavy collaboration mix consistent with use as a productive tool; substantial realised productivity gains; and larger expected gains still. This is the country the draft policy is addressed to.

Lessons from the withdrawal

The withdrawal of the draft on 26 April 2026 is the immediate occasion for a blank sheet. An independent audit conducted for this report identified at least eight fabricated academic references in the draft's 67-entry bibliography, and the Minister withdrew the draft on the same grounds and announced consequence management for those responsible (Department of Communications and Digital Technologies, 2026*b*). The episode raises a direct question about the state's readiness to oversee, audit and certify the AI systems any AI policy would be asked to regulate.

Several instincts in the withdrawn draft are worth identifying so the next draft can avoid them. The first was institutional overhang: the draft proposed eight new bodies before specifying the rules that would govern how each operated. The second was an unsourced economic case: the headline motivation was that AI "is projected to contribute approximately USD $19.9 trillion to the global economy by 2030" (Department of Communications and Digital Technologies, 2026*a*, p. 244), a consultancy projection used unqualified, with no discussion of how much would accrue to South Africa or the track record of comparable projections for past general-purpose

technologies. The third was a flawed analogy: an AI Insurance Superfund "modelled after the Road Accident Fund, to compensate individuals or entities harmed by AI-driven outcomes" (Department of Communications and Digital Technologies, 2026*a*, p. 320), without engagement with the RAF's accumulated deficit, backlog and governance difficulties. The fourth was the language of delegation: several interventions used wording that was, in the drafters' own words, "purposefully open-ended" (Department of Communications and Digital Technologies, 2026*a*, p. 43), deferring the actual content of those interventions to whichever officials and stakeholders happened to take part in the rule-writing process.

The fifth and broadest concern is one of orientation. The verbs of the draft were *establish*, *regulate*, *monitor*, *audit*, *certify*, *enforce* and *mandate*. The verbs of growth, diffusion and reallocation appeared less often. The next draft has the chance to be ordered around the second set of verbs, not the first.

What the literature actually says

The five subsections that follow draw on recent research in economics to identify the design principles a sensible AI policy would be built on. They move from positive findings on the spread of the technology and the knowledge base behind it, to cautionary findings on regulatory capture, the risks of committing to too much at once, and the frictions that prevent firms and workers from moving to where AI is most productive. The cumulative conclusion is that what most determines whether South Africans capture the benefits of AI lies outside an AI-specific document.

Diffusion, not deployment

A central error in AI policy design is to treat AI as a technology for the state to deploy. Agrawal et al. (2019) argue that AI is not the kind of technology a government regulates with a single AI law. AI is a sharp fall in the cost of prediction, and its effects pass through existing policy domains: privacy, trade,

liability, labour and competition. The right thing for a national policy to be about is the conditions under which AI tools spread through a real economy. Agrawal et al. (2026) make a complementary point: AI's productivity gains arise less from the technology itself than from the practical investments around it: skills, organisational change, measurement and new ways of working. Without those complements, AI generates limited measurable gains; with them, it reorganises how work is done. Babina (2026) reaches the same conclusion in stronger form, reviewing firm-level evidence that productivity gains from AI follow a J-curve: in studies of US manufacturing establishments, AI-adopting firms see lower productivity growth in the years immediately after adoption and higher productivity growth thereafter, as adjustment costs are absorbed and complementary investments mature. A second measurement problem compounds the first: Brynjolfsson et al. (2026) show that because most generative-AI use is currently free or low-priced, GDP and revenue figures understate the welfare gains from these tools by roughly an order of magnitude, so any policy designed against the value AI creates for its providers will be designed against the wrong baseline.

The South African data fit the pattern: reported productivity gains among BER survey respondents are widely distributed but uneven, with some respondents reporting large gains and others none. Anthropic's data show that the work being done with AI is concentrated in a handful of task families, and *SpendTrend26* shows AI-using consumers behaving differently across the population, with 42 per cent already substituting toward cheaper alternatives and 35 per cent switching providers. An AI policy that addresses deployment in priority sectors at length and diffusion only briefly will miss most of the welfare gain. A serious AI policy needs a national measurement capability for AI adoption, an expansion of StatsSA's business surveys to include the questions the BER is already asking, a link between AI use and the administrative tax and payroll data the National Treasury already holds, and a coordinated programme of rapid retraining for the middle- and lower-skilled workers for whom recent evidence suggests the largest gains lie. These interventions cost less than a new architecture of supervisory institutions and would generate the data the rest of the policy needs. Without them, South Africa would be regulating an activity it cannot measure.

Building the knowledge base, and keeping the best researchers

Every country that has captured a disproportionate share of a new general-purpose technology's gains has done so by investing in the underlying knowledge before the applications were obvious. Agrawal et al. (2026) document the gap that opens when a country tries to do it the other way around: the median foundation model used by frontier scientists in 2024 was 26 times smaller than the median model released that year. Access matters even at the frontier. For a country whose AI users today are mostly working professionals and consumers rather than research scientists, the gap is wider still.

The draft was rhetorically strong on education, AI hubs, supercomputing and the importance of research, and it was right to be. The question is whether the institutional arrangements it proposed lower or raise the cost of access to useful knowledge. Several provisions, taken individually, can be defended; taken together, they raise it. Mandatory watermarking of language-model outputs, sovereignty framing on data, licensing for high-risk deployment, preferential procurement and a cautious posture on cross-border data flows describe a more closed environment for the free movement of ideas than the one in which AI's underlying research has thrived. Gans (2026) makes the parallel point on intellectual property: broad rules set in advance to prohibit AI training on certain data perform worse than rules that allow training to proceed and compensate harmed parties after the fact, once the actual harm is identifiable. Pre-clearing every use of every dataset is costly in time and money, and the welfare cost of suppressing useful knowledge flows is larger than the welfare cost of letting the flows proceed and remedying harms when they appear.

South Africa should therefore fund a small number of research centres at scale and protect their autonomy across political cycles. It should keep visa, mobility and intellectual-freedom rules generous enough to retain and attract the small number of frontier researchers who will determine whether the country produces original AI work or only consumes imported tools. It should

default, on copyright, to tracking how content is used and compensating harmed creators after the fact, rather than blocking training in advance. And it should stop treating local data ownership as a sufficient strategy for AI sovereignty. Korinek and Vipra (2025) argue that processing power, talent and data are jointly the binding constraint on building a competitive domestic AI ecosystem; policies that secure only one of the three will not produce a competitive ecosystem.

Govern less, but govern it properly

The draft proposed at least eight new bodies. Their usefulness in principle is not the question; the question is whether their constitutional rules (the second-order rules that govern how each body operates: how members are appointed and removed, how budgets are set, when the body sunsets, and how its decisions can be appealed) are specified clearly enough for the bodies to be competent, accountable and resistant to capture once they exist. In the draft as published, most of these rules of operation were deferred. Metcalf (2026) shows that AI safety regulation is a domain in which the risk of regulatory capture (the well-documented pattern in which oversight bodies end up working in the interest of the firms they regulate rather than the public's) is structurally high, because the technical complexity of compliance gives well-resourced incumbents an in-built advantage in shaping the rules. He argues that anti-capture design must be built in from the start: rotating leadership, transparent agenda-setting, published meeting minutes, independent review of standards, and tiered compliance burdens that do not entrench scale advantages under the banner of safety. The draft included none of these.

Hadfield and Koh (2025) make a parallel point about institutional infrastructure. AI systems that act autonomously on a user's behalf, what the literature calls agentic AI, raise a basic accountability question: when such a system makes a decision, the chain of responsibility must be clear. An ethics board can declare principles, but only a legal regime that registers

each deployed AI system, ties it to a named principal, and provides for accountability in contract and tort makes those principles enforceable.

Competition policy bears on the same governance question. Korinek and Vipra (2025) and Athey and Scott Morton (2025) both argue that AI markets concentrate even without any anti-competitive behaviour, because of three features that go together: very large fixed costs of training a model, cost advantages that grow with size, and high costs for customers who switch providers. The empirical evidence on whether these structural features have actually produced higher market concentration in AI markets is, however, mixed. Babina (2026) reviews recent work showing that average AI adoption has not yet produced a systematic rise in markups, and that generative AI may give smaller and younger firms a means of challenging incumbents. The structural concern is therefore real, but the welfare case for new licensing regimes remains weak. Athey and Scott Morton (2025) add that workers displaced by AI can be harmed twice: first by being substituted out of their jobs and then by the high prices charged for the AI services that displaced them. The policy implication is more aggressive use of the tools the Competition Commission already has, the market-inquiry function and the existing Competition Act, at the upstream layers of AI, and public procurement that keeps the state's right to switch suppliers and leave contracts. New licensing regimes or new bodies are not required.

Sequence the commitments

A policy document for a fast-moving technology can compensate for uncertainty in two ways. The first is to reduce the number of commitments at the start, sequence them visibly, and add to them once each has been delivered. The second is to commit to many things at once and discover during implementation which can be supported. The draft chooses the second route. The historical record on this choice in South African policy is poor.

Jin et al. (2026) formalise the regulatory side of this. They show that when a regulator uses AI to police firms that are themselves using AI, more

monitoring only reduces harm if the regulator's capacity to recognise and act on patterns crosses a certain threshold. Below that threshold, the firms keep redesigning their behaviour faster than the regulator can keep up, and the regulator ends up with more cases on its books but no less harm in the world. The implication is general: in a domain where the state cannot keep up with the technology, building more state does not mean more enforcement. The Road Accident Fund analogy in the AI Insurance Superfund proposal raises immediate concerns. The binding constraint on South African AI policy is the state's implementation capacity. The draft's response was to build more state. A more useful response is to build a narrower state, sequence its commitments, and expand only after the narrow ones have demonstrably worked.

In practice that means committing to a small number of targeted, narrowly defined interventions, in the spirit of evidence-led industrial policy (Juhász et al., 2024). Where the government believes a genuine market failure justifies state action, the appropriate response is a focused project with a clear theory of change, a budget, a deadline and an independent evaluation, rather than a wide architecture of supervision. Examples of what that could look like include an AI-assisted clinical decision-support system in a public hospital; an isiZulu or Sesotho language resource built jointly with a university group that already has half of it on a hard drive; or an AI-assisted revenue-administration tool at SARS. None of these requires new legislation. Each, if carefully evaluated, generates the kind of evidence the rest of the policy needs. Each success makes the case for further commitments stronger than additional planning text in the policy document.

Reallocation, not aspiration

Most of the productivity gains from past general-purpose technologies came from reallocation: workers and capital moving from less productive uses to more productive ones, firms reorganising around the new tool, and incumbents being displaced by entrants that used the tool more effectively. In every case the measured contribution of the technology itself was smaller

than the contribution of the reorganisation it forced. Agrawal et al. (2019) argue that the policies that determine whether reallocation happens are labour, trade, liability, privacy and competition policy: the policies governing how easily things and people can move. Hadfield and Koh (2025) extend the point to agentic markets. AI agents can change market and firm boundaries, but only if the contracting and accountability infrastructure exists to let them do so.

South Africa's reallocation problem is not new. Exchange-control rules that penalise software intellectual property, spectrum allocation, the mobility of skilled labour across firms and provinces, labour laws that raise the cost of hiring and adjusting staff as work reorganises around new tools, the cost and speed of cross-border data transfer, the depth of local venture capital: each of these is a friction that suppresses the reallocation through which productivity gains usually arrive. The withdrawn draft mentioned some of these in passing while devoting most of its pages to the proposed institutions; the next draft can correct that imbalance.

Consider a small South African firm hiring its first AI-using employee, where the realised gains documented earlier suggest a productivity uplift of five per cent or more. For that gain to translate into a measurable economic outcome, several things must work: the firm needs the bandwidth to make AI tools usable in the office and on the road; the owner needs to be able to pay for foreign software subscriptions without a six-week exchange-control file at the Reserve Bank; and the firm needs to be able to hire and let go staff as work reorganises around the tool. None of these requirements is a regulatory novelty introduced by AI. Each is a long-standing feature of the South African economy that the draft policy does not propose to change.

The cost of getting the ratio wrong falls hardest on small firms: the same compliance regime that is modest for a multinational with a compliance department is, for a Cape Town start-up of three engineers, prohibitive. A policy that raises fixed compliance costs without thinking about scale incidence selects against the firms the policy is intended to support.

What an adoption-first AI policy looks like

The five threads of the previous section together yield a five-point programme that the next AI policy could put into practice without new legislation, drawing only on the implementation pathways that already exist.

1. Roll out the digital infrastructure AI adoption depends on. Competitively priced bandwidth, last-mile connectivity, exchange-control rules that do not penalise software and intellectual-property transactions, and spectrum policy aligned with mobile-first AI use. None of this is in the policy as a binding commitment; all of it is closer to the binding constraint on whether South African firms and households actually capture AI's gains.

2. Fewer, narrower institutions, with constitutional rules from day one. Halve the institutional count. Fold the ethics, safety and ombudsperson functions into existing bodies where the mandate exists. For each remaining body, publish in the primary policy document the appointment procedure, the removal procedure, the conflict-of-interest disclosure regime, the budget cap, the sunset clause and the performance standards against which the body will be evaluated (Metcalf, 2026). Do not defer these to subsequent regulations that the affected bodies will help to write. The legitimacy of any regulatory regime sits in those second-order rules; treating them as administrative detail leaves the regime open to capture.

3. Targeted interventions, evaluated like industrial policy. Where the government believes there is a real market failure (a situation in which markets, left alone, will not deliver a good outcome) that justifies state action, target it narrowly rather than across the economy. Commit the funding. Set the deadline. Publish the independent evaluation. Possible candidates include a clinical decision-support deployment in a public hospital,

an indigenous-language resource and an AI-assisted revenue tool at SARS (Agrawal et al., 2026).

4. A national measurement capability, not a new monitoring centre. Build on the BER survey approach by establishing a permanent national instrument for measuring AI adoption. Combine it with StatsSA business and household surveys, SARS payroll data and a small set of administrative-data linkages. Build the dashboards inside StatsSA. Three years of disciplined measurement would put South Africa in a stronger evidentiary position on AI than almost any other middle-income economy. Jin et al. (2026) warn against the alternative: when regulators deploy AI to monitor firms but cannot keep up with how those firms adapt, caseloads rise without firm behaviour changing. A measurement capability differs from a monitoring centre in that its job is to inform, not to enforce.

5. Embed AI into existing programmes, not new ones. The temptation in any technology policy is to create new courses, new qualifications and new accreditation regimes for the new tool. The historical record on this in South Africa is not encouraging: a new TVET or university qualification typically takes years to design, accredit and roll out, and AI is moving faster than any of those processes. The faster route is to incorporate AI use into every existing university and TVET programme, vocational stream and in-service apprenticeship, and to allow lecturers and curriculum designers wide latitude to experiment with how it is used. The problem here is not that markets are failing to deliver education; it is that the state's quality-assurance institutions, designed for a slower world, will themselves prevent the educational system from adapting at the speed of the technology its graduates will be expected to use (Agrawal et al., 2026). Pair this with public-sector pilots that are randomised and evaluated, public procurement that keeps the right to switch suppliers, and modest, calibrated support for the consultancies, integrators and small firms that help organisations reorganise around new tools.

A sixth recommendation, smaller in scope, follows from Korinek and Vipra (2025) and Athey and Scott Morton (2025): give the Competition Commission a clear mandate to watch for concentration among the firms

that build foundation models and supply cloud computing and processing power, conduct a market inquiry into access to those services, and write procurement rules that prevent the state from being locked into a single provider. The recommendation lives inside the existing Competition Act and the existing Commission. It does not require new bodies. The competition-policy argument in this report is therefore narrower than its strong form: use what already exists, well, before building more.

Conclusion

In the BER's 2026Q1 survey, 82 per cent of managers and professionals expected AI to add five per cent or more to their productivity over the next three years. Fifty-seven per cent of private individuals expected the same. These figures are already shaping behaviour: households are paying for AI subscriptions, firms are training staff, and 40 per cent of consumers are using AI weekly to make purchasing decisions. The policy's role is to ensure that those gains translate into broad-based prosperity, not to insert administrative friction between users and tools that are already producing measurable gains.

The drafters of the withdrawn policy took AI seriously, and the Minister's withdrawal of the draft was itself an act of accountability worth acknowledging. The next draft must keep what was right about the first (the seriousness, the governance instinct, the desire for fairness and protection) and let go of what was wrong. The lesson the withdrawal makes plain is that what really limits South Africa's ability to capture the value of AI is not the absence of new institutions to enforce ethical use, but the presence of older institutions that already throttle bandwidth, slow cross-border payments for software and starve the country of fair access to cloud computing and processing power. Govern less, but govern properly. Set out the rules of operation for each new body before it exists; commit narrowly and sequence visibly; measure before supervising. Empowerment, in this domain, is when the frictions that inhibit deployment of AI are reduced so that South Africans can actually use the technology that is already producing measurable gains for

them. The test for any further intervention is the proposition this report has defended throughout:

> *A national AI policy will be judged by whether it equips and empowers South Africans to put the defining general-purpose technology of this era to work to build a more prosperous future, or stands between them and it.*

References

Agrawal, A., Gans, J. S. and Goldfarb, A. (2019), Economic policy for artificial intelligence, *in* J. Lerner and S. Stern, eds, 'Innovation Policy and the Economy', Vol. 19, University of Chicago Press, Chicago, pp. 139–159.

Agrawal, A. K., McHale, J. and Oettl, A. (2026), AI in science, NBER Working Paper 34953, National Bureau of Economic Research.

Anthropic (2026), 'Anthropic Economic Index: Claude.ai usage data, week of 2026-02-05 to 2026-02-12', Data documentation accompanying the fourth AEI report.

Athey, S. and Scott Morton, F. (2025), Artificial intelligence, competition, and welfare, NBER Working Paper 34444, National Bureau of Economic Research. Forthcoming as a chapter in *The Economics of Transformative AI*.

Babina, T. (2026), Understanding firms' AI efforts and their economic impact, NBER Working Paper 35123, National Bureau of Economic Research.

Brynjolfsson, E., Collis, A., Eggers, F., Kazinnik, S. and Nguyen, D. (2026), What is generative AI worth?, Technical report, Stanford University Digital Economy Lab. Working paper, 13 April 2026.

Bureau for Economic Research (2026), 'Survey results: Outlook and AI impact, 2026Q1', Stellenbosch University.

Department of Communications and Digital Technologies (2026*a*), 'Draft South Africa National Artificial Intelligence (AI) Policy', Government Gazette No. 54477, Vol. 730, Notice 3880 of 2026. Approved by Cabinet 25 March 2026; gazetted for public comment 10 April 2026; withdrawn by the Minister on 26 April 2026.

Department of Communications and Digital Technologies (2026*b*), 'Statement by minister Solly Malatsi on the integrity of the Draft National Artificial Intelligence Policy', Media statement, Pretoria. 26 April 2026.

Gans, J. S. (2026), 'Copyright policy options for generative artificial intelligence', *Journal of Law and Economics* **69**(1), 1–19.

Hadfield, G. K. and Koh, A. (2025), An economy of AI agents, *in* 'The Economics of Transformative AI', University of Chicago Press / NBER, Chicago, chapter 5. Preliminary draft, October 2, 2025.

Jin, G. Z., Sokol, D. D. and Wagman, L. (2026), Adaptive enforcement with AI-augmented monitoring, NBER Working Paper 35010, National Bureau of Economic Research.

Juhász, R., Lane, N. and Rodrik, D. (2024), 'The new economics of industrial policy', *Annual Review of Economics* **16**(1), 213–242.

Korinek, A. and Vipra, J. (2025), 'Concentrating intelligence: Scaling and market structure in artificial intelligence', *Economic Policy* **40**(121), 225–256.

Metcalf, T. (2026), 'AI safety and regulatory capture', *AI & Society* **41**, 2451–2466. Published online 3 August 2025.

Visa and Discovery Bank (2026), 'SpendTrend26: A look at South African consumer spending', Discovery Bank, Sandton.