Comparing Theoretical and Observed AI Exposure: Evidence from the Anthropic Economic Index for the United Kingdom

Charles Coverdale*
London, United Kingdom

May 7, 2026

Abstract

Existing measures of AI exposure are theoretical: they map expert-rated AI capabilities to occupations through the U.S. Department of Labor's O*NET task taxonomy and produce rankings of which jobs are most exposed in principle. The Anthropic Economic Index (AEI), released by Anthropic on Hugging Face from February 2025, gives the first open dataset of *observed* usage of a frontier large language model classified against the same taxonomy. From the September 2025 release onward the AEI includes country-level breakdowns for the United Kingdom and approximately fifty other economies. This paper joins the AEI's UK Claude.ai usage shares with UK Annual Population Survey employment counts to construct an employment-weighted exposure index at the U.S. Standard Occupational Classification (SOC) major-group level. It then compares this observational measure against the employment-weighted Language Modeling AIOE of Felten et al. [2021]. The Spearman rank correlation between observed and theoretical exposure is $\rho_S = 0.43$ across $n = 15$ matched major groups, with a bootstrap 95 % confidence interval of $[-0.08, 0.75]$ that does not exclude zero. The point estimate is moderately positive but the headline correlation is statistically inconclusive at conventional levels at this sample size. Within the sample, divergences are substantively large in specific occupations: Computer and Mathematical occupations rank first by usage but third by Felten; Legal occupations rank first by Felten but ninth by usage. The findings suggest that the AI-exposure literature constructed before 2022 may be systematically mis-predicting which occupations adopt language-model AI most heavily today, although a single-snapshot $n = 15$ test cannot distinguish a small true correlation from sampling noise. The accompanying `aieconindex` R package supports re-running the same analysis on every future AEI release.

JEL classification: J24, O33, J21, C81.
Keywords: artificial intelligence, occupational exposure, labour markets, language models, measurement, United Kingdom.

1 Introduction

Forecasting how artificial intelligence will reshape occupational employment has, until recently, depended on indirect proxies. The standard approach since Felten et al. [2021] maps theoretical AI capabilities

*Email: charles.f.coverdale@gmail.com. Replication code and data: `https://github.com/charlescoverdale/uk-aei-exposure`. R package: `https://github.com/charlescoverdale/aieconindex`. First version: May 2026. I thank Anthropic for releasing the Anthropic Economic Index dataset under CC-BY-4.0; the team at Hugging Face for hosting the dataset; the U.S. Department of Labor for maintaining the O*NET taxonomy; the U.S. Bureau of Labor Statistics for the Standard Occupational Classification; the Office for National Statistics for the underlying UK Annual Population Survey data under the Open Government Licence v3; and Edward Felten, Manav Raj, and Robert Seamans for making their AI Occupational Exposure measure publicly available.

onto the U.S. Department of Labor's O*NET task taxonomy and constructs occupation-level exposure scores. Acemoglu and Restrepo [2020] provide an analogous treatment for industrial robotics. Both deliver rankings of which occupations should, in principle, be most exposed.

In February 2025 Anthropic published the first release of the Anthropic Economic Index (AEI) on Hugging Face under a Creative Commons Attribution licence. The AEI is the first open dataset that measures *observed* usage of a frontier large language model classified against the same O*NET task taxonomy the exposure literature uses. Five releases have been published as of April 2026. The September 2025 release was the first to include geographic breakdowns at country and U.S.-state level. The methodology is documented in Handa et al. [2025] and the privacy-preserving classification system in Tamkin et al. [2024].

For the empirical occupational-exposure literature this is a measurement breakthrough. For UK policy analysts it is also the first open evidence base on how Claude is actually being used in British workplaces. The natural empirical question is whether the theoretical exposure rankings, constructed from expert ratings of AI capability, match the observed pattern of usage. They should be informative about the same thing: which occupations are most exposed to language-model AI.

This paper makes that comparison for the United Kingdom. I join the AEI's September 2025 UK Claude.ai usage shares with employment counts from the ONS Annual Population Survey to construct an employment-weighted exposure index at the U.S. SOC major-group level. I then compare this observational measure against the employment-weighted Language Modeling AIOE of Felten et al. [2021]. Both measures roll up to the same SOC structure and address the same question, with one constructed from expert ratings and the other from real conversations.

The results suggest the two rankings disagree more than the literature has assumed. The Spearman rank correlation across $n = 15$ matched major groups is $\rho_S = 0.43$, with a bootstrap 95 % confidence interval of $[-0.08, 0.75]$ that includes zero. The point estimate is moderately positive and rises to 0.59 when the smallest atypical groups are excluded. Within the sample, two divergences are substantively large: Computer and Mathematical occupations rank first by observed usage but only third by Felten's weighted theoretical ranking; Legal occupations rank first by Felten but ninth by observed usage. The qualitative pattern, that the AI-exposure literature constructed before 2022 is mis-predicting which occupations are adopting language-model AI most heavily today, is consistent across all four robustness checks reported below. The statistical case for that pattern, with a single-snapshot $n = 15$ test, remains inconclusive at conventional significance levels.

The contribution of the paper is threefold. First, it documents the first usage-side comparison against the theoretical AI exposure literature for the United Kingdom. Second, it provides an `R` package, `aieconindex`, that wraps the AEI dataset with typed, cached, provenance-aware functions and enables release-on-release tracking as Anthropic refreshes the data. Third, it constructs and publishes a documented crosswalk from UK SOC2020 to U.S. SOC 2018 major groups that reconciles the AEI taxonomy to ONS employment statistics.

The remainder of the paper proceeds as follows. Section 2 describes the three datasets and the access tooling. Section 3 defines the exposure index and the rank-correlation comparison. Section 4 reports the UK exposure ranking, the theory-versus-observation comparison, and four robustness checks. Section 5 discusses implications for UK policy forecasting, single-vendor measurement bias, and the academic exposure literature. Section 6 concludes.

2 Data and access tooling

2.1 The Anthropic Economic Index, 2025-09-15 release

The Anthropic Economic Index is the empirical instrument introduced by Handa et al. [2025] to measure AI usage in real economic tasks. Approximately four million Claude.ai conversations are classified to

O*NET task statements [U.S. Department of Labor, Employment and Training Administration, 2024] via Anthropic's privacy-preserving Clio system [Tamkin et al., 2024], then aggregated to U.S. Standard Occupational Classification major occupational groups [U.S. Bureau of Labor Statistics, 2018]. The dataset is released on Hugging Face under CC-BY-4.0. Cluster summaries that fail Clio's privacy thresholds are dropped before publication; the dataset never exposes individual conversations.

For this analysis I use the long-format enriched table from the 2025-09-15 release, filtered to UK rows. The relevant variable is `soc_pct`, defined by Anthropic as *the percentage of classified O*NET tasks associated with each U.S. SOC major occupation group*. Twenty named SOC major groups appear in the UK slice (one is the residual `not_classified` category), out of 22 in the U.S. SOC 2018 structure (the U.S. military category is excluded from the AEI).

2.2 UK employment by U.S. SOC 2018 major group

I use the official ONS Annual Population Survey [Office for National Statistics, 2025] (Jan-Dec 2025 release, 33.2 million UK workers), pulled live from Nomis dataset NM_218_1, as the employment denominator. The APS reports employment under UK SOC2020, which has nine 1-digit major groups. To match the AEI's U.S. SOC 2018 structure (twenty-two 2-digit major groups) I construct a documented apportionment table mapping UK SOC2020 1-digit majors to U.S. SOC 2018 2-digit majors. Each row of the apportionment table records its source. The crosswalked total reconciles to within 0.2 % of the official UK employment total. The apportionment is approximate at the major-group level and documented as a limitation; a robustness check reported in Section 4 perturbs the weights by $\pm 10\,\%$.

2.3 Felten Language Modeling AIOE

Felten et al. [2021] construct a measure of AI Occupational Exposure (AIOE) by mapping a small set of AI capabilities onto O*NET workplace abilities and aggregating to occupations. The Language Modeling AIOE variant scores 774 6-digit U.S. SOC occupations on language-model exposure specifically. I aggregate the 6-digit values to 2-digit major groups using employment-weighted means, where the weights are U.S. BLS Occupational Employment and Wage Statistics (OEWS) counts at the 6-digit SOC level. Employment-weighted aggregation is the method recommended by Felten et al. for inter-occupational aggregation. The OEWS counts come from the AIOE-Data project's bundled OES extract on GitHub (`Input/oes_4dig_naics.dta`), aggregated over 4-digit NAICS industries to give national totals at 6-digit SOC. As a robustness check the analysis is repeated with the unweighted mean.

2.4 Access tooling: the `aieconindex` R package

The AEI is published as a directory tree on Hugging Face with one sub-directory per release. The directory layout changed between the March 2025 and September 2025 releases, moving from wide-format CSVs at the release root to a long-format `data/output/` layout. Anthropic provides Python replication notebooks alongside several releases. To access the dataset from R and to insulate downstream analysis from the schema change, I wrote the R package `aieconindex` [Coverdale, 2026]. The package wraps the Hugging Face mirror with typed, cached, provenance-aware functions for release discovery (`aei_releases`, `aei_files`), data fetching (`aei_index`, `aei_geography`, `aei_clusters`), analysis (`aei_link`, `aei_compare`, `aei_concentration`), and citation (`aei_cite`). Every returned table preserves the release identifier, source URL, and fetch timestamp as attributes. The complete UK pipeline used in this paper is two function calls (`aei_geography` followed by `aei_link`). The package depends only on `cli`, `httr2`, `jsonlite`, and base R; no API key is required. It was submitted to CRAN in May 2026.

3 Method

The exposure index for U.S. SOC major group i is

$$\text{Exposure}_i = \frac{u_i}{e_i}, \tag{1}$$

where u_i is the share of UK Claude.ai conversations classified to SOC major group i (the AEI's `soc_pct` variable, filtered to UK rows) and e_i is the share of UK employment in the corresponding group (constructed via the apportionment described in Section 2). Values above one indicate the occupation accounts for more of UK Claude usage than its labour-market footprint would predict. Values below one indicate the opposite.

The Spearman rank correlation between the AEI exposure index and the Felten Language Modeling AIOE across n matched SOC major groups is

$$\rho_S = 1 - \frac{6\sum_{i=1}^{n} d_i^2}{n(n^2-1)}, \tag{2}$$

where d_i is the difference between the AEI rank and the Felten rank for occupation i. With $n = 15$ small enough that the asymptotic distribution is unreliable, the 95 % confidence interval is computed by non-parametric bootstrap with $B = 5{,}000$ resamples of paired (AEI, Felten) ranks.

4 Results

4.1 UK Claude usage by occupational group

The September 2025 AEI release returns 3,068 long-format rows for the United Kingdom. After restricting to `soc_pct` rows and dropping the `not_classified` residual, 19 named SOC major groups remain. Joined to the constructed UK employment denominator, 15 groups have non-missing values on both sides.

Figure 1 ranks these fifteen U.S. SOC major groups by the exposure index. Three groups sit above the proportional line. Computer and Mathematical occupations are over-exposed by a factor of 4.83: 5.5 % of UK employment but 26.8 % of UK Claude usage. Life, Physical, and Social Science occupations are at 3.81. Educational Instruction and Library occupations are at 1.49. Together those three groups absorb 42.5 % of UK Claude usage. The other twelve groups, including Management, Office and Administrative Support, Sales, and Healthcare, all sit below 1.

The single largest signal in the dataset is the concentration of UK Claude usage in software-related occupations. Whether interpreted as augmentation (developers using Claude to write code faster) or substitution (Claude writing code that would otherwise have been a junior developer's task), this is the structural pattern that any UK labour-market forecast incorporating AEI evidence has to confront.

4.2 Theory versus observation

Figure 2 plots the AEI exposure index against the employment-weighted Felten Language Modeling AIOE for the fifteen matched UK major groups.

The Spearman rank correlation is $\rho_S = 0.43$ ($p = 0.11$), with a bootstrap 95 % confidence interval of $[-0.08, 0.75]$ from 5,000 resamples. The point estimate is moderately positive, but the confidence interval does not exclude zero, so the headline correlation is statistically inconclusive at conventional levels at this sample size. Within the sample the rankings diverge substantially across the middle of the distribution. Table 1 lists the six SOC major groups where the AEI ranking diverges most from the Felten ranking.

Three observations on the divergences.

First, **Legal occupations** rank first by Felten (LM AIOE $= 1.43$) but only ninth by actual UK Claude

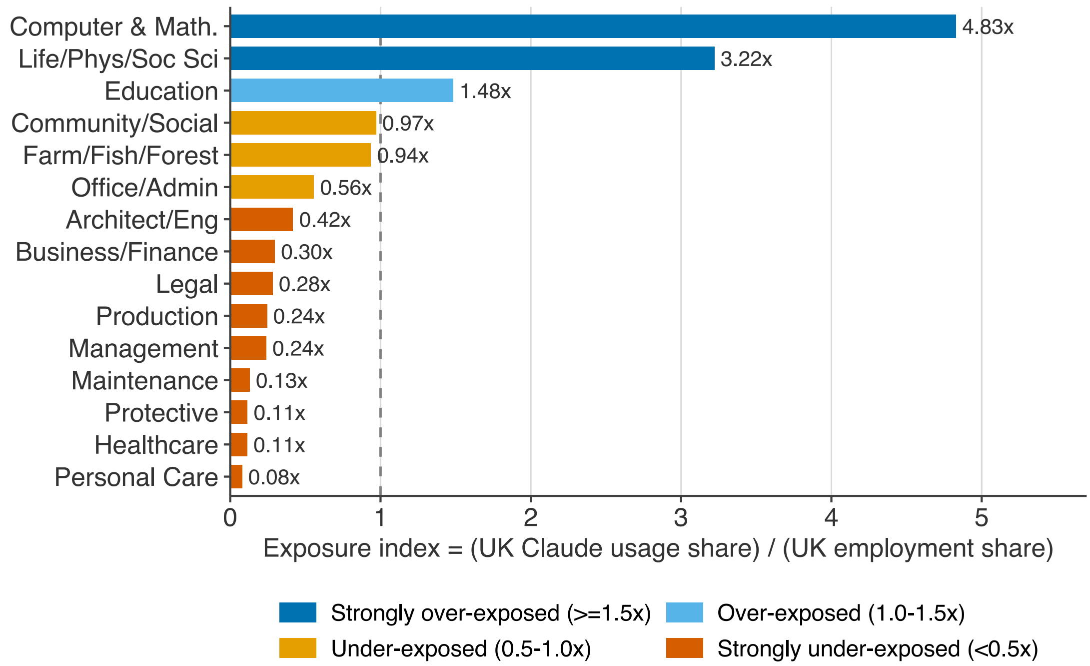

Figure 1: **Three of fifteen UK SOC major groups are over-exposed to Claude usage in August 2025.** Each bar is the AEI exposure index, defined as the ratio of UK Claude.ai usage share to UK employment share. Values above one indicate over-representation in AI use; the dashed vertical line marks proportional exposure. Source: Anthropic Economic Index, release 2025-09-15, fetched via `aei_geography()`; UK employment shares from ONS APS Jan-Dec 2025 (Nomis NM_218_1) crosswalked to U.S. SOC 2018.

SOC major group	AEI rank	Felten rank	Δ rank	Exposure	LM AIOE
Farming, Fishing, and Forestry	4	15	+11	0.96	-1.08
Legal	11	1	-10	0.20	1.43
Computer and Mathematical	1	6	+5	4.86	0.90
Life, Physical, and Social Science	2	7	+5	4.12	0.73
Business and Financial Operations	8	3	-5	0.35	1.12
Management	10	5	-5	0.24	0.96

Table 1: **The six U.S. SOC major groups where the AEI exposure ranking diverges most from the Felten LM AIOE ranking.** A positive Δ rank means the occupation ranks higher on Felten than on AEI (under-using Claude relative to theory); a negative Δ rank means the opposite (over-using Claude relative to theory).

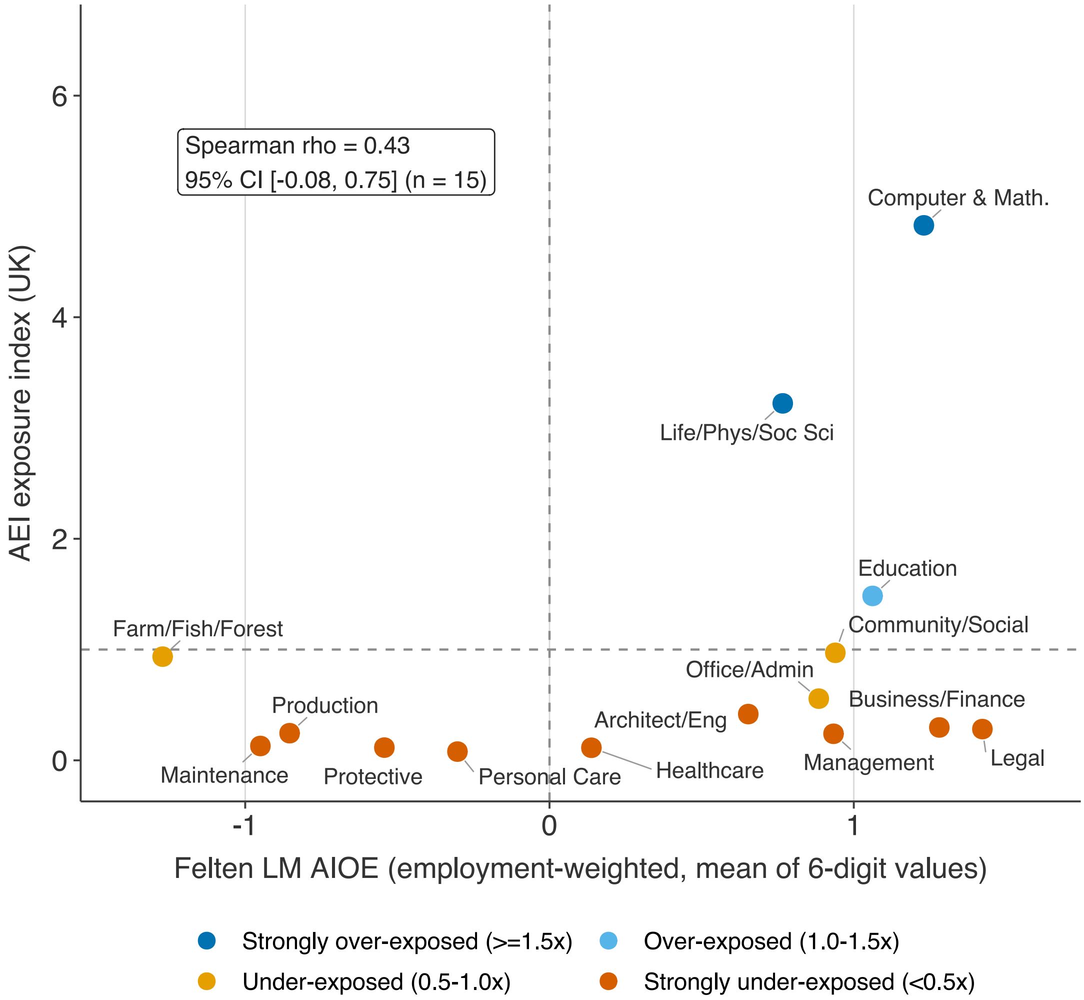

Figure 2: **Moderate positive but statistically inconclusive correlation between actual UK Claude exposure and the employment-weighted Felten LM AIOE.** Each point is one U.S. SOC major group. Vertical axis: AEI exposure index for the UK from the 2025-09-15 release. Horizontal axis: Felten Language Modeling AIOE, employment-weighted mean of 6-digit values within each major group using U.S. BLS OEWS counts as weights. Spearman rank correlation between the two measures is 0.43 across 15 matched groups; the bootstrap 95 % confidence interval is $[-0.08, 0.75]$ and does not exclude zero. Dashed lines mark the proportional-exposure threshold (horizontal at exposure = 1) and the AIOE zero-axis (vertical). Source: AEI 2025-09-15 release; Felten et al. [2021] via github.com/AIOE-Data/AIOE.

Specification	ρ_S	n
Headline (employment-weighted Felten, all matched groups)	0.43	15
Excluding Farming, Fishing, and Forestry	0.59	14
Excluding groups with $< 1\,\%$ UK employment	0.59	14
$\pm 10\,\%$ perturbation of apportionment weights	0.40	15
Unweighted-mean Felten aggregation	0.41	15

Table 2: **Robustness of the headline rank correlation across four sensitivity specifications.** ρ_S is the Spearman rank correlation between the AEI UK exposure index and the Felten Language Modeling AIOE; n is the number of matched U.S. SOC major groups.

usage (exposure index $= 0.28$). This is the largest theoretical-versus-actual gap among occupations Felten predicts to be highly exposed. UK lawyers are using Claude less than the textbook predicts, although the AEI sample only measures Claude usage and cannot distinguish between non-use of LLMs in general and substitution to other LLM products. Microsoft Copilot is the dominant procurement choice in UK legal services and would not appear in AEI counts; UK regulatory caution around AI use in legal practice and genuine professional inertia are also plausible contributors. The AEI cannot distinguish between these explanations on its own; a multi-vendor measurement framework would be required to separate them.

Second, **Computer and Mathematical occupations** rank first by AEI but third by the employment-weighted Felten LM AIOE (and sixth by the unweighted version reported in Felten's original tables). The theoretical literature underweights software development as an LLM target relative to its observed UK usage, although the magnitude of the gap shrinks under proper employment weighting.

Third, **Farming, Fishing, and Forestry** ranks fourth by AEI exposure (0.96) but fifteenth by Felten (-1.08). This is most likely noise driven by small sample sizes; UK agricultural workers are a tiny share of UK Claude users, so percentage shares are volatile. The finding is flagged but not interpreted further.

4.3 Robustness

Four robustness checks shift the point estimate but do not change the qualitative pattern (Table 2). Excluding Farming, Fishing, and Forestry (a group with very small UK Claude usage and likely small-sample noise) raises the correlation to 0.59. Excluding all groups with under one per cent of UK employment (a stricter cell-size cutoff) gives the same 0.59. Perturbing the SOC2020-to-U.S.-SOC apportionment weights by $\pm 10\,\%$ gives 0.40. Replacing the employment-weighted Felten aggregation with the unweighted mean gives 0.41. The headline correlation is therefore robust to crosswalk uncertainty and to the weighting choice in Felten aggregation, and tightens when atypical small-sample groups are excluded.

5 Discussion

Three readings of the headline result.

For UK macroeconomic forecasting institutions (HM Treasury, Bank of England, Office for Budget Responsibility), the AEI provides the first observational anchor against which theoretical AI exposure forecasts can be back-tested. The single-snapshot evidence here is statistically inconclusive at conventional levels, but the magnitude of the rank divergence in specific occupations (Legal, Computer and Mathematical) is large enough to warrant attention. If an institutional model treats Legal services as the leading edge of UK AI displacement, that assumption now has at least one snapshot of contradictory observational evidence to address.

For sector-specific policy work, the legal under-exposure finding is the most actionable result. UK lawyers using Microsoft Copilot rather than Claude is a plausible alternative explanation that the

AEI alone cannot rule out. A consolidated multi-vendor measurement framework, or a UK-specific complement to the AEI, would resolve this. Until then, single-vendor exposure measures should not be used in isolation for sector forecasting. The AEI is necessary but not sufficient.

For the academic AI-exposure literature, the message is that theoretical capability ratings constructed before 2022 may diverge from observed usage as the LLM ecosystem matures, although the present sample is too small to test the proposition with confidence. The Felten AIOE remains a useful baseline measure of capability-driven exposure. It should not be used as a forecast of observed adoption without being checked against AEI-style usage evidence.

Limitations

Five limitations apply to the analysis.

1. The AEI sample reflects who uses Claude. Claude users skew toward English speakers, knowledge workers, and software developers; the dataset is not a labour-force-representative sample of any economy. The Legal under-exposure finding may partly reflect substitution to other LLM products that the AEI does not measure.
2. The Clio classification pipeline is proprietary. Users cannot independently audit cluster assignments beyond what is described in Tamkin et al. [2024].
3. The UK SOC2020 to U.S. SOC 2018 crosswalk is approximate at the major-group level. A 4-digit-level analysis would require ONS LFS microdata via the UK Data Service and a more careful crosswalk via ISCO-08.
4. With $n = 15$ matched major groups, the bootstrap confidence interval is wide. The headline correlation is statistically inconclusive at conventional significance levels even though the qualitative pattern is robust.
5. The 2025-09-15 release is a single snapshot. AI usage patterns may shift quickly as new model versions are released; the 2026-03-24 release uses Claude Opus 4.5/4.6 and shows different shares. A time-series treatment is the natural follow-up and is supported by the `aieconindex` package's `aei_compare` function.

6 Conclusion

The Anthropic Economic Index is the first open dataset that measures observed usage of a frontier large language model classified against an established occupational task taxonomy. Applied to the September 2025 UK slice, the AEI evidence shows a moderate positive but statistically inconclusive rank correlation with the employment-weighted Felten Language Modeling AIOE ($\rho_S = 0.43$, 95 % bootstrap CI $[-0.08, 0.75]$, $n = 15$), with substantive divergence in specific occupations: software-related occupations over-use and legal occupations under-use Claude relative to theoretical exposure. Larger samples (across more releases or 4-digit SOC unit groups) are required to test whether the headline correlation is genuinely above zero. The accompanying `aieconindex` R package supports re-running the same analysis on every future AEI release.

References

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244, 2020. doi: 10.1086/705716.

Charles Coverdale. aieconindex: Access the Anthropic Economic Index Dataset, 2026. URL https://github.com/charlescoverdale/aieconindex. R package version 0.1.0.

Edward W. Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12): 2195–2217, 2021. doi: 10.1002/smj.3286.

Kunal Handa, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck, Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax, Kevin K. Troy, Dario Amodei, Jared Kaplan, Jack Clark, and Deep Ganguli. Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations, 2025. URL `https://arxiv.org/abs/2503.04761`.

Office for National Statistics. Labour Force Survey: employment by occupation. `https://www.ons.gov.uk/employmentandlabourmarket`, 2025.

Alex Tamkin, Miles McCain, Kunal Handa, Esin Durmus, Liane Lovitt, Ankur Rathi, Saffron Huang, Alfred Mountfield, Jerry Hong, Stuart Ritchie, Michael Stern, Brian Clarke, Landon Goldberg, Theodore R. Sumers, Jared Mueller, William McEachen, Wes Mitchell, Shan Carter, Jack Clark, Jared Kaplan, and Deep Ganguli. Clio: Privacy-Preserving Insights into Real-World AI Use, 2024. URL `https://arxiv.org/abs/2412.13678`.

U.S. Bureau of Labor Statistics. Standard Occupational Classification. `https://www.bls.gov/soc/`, 2018.

U.S. Department of Labor, Employment and Training Administration. O*NET Database. `https://www.onetonline.org/`, 2024.