

New Zealand Economic Opportunity Report 2026

Methodology Note

AI's impact on innovation

In order to estimate the overall economic opportunity from AI, we:

- Draw on the US [O*Net occupation database](#), which contains information on 51 different types of work activity for around ~800 types of occupations
- Use a series of LLM categorisations to classify the extent to which each work activity is exposed to automation or augmentation by generative AI. This allows us to estimate the proportion of tasks in each occupation that are susceptible to automation.
- Aggregate these results into broader economic categories based on each occupation's share of New Zealand's employment and average wage bill, and create our own crosswalk to convert the results from each O*Net occupation to the corresponding occupation in ISCO-08.
- Aggregate by wagebill, occupation and sector to produce an estimate of the total possible improvement in labour productivity.
- Assume capital intensity remains constant, and convert this labour productivity improvement into an overall improvement in GVA.

We then use this model to estimate the ability of AI to augment R&D-relevant tasks and occupations in particular, aggregating up to a global and country-specific level by adjusting for the different R&D workforce composition in each country and existing levels of R&D activity.

For our global baseline, we draw on [WIPO's Global Innovation Index](#) estimate of total R&D spending and rebase it into constant 2025 USD using [BLS CPI-U data](#), projecting it forward to 2040 at a 2.5% real CAGR. We apply a 60% labour share to isolate the wage-bill component of R&D expenditure that is susceptible to AI-driven productivity gains, anchored to [OECD MSTI](#) and [NSF/NCSSES](#) estimates of the labour share of business and total R&D.

In order to convert R&D output, we accumulate additional R&D output into a knowledge stock with a five-year innovation lag and a 13.7% annual depreciation rate — both blended from separate values for business R&D (four-year lag, 15% depreciation) and public R&D (7.5-year lag, 10% depreciation) using

OECD GERD sector shares, drawing on [Hall, Mairesse & Mohnen \(2010\)](#) and [Pakes & Schankerman \(1984\)](#).

For the country decomposition, our modelled gains are re-weighted using each country's own R&D workforce composition, drawn from [OECD MSTI](#) for OECD members and [UNESCO UIS](#) for the rest, supplemented by national statistical sources for Taiwan, Hong Kong, Thailand and Malaysia. Each country's GDP-lens value is split into a domestic-return component, controlled by a retention rate tiered from 50% to 75% by trade openness (anchored to [Coe & Helpman \(1995\)](#), [Eaton & Kortum \(1999\)](#), and [Myers & Lanahan \(2022\)](#)), and an absorbed-spillover component.

Value from encouraging SMB adoption

We estimate the unrealised value from SMBs adopting AI at the same rate as large enterprises by starting with the total AI opportunity in each market from our AI opportunity model.

We allocate a share of this opportunity to SMBs using their structural share of GVA. We draw on national and regional sources for SMB value-added shares, including APEC data and national statistical or government sources.

We then estimate the current AI adoption gap between SMBs and large enterprises. Where available, we use country-level survey data on AI adoption by firm size from IMDA, IBM, and the OECD. For markets without direct country-level data, we apply peer-group benchmark gaps.

To reflect the greater potential upside in less digitally mature markets, we add a digital catch-up premium. This uses a digital intensity index built from country-level indicators including internet access, digital payments, e-commerce, small-firm digital operations, and mobile connectivity. The premium is calibrated using evidence from [Brynjolfsson et al. \(2025\)](#), which finds that lower-skilled workers gained disproportionately from AI adoption. We use this as indicative evidence that businesses and markets starting further behind may have greater catch-up potential from adopting AI, while bounding the premium to reflect institutional and infrastructural constraints.

Finally, we convert the unrealised economic value into labour equivalents by dividing by GVA per worker, drawing on ILO employment data and World Bank GDP data.

ROI from AI adoption

We estimate the return on AI investment for each sector by comparing the realisable productivity opportunity from AI with projected levels of AI spending.

On the opportunity side, we estimate the labour productivity uplift from our AI opportunity model, distributing these gains across sectors using a Eurostat staffing patterns crosswalk on the relationship between occupations and industries.

We convert this into a sector-level productivity opportunity by applying each sector's GVA and labour share. We use national statistics, UN, and OECD data to estimate sector GVA, and draw on national



statistics, OECD, and World KLEMS data to estimate sector-specific labour shares. We then allow adoption to rise over time in our modelling, drawn on observed adoption rates over time from Eurostat and BCG industry adoption data.

On the spend side, we start with IDC's forecasts for APAC AI spending up to 2028, allocating it across markets using GDP and the IMF AI Preparedness Index score, and then across sectors based on their share of national GVA. Spending beyond 2028 is projected using historical patterns from previous technology waves, such as ecommerce and cloud computing.

Finally, cumulative ROI is calculated as the cumulative realisable AI opportunity between 2025 and 2035 minus cumulative AI spend, divided by cumulative AI spend.

New jobs from AI

We combine estimates for three sources of potential AI related employment:

- The first is the core AI industry (software and services). For AI services and integration, we use [Gartner's](#) global AI Services market sizing, an APAC demand share calibrated to [IDC's APAC AI spending track](#), and revenue-per-employee benchmarks of around \$50–60K from major IT services firms ([TCS](#), [Infosys](#)), with an adjustment for India. For AI software and platform companies, we use Gartner's parallel AI Software market and a higher revenue-per-employee benchmark calibrated to APAC AI product firms([SenseTime](#), [iFlytek](#), [4Paradigm](#)).
- The second is AI-related infrastructure spending in data centres and chips. We estimate this through a stock-flow capex model grounded in TSMC and Samsung Device Solutions capex disclosures, [SEMI global semiconductor equipment billings](#), and jobs-per-\$B benchmarks from [TSMC's Arizona project](#) and [IMF Working Paper WP/21/131](#) on infrastructure employment multipliers.
- The third is potential jobs at AI-enabled new businesses. We extrapolate this from the ratio between core services / infrastructure jobs and product jobs in the mature internet economy, applying a multiplier to the core AI industry total drawn from the [IAB/Harvard Digital Economy Study \(2025\)](#), which reports both direct and indirect digital economy jobs in the US, cross-checking this against [GSMA's APAC mobile economy](#) and [CompTIA's tech workforce](#) figures.

Our estimate is deliberately conservative. It does not include wider induced demand for other occupations in the economy, which is likely to be orders of magnitude larger.

Economic Impact of Google Products

We estimate the economic impact of Google's products and services as the direct economic activity generated for businesses, publishers, creators and developers through Google Search, Ads, AdSense, YouTube, Play, Cloud and Workspace.



We estimate the total impact by summing the impact of the following products.

Search and Ads. We multiply the size of the paid search advertising market by Google's share of the search engine market and by an estimated return on advertising spend. We draw on third-party estimates of paid search advertising market size from PwC's Global Entertainment, Media & Telecoms Outlook, Statista Market Insights, and eMarketer, and use StatCounter data for Google's search market share. We apply an 8x return on advertising spend, based on [Google's methodology](#).

Cloud and Workspace. We multiplied the total market size for cloud and workspace products by Google's share of the cloud market. We draw on Statista Market Insights for both the total cloud market size and Google Cloud's market share.

AdSense. We estimate publisher earnings from displaying Google-served ads on third-party websites and apps by taking Google Network revenues from Alphabet earnings reports and multiplying by the share paid to partner publishers through traffic acquisition costs. We then allocate this global publisher income to countries using each country's share of global display advertising spend, drawing on PwC's Global Entertainment, Media & Telecoms Outlook data.

YouTube. We estimate creator earnings from YouTube advertising by taking YouTube Ads revenue from Alphabet quarterly earnings reports and multiplying by YouTube's creator revenue share. We use a default creator share of 55%, based on [YouTube's creator economy documentation](#). We then allocate global creator earnings to countries using each country's share of global video advertising spend, drawing on PwC's Global Entertainment, Media & Telecoms Outlook data.

Google Play. We draw on Sensor Tower data for the app revenue earned by developers through the Play ecosystem.

Economic Impact of Google Products for SMBs

We multiply the economic impact of Google products in each market by an SMB share estimated by combining the economic contribution of SMBs to a market's economy with their digital intensity.

We draw on official national statistics from government or regional economic institutions for SMB contribution to national GVA. These include APEC Policy Support Unit data and OECD SME Policy Index in addition to SMB value-added shares cited by national labour ministries.

We also construct a digital intensity index for each market. This combines five country-level indicators: internet connectivity, broadband access, small-firm digital operations, digital payments usage, and mobile connectivity quality or affordability. We draw on World Bank WDI indicators for internet and broadband access, World Bank Enterprise Surveys for small-firm digital behaviour, Global Findex for digital payments, and the GSMA Mobile Connectivity Index for mobile connectivity.



Jobs from Google Products

We multiply the economic impact of Google products in each market by its labour share of income, drawing on Penn World Tables data to estimate the labour-attributable portion of this activity.

We then estimate GVA per worker in each market, drawing on World Bank and IMF World Economic Outlook data on GVA and GDP as well as ILO data on labour force.

Finally, we divide the labour-attributable impact by GVA per worker to estimate the number of jobs supported.

Exports from Google Products

We estimate the value of exports enabled by Google products by estimating the following export shares to multiply with the economic impact of the relevant Google products:

Search and Ads, Cloud and Workspace, AdSense. We define export share as commercial service exports divided by services GVA, drawn from World Bank data on service exports and services value added. This captures the share of a country's service economy that is outward-facing, and is used as a proxy for the share of Google-enabled activity likely to generate revenue from overseas customers.

We use this same service export ratio for Search and Ads, AdSense and Cloud because all three products primarily support digitally enabled service exporters, including e-commerce sellers, digital marketers, SaaS firms, IT outsourcers, publishers and cloud-enabled platforms. The products differ in the size of their activity base, but use the same export-intensity proxy.

YouTube. We draw on data from the YouTube API to estimate the share of a country's YouTube activity attributable to overseas viewers of domestically produced content, using this as a proxy for export share.

Play. We draw on Sensor Tower data for the proportion of app revenue earned by developers from consumers in other countries.

Consumer Surplus from Google Products

Following [Brynjolfsson et al \(2019\)](#), we used a 'willingness to accept' framing to model the current default hypothetical consumers face. As part of our polling, we asked participants a single discrete binary choice question of "Would you prefer to keep access to [product] or go without access to [product] for one month and get paid [price]" with the price offered randomised between set levels per country.

We regressed the results of this poll to derive a demand curve and used this to calculate total consumer surplus per user. We also scaled this estimate by third party estimates of Internet prevalence and polling information on product usage by country. In order to reduce noise, we regressed our country estimates against wider macroeconomic indicators like GDP per capita, and used this to produce our final smoothed per country estimate.



Potential lives saved from AI discovered drugs

We estimate the time saved by AI in drug discovery by taking an estimate from the literature for the share of the development lifecycle that AI can compress, drawing primarily on [Agarwal et al. \(2026\)](#). We apply this to region-specific development timelines from [Paul et al. \(2010\)](#), with a longer baseline for markets with a neglected tropical disease burden, based on [Stefanakis et al. \(2012\)](#)'s evidence on longer NTD (Neglected Tropical Disease) development timelines.

We then estimate the number of additional treatments that could reach market by 2040. We start from IQVIA data on recent global novel active substance launches, and apply an output multiplier that reflects shorter development cycles. The multiplier is adjusted over time using AI adoption assumptions drawn from Nvidia's healthcare AI adoption data and evidence from Tufts on AI implementation in clinical trials.

Finally, we convert additional treatments into lives saved. We use [Lichtenberg's \(2019\)](#) estimate of life-years gained per new drug, then convert life-years into lives using a range of assumptions for life-years saved per death averted, drawing on estimates from JAMA Health Forum and The Lancet.

Acceleration of product development

We estimate how much faster start-ups could release new product features by modelling AI's impact across the software development lifecycle.

We first divide the lifecycle into development, QA and testing, design, and project management, using effort-allocation benchmarks from Hypersense Software. We then assign AI speedup assumptions to each stage. For development and QA, we use our AI coding agents methodology, which draws on developer task-share evidence from Atlassian and TideLift / New Stack, with lower-bound adjustments informed by METR. For design, we draw on evidence from Cieden, IDEO / Parallel and Startup House on AI-assisted prototyping and design workflows, while reflecting Nielsen Norman Group's caution that AI-generated designs are not always production-ready. For project management and requirements work, we draw on [Noy and Zhang \(2023\)](#), [Dell'Acqua et al. \(2026\)](#), and Startup House evidence on AI-assisted drafting, synthesis and discovery workflows.

We combine these stage-level effects using [Amdahl's Law](#), so the overall speedup reflects both the size of each stage and the AI speedup within it.

Finally, we adjust the base speedup by country using projected AI adoption in 2035 from the Public First AI adoption model, which combines Google Patents data on AI innovation, Sensor Tower data on consumer AI app usage, and data from various AI model-builders on the use of their consumer and enterprise AI tools. This reflects the extent to which firms in each market are likely to realise the potential speedup from AI tools.

Economic value of generative media for the creative sector

We use our AI opportunity model to estimate the productivity impact of generative media tools, drawing on O*NET task data across 26 core creative occupations as well as creative tasks performed throughout

the wider economy. We use a series of LLM classifications to assess both the potential for AI augmentation and constraining factors such as sensitivity, creativity requirements, or intrinsic human value, then aggregate up to sectors and occupations weighted by wagebill, mapping across to international ISCO-08 codes and applying an S-curve adoption model through to 2035. In the second stage, we layer in potential demand effects from new forms, improved distribution and wider economic growth, accounting for price effects, income elasticity, and viability effects as lower costs make marginal creative projects more feasible. We triangulate between historical analogues such as UGC and mobile gaming and a bottom-up estimate of new forms to give a conservative overall figure.

AI-powered personalised advertising

We estimate the additional brand value from AI-powered personalised advertising by applying channel-specific AI performance uplifts to digital advertising spend.

We start with digital advertising spend data from Statista, categorising this into four channels: Search, Display, Direct Messaging and Other. We then assign an AI-driven uplift to each channel from the closest available academic evidence. For Search, we use [Relsenbichler et al. 's \(2025\)](#) estimate of the conversion uplift from applying LLMs to search advertising. For Display, we use [Lee et al. 's \(2025\)](#) estimate of the click-through uplift from visual generative AI, using the lower-bound estimate. For Direct Messaging, we use [Kapoor and Kumar's \(2025\)](#) estimate of engagement gains from personalised WhatsApp video ads. For Other channels, we use [Tian et al. 's \(2025\)](#) estimate of the ROI uplift from AI targeting in e-commerce.

We calculate the overall uplift in effectiveness in each market by weighting these channel-level uplifts by each channel's share of digital ad spend. We then convert the uplift into additional brand revenue using channel-level ROAS, drawing on benchmark evidence from Nielsen, Triple Whale, and WhatConverts.

Saving time for doctors

We draw on our AI opportunity model to estimate the country-specific productivity uplift for medical doctors from AI. We take the automatability of doctors' tasks at the occupation level (ISCO-08 group 221, generalist and specialist medical practitioners), wage-bill-weighted using New Zealander employment and wages, focusing on the administrative and documentation tasks that AI can take on rather than direct clinical care.

We then draw on ILO data on hours worked and employment in the health sector, isolating doctors' hours using the share of health-sector employment accounted for by medical practitioners.

We multiply these together and adjust by the health-sector labour share of income (ILO), so the estimate captures the share of AI-enabled productivity gains that accrue as freed clinician time rather than being captured as reduced costs. Finally, we convert the aggregate into hours saved per doctor per week. We cross-check the scale of the administrative burden against time-and-motion evidence on physician documentation load from [Sinsky et al. \(2016\)](#) and [Woolhandler & Himmelstein \(2014\)](#).



AI-enabled diagnosis

We identify a fully-sourced core of high-mortality chronic disease groups, type 2 diabetes, cardiovascular disease, and four cancers (lung, colorectal, breast and prostate), selected for high national mortality burden, avoidable late-stage detection, and a published mortality-reduction figure the AI can be shown to act on. For each disease we take the annual mortality pool from Health New Zealand Mortality Collection and Cancer Web Tool data, then apply an AI clinical shift rate (drawn from deployed AI clinical decision support and intervention trials) and a per-disease mortality gain ratio anchored to a published AI-deployment study, with peer-reviewed surrogate-to-mortality mappings used where mortality endpoints are not directly measured.

Boosting public sector efficiency

We estimate the quantity and value of public sector time that AI could free up by modelling the occupational composition of the public sector and combining this with our AI opportunity model. We first establish for each occupation the share of workers employed in the public sector, using the [US BLS OEWS](#) NAICS-by-SOC crosswalk to compute government employment as a fraction of total employment per occupation, crosswalking SOC to ISCO-08 and applying those shares to the country's occupation-level wage-bill and employment data. Where public sectors provide different offers relative to the US, e.g., healthcare, adjust using sector-specific datasets like the [WHO Global Health Expenditure Database](#). We then draw on our AI opportunity model to estimate the potential productivity uplift from AI in each occupation, summing to express the result as a share of the public-sector wage bill and as a monetary figure, with hours freed computed in parallel from public-sector worker counts and ILO annual hours.

Saving time for teachers

We draw on our AI opportunity model to estimate the country-specific productivity uplift for primary and secondary school teachers from AI.

We then draw on Ministry of Manpower and ILO data on hours worked and employment in the education sector, isolating primary and secondary teacher hours using the share of education-sector employment accounted for by these professions.

We multiply these together and adjust by labour share of income using ILO data so the estimate captures the share of AI-enabled productivity gains that accrue as worker-side benefit rather than being captured as reduced costs or other capital-side gains. Finally, we convert aggregate freed hours into two outputs: total annual hours saved across teachers, and hours saved per teacher per week.

Easing cost of living pressures

We estimate the annual saving AI tools could unlock for the average New Zealand household by building the figure bottom-up from eight household-level deal-finding and switching channels, spanning

insurance, mortgages, groceries, fuel, electricity, subscriptions, telecoms and smart-meter energy use. We select these channels by walking New Zealand's own largest household costs in the Stats NZ Household Expenditure Statistics and keeping only those with a published, measurable AI savings effect. For each channel we multiply a published saving per affected household by the share of households it applies to and, where available, the observed annual rate at which households act. We sum these into a total annual saving per household and express it as a share of median household disposable income from Stats NZ.

AI-enabled energy efficiency for households and businesses

We estimate the gain in efficiency that AI could deliver across New Zealand's household and business energy use by scoping our claim to buildings rather than the whole economy. We take the IEA's projected AI-driven savings rates for residential and non-residential (commercial) buildings by 2035, under its widespread-adoption case, from its Energy and AI report, and weight them by New Zealand's actual residential and commercial energy consumption from MBIE's Energy in New Zealand 2025. This gives a consumption-weighted household-and-business efficiency gain, which we also express as absolute energy saved by converting from petajoules to terawatt-hours. Finally, we translate the residential and commercial savings rates into indicative annual electricity-bill savings for a typical household and a small business, using average New Zealand household and business electricity spend.

AI- improvements to farming productivity

We estimate the productivity gain AI could deliver to New Zealand's pastoral-livestock sector, scoping the claim to dairy and red meat, which make up the large majority of New Zealand agriculture by value. We set a per-adopter productivity uplift by reviewing the peer-reviewed and industry evidence on animal health, feed, breeding, grazing and loss reduction, taking the median of the comparable output-uplift estimates as a base rate and anchoring on the cleanest causal whole-farm study. We apply the uplift to current production, taken from FAOSTAT via Our World in Data and from DairyNZ for milk-solids, to give the additional annual output. As supporting context we compare the uplift to the gap between the average and top-performing New Zealand farms, from the DairyNZ Economic Survey and the Beef and Lamb New Zealand Farm Survey, to show what share of that management gap it would close, and we cross-check the uplift against our AI labour-productivity model applied to livestock-producer occupations.

Polling claims

Polling claims are derived from a survey of 1,036 online adults based in New Zealand in June and July 2026, conducted in English. All results are weighted using Iterative Proportional Fitting, or 'Raking'. The results are weighted by age group, gender, and education level to nationally representative proportions.

We used a range of different panel providers who contacted respondents on our behalf. In return for their participation in our survey, respondents were provided with a financial incentive.



Like all polling data, market research is susceptible to poor memory or consumers not answering truthfully. In order to reduce the risk of this, we completed a number of standard quality checks on the polling data to help ensure that respondents are paying attention:

- Excluding respondents who take too long to answer;
- Excluding respondents who 'straight-line', eg. always picking the top or left most option to every question;
- Excluding respondents who fail an attention check, eg in the middle of a longer question, we ask them to pick a particular option if they are reading;
- Excluding respondents whose answers all perfectly match another;
- Excluding respondents whose open text answers are incoherent or look like they have been generated by a computer bot.

