AI at Home, AI at Work

Fatih Kansoy*
University of Oxford

This version: November 2026
Latest version: fatih.ai/ai_home_work.pdf

Abstract

Does generative AI serve mainly households or workplaces? Early studies disagree because each observes one channel or scores hypothetical task capability. We combine public logs from ChatGPT, Claude, and Google Gemini with employment and hourly pay surveys across seven Western economies. Wherever privacy rules withhold joint tables, we pair baseline accounts with partial-identification bounds over all admissible allocations. Free consumer surfaces handle mainly home life and study on all three platforms while paid plans and commercial APIs shift strongly towards paid work. Within each provider, task mix tracks access channels and personal-service intensity far more closely than national borders. In paid employment, broad job families and pre-AI task ratings explain most occupational pay gradients on both Claude and directly reported US Gemini records. Coworker wage dispersion lifts clerical task reach in the bottom fifth even as hourly pay weights concentrate labour-cost exposure at the top. Commercial APIs execute mostly direct instructions whereas human chat stays interactive in commercial and regulatory work. Their bounds overlap only where technical teams rely on feedback loops. Because public logs match prompts to tasks rather than workers, these estimates capture allocated incidence at current prices rather than causal effects.

Keywords: Generative AI, task content, household production, occupational pay, partial identification, workplace delegation
JEL classification: J24, J31, D13, O33, C14

*University of Oxford. Email: fatih.kansoy@economics.ox.ac.uk.

1. Introduction

Is a language model mainly a tool for work? Our answer depends on where we look. On free search pages and consumer chat apps, most prompts concern home life such as meals, health, shopping, and schoolwork. Paid work rises in subscription plans and exceeds four fifths of calls on Anthropic's commercial application programming interface (API), though longer work tasks bring it near half of estimated hours saved in chat. In earlier waves of office computerisation, employers bought the machines and put them on office desks (Autor, Levy, and Murnane 2003). Language models entered homes, jobs, and company software at the same time. If analysts treat every chatbot prompt as workplace labour, they fold unpaid domestic time savings (Becker 1965) into market output (Acemoglu 2025, Jones 2026). If they look only at company software, they miss everyday use.

The early evidence on generative AI points in three different directions. In our view, this disagreement is easy to explain because each study observes one channel or scores what a model could do rather than how people use it. Consumer logs from ChatGPT and Google Gemini show that personal and educational queries outnumber workplace prompts in almost every economy (Chatterji et al. 2025, OpenAI 2026, Iscenko et al. 2026). Occupational exposure indices instead rank which tasks a model could speed up in principle and point to well-paid professions (Felten, Raj, and Seamans 2021, Eloundou et al. 2024). Worker adoption surveys (Humlum and Vestergaard 2025, Bick, Blandin, and Deming 2026) and raw platform pay regressions (Handa et al. 2025, Iscenko et al. 2026) echo that same upper-tier pattern. Commercial API reports depict back-office software driven by direct instructions (Appel et al. 2026, Massenkoff et al. 2026). Read alone, each source shows a different technology until we track how users sort across channels.

This paper asks three linked questions about how generative AI is used. We examine how demand divides between unpaid life and paid work across channels, which occupations and employee pay bands encounter AI tasks at work, and whether human chat differs from firm software. A simple choice between saved human time and fixed access costs links all three questions (Becker 1965, Acemoglu and Restrepo 2018). Free search and chat charge nothing at entry and attract everyday household tasks. Paid plans and commercial APIs offer higher capacity, wider context windows, or software integration for fees or setup costs that frequent workplace demand spreads over many requests. Once flat-rate access is paid, extra prompts carry no monetary price and workplace intensity should follow codifiable desk tasks and office hierarchies rather than pay alone (Garicano 2000, Autor and Thompson 2025). Wide pay dispersion inside narrow job titles then separates task reach across pay bands from hourly wage costs (Autor and Handel 2013).

No public file links prompts to a user's job, wage, and employer. We therefore combine public usage releases from all three major model providers with official employment and hourly earnings surveys across seven Western economies (the United States, the United Kingdom, Germany, France,

Spain, Italy, and the Netherlands). OpenAI Signals provides twenty-seven monthly waves of adult ChatGPT messages (July 2024 to September 2026) across 160 economies and four plans (OpenAI 2026). Google Gemini ATLAS adds 14.65 million interactions across 218 economies in April 2026 alongside household time-use categories and directly reported US workplace shares across 335 Broad occupations (Iscenko et al. 2026). Anthropic's Economic Index spans 121 economies across five waves and applies one shared task and interaction-mode rubric across consumer Claude chat and commercial first-party API calls (Massenkoff et al. 2026). Official surveys connect these logs to 184 million American and British payroll jobs with internal pay percentiles, 124 million Continental European workers, pre-ChatGPT task descriptors, and earlier capability scores.

Public platform releases protect privacy through small-cell suppression and omitted cross-tabulations. Providers report work shares and modes without overlap. Anthropic withholds national six-digit work shares while official surveys publish pay percentiles inside occupations without coworker prompt ranks. Our strategy pairs baseline accounts with partial-identification bounds (Fréchet 1951, Charnes and Cooper 1962, Manski 2003). For channel and task sorting, we combine country-fixed-effects regressions with clustered bootstrap inference, zero-parameter cross-country forecasts, and an API parity threshold. Employment-weighted Poisson and shifted-log regressions estimate pay gradients across all three panels while Google's US shares check raked Claude allocations and continuous pay grids separate task shares from wage costs. Wherever joint cells are unobserved, linear programmes, Fréchet bounds, and breakdown thresholds span every value consistent with published totals and gauge how strong unobserved selection must be to overturn a result (Cross and Manski 2002, Masten and Poirier 2020).

Consumer AI use is mainly domestic across all three ecosystems while work shares rise sharply in paid channels. Non-work queries form the majority across almost all economies on Gemini, ChatGPT, and Claude as households ask about leisure, study, cooking, health, and shopping. Within each provider, work-related shares climb across OpenAI's consumer plans and paid-work shares jump from consumer chat to Anthropic's commercial API. Plan gaps reflect user and employer selection alongside fees rather than pure price elasticities (Datta, Lambert, and Schneider 2026). Across our seven Western economies, national work shares stay tight within each provider yet vary widely across providers and tasks. Personal-service tasks such as health care, food, and social support stay overwhelmingly domestic when families seek free advice at home and in-person delivery slows workplace use (Cerina, Moro, and Rendall 2021). Six countries predict the seventh cleanly.

Inside paid work, steep occupational pay slopes mostly reflect differences across broad job families and task capabilities rather than pay gaps within one field. Unconditional workplace AI intensity rises strongly with median pay on American and British Claude records and directly reported US Gemini records. Once we control for two-digit job families, pre-ChatGPT task descriptors, and capability scores, most of that slope disappears in nearly every specification and worst-case

level bounds span zero. Only Google's directly matched broad-occupation subset keeps a modest positive slope. Single-digit aggregation reproduces the steep slope in coarse Continental European tables (Cross and Manski 2002). Internal pay dispersion also matters. Because many clerical workers earn below their occupation's median, continuous pay grids under equal within-occupation use lift the bottom fifth's allocated share severalfold on Claude and Gemini. That gain stays strictly positive across British joint bounds and eligible-title Gemini allocations while US Claude joint bounds span zero. Even so, steep hourly wage gaps keep most wage-weighted exposure in the top fifth.

When a workplace task reaches a model, commercial APIs receive direct instructions far more often than chat windows. Full-population bounds under Anthropic's five-mode rubric (Massenkoff et al. 2026) do not overlap between chat and first-party APIs even under worst-case suppressed and unclassified prompts. Individual workflows show where human dialogue stays central. In software development, infrastructure operations, and data analysis, multi-turn feedback loops let chat bounds overlap API bounds as engineers paste error logs and tables. Their ordering then becomes unsigned. Chat stays strictly below commercial API floors in regulatory compliance, copywriting, and sales operations even after feedback loops enter. Firms thus automate standardised bulk tasks on APIs and keep bespoke persuasion and policy judgement in chat.

In our view, these findings reconcile three strands of AI labour research. Platform diffusion studies (Chatterji et al. 2025, Iscenko et al. 2026) find mainly non-work traffic whereas worker surveys (Humlum and Vestergaard 2025, Bick, Blandin, and Deming 2026) report rapid workplace adoption. Both hold once we separate prompt volume from person-level use and see that channels and tasks outweigh national borders. Next to exposure indices (Felten, Raj, and Seamans 2021, Eloundou et al. 2024), platform pay studies (Handa et al. 2025, Iscenko et al. 2026, Datta, Lambert, and Schneider 2026), and macro cost models (Acemoglu 2025, Jones 2026), our common-support comparison of Claude and Gemini shows that raw pay gradients mostly separate broad white-collar families from manual work. Our internal wage grids also show that clerical task reach and top-heavy wage-bill exposure are two sides of one pay distribution. Compared with workplace trials (Noy and Zhang 2023, Brynjolfsson, Li, and Raymond 2025, Cui et al. 2026) and API reports (Appel et al. 2026), our assumption-free bounds (Manski 2003, Masten and Poirier 2020) separate routine API execution from workflows where human iteration persists. To assess how AI reshapes work, economists must examine consumer apps, internal pay bands, and firm software together.

Our estimates measure task and wage incidence at current prices rather than causal wage or employment effects. Public logs also match prompts to tasks rather than workers and omit enterprise workspaces, coding agents, and human review after an API call. Section 2 details the datasets, surveys, framework, and bounds. Section 3 examines how AI use divides between home and work across channels, tasks, and countries. Section 4 analyses pay gradients and pay bands before Section 5 compares chat with APIs and Section 6 concludes.

2. Data and Method

To study how generative AI enters homes and workplaces at population scale, we pair public usage releases from three model providers with official labour surveys. No single company covers the whole market. No public file links prompts to wages. Each provider also serves a distinct mix of casual consumers, paid subscribers, and business software clients (Yin and Ogut 2026). Between 2024 and 2026, OpenAI, Anthropic, and Google released privacy-screened aggregate datasets built from millions of real interactions so researchers could measure actual AI demand while private prompt text stayed on secure servers (Table 1, Panel A). In each release, automated classifiers assign prompts to work context, job tasks, or interaction modes before privacy rules add noise and suppress small cells. Together, the three ecosystems give us a twenty-seven-month panel of consumer plans from OpenAI, a shared chat-and-API rubric from Anthropic, and global household time-use tables plus directly reported US workplace occupation shares from Google.

OpenAI Signals version 3.0 provides our longest view of mass consumer demand through 27 monthly waves of adult ChatGPT messages from July 2024 to September 2026 across 160 economies (149 in September 2026) and the Free, Go, Plus, and Pro consumer plans (OpenAI 2026). Enterprise seats and Codex are excluded (Chatterji et al. 2026). OpenAI samples and reweights messages, adds Gaussian differential-privacy noise, and drops small cells (Chatterji et al. 2025, OpenAI 2026). We keep OpenAI's two context labels separate. Its binary label records whether a message relates to work across countries, months, and September 2026 plans without a separate school category. Its global three-way series divides messages into paid work, schoolwork, and personal queries alongside monthly topic shares. A separate September 2026 table splits work-related messages across 126 economies into output-oriented *Doing*, query-based *Asking*, and *Expressing* shares. This dataset lets us test whether non-work consumer traffic stabilised over two years, how work-related messages climb across paid consumer plans, and how ChatGPT's workplace task rankings align with Claude.

The Anthropic Economic Index adds 5 waves between August 2025 and May 2026 and provides the only public release that applies one shared task and interaction-mode rubric across both consumer chat and paid commercial APIs (Handa et al. 2025, Appel et al. 2026, Massenkoff et al. 2026). Four waves from November 2025 onward classify paid work, personal use, coursework, and five interaction modes with small-cell suppression. Consumer tables pool Claude.ai web, desktop, and mobile conversations with Cowork across 121 economies, 52 American states and territories, 652 subregions, 1,007 May 2026 task clusters, 718 O*NET occupation work-share priors, and model-estimated human task durations. Commercial first-party API files cover twenty workflows and 860 May 2026 clusters. Claude Code and third-party cloud hosts such as AWS Bedrock and Google Vertex AI are omitted. Pipeline design changed in April 2026. Anthropic moved from one-week snapshots to full-month hourly stratified samples and widened the classifier's text window on long

TABLE 1. Data Releases, Official Surveys, and What Public Files Identify

Panel A. Platform Releases, Official Labour Surveys, and Constructed Panels

Source and window	Scale and unit	Published content and exclusions
OpenAI Signals v3.0 (Jul 2024–Sep 2026)	160 economies (149 in Sep 2026); 27 waves; adult messages on 4 plans	Binary work share by country, month, and plan; global work/school/other; topics; Doing within work (126 econ.); omits Enterprise and Codex
Anthropic Consumer (Aug 2025–May 2026)	May 2026: 121 econ., 52 US states/terr., 652 subregions; 5 waves	Claude.ai and Cowork conversations (4 context waves); work/personal/coursework; 5 modes; 1,007 May clusters; 718 O*NET work shares; durations
Anthropic 1P API (Aug 2025–May 2026)	Global first-party API calls; 5 waves (full-month from Apr 2026)	Global context and 5 interaction modes across 20 workflows and 860 May clusters; omits Claude Code and third-party cloud hosts
Google Gemini ATLAS v1.0 (Apr 2026)	218 economies, 51 US jurisdictions; 14.65M Search & App interactions	Country/state work shares; ATUS household categories and night/weekend blocks; 335 direct US Broad SOC shares; omits paid API and Workspace
US BLS OEWS (May 2025)	830 6-digit SOCs (155.48M jobs); 334 matched Broad SOCs (143.62M)	437 active and 393 zero or suppressed 6-digit titles retained; mean hourly pay and 5 hourly percentiles
UK ONS APS (2025) & ASHE (2026)	402 paid 4-digit SOC2020 groups (33.07M APS; 28.85M ASHE jobs)	26 job families (29 zeros); APS workforce weights incl. self-employed; ASHE mean pay and 11 hourly percentiles
Eurostat SES (2022) & EU-LFS (2025)	DE, FR, ES, IT, NL; 40 country $\times$ 1-digit ISCO cells (123.91M workers)	Within-country LFS shares; SES mean hourly pay in firms ≥ 10 staff; 8 major groups per country; no within-group wage percentiles
Constructed Panels (2025–2026)	154 country $\times$ 2-digit SOC cells (7 econ.); 184.33M US/UK pay-grid jobs	Multi-provider context shares, 22-group task sorting, pre-AI task indices (R_g, S_g), capability scores, and continuous employee wage grids

Panel B. What Public Files Identify and Where Bounds Apply

Target quantity	Published inputs vs. withheld cells	Identification status and section
Work share by country, plan, and task group	Country shares (3 platforms), OpenAI plans, Gemini ATUS blocks, 7-economy 2-digit shares	Directly observed; trends, channel sorting, task regressions, and held-out prediction (§3)
Combined chat and API work share	National consumer shares and global API work shares published; national API volumes withheld	Thresholded; parity frontier transporting global API work share and hours-saved ratio (§3)
Workplace AI by occupation and pay	Direct US Gemini Broad shares (334); Claude 2-digit totals and 718 O*NET priors	Direct Broad-SOC regressions (Gemini); baseline raking and LP level-slope bounds (Claude) (§4)
Workplace AI by employee pay fifth	US 5-knot and UK 11-knot pay grids (184.33M jobs); coworker prompt ranks unobserved	Equal-prompting traffic H_q and wage-weighted B_q (Claude & Gemini); grid/LP bounds and tilts (§4)
Directive share of workplace requests	Anthropic work and 5-mode shares published separately (ChatGPT Doing; Gemini automation)	Bounded; Fréchet–Hoeffding intervals retaining suppressed and unclassified prompts (§5)
Request mode by individual worker pay	Platform interaction logs and survey earnings are unlinked across individuals	Unidentified; motivates privacy-safe aggregate tabulations (§6)

Notes: Panel A lists the three platform releases (OpenAI 2026, Appel et al. 2026, Massenkoff et al. 2026, Iscenko et al. 2026; see also Handa et al. 2025), US and UK surveys (U.S. Bureau of Labor Statistics 2026, Office for National Statistics 2025, 2026), and Continental European surveys (Eurostat 2025b,a). OpenAI excludes Enterprise and Codex; Anthropic consumer pools Claude.ai chat and Cowork while its commercial 1P API excludes Claude Code and third-party cloud hosts; Google ATLAS conversational tables pool Search AI Mode and Gemini app while occupational tables add free-tier Gemini API and omit paid API text, Vertex AI, and Workspace. Panel B maps empirical targets by identification status.

prompts. Unclassified API modes then fell from 15.2 to 5.5 per cent. Because the rubric stays fixed across both channels, Anthropic's data anchor our seven-economy task regressions, our detailed American and British occupational panels, and our chat-versus-API delegation bounds.

Google's AI and Economy ATLAS version 1.0 covers two weeks in April 2026 (Iscenko et al. 2026). Google builds these public files from 14.65 million de-identified interactions across 218 economies and 51 US jurisdictions in clusters of at least ten distinct users after contribution bounds, differential-privacy noise, and small-cell suppression. Geographic and household tables pool conversational Gemini app and Google Search AI Mode interactions and map non-work queries onto American Time Use Survey categories across weekday business hours, weekday nights, and weekends. Separate occupational files combine those conversational surfaces with content-eligible free-tier Gemini API and AI Studio calls across 335 US-reported Broad SOC occupations and twenty-two Major groups. Military codes, "All Other" residuals, paid API text, Vertex AI, and Workspace are omitted. Google's release plays two roles here. Its global geographic and time-use files check whether consumer non-work majorities hold on search-led surfaces while its directly reported US Broad SOC shares test our American pay gradients without global-to-national adjustment.

Official employment and hourly earnings surveys in seven Western economies connect these platform task logs to real jobs and wage distributions. Prompts are not workers. Because text classifiers across all three providers assign each prompt to an O*NET or SOC task statement without access to the user's true job title or employer, our occupational panels measure task incidence rather than linked worker-level use. That distinction guides Section 4. We study the United States, the United Kingdom, Germany, France, Spain, Italy, and the Netherlands because they combine over 308 million workers with a sharp contrast in statistical detail. The United States and the United Kingdom publish fine occupations alongside internal hourly pay percentiles while Eurostat publishes only coarse single-digit occupational earnings in Continental Europe. In the United States, the May 2025 Occupational Employment and Wage Statistics (OEWS) cover 155.48 million employees across 830 six-digit SOC occupations with mean hourly pay and five hourly pay percentiles (U.S. Bureau of Labor Statistics 2026). We keep all 393 zero-coded or suppressed titles alongside 437 active titles so that manual jobs stay in the sample. When we match Google's 335 US Broad SOC codes to civilian OEWS wages, we retain 334 Broad occupations (all except legislators, who lack a published median hourly wage) that employ 143.62 million workers.

British and European surveys complete the panel. In the United Kingdom, the 2025 Annual Population Survey (APS) counts 33.07 million workers across 402 paid four-digit SOC2020 unit groups in twenty-six job families and covers the self-employed (Office for National Statistics 2025). APS employment counts weight our O*NET-to-British crosswalk so that tiny sub-occupations cannot distort unit-group averages. The 2026 Annual Survey of Hours and Earnings (ASHE) adds employer-reported hourly pay without overtime across eleven percentiles and means for 28.85 million matched payroll jobs (Office for National Statistics 2026). Together, the American and British surveys

give us 184.33 million payroll jobs for continuous within-occupation wage grids. For Germany, France, Spain, Italy, and the Netherlands, Eurostat pairs 2022 Structure of Earnings Survey mean hourly pay (firms of ten or more staff) with 2025 Labour Force Survey employment across 40 single-digit ISCO-08 cells that cover 123.91 million workers (Eurostat 2025b,a). Internal wage percentiles are unavailable. That contrast with the American and British files lets us measure directly how single-digit aggregation alters occupational pay slopes.

To test whether workplace AI tracks job content rather than hourly pay alone, we need task measures fixed before our 2025–2026 usage logs. We build two standardised task indices from pre-2023 O*NET and PIAAC descriptors (Kansoy 2026). Structured desk intensity R_g captures computer use, documentation, and numerical work (Autor, Levy, and Murnane 2003). Personal-service intensity S_g measures face-to-face social interaction minus half of problem-solving demand and half of computer use. It peaks in health support, care, food preparation, and teaching where families can substitute free advice at home for paid market services (Cerina, Moro, and Rendall 2021). Because R_g and S_g load on computer use with opposite signs, they enter separate country-fixed-effects regressions across our 154 seven-economy two-digit SOC cells in Section 3. Detailed occupational regressions in Section 4 also add the occupational AI capability scores of Felten, Raj, and Seamans (2021) and Eloundou et al. (2024). Both scores predate our usage window and omit platform conversation logs.

A simple choice between saved human time and fixed access costs organises how households, workers, and firms sort across channels (Becker 1965, Acemoglu and Restrepo 2018). Suppose task g requires t_g human hours valued at v per hour, equal to the wage inside firms and the shadow value of unpaid time at home. Channel k carries a fixed access or setup cost F_k, a marginal query price p_k, and an effective net time-saving rate σ_{gk} set by model capacity, rate limits, context length, or software integration ($\sigma_{gk} > \sigma_{g0}$ when $F_k > 0$). Once F_k is paid, a user brings task g to channel k whenever $v\sigma_{gk}t_g > p_k$. Free search and chat set $F_0 = p_0 = 0$ at entry while flat-rate plans set $p_k = 0$ after a monthly fee. Over N repeated requests, the net gain from a paid channel equals $Nv(\sigma_{gk} - \sigma_{g0})t_g - F_k$. Heavy workplace users and high-value firm workflows therefore benefit more from paid plans or API integration while occasional domestic queries stay on free surfaces. Across consumer tasks, paid workplace volume $M_{cg}^{W}(\alpha_c, R_g)$ rises with national corporate adoption α_c and desk content R_g while non-work volume $M_{cg}^{\mathrm{NW}}(S_g)$ expands in personal-service tasks S_g. Whenever workplace volume is positive and household volume rises strictly with S_g, the consumer paid-work share satisfies

$$w_{cg}(S_g, R_g; \alpha_c) = \frac{M_{cg}^{W}(\alpha_c, R_g)}{M_{cg}^{W}(\alpha_c, R_g) + M_{cg}^{\mathrm{NW}}(S_g)}, \qquad \frac{\partial w_{cg}}{\partial S_g} < 0. \tag{1}$$

Slower employer adoption in services delivered in person also lowers M_{cg}^{W} and reinforces the negative sign in Equation (1). As a result, Section 3 estimates the net association between S_g and

w_{cg}. When workplace and non-work volumes carry common country multipliers, log odds separate into country and task terms that preserve task rankings across borders.

Inside paid employment, flat-rate access and organisational hierarchies shape how workplace AI relates to pay and software integration. Higher wages raise the value of saved time v and encourage adoption. Yet zero marginal prompt prices on flat-rate plans imply that conditional workplace intensity should follow codifiable task content rather than hourly pay alone. Knowledge-hierarchy and social-skill models (Garicano 2000, Deming 2017, Autor and Thompson 2025) predict that workplace prompt volume flattens near the top of white-collar families because mid-paid specialists handle much of the codified text and analysis while senior executives spend more hours on negotiation, supervision, and sign-off. Inside narrow job titles, wide internal wage dispersion (Autor and Handel 2013) allows lower-paid clerical staff to lift bottom-fifth allocated traffic shares even as steep hourly pay gaps keep hourly-wage-weighted exposure concentrated in the top fifth under cost-share accounting (Acemoglu 2025, Jones 2026). API integration also requires upfront software work ($F_{\text{API}} > 0$) that repeatable direct instructions amortise more easily than open-ended chat drafts (Handa et al. 2025, Brynjolfsson, Li, and Raymond 2025).

Privacy rules withhold the interior cross-tabulations in these public files (Table 1, Panel B). That omission matters. Suppose a task group reports that 60 per cent of consumer conversations concern paid work and 50 per cent are direct instructions without a published overlap. Anywhere between 10 and 50 per cent of all conversations can belong to both groups. The directive share of workplace conversations can therefore lie anywhere from one sixth to five sixths. The same observability gap appears when national Anthropic files publish two-digit job-family work totals and all-context caps without national six-digit work shares, or when official wage surveys publish hourly pay percentiles inside occupations without coworker prompt ranks. To avoid unverified independence assumptions, we pair transparent baseline accounts (logit-raked occupational allocations and equal-use continuous pay grids) with partial-identification bounds that optimise over every joint table consistent with the published margins (Fréchet 1951, Charnes and Cooper 1962, Manski 2003). We also compute breakdown thresholds that show how large unrecorded national API volumes or within-occupation pay tilts would have to be to overturn a headline result (Cross and Manski 2002, Masten and Poirier 2020). Online Appendix A documents all crosswalks and wage-tail rules while Online Appendix B derives each bounding programme.

3. AI at Home and Across Access Channels

Most consumer interactions with language models happen outside paid work on every platform (Figure 1, Panel A; Table 2, Panel A). In Google's April 2026 ATLAS release across 218 economies, non-work queries hold the majority in every economy around a median of 88.66 per cent (Iscenko

et al. 2026). OpenAI's adult consumer plans show how quickly mass adoption shifted demand towards everyday life. The cross-economy median binary non-work share rose from 51.45 per cent in July 2024 to 70.40 per cent a year later as casual users arrived, then held inside a narrow plateau for fourteen straight months at 70.40 per cent across all 149 economies in September 2026 (Chatterji et al. 2025, OpenAI 2026). Consumer Claude confirms the same pattern under an independent classifier, with personal use and coursework above half in 113 of 121 economies (median 59.18 per cent), in every American state and territory, and in 587 of 652 global subregions (Massenkoff et al. 2026). Yet the three ecosystems sit nearly thirty percentage points apart even across the exact 121 economies reported by all three providers (Online Appendix Table A5). That gap reflects product design and user mix rather than geography. Brief search queries on Google carry far less workplace context than longer multi-turn workspace chats on Claude, while ChatGPT bridges both audiences.

Context splits and household time-use tables show that everyday domestic needs and education account for most consumer demand, while label design shapes how large work appears. On ChatGPT, OpenAI's binary work-related share (32.2 per cent in September 2026, down from 51.6 per cent in July 2024) sits well above its three-way paid-work share (21.0 per cent, down from 27.3 per cent) because the binary label folds in schoolwork. Outside paid jobs and study, personal queries climbed from 48.7 to 65.5 per cent of global ChatGPT messages, led by practical guidance, information lookup, and everyday writing (69.1 per cent of non-work traffic; OpenAI 2026). Anthropic separates all three contexts directly and finds 40.20 per cent personal use and 16.45 per cent coursework globally against 43.36 per cent paid work. Google's American Time Use Survey tables confirm the same domestic mix in the United States (Online Appendix Table A5). Leisure leads at 32.0 per cent alongside homework and research (22.5 per cent), household chores and cooking (15.6 per cent), shopping and personal services (13.6 per cent), and health self-care (9.0 per cent). Almost two thirds of those household prompts (63.9 per cent in America and 69.0 per cent globally) occur on weekday nights and weekends. Because 64.3 per cent of 112 awake weekly hours fall outside a standard forty-hour workweek, off-hours timing alone cannot prove a user sat at home rather than at a desk. Still, it shows that domestic demand spreads evenly across the full week.

In our view, these household numbers resolve the apparent contradiction between platform logs and worker adoption surveys (Bick, Blandin, and Deming 2026, Humlum and Vestergaard 2025). Prompts and workers are different units. A labour survey classifies an employee as a workplace adopter after even occasional use on the job. Platform logs instead count total message or conversation volume. One worker who drafts a weekly memo at the office and asks daily questions about meals, symptoms, and children's homework at home enters a survey as a workplace user. Most of that person's prompts still produce unpaid household services outside gross domestic product (Becker 1965, Cerina, Moro, and Rendall 2021). To separate market work from domestic self-service, we therefore turn to paid plans and commercial software.

TABLE 2. Work and Non-Work Use Across Three Ecosystems, Plans, and Seven Western Economies

Panel A. Global Context Shares, Subscription Plans, and Parity Thresholds

Platform and channel	Coverage	Non-work (%)	Work (%)	$E^\star$ (%)	$\kappa^\star$ ($\hat{\kappa}$)
OpenAI ChatGPT consumer plans (Jul 2024–Sep 2026; binary work-related unless stated)					
Economy median (Jul 2024)	148 (85 maj.)	51.45	48.55	–	–
Economy median (Sep 2026)	149 (149 maj.)	70.40	29.60	–	–
Global 3-way (Sep 2026)	Global	79.0	21.0	85.22	3.76 (–)
Free plan (Sep 2026)	Global	75.3	24.7	–	–
Go plan (Sep 2026)	Global	57.6	42.4	–	–
Plus plan (Sep 2026)	Global	49.2	50.8	–	–
Pro plan (Sep 2026)	Global	31.9	68.1	–	–
Anthropic Claude consumer chat and commercial first-party API (May 2026)					
Consumer median	121 (113 maj.)	59.18	40.82	26.98	1.45 (–)
Consumer global	Global	56.64	43.36	19.51	1.31 (1.49)
Commercial 1P API	Global	15.97	84.03	–	–
Google Gemini ATLAS conversational Search AI Mode and Gemini App (Apr 2026)					
Economy median	218 (218 maj.)	88.66	11.34	113.6	7.82 (–)
Global Search & App	Global	86.5	13.5	107.3	6.41 (–)

Panel B. Seven Western Economies: Cross-Platform Shares and Held-Out Fit

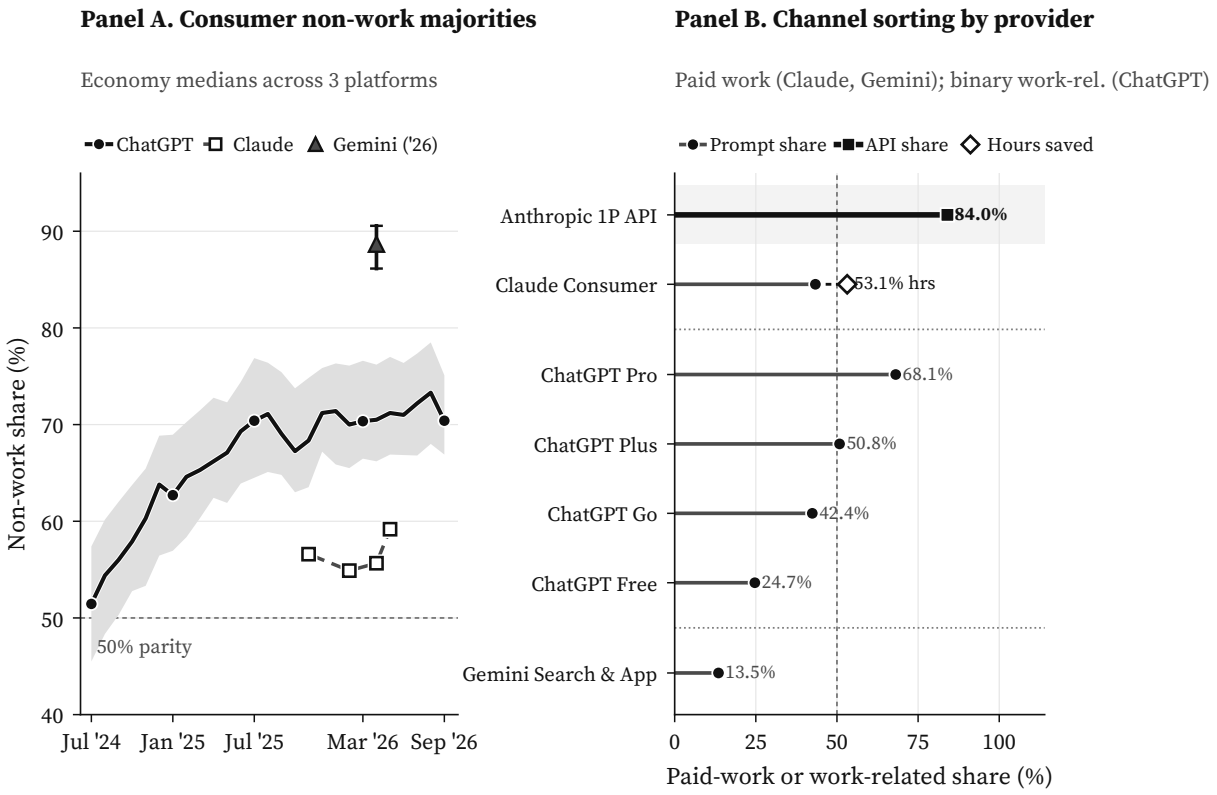
Economy	Claude work	Claude non-work	ChatGPT non-work	Gemini non-work	LOCO R^2
United States	41.32%	58.68%	66.9%	88.95%	96.0%
United Kingdom	39.02%	60.98%	69.1%	89.37%	94.9%
Germany	37.62%	62.38%	73.7%	89.14%	93.4%
France	34.23%	65.77%	72.0%	89.56%	80.6%
Spain	42.28%	57.72%	71.4%	89.38%	92.0%
Italy	42.89%	57.11%	67.5%	89.45%	94.0%
Netherlands	40.82%	59.18%	69.3%	88.27%	96.2%

Panel C. Pre-AI Task Regressions (154 Cells, Country FE, 22 SOC2 Clusters)

Pre-AI task predictor	Slope (pp/SD)	Cluster SE	WCB p	RI-β p	Within R^2
Personal service S_g	−9.94	(3.23)	0.0145	0.0130	0.292
Structured desk R_g	+2.27	(3.34)	–	–	0.037

Notes: Panel A reports OpenAI Signals v3.0 consumer plans (OpenAI 2026; excluding Enterprise and Codex; binary work-related shares include schoolwork whereas the global three-way row separates paid work 21.0%, personal 65.5%, and schoolwork 13.5%), Anthropic May 2026 consumer and commercial 1P API releases (Massenkoff et al. 2026; excluding Claude Code and third-party clouds), and April 2026 Google Gemini ATLAS conversational Search AI Mode and app (Iscenko et al. 2026). Parity thresholds $E^\star$ and $\kappa^\star$ under Equation (2) (transporting $w_E = 0.8403$) apply only to school-free series; $\hat{\kappa}$ is Claude's estimated work-to-non-work duration ratio. Panel B reports May 2026 Claude, September 2026 ChatGPT, and April 2026 Gemini shares alongside zero-parameter six-country leave-one-country-out (LOCO) level fit across 22 SOC2 groups (mean $r = 0.983$, Spearman $\rho = 0.943$, level $R^2 = 92.4\%$, shape $R^2 = 96.0\%$; full folds in Online Appendix Table A1). Panel C reports unweighted country-FE regressions across 154 cells with 22-cluster SEs, wild cluster bootstrap (WCB) p-values (Cameron, Gelbach, and Miller 2008), and randomisation-inference (RI-β) p-values.

FIGURE 1. Consumer Non-Work Majorities Across Three Platforms and Within-Provider Channel Sorting



Notes: Panel A plots cross-economy median binary non-work message shares and interquartile bands across 27 monthly waves of OpenAI Signals consumer plans from July 2024 to September 2026 (OpenAI 2026), cross-economy median non-work shares across four Anthropic consumer waves between November 2025 and May 2026 (hollow squares; Massenkoff et al. 2026), and the April 2026 Google Gemini ATLAS median and interquartile range across 218 economies (filled triangle; Iscenko et al. 2026) relative to 50% parity. Panel B plots channel shares grouped within each provider using each release's work definition across Anthropic paid-work shares (commercial 1P API 84.03% excluding Claude Code and cloud hosts; consumer chat 43.36% of conversations and hollow diamond at 53.15% of model-estimated task-hours in complete-duration cells), September 2026 ChatGPT binary work-related message shares including schoolwork (Free 24.7%, Go 42.4%, Plus 50.8%, and Pro 68.1%, excluding Enterprise and Codex), and April 2026 conversational Gemini paid-work share (13.5%).

Within each provider, paid work climbs sharply in channels that charge subscription fees or require software integration (Figure 1, Panel B; Table 2, Panel A). Because measurement units, work labels, and collection dates differ across companies, we compare channels strictly inside one provider at a time. Across OpenAI's September 2026 consumer plans, work-related message shares rise step by step from 24.7 per cent on Free to 42.4 per cent on Go, 50.8 per cent on Plus, and 68.1 per cent on Pro even before excluded Enterprise seats enter. Under Anthropic's single shared rubric in May 2026, paid work nearly doubles from 43.36 per cent of consumer conversations to 84.03 per cent of commercial first-party API calls (up from 73.89 per cent in November 2025). Iscenko et al. (2026) likewise report that paid work reaches 98.8 per cent on the Gemini API against 13.5 per cent on conversational surfaces. We leave that Gemini API bar out of Figure 1 since Google lacked paid

API prompt text, coded all paid API calls as work by assumption, and excluded Europe. Even inside OpenAI and Anthropic, these channel gaps are descriptive associations rather than causal price elasticities. Monthly fees do screen casual queries out, yet higher earners also purchase personal plans more often and employers reimburse staff subscriptions (Datta, Lambert, and Schneider 2026).

Because no provider publishes commercial API traffic by country, whether paid work forms a majority of total use depends on both unobserved API scale and relative task length. Set consumer volume equal to one with non-work share s_{NW}. Let $E \geq 0$ be the ratio of commercial API calls to consumer conversations, let w_E be the transported global API paid-work share, and let $\kappa > 0$ be the ratio of human hours saved per workplace request to a non-work request. Whenever $\kappa < \kappa^\star = s_{\text{NW}}/(1 - s_{\text{NW}})$ and $\kappa w_E > 1 - w_E$, paid work reaches half of combined duration-weighted use once E reaches

$$E^\star(\kappa, w_E) = \frac{s_{\text{NW}} - \kappa(1 - s_{\text{NW}})}{\kappa w_E - (1 - w_E)}. \tag{2}$$

Online Appendix B derives every boundary case. Under equal task durations ($\kappa = 1$) and Anthropic's global API work share ($w_E = 0.8403$), unobserved API volume must reach 19.51 per cent of consumer volume globally (26.98 per cent in the median economy) for total Claude requests to reach parity. On more household-heavy platforms, the hurdle rises steeply. Equal-duration parity requires commercial API traffic equal to 85.22 per cent of consumer volume under September's three-way ChatGPT split (or 59.95 per cent in the median economy on the school-inclusive binary series) and 107.3 per cent on global conversational Gemini (Table 2, Panel A).

Task duration changes the consumer accounting unit even before commercial API volume enters (Figure 1, Panel B). Across global consumer Claude clusters with complete duration estimates, Anthropic's model puts the average workplace task at 5.02 human hours against 3.37 hours outside work ($\hat{\kappa} = 1.49 > \kappa^\star = 1.31$). Those longer workplace tasks lift paid work from 43.36 per cent of consumer conversations to 53.15 per cent of estimated time saved globally, 50.54 per cent in America, and 47.00 per cent in Britain (Online Appendix Table A1). We read that shift cautiously. Model-estimated durations rest on classifier judgements of text complexity rather than experimental stopwatch measurements (Noy and Zhang 2023), and they do not record whether a user kept the output. Prompt counts and estimated hours therefore capture two distinct margins of AI use and should be evaluated side by side.

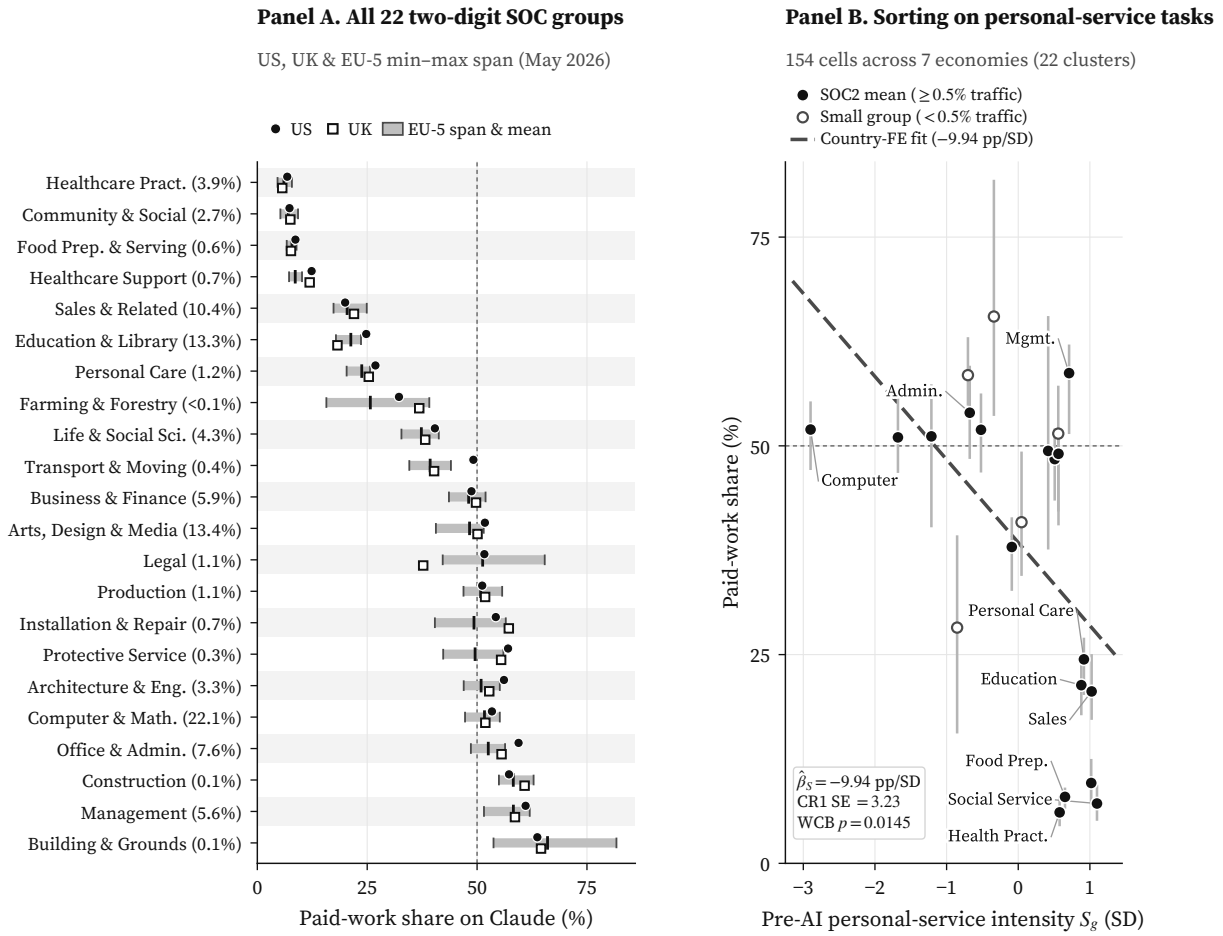
Across our seven Western economies, work shares vary far more across platforms and tasks than across national borders (Figure 2; Table 2, Panel B). Within any one provider, national shares across the United States, the United Kingdom, Germany, France, Spain, Italy, and the Netherlands sit inside a narrow band. Paid work spans just 10.4 to 11.7 per cent on Gemini and 34.23 to 42.89 per cent on consumer Claude while binary work-related messages on ChatGPT stay between 26.3 and 33.1 per cent. Platform gaps are far wider. In those same seven economies, Claude's lowest national

paid-work share (France) sits more than twenty-two percentage points above Gemini's highest (the Netherlands). Task differences inside consumer Claude dwarf national gaps as well (Figure 2, Panel A). Seven-economy mean paid-work shares fall below one tenth in healthcare practitioners (6.1 per cent), community and social service (7.2 per cent), and food preparation (8.0 per cent), and none of those twenty-one country-by-task cells exceeds 9.3 per cent. Office and managerial categories sit roughly six to ten times higher, from 52.0 per cent in computer and mathematical tasks (which draw 22.1 per cent of all consumer traffic) to 54.0 per cent in office administration and 58.7 per cent in management.

Pre-AI task structure explains this sharp occupational divide across all 154 country-by-occupation cells (Figure 2, Panel B; Table 2, Panel C). In country-fixed-effects regressions clustered on the twenty-two occupation groups, a one-standard-deviation higher pre-AI personal-service intensity S_g predicts a 9.94 percentage point lower paid-work share, a drop that remains statistically precise under both wild cluster bootstrap and randomisation-inference tests (Table 2, Panel C). That negative relationship matches both margins of Equation (1). Families substitute free AI guidance for costly or slow market services at home while in-person delivery limits screen use and slows employer rollout. Entered separately, structured desk intensity R_g yields a small slope of +2.27 percentage points that is statistically indistinguishable from zero. That weak slope is itself informative. Students and households also draft documents, build spreadsheets, and write code outside paid employment, so desk-heavy fields still draw large non-work volumes. The measurement cost for labour research is severe. Because every country-by-task cell in healthcare practitioners, social service, and food preparation sits below 9.3 per cent paid work, analysts who treat unfiltered consumer chat as workplace labour inflate task counts more than tenfold across all seven economies in those personal-service fields.

Zero-parameter out-of-sample forecasts confirm that national task profiles align tightly across all seven Western economies (Table 2, Panel B; Online Appendix Table A1 and Figure A1). Anthropic computes each country's twenty-two two-digit work shares directly from national conversations rather than from a scaled global template. Cross-country alignment is therefore a genuine empirical feature of actual demand rather than an artefact of our code. When we predict each held-out economy's twenty-two shares from the simple unweighted mean of the other six, the mean correlation reaches $r = 0.983$ and out-of-sample fit averages 92.4 per cent in levels and 96.0 per cent once country means are removed. France is the sole level outlier (80.6 per cent level fit against 94.7 per cent shape fit) because French paid-work shares sit 6.3 percentage points lower across the board yet keep nearly the same task order. National institutions, language, and workplace norms can shift the overall level of workplace AI use across countries yet preserve which tasks belong at home and which belong at work. Cross-country exposure studies routinely assume that occupational task structure travels across borders (Cazzaniga et al. 2024). Realised consumer task splits confirm that assumption across seven Western economies.

FIGURE 2. Paid-Work Shares Across All Twenty-Two Occupation Groups and Pre-AI Personal-Service Intensity



Notes: Panel A plots May 2026 consumer Claude paid-work shares across all 22 two-digit SOC occupation groups sorted by seven-economy mean (with unweighted seven-economy mean conversation shares in parentheses) for the United States (filled circles), the United Kingdom (hollow squares), and the five Continental European economies (shaded min–max bars and vertical mean ticks; Massenkoff et al. 2026). Panel B plots seven-economy mean paid-work shares (with seven-country min–max vertical whiskers) against standardised pre-AI personal-service intensity S_g alongside the country-fixed-effects slope ($\hat{\beta}_S = -9.94$ percentage points per SD, $N = 154$, wild cluster bootstrap $p = 0.0145$).

4. Workplace AI Across Occupations and Pay

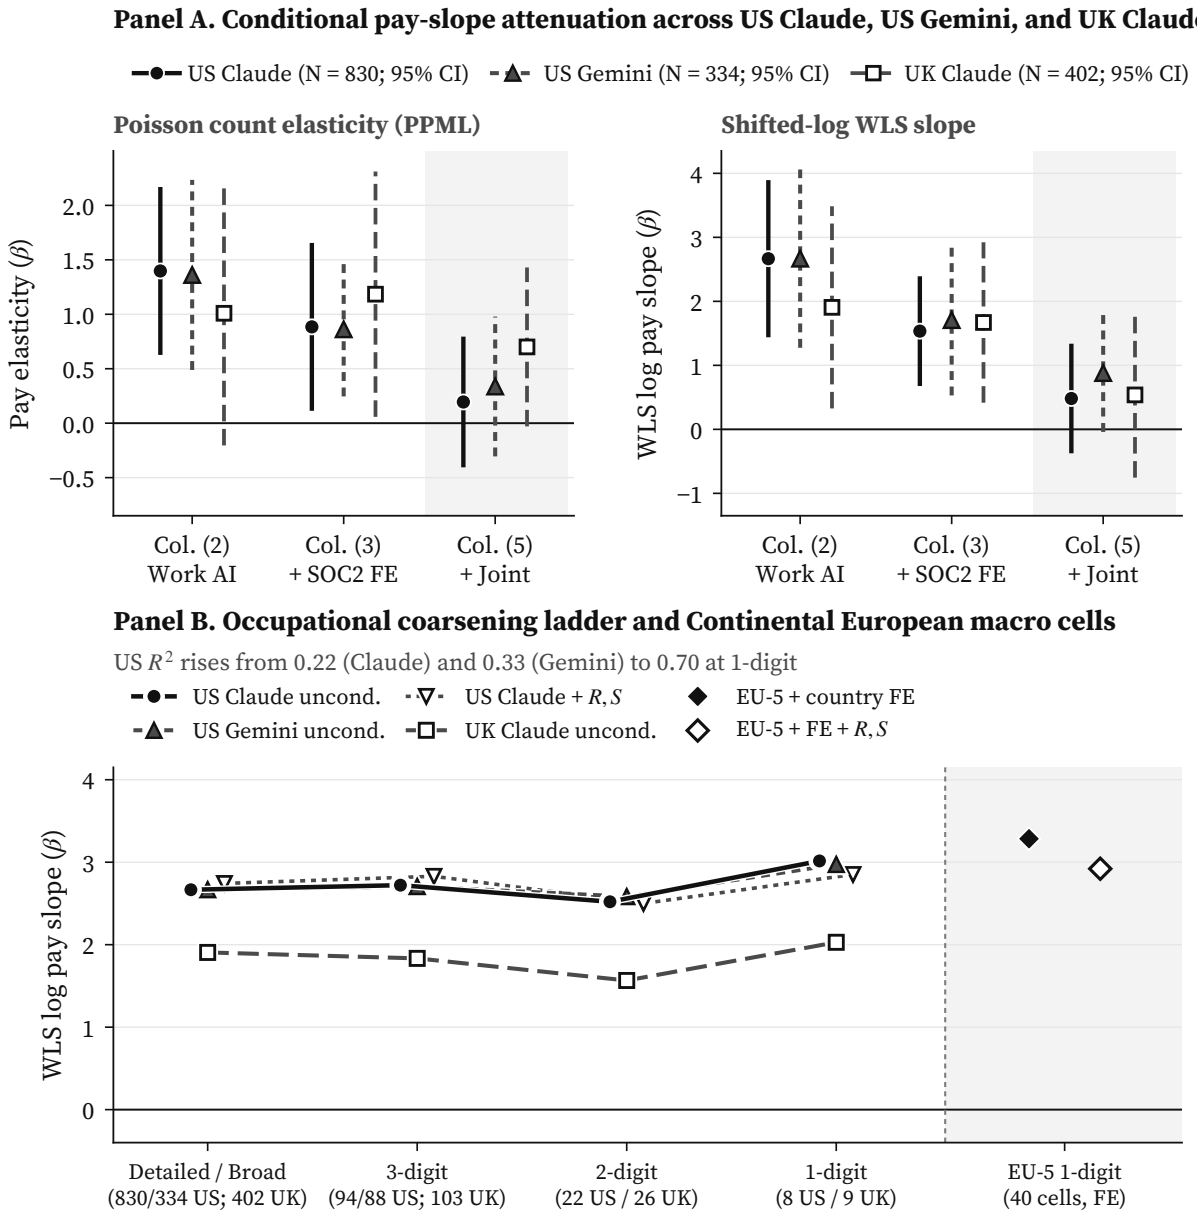
Once household self-service is removed from occupational task counts, workplace AI intensity rises steeply with median hourly pay across both the United States and the United Kingdom (Figure 3, Panel A; Table 3, Columns 1 and 2). Unfiltered traffic assigns domestic questions about health, meals, and homework onto service and care jobs even though no professional is on the clock. Once we strip out non-work conversations, lower-paid food and healthcare-support groups shrink far more than desk professions, and the raw pay gradient steepens sharply. Employment-weighted Poisson elasticities of workplace intensity $I_i = s_i/e_i$ on log median hourly pay climb from $+1.19$ to

+1.40 across 830 American six-digit occupations and from $+0.71$ to +1.01 across 402 British four-digit unit groups, while shifted-log slopes reach +2.67 in America and +1.91 in Britain. Worst-case linear-programme bounds stay strictly positive in both countries (Table 3). A natural worry is that our global-to-national allocation step created this steep gradient mechanically. Google's ATLAS release rules out that concern. Across 334 directly reported US Broad occupations that employ 143.62 million workers, Gemini produces an almost identical raw elasticity of +1.36 in Poisson and +2.67 in shifted-log WLS and closely reproduces Iscenko et al. (2026)'s published $+2.68$ annual-pay slope. When we aggregate US Claude onto those exact 334 Broad codes, the two independent platforms align closely in both raw slopes and cross-occupation ranks ($r = 0.76$ across active occupations and Spearman $\rho = 0.91$ across job families; Online Appendix Table A5 and Figure A3).

Most of that steep pay gradient reflects differences across broad job families and pre-AI task capabilities rather than pay gaps within one occupation group (Figure 3, Panel A; Table 3, Columns 3 to 5). Just as earlier waves of office computerisation concentrated keyboard work in clerical and analytical fields rather than manual trades (Autor, Levy, and Murnane 2003, Goos, Manning, and Salomons 2014), broad occupational structure explains the bulk of workplace AI variation. Two-digit job-family fixed effects alone lift shifted-log explanatory power from 0.222 to 0.525 on US Claude, from 0.329 to 0.618 on US Gemini, and from 0.108 to 0.552 on UK Claude. Those job-family indicators account for 70 to 89 per cent of explained variance under a Shapley decomposition and push American and British linear-programme level bounds across zero (Table 3, Column 3). Once we add pre-ChatGPT task descriptors (R_i, S_i) alongside the capability scores of Felten, Raj, and Seamans (2021) and Eloundou et al. (2024) in Column (5), the conditional pay slope falls by two thirds to more than four fifths in five of our six American and British specifications. On US Claude, the Poisson elasticity drops to +0.19 and the shifted-log slope drops to +0.48. On US Gemini they fall to +0.34 and +0.88, and in Britain shifted-log WLS drops to +0.54. None of these five estimates is statistically distinguishable from zero once standard errors are clustered by job family (Table 3; standalone capability checks appear in Online Appendix Table A2).

Two specific pockets still retain a modest positive slope, and both deserve careful evaluation. On Google Gemini, attenuation holds in every one of twenty-two leave-one-family-out tests and across all 454 OEWS Broad codes when unpublished titles are set to zero (Online Appendix Table A5 and Figure A3). Yet inside the 148 multi-title Broad codes matched directly to official OEWS Broad wage rows, a modest positive conditional slope survives ($+0.79$ in Poisson and $+1.02$ in WLS). In Britain, full controls cut the Poisson elasticity by nearly a third yet leave a borderline positive estimate of +0.70 (Table 3, Panel C). In our view, two forces sit behind this pattern. Higher hourly wages raise the value of saved time v and help workers or employers justify paid plans, which can sustain a mild positive slope inside selected groups. At the same time, pre-2023 capability scores correlate closely with education and wages yet omit newer model skills. Column (5) therefore

FIGURE 3. Occupational Pay-Slope Attenuation Across Two Providers and the Coarsening Ladder



Notes: Panel A plots employment-weighted PPML count elasticities (left) and shifted-log WLS slopes (right) with 95% two-digit SOC2 cluster confidence intervals (t_{21} in the US, t_{25} in the UK) across Column (2) (Work AI), Column (3) (+ SOC2 FE), and Column (5) (+ Joint) of Table 3 for 830 US six-digit Claude occupations (filled circles), 334 US Broad Gemini occupations (filled triangles), and 402 paid UK four-digit Claude unit groups ($N = 398$ in Column 5, hollow squares). Panel B traces WLS log/shifted-log pay slopes as detailed or Broad occupations are aggregated through three-digit, two-digit, and single-digit ISCO-08/SOC1 groups across US Claude, US Gemini, and UK Claude alongside 40 within-country EU-5 ISCO-08 cells ($+3.28$ unconditional filled diamond; $+2.92$ conditional hollow diamond).

shows that hourly pay adds limited predictive power once job family and pre-AI task feasibility are known, and does not prove that wages play zero causal role in adoption.

TABLE 3. Occupational Pay Gradients Across US Claude, US Gemini, UK Claude, and European Cells

	Unconditional		Job-family FE		Full
Estimator	(1) All AI	(2) Work AI	(3) SOC2 FE	(4) $+R, S$	(5) Joint
Panel A. United States — Anthropic Claude *(830 6-digit SOCs, 155.48M employees, 22 clusters)*					
PPML elasticity	+1.185	+1.398	+0.884	+1.073	+0.195
Cluster SE	(0.344)	(0.371)	(0.371)	(0.327)	(0.289)
Shifted-log WLS	+2.044	+2.666	+1.534	+1.555	+0.482
Cluster SE	(0.509)	(0.590)	(0.412)	(0.385)	(0.412)
Level LP $[L, U]$	–	[0.80, 2.60]	[−2.12, 3.69]	[−2.16, 4.16]	[−3.58, 3.68]
Shifted-log R^2	0.132	0.222	0.525	0.532	0.572
Panel B. United States — Google Gemini ATLAS *(334 direct Broad SOCs, 143.62M employees, 22 clusters)*					
PPML elasticity	–	+1.361	+0.862	+1.341	+0.337
Cluster SE	–	(0.420)	(0.296)	(0.425)	(0.309)
Shifted-log WLS	–	+2.667	+1.707	+1.839	+0.880
Cluster SE	–	(0.670)	(0.566)	(0.519)	(0.443)
Shifted-log R^2	–	0.329	0.618	0.622	0.656
Panel C. United Kingdom — Anthropic Claude *(402 4-digit SOC2020 codes, 33.07M APS workers, 26 clusters)*					
PPML elasticity	+0.714	+1.009	+1.185	+0.880	+0.700
Cluster SE	(0.507)	(0.589)	(0.547)	(0.422)	(0.354)
Shifted-log WLS	+1.360	+1.906	+1.670	+1.275	+0.538
Cluster SE	(0.664)	(0.767)	(0.609)	(0.582)	(0.628)
Level LP $[L, U]$	–	[0.15, 1.95]	[−0.11, 2.37]	[−0.69, 2.11]	[−0.85, 2.14]
Shifted-log R^2	0.076	0.108	0.552	0.568	0.606

Panel D. Single-Digit Aggregation Across US, UK, and Continental European Cells

Single-digit sample	Cells (N)	Uncond. log	$+R, S$ (no FE)	Uncond. PPML	Uncond. R^2
US Claude (1-digit ISCO-08)	8	+3.02	+2.85	+2.33	0.70
US Gemini (1-digit ISCO-08)	8	+2.97	+3.04	+2.00	0.70
UK Claude (1-digit SOC1)	9	+2.03	+1.65	+1.07	0.39
EU-5 Eurostat (1-digit ISCO-08, FE)	40	+3.28	+2.92	+2.09	0.72

Notes: Employment-weighted regressions relate workplace AI intensity $I_i = s_i/e_i$ to log median hourly pay (log mean pay in EU-5 cells) across 830 US 6-digit SOCs on May 2026 Claude (155.48M employees, 393 zeros; Panel A), 334 US-reported Broad SOCs on April 2026 Gemini ATLAS (143.62M employees, 17 zeros; Panel B, where Column 1 is blank since ATLAS publishes workplace shares directly), 402 paid UK 4-digit SOC2020 unit groups weighted by 2025 APS employment (29 zeros; $N = 398$ in Column 5; Panel C), and 40 within-country EU-5 ISCO-08 cells (123.91M LFS workers, 8 clusters; Panel D). Column (5) adds pre-AI capability scores (Felten, Raj, and Seamans 2021, Eloundou et al. 2024) to SOC2 fixed effects and R_i, S_i; standalone capability-only models appear in Online Appendix Table A2. Level LP $[L, U]$ bounds the level WLS slope over Ω via Online Appendix Equation (8). Standard errors cluster on two-digit job families (ISCO-08 in Panel D).

A two-part split between active presence and intensity among active occupations clarifies why conditional pay gradients flatten inside white-collar fields (Online Appendix Tables A2 and A5). Full

controls cut the active log-pay slope by three quarters on US Gemini, where 317 of 334 Broad codes record positive workplace use, from $+2.77$ to an imprecise +0.70. In Britain, where 373 of 402 unit groups are active, full controls halve the active slope from +1.66 to an imprecise +0.88. US Claude shows a slightly negative extensive-margin slope across its 393 zero-coded or suppressed titles alongside a 45 per cent drop among active titles (from +2.27 to +1.24). When we impute small positive shares to censored codes, the joint shifted-log slope remains statistically indistinguishable from zero. Once flat-rate access is paid, extra prompts carry zero marginal monetary price. Knowledge-hierarchy and social-skill models (Garicano 2000, Deming 2017, Autor and Thompson 2025) predict that workplace prompt volume then plateaus near the top of white-collar hierarchies. Mid-paid developers, analysts, and paralegals spend much of their workday on codified text and code. Senior partners, executives, and clinical directors spend more hours on face-to-face persuasion, supervision, and legal sign-off that seldom enter a prompt window.

Coarse Continental European statistics cannot test this within-family pattern because single-digit occupational classes pool white-collar professionals with manual workers (Figure 3, Panel B; Table 3, Panel D). Across 40 single-digit ISCO-08 cells that cover 123.91 million workers in Germany, France, Spain, Italy, and the Netherlands, log pay slopes sit near $+3$ both unconditionally (+3.28) and with R_i and S_i included (+2.92). That steeper European macro slope is not proof of sharper pay selection inside European jobs. When we aggregate America's 830 Claude codes or 334 Gemini Broad codes up to those exact eight single-digit ISCO-08 groups, explanatory power jumps to 0.70 and single-digit pay slopes match Europe closely at +3.02 and $+2.97$ (Table 3, Panel D). Britain's 402 unit groups show the same ecological lift at the nine-group single-digit rung, where R^2 jumps from 0.108 to 0.39 at a flatter slope ($+2.03$). Once narrow job codes are collapsed into single-digit classes, no within-family variation remains to condition on. Steep Continental macro slopes therefore reflect ecological aggregation across broad manual and white-collar groups (Cross and Manski 2002) alongside any European wage compression.

Even fine occupational regressions assign every employee in a job title its median hourly pay and hide wide wage gaps among coworkers (Autor and Handel 2013, Atalay et al. 2020). Eurostat publishes no internal occupational wage percentiles across our European cells. We therefore build continuous within-occupation pay distributions across 184.33 million payroll jobs from five-knot US OEWS grids (155.48 million employees) and eleven-knot UK ASHE grids (28.85 million matched jobs across 382 unit groups; Table 4). Internal overlap across national wage cutoffs is pervasive. Slightly over half of employees (50.8 per cent in America and 52.2 per cent in Britain) work in occupations whose interior pay spans the bottom-fifth cutoff of $17.49 or £13.89 per hour. Nearly every employee (98.0 and 99.7 per cent) works in an occupation that straddles at least one quintile cutoff. Customer-service representatives, general office clerks, administrative assistants, and bookkeepers all earn above the bottom fifth at their occupation medians yet employ large numbers of staff below it. Let ϕ_{iq} denote the share of occupation i's employees paid in national fifth q at cell mean hourly

wage $\bar{w}_{iq}$. Under equal within-occupation use, allocated workplace traffic shares H_q and baseline hourly-wage-weighted shares B_q satisfy

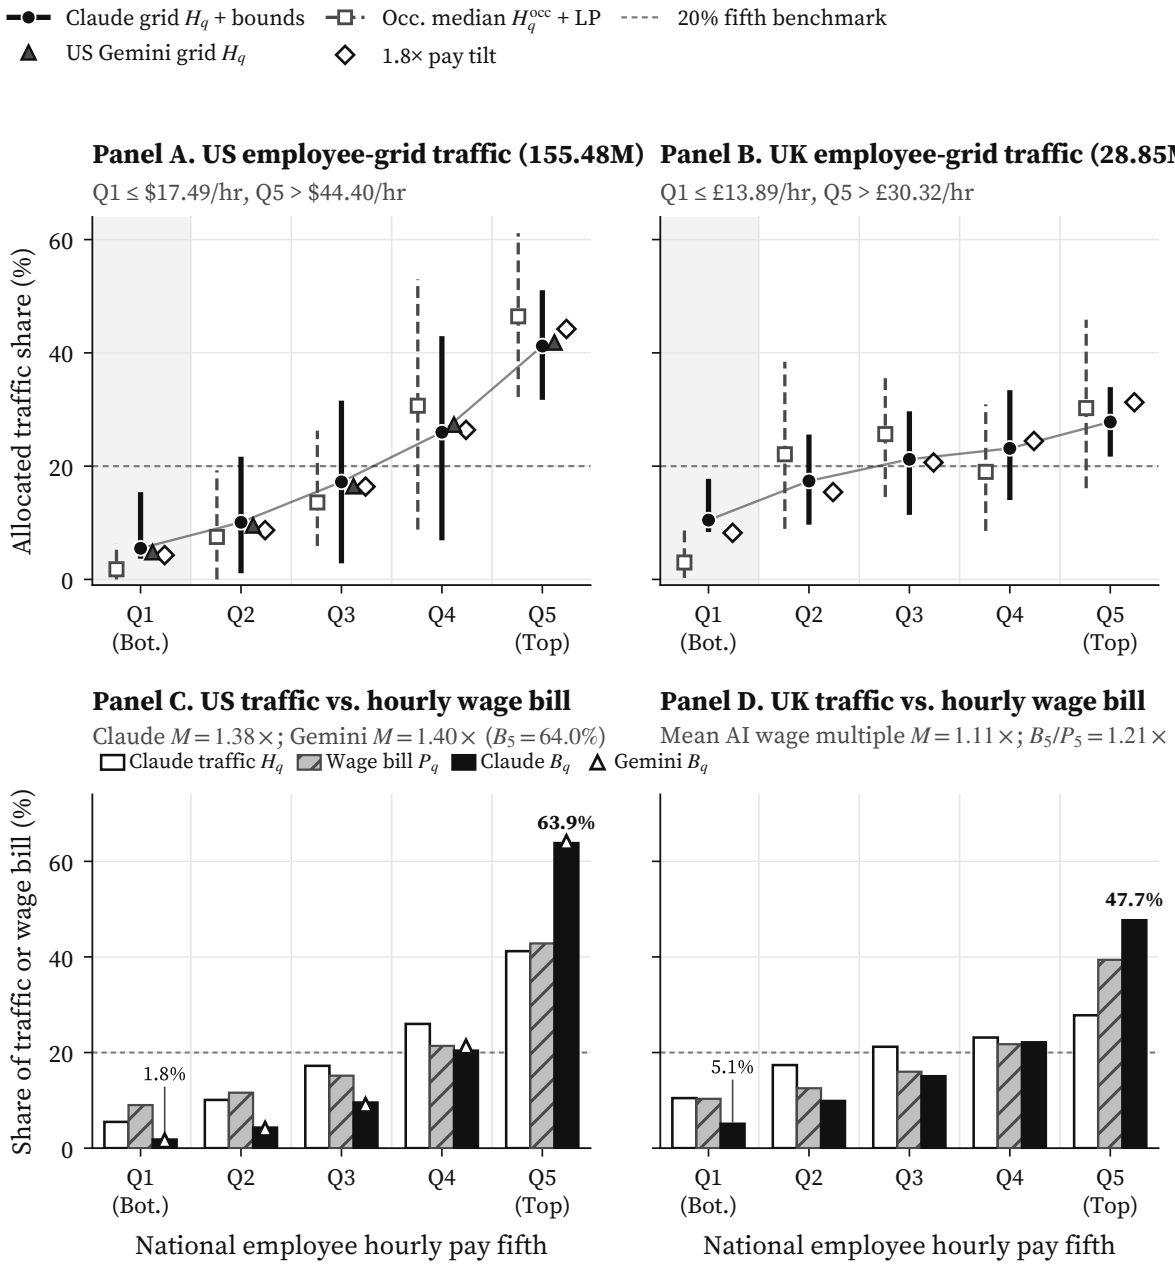
$$H_q(\mathbf{m}^W, \boldsymbol{\phi}) = \frac{\sum_i m_i^W \phi_{iq}}{\sum_i m_i^W}, \qquad B_q(\mathbf{m}^W, \boldsymbol{\phi}) = \frac{\sum_i m_i^W \phi_{iq} \bar{w}_{iq}}{\sum_{r=1}^{5} \sum_i m_i^W \phi_{ir} \bar{w}_{ir}}. \tag{3}$$

Internal pay grids lift bottom-fifth allocated shares sharply across both countries and both platforms (Figure 4, Panels A and B; Table 4). Relative to occupation-median assignment, employee-grid quintiles raise the bottom fifth's share more than threefold on US Claude (from 1.79 to 5.49 per cent) and UK Claude (from 2.99 to 10.47 per cent). On US Gemini, the bottom fifth's share rises more than fivefold from 0.89 to 4.82 per cent. Because public files observe neither exact within-family task shares nor coworker prompt ranks, we test how resilient this gain is along both margins. Across all admissible wage-grid interpolations, bottom-fifth gains stay strictly positive in every series, and eligible detailed-title Gemini bounds stay strictly positive as well (Online Appendix Table A5). When we relax occupational allocation and wage interpolation jointly, the British paired gain stays strictly positive ($[+0.16, +18.69]$ percentage points) while US Claude joint bounds span zero under adversarial within-family assignment (Table 4). A realistic 1.8-fold tilt towards higher earners inside identical job codes preserves most of the lift (4.29 per cent in America and 8.22 per cent in Britain). To erase the bottom-fifth gain completely, top-fifth earners inside the exact same occupation would have to prompt 7.66 times as heavily as bottom-fifth coworkers on US Claude, 8.95 times on UK Claude, and 14.21 times on US Gemini. Existing adoption surveys and bank subscription records find far milder within-group earnings gradients (Humlum and Vestergaard 2025, Datta, Lambert, and Schneider 2026), while workplace trials report the largest gains among junior support agents and lower-ranked writers (Noy and Zhang 2023, Brynjolfsson, Li, and Raymond 2025).

Baseline hourly wage weights concentrate cost-weighted exposure heavily in the top fifth (Figure 4, Panels C and D; Table 4). Wage weights shift the accounting unit from task volume to hourly labour cost (Acemoglu 2025, Jones 2026). In the United States, the highest-paid fifth holds nearly two thirds of hourly-wage-weighted traffic (63.85 per cent on Claude and 64.04 per cent on Gemini against 42.83 per cent of the baseline wage bill, almost 1.5 times its payroll share). The bottom fifth holds under two per cent (1.83 and 1.61 per cent against 9.01 per cent of baseline wages). Two independently measured platforms thus produce virtually identical American AI mean-wage multiples ($M = 1.38$ on Claude and 1.40 on Gemini) because steep US hourly pay dispersion dominates both allocations. In Britain, top-fifth wage-weighted concentration is noticeably milder at 47.71 per cent against 39.40 per cent of baseline wages (1.21 times baseline, $M = 1.11$). Britain's smaller multiple reflects both its flatter pay structure and its broader clerical AI reach.

Allocated traffic shares and wage-weighted shares therefore describe one workplace task distribution from two distinct accounting angles. Each employee fifth contains one fifth of matched payroll jobs while hourly pay rises steeply from bottom to top. Under equal within-occupation use,

FIGURE 4. Allocated Workplace AI Traffic and Hourly-Wage-Weighted Incidence Across Employee Pay Fifths



Notes: Top panels plot allocated workplace AI traffic shares across national employee hourly pay fifths (Q1–Q5) in US May 2025 OEWS (155.48M employees; Panel A) and UK 2026 ASHE (28.85M matched employee jobs; Panel B) under Claude occupation-median assignment H_q^{occ} (hollow squares with dashed LP intervals), Claude employee-grid traffic H_q under equal within-occupation prompting (filled circles with solid quantile-grid intervals), a $1.8\times$ linear within-occupation pay tilt (hollow diamonds), and eligible-title US Gemini employee-grid traffic H_q (filled triangles in Panel A) relative to the 20% fifth benchmark. Bottom panels compare Claude allocated traffic shares H_q, baseline hourly-wage-bill shares P_q, and hourly-wage-weighted traffic shares B_q (with hollow triangles marking US Gemini B_q in Panel C) across Q1–Q5 in the United States (Panel C) and the United Kingdom (Panel D).

TABLE 4. Allocated Workplace AI Traffic and Hourly-Wage-Weighted Concentration Across US and UK Pay Fifths

Measure	National employee hourly pay fifth					Summary
	Q1 (Bot.)	Q2	Q3	Q4	Q5 (Top)	
Panel A. United States OEWS *(155.48M employees; Q1 ≤ $17.49, Q5 > $44.40/hr)*						
1. Occupation-median benchmark						
Claude median H_q^{occ} (%)	1.79	7.50	13.59	30.67	46.45	Q5/Q1 = 25.9×
Claude occ. LP $[L, U]$	[0.00, 6.30]	[0.00, 19.26]	[5.89, 26.26]	[8.80, 52.97]	[32.21, 61.14]	LP Q5/Q1 ≥ 5.11
Gemini median H_q^{occ} (%)	0.89	7.75	14.53	32.87	43.97	Q5/Q1 = 49.4×
2. Continuous within-occupation employee pay grid						
Claude grid H_q (%)	5.49	10.09	17.22	25.99	41.21	$\rho^\star = 7.66\times$
Claude grid $[L, U]$	[3.66, 15.41]	[1.09, 21.66]	[2.84, 31.56]	[6.91, 42.95]	[31.71, 51.06]	$\Delta H_1^{\text{grid}} \geq +1.87$
Claude joint LP $(\Omega \times \mathcal{G})$	[0.68, 21.93]	[0.01, 29.46]	[0.56, 39.59]	[0.73, 51.48]	[21.96, 61.77]	$\Delta H_1^{\text{joint}}$: [−2.23, +19.90]
Claude tilt 1.8× (%)	4.29	8.70	16.39	26.40	44.22	–
Gemini grid H_q (%)	4.82	9.54	16.46	27.34	41.84	$\rho^\star = 14.21\times$
3. Baseline hourly-wage bill vs. wage-weighted traffic						
Wage bill P_q (%)	9.01	11.60	15.16	21.39	42.83	$\sum P_q = 100\%$
Claude wage-wt. B_q (%)	1.83	4.33	9.57	20.42	63.85	$M = 1.38\times$
Gemini wage-wt. B_q (%)	1.61	4.04	9.06	21.25	64.04	$M = 1.40\times$
Panel B. United Kingdom ASHE 2026 *(28.85M matched jobs; Q1 ≤ £13.89, Q5 > £30.32/hr)*						
1. Occupation-median benchmark						
Claude median H_q^{occ} (%)	2.99	22.10	25.66	18.99	30.25	Q5/Q1 = 10.1×
Claude occ. LP $[L, U]$	[0.25, 9.11]	[8.92, 38.43]	[14.56, 36.54]	[8.56, 30.91]	[16.10, 45.86]	LP Q5/Q1: [2.00, 169.3]
2. Continuous within-occupation employee pay grid						
Claude grid H_q (%)	10.47	17.38	21.21	23.14	27.80	$\rho^\star = 8.95\times$
Claude grid $[L, U]$	[8.37, 17.76]	[9.67, 25.58]	[11.40, 29.68]	[14.02, 33.41]	[21.68, 33.95]	$\Delta H_1^{\text{grid}} \geq +5.38$
Claude joint LP $(\Omega \times \mathcal{G})$	[4.38, 22.78]	[5.41, 31.15]	[8.80, 32.96]	[9.96, 38.04]	[14.57, 40.27]	$\Delta H_1^{\text{joint}}$: [+0.16, +18.69]
Claude tilt 1.8× (%)	8.22	15.42	20.65	24.44	31.27	–
3. Baseline hourly-wage bill vs. wage-weighted traffic						
Wage bill P_q (%)	10.32	12.53	15.99	21.75	39.40	$\sum P_q = 100\%$
Claude wage-wt. B_q (%)	5.13	9.90	15.08	22.17	47.71	$M = 1.11\times$

Notes: Columns report task-allocated workplace AI shares across fixed national employee hourly pay fifths Q1 (lowest-paid) through Q5 (highest-paid) in US May 2025 OEWS (155.48M employees across 830 codes; Panel A) and UK 2026 ASHE (28.85M matched jobs across 382 unit groups; Panel B). Employee-grid cutoffs are $17.49, $22.13, $29.57, and $44.40/hr in the US (£13.89, £16.67, £21.75, and £30.32/hr in the UK); occupation-median cutoffs are $17.86, $22.00, $29.70, and $40.23/hr in the US (£14.34, £16.90, £22.55, and £28.90/hr in the UK). Rows report occupation-median shares H_q^{occ} with LP bounds, Equation (3) equal-use grid shares H_q with quantile-grid bounds $[L, U]$ and joint LP bounds over $\Omega \times \mathcal{G}$, 1.8× linear pay tilts, eligible-title US Gemini comparisons (full bounds in Online Appendix Table A5), break-even Q5/Q1 tilts $\rho^\star$, baseline hourly-wage-bill shares P_q, wage-weighted shares B_q, and AI mean-wage multiples M.

H_q captures how broadly AI tasks reach lower-paid clerical bands and makes no claim about observed worker-level use. By contrast, B_q translates those same allocated tasks into baseline hourly labour-cost shares for macroeconomic work and does not measure realised firm cost savings or post-adoption wage changes.

5. Human Iteration in Chat and Direct Instructions on APIs

Even once we map workplace AI onto occupations, an open question remains about how humans and firm software instruct a model. Field experiments report substantial productivity gains when writers revise drafts in chat (Noy and Zhang 2023), support agents read live reply suggestions (Brynjolfsson, Li, and Raymond 2025), and developers work with code assistants (Cui et al. 2026). Anthropic applies one shared five-mode rubric across consumer conversations (Claude.ai chat and Cowork) and commercial first-party API calls (Appel et al. 2026, Massenkoff et al. 2026). Pure directive requests hand over task completion with minimal human exchange while collaborative modes span multi-turn feedback loops, task iteration, study, and validation. Because public files report each cluster's workplace share $s_g \in (0, 1]$ separately from its observed target-mode share d_g^{obs} and unclassified share $u_g \geq 0$, and withhold small clusters with workplace share α_{rem}, probability conservation $d_g = s_g\theta_g + (1 - s_g)b_g$ (Fréchet 1951, Manski 2003) bounds the true cluster workplace rate θ_g and full-population share Θ_{full} by

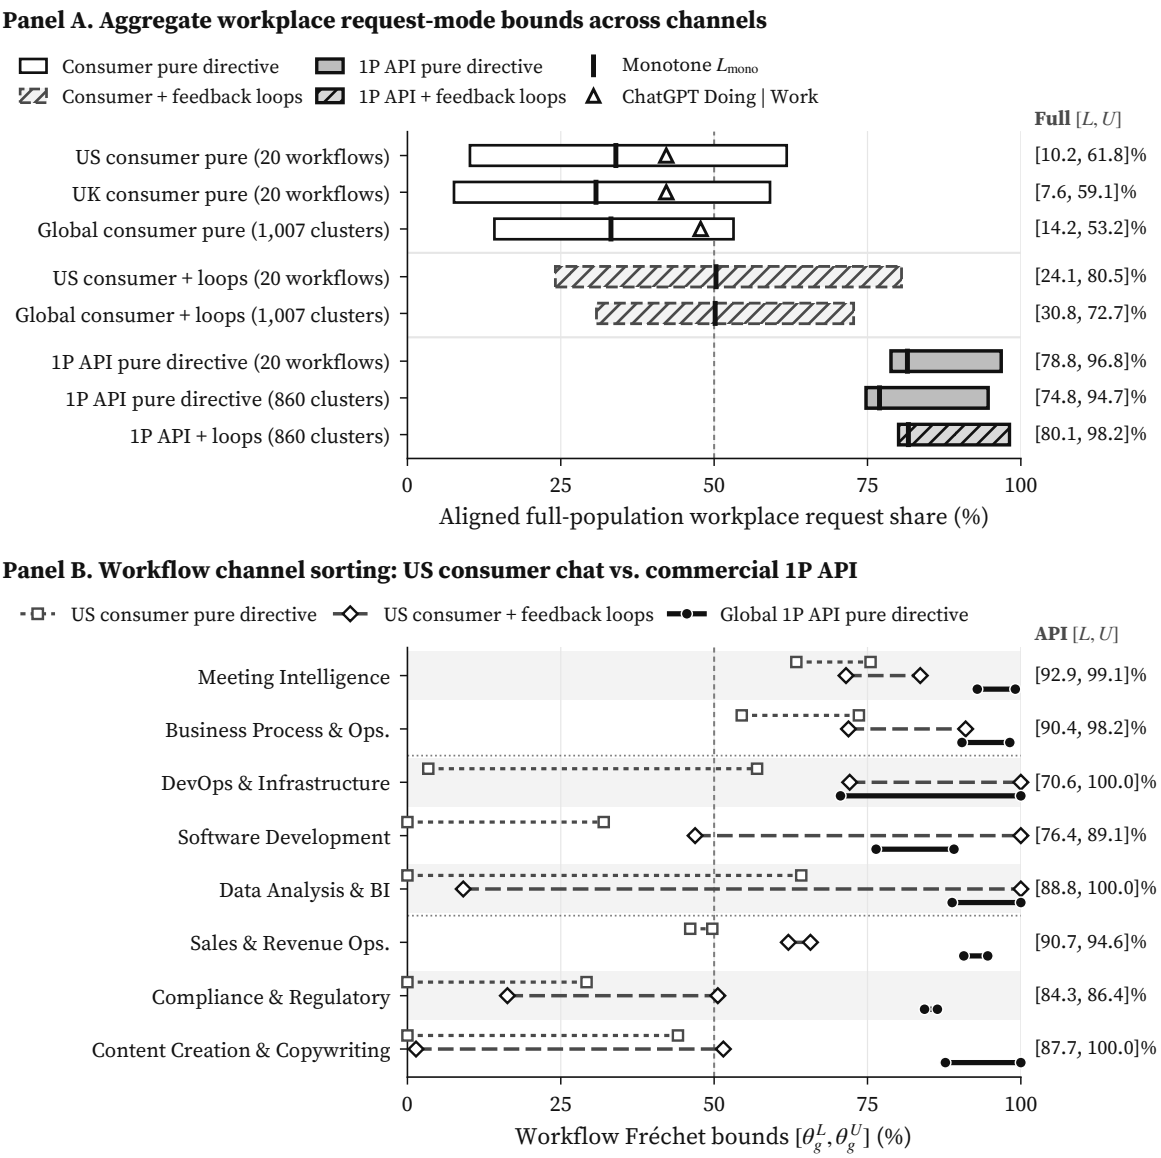
$$\theta_g^L = \frac{\max(0,\, d_g^{\text{obs}} + s_g - 1)}{s_g}, \qquad \theta_g^U = \frac{\min(d_g^{\text{obs}} + u_g,\, s_g)}{s_g}, \tag{4}$$
$$\Theta_{\text{full}} \in \left[(1 - \alpha_{\text{rem}})\Theta_{\text{ret}}^L,\ (1 - \alpha_{\text{rem}})\Theta_{\text{ret}}^U + \alpha_{\text{rem}}\right],$$

where $\Theta_{\text{ret}}^{L,U} = \sum_{g \in \mathcal{R}} \tilde{\alpha}_g \theta_g^{L,U}$ weights retained clusters by workplace volume $\tilde{\alpha}_g$. Online Appendix B shows that finer subclusters cannot widen retained bounds and defines the monotone-selection floor L_{mono} alongside the parity-breakdown index $\psi_g^\star$ (Cross and Manski 2002, Masten and Poirier 2020).

In aggregate, commercial first-party APIs receive direct instructions far more often than consumer chat windows, and their bounds never overlap (Figure 5, Panel A; Table 5, Panel A). Across 1,007 global consumer task clusters (98.2 per cent of May 2026 workplace conversations), full-population pure directive bounds span $[14.2, 53.2]$ per cent, with almost six tenths of retained consumer work strictly below half. On the 860 commercial API clusters, full-population pure directive bounds sit more than twenty-one percentage points higher at $[74.8, 94.7]$ per cent, with 95.8 per cent of retained API work strictly above half. Narrower API brackets do not create this separation. Every Fréchet interval contains its true rate regardless of width, which confirms that disjoint intervals establish a strict channel gap even under worst-case assignments for suppressed and unclassified prompts. The same separation holds in the April 2026 full-month wave and across

every one of our seven Western economies (Table 5, Panel A). Even when multi-turn feedback loops are added to the target mode across global consumer clusters, the global consumer ceiling (72.7 per cent) still stops below the commercial API pure directive floor (74.8 per cent).

FIGURE 5. Channel and Workflow Bounds on Workplace Request Modes



Notes: Panel A plots May 2026 full-population identification intervals $[\Theta_{\text{full}}^L, \Theta_{\text{full}}^U]$ from Equation (4) for pure directive consumer conversations (open white bars), consumer pure directive plus feedback-loop requests (pale hatched bars), commercial 1P API pure directive calls (solid grey bars), and API calls adding feedback loops (dark hatched bar). Vertical ticks mark monotone-selection floors L_{mono} under $\theta_g \geq b_g$; hollow triangles mark September 2026 ChatGPT *Doing*|*Work* message benchmarks. Panel B plots workflow intervals $[\theta_g^L, \theta_g^U]$ across the eight focal workflows of Table 5, Panel B, grouped into operational, technical, and commercial-regulatory workflows.

Workflow comparisons reveal that this aggregate channel gap divides into three distinct economic patterns (Figure 5, Panel B; Table 5, Panel B; Online Appendix Table A3). In all ten workflows

where paid work forms a majority of American consumer traffic, the consumer pure directive ceiling stays below the commercial API floor. Only meeting intelligence ($[63.4, 75.5]$ per cent in America) and business-process operations ($[54.5, 73.6]$ per cent) exceed half pure directive inside American and British chat without extra assumptions ($\psi_g^\star = 0.00$) before both climb past nine tenths on the API. Technical workflows show why feedback loops matter. Pure one-shot delegation is low in consumer chat for software development (at most 32.0 per cent), DevOps, and data analysis. Once multi-turn feedback loops enter, all three consumer intervals overlap their commercial API intervals and make their channel order unsigned. In practice, engineers paste back error logs, test results, and intermediate tables across turns rather than rely on a single prompt. Commercial and regulatory workflows stay sharply split. Sales operations, copywriting, and regulatory compliance sit strictly below half pure directive in American and British chat ($\psi_g^\star = \infty$) and stay at least twenty-five percentage points below commercial API floors (84.3 to 90.7 per cent) even after every feedback-loop turn is counted (Table 5, Panel B).

Three mechanisms can explain why commercial APIs look more directive, and our workflow contrasts show that a pure counting artefact cannot explain the full pattern. Under the fixed-cost framework in Section 2, API integration requires upfront software work, prompt templates, and maintenance ($F_{\text{API}} > 0$). Standardised bulk jobs such as transcript extraction, catalogue copy, and compliance checks spread those setup costs far more easily than bespoke client persuasion or open-ended policy judgement (Handa et al. 2025, Brynjolfsson, Li, and Raymond 2025, Autor and Thompson 2025). A second explanation is mechanical. Consumer logs bundle a multi-turn thread into one conversation whereas stateless API endpoints score each HTTP request on its own. That unit difference pushes API directive shares upward whenever software pipelines split a larger job into standalone calls, though chat wrappers that resend full conversation history on every turn pull in the opposite direction. A third explanation is selection across firms and sub-tasks inside each broad workflow label. Public margins cannot isolate a single structural parameter. Yet a uniform call-counting bias across workflows would shift all tasks in tandem and cannot explain why feedback loops move software development and DevOps bounds into overlap while compliance, copywriting, and sales stay at least twenty-five percentage points apart.

Independent rubrics from OpenAI and Google show consistent task and cross-country regularities inside each ecosystem, although they do not establish a single cross-provider delegation rank (Table 5, Panel A; Online Appendix Tables A4 and A5 and Figure A2). Anthropic scores whether an exchange runs in one turn or interactively. OpenAI tracks the share of work-related messages that ask for a concrete output (*Doing*) rather than information or expression. Google estimates the share of prompts on conversational and free-tier surfaces that ask for complete task automation across seventeen US Major SOC groups. Because these three rubrics ask different questions on different surfaces, raw percentages cannot be pooled across providers. Within each release, however, shared O*NET Generalized Work Activities in our seven Western economies align closely

TABLE 5. Coverage, Channel Bounds, and Workflow Sorting for Workplace Request Modes

Panel A. Full-Population Workplace Request Bounds Across Channels and Economies

		Anthropic bounds (%)		OpenAI	Retained work in cells	
Channel, economy, or partition	$1-\alpha_{\text{rem}}$	Full-pop. $[L,U]$	L_{mono}	Doing	$\theta_g^L > 50\%$	$\theta_g^U < 50\%$
Consumer channel: pure directive (20 workflows, 7 Western economies)						
United States (USA)	98.6%	[10.2, 61.8]	34.0%	42.2%	14.0%	52.2%
United Kingdom (GBR)	98.5%	[7.6, 59.1]	30.7%	42.2%	11.3%	59.7%
Germany (DEU)	98.6%	[5.7, 61.1]	29.2%	41.6%	0.4%	60.8%
France (FRA)	98.8%	[5.2, 62.9]	29.9%	40.5%	0.5%	26.5%
Spain (ESP)	98.7%	[6.2, 57.7]	30.2%	42.2%	0.4%	61.4%
Italy (ITA)	98.3%	[5.0, 56.5]	29.1%	42.0%	0.4%	68.3%
Netherlands (NLD)	98.6%	[7.3, 57.8]	29.8%	43.8%	11.3%	55.9%
Consumer task clusters and feedback-loop revisions						
Global pure (1,007 clusters)	98.2%	[14.2, 53.2]	33.2%	47.8%	6.1%	59.9%
US + loops (20 workflows)	98.6%	[24.1, 80.5]	50.3%	–	22.3%	0.0%
Global + loops (1,007 clusters)	98.2%	[30.8, 72.7]	50.2%	–	25.8%	21.2%
Commercial first-party API (global, excluding Claude Code)						
1P API pure (20 workflows)	95.1%	[78.8, 96.8]	81.5%	–	97.0%	0.0%
1P API pure (860 clusters)	89.2%	[74.8, 94.7]	77.0%	–	95.8%	0.0%
1P API + loops (860 clusters)	89.2%	[80.1, 98.2]	81.7%	–	98.5%	0.0%

Panel B. Workflow Channel Sorting Across Consumer Chat, Feedback Loops, and Commercial API

	Consumer channel (Claude.ai + Cowork)				Commercial 1P API	
Standardised workflow	US pure $[L,U]$	US $\psi_g^\star$	UK pure $[L,U]$	US + loops	API s_g	Pure $[L,U]$
Operational workflows above 50% pure directive in consumer channel						
Meeting Intelligence	[63.4, 75.5]	0.00	[62.3, 76.1]	[71.5, 83.6]	97.2%	[92.9, 99.1]
Business Process & Ops.	[54.5, 73.6]	0.00	[52.6, 72.6]	[71.9, 91.0]	95.3%	[90.4, 98.2]
Technical workflows with overlapping intervals after feedback loops enter						
DevOps & Infrastructure	[3.4, 57.0]	6.66	[0.0, 48.0]	[72.1, 100.0]	87.2%	[70.6, 100.0]
Software Development	[0.0, 32.0]	∞	[0.0, 29.7]	[46.9, 100.0]	90.2%	[76.4, 89.1]
Data Analysis & BI	[0.0, 64.2]	5.89	[0.0, 58.3]	[9.1, 100.0]	87.2%	[88.8, 100.0]
Commercial and regulatory workflows remaining separated from API floors						
Sales & Revenue Ops.	[46.1, 49.7]	∞	[45.1, 48.0]	[62.1, 65.7]	99.6%	[90.7, 94.6]
Compliance & Regulatory	[0.0, 29.2]	∞	[0.0, 29.6]	[16.3, 50.6]	98.8%	[84.3, 86.4]
Content Creation & Copywriting	[0.0, 44.1]	∞	[0.0, 42.6]	[1.4, 51.5]	87.6%	[87.7, 100.0]

Notes: Panel A reports May 2026 workplace coverage $1-\alpha_{\text{rem}}$, Equation (4) full-population bounds $[\Theta_{\text{full}}^L, \Theta_{\text{full}}^U]$, monotone floors L_{mono} under $\theta_g \geq b_g$, September 2026 ChatGPT *Doing*|*Work* shares (measured on a separate output-request rubric), and retained workplace shares in cells strictly above ($\theta_g^L > 50\%$) or below ($\theta_g^U < 50\%$) parity; complete twenty-workflow bounds appear in Online Appendix Table A3. API series are global. Panel B reports eight focal workflows across consumer chat (Claude.ai and Cowork), feedback-loop revisions, parity-breakdown index $\psi_g^\star$, and commercial 1P API workplace shares s_g and pure directive bounds excluding Claude Code.

across work-filtered Claude and work-related ChatGPT traffic (Spearman ρ from 0.769 to 0.842 in Q2 2026 and 0.776 to 0.845 in Q3). National ChatGPT *Doing* shares also sit inside a narrow 40.5 to 43.8 per cent band. On Google Gemini, requested full-automation shares correlate strongly with pre-AI structured desk intensity R_g ($r = +0.72$). Those rates reach 35.65 per cent in Office and Administrative Support and 20.43 per cent in Computer and Mathematical work against just 1.68 per cent in Management and 2.04 per cent in Legal occupations (Online Appendix Table A5), a split that echoes the knowledge-hierarchy pattern in Section 4.

Our empirical evidence carries clear boundaries that govern how every result in Sections 3 to 5 should be read. Every occupational table matches prompts to O*NET or SOC tasks from conversation text without access to the user's true job title, employer, or wage, so task incidence need not equal individual worker use when managers write code or households ask medical questions. Coverage across platforms is also partial. The three public releases span non-overlapping 2026 windows from April to September and omit dedicated enterprise workspaces, coding agents, and cloud hosts, such as ChatGPT Enterprise, Codex, Claude Code, Google Workspace, paid Gemini API text, third-party clouds, and national API volumes. Measurement units likewise differ across messages, conversations, search queries, and stateless API calls. Because no provider publishes cell denominators or sampling variances, we treat our linear-programme and Fréchet intervals as deterministic identification sets conditional on published margins rather than confidence intervals for sampling or classifier error (Imbens and Manski 2004, Stoye 2009). Prompt classifiers also stop at the model window and cannot see whether an API response triggers an automated action or passes through software checks and human review inside the firm (Brynjolfsson, Li, and Raymond 2025, Tomlinson et al. 2025). Above all, our estimates measure where task and baseline-wage incidence falls at current prices rather than causal effects on wages, hours, or employment.

6. Conclusion

Whether generative AI looks like a consumer appliance or a workplace engine depends on which channel we observe. The channel decides the answer. Across ChatGPT, Claude, and Google Gemini in seven Western economies, free search and chat surfaces serve mainly unpaid life at home and study while paid subscriptions and commercial APIs shift steeply towards paid work. Within any one provider, national borders matter far less than access channels and task structure, and six economies predict the seventh out of sample with almost no loss of fit. Pre-AI personal-service intensity explains why tasks such as health care, food preparation, and social support stay overwhelmingly domestic in consumer chat. Analysts who treat unfiltered consumer traffic as workplace labour therefore fold unpaid household output into market accounts and overstate workplace task volume more than tenfold in personal-service fields.

Inside paid employment, our evidence reframes how generative AI relates to wages and firm organisation. Steep occupational pay gradients across both Claude and directly reported US Gemini occupations arise mainly between broad white-collar families and manual trades rather than within one occupation group. Coarse European tables hide that pattern. Single-digit statistics in Continental Europe reproduce steep macro slopes simply because no within-family variation remains to condition on. Inside narrow job titles across 184 million American and British payroll jobs, wide coworker wage dispersion pulls lower-paid clerical staff into workplace AI reach even as steep hourly wage gaps concentrate hourly-wage-weighted exposure in the top fifth. Traffic shares and wage-bill shares answer two different questions on one distribution. When workplace tasks reach a model, assumption-free bounds show that commercial APIs run mostly on direct instructions whereas human chat stays interactive. Multi-turn feedback loops bring chat and API bounds into overlap in technical work such as software development and infrastructure operations, while regulatory compliance, copywriting, and sales stay sharply separated between bespoke human judgement in chat and standardised bulk execution on APIs.

Privacy need not block better measurement. Model providers do not need to expose a single private prompt to collapse our partial-identification intervals. Providers could publish matched national consumer and commercial API volumes on one common unit, report national workplace shares across detailed occupations, and release thresholded cross-tabulations of work context against interaction mode inside task clusters. Statistical agencies could pair those releases with finer hourly wage percentiles, harmonised detailed occupational earnings across Continental Europe, and short labour-force survey modules that record weekly workplace AI volume alongside job title and hourly pay. Those public tables would replace wide identification bounds with direct national accounts of workplace AI.

Our estimates stop short of causality. Every figure in this paper measures where task and baseline-wage incidence falls at current prices rather than causal effects on wages, hours, or employment. Causal evidence on how chat assistants and automated API pipelines alter staffing, task boundaries, career progression, and unpaid domestic welfare will require longitudinal employer-employee panels linked to enterprise AI rollouts alongside consented household time-use diaries. Until economists study unpaid home use, internal pay bands, and firm software together, single-channel snapshots will continue to confuse one slice of generative AI with the whole technology.

References

Acemoglu, Daron. 2025. “The Simple Macroeconomics of AI.” *Economic Policy* 40 (121): 13–58.

Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review* 108 (6): 1488–1542.

Appel, Ruth, Maxim Massenkoff, Peter McCrory, Miles McCain, Ryan Heller, Tyler Neylon, and Alex Tamkin. 2026. "Anthropic Economic Index Report: Economic Primitives." research report, Anthropic. https://www.anthropic.com/research/anthropic-economic-index-january-2026-report.

Atalay, Enghin, Phai Phongthiengtham, Sebastian Sotelo, and Daniel Tannenbaum. 2020. "The Evolution of Work in the United States." *American Economic Journal: Applied Economics* 12 (2): 1–34.

Autor, David H., and Michael J. Handel. 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics* 31 (S1): S59–S96.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): 1279–1333.

Autor, David, and Neil Thompson. 2025. "Expertise." *Journal of the European Economic Association* 23 (4): 1203–1271.

Becker, Gary S. 1965. "A Theory of the Allocation of Time." *Economic Journal* 75 (299): 493–517.

Bick, Alexander, Adam Blandin, and David J. Deming. 2026. "The Rapid Adoption of Generative AI." *Management Science* 72 (10): 8440–8456.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *Quarterly Journal of Economics* 140 (2): 889–942.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–427.

Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma J. Rockall, and Marina Mendes Tavares. 2024. "Gen-AI: Artificial Intelligence and the Future of Work." Staff Discussion Note SDN/2024/001, International Monetary Fund.

Cerina, Fabio, Alessio Moro, and Michelle Rendall. 2021. "The Role of Gender in Employment Polarization." *International Economic Review* 62 (4): 1655–1691.

Charnes, Abraham, and William W. Cooper. 1962. "Programming with Linear Fractional Functionals." *Naval Research Logistics Quarterly* 9 (3-4): 181–186.

Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. "How People Use ChatGPT." Working Paper 34255, National Bureau of Economic Research.

Chatterji, Aaron, David Holtz, Neel Rakholia, Prasanna Tambe, and Gawesha Weeratunga. 2026. "How Organizations Use AI: Evidence from ChatGPT." working paper, OpenAI. https://cdn.openai.com/pdf/how-organizations-use-chatgpt.pdf.

Cross, Philip J., and Charles F. Manski. 2002. "Regressions, Short and Long." *Econometrica* 70 (1): 357–368.

Cui, Kevin Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2026. "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers." *Management Science* 72 (11): 9408–9420.

Datta, Nikhil, Peter John Lambert, and Patrick Schneider. 2026. "Who Pays for AI?" Working Paper 828/2026, University of Warwick, CAGE Research Centre. https://warwick.ac.uk/fac/soc/economics/research/centres/cage/publications/workingpapers/2026/who_pays_for_ai/.

Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *Quarterly Journal of Economics* 132 (4): 1593–1640.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–1308.

Eurostat. 2025a. "Employment by Detailed Occupations (ISCO-08, Dataset lfsa_egai2d)." European Union Labour Force Survey. https://ec.europa.eu/eurostat/databrowser/view/lfsa_egai2d/.

Eurostat. 2025b. "Structure of Earnings Survey 2022: Mean Hourly Earnings by Sex, Occupation and Economic Activity (earn_ses22_47)." Dissemination API Dataset. https://ec.europa.eu/eurostat/databrowser/view/earn_ses22_47/.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12): 2195–2217.

Fréchet, Maurice. 1951. "Sur les tableaux de corrélation dont les marges sont données." *Annales de l'Université de Lyon, Section A: Sciences Mathématiques et Astronomie* 14: 53–77.

Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy* 108 (5): 874–904.

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104 (8): 2509–2526.

Handa, Kunal, 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. 2025. "Which Economic Tasks Are Performed with AI? Evidence from Millions of Claude Conversations." arXiv Preprint 2503.04761, Anthropic. https://arxiv.org/abs/2503.04761.

Humlum, Anders, and Emilie Vestergaard. 2025. "The Unequal Adoption of ChatGPT Exacerbates Existing Inequalities among Workers." *Proceedings of the National Academy of Sciences* 122 (1): e2414972121.

Imbens, Guido W., and Charles F. Manski. 2004. "Confidence Intervals for Partially Identified Parameters." *Econometrica* 72 (6): 1845–1857.

Iscenko, Zanna, Scott Strand, Yiyuan Chen, Guillaume Aimard, Mihai Codreanu, Vivek Sampathkumar, Alex Imas, Julian Jacobs, Evalyne Muiruri, Jia Jen Ng, Samirah Javed, Juan Mateos-Garcia, Denis Calin, Andrew Kim, Josh Martin, Omar Ajmeri, Fabien Curto Millet, and James Manyika. 2026. "Google's AI & Economy ATLAS v1.0: Mapping Gemini Usage in the Economy." Technical Report arXiv:2608.00038, Google and Google DeepMind. https://arxiv.org/abs/2608.00038.

Jones, Charles I. 2026. "A.I. and Our Economic Future." *Journal of Economic Perspectives* 40 (3): 3–22.

Kansoy, Fatih. 2026. "AI Use and Occupations: Task-Weighted International Crosswalks and Distributional Benchmarks." research dataset and working paper, University of Oxford. https://fatih.ai/ai-use-and-occupations/.

Manski, Charles F. 2003. *Partial Identification of Probability Distributions*. New York: Springer.

Massenkoff, Maxim, Eva Lyubich, Szymon Sacher, Zoe Hitzig, Shaoyi Zhang, Ryan Heller, and Peter McCrory. 2026. "Anthropic Economic Index Report: Cadences." research report, Anthropic. https://www.anthropic.com/research/economic-index-june-2026-report.

Masten, Matthew A., and Alexandre Poirier. 2020. "Inference on Breakdown Frontiers." *Quantitative Economics* 11 (1): 41–111.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381 (6654): 187–192.

Office for National Statistics. 2025. "Annual Population Survey (APS) Employment by SOC2020 Occupation and ITL1 Region, Jan–Dec 2025." Nomis Dataset NM_218_1. https://www.nomisweb.co.uk.

Office for National Statistics. 2026. "Employee Earnings in the UK: ASHE Table 14 (4-Digit SOC2020), 2026 Provisional and 2025 Revised." Dataset under Open Government Licence v3.0. https://www.ons.gov.

uk.

OpenAI. 2026. “OpenAI Signals: Version 3.0 Release and Data Dictionary (July 2024–September 2026).” Dataset under CC BY 4.0. https://cdn.openai.com/signals/data-download-csv.zip.

Stoye, Jörg. 2009. “More on Confidence Intervals for Partially Identified Parameters.” *Econometrica* 77 (4): 1299–1315.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. 2025. “Working with AI: Measuring the Applicability of Generative AI to Occupations.” arXiv Preprint 2507.07935, Microsoft Research. https://arxiv.org/abs/2507.07935.

U.S. Bureau of Labor Statistics. 2026. “Occupational Employment and Wage Statistics (OEWS), May 2025 National Cross-Industry Estimates.” Time-Series Release oe.data.0.Current. https://download.bls.gov/pub/time.series/oe/.

Yin, Michelle, and Burhan Ogut. 2026. “Who Uses AI? Platform Selection and the Measurement of Occupational AI Exposure.” arXiv Preprint 2605.21743, arXiv. https://arxiv.org/abs/2605.21743.

Online Appendix

A. Multi-Channel Data Construction, Harmonised Crosswalks, and Pay Grids

Global context series combine 27 monthly waves of OpenAI Signals version 3.0 from July 2024 to September 2026 across adult consumer accounts, excluding ChatGPT Enterprise and Codex (OpenAI 2026). Those tables cover 160 economies in total and 149 in September 2026, and we preserve Namibia's two-letter ISO code NA explicitly rather than treating it as missing text. OpenAI publishes binary work-related shares separately from its global three-way Work, Schoolwork, and Other taxonomy, monthly topic shares by work status, and September 2026 consumer plan splits (Free, Go, Plus, and Pro), alongside output-oriented *Doing* shares among work-related messages across 126 economies. We pair those records with 5 waves of the Anthropic Economic Index from August 2025 to May 2026, including 4 work-context waves from November 2025 onward (Appel et al. 2026, Massenkoff et al. 2026). Anthropic separates consumer conversations (web, desktop, and mobile Claude.ai chat pooled with Cowork across 121 economies, 52 US states and territories, 652 global subregions, 718 O*NET occupation workplace priors, and 1,007 fine task clusters) from commercial first-party API calls. Commercial API tables span twenty standardised workflows and 860 fine clusters under full-month expanded-context sampling from April 2026 onward, while omitting dedicated Claude Code sessions and third-party cloud hosts.

Pre-AI task predictors come from O*NET and PIAAC descriptors fixed before our 2025–2026 usage window (Kansoy 2026), paired with occupational AI capability rubrics from Felten, Raj, and Seamans (2021) and Eloundou et al. (2024). Structured desk intensity R_g standardises computer use, documentation, and numerical work. Personal-service intensity S_g standardises social interaction demand net of half problem-solving demand and half computer use, so it peaks in care, health support, food preparation, and teaching. In the United States, May 2025 BLS OEWS tables cover 830 six-digit SOC occupations that employ 155.48 million payroll workers (U.S. Bureau of Labor Statistics 2026). That panel retains 437 active matched titles alongside 393 zero-coded or suppressed titles (254 censored or zero cells, 92 sub-threshold rounded cells, and 47 unmapped residuals). Normalised employment-conserving weights map all 830 American codes onto eight non-military, non-agricultural single-digit ISCO-08 groups.

In the United Kingdom, the 2025 ONS Annual Population Survey (APS) counts 33.07 million workers across four-digit SOC2020 unit groups (Office for National Statistics 2025). Employer payroll records in 2026 ASHE Table 14 cover 28.92 million employee jobs, of which 28.85 million match across the 382 unit groups with positive ASHE employment out of 402 paid unit groups (including 29 zero-use groups) in 26 two-digit SOC2020 job families (Office for National Statistics 2026). To map

US O*NET task volumes onto UK four-digit unit groups without letting tiny sub-occupations distort unit-group averages, multi-match cells use observed UK APS unit-group employment weights (w_{emp}) rather than equal-split weights (w_{norm}). Table A2 compares both weighting rules and contrasts 26 two-digit SOC2 fixed effects with 9 single-digit SOC1 fixed effects. Four narrow UK unit groups lack crosswalked capability scores in Column (5) of Table 3 and the Felten-only row of Table A2 ($N = 398$), and three lack scores in the Eloundou-only row of Table A2 ($N = 399$). Occupation-level pre-AI R_i and S_i remain complete across all 402 unit groups through sample-median descriptor imputation.

In Continental Europe, Eurostat 2022 Structure of Earnings Survey hourly pay in enterprises with at least ten staff and 2025 Labour Force Survey employment yield 40 non-military, non-agricultural country-by-ISCO-1 cells covering 123.91 million workers (Eurostat 2025b,a). Within-country employment shares weight each country equally on Section 4's single-digit rung. Because Eurostat publishes no fine within-occupation wage percentiles across these cells, Continental Europe cannot enter our continuous pay-grid analysis.

Across the 184.33 million American and British employee jobs with published pay quantiles, we construct piecewise linear cumulative distribution functions inside every occupation. American grids integrate five OEWS wage knots above a $7.25 federal minimum-wage floor with mean-calibrated upper tails clipped to $[1.02p_{90}, 3.50p_{90}]$, tested against restricted $[0.95p_{10}, 1.15p_{90}]$ tails. British grids integrate eleven ASHE wage knots alongside occupational means. Suppressed British p_{10} values are imputed from single-digit SOC1 median p_{10}/p_{20} ratios floored at £8.00 and checked against the £12.71 adult floor.

B. Methods and Derivations of the Bounds

This appendix derives the expressions and bounding programmes used in Sections 2 to 5. Throughout, w with a country, task, or channel subscript denotes a paid-work share while $\bar{w}_{iq}$, w_i^{med}, and w_i^{mean} denote hourly pay in occupation i. In Section 5, s_g denotes the workplace share of task cluster g. Differentiating the consumer paid-work share in Equation (1) with respect to personal-service intensity S_g gives

$$\frac{\partial w_{cg}}{\partial S_g} = -\frac{M_{cg}^{W}(\alpha_c, R_g)\,(\partial M_{cg}^{\text{NW}}/\partial S_g)}{\left(M_{cg}^{W} + M_{cg}^{\text{NW}}\right)^2} < 0, \qquad \log\frac{w_{cg}}{1-w_{cg}} = \log\frac{\alpha_c}{\gamma_c} + \log\frac{M_0^{W}(R_g)}{M_0^{\text{NW}}(S_g)}. \tag{5}$$

The derivative is negative whenever workplace volume is positive and household volume rises strictly with S_g. Slower employer adoption in services delivered in person, $\partial M_{cg}^{W}/\partial S_g \le 0$, adds a nonpositive numerator term that reinforces the same sign. Multiplicative country factors $\alpha_c, \gamma_c > 0$ with $M_{cg}^{W} = \alpha_c M_0^{W}(R_g)$ and $M_{cg}^{\text{NW}} = \gamma_c M_0^{\text{NW}}(S_g)$ produce the additive log-odds split in Equation (5) and preserve task rankings across countries. Across the 154 cells of our seven-economy panel,

Section 3 estimates unweighted country-fixed-effects regressions of percentage-point paid-work shares on S_g and separately on R_g. Because S_g and R_g vary across 22 two-digit groups, both regressions use small-sample clustered standard errors on the 22 groups, while the personal-service regression adds restricted Rademacher wild cluster bootstrap tests (Cameron, Gelbach, and Miller 2008) and 22-group randomisation inference. Leave-one-country-out validation predicts each held-out country's 22 shares from the unweighted six-country mean with zero fitted parameters to test level fit alongside rank stability.

Consumer non-work majorities alone cannot prove that non-work accounts for most total AI use because national commercial API volumes are withheld. We normalise national consumer volume to one with non-work share $s_{\text{NW}} \in (0,1)$ and write $E \geq 0$ for the unobserved ratio of national commercial API calls to consumer conversations. Country-specific API context shares are also unpublished, so every national threshold transports Anthropic's global API paid-work share $w_E \in (0,1)$. When workplace requests save $\kappa > 0$ times as many human hours as non-work queries, duration-weighted paid work reaches parity whenever $W(E,\kappa) \geq 1/2$, with interior parity ratio $E^\star(\kappa, w_E)$:

$$\begin{aligned} W(E,\kappa) &= \frac{\kappa\left[(1-s_{\text{NW}}) + Ew_E\right]}{\kappa\left[(1-s_{\text{NW}}) + Ew_E\right] + s_{\text{NW}} + E(1-w_E)} \geq \frac{1}{2}, \\ E^\star(\kappa, w_E) &= \frac{s_{\text{NW}} - \kappa(1-s_{\text{NW}})}{\kappa w_E - (1-w_E)} \qquad \text{for } \kappa < \kappa^\star \text{ and } \kappa w_E > 1 - w_E. \end{aligned} \tag{6}$$

Three regimes govern Equation (6). Consumer traffic already reaches parity at $E = 0$ once $\kappa \geq \kappa^\star = s_{\text{NW}}/(1-s_{\text{NW}})$, though large API volume could reverse parity if $\kappa w_E < 1 - w_E$. When $\kappa < \kappa^\star$ and $\kappa w_E > 1 - w_E$, $E^\star(\kappa, w_E)$ is the smallest positive API-to-consumer ratio that brings paid work to one half (Masten and Poirier 2020). When $\kappa < \kappa^\star$ and $\kappa w_E \leq 1 - w_E$, no finite API volume achieves parity. Section 3 evaluates equal-duration prompt parity $E^\star(1, w_E)$ and duration parity $\kappa^\star$ at $w_E = 0.8403$ alongside the one-sided consumer misclassification margin $m^\star = \max(0,\, s_{\text{NW}} - 1/2)$.

Section 4 relates workplace AI intensity $I_i = s_i/e_i$ to log hourly pay x_i across 830 American six-digit occupations, 402 paid British four-digit unit groups, and 40 European single-digit ISCO-08 cells (Autor, Levy, and Murnane 2003, Autor and Thompson 2025). Here x_i is log median hourly pay in America and Britain and log mean hourly pay in Europe, where within-country employment shares weight each country equally. Within each estimation sample $\mathcal{S}$, workplace traffic shares $s_i = m_i^W / \sum_{j \in \mathcal{S}} m_j^W$ and employment shares e_i sum to one so that employment-weighted mean intensity equals one. National Anthropic files publish all-context source-occupation conversation shares m_j^{all} and two-digit job-family paid-work shares w_f, which fix family workplace targets $T_f^W = w_f \sum_{j \in f} m_j^{\text{all}}$. Baseline allocations clip Anthropic's 718 global detailed O*NET work-share priors π_j to $[10^{-4}, 1 - 10^{-4}]$ after filling missing priors with Generalized Work Activity and two-digit shares in America and with two-digit shares in Britain. Inside each two-digit family f, a 60-step bisection shifts log odds by one constant ζ_f until source workplace volumes $m_j^W = m_j^{\text{all}}/[1 + \exp(-(\text{logit}\,\pi_j + \zeta_f))]$

match T_f^W, and fixed APS employment weights then map source O*NET allocations onto British unit groups. Keeping all zero-coded titles, our main regressions with family fixed effects $\xi_{f(i)}$ and pre-AI task and capability controls $\mathbf{z}_i$ (Felten, Raj, and Seamans 2021, Eloundou et al. 2024) are

$$\mathbb{E}[s_i \mid x_i, f(i), \mathbf{z}_i] = e_i \exp(\beta_0 + \beta x_i + \xi_{f(i)} + \mathbf{z}_i^\top \boldsymbol{\gamma}), \qquad \log(I_i + \epsilon) = \beta_0 + \beta x_i + \xi_{f(i)} + \mathbf{z}_i^\top \boldsymbol{\gamma} + \eta_i. \quad (7)$$

We estimate the first equation by Poisson pseudo-maximum likelihood (PPML) with offset $\log e_i$ (Santos Silva and Tenreyro 2006) and the second by employment-weighted least squares (WLS). The shift ϵ equals half the smallest positive I_i when zero-coded titles exist and zero in single-digit cells. Because shifted-log slopes depend on ϵ (Chen and Roth 2024), Table A2 also estimates linear probability models for $s_i > 0$ and pure log-WLS on active titles after re-normalising s_i and e_i. Standard errors condition on the baseline allocation and cluster on 22 American families, 26 British families, or eight European ISCO-08 groups with country fixed effects.

Dropping the logit-raking rule defines the nonempty compact polytope Ω of nonnegative source-occupation workplace vectors $\mathbf{m}^W = (m_j^W)$ that obey national all-context caps $0 \le m_j^W \le m_j^{\text{all}}$, hold baseline zero-traffic codes at zero, and match each two-digit family target $\sum_{j \in f} m_j^W = T_f^W$. Detailed occupation and source indices coincide in America, whereas British unit-group volumes equal the fixed APS employment-weighted crosswalk image of $\mathbf{m}^W \in \Omega$. Subsamples that omit active titles replace affected family equalities with two-sided bounds. Online Appendix D also tests a relaxed American polytope that caps censored and sub-threshold rounded codes at 0.005 per cent under two-sided family bounds alongside a point WLS check at 0.0025 and 0.0010 per cent raw shares. Let $\check{x}_i$ be the employment-weighted residual from projecting x_i onto the remaining regressors in $\mathcal{S}$ and set $D_x = \sum_{i \in \mathcal{S}} e_i \check{x}_i^2$. By the Frisch–Waugh–Lovell theorem, the level-WLS slope $\beta^{\text{lev}}(\mathbf{m}^W) = (\sum_{i \in \mathcal{S}} \check{x}_i m_i^W)/(D_x \sum_{i \in \mathcal{S}} m_i^W)$, occupation-median fifth shares H_q^{occ}, allocated employee fifth shares H_q, bottom-fifth gains $\Delta H_1 = H_1 - H_1^{\text{occ}}$, and fifth-share ratios all take the linear-fractional form $\mathbf{c}^\top \mathbf{m}^W / (\mathbf{d}^\top \mathbf{m}^W)$. Even though $\sum_{i \in \mathcal{S}} m_i^W$ is constant on the full American baseline polytope, the denominator $\mathbf{d}^\top \mathbf{m}^W$ varies across Ω in Britain, in restricted subsamples, on the relaxed < 0.005 per cent polytope, and in fifth-share ratios. Setting $z = 1/(\mathbf{d}^\top \mathbf{m}^W) > 0$ and $\tilde{\mathbf{m}} = z\mathbf{m}^W$ (Charnes and Cooper 1962) turns both extrema over Ω into linear programmes:

$$\min_{\tilde{\mathbf{m}},\, z \ge 0} \text{ and } \max_{\tilde{\mathbf{m}},\, z \ge 0} \mathbf{c}^\top \tilde{\mathbf{m}} \quad \text{subject to} \quad \tilde{\mathbf{m}} \in z\Omega, \;\; \mathbf{d}^\top \tilde{\mathbf{m}} = 1. \quad (8)$$

Fifth-share ratios report ∞ when the denominator's lower bound reaches zero while the numerator stays positive. Job-family fixed effects remove between-family variation from the slope numerator while family totals remain in Ω, which widens conditional level bounds. Nonlinear PPML and shifted-log slopes are not bounded here because they are not linear ratios.

Wide coworker pay dispersion inside occupations separates traffic shares across pay bands from hourly-wage-weighted shares (Autor and Handel 2013). Let $G_i \in \mathcal{G}_i$ denote the monotone hourly wage distribution function in occupation i, with $\mathcal{G} = \prod_i \mathcal{G}_i$, and let $0 = \tau_0 < \tau_1 < \cdots < \tau_4 <$

$\tau_5 = \infty$ be baseline national employee quintile cutoffs with $G_i(0) = 0$ and $G_i(\infty) = 1$. Cell overlap $\phi_{iq} = G_i(\tau_q) - G_i(\tau_{q-1})$ is the share of occupation i's employees paid in fifth q at cell mean hourly wage $\bar{w}_{iq}$. Across the 184.33 million American and British payroll jobs, allocated workplace traffic shares H_q under equal within-occupation use and baseline hourly-wage-weighted shares B_q satisfy

$$H_q(\mathbf{m}^W, \boldsymbol{\phi}) = \frac{\sum_i m_i^W \phi_{iq}}{\sum_i m_i^W}, \qquad B_q(\mathbf{m}^W, \boldsymbol{\phi}) = \frac{\sum_i m_i^W \phi_{iq} \bar{w}_{iq}}{\sum_{r=1}^5 \sum_i m_i^W \phi_{ir} \bar{w}_{ir}}, \qquad \sum_{q=1}^5 H_q = \sum_{q=1}^5 B_q = 1. \quad (9)$$

We evaluate B_q on the baseline continuous grid and compare it with baseline hourly-wage-bill shares P_q at one paid hour per matched job under macroeconomic cost-share accounting rather than realised savings (Acemoglu 2025, Jones 2026). Occupation-median benchmarks H_q^{occ} instead assign each occupation's traffic using employment-weighted median-wage cutoffs. Holding each benchmark's cutoffs fixed keeps every traffic-share objective linear-fractional with fixed bracket endpoints, while restricted-tail checks in Online Appendix D recalculate American cutoffs. Published wage percentiles define step bounds $\underline{G}_i(\tau)$ and $\overline{G}_i(\tau)$, equal to $[0, p_{\min}]$ below the lowest published percentile $p_{\min}$, $[p_{\max}, 1]$ above the highest published percentile $p_{\max}$, and $[0, 1]$ when fewer than three percentiles are published. Each single-fifth overlap therefore lies between $\phi_{iq}^{\text{lb}} = \max(0, \underline{G}_i(\tau_q) - \overline{G}_i(\tau_{q-1}))$ and $\phi_{iq}^{\text{ub}} = \min(1, \overline{G}_i(\tau_q) - \underline{G}_i(\tau_{q-1}))$. Because unit-group workplace volumes are nonnegative on Ω, single-fifth brackets vary independently across occupations, and neither $\sum_i m_i^W$ nor H_1^{occ} depends on $\boldsymbol{\phi}$, the bilinear joint extrema of each fifth share H_q and of ΔH_1 over $\Omega \times \mathcal{G}$ occur fifth by fifth at ϕ_{iq}^{lb} and ϕ_{iq}^{ub} and solve via Equation (8). Relaxing equal use inside occupations for $\rho > 0$ replaces ϕ_{iq} with $\tilde{\phi}_{iq}(\rho) = \phi_{iq}\lambda_q(\rho) / \sum_{r=1}^5 \phi_{ir}\lambda_r(\rho)$ under linear weights $\lambda_q^{\text{lin}}(\rho) = 1 + (\rho - 1)(q - 1)/4$ or geometric weights $\lambda_q^{\text{geo}}(\rho) = \rho^{(q-1)/4}$ and solves $H_1(\rho^\star) = H_1^{\text{occ}}$ for break-even $\rho^\star$ whenever $H_1(1) > H_1^{\text{occ}}$.

Section 5 bounds the workplace rate $\theta_g \in [0, 1]$ of a target interaction mode against the non-work rate $b_g \in [0, 1]$ across retained task clusters $g \in \mathcal{R}$. In that section, g indexes workflows or fine Anthropic clusters rather than two-digit occupational groups. Public files report cluster traffic share N_g, workplace share $s_g \in (0, 1]$, observed target-mode share $d_g^{\text{obs}} \geq 0$, other classified share $c_g^{\text{obs}} \geq 0$, and unclassified share $u_g = 1 - d_g^{\text{obs}} - c_g^{\text{obs}} \geq 0$, which confines the true target share to $d_g \in [d_g^{\text{obs}}, d_g^{\text{obs}} + u_g]$. Probability conservation $d_g = s_g\theta_g + (1 - s_g)b_g$ (Fréchet 1951, Manski 2003), retained workplace weights $\tilde{\alpha}_g = N_g s_g / \sum_{h \in \mathcal{R}} N_h s_h$, and suppressed workplace share $\alpha_{\text{rem}} = \max(0, N_{\text{tot}}^W - \sum_{h \in \mathcal{R}} N_h s_h)/N_{\text{tot}}^W$ under channel workplace share N_{tot}^W give

$$\theta_g^L = \frac{\max(0, d_g^{\text{obs}} + s_g - 1)}{s_g}, \qquad \theta_g^U = \frac{\min(d_g^{\text{obs}} + u_g, s_g)}{s_g}, \qquad (10)$$
$$\Theta_{\text{full}} \in [(1 - \alpha_{\text{rem}})\Theta_{\text{ret}}^L, \ (1 - \alpha_{\text{rem}})\Theta_{\text{ret}}^U + \alpha_{\text{rem}}],$$

where $\Theta_{\text{ret}}^{L,U} = \sum_{g \in \mathcal{R}} \tilde{\alpha}_g \theta_g^{L,U}$ and suppressed traffic takes worst-case rates in $[0, 1]$. For $0 < s_g < 1$, θ_g^L is attained at $d_g = d_g^{\text{obs}}$ and $b_g = \min(1, d_g^{\text{obs}}/(1 - s_g))$, while θ_g^U is attained at $d_g = d_g^{\text{obs}} + u_g$ and $b_g = \max(0, (d_g^{\text{obs}} + u_g - s_g)/(1 - s_g))$. Those endpoints reproduce the one-sixth to five-sixths

example in Section 2 when $s_g = 0.6$, $d_g^{\text{obs}} = 0.5$, and $u_g = 0$. At $s_g = 1$ the bounds collapse to $[d_g^{\text{obs}}, \min(d_g^{\text{obs}} + u_g, 1)]$, and higher s_g narrows $[\theta_g^L, \theta_g^U]$ mechanically. For nested partitions with identical retained workplace traffic and consistent subcluster margins, finer clusters cannot widen $\Theta_{\text{ret}}^{L,U}$ because $\max(0, x_1) + \max(0, x_2) \geq \max(0, x_1 + x_2)$ and $\min(a_1, b_1) + \min(a_2, b_2) \leq \min(a_1 + a_2, b_1 + b_2)$. Monotone selection assumes $\theta_g \geq b_g$, which implies $\theta_g \geq d_g^{\text{obs}}$ and yields the floor $L_{\text{mono}} = (1 - \alpha_{\text{rem}}) \sum_{g \in \mathcal{R}} \tilde{\alpha}_g d_g^{\text{obs}}$. Our parity-breakdown index sets $\psi_g^\star = 0$ when $\theta_g^L > 1/2$, $\psi_g^\star = \infty$ when $\theta_g^U < 1/2$ or $d_g^{\text{obs}} + u_g \leq s_g/2$, and $\psi_g^\star = (1 - d_g^{\text{obs}} - s_g/2)/(d_g^{\text{obs}} + u_g - s_g/2)$ otherwise. For $0 < s_g < 1$, dropping $-u_g$ from the numerator makes $\psi_g^\star$ an upper bound on the minimum workplace-to-non-work odds ratio compatible with $\theta_g \geq 1/2$ at the most favourable unclassified assignment $d_g = d_g^{\text{obs}} + u_g$, with equality whenever $u_g = 0$ (Cross and Manski 2002, Masten and Poirier 2020).

C. Subnational Distributions, Parity Frontiers, and Zero-Parameter LOCO

Table A1 details country-by-country parity thresholds and out-of-sample task fit across our seven focal Western economies. Panel A reports consumer non-work shares, misclassification margins $m^\star$, equal-duration API parity ratios $E^\star(1, 0.8403)$ with implied combined API shares, consumer duration-parity thresholds $\kappa^\star$, and task-mix-weighted estimated duration ratios $\hat{\kappa}$ across complete-duration cells. Estimated $\hat{\kappa}$ spans 1.35 in Spain to 1.45 in America and Italy. That ratio clears the duration-parity hurdle $\kappa^\star$ in America ($1.45 > 1.42$) and Italy ($1.45 > 1.33$), yet stays below parity in Britain ($1.39 < 1.56$). Paid work accordingly represents 50.54 per cent of estimated consumer task-hours saved in America against 47.00 per cent in Britain. Panel B lists all 8 economies out of 121 where May 2026 consumer Claude paid-work traffic exceeds half alongside all seven folds of our zero-parameter leave-one-country-out check.

Figure A1 shows that consumer non-work majorities hold well below the national level and that occupational work splits travel cleanly across Western borders. In Panel A, every one of the 52 American states and territories exceeds 50 per cent non-work on May 2026 consumer Claude (median 59.61 per cent; California 58.48 per cent), as do 587 of 652 global subregions (median 58.30 per cent). Panel B plots all 154 country-by-occupation paid-work shares against unweighted held-out six-country means. Those zero-parameter forecasts reach mean Pearson $r = 0.983$ and mean level fit of 92.4 per cent. France posts a lower level fit (80.6 per cent versus 94.7 per cent demeaned shape fit) because French paid-work shares lie 6.31 percentage points lower on average while keeping almost the same task ranking (Spearman $\rho = 0.935$).

TABLE A1. Seven-Economy Joint Parity Frontiers, Majority-Work Claude Economies, and Zero-Parameter LOCO Folds

Panel A. Seven-Economy Joint Parity Frontiers ($w_E = 84.03\%$ Commercial API) and Task-Duration Ratios

	Anthropic consumer & commercial API (May 2026)				OpenAI Signals (Sep 2026)	
Economy	Non-work %	$m^\star$ (pp)	$E^\star$ (comb. %)	$\kappa^\star$ ($\hat{\kappa}$)	Non-work %	$E^\star$
United States (USA)	58.68%	8.68	25.5% (20.3%)	1.42 (1.45)	66.9%	49.7%
United Kingdom (GBR)	60.98%	10.98	32.3% (24.4%)	1.56 (1.39)	69.1%	56.1%
Germany (DEU)	62.38%	12.38	36.4% (26.7%)	1.66 (1.43)	73.7%	69.6%
France (FRA)	65.77%	15.77	46.3% (31.7%)	1.92 (1.44)	72.0%	64.7%
Spain (ESP)	57.72%	7.72	22.7% (18.5%)	1.37 (1.35)	71.4%	62.9%
Italy (ITA)	57.11%	7.11	20.9% (17.3%)	1.33 (1.45)	67.5%	51.4%
Netherlands (NLD)	59.18%	9.18	27.0% (21.2%)	1.45 (1.42)	69.3%	56.7%

Panel B. Majority-Work Claude Economies (8 of 121) and Zero-Parameter 7-Fold LOCO Fit

Majority-work Claude economies (paid-work %): BRA (57.35%), ARE (54.63%), ISR (53.32%), VNM (53.13%), THA (53.12%), CYP (52.16%), KEN (51.15%), QAT (50.85%).

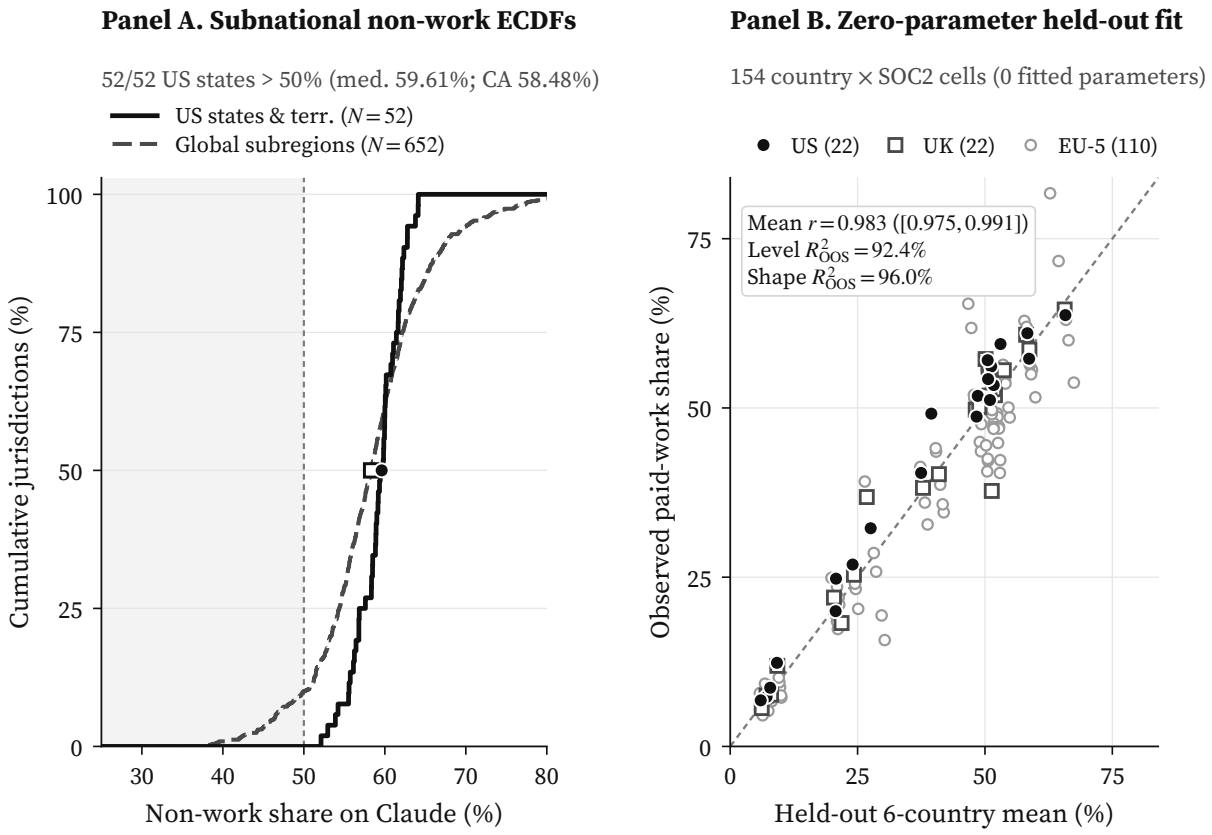
Held-out economy (22 SOC2 cells)	Training set	Pearson r	Spearman ρ	Level R^2_{OOS}	Shape R^2_{OOS}	Mean gap
United States (USA)	6 other	0.990	0.970	96.0%	97.9%	+2.60 pp
United Kingdom (GBR)	6 other	0.975	0.936	94.9%	95.1%	+0.78 pp
Germany (DEU)	6 other	0.991	0.959	93.4%	97.7%	−3.69 pp
France (FRA)	6 other	0.984	0.935	80.6%	94.7%	−6.31 pp
Spain (ESP)	6 other	0.979	0.910	92.0%	95.0%	+3.60 pp
Italy (ITA)	6 other	0.981	0.915	94.0%	95.6%	+2.60 pp
Netherlands (NLD)	6 other	0.981	0.973	96.2%	96.3%	+0.42 pp

Notes: Panel A reports May 2026 consumer Claude non-work shares, one-sided misclassification margins $m^\star = s_{\text{NW}} - 0.50$, Equation (2) API-to-consumer parity ratios $E^\star(1, 0.8403)$ (with combined API share $E^\star/(1+E^\star)$ in parentheses), consumer duration-parity ratios $\kappa^\star$ alongside task-mix-weighted estimated savings ratios $\hat{\kappa}$ across complete-duration cells in parentheses, and September 2026 ChatGPT non-work shares and API parity ratios across all seven focal economies. Panel B lists all 8 majority-work Claude economies out of 121 and reports held-out Pearson r, Spearman ρ, pure level R^2_{OOS}, demeaned shape R^2_{OOS}, and mean level gaps when each country's 22 SOC2 work shares are predicted from the unweighted six-country mean with zero estimated parameters.

D. Two-Part Audits, Crosswalk Diagnostics, and Payroll Cutpoint Straddling

Table A2 audits the occupational pay gradients of Table 3 across extensive versus intensive margins, sub-threshold censoring rules, standalone capability controls, crosswalk weights, and alternative wage bases. Panel A splits American and British regressions into active-presence linear probability models ($\mathbf{1}\{s_i > 0\}$) and pure log-intensity WLS among active occupations ($N = 437$ in America; $N = 373$ in Britain). In America, the extensive-margin slope stays small under SOC2 fixed effects and turns slightly negative under full controls, while the active intensive-margin slope falls by nearly half from +2.27 unconditionally to +1.24 under joint controls. In Britain, where more than nine tenths of unit groups are active, full controls cut the active slope in half from +1.66 to an imprecise +0.88. Panel B confirms that American attenuation does not hinge on zero-coded titles or wage

FIGURE A1. Subnational Non-Work Distributions Across 52 US Jurisdictions and 652 Global Subregions and Held-Out LOCO Predictions



Notes: Panel A plots exact step empirical cumulative distribution functions of May 2026 consumer Claude non-work shares across 52 US states and territories (solid black curve with filled circle median, 59.61%) and 652 global subregions (dashed dark-grey curve with hollow square median, 58.30%). Panel B plots observed May 2026 consumer Claude paid-work shares against zero-parameter held-out six-country unweighted means across all 154 cells of the seven-economy panel (filled black circles for the United States, hollow squares for the United Kingdom, and hollow grey circles for the five Continental European economies) around the 45-degree line.

choices. Imputing positive raw shares (0.0025 and 0.0010 per cent) to censored and sub-threshold rounded titles keeps the joint shifted-log slope statistically indistinguishable from zero (+0.555) and leaves relaxed level LP bounds ($[-3.64, 3.77]$) close to baseline. Standalone capability controls attenuate raw pay slopes strongly even without job-family fixed effects, and SOC2-conditioned slopes stay tightly aligned across GWA workplace priors, constant two-digit shares, hourly-reported occupations, educator exclusions, FTE-1440 wages, strict six-digit matches, and annual mean pay.

Panel C of Table A2 shows why harmonised 26-group SOC2 fixed effects and APS employment-weighted crosswalks (w_{emp}) matter for transatlantic comparability in Britain. Replacing w_{emp} with naive equal-split weights (w_{norm}) pulls the SOC2-conditioned PPML elasticity sharply down from $+1.185$ to $+0.238$, and pairing equal-split weights with coarse 9-group SOC1 fixed effects pushes PPML slightly negative (-0.082). Under an exact SOC2020 crosswalk, the joint shifted-log slope still

attenuates to an imprecise $+0.659$ while the joint PPML elasticity stays modestly positive ($+0.966$) as in our baseline Table 3. Panel D shows that Continental European single-digit slopes remain near $+3$ across raw headcount weights, Purchasing Power Standard earnings, public-administration exclusions, and 24-cell all-firm-size subsamples, and sit at $+2.510$ under fixed American ISCO-1 employment weights.

Payroll overlap diagnostics confirm how widely coworker wage dispersion cuts across national hourly pay cutoffs in both America and Britain. Occupations that span at least one national quintile cutoff employ 98.0 per cent of US OEWS employees (98.02 per cent under restricted $[0.95p_{10}, 1.15p_{90}]$ tails) and 99.7 per cent of matched UK ASHE employee jobs. Roughly half of each workforce (50.8 and 52.2 per cent) sits in occupations that directly straddle the Q1–Q2 threshold. Switching from linear within-occupation tilt weights $\lambda_q^{\text{lin}}(\rho) = 1 + (\rho - 1)(q - 1)/4$ to geometric weights $\lambda_q^{\text{geo}}(\rho) = \rho^{(q-1)/4}$ raises the top-to-bottom break-even tilt needed to erase the bottom-fifth gain to 15.39 in America (7.02 linear under restricted tails) and 17.53 in Britain.

E. Complete 20-Workflow Request-Mode Bounds and Cross-Platform Concordance

Table A3 reports all twenty standardised request workflows across US and UK consumer conversations alongside US feedback-loop revisions and global commercial first-party API bounds. Before worst-case suppressed traffic enters Table 5, retained-support commercial API pure directive bounds sit even tighter at $[82.9, 96.6]$ per cent across twenty workflows and $[83.9, 94.0]$ per cent across 860 fine clusters. Across individual workflows in Table A3, household-heavy and open-ended consumer groups carry wide $[0, 100]$ consumer intervals. On the commercial API, unrounded pure directive lower bounds exceed 50 per cent in fourteen of the twenty workflows, exceed 60 per cent in thirteen, and exceed 84 per cent in ten. That gap confirms that one-shot directive tasks sort sharply into programmatic interfaces.

Table A4 and Figure A2 test cross-platform task alignment across all seven focal Western economies without ad hoc zero pseudocounts. Across 25 to 33 shared O*NET Generalized Work Activities, work-filtered Claude shares match work-related ChatGPT shares closely, with Spearman rank correlations from 0.769 to 0.842 in Q2 2026 and 0.776 to 0.845 in Q3 2026. Across 102 to 136 finer Intermediate Work Activities, correlations run from 0.517 to 0.605 in Q2 and 0.533 to 0.609 in Q3. Stripping non-work traffic reduces cross-platform dissimilarity by 3.54 to 7.60 percentage points across the seven economies. Within the 33 strictly positive shared US GWAs in Q2 2026, the unadjusted log work-share ratio $\log(s_{\text{Claude}}/s_{\text{ChatGPT}})$ rises with pre-AI structured desk intensity R_g (Spearman $\rho = +0.438$, $p = 0.0108$), consistent with heavier structured desk traffic on Claude.

TABLE A2. Two-Part Audits, Crosswalk Weighting Diagnostics, and Wage Sensitivity Across US, UK, and EU-5 Panels

Specification, crosswalk or two-part audit	N	WLS β (SE)	PPML β (SE)	Level β	Level LP $[L, U]$
Panel A. Two-Part Extensive and Active-Intensive Margin Decomposition					
US extensive margin LPM ($\mathbf{1}\{s_i > 0\}$) + SOC2 FE	830	+0.087 (0.081)	–	–	–
US extensive margin LPM ($\mathbf{1}\{s_i > 0\}$) + Joint controls	830	−0.163 (0.111)	–	–	–
US active intensive margin ($s_i > 0$), unconditional	437	+2.267 (0.599)	–	–	–
US active intensive margin ($s_i > 0$) + SOC2 FE	437	+1.383 (0.416)	–	–	–
US active intensive margin ($s_i > 0$) + Joint controls	437	+1.237 (0.469)	–	–	–
UK active intensive margin ($s_i > 0$), unconditional	373	+1.663 (0.753)	–	–	–
UK active intensive margin ($s_i > 0$) + SOC2 FE	373	+1.864 (0.593)	–	–	–
UK active intensive margin ($s_i > 0$) + Joint controls	373	+0.877 (0.610)	–	–	–
Panel B. United States Six-Digit OEWS Censoring and Wage Diagnostics					
Felten AIOE capability only (no FE)	830	+0.717 (0.403)	+0.594 (0.377)	+0.895	[−0.38, 1.99]
Eloundou GPT capability only (no FE)	830	+0.987 (0.465)	+0.666 (0.360)	+0.412	[−0.97, 1.67]
Sub-threshold censoring ($< 0.005\%$ cap) + Joint	830	+0.555 (0.472)	−0.138 (0.423)	+0.006	[−3.64, 3.77]
GWA work prior + SOC2 FE	830	+1.569 (0.525)	+0.842 (0.391)	+0.873	[−2.12, 3.69]
Constant SOC2 work share + SOC2 FE	830	+1.346 (0.503)	+0.668 (0.380)	+0.697	[−2.12, 3.69]
Hourly-reported SOC6s ($N = 772$) + SOC2 FE	772	+1.824 (0.465)	+0.983 (0.396)	+1.028	[−2.19, 3.92]
Excl. SOC 25 Educators ($N = 766$) + SOC2 FE	766	+1.659 (0.445)	+0.935 (0.414)	+0.965	[−1.98, 3.68]
FTE-1440 wage basis + SOC2 FE	830	+1.325 (0.472)	+0.592 (0.453)	+0.623	[−2.30, 3.43]
Strict un-consolidated SOC6 + SOC2 FE	830	+1.706 (0.447)	+0.886 (0.371)	+0.919	–
Log annual mean pay ($N = 825$) + SOC2 FE	825	+1.453 (0.353)	+0.728 (0.348)	+0.762	[−2.06, 3.30]
Panel C. United Kingdom Four-Digit SOC2020 Crosswalk and Wage Diagnostics					
Felten AIOE capability only (no FE)	398	−0.422 (0.573)	−0.002 (0.572)	+0.088	[−1.14, 1.17]
Eloundou GPT capability only (no FE)	399	−0.185 (0.293)	−0.042 (0.268)	−0.325	[−1.39, 0.62]
Equal-split crosswalk (w_{norm}) + SOC2 FE	402	+1.615 (0.611)	+0.238 (0.607)	+0.220	[−1.13, 1.60]
Coarse 9 SOC1 FE (emp.-weighted xwalk)	402	+1.424 (0.466)	+0.948 (0.405)	+1.003	[−0.38, 2.16]
Coarse 9 SOC1 FE + Equal-split (w_{norm})	402	+0.890 (0.516)	−0.082 (0.435)	−0.081	[−1.36, 1.22]
GWA work prior + SOC2 FE	402	+1.802 (0.559)	+1.197 (0.542)	+1.169	[−0.11, 2.37]
Full-time hourly wage + SOC2 FE	389	+1.946 (0.502)	+1.480 (0.537)	+1.394	[0.07, 2.50]
Exact SOC2020 crosswalk + SOC2 FE	402	+1.780 (0.756)	+1.552 (0.534)	+1.639	–
Exact SOC2020 crosswalk + Joint controls	398	+0.659 (0.809)	+0.966 (0.427)	+1.003	–
Panel D. Continental European Single-Digit ISCO-08 Sensitivity					
Raw headcount e_{ci} weights + Country FE	40	+3.263 (0.749)	+2.087 (0.597)	+2.397	–
Purchasing Power Standard (PPS) + FE	40	+3.284 (0.727)	+2.087 (0.597)	+2.402	–
SES NACE B–S excl. Public Admin. (O)	40	+3.186 (0.706)	+2.055 (0.581)	+2.347	–
All-firm-size wage (DEU, ESP, NLD; $N = 24$)	24	+3.339 (0.543)	+2.346 (0.560)	+2.402	–
All-firm-size + R, S ($N = 24$)	24	+2.999 (0.415)	+2.458 (0.434)	+2.278	–
Fixed US employment weights + R, S	40	+2.510 (0.802)	+1.923 (0.623)	+1.754	–

Notes: Sensitivity and diagnostic specifications correspond to Table 3. Panel A decomposes US and UK unconditional and conditional slopes into US extensive-margin LPM ($\mathbf{1}\{s_i > 0\}$, $N = 830$) and US/UK active intensive-margin log WLS ($s_i > 0$ with complete regressors; $N = 437$ US, $N = 373$ UK). Panel B reports standalone US capability-only models, the sub-threshold censoring envelope ($< 0.005\%$ rounding cap across censored and rounded titles), and alternative priors and wage bases. Panel C reports standalone UK capability-only models and compares baseline UK APS employment-weighted crosswalks (w_{emp}) with 26 SOC2 fixed effects against equal-split crosswalks (w_{norm}), coarse 9-group SOC1 fixed effects, full-time hourly wages, and exact SOC2020 crosswalks. Panel D examines EU-5 raw headcount weights, PPS earnings, NACE Section O exclusions, 24-cell all-firm-size records (DEU, ESP, NLD), and fixed US ISCO-1 employment weights.

TABLE A3. Complete 20-Workflow Request-Mode Bounds Across US/UK Consumer Chat, Feedback Loops, and Commercial 1P API

	US margins (%)			Consumer channel bounds				1P API
Standardised workflow (all 20)	Share	Work	Dir.	US $[L, U]$	$\psi_g^\star$	UK $[L, U]$	+ loops	Pure $[L, U]$
Sign-identified above 50% pure directive in US consumer channel ($\theta_g^L > 50\%$)								
Conversation & Meeting Intel.	0.37	92.2	66.2	[63.4, 75.5]	0.00	[62.3, 76.1]	[71.5, 83.6]	[92.9, 99.1]
Business Process & Operations	6.26	85.5	61.1	[54.5, 73.6]	0.00	[52.6, 72.6]	[71.9, 91.0]	[90.4, 98.2]
Interior Fréchet intervals spanning 50% in US consumer channel ($\theta_g^L \leq 50\% \leq \theta_g^U$)								
Document Processing & Extraction	3.01	60.1	49.3	[15.7, 85.7]	0.96	[5.4, 93.6]	[25.5, 95.5]	[95.6, 100.0]
DevOps & Infra. Ops.	2.75	66.2	36.0	[3.4, 57.0]	6.66	[0.0, 48.0]	[72.1, 100.0]	[70.6, 100.0]
Companionship & General Conversation	1.60	3.1	14.3	[0.0, 100.0]	1.91	[0.0, 100.0]	[0.0, 100.0]	[0.0, 100.0]
Customer Support & Service Ops.	0.52	35.5	32.4	[0.0, 100.0]	2.43	[0.0, 100.0]	[39.5, 100.0]	[60.0, 94.9]
Education & Learning	9.28	15.6	26.9	[0.0, 100.0]	3.27	[0.0, 100.0]	[0.0, 100.0]	[24.6, 100.0]
Existential, Relational, and Emotional Support	4.50	8.5	7.5	[0.0, 100.0]	6.74	[0.0, 100.0]	[0.0, 100.0]	[0.0, 100.0]
Hobbies & Lifestyle	13.59	2.6	28.7	[0.0, 100.0]	2.38	[0.0, 100.0]	[0.0, 100.0]	[0.0, 100.0]
Knowledge Retrieval & Enterprise Search	3.97	25.4	38.6	[0.0, 100.0]	1.59	[0.0, 100.0]	[0.0, 100.0]	[76.5, 100.0]
Other / Unclear	0.33	10.0	18.4	[0.0, 100.0]	0.85	[0.0, 100.0]	[0.0, 100.0]	[0.0, 100.0]
Personal AI Assistant	4.19	11.1	49.8	[0.0, 100.0]	0.96	[0.0, 100.0]	[0.0, 100.0]	[25.3, 100.0]
Research & Intelligence	13.46	31.7	31.7	[0.0, 100.0]	2.98	[0.0, 97.0]	[0.0, 100.0]	[84.7, 100.0]
Trust & Safety / Platform Integrity	0.08	29.5	36.5	[0.0, 100.0]	2.03	[0.0, 100.0]	[0.0, 100.0]	[92.0, 97.3]
Data Analysis & BI	3.76	51.1	31.5	[0.0, 64.2]	5.89	[0.0, 58.3]	[9.1, 100.0]	[88.8, 100.0]
Sign-identified below 50% pure directive in US consumer channel ($\theta_g^U < 50\%$)								
Sales & Revenue Operations	1.62	97.8	47.3	[46.1, 49.7]	∞	[45.1, 48.0]	[62.1, 65.7]	[90.7, 94.6]
Cybersecurity & Threat Detection	0.27	58.4	26.5	[0.0, 49.0]	∞	[0.0, 44.2]	[36.2, 100.0]	[87.9, 93.5]
Content Creation & Copywriting	19.99	67.2	28.6	[0.0, 44.1]	∞	[0.0, 42.6]	[1.4, 51.5]	[87.7, 100.0]
Software Development	8.13	65.8	19.2	[0.0, 32.0]	∞	[0.0, 29.7]	[46.9, 100.0]	[76.4, 89.1]
Compliance & Regulatory	0.96	75.7	20.5	[0.0, 29.2]	∞	[0.0, 29.6]	[16.3, 50.6]	[84.3, 86.4]

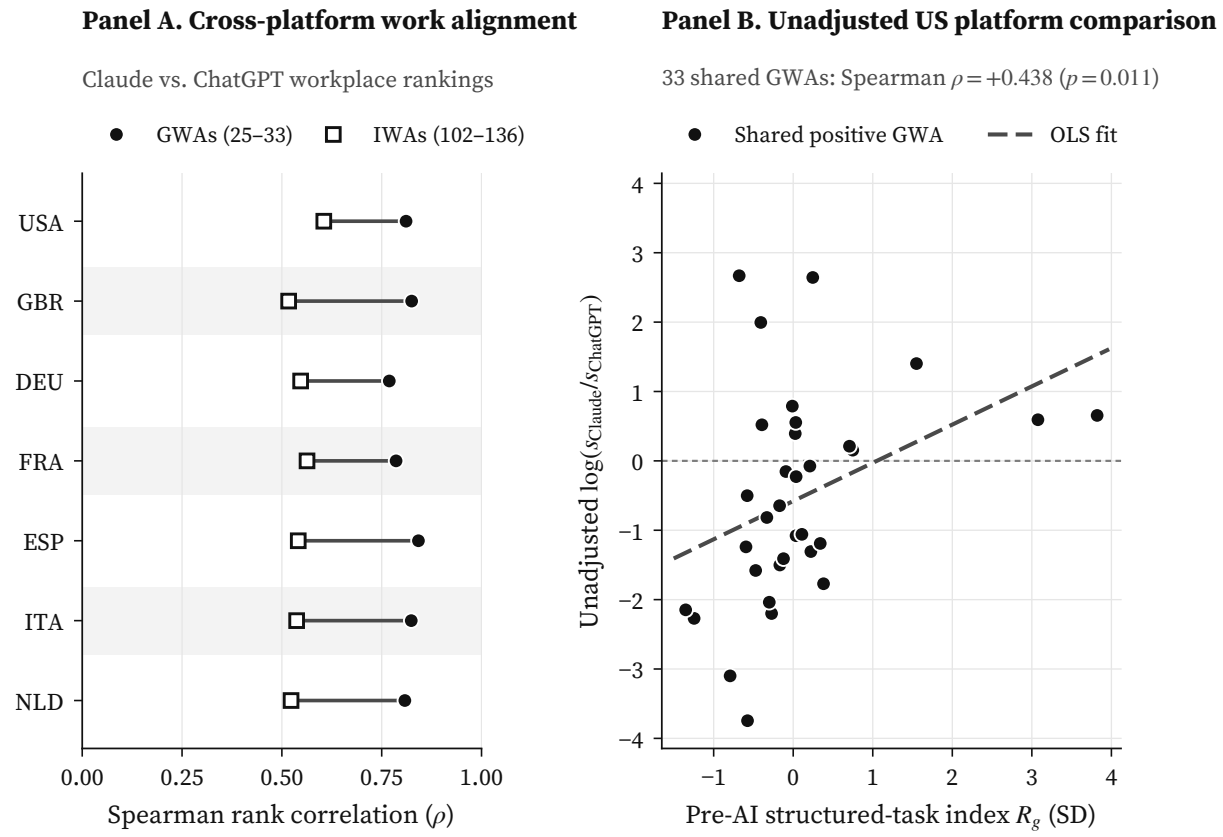
Notes: Complete 20-workflow decomposition corresponding to Table 5 Panel B, grouped by US consumer-channel parity regime. Columns report US consumer traffic share, marginal US workplace and pure directive shares (%), US and UK consumer-channel pure directive Fréchet bounds $[\theta_g^L, \theta_g^U]$, US conservative parity-breakdown index $\psi_g^\star$ defined after Equation (10), US consumer bounds with feedback loops added, and global commercial 1P API pure directive bounds. Consumer-channel columns pool Claude.ai chat and Cowork.

TABLE A4. GWA Structured-Desk Benchmarks and Cross-Platform Work-Task Concordance Across Seven Economies

	Shared cells		Structured desk R_g		GWA Spearman ρ		IWA Spearman ρ	
Economy (filter Δ pp)	GWAs	IWAs	β_R^{indep}	Fréchet $[\underline{\beta}_R, \overline{\beta}_R]$	Q2	Q3	Q2	Q3
United States (USA; 7.05 pp)	33	136	+1.69	[−18.3, +15.6]	0.811	0.810	0.605	0.609
United Kingdom (GBR; 6.60 pp)	31	112	+1.77	[−17.5, +15.0]	0.825	0.786	0.517	0.533
Germany (DEU; 6.53 pp)	28	114	+1.26	[−17.6, +14.3]	0.769	0.776	0.547	0.551
France (FRA; 4.67 pp)	29	119	+0.71	[−19.7, +14.4]	0.786	0.776	0.563	0.565
Spain (ESP; 3.54 pp)	29	108	+0.86	[−16.0, +13.0]	0.842	0.822	0.541	0.572
Italy (ITA; 4.87 pp)	27	103	+0.42	[−16.0, +12.0]	0.824	0.818	0.537	0.591
Netherlands (NLD; 7.60 pp)	25	102	+0.85	[−16.7, +12.9]	0.808	0.845	0.523	0.570

Notes: Columns report shared O*NET Generalized Work Activities (GWAs) and Intermediate Work Activities (IWAs) across all seven focal Western economies (with percentage-point dissimilarity reductions after non-work traffic is removed in parentheses), workplace-weighted WLS slopes of Anthropic consumer GWA pure directive rates on standardised pre-AI structured desk intensity R_g under conditional independence (β_R^{indep}) and Fréchet bounds ($[\underline{\beta}_R, \overline{\beta}_R]$), and cross-platform Spearman rank correlations across Q2 and Q3 2026 OpenAI waves.

FIGURE A2. Cross-Platform Workplace Task Alignment Across Seven Economies and Unadjusted US Structured-Task Comparison



Notes: Panel A compares cross-platform Spearman rank correlations between work-filtered Claude and work-related ChatGPT shares across shared Generalized Work Activities (25–33 GWAs, filled black circles) and Intermediate Work Activities (102–136 IWAs, hollow squares) in all seven focal Western economies. Panel B scatters the unadjusted log ratio $\log(s_{\text{Claude}}/s_{\text{ChatGPT}})$ across shared strictly positive US GWAs against standardised pre-AI structured desk intensity R_g (Spearman $\rho = +0.438$, $p = 0.0108$).

F. Google Gemini ATLAS Triangulation and Common-Support Diagnostics

Google's AI and Economy ATLAS version 1.0 (Iscenko et al. 2026) provides an independent public CC-BY 4.0 aggregate release built from 14,653,926 de-identified interactions sampled between 6 and 19 April 2026. Clusters require at least ten distinct users and apply contribution limits, differential-privacy noise, and cell suppression. Geographic intensity tables (218 countries and territories plus 51 US jurisdictions) and household time-use tables pool conversational Gemini app and Google Search AI Mode interactions and exclude API traffic. Occupational tables combine those conversational surfaces with content-eligible free-tier Gemini API and AI Studio calls across 335 US-reported five-digit Broad SOC occupations and twenty-two Major groups. Those occupational files omit military occupations (55-0000), detailed "All Other" residuals, paid Gemini API prompt

text, Vertex AI, and Google Workspace. When we match Google's 335 Broad codes to civilian May 2025 BLS OEWS wages, we retain 334 Broad occupations and drop only code 11-1030 Legislators, who lack a published median hourly wage. That sample links 148 multi-title Broad codes directly to OEWS Broad wage rows and 186 single-constituent Broad codes to their sole detailed OEWS title. Together, those matched codes cover 143.62 million employees (92.4 per cent of national OEWS employment and 99.93 per cent of published US Gemini workplace traffic).

Table A5 and Figure A3 assemble the full Google Gemini verification checks referenced in Sections 3 to 5. Across the 121 economies observed in all three releases, Panel A verifies that the thirty-percentage-point platform ordering holds unchanged (88.2 per cent non-work on Gemini, 69.8 per cent on ChatGPT, and 59.2 per cent on Claude). Panel A also breaks down US non-work requests across American Time Use Survey domains and off-hours shares. In Panel B, we compare US Gemini and aggregated US Claude on the exact same 334 Broad SOC support with t_{21} confidence intervals, alongside the 148-code direct OEWS Broad subsample, 430-code non-residual and 454-code full-OEWS zero-imputation checks, and all twenty-two leave-one-family-out runs. For internal wage bands, Panel C reports the eligible non-residual OEWS pay-grid distribution H_q, quantile-grid bounds, and analytical bounds across every eligible detailed-title assignment within each Broad code (confirming a strictly positive paired bottom-fifth gain $\Delta H_1 \in [+2.15, +6.12]$ percentage points). Panel D summarises published US Major-SOC full-automation request and non-routine cognitive rates across the seventeen reporting occupation groups.

TABLE A5. Google Gemini ATLAS Common-Economy, Common-Support Wage, Pay-Grid, and Task-Intent Diagnostics

Panel A. Three-Platform Common-Economy Check ($N = 121$) and US Gemini ATUS Domains

Measure or US ATUS domain	Share (%)	Off-hrs (%)	Benchmark	Maj. non-work	Scope
Common-economy: Gemini	88.15	–	ChatGPT 69.80%	121 / 121	Search & App
Common-economy: Claude	59.18	–	3-platform overlap	113 / 121	Claude.ai + Cowork
US ATUS: Leisure & socialising	32.01	68.8	Global 31.40%	–	US non-work
US ATUS: Education & homework	22.53	50.2	Global 26.37%	–	US non-work
US ATUS: Household & cooking	15.62	67.7	Global 13.27%	–	US non-work
US ATUS: Purchases & services	13.63	65.9	Global 11.49%	–	US non-work
US ATUS: Personal care & health	8.95	70.0	Global 10.19%	–	US non-work

Panel B. Common-Support US Occupational Pay-Slope Diagnostics (22 SOC2 Clusters; t_{21} 95% CIs)

Sample and estimator	N	Uncond.	+ SOC2 FE	+ Joint	Joint 95% t_{21} CI
Gemini Broad SOCs (PPML)	334	+1.361 (0.420)	+0.862 (0.296)	+0.337 (0.309)	[−0.31, +0.98]
Claude on 334 Broad (PPML)	334	+1.428 (0.407)	+1.144 (0.323)	+0.197 (0.314)	[−0.46, +0.85]
Gemini Broad SOCs (WLS)	334	+2.667 (0.670)	+1.707 (0.566)	+0.880 (0.443)	[−0.04, +1.80]
Claude on 334 Broad (WLS)	334	+2.741 (0.623)	+1.497 (0.389)	+0.570 (0.488)	[−0.45, +1.58]
Gemini direct Broad (PPML)	148	+1.599 (0.408)	+1.396 (0.359)	+0.790 (0.348)	[+0.07, +1.51]
Gemini direct Broad (WLS)	148	+3.451 (0.675)	+2.283 (0.732)	+1.022 (0.430)	[+0.13, +1.92]
Gemini non-resid. zeros (PPML)	430	+1.432 (0.417)	+0.887 (0.300)	+0.468 (0.280)	[−0.11, +1.05]
Gemini all Broad zeros (PPML)	454	+1.439 (0.421)	+0.883 (0.300)	+0.458 (0.345)	[−0.26, +1.18]
Gemini 22-fold LOFO (PPML)	297–327	[+1.24, +1.52]	[+0.64, +0.97]	[+0.16, +0.46]	22 / 22 attenuate

Panel C. Eligible-Title US OEWS Pay-Grid Allocation Across National Hourly Pay Fifths

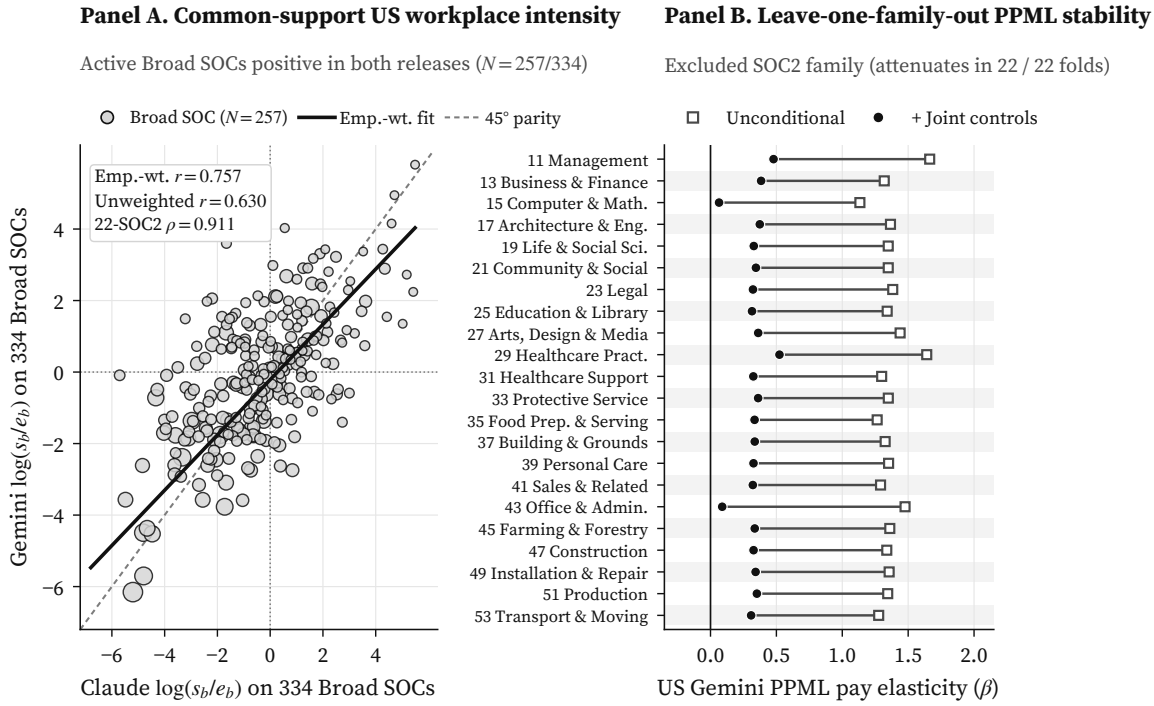
Distributional measure	Q1 (Bot.)	Q2	Q3	Q4	Q5 (Top)
Gemini grid H_q (%)	4.82	9.54	16.46	27.34	41.84
Quantile-grid $[L, U]$ (%)	[2.94, 14.49]	[1.31, 20.58]	[4.50, 31.96]	[7.03, 42.83]	[32.58, 52.09]
Within-Broad H_1 & summary	[3.06, 6.79]	ΔH_1: [+2.15, +6.12] pp		$\rho^\star = 14.21\times$	$B_5 = 64.0\%$, $M = 1.40$

Panel D. Published US Major-SOC Task-Intent Rates (17 Reporting Groups)

Task-intent measure	Mgmt.	Comp. & Math	Admin. Supp.	Legal	Summary
Full-automation request (%)	1.68	20.43	35.65	2.04	$r(R_g) = +0.72$
Non-routine cognitive (%)	55.9	87.9	6.3	94.3	Health 37.4%

Notes: Panel A compares non-work medians across the 121 economies observed in all three releases and reports US Gemini ATUS non-work domains and weekday-night/weekend shares (Iscenko et al. 2026). Panel B reports employment-weighted PPML and shifted-log WLS log-pay slopes (22 SOC2 clusters, t_{21} 95% CIs) on the common 334-code Broad SOC sample (143.62M employees) for Gemini and aggregated US Claude, on the 148 direct OEWS Broad codes (82.0M employees), under zero imputation across 430 non-residual (152.2M) and 454 total OEWS Broad codes (155.5M), and across 22 leave-one-SOC2-out (LOFO) folds. Panel C reports eligible non-residual OEWS pay-grid traffic shares H_q, quantile-grid intervals, exact min–max bounds across eligible detailed titles inside each Broad SOC, break-even tilt $\rho^\star$, and top-fifth wage-weighted share B_5 with mean-wage multiple M. Panel D reports published US Major-SOC full-automation request and non-routine cognitive rates across the 17 reporting groups.

FIGURE A3. Common-Support Workplace AI Intensity Concordance and Leave-One-Family-Out Stability on Google Gemini



Notes: Panel A scatters US Gemini log relative workplace intensity $\log(s_b/e_b)$ against aggregated US Claude log relative workplace intensity across the 257 Broad SOCs active in both releases out of 334 common Broad codes (circle area proportional to May 2025 OEWS employment; employment-weighted $r = 0.76$). Panel B plots unconditional and joint-control PPML pay elasticities across all twenty-two leave-one-SOC2-family-out jackknife subsamples on US Gemini.

References for the Online Appendix

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40 (121): 13–58.

Appel, Ruth, Maxim Massenkoff, Peter McCrory, Miles McCain, Ryan Heller, Tyler Neylon, and Alex Tamkin. 2026. "Anthropic Economic Index Report: Economic Primitives." research report, Anthropic. https://www.anthropic.com/research/anthropic-economic-index-january-2026-report.

Autor, David H., and Michael J. Handel. 2013. "Putting Tasks to the Test: Human Capital, Job Tasks, and Wages." *Journal of Labor Economics* 31 (S1): S59–S96.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): 1279–1333.

Autor, David, and Neil Thompson. 2025. "Expertise." *Journal of the European Economic Association* 23 (4): 1203–1271.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–427.

Charnes, Abraham, and William W. Cooper. 1962. "Programming with Linear Fractional Functionals." *Naval Research Logistics Quarterly* 9 (3-4): 181–186.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with Zeros? Some Problems and Solutions." *Quarterly Journal of Economics* 139 (2): 891–936.

Cross, Philip J., and Charles F. Manski. 2002. "Regressions, Short and Long." *Econometrica* 70 (1): 357–368.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–1308.

Eurostat. 2025a. "Employment by Detailed Occupations (ISCO-08, Dataset lfsa_egai2d)." European Union Labour Force Survey. https://ec.europa.eu/eurostat/databrowser/view/lfsa_egai2d/.

Eurostat. 2025b. "Structure of Earnings Survey 2022: Mean Hourly Earnings by Sex, Occupation and Economic Activity (earn_ses22_47)." Dissemination API Dataset. https://ec.europa.eu/eurostat/databrowser/view/earn_ses22_47/.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12): 2195–2217.

Fréchet, Maurice. 1951. "Sur les tableaux de corrélation dont les marges sont données." *Annales de l'Université de Lyon, Section A: Sciences Mathématiques et Astronomie* 14: 53–77.

Iscenko, Zanna, Scott Strand, Yiyuan Chen, Guillaume Aimard, Mihai Codreanu, Vivek Sampathkumar, Alex Imas, Julian Jacobs, Evalyne Muiruri, Jia Jen Ng, Samirah Javed, Juan Mateos-Garcia, Denis Calin, Andrew Kim, Josh Martin, Omar Ajmeri, Fabien Curto Millet, and James Manyika. 2026. "Google's AI & Economy ATLAS v1.0: Mapping Gemini Usage in the Economy." Technical Report arXiv:2608.00038, Google and Google DeepMind. https://arxiv.org/abs/2608.00038.

Jones, Charles I. 2026. "A.I. and Our Economic Future." *Journal of Economic Perspectives* 40 (3): 3–22.

Kansoy, Fatih. 2026. "AI Use and Occupations: Task-Weighted International Crosswalks and Distributional Benchmarks." research dataset and working paper, University of Oxford. https://fatih.ai/ai-use-and-occupations/.

Manski, Charles F. 2003. *Partial Identification of Probability Distributions*. New York: Springer.

Massenkoff, Maxim, Eva Lyubich, Szymon Sacher, Zoe Hitzig, Shaoyi Zhang, Ryan Heller, and Peter McCrory. 2026. "Anthropic Economic Index Report: Cadences." research report, Anthropic. https://www.anthropic.com/research/economic-index-june-2026-report.

Masten, Matthew A., and Alexandre Poirier. 2020. "Inference on Breakdown Frontiers." *Quantitative Economics* 11 (1): 41–111.

Office for National Statistics. 2025. "Annual Population Survey (APS) Employment by SOC2020 Occupation and ITL1 Region, Jan–Dec 2025." Nomis Dataset NM_218_1. https://www.nomisweb.co.uk.

Office for National Statistics. 2026. "Employee Earnings in the UK: ASHE Table 14 (4-Digit SOC2020), 2026 Provisional and 2025 Revised." Dataset under Open Government Licence v3.0. https://www.ons.gov.uk.

OpenAI. 2026. "OpenAI Signals: Version 3.0 Release and Data Dictionary (July 2024–September 2026)." Dataset under CC BY 4.0. https://cdn.openai.com/signals/data-download-csv.zip.

Santos Silva, J. M. C., and Silvana Tenreyro. 2006. "The Log of Gravity." *Review of Economics and Statistics* 88 (4): 641–658.

U.S. Bureau of Labor Statistics. 2026. "Occupational Employment and Wage Statistics (OEWS), May 2025 National Cross-Industry Estimates." Time-Series Release oe.data.0.Current. https://download.bls.gov/pub/time.series/oe/.