

---

# The Jagged Global Economy: Frontier AI Unevenly Exposes National Economies

---

**Arul Murugan**  
UC Berkeley

**Tomás Aguirre**  
University of São Paulo

**Abhishek Nagaraj**  
UC Berkeley

**Rishi Bommasani**  
Stanford University

## Abstract

Frontier AI’s labor-market effects matter to workers, firms, and policymakers, but current evidence generally comes from a handful of high-income economies. The capabilities of frontier AI are jagged across work tasks and national economies diverge in how they allocate human labor. We introduce a national AI exposure metric that combines occupation-level exposure scores and international employment data for 141 countries. We find that high income countries are substantially more exposed than low income countries and that Europe & Central Asia are 50% more exposed than Sub-Saharan Africa. We also find a gender gap: women are more exposed than men in 91% of countries, driven by their concentration in white-collar and sales occupations. The exceptions are countries where women’s employment remains concentrated in agriculture and household enterprises. We validate our national AI exposure estimates by showing they predict national AI adoption statistics published by Anthropic, Microsoft, and OpenAI. Beyond direct exposure, we identify a new mechanism for indirect exposure due to cross-country income dependencies. Some nations such as Tajikistan depend heavily on foreign workers remitting money back to their home countries: Tajikistan’s direct exposure to frontier AI is below-average but because 37% of Tajikistan GDP is Russian remittance and Russia is very exposed, Tajikistan’s remittance-accounted exposure becomes above-average. Our research shows that national variation in exposure is large enough that policy responses calibrated to U.S. or European labor markets will not generalize.

## 1 Introduction

How frontier AI will reshape national economies is central to workers, firms, and governments. Around the world, the public already expects AI to reshape work: a majority of respondents in the 2025 Stanford AI Index report that AI will change how people do their jobs within five years, while the 2025 Ipsos AI Monitor finds that more people think AI will worsen than improve their local job market [Stanford HAI, 2025, Ipsos, 2025]. Policymakers frame frontier AI as a strategic priority [UK DSIT, 2025, Canada ISED, 2025, European Commission, 2025] and, reciprocally, frontier AI firms target national AI sovereignty agendas [OpenAI, 2025]. AI investment is seen as critical for future growth: the UN projects that the global AI market will expand to \$4.8 trillion by 2033 [UN, 2025].

These effects are likely to be very unevenly distributed across countries. Frontier AI capabilities are jagged across tasks [Dell’Acqua et al., 2026], and national economies differ substantially in how they allocate workers to jobs. Consider call center work. Brynjolfsson et al. [2025a] find that deploying generative AI in a large call center increases issues resolved per hour by 15% with larger gains for less experienced and lower-skilled workers. A productivity shock of this level at scale would disproportionately disrupt nations like India or the Philippines where a large share of the national labor force is tied to call center work. The same frontier AI technologies could complement workers in one nation, reorganize entry-level labor markets in another, and have little effect in a third.

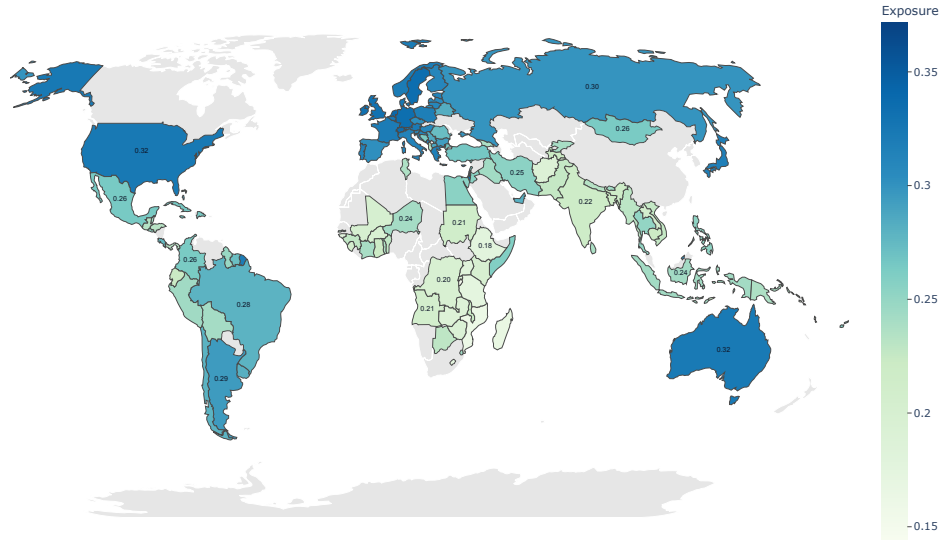


Figure 1: **National AI exposure varies substantially across countries.** Darker countries are more exposed to frontier AI. Countries without 2-digit ISCO-08 employment data are shown in gray.

We introduce a national AI exposure metric to capture this heterogeneity. Our approach combines occupational exposure estimates from the economics of frontier AI literature with internationally comparable statistics on how countries allocate human labor across occupations.<sup>1</sup> The metric summarizes how strongly a national economy’s current labor allocation aligns with the tasks that frontier AI can already accelerate or transform. The resulting object is not a forecast of adoption, wages, or net employment losses. Rather, it is a comparable measure of exposure at the national level: a way to quantify which countries are more exposed, through the composition of their labor markets, to the jagged capabilities of frontier AI. This exposure is not necessarily positive or negative: it signals opportunities for productivity gains but these gains are not certain, they may contribute to economic growth, and they may reduce the comparative advantage of human labor but could also increase it via certain patterns of augmentation.

We find substantial variation in national exposure (Figure 1): the most exposed nation (Luxembourg) is  $2.6\times$  more exposed than the least exposed nation (Burundi). The most exposed regions are North America and Europe & Central Asia, which are at least 50% more exposed than the least-exposed region of Sub-Saharan Africa. We study different explanations for this result and find that disparities in the prevalence of white collar work explain much of the observed variation.

Given the overall trends in national exposure, we disaggregate to study specific country and worker characteristics. At the country level, we find that higher income countries are more exposed. At the worker level, we find a pervasive gender gap: women in 91% of studied nations are more exposed than men because of their occupational composition. In many countries, women already face barriers to labor force entry and receive lower wages for comparable work product [International Labour Organization, 2024, Blau and Kahn, 2017]. The groups we find to be more exposed (i.e. richer countries, women) also poll less optimistically about AI work [McClain et al., 2025], so our results may explain why these groups are less optimistic.

We test the predictive validity of our national AI exposure metric by predicting national AI adoption according to usage statistics published by Anthropic, Microsoft, and OpenAI. We find strong monotonic relationships between our exposure estimates and all of these adoption statistics. We estimate

<sup>1</sup> Concurrent with our work, Gmyrek et al. [2026] introduce a national AI exposure metric with a similar methodology but do not publish per-country exposure estimates at the time of writing. Garg et al. [2026] introduce the Global Automation Atlas (<https://automationatlas.org>), a complementary task-country automation measure that allows exposure labels to vary across national contexts.

that 0.10 increase in national AI exposure corresponds to  $12\times$  growth in national per-capita Claude usage and 19 percentage-point increase in the national generative AI adoption rate.

Since the economic impacts of frontier AI are not only determined by domestic activity but also cross-country relationships, we study indirect exposure channels. We identify a novel mechanism for national exposure: foreign workers sending remittance back to their families in their home countries. We show that this channel substantially increases national exposure for several countries where foreign remittance is a sizable fraction of national GDP. For example, 25% of the GDP of Honduras, Guatemala, and El Salvador is remittance and more than 80% for all three countries is remitted from the United States. Since the national AI exposure of these three Central American countries is below average but the United States is highly exposed, accounting for this indirect effect substantially raises national AI exposure in Honduras, Guatemala, and El Salvador. National policymakers in these countries should track how the foreign labor economies they depend on are changing because of frontier AI adoption and the resulting second-order effects on their national income. Overall, our research identifies critical heterogeneity that should inform national AI and economic policy.

## 2 Data

To estimate national exposure to frontier AI, we combine occupational employment statistics from the International Labour Organization (ILO) with occupational exposure estimates from Gmyrek et al.

### 2.1 Occupational employment statistics

To study the impact of AI on labor economies around the world, we use country-level data compiled by the ILO. The ILO publishes employment statistics based on the International Standard Classification of Occupations (ISCO-08), which is a four-level hierarchy that spans 436 occupational categories. As an example, one of the 436 occupational categories is “Generalist Medical Practitioners” (code 2211; code 2 = Professionals, code 22 = Health Professionals, code 221 = Medical doctors).<sup>2</sup>

We use the most recent annual statistics for each country’s employment counts for the 43 2-digit (sub-major group) ISCO-08 occupational categories, along with available ILO wage statistics for wage-weighted extensions. The counts are produced by national statistical offices through country-level data collection (e.g. labor force surveys, census, administrative records) and then compiled and disseminated by the ILO. We describe data processing in Appendix A, including the country filters and the exclusion of armed-forces, aggregate, and residual occupational rows. We end up with employment data for 40 occupational categories across 141 countries. This provides strong coverage of the world’s nations, with most missing nations being very small countries with less mature government labor data infrastructure.<sup>3</sup>




Figure 2: Illustrative cross-country differences in labor composition.

Figure 2 visualizes how employment composition varies around the world. For example, office- and knowledge-intensive work (ISCO major groups 1–4) accounts for 61% of employment in the United States, but only 14% in India and 5% in Mozambique. And agriculture (ISCO 6) accounts for 73.0% of employment in Mozambique and 35% in India but only 0.4% in the United States. Beyond labor

<sup>2</sup>For brevity, we will sometimes refer to occupational categories as “occupations” where the distinction is immaterial.

<sup>3</sup>We lack 2-digit ISCO-08 data for a few major economies (e.g. China, Canada, Saudi Arabia).

compositional differences, wages also vary greatly: the average monthly white-collar wage is PPP \$7,778 in the United States, compared with PPP \$1,740 in India.<sup>4</sup> White-collar work earns 74.2% of the wage bill in the United States to 30.4% in India.

## 2.2 Occupational exposure estimates

AI capabilities can impact labor through multiple mechanisms. A critical mechanism is *productivity*. AI could reduce the time spent performing tasks. *Exposure* measures the time spent by a worker performing a task (or occupation) that *could* be saved by current technological capabilities. Exposure estimates bridge technological progress with economic task-based models, a standard framework in the economics of technological change that studies how technologies substitute for, complement, and create tasks rather than affecting occupations as undifferentiated bundles [Autor et al., 2003, Acemoglu and Autor, 2011, Acemoglu and Restrepo, 2019].

Eloundou et al. [2024] popularized the application of exposure to the economic impacts of frontier AI. To estimate occupational exposure, they average across an occupation’s tasks based on whether AI capabilities would save half the time spent on the task while preserving quality. To encode the level of capability at the time, they specified primitives that models could do (e.g. summarize medium-length documents) and could not do (e.g. repair physical equipment). Their empirical estimates use human and GPT-4 annotations to determine task exposure, leading to an overall estimate that “80% of the U.S. workforce could have at least 10% of their work tasks affected by the introduction of LLMs, while approximately 19% of workers may see at least 50% of their tasks impacted.” Prior to the work of Eloundou et al. [2024], earlier works estimated occupational exposure to machine learning and AI based on occupational abilities [Felten et al., 2018, 2021], patents [Webb, 2020], and by mapping cognitive abilities required for tasks to benchmarks [Tolan et al., 2021]. Following their work, others have expanded exposure estimates to multimodal capabilities [Felten et al., 2023], complementarity rather than pure substitution [Pizzinelli et al., 2023], and to integrate evidence from observed usage [Massenkoff and McCrory, 2026].

While exposure pervades empirical research on frontier AI, critiques target (i) the aggregation from tasks to occupations or (ii) the divergence between theorized and observed exposure. The standard aggregation of averaging task-level exposure to yield occupation-level exposure may incorrectly imply a linear production structure: occupations with the same average exposure may face different impacts when exposed tasks appear in chains or when tasks are quality complements and bottlenecks [Demirer et al., 2026, Gans and Goldfarb, 2026]. Further, the impact on wages or employment may be better predicted not by the tasks that are automated, but by the expertise requirements of the residual tasks [Autor and Thompson, 2025]. Eloundou et al. [2024] emphasize that exposure quantifies technical feasibility rather than observed productivity: Narayanan and Kapoor [2025] posit a variety of diffusion bottlenecks and Massenkoff and McCrory [2026] quantify large gaps between theoretical exposure and observed usage in several occupations.

Most AI exposure estimates use the U.S. occupational classification system. Our focus is global and not limited to the U.S., so we use exposure estimates for the international ISCO occupational classification system prepared by researchers at the ILO [Gmyrek et al., 2025a]. Since the ISCO occupations do not have rich task-level decompositions, they first use Polish occupational data that maps occupations to 29,753 distinct tasks and gather human judgments about the automation potential of a representative sample of 2,861 tasks from 1,640 workers. Experts review these worker assessments and they use the resulting expert-adjusted annotations to train a model to predict task-level exposure, which is then aggregated for all 4-digit ISCO-08 occupations.

## 3 National AI Exposure Estimates

We quantify national exposure to frontier AI based on nation-agnostic occupational exposure estimates and nation-specific occupational employment statistics.<sup>5</sup> Prior work applies similar approaches using different estimates for occupational AI exposure [Cazzaniga et al., 2024, Lewandowski et al., 2025,

---

<sup>4</sup>PPP denotes purchasing-power-parity-adjusted U.S. dollars to account for local price levels across countries.

<sup>5</sup>This neglects heterogeneity (e.g. task composition) in the same occupation across countries. However, intra-occupation heterogeneity may be dominated by intra-category cross-occupational heterogeneity at the 2-digit ISCO level. Gmyrek et al. [2026] attempt to address this via a more complex nation-sensitive estimation procedure for occupation exposure.

Gmyrek et al., 2026]. We define national AI exposure as the employment-share-weighted average exposure of a country’s occupational structure.

For an occupation  $j$ , let  $e_j$  quantify its exposure and  $c_{i,j}$  count its employment in nation  $i$ .

$$n_i = \frac{\sum_j c_{ij} e_j}{\sum_j c_{ij}} \quad (1)$$

National AI exposure varies substantially across countries (Figure 1): the most-exposed country (Luxembourg,  $n = 0.37$ ) is  $2.6\times$  more exposed than the least-exposed country (Burundi;  $n = 0.14$ ). Luxembourg’s unusually large finance-and-insurance sector drives its high exposure: business and administration professionals and associate professionals together account for 38.7% of total exposure, with legal, social and cultural professionals contributing another 10.0%. This pattern is consistent with Luxembourg’s unusually large finance-and-insurance sector.<sup>6</sup> Singapore also demonstrates high exposure through a similar but more technical/managerial mechanism: national exposure is 0.37, the white-collar share is 73.8%, and business, administration, and managerial occupations together account for 45.1% of total exposure [Ministry of Manpower, Singapore, 2026]. At the bottom, Burundi shows how strongly the index is pulled down by agrarian employment structure. It has the sample’s lowest exposure (0.1440), the lowest white-collar share (3.1%), and a strongly negative white-collar residual (-0.0468). Nearly 69.3% of Burundi’s total exposure comes from subsistence farmers, fishers, hunters and gatherers alone [World Bank, 2024]. §B.1 shows that the results are robust to the choice of occupational exposure estimates.

On top of the large variation in national exposure, we find a persistent *gender gap*: women are more exposed than men in 126 of the 138 countries we study. The median gender gap is 13.7% of national exposure and is driven by occupational composition.<sup>7</sup> Cross-country differences in male and female white-collar shares explain about half of the variation in the gender gap (50.3%), and including sales raises that explanatory power to 87.8%, indicating that women’s greater concentration in white-collar and sales work is the dominant structural mechanism. The pattern is visible in concrete cases. Women are substantially more exposed in the Philippines (0.2964 for women versus 0.2214 for men) and Jamaica (0.3173 versus 0.2276), where official statistics show women concentrated in service-and-sales, clerical, and professional work [Philippine Statistics Authority, 2024, Statistical Institute of Jamaica, 2020]. By contrast, women are less exposed in Pakistan (0.1970 for women versus 0.2338 for men) and India (0.2001 versus 0.2240), where women’s employment remains more concentrated in agriculture and household-enterprise work [Pakistan Bureau of Statistics, 2026, Ministry of Statistics and Programme Implementation, 2023]. This result that women may be more prone to AI-mediated labor disruption warrants concern given preexisting challenges for women to find gainful employment. In many countries, women experience barriers to labor force participation [International Labour Organization, 2024] and receive lower wages than men, including within occupations and after accounting for standard observable characteristics [Blau and Kahn, 2017]. And on the topic of AI, women are less likely to adopt the technology for work [Humlum and Vestergaard, 2025, Otis et al., 2024] and have less positive impressions of the technology [McClain et al., 2025]. Consequently, women workers may be more disrupted by AI because of men adopting AI in work.

### 3.1 Explaining national AI exposure

Cross-national variation in labor composition yields substantial differences in national AI exposure. Can these differences be explained by more basic properties of the countries or their workforces? Figure 3 shows national exposure as a function of the white collar share of the country’s labor force. We find that the white collar share is highly predictive of exposure ( $R^2 = 0.91$ ). We also find that gross income is predictive of exposure, but less so than white collar share: log gross national income (log GNI) explains 77% of cross-country variation in national AI exposure. Together, the two predictors have an  $R^2$  of 0.93, indicating that the workforce-specific white collar share statistic is more explanatory than the broader national-level income statistic. We further explore how the exact classification of occupations as white collar or blue collar affects this result: changing the label of a single occupational category by classifying sales occupations (ISCO-08 code 52) as white collar instead of blue collar substantially increases explanatory power ( $R^2 = 0.96$ ). In §B.4, we test

<sup>6</sup>Luxembourg had the European Union’s highest employment share in finance and insurance [Eurostat, 2024].

<sup>7</sup>Manning and Aguirre [2026] observe this pattern for the U.S. labor economy.

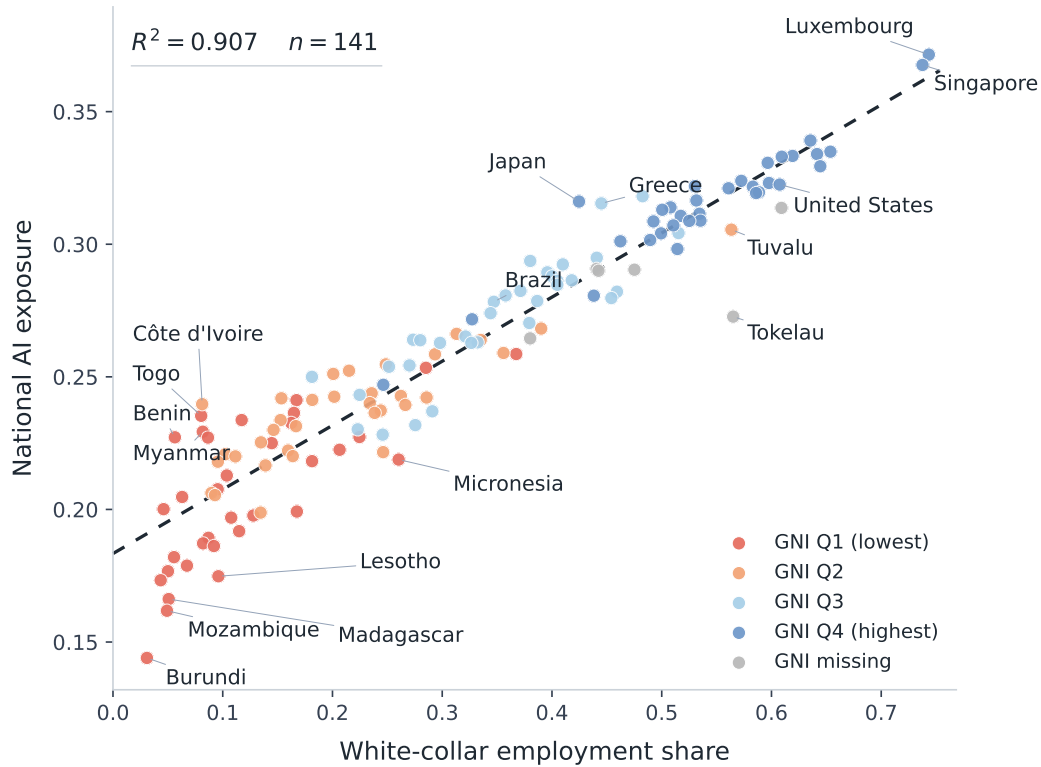


Figure 3: **White-collar employment share predicts national AI exposure.**

other predictors, generally finding that measures related to computer use, written comprehension, administrative work, and information processing are good predictors.

The correlation between national AI exposure, white collar work, and national income manifests via other channels. Cross-country surveys find substantial heterogeneity in public opinion on AI on matters spanning trust in the technology, confidence in its benefits, concerns about its risks, and desire for regulation. For example, respondents in advanced economies generally report less trust and perceive less benefit from frontier AI relative to respondents from developing economies [Pew Research Center, 2025, Gillespie et al., 2025, Ipsos, 2025]. Anthropic interviewed 81k Claude users, finding that workers in occupations with more observed exposure are more likely to express concern about job displacement [Massenkoff and Huang, 2026]. Beyond heterogeneity, many of these studies do report an overarching concern that the public fears and expects AI-driven job loss [Ipsos, 2024], which reflects a common interpretation of exposure as displacement risk even if exposure need not imply worse labor-market outcomes [Imas and Shukla, 2026].

#### 4 Predicting National AI Adoption

While exposure estimates translate from technological capabilities to economic quantities, they can be prone to misunderstanding [Eloundou et al., 2024, Manning and Aguirre, 2026, Imas and Shukla, 2026]. Exposure identifies where AI could plausibly increase productivity, not which jobs will necessarily disappear. For past technologies, economics research substantiates that task- and occupation-level exposure estimates predict downstream labor outcomes [Webb, 2020, Autor and Thompson, 2025]. For frontier AI, emerging evidence also shows that exposure estimates predicted observed frontier AI adoption [Tomlinson et al., 2025, Chatterji et al., 2025, Handa et al., 2025] and AI-mediated employment reduction [Brynjolfsson et al., 2025b, Tucker, 2026].

We test whether national AI exposure estimates predict observed national outcomes. As Chandar [2025] notes, current data is insufficient to study broad cross-national wage or employment changes

attributable to AI. Therefore, we focus our attention on predicting national AI adoption using statistics published by frontier AI companies based on chatbot usage data. We use data from Anthropic’s Anthropic Economic Index for Claude usage [Massenkoff et al., 2026], Microsoft’s AI Diffusion adoption measures [Microsoft AI Economy Institute, 2026a,b], and OpenAI’s ChatGPT usage statistics [Chatterji et al., 2025, OpenAI, 2026] with details in §A.5.

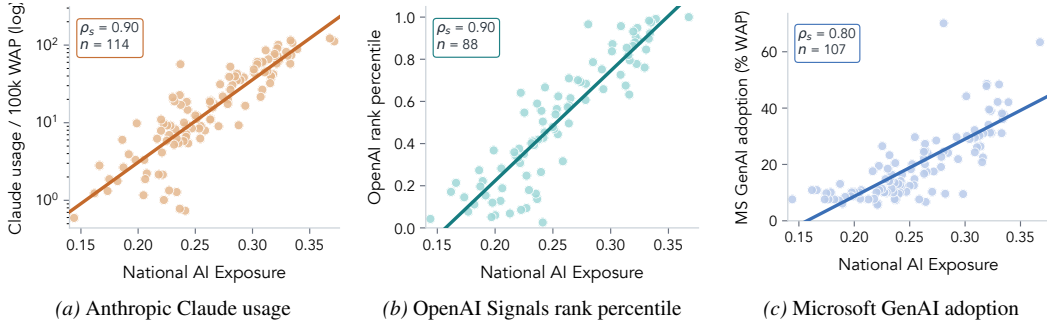


Figure 4: **National AI exposure predicts national AI adoption.** (a) Anthropic Claude usage per 100,000 working-age people, using the Anthropic Economic Index Claude.ai country release for February 5–12, 2026, and plotted on a log scale. (b) OpenAI Signals country-rank percentile for calendar-year 2025, based on OpenAI’s population-normalized ranking of countries by per-capita share of sampled consumer ChatGPT messages; OpenAI publishes ranks rather than country-level usage rates. (c) Microsoft AI Diffusion adoption rate for Q1 2026, measured as the percent of the working-age population using any of the covered generative-AI tools.

We find that our national AI exposure estimates predict all three observed AI usage measures, although the three companies publish different types of statistics. Using the Anthropic data, we find that national Claude usage scales approximately exponentially in national AI exposure: in the fitted relationship, a 0.10 increase in national AI exposure yields a  $12\times$  increase in national Claude usage per 100,000 working-age people. Since OpenAI ranks countries rather than publishing per-country statistics, we study ranking correlation within the 88 countries shared across our data and OpenAI Signals. The top quintile by exposure (average exposure percentile = 90%) has high ChatGPT usage with an average OpenAI percentile of 86.5%. Conversely, the bottom quintile by exposure (average exposure percentile = 10%) has low ChatGPT usage with an average OpenAI percentile of 14.9%. Using the Microsoft data, we find that national generative AI adoption rate<sup>8</sup> scales linearly in national AI exposure: in the fitted relationship, a 0.10 increase in national AI exposure yields an additional 203 national adopters per 1,000 working-age people.

Given our previous results (§3.1) that national AI exposure is itself predictable from national white-collar share and gross national income (GNI), we test whether national AI exposure provides additional explanatory power over these standard national metrics in predicting national AI adoption. Exposure alone strongly predicts all three adoption measures, with  $R^2 = 0.77$  for Anthropic Claude usage,  $R^2 = 0.81$  for OpenAI Signals, and  $R^2 = 0.61$  for Microsoft AI Diffusion. However, once white-collar share and log GNI are included, adding national AI exposure contributes little additional explanatory power: the full-model  $R^2$  is unchanged for Anthropic and Microsoft after rounding, and increases by only 0.001 for OpenAI (see Table 6).

## 5 Indirect Channels for National AI Exposure

While the direct impact of AI on domestic workers is central to determining how national labor economies change because of AI, it is not the sole determinant. National economic outcomes depend on the activities of other nations (e.g. trade, tariffs, outsourcing). Given concerns that AI-driven automation may reduce labor demand and domestic wages [Acemoglu, 2024, Autor and Thompson, 2025, Brynjolfsson et al., 2025b], we consider how this may be mediated by the impacts of AI on foreign countries.

<sup>8</sup>Adoption counts visits to generative AI sites/apps: Alice, ChatGPT, Character.ai, Claude, CIOVA X, DeepSeek, ERNIE Bot (Yiyan.baidu), GigaChat, Google Gemini, Grok, Khanmigo.ai, Meta.ai, Microsoft Copilot, Midjourney, Mistral.ai, NanoSemantics AI Assistant, Perplexity, Tongyi Qianwen, and Xiaowei.

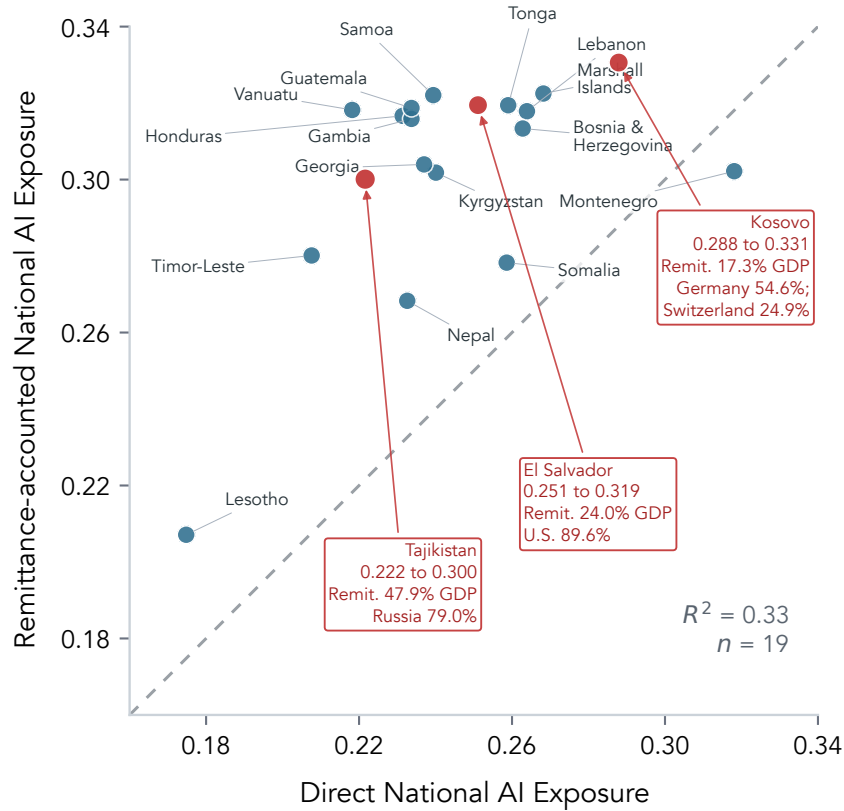


Figure 5: **Remittance-accounted vs. direct national AI exposure.** Countries shown have remittance of at least 10% of national GDP.

We observe that while the concern is often expressed as AI suppressing wages, the underlying concern is that individuals will not be able to generate the income required to sustain their lives. In most countries, this can be restated in terms of wages, because individual income and individual wages are essentially equal. However, we observe that in some nations, a second mechanism meaningfully contributes to income: *remittance*. Remittance is the transfer of money by a foreign worker (usually a migrant) to their home country to serve as income for their families. We hypothesize that remittance introduces a second channel for national exposure: if income in country A is sourced via remittance from country B, then national exposure in country B indirectly exposes country A.

To study this hypothesis, we use the KNOMAD database from the World Bank on global migration and development (§A.4; World Bank, 2026a,b). The KNOMAD country-country remittance matrix estimates annual remittance flows between source and target countries by partitioning total target-side remittance inflow proportional to bilateral migration distributions [Ratha et al., 2022]. For example, since Mexico received \$52 billion from the U.S. in remittance during 2021, a decline in U.S. labor demand could indirectly reduce income in Mexico via remittance reduction.

We find that the remittance channel generally raises measured exposure among highly remittance-dependent countries: 18 of the 19 countries with remittances above 10% of GDP move above the diagonal in Figure 5. Tajikistan is the clearest example. It received \$6.80 billion in personal remittances in 2024, equal to 47.9% of GDP. The latest KNOMAD bilateral matrix is for 2021; in that matrix, Russia accounts for \$2.31 billion of Tajikistan’s \$2.92 billion in measured bilateral remittance inflows, or 79.0%. Tajikistan’s direct national AI exposure is low ( $n = 0.222$ , 21st percentile among countries with measured exposure), but its remittance inflows come primarily from more exposed economies, led by Russia ( $n = 0.298$ , 75th percentile) and Germany ( $n = 0.333$ , 96th percentile). As a result, Tajikistan’s remittance-accounted exposure rises to  $n = 0.300$ , comparable to a 75th-percentile country in the direct exposure distribution.

This result illustrates that even if countries are not directly exposed to frontier AI, or that theoretical exposure is unlikely to convert into observed productivity (e.g. due to the digital divide and lack of internet access), they should prepare for the ripple effects of AI adoption elsewhere. In particular, while policymakers in these countries currently account for remittance in strategic planning given its scale, they may not account for the second-order remittance-mediated effect of AI on their economy.

## 6 Discussion

Frontier AI development is concentrated in a few nations like the United States and China, but its economic impacts are globally diffuse. The national AI exposure metric we introduce provides an initial lens for clarifying how frontier AI capabilities may impact countries around the world in different ways and the specific occupational categories within those labor markets. This encodes two forms of heterogeneity: occupations are differentially exposed and national labor markets are differentially composed of occupations. To make the most of these metrics, we believe they should be integrated with measures that operate at other levels of abstraction in reasoning about the impacts of (frontier) AI on the economy. In turn, this introduces a third form of heterogeneity: data availability varies substantially across the world.

**Heterogeneity in data availability.** As we also consider in this work, usage data is especially useful in the study of frontier AI’s economic impacts because it directly reflects the actual tasks and workflows where AI is adopted [Manning, 2025]. From the international perspective, usage data is attractive because it centralizes data pertinent to many different countries at a single entity, namely the AI provider. Notably, several frontier AI companies signed onto the New Delhi Frontier AI Commitments at the 2026 India AI Impact Summit to publish usage data reports.<sup>9</sup> For similar reasons, other private datasets such as payroll or payments information may help characterize frontier AI adoption while being naturally consolidated [e.g. Brynjolfsson et al., 2025b]. On the other hand, many forms of traditional labor economics data hinge on government data collection infrastructure. Therefore, data quality and quantity may vary greatly across the world, generally leading to greater availability for advanced economies. As our results show, frontier AI’s impacts on advanced economies are unlikely to resemble impacts on developing nations. Therefore, making the most of other data sources and sensibly imputing data in cases of missingness will be essential for characterizing the impacts of frontier AI on developing economies.

**Limitations.** The primary limitations of our work relate to limits of the underlying data we rely upon for occupational exposure estimates and national employment statistics. First, we treat an occupation as identically exposed for every country. While most work neglects within-country variation in an occupation (e.g. software engineers in San Francisco may differ from those in New York City), we more strongly neglect cross-country intra-occupation variation (e.g. software engineers in San Francisco may differ from those in New Delhi). A concurrent task-country approach relaxes this assumption by evaluating task exposure within national contexts [Garg et al., 2026]; our measure is complementary because it focuses on employment-weighted frontier-AI exposure using harmonized ISCO occupational employment data. Second, we operate at the ISCO 2-digit level which collapses many distinct occupations into a single category. For example, group 26 combines lawyers (2611) with musicians, singers, and composers (2652), even though their task bundles are qualitatively different. Third, the underlying occupational classification system (ISCO-08) was developed in 2008 and may not capture occupational emergence in the past two decades. Looking forward, this extends to occupations that emerge due to frontier AI (e.g. prompt engineers). Fourth, we do not have complete coverage of the world’s nations and we inherit the ways in which the ILO and World Bank handle national data for contentious regions of the world. Perhaps most critically, we do not have national AI exposure estimates for China, Canada, and Saudi Arabia as a result.<sup>10</sup>

---

<sup>9</sup>The commitment to *Advance Analysis on Real-World AI Usage* states that “Participating organizations will work to enhance analysis regarding global AI adoption for economic purposes ... Publish — by the next AI Summit — statistical insights derived from anonymised, aggregated and taxonomized usage data, either directly or (where relevant) through contributions to international efforts.”

<sup>10</sup>We encourage future work to explore data imputation for these countries from domestic data surveys if data via the ILO is unavailable, though we refrained from doing this for analytic simplicity.

## Acknowledgements

We thank Alex Spangher, Alexander Wan, Bharat Chandar, Diyi Yang, Dylan Clement, Erik Brynjolfsson, Sam Manning, Sarah Bana, Satya Krishna, Stephane Hatgis-Kessell, and Tom Cunningham for helpful discussion.

## References

- Stanford HAI. Artificial intelligence index report 2025, 2025. URL [https://hai.stanford.edu/assets/files/hai\\_ai\\_index\\_report\\_2025.pdf](https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf). Accessed 2026-04-14.
- Ipsos. Ipsos ai monitor 2025, 2025. URL <https://www.ipsos.com/sites/default/files/ct/publication/documents/2025-06/Ipsos-AI-Monitor-2025.pdf>. Accessed 2026-04-14.
- UK DSIT. Ai opportunities action plan. Government of the United Kingdom, January 2025. URL <https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan>. Accessed 2026-04-14.
- Canada ISED. Canadian sovereign ai compute strategy, October 2025. URL <https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy>. Accessed 2026-04-14.
- European Commission. Ai continent action plan, April 2025. URL <https://digital-strategy.ec.europa.eu/en/factpages/ai-continent-action-plan>. Accessed 2026-04-14.
- OpenAI. Introducing openai for countries, May 2025. URL <https://openai.com/global-affairs/openai-for-countries/>. Accessed 2026-04-14.
- UN. Technology and innovation report 2025: Inclusive artificial intelligence for development, 2025. URL [https://unctad.org/system/files/official-document/tir2025\\_en.pdf](https://unctad.org/system/files/official-document/tir2025_en.pdf). Accessed 2026-04-14.
- Fabrizio Dell’Acqua, Edward McFowland III, Ethan Mollick, Hila Lifshitz, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2):403–423, 2026. doi: 10.1287/orsc.2025.21838. URL <https://pubsonline.informs.org/doi/10.1287/orsc.2025.21838>.
- Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942, 2025a. doi: 10.1093/qje/qjae044. URL <https://academic.oup.com/qje/article/140/2/889/7990658>.
- Paweł Gmyrek, Mariana Viollaz, and Hernan Winkler. Disruption without dividend?: How the digital divide and task differences split GenAI’s global impact. ILO Working Paper 166, International Labour Organization, Geneva, 2026. URL [https://www.ilo.org/sites/default/files/2026-03/wp166\\_web.pdf](https://www.ilo.org/sites/default/files/2026-03/wp166_web.pdf).
- Prashant Garg, Tommaso Crosta, and Jasmin Baier. Global Automation Atlas, May 2026. URL <https://automationatlas.org/paper/>.
- International Labour Organization. The impact of care responsibilities on women’s labour force participation. Statistical brief, International Labour Organization, October 2024. URL [https://www.ilo.org/sites/default/files/2024-10/GEDI-STAT%20brief\\_formatted\\_28.10.24\\_final.pdf](https://www.ilo.org/sites/default/files/2024-10/GEDI-STAT%20brief_formatted_28.10.24_final.pdf).
- Francine D. Blau and Lawrence M. Kahn. The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):789–865, 2017. doi: 10.1257/jel.20160995. URL <https://www.aeaweb.org/articles?id=10.1257/jel.20160995>.
- Colleen McClain, Brian Kennedy, Jeffrey Gottfried, Monica Anderson, and Giancarlo Pasquini. How the u.s. public and ai experts view artificial intelligence. Technical report, Pew Research Center, April 2025. URL <https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/>.

- David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333, 2003.
- Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In David Card and Orley Ashenfelter, editors, *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier, 2011. doi: 10.1016/S0169-7218(11)02410-5.
- Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019. doi: 10.1257/jep.33.2.3.
- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308, 2024. doi: 10.1126/science.adj0998. URL <https://www.science.org/doi/abs/10.1126/science.adj0998>.
- Edward W. Felten, Manav Raj, and Robert Seamans. A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108:54–57, 2018. doi: 10.1257/pandp.20181021.
- 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.
- Michael Webb. The impact of artificial intelligence on the labor market. Working paper, 2020.
- Songül Tolan, Annarosa Pesole, Fernando Martínez-Plumed, Enrique Fernández-Macías, José Hernández-Orallo, and Emilia Gómez. Measuring the occupational impact of ai: Tasks, cognitive abilities and ai benchmarks. *Journal of Artificial Intelligence Research*, 71:191–236, 2021. doi: 10.1613/jair.1.12647.
- Edward W. Felten, Manav Raj, and Robert Seamans. Occupational heterogeneity in exposure to generative ai. *SSRN Electronic Journal*, 2023. doi: 10.2139/ssrn.4414065.
- Carlo Pizzinelli, Augustus Panton, Marina M. Tavares, Mauro Cazzaniga, and Longji Li. Labor market exposure to ai: Cross-country differences and distributional implications. Technical Report WP/23/216, International Monetary Fund, 2023.
- Maxim Massenkoff and Peter McCrory. Labor market impacts of ai: A new measure and early evidence, 2026. URL <https://www.anthropic.com/research/labor-market-impacts>.
- Mert Demirer, John J. Horton, Nicole Immorlica, Brendan Lucier, and Peyman Shahidi. Chaining tasks, redefining work: A theory of ai automation. Technical Report 34859, National Bureau of Economic Research, 2026.
- Joshua S. Gans and Avi Goldfarb. O-ring automation. Technical Report 34639, National Bureau of Economic Research, 2026.
- David Autor and Neil Thompson. Expertise. Working Paper 33941, National Bureau of Economic Research, June 2025. URL <http://www.nber.org/papers/w33941>.
- Arvind Narayanan and Sayash Kapoor. Ai as normal technology: An alternative to the vision of ai as a potential superintelligence. Technical report, Knight First Amendment Institute, April 2025. URL <https://knightcolumbia.org/content/ai-as-normal-technology>.
- Paweł Gmyrek, Janine Berg, Karol Kamiński, Filip Konopczyński, Agnieszka Ładna, Balint Nafradi, Konrad Rosłaniec, and Marek Troszyński. Generative ai and jobs: A refined global index of occupational exposure. Technical Report 140, International Labour Office, Geneva, 2025a. URL <https://doi.org/10.54394/HETP0387>.
- Mauro Cazzaniga, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma J. Rockall, and Marina Mendes Tavares. GenAI: Artificial intelligence and the future of work. IMF Staff Discussion Note SDN/2024/001, International Monetary Fund, Washington, DC, January 2024. URL <https://www.imf.org/en/publications/staff-discussion-notes/issues/2024/01/14/gen-ai-artificial-intelligence-and-the-future-of-work-542379>.

- Piotr Lewandowski, Karol Madoń, and Albert Park. Workers' exposure to AI across development. IBS Working Paper 02/2025, Instytut Badan Strukturalnych, 2025. URL [https://ibs.org.pl/wp-content/uploads/2025/03/Workers\\_AI\\_exposure\\_across\\_development\\_IBS\\_WP\\_202502.pdf](https://ibs.org.pl/wp-content/uploads/2025/03/Workers_AI_exposure_across_development_IBS_WP_202502.pdf).
- Eurostat. 4.9 million employed in finance and insurance sector. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20240610-2>, June 2024. Accessed 2026-04-20.
- Ministry of Manpower, Singapore. Labour force in singapore 2025. Technical report, Ministry of Manpower, Singapore, February 2026. Accessed 2026-04-20.
- World Bank. Burundi: Supporting youth-at-risk through the creative sector (p505887) project information document (pid). Technical Report PIDDC00690, World Bank, November 2024. Accessed 2026-04-20.
- Sam J Manning and Tomás Aguirre. How adaptable are american workers to ai-induced job displacement? Working Paper 34705, National Bureau of Economic Research, January 2026. URL <http://www.nber.org/papers/w34705>.
- Philippine Statistics Authority. 2024 fact sheet on women and men in the philippines: Work and economic participation. Technical report, Philippine Statistics Authority, 2024. Accessed 2026-04-20.
- Statistical Institute of Jamaica. Jamaican labour market: Impact of covid-19. Technical report, Statistical Institute of Jamaica, July 2020. Accessed 2026-04-20.
- Pakistan Bureau of Statistics. Labour force survey 2024–25. Technical report, Pakistan Bureau of Statistics, 2026. Accessed 2026-04-20.
- Ministry of Statistics and Programme Implementation. Women and men in india 2023: Participation in economy. Technical report, Ministry of Statistics and Programme Implementation, Government of India, 2023. Accessed 2026-04-20.
- Anders Humlum and Emilie Vestergaard. The unequal adoption of chatgpt exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences of the United States of America*, 122(1):e2414972121, 2025. doi: 10.1073/pnas.2414972121. URL <https://doi.org/10.1073/pnas.2414972121>.
- Nicholas G. Otis, Solène Delecourt, Katelyn Cranney, and Rembrand Koning. Global evidence on gender gaps and generative ai. Working Paper 25-023, Harvard Business School, 2024. URL <https://www.hbs.edu/ris/download.aspx?name=25023.pdf>.
- Pew Research Center. How people around the world view ai. Technical report, Pew Research Center, October 2025. URL <https://www.pewresearch.org/global/2025/10/15/how-people-around-the-world-view-ai/>.
- Nicole Gillespie, Steve Lockey, T. Ward, A. Macdade, and G. Hassed. Trust, attitudes and use of artificial intelligence: A global study 2025. Technical report, The University of Melbourne and KPMG, 2025. URL [https://figshare.unimelb.edu.au/articles/report/Trust\\_attitudes\\_and\\_use\\_of\\_artificial\\_intelligence\\_A\\_global\\_study\\_2025/28822919](https://figshare.unimelb.edu.au/articles/report/Trust_attitudes_and_use_of_artificial_intelligence_A_global_study_2025/28822919). Accessed 2026-04-17.
- Maxim Massenkoff and Saffron Huang. What 81,000 people told us about the economics of ai, 2026. URL <https://www.anthropic.com/research/81k-economics>.
- Ipsos. Ipsos predictions survey 2025: Positivity about how this year has gone highest since before the pandemic, December 2024. URL <https://www.ipsos.com/en-us/ipsos-predictions-2025>. Reports that globally 65% expect AI-driven job losses and 43% expect AI-driven job creation. Accessed: 2026-04-09.
- Alex Imas and Soumitra Shukla. How will ai-driven automation actually affect jobs?, March 2026. URL <https://aleximas.substack.com/p/how-will-ai-driven-automation-actually>.

- Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. Working with ai: Measuring the applicability of generative ai to occupations. *arXiv preprint arXiv:2507.07935*, 2025. doi: 10.48550/arXiv.2507.07935. URL <https://arxiv.org/abs/2507.07935>.
- Aaron Chatterji, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. How people use chatgpt. Working Paper 34255, National Bureau of Economic Research, September 2025. URL <https://www.nber.org/papers/w34255>.
- 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. *arXiv preprint arXiv:2503.04761*, 2025. doi: 10.48550/arXiv.2503.04761. URL <https://arxiv.org/abs/2503.04761>.
- Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Technical report, Stanford Digital Economy Lab, November 2025b. URL [https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine\\_Nov25.pdf](https://digitaleconomy.stanford.edu/app/uploads/2025/11/CanariesintheCoalMine_Nov25.pdf).
- Lee C. Tucker. You're (not) hired: Artificial intelligence and early career hiring in the quarterly workforce indicators. Technical report, April 2026. URL [http://leetucker.net/docs/Youre\\_not\\_hired\\_Tucker\\_20260417.pdf](http://leetucker.net/docs/Youre_not_hired_Tucker_20260417.pdf).
- Bharat Chandar. Ai and labor markets: What we know and don't know, October 2025. URL <https://digitaleconomy.stanford.edu/news/ai-and-labor-markets-what-we-know-and-dont-know/>. Accessed: 2026-04-09.
- Maxim Massenkoff, Eva Lyubich, Peter McCrory, Ruth Appel, and Ryan Heller. Anthropic economic index report: Learning curves, 2026. URL <https://www.anthropic.com/research/economic-index-march-2026-report>.
- Microsoft AI Economy Institute. Global ai adoption in 2025: A widening digital divide. Technical report, Microsoft, 2026a. URL <https://www.microsoft.com/en-us/corporate-responsibility/topics/AI-Economy-Institute/reports/Global-AI-Adoption-2025/>.
- Microsoft AI Economy Institute. Ai diffusion report data. <https://github.com/microsoft/ai-diffusion-report>, 2026b. Q1 2026 update; MIT licensed data repository; accessed May 19, 2026.
- OpenAI. Openai signals: Measuring ai adoption, protecting privacy, and empowering decisions. Technical report, February 2026. URL <https://cdn.openai.com/signals/openai-signals-global-report.pdf>. Global affairs report.
- Daron Acemoglu. The simple macroeconomics of ai. Working Paper 32487, National Bureau of Economic Research, May 2024. URL <https://www.nber.org/papers/w32487>.
- World Bank. Bilateral remittance estimates using migrant stocks, host country incomes, and origin country incomes (US\$ million). [https://data360.worldbank.org/en/indicator/WB\\_KNOMAD\\_BRE](https://data360.worldbank.org/en/indicator/WB_KNOMAD_BRE), 2026a. KNOMAD/World Bank Data 360 indicator, accessed April 19, 2026.
- World Bank. Remittance inflows (US\$ million). [https://data360.worldbank.org/en/indicator/WB\\_KNOMAD\\_MRI](https://data360.worldbank.org/en/indicator/WB_KNOMAD_MRI), 2026b. KNOMAD/World Bank Data 360 indicator, accessed April 19, 2026.
- Dilip Ratha, Sonia Plaza, and Eung Ju Kim. Bilateral remittance matrix (new). <https://blogs.worldbank.org/en/peoplemove/bilateral-remittance-matrix-new>, December 2022. World Bank People Move blog, December 19, 2022.
- Sam Manning. Understanding ai's labor market impacts, 2025. URL <https://cdn.sanity.io/files/d81r1a4f/staging/6c5cbe4cbe50f97034173d7f4b55ae1f1377b655.pdf>.

Pawel Gmyrek, Janine Berg, Karol Kamiński, Filip Konopczyński, Agnieszka Ładna, Balint Nafradi, Konrad Rosłaniec, and Marek Troszyński. Generative ai and jobs: A refined global index of occupational exposure. Technical Report 140, International Labour Organization, 2025b.

Ruth Appel, Peter McCrory, Alex Tamkin, Michael Stern, Miles McCain, and Tyler Neylon. Anthropic economic index report: Uneven geographic and enterprise ai adoption, September 2025. URL <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>.

International Labour Organization. Ilostat database, 2026. URL <https://ilostat ilo.org/data/>. Accessed 2026-04-14.

O\*NET Resource Center. O\*net 30.2 data dictionary. <https://www.onetcenter.org/dictionary/30.2/text/>, April 2026. Site updated 2026-04-14; accessed 2026-04-20.

Seyed Mahdi Hosseini Maasoum and Guy Lichtinger. Generative ai, expertise, and effective labor supply. Technical report, SSRN, January 2026. URL <https://ssrn.com/abstract=6059674>.

## A Data

### A.1 ILO employment

Table 1: Country coverage in the ILOSTAT employment pipeline

Status	Countries
Included in baseline sample (medium/high reliability)	141
Observed but excluded for low reliability	6
Excluded: TOTAL/NEC only	2
Excluded: ISCO-88 only	9
Excluded: major-group only	9

Baseline sample counts reflect the current merged ILO refresh and include countries with medium or high reliability only.

We use employment data published by the International Labour Organization (ILO), taking the most recent annual country observation available in the local ILO mirror accessed in April 2026. The total-employment country-years in the underlying panel range from 2009 to 2025, with most observations from 2024 or 2025. The employment data used in the main analysis are reported at the 2-digit (sub-major group) level of ISCO-08, yielding up to 40 occupation groups per country once armed-forces, aggregate, and residual categories are excluded. The three excluded armed-forces sub-major groups are ISCO-08 codes 01, 02, and 03; we also drop aggregate TOTAL rows and residual X rows rather than redistributing them. In the ILO mirror EMP\_TEMP\_SEX\_OC2\_NB\_A, 113 countries report nonzero employment in the armed-forces groups. Among those countries, the excluded armed-forces share is usually small (median 0.42%) but reaches 5.11% in Iraq, 4.61% in Palestine, 4.46% in Lebanon, and 2.95% in Sudan. While the data nominally covers 167 nations, we analyze the 141 where the ILO ascribes greater confidence to the underlying data quality and data collecting procedures. Gender-disaggregated exposure estimates use the same ILO occupational-employment pipeline restricted to sex-disaggregated country-occupation cells. For the main white-collar measure, we classify ISCO-08 major groups 1–4 (managers, professionals, technicians and associate professionals, and clerical support workers) as white-collar work. Sales occupations are treated separately in the mechanism analysis because ISCO-08 code 52 is a boundary case: including sales with white-collar work substantially increases the explanatory power of the simple labor-composition split.

### A.2 ISCO exposure estimates

Table 2: Hierarchy of occupational exposure estimates built from Gmyrek et al. (2025)

Level	Count	Aggregation rule	Use in this paper
Task-by-occupation rows in the Gmyrek workbook	3,265	Exposure score assigned to each task row within a 4-digit ISCO-08 occupation	Raw exposure source layer
4-digit occupations (ISCO-08)	427	Mean task-level exposure score within each 4-digit occupation	Occupation-level exposure estimates
2-digit occupations	40	Mean 4-digit occupation score within each 2-digit ISCO-08 group	Matched to ILO employment shares in the main cross-country analysis
1-digit occupation groups	9	Aggregation of 2-digit occupations to ISCO major groups	Descriptive summaries and wage-weighted extensions

*Note:* Counts are based on the Gmyrek et al. (2025) exposure workbook used in our pipeline. The workbook contains 3,219 predicted rows and 46 reconciled rows.

We use the occupational exposure estimates for ISCO-08 [Gmyrek et al., 2025b]. We describe these estimates in terms of their relationship to the ISCO-08 multi-level hierarchy in Table 2. Since we only have employment data at the 2-digit level of ISCO-08, we average the occupational exposure estimates for all occupations within the same 2-digit category. Of the 436 ISCO-08 4-digit occupations, the Gmyrek et al. data score 427. The nine unscored 4-digit codes are six “not elsewhere classified” residual codes (1439, 3139, 3435, 5249, 7319, 8189) and three armed-forces codes (0110, 0210, 0310). We do not impute exposure scores for these unscored residual or armed-forces occupations.

### **A.3 World Bank national statistics**

We merge the occupational exposure panel with World Bank national statistics used in the descriptive and regression analyses, including GNI per capita in PPP terms, internet penetration, population, and remittance indicators. GNI and internet penetration enter the appendix predictor regressions as country-level development and infrastructure controls.

### **A.4 World Bank KNOMAD remittance**

We use World Bank WDI data on personal remittances received to measure remittance inflows in current U.S. dollars and as a share of GDP, using reported WDI values directly where available. We use the KNOMAD bilateral remittance matrix (`WB.KNOMAD.BRE`) to estimate source-country shares for remittance-receiving countries, again using the reported matrix values directly; the latest bilateral matrix available in our data is for 2021. The remittance examples therefore combine the latest available WDI remittance totals or remittance-to-GDP shares, generally 2024, with 2021 KNOMAD bilateral source-country shares. We do not impute missing bilateral corridors, extrapolate bilateral shares to later years, or assign synthetic exposure values.

### **A.5 Usage data statistics**

The observed adoption analysis uses three external usage statistics: Anthropic Economic Index country-level Claude usage, OpenAI Signals country rankings for population-normalized ChatGPT use, and Microsoft AI Diffusion country-level generative AI adoption rates [Appel et al., 2025, Handa et al., 2025, OpenAI, 2026, Microsoft AI Economy Institute, 2026a,b]. We match these external country statistics to the 141-country exposure panel using ISO-3 country codes where publishers provide them and otherwise using unambiguous country or economy names. We do not adjudicate contested political status or reassign entities across source definitions: labels such as Taiwan, China, Korea, or other country/economy groupings are retained as reported by the source, and an observation is used only when it maps unambiguously to an entity in the exposure panel. If a source reports an entity that does not cleanly match the exposure panel, we omit it rather than imputing or reallocating it. Each adoption regression uses the complete-case country sample available for the corresponding outcome and covariates. The resulting complete-case samples are 114 countries for Anthropic Claude usage, 88 for OpenAI Signals, and 106 for Microsoft AI Diffusion.

### **A.6 Reproducibility and asset provenance**

The public release for reviewers has two parts. The dataset release contains the derived measured CSV tables used in the paper, Croissant metadata, a data dictionary, and a source-data manifest. The code package contains validation scripts, figure and table builders, and two clean notebooks (`01_dataset_tour.ipynb` and `02_reproduce_main_results.ipynb`) that reproduce paper-facing values, regression tables, and release-supported figures from the released derived tables where redistribution permits. The analysis consists of deterministic data validation, aggregation, crosswalks, plotting, and ordinary least squares regressions run in Python on CPU; no model training or GPU computation is required. The released derived tables freeze the values used in this submission. Raw-source rebuilds are deterministic conditional on the source snapshots listed in Table 3, but some upstream ILO and World Bank mirrors may be revised over time; for this reason, the release workflow does not require live API pulls and instead reproduces results from the released tables. Where source data cannot be redistributed under the provider’s terms, the release records the original source, access date or version, redistribution status, and release tables supported in the source-data manifest. Raw third-party files and private working notebooks are not redistributed; the Microsoft AI Diffusion country/economy adoption table is included from Microsoft’s official MIT-licensed GitHub data release after ISO alpha-3 matching.

## **B Exposure Analysis**

### **B.1 Alternative Occupational Exposure Estimates**

Throughout our work, we use the occupational exposure estimates of Gmyrek et al. [2025b] because they are naturally designed for the ISCO-08 occupational categories. However, other exposure

Table 3: Source data and provenance for the main analyses.

Asset	Use in paper	Version or date	Credit and license/terms source
ILOSTAT employment and wage statistics	National occupational employment and wage composition	Most recent annual country statistics; accessed 2026-04-14	International Labour Organization [2026]; <a href="https://ilostat.ilo.org/data/">https://ilostat.ilo.org/data/</a>
Gmyrek et al. ISCO exposure estimates	Primary occupation-level exposure scores	2025 ILO Working Paper 140 release	Gmyrek et al. [2025b]
World Bank WDI and KNOMAD	GNI, remittance totals, and bilateral remittance shares	WDI 2024 where available; KNOMAD bilateral matrix latest year 2021	World Bank [2026a,b]; <a href="https://databank.worldbank.org/">https://databank.worldbank.org/</a>
Anthropic Economic Index	Country-level Claude usage validation	Country usage release used in Figure 4	Massenkoff et al. [2026]
OpenAI Signals	Population-normalized ChatGPT usage rank validation	Calendar-year 2025 country ranks	OpenAI [2026]
Microsoft AI Diffusion	Country-level generative AI adoption validation	Q1 2026 GitHub data update	Microsoft AI Economy Institute [2026a,b]; MIT license at <a href="https://github.com/microsoft/ai-diffusion-report">https://github.com/microsoft/ai-diffusion-report</a>
O*NET 30.2 and alternative exposure estimates	Occupational-property and robustness analyses	O*NET 30.2; source releases cited in appendix	O*NET Resource Center [2026], Eloundou et al. [2024], Hosseini Maasoum and Lichtinger [2026]

estimates are more popular in the broader economics of AI literature, which often derive from the U.S. Department of Labor such as the O\*NET work taxonomy and the SOC occupational list. Via a crosswalk that aligns the DOL resources and the SOC list with the ILO resources and the ISCO list, we can compare the occupational exposure estimates we use to other popular choices.

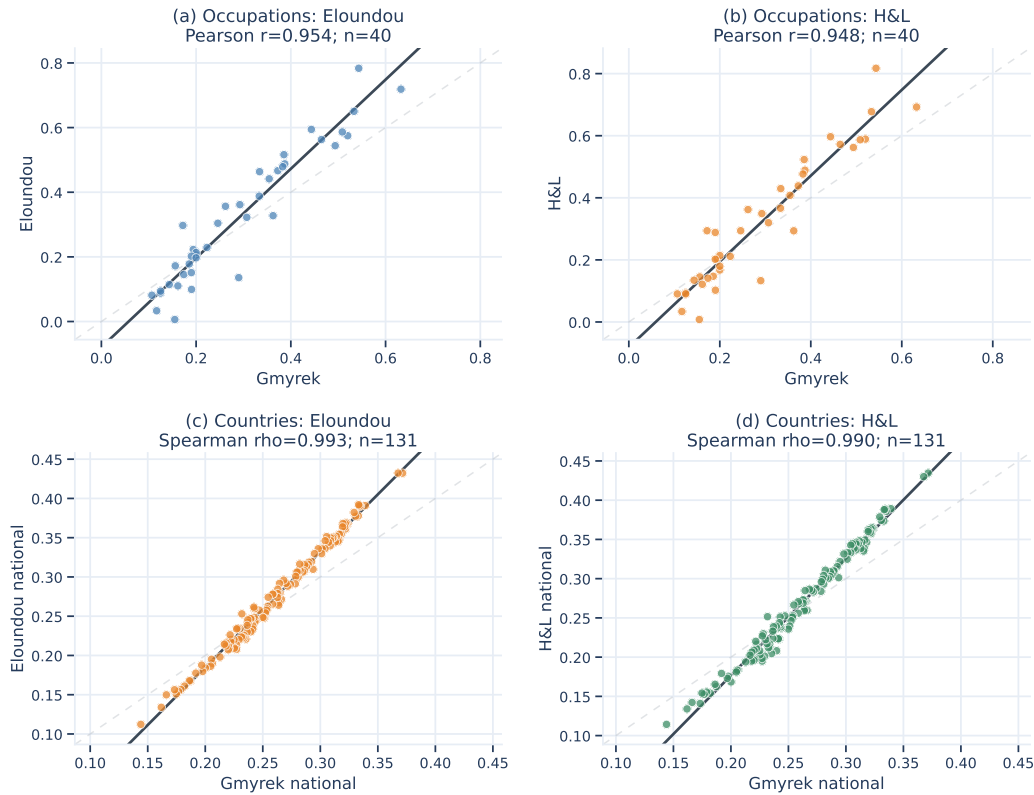
We compare our primary Gmyrek et al. exposure scores to two prominent alternatives: the LLM-task exposure measure from Eloundou et al. [2024] and the occupational exposure measure from Hosseini Maasoum and Lichtinger [2026]. Because these alternatives are defined over O\*NET/SOC occupations rather than ISCO-08 occupations, we use the SOC-to-ISCO crosswalk described above and aggregate scores to ISCO-08 2-digit occupations before recomputing national exposure scores.<sup>11</sup>

The estimates are highly correlated at the occupational and national levels (see Figure 6). At the occupation level, the resulting measures are highly correlated with the Gmyrek et al. [2025a] scores: Pearson correlations are  $r = 0.954$  for Eloundou et al. [2024] and  $r = 0.948$  for Hosseini Maasoum and Lichtinger [2026]. At the national level, rankings are even more stable, with Spearman correlations of  $\rho = 0.993$  and  $\rho = 0.990$ , respectively. The main disagreements are concentrated in a small number of occupations, such as ICT professionals and assemblers, but these differences do not materially change the cross-country ranking of national exposure.

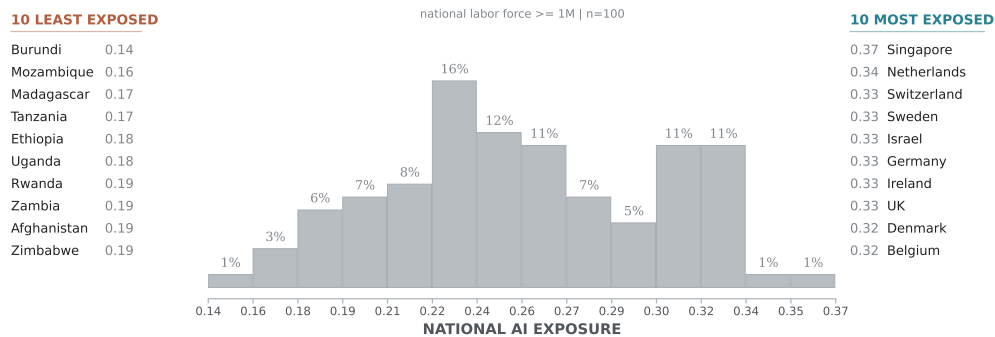
## B.2 National Exposure Distribution and Country Scores

The main text visualizes national AI exposure geographically in Figure 1. For completeness, Figure 7 shows the corresponding distributional view, including the least- and most-exposed countries among countries with at least one million workers. Table 4 reports the full 141-country exposure panel used in the map.

<sup>11</sup>The robustness pipeline maps O\*NET-SOC 8-digit occupations to SOC-2018 6-digit occupations, then to SOC-2010 6-digit occupations using the BLS 2018 reclassification crosswalk, then to ISCO-08 4-digit occupations using the BLS 2012 SOC-ISCO crosswalk, and finally to ISCO-08 2-digit occupations. At each merge, we take unweighted means across crosswalk rows. The BLS SOC-ISCO crosswalk is many-to-many and includes partial-overlap flags, which we do not weight separately. As a sensitivity check, assigning each SOC-to-ISCO mapping a part-share weight of  $1/k$  for a SOC code that maps to  $k$  ISCO codes changes the Gmyrek-vs-Eloundou Pearson correlation from 0.954 to 0.957 and the Spearman correlation from 0.950 to 0.955, leaving the robustness conclusion unchanged.



**Figure 6: National exposure estimates are robust to alternative occupational exposure measures.** The top row compares ISCO-08 2-digit occupational exposure scores. The bottom row compares national exposure scores after reweighting each occupational index by country-level occupational employment. The alternative indices differ in their exact task definitions and source occupational taxonomies, but they produce very similar occupational scores and highly stable national rankings.



**Figure 7: Distribution of national AI exposure estimates.** The side panels list the ten least- and most-exposed countries subject to the labor-force threshold used for visual readability; the histogram summarizes the full measured sample.

Table 4: National AI exposure estimates for all measured countries.

Country	<i>n</i>	Country	<i>n</i>	Country	<i>n</i>
Afghanistan	0.189	Germany	0.333	Norway	0.320
Albania	0.230	Ghana	0.220	Pakistan	0.225
Angola	0.206	Greece	0.315	Palau	0.295
Argentina	0.294	Guatemala	0.234	Palestine (State of)	0.253
Australia	0.322	Guinea	0.227	Panama	0.232
Austria	0.321	Guinea-Bissau	0.179	Papua New Guinea	0.236
Bahamas	0.282	Guyana	0.247	Peru	0.243
Bangladesh	0.218	Honduras	0.231	Philippines	0.252
Barbados	0.286	Hungary	0.309	Poland	0.317
Belarus	0.282	India	0.217	Portugal	0.314
Belgium	0.323	Indonesia	0.242	Romania	0.272
Belize	0.242	Iran (Islamic Republic of)	0.254	Russian Federation	0.298
Benin	0.227	Iraq	0.243	Rwanda	0.186
Bhutan	0.242	Ireland	0.331	Samoa	0.239
Bolivia (Plurinational State of)	0.241	Israel	0.333	Sao Tome and Principe	0.230
Bosnia and Herzegovina	0.263	Italy	0.313	Serbia	0.285
Botswana	0.228	Jamaica	0.266	Sierra Leone	0.205
Brazil	0.278	Japan	0.316	Singapore	0.368
Brunei Darussalam	0.307	Jordan	0.255	Slovenia	0.304
Bulgaria	0.292	Kenya	0.199	Solomon Islands	0.199
Burkina Faso	0.198	Kiribati	0.227	Somalia	0.259
Burundi	0.144	Kosovo	0.288	Spain	0.301
Cabo Verde	0.220	Kyrgyzstan	0.240	Sri Lanka	0.237
Cambodia	0.221	Lao People's Democratic Republic	0.205	Sudan	0.213
Chile	0.279	Latvia	0.304	Suriname	0.270
Colombia	0.263	Lebanon	0.264	Sweden	0.334
Congo, Democratic Republic of the	0.197	Lesotho	0.175	Switzerland	0.335
Cook Islands	0.290	Lithuania	0.309	Tajikistan	0.222
Costa Rica	0.274	Luxembourg	0.371	Tanzania, United Republic of	0.173
Croatia	0.311	Madagascar	0.166	Thailand	0.250
Cuba	0.265	Malawi	0.223	Timor-Leste	0.208
Cyprus	0.322	Maldives	0.280	Togo	0.235
Czechia	0.302	Mali	0.200	Tokelau	0.273
Côte d'Ivoire	0.240	Marshall Islands	0.268	Tonga	0.259
Denmark	0.324	Mauritius	0.289	Tunisia	0.236
Dominican Republic	0.264	Mexico	0.264	Tuvalu	0.305
Ecuador	0.222	Micronesia (Federated States of)	0.219	Türkiye	0.265
Egypt	0.254	Mongolia	0.263	Uganda	0.182
El Salvador	0.251	Montenegro	0.318	United Arab Emirates	0.281
Estonia	0.311	Montserrat	0.290	United Kingdom of Great Britain and Northern Ireland	0.329
Eswatini	0.244	Mozambique	0.162	United States of America	0.323
Ethiopia	0.177	Myanmar	0.229	Uruguay	0.281
Fiji	0.259	Nepal	0.233	Vanuatu	0.218
Finland	0.309	Netherlands	0.339	Viet Nam	0.225
France	0.319	Niger	0.241	Wallis and Futuna	0.291
Gambia	0.234	Niue	0.314	Zambia	0.187
Georgia	0.237	North Macedonia	0.286	Zimbabwe	0.192

### B.3 Gender Differences in National AI Exposure

Figure 8 visualizes the country-level gender exposure gap discussed in the main text. The x-axis reports the female-minus-male exposure gap as a percentage of total national exposure, so positive values indicate that female exposure exceeds male exposure. The pattern is broad rather than driven by a handful of countries: women are more exposed than men in most measured countries, while the largest negative gaps appear in countries where women’s employment remains more concentrated in agriculture and other less-exposed occupations. Among the 141 countries in the main exposure panel, 138 have reliable sex-disaggregated occupational comparisons. Egypt, Micronesia, and Wallis and Futuna remain in the national exposure panel, but are excluded from this gender-gap analysis because their sex-disaggregated occupational cells do not meet the reliability criteria used for the appendix figure.

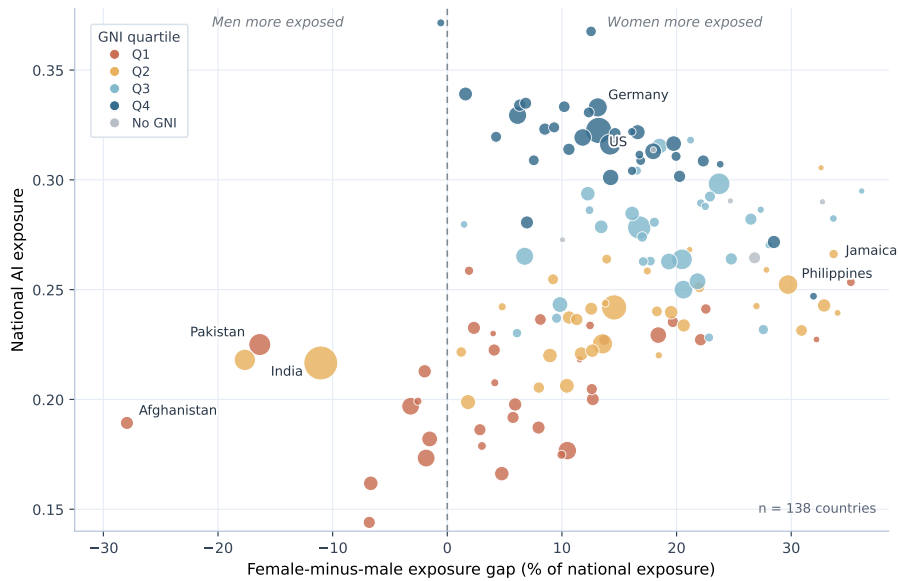


Figure 8: **Gender differences in national AI exposure.** Positive values indicate that female exposure exceeds male exposure, measured as  $(n_{female} - n_{male})/n_{total}$ . Points are countries with sex-disaggregated occupational employment data; colors indicate GNI per capita quartiles and bubble sizes scale with national employment.

### B.4 Predicting National AI Exposure

Beyond the main white-collar specification, we test whether national exposure can be predicted from broader national and occupational covariates. Among basic national covariates, white-collar share is the strongest single predictor; log GNI and internet penetration are also positively associated with exposure, but explain less cross-country variation on their own. We construct the cognitive-minus-physical (CMP) score as a task-composition diagnostic from O\*NET 30.2. At the ISCO-08 2-digit occupation level, we map O\*NET SOC occupations to ISCO-08 and compare O\*NET occupational properties with the Gmyrek et al. exposure scores.

The cognitive side is the average of the five O\*NET properties most positively associated with exposure: Working with Computers, Written Comprehension, Reading Comprehension, Administrative, and Computers and Electronics. The physical side is the average of the five properties most negatively associated with exposure: Trunk Strength, Stamina, Performing General Physical Activities, Static Strength, and Gross Body Coordination. For each occupation, CMP is the cognitive average minus the physical average; for each country, national CMP is the employment-share-weighted average CMP across its measured ISCO-08 occupations. Higher CMP therefore means that a country’s employment is concentrated in occupations with more computer/document/information-processing content and less physical/manual content. This O\*NET-derived CMP score provides a more task-grounded measure of labor composition and predicts exposure more strongly than the coarse white-collar split,

indicating that the underlying cognitive and physical organization of work is the main driver rather than the white-collar label itself.

Table 5: Predicting national AI exposure.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Specification	WC	GNI	Internet	CMP	WC+GNI	WC+Internet	WC+CMP	Full
White-collar share	0.242*** (0.007)				0.196*** (0.011)	0.212*** (0.008)	-0.046*** (0.017)	-0.038** (0.015)
log GNI		0.036*** (0.002)			0.009*** (0.002)			0.004* (0.002)
Internet penetration			0.157*** (0.010)			0.031*** (0.009)		-0.007 (0.007)
Cognitive-minus-physical score				0.105*** (0.001)			0.124*** (0.007)	0.117*** (0.007)
Observations	141	135	130	134	135	130	134	124
R <sup>2</sup>	0.907	0.772	0.669	0.973	0.928	0.925	0.975	0.979

Note: Entries report OLS coefficients with HCl1 robust standard errors in parentheses. The dependent variable is national AI exposure. Internet penetration is measured as a 0–1 share. Each column uses the complete-case sample for the variables in that specification. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

## B.5 Predicting National AI Adoption

We also estimate the full set of models that include or exclude white-collar share, log GNI, and national AI exposure for each of the three observed adoption outcomes. National AI exposure strongly predicts adoption on its own across all three datasets. However, once white-collar share and log GNI are included, adding national AI exposure contributes essentially no additional explanatory power and the full-model exposure coefficient is not statistically distinguishable from zero. This suggests that exposure is useful as a structural summary, while the observed cross-country adoption relationship is largely aligned with broader national labor composition and income.

Table 6: Predicting national AI adoption: full regression specifications.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: Anthropic Claude</i> (dependent variable: log10(Claude usage / 100k WAP))								
	Intercept	White-collar	log GNI	Exposure	White-collar + log GNI	White-collar + exposure	log GNI + exposure	Full model
Specification	only	only	only	only	only	only	only	only
White-collar share		2.761*** (0.123)			1.123*** (0.270)	1.878*** (0.545)		1.112*** (0.423)
log GNI			0.471*** (0.020)		0.305*** (0.051)		0.336*** (0.061)	0.304*** (0.059)
National AI exposure				10.643*** (0.453)		3.602* (2.142)	3.561*** (1.296)	0.056 (1.986)
Observations	114	114	114	114	114	114	114	114
R <sup>2</sup>	0.000	0.792	0.837	0.769	0.863	0.799	0.854	0.863
<i>Panel B: OpenAI Signals</i> (dependent variable: OpenAI rank percentile)								
	Intercept	White-collar	log GNI	Exposure	White-collar + log GNI	White-collar + exposure	log GNI + exposure	Full model
Specification	only	only	only	only	only	only	only	only
White-collar share		1.343*** (0.056)			0.571*** (0.105)	0.834*** (0.212)		0.686*** (0.144)
log GNI			0.215*** (0.009)		0.137*** (0.017)		0.152*** (0.025)	0.144*** (0.021)
National AI exposure				5.215*** (0.219)		2.084*** (0.777)	1.761*** (0.590)	-0.634 (0.724)
Observations	88	88	88	88	88	88	88	88
R <sup>2</sup>	0.000	0.828	0.874	0.812	0.908	0.839	0.891	0.909
<i>Panel C: Microsoft AI Diffusion</i> (dependent variable: MS GenAI adoption Q1 2026 (% WAP))								
	Intercept	White-collar	log GNI	Exposure	White-collar + log GNI	White-collar + exposure	log GNI + exposure	Full model
Specification	only	only	only	only	only	only	only	only
White-collar share		53.981*** (4.045)			37.123*** (8.541)	45.650*** (11.606)		39.374*** (11.082)
log GNI			8.041*** (0.752)		2.990** (1.381)		3.908** (1.627)	3.094* (1.639)
National AI exposure				203.888*** (16.654)		34.030 (38.948)	116.761*** (38.042)	-11.584 (50.533)
Observations	106	106	106	106	106	106	106	106
R <sup>2</sup>	0.000	0.651	0.599	0.611	0.670	0.653	0.641	0.671

Note: Entries report OLS coefficients with HCl1 robust standard errors in parentheses. The table contains 24 models: three adoption outcomes times eight predictor specifications. Each panel uses the outcome's complete-case country sample across white-collar share, log GNI, and national AI exposure. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .