AI Career Risk *Index 2026*

The variance inside your job is bigger than the variance between jobs.

OECD, ILO, and Anthropic all agree AI's impact happens at the **task level**, not the job level. The Anthropic Economic Index publishes adjacent data — observed Claude usage by task — but no source publishes per-occupation **task-replaceability** scores under open licence. This report fills that gap: 334 occupations, 2,563 tasks, every score downloadable under CC BY 4.0. The macro-economists agree on the framework. Nobody publishes the per-task drill-down. We do.

45.0	***36.53%***	***64.12***
Median exposure score (0–100)	Occupations with bimodal task distribution	Mean within-occupation variance

Most jobs are reshaped, not replaced. The median occupation scores **45 out of 100** on AI exposure — and the spread of task-level capability scores inside that median job is **64 percentage points** wide. That single number is the report's core finding: when you break a job into its actual tasks, the within-job variance is bigger than the variance between jobs. In **36.5% of audited occupations** — roughly one in three — that spread is sharp enough to count as bimodal: at least one task scoring above 80% on AI exposure sitting alongside at least one task below 20%. Those are the jobs splitting cleanly in two: a set of tasks AI already does well, and a set it can't touch — with not much in between.

What's in this report: the methodology behind the scores, the variance pattern with worked examples, the 119 tasks already in the AI bullseye (≥85% replaceability), where the 334 occupations sit by tier, sector-level breakdowns, and time-to-impact windows. The full task-level scoring is downloadable as a CC BY 4.0 dataset.

The frame: why tasks, not jobs

Frey & Osborne changed the conversation in 2013 with one number — "about 47 per cent of total US employment at risk." They scored whole occupations. Every serious follow-up since has argued they were measuring the wrong unit.

The first major refinement came from Arntz, Gregory and Zierahn (OECD Working Paper 189, 2016). Their critique, in their own words: "the occupation-based approach used by Frey and Osborne might lead to an overestimation of job automatibility, as occupations labelled as high-risk occupations often still contain a substantial share of tasks that are hard to automate." Re-running the analysis at the task level rather than the occupation level dropped the share of US jobs at "high automation potential" from about 38% to roughly 9%. Same data, finer unit of analysis, different number.

That shift — from occupations to tasks as the unit of analysis — is now mainstream. The OECD Employment Outlook 2023 (chapter 3, "Artificial Intelligence and the Labour Market") adopts task-based methodology explicitly and finds that AI has progressed most in non-routine cognitive tasks: information ordering, memorisation, perceptual speed, deductive reasoning. The ILO 2025 Refined Global Index (joint with NASK) goes further: it scored 2,861 tasks across 1,640 surveyed workers, concluded "1 in 4 jobs globally are exposed to generative AI," and stated plainly that *"transformation of jobs is the most likely impact of GenAI."* Not replacement — transformation. Tasks getting reshuffled inside the same job title.

Anthropic's Economic Index (March 2026 release) works the same way from a different direction. Rather than scoring tasks for capability, it observes which tasks Claude is actually being used for, mapped onto O*NET's task taxonomy. Headline finding: "49% of jobs have seen at least a quarter of their tasks performed using Claude." Up from roughly 36% in the February 2025 report thirteen months earlier. The unit of measurement is, again, the task — not the job title. By 2026, the task-not-job framework is settled science. What is still missing — and what this report adds — is the per-occupation drill-down underneath it.

So where's the gap, and what's left to add? The consensus exists at the methodology level. What it does not yet include — anywhere we could find under a free, open licence — is the per-occupation, per-task drill-down: 334 specific job titles, each broken into its actual day-to-day tasks, each task scored individually. Anthropic's Hugging Face dataset comes closest — it publishes ~18,000 O*NET task scores under CC BY — but the metric is *observed Claude usage frequency*, not capability-based replaceability. It tells you how often Claude is being asked to do a task; it does not tell you whether current AI *can* do that task well. Those are different questions, and an individual worker trying to plan ahead needs the second one. OECD and ILO publish their analyses aggregated to the occupation level; the per-task scoring sits in working papers and internal datasets, not in a downloadable open-licence release.

The Index 2026 complements the macro work; it doesn't replace it. The aggregate work — 47%, 9%, 22%, 1-in-4, 49% — answers a policy question: what is AI doing to the labour market overall? This report answers a different one: what is AI doing to your job, this year, task by task? That is per-occupation, per-task, and it needs a per-occupation, per-task answer. That is what this Index publishes.

What the data shows: variance within jobs

The headline number from the audit is the within-job spread. Across all 334 fully-scored occupations, the mean gap between an occupation's most-replaceable task and its least-replaceable task is **64 percentage points**. The median is the same — 64 pp. That is the report's core finding in one number: when you break a job down into its actual tasks, the spread inside the job is, on average, almost two-thirds of the entire 0–100 scale.

Six worked examples, drawn straight from the audit. Same job title. Same person. Tasks scoring 80 percentage points apart on the same scale:

- **Bank Teller** — "Balance Enquiries & Statements" scores 95% replaceable; "Vulnerable Customer Support" scores 10%. Same job. Same person. **85 percentage points apart.**

- **Bookkeeper** — "Transaction Data Entry" 95% vs "Client Briefing & Advisory" 11%. 84 pp.
- **Paralegal** — "Draft legal documents" 91% vs "Court appearances & advocacy" 8%. 83 pp.
- **Accountant** — "Bank Reconciliation" 92% vs "Client Advisory & Relationship Management" 9%. 83 pp.
- **Photographer** — "Stock & Generic Product Photography" 90% vs "Event & Wedding Photography" 8%. 82 pp.
- **Customer Service Agent** — "Routine enquiry handling" 92% vs "Empathic & vulnerable customer support" 12%. 80 pp.

These are not cherry-picked outliers. They are the top of a distribution where the typical occupation has a 64 pp internal spread. "Is my job safe?" is the wrong question. The right one: which of my tasks survive the next five years, and what does the residual job look like once the high-replaceability tasks go to software? For most of the roles above, the answer is that the job itself continues to exist — but compressed, re-bundled toward the human-only end, and probably with fewer entry-level openings because the on-ramp tasks are precisely the ones that automate first.

The 1-in-3 bimodal split

A subset of the variance pattern is what we call **bimodal**: the occupation contains at least one task scoring ≥80% AI-replaceable *and* at least one task scoring ≤20%. **122 of the 334 audited occupations — 36.5%, or roughly one in three — meet that threshold.** These are the jobs where the internal split is sharp rather than smooth: there is a band of work that current AI can clearly do, and a band of work it clearly cannot, with relatively little in between. Bank Teller, Paralegal, Accountant, Bookkeeper, Customer Service Agent and Photographer (all above) are bimodal cases.

> **Threshold sensitivity.** The 80/20 cut is the headline definition. The bimodal share at adjacent thresholds, recomputed across the same 334 occupations: **282 occupations (84.4%) at 70/30**; **122 (36.5%) at 80/20**; **7 (2.1%) at 90/10**. The pattern is robust to the cut at the shoulders and tightens rapidly at the extreme. We use 80/20 because it identifies the jobs where the split is sharp without requiring near-perfect polarisation; the underlying variance finding (mean 64 pp spread) does not depend on this choice.

One thing to flag, because it matters: **the average occupation is not bimodal**. Most of the 334 jobs have a smooth distribution of task scores — some high, some low, with a continuous middle and no sharp gap. The 1-in-3 figure is meaningful because it identifies the roles where AI's effect on day-to-day work will feel most discontinuous: one set of tasks goes away quickly, another doesn't budge, and the experience of doing the job changes in a way that is harder to absorb gradually. But it is not the report's headline finding. The headline is the variance number — 64 pp on average — because that pattern holds across the whole dataset, not just the sharply-split third of it. Bimodal occupations are a specific (and important) expression of variance, not a separate phenomenon.

Two cautions before reading the rest of the report. First, "replaceable" here measures *capability*: whether current AI can perform a task to a useful standard. It does not measure whether your specific employer has deployed AI to do it, or whether they will. Cost, regulation, organisational change resistance, and customer preference all sit between capability and adoption — and they vary by industry, by employer, and by country.

Second, the scores are point-in-time. AI capability has shifted noticeably even in the thirteen months between Anthropic's Feb 2025 and March 2026 Economic Index releases. We will re-score and republish annually.

Top tasks at risk: where AI's actual reach lives

119 tasks across 334 occupations score at or above 85% AI-replaceability. That is one task in 22 — and they cluster in roles where the working day is structured around documents, rules, or scripted interactions.

The ranking sits at the task level, not the occupation level, for a methodological reason: occupation scores cluster tightly (most of the 334 audited occupations land within 2-3 points of their tier neighbours), but the task layer is genuinely fine-grained — 83 distinct replaceability percentages across 2,563 scored tasks. The three tier groups below walk through the high-exposure cluster from highest to lowest replaceability.

The 97-95% band: transactional, scripted, document-bound

The ceiling of the dataset. These are tasks where current AI matches or exceeds competent human performance on standard inputs: form transcription, balance enquiries, payroll calculation, scripted outbound calls, transaction data entry, password resets, CV screening. Most are characterised by structured input, rule-driven processing, and a defined output format. These are also the tasks that anchor junior and entry-level positions in clerical, banking and customer-service work — which is where the labour market has already started to feel it.

97–95% AI-replaceable (7 tasks)

TASK	OCCUPATION	SECTOR	%
Form & document transcription	Data Entry Clerk	Administration	97%
Balance Enquiries & Statements	Bank Teller	Finance & Banking	95%
Transaction Data Entry	Bookkeeper	Finance & Accounting	95%
Password Reset & Account Management	IT Support Analyst	Technology	95%
Payroll processing & calculation	Payroll Administrator	Human Resources	95%
CV screening & shortlisting	Recruiter	Human Resources	95%
Scripted Outbound Calling	Telemarketer	Sales & Customer	95%

The 94-90% band: routine knowledge work

One step in from the ceiling, and substantially broader. This band includes routine claims processing, standard risk assessment, case-law research, contract drafting from templates, programmatic media buying, regression test execution, and the document production end of legal and medical work. The pattern: tasks that require some domain knowledge but where the underlying work is pattern-matching against a corpus (legal precedent, prior cases, market data, a style guide). These are the tasks where large language models combined with retrieval are demonstrably good — and where the economic case for AI adoption is already being made, not waiting on capability.

94–90% AI-replaceable (38 tasks)

TASK	OCCUPATION	SECTOR	%
Item Scanning & Payment	Cashier	Hospitality & Personal Service	94%
Database record creation & update	Data Entry Clerk	Administration	94%
Bank Reconciliation	Bookkeeper	Finance & Accounting	93%
Invoice & order data processing	Data Entry Clerk	Administration	93%
Trade Execution — Retail	Stockbroker	Financial Services	93%
Bank Reconciliation	Accountant	Finance & Accounting	92%
Cash Deposits & Withdrawals	Bank Teller	Finance & Banking	92%
Tier-1 Query Resolution	Call Centre Agent	Sales & Customer	92%
Routine claims processing & settlement authority	Claims Adjuster	Financial Services	92%
Product & Category Page Copy	Content Writer	Creative & Design	92%
Automated Payment Reminders	Credit Controller	Finance & Accounting	92%
Routine enquiry handling (FAQs, account info)	Customer Service Agent	Sales & Customer	92%
Stock & Generic Illustration	Illustrator	Creative & Design	92%
Case law research & synthesis	Legal Researcher	Legal	92%
Document production & word processing	Legal Secretary	Legal	92%
Medical transcription & clinical note creation	Medical Secretary	Healthcare	92%
Tax & National Insurance calculations	Payroll Administrator	Human Resources	92%
Bid Strategy & Optimisation	PPC Specialist	Marketing	92%
Data Entry & Call Logging	Telemarketer	Sales & Customer	92%
Small Business & Template Website Design	Web Designer	Creative & Design	92%
Photographic vehicle & property damage assessment	Insurance Claims Adjuster	Financial Services	91%
Draft legal documents	Paralegal	Legal	91%
Fund Transfers & Bill Payments	Bank Teller	Finance & Banking	90%
Password Resets & Account Updates	Call Centre Agent	Sales & Customer	90%
Advertising & PPC Copy	Copywriter	Marketing	90%
Data entry & case tracking	Court Clerk	Legal	90%
Spreadsheet data entry & formatting	Data Entry Clerk	Administration	90%
Standard risk assessment & scoring	Insurance Underwriter	Financial Services	90%

TASK	OCCUPATION	SECTOR	%
Safety stock & reorder point calculation	Inventory Analyst	Supply Chain & Operations	90%
Breaking News & Wire Reporting	Journalist	Creative & Design	90%
Programmatic Digital Buying & Optimisation	Media Buyer	Marketing	90%
Stock / Library Music	Musician	Creative & Media	90%
Member Record Maintenance	Pension Administrator	Financial Services	90%
Stock & Generic Product Photography	Photographer	Creative & Design	90%
Regression Test Execution	QA Engineer	Technology	90%
Checkout & Payment Processing	Retail Assistant	Sales & Customer	90%
Lead Qualification	Telemarketer	Sales & Customer	90%
Personal Lines Risk Assessment	Underwriting Analyst	Financial Services	90%

The 89-85% band: structured but judgement-adjacent

The lower edge of the high-exposure cluster, and the broadest band. Tasks here still score very high but typically involve some judgement, edge-case handling, or stakeholder coordination that current AI handles well on average but inconsistently at the margin: copy editing, e-discovery review, network monitoring, salary benchmarking, social-media graphics, statutory analysis, scheduled reporting. Many of these are the tasks where the workflow is shifting toward human-supervised AI rather than replacement — the human is still in the loop, but doing review and exception handling rather than first-pass production.

89–85% AI-replaceable (74 tasks)

TASK	OCCUPATION	SECTOR	%
Transaction Categorisation	Accountant	Finance & Accounting	89%
Order status & delivery tracking queries	Customer Service Agent	Sales & Customer	89%
Background / Extra Work	Actor	Creative & Media	88%
Invoice Processing (AP/AR)	Bookkeeper	Finance & Accounting	88%
Transcription & Caption Production	Broadcast Journalist	Creative & Design	88%
Call Routing & Triage	Call Centre Agent	Sales & Customer	88%
Damage image & video assessment	Claims Adjuster	Financial Services	88%
Salary benchmarking & market data analysis	Compensation & Benefits Manager	Human Resources	88%
Contract drafting & templating	Contracts Manager	Legal	88%
Property searches & due diligence	Conveyancer	Legal	88%
Social Media Copy	Copywriter	Marketing	88%
Document filing & records management	Court Clerk	Legal	88%
Financial Statement Spreading	Credit Analyst	Finance & Accounting	88%

TASK	OCCUPATION	SECTOR	%
Data validation & accuracy checking	Data Entry Clerk	Administration	88%
Backup and Recovery Automation	Database Administrator	Technology	88%
Email Copywriting & Content Creation	Email Marketing Manager	Marketing	88%
Calendar & scheduling management	Executive Assistant	Administration	88%
Order Execution & Routing	Financial Trader	Financial Services	88%
Social Media Graphics & Templates	Graphic Designer	Creative & Design	88%
Policy pricing & premium rating	Insurance Underwriter	Financial Services	88%
Basic Troubleshooting via Scripts	IT Support Analyst	Technology	88%
Document Verification & ID Checking	KYC Analyst	Finance & Banking	88%
Statutory & regulatory analysis	Legal Researcher	Legal	88%
Campaign Pacing & Budget Management	Media Buyer	Marketing	88%
Appointment booking & patient scheduling	Medical Secretary	Healthcare	88%
Legal research	Paralegal	Legal	88%
Statutory payments (SSP, SMP, SPP)	Payroll Administrator	Human Resources	88%
Benefit Calculations	Pension Administrator	Financial Services	88%
Drug interaction & contraindication checking	Pharmacist	Healthcare	88%
Transcription & Show Notes	Podcast Producer	Creative & Design	88%

+44 more tasks in this band — see the tasks CSV for the full list.

One important caveat for anyone reading these scores as a forecast of their own working life: **"replaceable" measures capability, not adoption.** A task scoring 92% means current AI can perform it to a useful standard against the typical version of the task. It does not mean your specific workplace has deployed AI to do it, or that they will. Cost, regulation (especially in healthcare, financial services, and legal work), customer preference, organisational inertia, and the pace of procurement cycles all sit between "AI can do this" and "AI is doing this here." Those factors are outside the scope of this Index and vary too much by employer to score generically. What we measure is the underlying capability gradient — the conditions under which adoption becomes plausible — not the rate of adoption itself.

The full task-level scoring — all ~2,500 tasks across 334 occupations, with task name, occupation, sector, replaceability percentage, and time window — is downloadable as the tasks CSV under CC BY 4.0. Reuse is welcome with attribution to JobForesight.

Where the 334 occupations land

Of 334 occupations, **only two land in the Critical band (≥85)**: Data Entry Clerk and Telemarketer, both structured around transactional, scripted work. **30 sit in Very High (70-84)**, **43 in High (60-69)**, **148 in Moderate (40-59)**, and **111 in Low (<40)**. More than three-quarters of the labour market — 259 of 334 occupations, or 77.5% — sits in the Moderate or Low bands. This is not a story about mass replacement. It is a story of reshuffled day-to-day work, with a small leading edge of clerical and transactional roles where within-occupation variance has already widened sharply enough that the job description changes before the job title does.

The five-band structure — Critical, Very High, High, Moderate, Low — is a deliberate methodological choice, not a presentation one. Neighbouring jobs typically sit within 2-3 points of each other; the differences between a "ranked 31st" and "ranked 38th" occupation are mostly noise from small task-set differences, not signal worth acting on. Tiering at roughly 10-15 point intervals is where the statistically honest line sits. A precise occupation-level ranking would imply a precision the underlying scoring does not support — that has been the field's most consistent critique of the original Frey-Osborne 2013 rankings ever since Arntz, Gregory and Zierahn (2016).

Within tiers, the rosters below are ordered alphabetically. We do not publish per-occupation ranks. If you want the underlying numbers anyway, every occupation's full score, sector, and time window is in the occupations CSV — anyone is free to re-rank.

Critical Exposure (≥85) — 2 occupations

Data Entry Clerk, Telemarketer.

Very High Exposure (70–84) — 30 occupations

Accountant, Bank Teller, Bookkeeper, Call Centre Agent, Cashier, Claims Adjuster, Content Writer, Conveyancer, Copywriter, Credit Analyst, Customer Service Agent, Executive Assistant, Insurance Claims Adjuster, Insurance Underwriter, KYC Analyst, Legal Researcher, Legal Secretary, Media Buyer, Medical Secretary, Paralegal, Payroll Administrator, Pension Administrator, PPC Specialist, Reporting Analyst, SEO Specialist, Stockbroker, Translator, Travel Agent, Underwriting Analyst, Web Designer.

High Exposure (60–69) — 43 occupations

AML Analyst, Auditor, Business Intelligence Analyst, Business Process Analyst, Cost Accountant, Court Clerk, Credit Controller, Data Analyst, Data Quality Analyst, Database Administrator, Demand Planner, Email Marketing Manager, Facilities Coordinator, Financial Analyst, Graphic Designer, Group Accountant, Growth Analyst, Healthcare Administrator, Hedge Fund Analyst, Illustrator, Insight Analyst, Interpreter, Inventory Analyst, IT Support Analyst, Management Accountant, Market Research Analyst, Medical Writer, Mortgage Advisor, Mortgage Underwriter, Musician, Notary, Office Manager, Operations Analyst, Photographer, Procurement Specialist, Recruiter, Retail Assistant, Revenue Analyst, School Administrator, Social Media Manager, Supply Chain Analyst, Systems Analyst, Tax Advisor.

Moderate Exposure (40–59) — 148 occupations

Account Manager, Actor, Actuary, Advertising Strategist, Affiliate Manager, Analytics Engineer, Analytics Manager, Animator, Architect, Audit Manager, Backend Developer, Biostatistician, Brand Designer, Building Manager, Building Surveyor, Business Analyst, Business Development Manager, Catastrophe Modeller, Chemist, Claims Manager, Climate Scientist, Clinical Trials Manager, Cloud Architect, Commercial Property Manager, Commercial Property Surveyor, Compensation & Benefits Manager, Compliance Analyst, Compliance Officer, Content Strategist, Contracts Manager, Conversion Rate Optimiser, Corporate Lawyer, Corporate Tax Specialist, Credit Risk Manager, Customer Success Manager, Customs Broker, Data Engineer, Data Governance Manager, Data Scientist, Decision Scientist, DevOps Engineer, Digital Marketing Manager, Economist, Editor, Education Consultant, Employment Lawyer, Environmental Scientist, ERP Consultant, ESG Analyst, Estate Agent, Event Planner, Facilities Manager, Financial Advisor, Financial Controller, Financial Trader, Fleet Manager, Forensic Accountant, Forensic Scientist, Frontend Developer, Full-Stack Developer, Fund Accountant, Fund Manager, General Insurance Broker, Geneticist, Geologist, Hotel Manager, HR Business Partner, HR Manager, Immigration Lawyer, Import-Export Manager, In-House Counsel, Industrial Engineer, Influencer Marketing Manager, Insolvency Practitioner, Insurance Broker, Interior Designer, Internal Auditor, Investment Analyst, Investment Banker, IT Consultant, Journalist, Lab Manager, Learning & Development Manager, Legal Compliance Officer, Legal Operations Manager, Letting Agent, Librarian, Life Insurance Adviser, Logistics Manager, Loss Adjuster, Management Consultant, Manufacturing Engineer, Marine Insurance Underwriter, Marketing Manager, Materials Scientist, Mobile Developer, Network Engineer, Neuroscientist, NHS Project Manager, Operations Manager, Pension Actuary, People Analytics Manager, Performance Marketing Manager, Petroleum Engineer, Pharmacist, Podcast Producer, Policy Analyst, Portfolio Manager, Privacy Lawyer, Private Equity Analyst, Product Designer, Production Planner, Project Manager, Property Lawyer, Property Manager, Public Relations Manager, Purchasing Manager, QA Engineer, Quantitative Analyst, Quantity Surveyor, Radiographer, Radiologist, Real Estate Agent, Regulatory Affairs Specialist, Reinsurance Analyst, Retail Manager, Risk Analyst, Sales Representative, Scrum Master, Social Media Strategist, Solicitor, Statistician, Strategy Consultant, Supply Chain Manager, Talent Acquisition Specialist, Technical Program Manager, Training Coordinator, Transfer Pricing Specialist, Treasury Analyst, Treasury Manager, Tutor, Underwriting Manager, Urban Planner, UX Designer, UX Researcher, Video Editor, Warehouse Manager, Workforce Planner.

Low Exposure (<40) — 111 occupations

Account Director, Aerospace Engineer, AI Product Manager, Airline Pilot, Application Architect, Art Director, Barber, Barrister, Biomedical Engineer, Branch Manager, Brand Strategist, Broadcast Journalist, Business Unit Director, Care Worker, Carpenter, Change Management Consultant, Chef, Chemical Engineer, Chief Data Officer, Chief Executive Officer, Chief Financial Officer, Chief Marketing Officer, Chief Operating Officer, Chief People Officer, Chief Technology Officer, Civil Engineer, Civil Servant, Clinical Psychologist, Cloud Engineer, Communications Director, Community Manager, Creative Director, Creative Strategist, Criminal Defence Lawyer, Cybersecurity Analyst, Cybersecurity Engineer, Data Architect, Dental Hygienist, Dentist, Developer Advocate, Dietitian, Diversity & Inclusion Manager, Doctor, Drug Regulatory Affairs Manager, Early Years Educator, Ecologist, Electrical Engineer, Electrician, Employee Relations Manager, Energy Engineer, Enterprise Architect, Environmental Engineer, Epidemiologist, Family Lawyer, Finance Director, Financial Planner, Firefighter, Game Designer, General Manager, Head Teacher, IP Lawyer, IT Director, IT Manager, Judge, Landscape Architect, Legal Consultant, Machine Learning Engineer, Managing Director, Marine Biologist, Mechanical Engineer, Non-Executive Director, Nurse, Occupational Health Advisor, Occupational Therapist, Optometrist, Organisational Psychologist, Pharmaceutical Scientist, Physicist, Physiotherapist, Platform Engineer, Plumber, Police Officer, Primary School Teacher, Private Banker, Product Manager, Programme Director, Property Developer, Psychiatrist, Psychologist, Relationship Manager, Research Scientist, Restaurant Manager, Risk Manager, Secondary School Teacher, Security Architect, Service Designer, Site Reliability Engineer, Social Worker, Software Developer, Solutions

Architect, Solutions Engineer, Speech & Language Therapist, Strategy Director, Structural Engineer, Surgeon, Teaching Assistant, Transformation Director, University Lecturer, Utility Engineer, Veterinarian, Wealth Manager.

One caveat to read these tiers with. The tier an occupation sits in tells you the weighted average exposure of its tasks. It does *not* tell you how that exposure is distributed inside the job. A Bookkeeper (Very High) and a Stockbroker (Very High) sit in the same band on aggregate, but their internal distributions are quite different. The variance section above, and the per-occupation pages linked from each tier, are the places to read distribution rather than aggregate. Within tiers, the data does not support precise ranking — full per-occupation scores remain available in the downloadable CSV for anyone who needs them.

Sector-level breakdown

Aggregating occupation scores by sector is the level at which most of the AI-and-jobs press coverage operates — "finance is at risk", "healthcare is safe", "the trades are immune". The sector view in this Index supports that shape — and then complicates it. First, the sector means hide enormous within-sector variance. Second, the sector-level number is the wrong unit of analysis for an individual planning their career — task-level is where the answer actually lives. Both points come straight from the audit.

The sector-level ranking confirms the headlines — and then immediately undercuts them. **Finance & Accounting** has the highest mean exposure score in the dataset at **60** (precise: 59.63) across 19 occupations, which is what you get when most of the work is structured, document-bound, and rule-driven — bookkeeping, payroll, credit analysis, reconciliation, revenue recognition. It is the sector where the variance and exposure findings overlap most cleanly with the flagship-tasks ranking above. **Sales & Customer** follows closely (mean **59**, n=11), pulled up by telemarketing, call-centre work and customer-service roles where scripted handling makes up a large share of the day. At the other end, **Skilled Trades** has the lowest mean among multi-occupation sectors at **20** across plumbing, carpentry and electrical work — tasks that combine on-site perception, manual dexterity, and live diagnostic judgement, which is exactly the cluster the OECD's 2023 Employment Outlook identifies as where AI has progressed least.

That much is consensus. What sector means understate is internal variance. Four sectors — **Sales & Customer (70.36 pp), Financial Services (70.32), Finance & Accounting (70.05) and Legal (69.77)** — all sit at roughly **70 percentage points** of mean within-occupation task variance, the widest band in the dataset (among sectors with n≥5). The ranking among them is a virtual tie inside the methodology's ±5-10 pp scoring tolerance. Compare two occupations from Sales & Customer: a Customer Service Agent scores 74 (Very High) on aggregate, with "Routine enquiry handling" at 92% and "Empathic & vulnerable customer support" at 12% — an 80 pp internal spread. An Account Manager sits in the same sector at a much lower aggregate of 45 — Moderate band, with the role centred on relationship work that current AI handles poorly. The "sales jobs are at risk" headline treats the two as one thing. The data treats them as opposites.

The clearest illustration of why sector-level analysis is too coarse comes from **Healthcare**. The sector mean is **37** (precise: 37.47; n=30), which puts it in the lower half of the table — "low exposure", in casual reading. But the sector range goes from Surgeon at 11 (the most resilient occupation in the entire dataset) to Medical Secretary at 77 (Very High band, with "Medical transcription & clinical note creation" at 92%). Two occupations in the same sector, sitting 66 points apart on aggregate. If a hospital administrator reads the sector mean and concludes "we are fine", they will miss that an entire layer of clerical and transcription work inside their building has tasks already in the AI bullseye. If a junior doctor reads the same number and worries, they will miss that surgical, diagnostic and bedside care tasks score in single digits. The sector is the wrong unit.

> *Sector-level analysis is too coarse to be useful for individuals — task-level is where the answer lives.*

The same pattern holds across the other large sectors. **Legal** (n=26, mean **49** [precise: 49.19]) ranges from Judge at 16 to Legal Secretary at 83 — a 67-point sector range, with paralegal, conveyancing, and contracts work clustering high while advocacy and bench work cluster low. **Technology** (n=49, mean **45** [precise: 44.51]) splits between exposed reporting/QA/IT-support roles in the 70s and architect/senior roles in the 20s. **Education** (n=12, mean **42**) ranges from Translator at 79 down to Early Years Educators at 12. The 67-point spread inside Legal is wider than the spread between any two sector means.

Sector summary table, ordered by median exposure score:

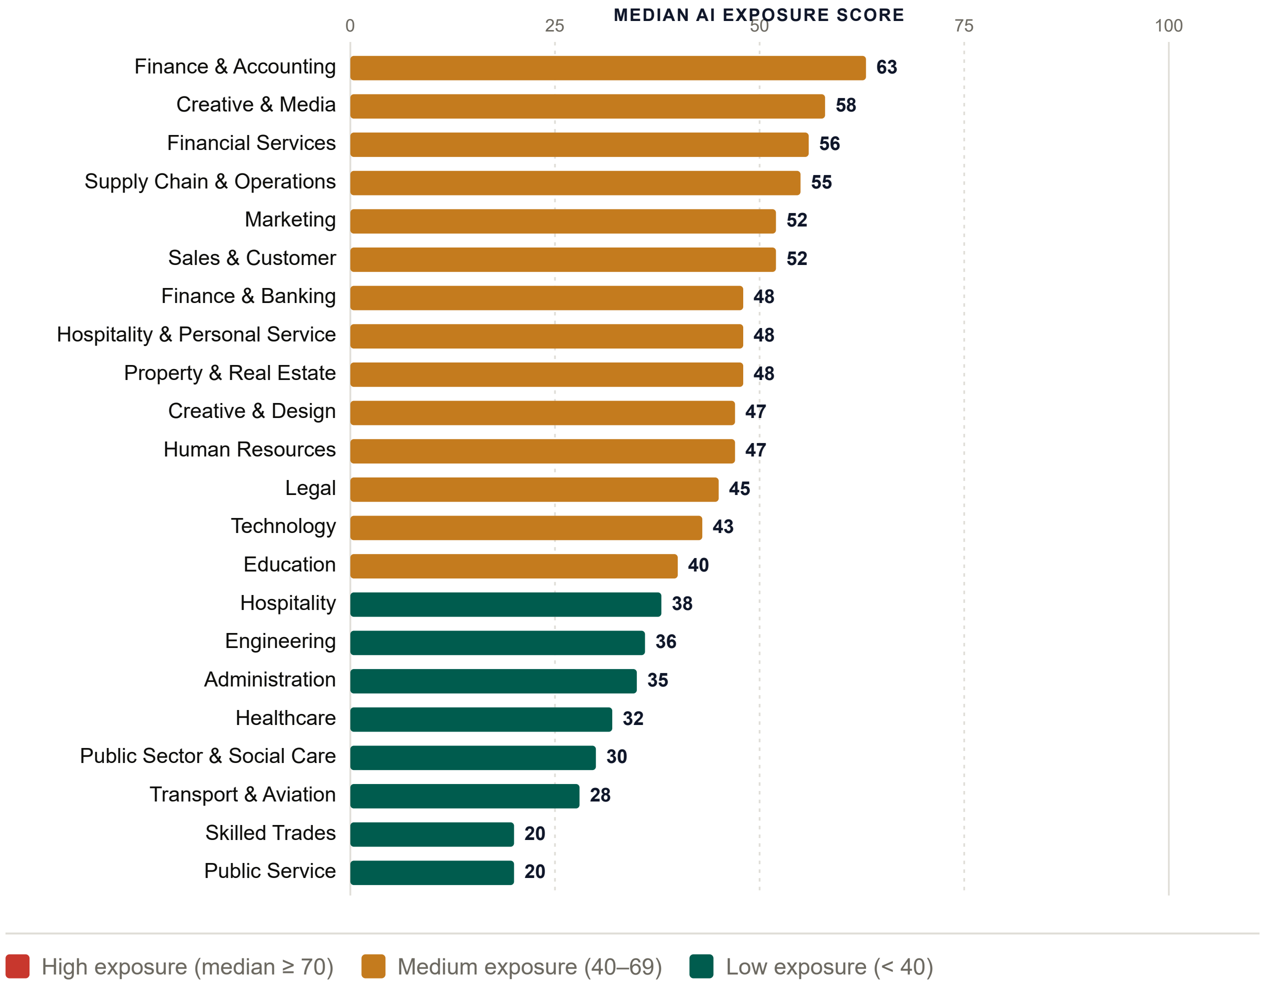

n = 22 sectors, sorted by median score

SECTOR	OCCUPATIONS	MEDIAN SCORE	HIGH-EXPOSURE SHARE
Finance & Accounting	19	63	16%
Creative & Media	2	58	0%
Financial Services	25	56	24%
Supply Chain & Operations	19	55	0%
Marketing	25	52	16%
Sales & Customer	11	52	36%
Finance & Banking	18	48	11%
Hospitality & Personal Service	2	48	50%
Property & Real Estate	12	48	0%
Creative & Design	20	47	10%
Human Resources	15	47	7%
Legal	26	45	15%
Technology	49	43	2%
Education	12	40	8%
Hospitality	3	38	0%
Engineering	20	36	0%
Administration	15	35	13%
Healthcare	30	32	3%
Public Sector & Social Care	6	30	0%
Transport & Aviation	1	28	0%
Skilled Trades	3	20	0%
Public Service	1	20	0%

The bottom line on the sector view: it is useful for rough orientation and for press-release-grade comparisons between industries, but it is not the layer at which an individual reader's question gets answered. The right next move from this section is to read the per-occupation page for whichever role you actually do, and look at the task-level scores there. The aggregate is a starting point; the residue of human-only tasks inside your job is the planning problem.

Time windows: when, not just what

Most "AI by 2030" headlines collapse a range of multi-year uncertainty into a single point on a calendar. The World Economic Forum's Future of Jobs Report 2025 forecasts "22% of jobs disrupted by 2030"; consultancies routinely publish single-date forecasts of varying ambition. **Single-date forecasts make for tidy headlines and useless planning.** We report time windows in **months, with a low-end and high-end**

range, because the underlying uncertainty is genuinely a range — and because individuals planning the next stretch of their working life need to know whether a task is being automated this calendar year or, more likely, somewhere between now and the early 2030s.

The shape of the distribution across the 334 audited occupations: **median window 18 to 36 months, mean 20 to 38**. About **70% of occupations have a high-end window between 24 and 72 months** — that is the broad central tendency, and it is consistent with the rate of adoption-not-just-capability change visible between the Anthropic Economic Index's February 2025 and March 2026 releases, which moved the "share of jobs with at least 25% of tasks performed by Claude" figure from roughly 36% to 49% in thirteen months. Rapid, but not year-zero rapid for most jobs.

The extremes at both ends are worth naming explicitly. **32 occupations have a low-end window of 4 months or fewer** — these are the fully transactional roles where current AI clearly already performs the task to a useful standard and the only remaining lag is procurement and integration: data entry, scripted telemarketing, payroll processing, document transcription, basic legal secretarial work. **58 occupations have a high-end window of 60 months or more**, with a long tail extending to 180 months: this is where regulated, embodied, or high-trust work concentrates — surgeons, airline pilots, plumbers, electricians, firefighters, primary-school teachers, judges. The wide outer bound on those reflects the multi-decade horizon over which capability could plausibly arrive, not a confident prediction it will.

One important methodological caveat that needs reading before any of these windows is acted on: **time windows measure capability readiness, not labour-market displacement.** A 12-month window means current AI is on track to perform that occupation's exposed tasks to a useful standard within the next year. It does not mean the displacement of workers performing those tasks will occur in 12 months — adoption typically lags capability by years, and the gap is set by cost-of-deployment, regulatory friction (especially in healthcare, financial services and legal work), organisational change resistance, and customer preference. Klarna's 2024-25 customer-service automation and subsequent 2026 partial reversal is the cleanest documented example of capability arriving ahead of organisational readiness, then re-calibrating. Read the windows as "when AI gets good enough to do the task", not "when your job disappears". Adoption lags capability by years.

Methodology

The Index 2026 covers **334 occupations and 2,563 individually-scored tasks**. The unit of analysis is the task, not the job title. Source taxonomy is the O*NET database maintained by the US Department of Labor — the same task taxonomy used by the original Frey-Osborne 2013 paper, by Arntz, Gregory and Zierahn (2016), by the OECD Employment Outlook 2023, and by Anthropic's Economic Index Hugging Face dataset. Using the same backbone keeps this Index comparable with the existing literature rather than a parallel taxonomy nobody else maps to. O*NET is in the public domain.

How tasks are scored

Each task receives a **0-100 AI-replaceability score**. The score measures the degree to which current frontier AI (as of early 2026) can perform the task to a useful working standard against a typical version of the task, without unusual scaffolding. Scoring is calibrated against two reference signals: **published AI capability evaluations** (LLM benchmark performance, published productivity studies including Brynjolfsson, Li & Raymond (2023) on customer support, the Noy & Zhang (2023) writing-task RCT, and Choi et al. (2024) on legal analysis), and **observed Claude usage** by task category from Anthropic's Economic Index March 2026 release. A task that Claude is observably doing at scale already, and that capability evaluations show frontier models handling competently, scores high. A task that requires embodied perception, live multi-stakeholder judgement, or regulated accountability scores low.

Per-occupation aggregate score is a weighted task-replaceability average. Weights are time-share — how much of a typical worker's day each task accounts for — drawn from O*NET work-activity weights and adjusted where UK practice differs materially from the US baseline. Each occupation has between **six and eleven tasks (mean 7.7)**. The aggregate score is the figure quoted at occupation level throughout this report and on the per-occupation pages.

Authorship and inter-rater reliability

The Index 2026 was scored by a single author (Robiul Islam) calibrated against named public benchmarks: Frey-Osborne (2013), OECD WP 189 (2016), Brynjolfsson, Li & Raymond (2023), Noy & Zhang (2023), Choi et al. (2024), the Anthropic Economic Index, and the ILO 2025 Refined Global Index. No second-rater inter-rater reliability number is published for this edition. Two careful analysts looking at the same task can land 5-10 points apart; assume ±5-10 percentage points of legitimate variation at the task level. A second-rater pass on a stratified subsample is committed for the 2027 edition. The audit script, source data, and per-occupation reasoning are public so any reader can re-score and check.

Why we publish in tiered bands

Occupation-level scores are calibrated in approximate bands (Low / Moderate / High / Very High / Critical). Within a band, the data does not support precise ranking. Task-level scores have finer differentiation (83 distinct values across 2,563 tasks) and are the primary unit of analysis in this Index.

This is the methodology section's most important constraint to flag up front. When occupations are re-scored against the same task list and the same calibration anchors, neighbouring occupations typically move within ±3-5 points in either direction. Publishing a numbered ranking — "Bookkeeper #6, Claims Adjuster #7" — would imply a precision the underlying signal does not carry. The Frey-Osborne ranking has been criticised on exactly this ground for over a decade. We chose to band rather than rank from the start. The task layer, where the underlying scoring is fine-grained (83 distinct percentages across 2,563 tasks), is where the precision actually lives — and that is where we publish ranked output.

What this Index does not measure

Three things, explicitly. **Cost of deployment.** A task scoring 92% replaceable does not mean automating it is cheap; integration, change management, and retraining are the dominant cost line in most enterprise AI projects, and they do not show up in a capability score. **Regulatory friction.** Healthcare, financial services, and legal work all have statutory or professional-body constraints on who can perform certain tasks, and those constraints lag capability by years. Over 1,000 FDA-cleared AI radiology tools exist as of late 2025, but the regulated workflow around them sets the actual adoption pace, not the underlying capability. **Organisational change resistance.** Klarna's reversal of its 2024-25 customer-service AI rollout after customer satisfaction declined is the cleanest published case study of organisational/customer push-back resetting an aggressive automation timeline. None of the three friction layers is in the score.

Other limitations

This Index uses the UK labour market as its referent — sector taxonomy follows ONS conventions, regulatory context references UK frameworks. Occupation coverage is currently 334 roles, biased toward white-collar information work; we under-cover some skilled-trade and physical-service occupations relative to their share of employment. The scores are point-in-time as of early 2026 and will shift as capability shifts; we plan to re-score and republish annually. Capability evaluations cited are public; calibration is judgement-based on those signals, not a closed-form model. Two careful analysts looking at the same task can land 5-10 points apart. That is part of why occupation-level output is banded, not ranked.

Reproducibility

Every task and score is in the tasks CSV; every occupation aggregate is in the occupations CSV. Both are CC BY 4.0. The audit script that produces the headline numbers quoted throughout this report (median, variance, bimodal share, sector rollups, tier counts) is published in the public repository as `audit_index.py` and is idempotent — anyone can re-run it against the published data and reproduce the figures.

About JobForesight

JobForesight scores AI exposure at the task level for 334 occupations. The full dataset is open. The methodology is documented. The audit script is in the public repo. Every number in this report is reproducible.

> **Author.** Robiul Islam, founder · LinkedIn
>
> **Publisher.** Solid Computing Ltd, registered in England & Wales (Company No. 07795981).
>
> **Press & general contact.** Use the contact form on jobforesight.com — flag "press" in the subject and I'll prioritise.
>
> **Update cadence.** We plan to update this Index annually.

How to cite this report

If you use the Index in research, journalism, or product work, please cite as:

Islam, R. (2026). *AI Career Risk Index 2026*. JobForesight. https://jobforesight.com/ai-career-risk-index-2026/

- Full report (PDF)
- Occupation-level scores (CSV)
- Task-level scores (CSV)