Early Signs of the Generative AI Labor Market Shock in the Global South: Evidence from Philippine Business Process Outsourcing*

Nico Ravanilla†

July 22, 2026

Abstract

Business process outsourcing is among the developing world's largest employers, and generative AI performs the tasks it sells. I examine the Philippines, the economy most dependent on it. I construct a sub-national exposure index from pre-ChatGPT occupational composition and observed usage in the Anthropic Economic Index, and trace employment as US clients adopt AI and automate what they bought. Employment growth slows in exposed provinces, industries, and municipalities, while wages and hours remain unchanged. Would-be workers remain unemployed, not yet absorbed at home or abroad. Displacement is beginning at the hiring margin — early, measurable, and concentrated where exposure is highest.

JEL codes: J23, O33, F16

Keywords: artificial intelligence, labor demand, occupational exposure, outsourcing

*Replication files construct all measures from public sources: the Anthropic Economic Index (Anthropic, 2026), Philippine Labor Force Survey public use files (Philippine Statistics Authority, 2025), the registries of the Philippine Economic Zone Authority (PEZA) (Philippine Economic Zone Authority, 2021), Philippine Presidential Proclamations (Official Gazette of the Republic of the Philippines, 2021), the Department of Migrant Workers' deployment statistics (Department of Migrant Workers, 2026), and the US Census Bureau's Business Trends and Outlook Survey (US Census Bureau, 2025).

†Associate Professor, School of Global Policy and Strategy, University of California San Diego. nravanilla@ucsd.edu

1 Introduction

For labor-abundant countries, exporting services has become the development path that manufacturing once was. The Philippines and its peers sell customer service, records processing, and back-office operations to rich-country firms, and a recent line of work argues that these newly tradable services can carry the employment growth that factories no longer provide (Nayyar et al., 2021; Baldwin and Forslid, 2023; Rajan and Lamba, 2024; Rodrik and Sandhu, 2025). Generative AI puts that path in question, because the tasks these exports consist of are the tasks the technology performs (Korinek and Stiglitz, 2021). The IMF puts global exposure at about 40 percent of employment, concentrated in the cognitive and clerical work that developing-country service exporters sell (Cazzaniga et al., 2024). Acemoglu and Restrepo warn that AI directed at automating tasks rather than creating them delivers displacement without much productivity growth (Acemoglu and Restrepo, 2019, 2020), and macroeconomic accounting in the same framework puts the aggregate US productivity gains from generative AI below one percent over ten years (Acemoglu, 2025). Workplace evidence is only now arriving, on both sides of this ledger. On the productivity side, the largest deployment study ran on customer support agents mostly in the Philippines and found gains of 15 percent, concentrated among novices (Brynjolfsson et al., 2025b). On the displacement side, the first estimates come from high-income economies: US payroll records (Brynjolfsson et al., 2025a), Danish administrative data (Humlum and Vestergaard, 2025), online labor platforms (Hui et al., 2024), and asset prices (Borri et al., 2026). Whether displacement has begun in the countries that sell the exposed work is an empirical question, and the natural place to look is the economy most dependent on that work.

That economy is the Philippines. The country has long exported workers on a large scale, and its business process outsourcing (BPO) industry, which employs 1.9 million workers serving mostly US clients (IT and Business Process Association of the Philippines, 2025), exports the work instead of the workers. India's technology industry is larger in absolute terms, at 5.4 million workers (NASSCOM, 2024), but those workers are below one percent of India's workforce (Government of India, 2024), against about four percent in the Philippines, so dependence is greatest there. In 2016 the International Labour Organization estimated that 89 percent of Philippine BPO workers held jobs at high risk of automation (Chang and Huynh, 2016). In the years that followed, the industry kept growing, adding several hundred thousand workers since the prediction, and is now the largest source of formal private employment in the country. The interesting question is not whether the predicted catastrophe arrived, because so far it has not. The more useful question is whether the prediction was early rather than wrong, now that generative AI can actually perform the tasks these jobs consist of, and how one would detect its arrival early enough for the answer to matter.

This paper builds a detection apparatus for that question and reports its first readings. The measurement strategy has two parts. The first is a sub-national AI exposure index for the Philippines. To construct it, I combine each province's occupational composition before the release of ChatGPT in November 2022, the start of mass-market generative AI, measured from labor force survey microdata by province of workplace, with occupation-level measures of *observed* AI usage from the Anthropic Economic Index (Handa et al., 2025; Massenkoff and McCrory, 2026). Whereas the existing exposure literature relies on predictions of what AI could do (Felten et al., 2021; Webb, 2020; Eloundou et al., 2024; Hampole et al., 2025), including a mapping of potential exposure for the Philippines itself (Cucio and Hennig, 2025), the weights I use record what AI is actually used to do. By this measure, customer service representatives, the core occupation of Philippine BPO, rank second in observed exposure among roughly 800 US occupations, with data entry keyers and medical records specialists close behind (Massenkoff and McCrory, 2026). The index locates the country's dependence on these occupations across the 87 province-level areas of the Labor Force Survey's workplace coding, which I refer to as provinces throughout.

The second part is a measure of the treatment itself. Generative AI did not arrive on a single date: capability accumulated across successive model generations, and adoption diffused gradually into the firms that buy Philippine services. I therefore measure the shock as a dose, the share of US firms using AI in production, from the US Census Bureau's Business Trends and Outlook Survey (Bonney et al., 2024). This share rises from zero before late 2022 to more than seven percent of firms by early 2025. Because BPO demand is overwhelmingly American, US adoption is the relevant dose, external to Philippine provinces by construction. The estimating equation interacts predetermined provincial exposure with lagged US adoption while controlling for exposure-specific linear trends, so that identification comes from the curvature of the adoption path rather than from smooth differential trends in exposed places.

The mechanism that connects these two parts is simple. US firms from banks and insurers to hospitals and retailers buy customer service, data entry, and records processing from Philippine providers. When such a firm adopts generative AI, the tasks it automates are the tasks it used to buy, so demand for new Philippine seats grows more slowly. Providers respond as a high-turnover industry can, freezing hiring and letting attrition shrink headcount within a quarter or two. The result is displacement that arrives as foregone hiring in the provinces, industries, and municipalities where the exposed occupations concentrate. The displaced can then shift toward less exposed work within the sector, look elsewhere in the domestic labor market, or go abroad, and the paper follows each margin. The index measures where the pressure lands, and the adoption series measures when and how strongly it arrives.

The analysis yields five main findings. First, employment responds to the dose in every design that I examine. As US adoption rises by one point, a province that is one standard deviation more exposed loses 0.15 percentage points of BPO employment share and 3.7 percent of BPO employment relative to less exposed ones. The same gradient appears within provinces across industries, in a specification where province-by-round fixed effects absorb every province-level shock, and it appears again across municipalities according to the intensity of their IT-zone infrastructure.

Second, the timing of the response is the timing of hiring rather than of firing. The response peaks at lags of one to two quarters and decays after that, as one would expect in an industry where attrition runs 30 to 40 percent a year, so a hiring freeze shrinks headcount within months without any contract being broken.[1] Consistent with this mechanism, wages and hours do not respond at all. The industry is adjusting by not hiring rather than by cutting pay, the same entry-margin pattern found in US payroll data (Brynjolfsson et al.,

[1] The Contact Center Association of the Philippines reports voluntary attrition of 31 to 36 percent in 2021–2022 and total attrition up to 45 percent (Contact Center Association of the Philippines, 2024). Earlier estimates range from 19 to 47 percent (Errighi et al., 2016).

2025a).

Third, the magnitudes are already economically meaningful. Between 2021 and mid-2025, BPO employment in the most exposed tercile of provinces grew by 42 percent, while the least exposed tercile grew by 67 percent. Had the most exposed tercile matched that pace, it would hold a quarter of a million more workers today, and cumulating the dose coefficient over the observed adoption path accounts for the whole of this gap.

Fourth, the paper traces where the displaced adjust. Within the sector, the occupation mix rotates away from the most exposed work, consistent with incumbents reassigned while hiring closes on entrants. Outside it, the would-be workers are so far absorbed by unemployment rather than by other employers: employment in other sectors of exposed provinces shows no response at lags of up to a year, the regional unemployment rate rises with the dose, and participation does not fall.

Fifth, the earliest overseas evidence is inconclusive but indicative. The stock of overseas workers, long the country's outlet for labor that the domestic economy cannot absorb, has not responded, but administrative records show first-time departures beginning to rotate upward in exposed regions. The paper states the argument these flows bear on: generative AI closes the channel through which workers earned foreign wages at home, and takes a first step toward testing whether the displaced reroute abroad.

The paper contributes a measurement strategy that detects the labor market effects of AI before they appear in aggregate statistics. The strategy crosses observed usage with sub-national labor microdata to establish where exposure sits, then doses the index with client-country adoption to establish when and how strongly the pressure arrives. Applied to the Philippines, the strategy extends the realized-effects evidence, so far confined to high-income economies, to a developing labor market, with the timing of the effects and the margin of adjustment. The remainder of the paper proceeds as follows. Section 2 builds the index, Section 3 states the empirical framework, Section 4 reports the results, and Section 5 states what the apparatus predicts next.

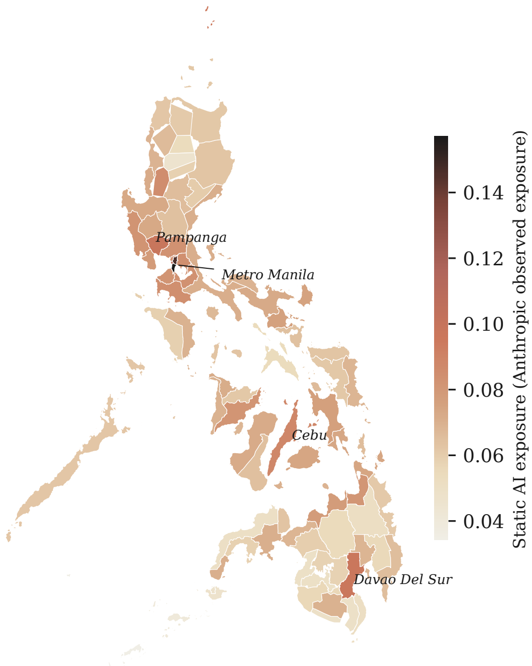

(a) Provincial exposure

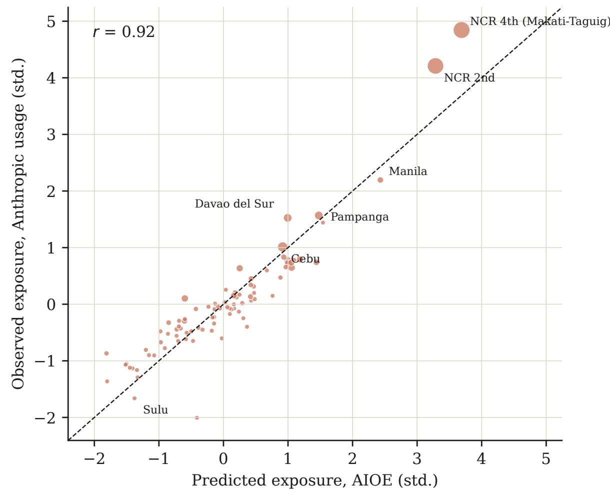

(b) Observed vs. predicted exposure

Figure 1: The provincial AI exposure index. Panel A: the static index, combining pre-ChatGPT (2021–2022) occupational employment by province of workplace (Labor Force Survey) with occupation-level observed AI exposure (Anthropic Economic Index). Panel B: the same index against a capability-based alternative that replaces usage weights with the AIOE of Felten et al. (2021). Both axes are standardized. Marker area is proportional to pre-period BPO employment. The dashed line is the 45-degree line.

2 An Observed-Usage Exposure Index for the Philippines

The index is a standard shift-share exposure measure with a nonstandard ingredient. For province p,

$$\text{Exposure}_p = \sum_o s_{op}\, e_o, \tag{1}$$

where s_{op} is occupation o's share of employment among jobs located in province p, measured over the seven quarterly Labor Force Survey (LFS) rounds before the release of ChatGPT (2021Q2–2022Q4), and e_o is occupation o's observed AI exposure. The shares come from the

LFS public use files (Philippine Statistics Authority, 2025). Less well known is that these files record not only each employed person's occupation (2-digit PSOC) but also the province and municipality of the person's workplace, which means that the index measures the exposure of jobs located in a province rather than of workers residing there. The weights e_o come from the Anthropic Economic Index's labor market release (Massenkoff and McCrory, 2026), which scores occupations by how much of their work AI is observed performing in millions of actual usage records, with the scoring concentrated on automation-type use in work contexts. I crosswalk these scores from US occupation codes to the Philippine classification through ISCO-08. Supplemental Appendix Section S1 documents the crosswalk, whose coverage exceeds 99 percent of employment, together with an industry-level version of the index that is built from each industry's national occupation mix and that will carry the within-province analysis below. By that industry measure, business-support services, the division that contains the call centers, is the most exposed of all 86 Philippine industries.

Using these weights for Philippine provinces requires one assumption. The relative ordering of occupations by AI substitutability must transfer from the usage data, which are global, concentrated in high-income countries, and organized on the US task taxonomy, to Philippine workplaces. In this setting the assumption is mild, because the occupations at the top of the ranking belong to an export industry whose jobs are American processes performed offshore: a customer-service job in Manila serves the same clients, follows the same scripts, and runs the same software as its US counterpart, often within the same firm. The transfer is weakest for non-traded occupations further down the ranking, where Philippine task content may depart from the US template. Measurement error of that kind attenuates the exposure gradient and works against the findings below. Panel B of Figure 1 provides the corresponding empirical check, since reweighting the index with an entirely different, capability-based scoring of the same occupations leaves the provincial ranking almost unchanged.

Panel A of Figure 1 maps the index. Exposure concentrates where the BPO industry sits, namely in the Metro Manila districts, led by the Makati–Taguig corridor, in the Clark corridor in Pampanga, and in the regional hubs of Cebu, Davao, Iloilo, and Bacolod, while the agricultural and maritime periphery is lightly exposed. Panel B addresses the concern that this geography is an artifact of the usage-based weights. When I replace e_o with the capability-based AIOE of Felten et al. (2021) and rebuild the index (Supplemental Appendix Section S8), the resulting ranking is nearly identical ($r = 0.92$). The geography of exposure is thus a fact about where clerical and customer-service work is done, and it is robust to how the occupation weights are constructed. The informative deviation is at the top of the distribution, where the usage-based index ranks the major hubs above the 45-degree line, so realized usage is more concentrated in their occupations than capability alone would predict.

Two further checks speak to whether the index measures what it should. The first is that the industry-level version, which uses no geographic information at all, ranks office administrative and business support services, the division that contains the call centers, as the most exposed of all 86 Philippine industries, and it places gambling and betting second, because the Philippine online-gaming industry staffs the same customer-service and back-office occupations (the full ranking is in Supplemental Appendix Table S7). The second is that the highly exposed provinces outside Metro Manila are the documented regional hubs of the industry, including the cities that the government's own digital-cities program designated as the sector's next locations, not an arbitrary collection of urban areas.

Panel A of Figure 2 shows the first descriptive result that the index delivers. I divide provinces into employment terciles by exposure and track BPO employment in each group. The three groups grow at almost the same rate through 2022, during the sector's broad post-pandemic boom, and only then do they separate. By the first half of 2025 the least and middle exposed terciles stand at 167 and 164 percent of their 2021 levels, while the most exposed tercile stands at 142. The divergence begins between 2022 and 2023 and widens each year after that. Panel B plots the AI adoption of US client firms on the same timeline, and the rest of the paper tests whether that adoption explains the divergence above it.

3 Empirical Framework

The estimating equation is

$$y_{pt} \;=\; \alpha_p + \gamma_t + \beta \left(\text{Exposure}_p \times D_{t-k}\right) + \delta \left(\text{Exposure}_p \times t\right) + \varepsilon_{pt}, \tag{2}$$

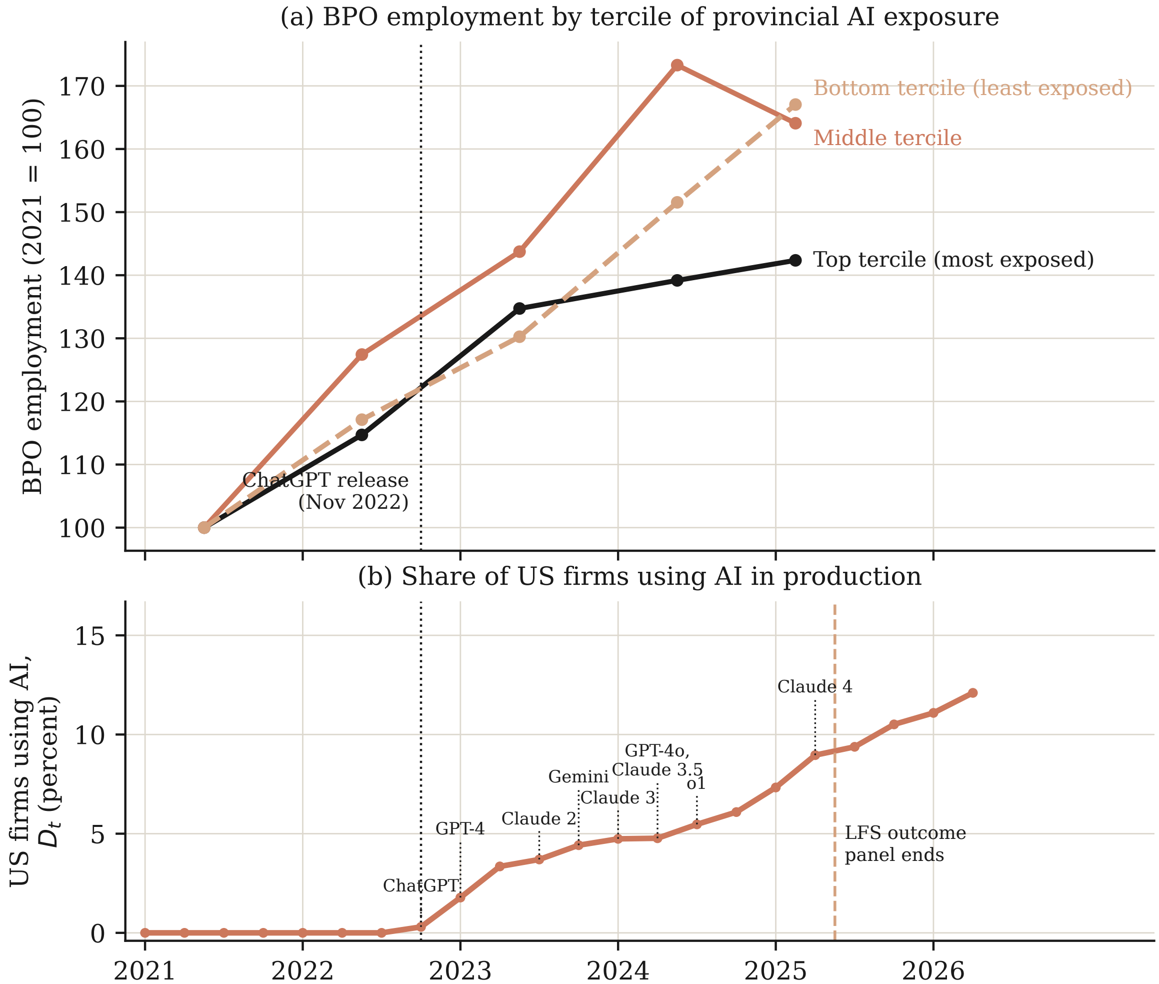

Figure 2: BPO employment by exposure tercile and the US adoption dose, on a shared timeline. Panel A: BPO covers PSIC divisions 62, 63, and 82, by province of workplace. Annual values average the quarterly LFS rounds (two rounds for 2025) and sit at year midpoints. Terciles split pre-period employment into thirds by the static index. Each series is indexed to its 2021 average. Panel B: D_t is the share of US firms using AI in production, from the Business Trends and Outlook Survey, with the early quarters spliced as described in Supplemental Appendix Table S3. Dotted markers indicate the release dates of major model generations. The dashed line marks the end of the LFS outcome panel in 2025Q2.

where y_{pt} is a BPO employment outcome for province p in quarterly LFS round t (2021Q2–2025Q2), Exposure_p is the index of equation (1), standardized to mean zero and unit variance across provinces weighted by pre-period employment, and D_t is the share of US firms using AI in production. The series is the biweekly Business Trends and Outlook Survey aggregated to quarters (Bonney et al., 2024). I splice the first three post-release quarters with paymentplatform and web-traffic data and set the series to zero before 2022Q4 (construction in Supplemental Appendix Table S3). Alternative splices and client-weighted doses, by sector and by geography, leave the estimates unchanged (Tables S1 and S2). Worker-side surveys record the same rapid diffusion (Bick et al., 2024).

Panel B of Figure 2 plots the series below the tercile employment paths, together with the release dates of the major model generations, and it shows why the paper treats the shock as a dose. Capability arrived in many steps, from ChatGPT in late 2022 through GPT-4, Claude 2, Gemini, Claude 3, and their successors, and no single one of these releases is the treatment. Adoption, in contrast, accumulated smoothly through the whole sequence, and it is this accumulation, rather than any one release date, that the estimating equation asks employment to track. Adoption was still accelerating when the outcome panel ends, which gives the paper's predictions their bite.

Each element of equation (2) does one job. US adoption, rather than Philippine usage, measures the shock because BPO demand is American and arrives through clients' decisions, and no Philippine province can move aggregate US adoption, since Philippine usage is below one percent of the global total. The two sources play separate roles, and neither is measured from Philippine outcomes: the Anthropic data supply the cross-sectional ranking and never enter the time series, while the adoption survey supplies the time variation and never enters the cross section. Adoption enters as a continuous dose rather than a post-ChatGPT indicator because capability accumulated across model generations and diffused gradually into client firms. The exposure-specific linear trend ensures that β is identified by the curvature of the adoption path, so that any confound that trends smoothly in more exposed places, such as urbanization, sectoral drift, or the maturation of the industry, loads onto δ rather than onto β. Adoption enters at a lag k because it reaches Philippine payrolls with delay, through hiring freezes, attrition, and contract renewals. The baseline uses $k = 1$, and the results trace the full lag profile.

The identifying assumption at the province level is that, conditional on exposure-specific linear trends, no omitted shock tracks the shape of the adoption curve differentially across exposure, and two further designs relax exactly this assumption. In the first design, I estimate equation (2) across industries within provinces, replacing Exposure_p with the industry-level index and including province-by-round and province-by-industry fixed effects. This specification absorbs every shock at the province level, including local demand and weather, so identification comes only from more versus less exposed industries inside the same province and quarter. Shocks specific to the exposed industries themselves, such as the 2023 return-to-office mandate for tax-zone firms, are not absorbed by construction. Supplemental Appendix Section S4 addresses the mandate: its step-like path does not track the adoption curve, its late-2024 relaxation coincides with growing rather than shrinking effects, and it gives no reason for the within-sector occupation rotation reported below. In the second design, I estimate the equation across municipalities, interacting the count of a municipality's IT-zone Presidential Proclamations through 2021 with the dose and including municipality and province-by-round fixed effects. Supplemental Appendix Table S4 tests the identifying assumption directly in the pre-release rounds, with a joint curvature test and a placebo dose placed where none existed. The within-province and municipal designs pass both tests, while the province-level gradient does not, because the pre-release rounds span the reopening from pandemic restrictions that bound hardest in the capital. The province-level estimates therefore carry the descriptive weight, and the within-province and municipal designs carry the identification. Standard errors are clustered at the treated unit throughout. The regressions are weighted by pre-period employment, fixed before the release. Weighting makes the estimates averages over workers rather than areas, the object the magnitude calculations require, and downweights the smallest, noisiest survey cells. The industry design also reports the unweighted estimate, because a gap between the two would signal heterogeneous effects (Solon et al., 2015), and the two agree. Shock-level inference in the sense of Adão et al. (2019), implemented through the equivalence of Borusyak et al. (2022), yields smaller standard errors than the ones I report (Supplemental Appendix Table S5).

Table 1: Employment responses to AI adoption across designs

	(1) BPO share	(2) log BPO	(3) log emp	(4) log emp	(5) log BPO	(6) log pay	(7) log hours
Exposure × adoption (lag 1)	-0.1495	-0.0366				0.0015	0.0021
	(0.0275)	(0.0119)				(0.0040)	(0.0019)
Exposure × adoption (lag 1)			-0.0060	-0.0123			
			(0.0025)	(0.0027)			
Exposure × adoption (lag 1)					-0.0595		
					(0.0258)		
Observations	1,479	1,312	67,290	67,832	1,170	594	594
R^2	0.977	0.953	0.960	0.902	0.894	0.655	0.378

Notes: The reported regressor is exposure interacted with US AI adoption lagged one quarter. All columns include exposure-by-linear-trend interactions and unit and time fixed effects. Columns 1–2: province panel (87 provinces, 17 rounds), weighted by pre-period employment, standard errors clustered by province. Columns 3–4: province-by-industry panel (86 industries) with province-by-round and province-by-industry fixed effects. Column 3 is weighted by pre-period cell employment. Standard errors are clustered by province-industry. Column 5: municipality panel with IT-zone intensity (log of one plus pre-2022 IT-zone proclamations) as exposure, municipality and province-by-round fixed effects, standard errors clustered by municipality, unweighted (weighted and clustering variants in Supplemental Appendix Table S6). Columns 6–7: province-level BPO daily pay and weekly hours (cells with at least ten sampled BPO workers). Standard errors in parentheses.

4 Results

4.1 Employment

Table 1 contains the paper's central estimates. Columns 1 and 2 report the province-level gradient. A province that is one standard deviation more exposed loses 0.149 percentage points of BPO employment share (SE 0.027) and 3.7 percent of BPO employment (SE 1.2) for each point of lagged US adoption, conditional on exposure-specific trends. These columns describe where the slowdown sits and how steeply it scales with exposure, and by the pre-release tests in Supplemental Appendix Table S4 they partly reflect the capital's reopening path, so the causal weight falls on the designs that follow. Within provinces, an industry that is one standard deviation more exposed loses 0.6 percent of employment per adoption point in the weighted specification and 1.2 percent in the unweighted one (columns 3 and 4). These estimates are identified only from comparisons across industries inside the same province and quarter, where province-by-round fixed effects absorb the reopening cycle and every other province-level shock. Municipalities tell the same story at the finest grain: employment in municipalities with IT-zone infrastructure falls relative to same-province neighbors as adoption accumulates (column 5, variants in Supplemental Appendix Table S6). Across three geographic designs with different fixed-effect structures, the sign is the same and the estimates are distinguishable from zero.

The magnitudes are economically meaningful, and Panel A of Figure 2 provides a natural benchmark for them. A province at the 75th percentile of exposure and one at the 25th are separated by 0.71 standard deviations. The column 2 coefficient, cumulated over the observed adoption path, implies that the more exposed province holds 19 log points less BPO employment by mid-2025 through the adoption channel alone. Scaled up to the tercile comparison, the dose term accounts for the entire 16-log-point growth gap between the most and least exposed terciles, and the point estimates, if anything, overshoot it, so the magnitudes are approximate. Part of the overshoot likely reflects the pre-release reopening curvature documented in Supplemental Appendix Table S4, which pushes the province-level coefficients toward larger negative values. The design that passes the pre-release tests delivers the same arithmetic: the BPO divisions sit 3.6 standard deviations above the employment-weighted mean industry, so the within-province coefficient cumulates to a 16-log-point shortfall for those divisions by mid-2025, the same size as the tercile gap. In levels, had the most exposed tercile grown at the least exposed tercile's rate from 2021 onward, it would hold roughly 240,000 more BPO workers today. That is about an eighth of the sector, and it takes the form of foregone growth rather than layoffs.

4.2 Timing

If the estimates reflect the transmission of client adoption to Philippine hiring, the response should arrive with a lag of a quarter or two rather than contemporaneously, because the industry's structure makes slowed hiring the cheap adjustment margin on both sides of the contract. On the commercial side, outsourcing agreements price capacity in seats and let clients scale volumes down without terminating the relationship, so an adopting client requests fewer seats at the next adjustment. On the labor side, Philippine law makes dismissals costly, because declaring positions redundant requires advance notice and separation pay,

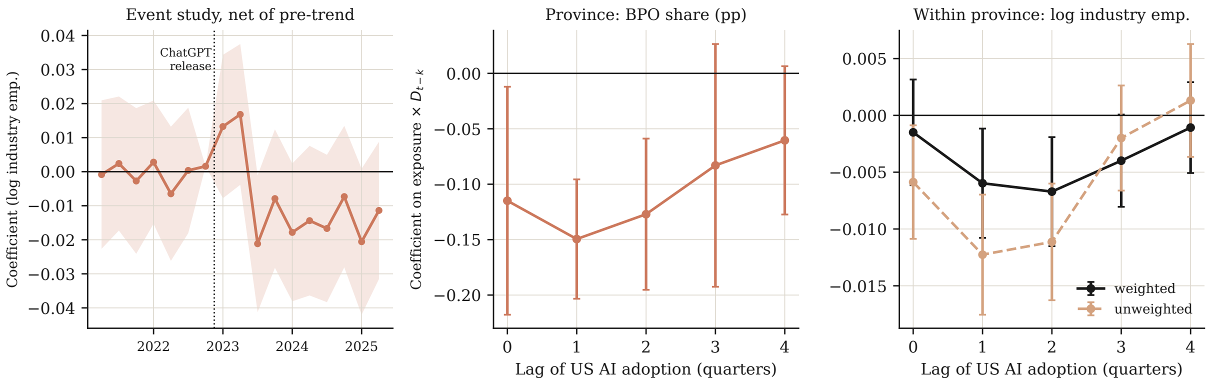

Figure 3: Dynamics of the employment response. Panel (a): within-province industry event-study coefficients, net of a linear trend fitted to the six pre-release points (specification of Table 1, column 3). Panels (b) and (c): each point is a separate regression with US adoption at the indicated lag, specifications of Table 1, columns 1 and 3–4. Bars: 95 percent confidence intervals. Panel (a) deviations are individually imprecise, and the powered test is the dose specification of Table 1.

while voluntary attrition of 30 to 40 percent a year costs the provider nothing.[2] A provider that stops backfilling sheds headcount within months without dismissing anyone. Panels (b) and (c) trace the coefficient at lags of zero to four quarters, and the pattern matches this prediction. The response peaks at one to two quarters in every design and decays by lag four. The province coefficient is −0.115 contemporaneously and −0.150 at a one-quarter lag, and the weighted within-province estimate turns on at lags one and two. The profile is also difficult to reconcile with confounding, because a smooth confound trending with adoption would not produce a peaked and decaying lag pattern. Panel (a) of Figure 3 shows the pattern quarter by quarter, net of trend. The detrended coefficients sit at zero before the release, stay near zero in the first two quarters after it, when lagged adoption was still close to zero, and turn negative in nearly every quarter from mid-2023 onward. Individual quarters are imprecise, because a trend extrapolated from six pre-period points carries wide uncertainty, so the powered test remains the dose specification of Table 1. In levels the exposed industries keep growing, so the pattern is deceleration rather than reversal.

[2]Labor Code of the Philippines, articles 296 to 298: probationary employment is capped at six months, and redundancy requires thirty days' written notice and separation pay of at least one month per year of service.

4.3 The Hiring Margin

Columns 6 and 7 of Table 1 apply the same specification to BPO daily pay and weekly hours, and neither outcome responds. The pay coefficient of 0.001 (SE 0.004) rules out wage responses larger than one percent per adoption point. Foregone entry also removes junior workers from the mix, raising observed pay, so the bound understates any true decline. The sector is adjusting through the quantity of new hiring rather than the price of incumbent labor. US payroll data locate the early effects on the same margin, among would-be entrants rather than incumbents (Brynjolfsson et al., 2025a). For a workforce as young as this one, foregone entry is where displacement begins.

4.4 Where the Displaced Workers Go

The displacement estimates raise the question of where the displaced adjust, and the first margin lies inside the sector. Within the BPO divisions, the occupation mix rotates away from high-exposure occupations as adoption rises, by about 7 percent per exposure standard deviation per adoption point (Supplemental Appendix Table S10), consistent with incumbents reassigned while hiring closes on entrants, and the aggregate estimates net the two. For the workers the sector no longer takes, the LFS follows the remaining margins: other sectors, unemployment and participation, and overseas work. A person's industry of employment is recorded at the province of the workplace, while unemployment, participation, and overseas membership belong to the residence record, which the files code at the region level. I estimate the sector margin across the 87 provinces and the remaining margins across the 17 regions, with exposure replaced by its employment-weighted regional average. The aggregate background rules out one competing reading: the national unemployment rate fell from 8.8 percent in early 2021 to 4.1 percent by mid-2025 while participation held between 60 and 65 percent, so exposed provinces are not running out of job seekers.

Panel (a) of Figure 4 applies the specification of equation (2) to log employment in nine destination sectors, BPO divisions excluded. No sector absorbs the displaced workers. Every coefficient in the panel lies within 0.02 of zero. The only positive estimate that is distinguishable from zero, in trade, is 0.008 (SE 0.003), fades by a one-year lag, and does

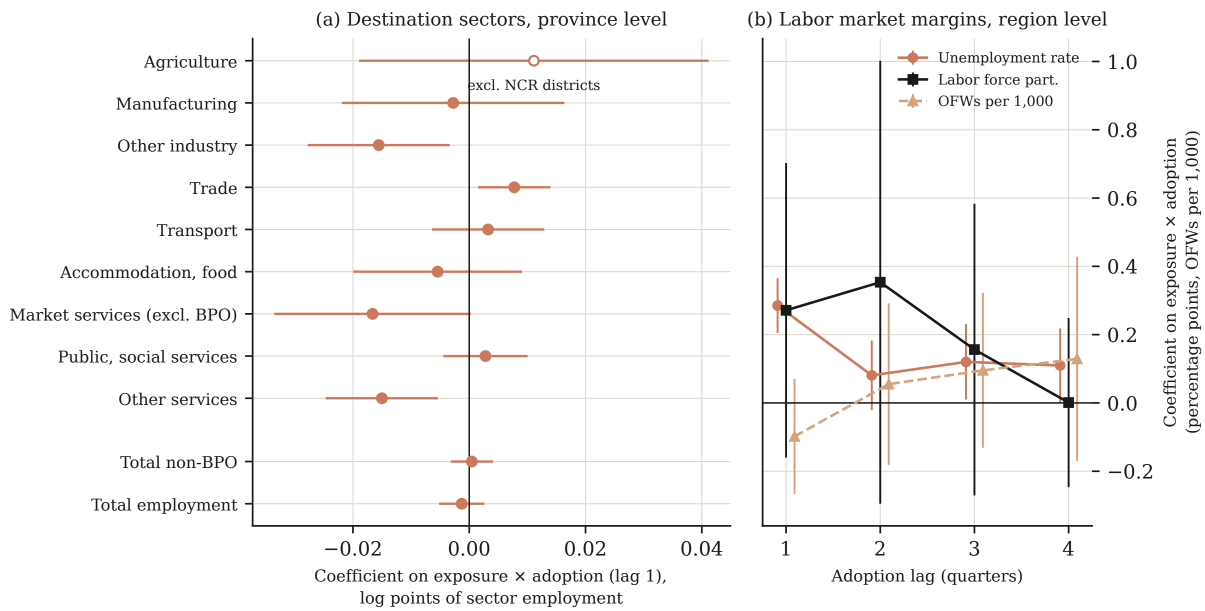

Figure 4: Where the adjustment goes. Panel (a): coefficients on exposure interacted with US adoption at a one-quarter lag, from equation (2) estimated on log employment in each destination sector across all 87 province-level areas, with BPO divisions 62, 63, and 82 excluded throughout. The agriculture estimate (open marker) instead excludes the four NCR districts, whose small agricultural cells otherwise dominate it. Both samples for every sector are in Supplemental Appendix Table S8. Panel (b): the same specification across the 17 regions, replacing provincial exposure with its employment-weighted regional average, for the unemployment rate and the labor force participation rate in percentage points and for overseas Filipino workers per 1,000 residents, with adoption entering at lags of one to four quarters. Bars: 95 percent confidence intervals. Standard errors clustered by province in Panel (a) and by region in Panel (b).

not survive excluding the NCR districts. The remaining distinguishable estimates are small negatives, in other industry and other services, consistent with mild local spillovers rather than absorption. The agriculture row's full-sample estimate, 0.117 (SE 0.024), comes from the four NCR districts' few-hundred-worker agricultural cells, so the panel plots the estimate excluding them, 0.011 (SE 0.015). Total employment summarizes the sector margin in one number. BPO holds about 4 percent of employment in the panel, so if every foregone BPO job disappeared with no local absorption, total employment would move by −0.0015 log points per adoption point, and the estimate is −0.0013 (SE 0.0020). The precision cannot separate full absorption from none, but the point estimate sits where the no-absorption arithmetic puts it. Reemployment takes time, yet the sector coefficients stay near zero at lags of up to a year.

Panel (b) traces the remaining margins across adoption lags of one to four quarters. The unemployment rate in exposed regions rises with adoption, by 0.29 percentage points per adoption point per unit of regional exposure at a one-quarter lag (SE 0.04, wild cluster bootstrap $p = 0.013$), and the point estimates stay positive at every horizon out to a year, though longer lags are imprecise under the bootstrap. In the United States the same exposure measure shows hiring slowdowns for young workers without a detectable unemployment rise (Massenkoff and McCrory, 2026), so the queue in exposed Philippine regions is the developing-country face of the same shock. Participation does not fall at any lag, so the missing workers are not leaving the labor force. The unemployment magnitude matches the foregone-jobs arithmetic of the previous section. Cumulated over the observed adoption path, the coefficient implies that the National Capital Region, which stands about two units of regional exposure above the other regions, carries an unemployment rate four to five percentage points higher by mid-2025 than the other regions' exposure would imply, while assigning the capital roughly half of the quarter million foregone jobs against its labor force of about seven million implies close to two points. Both calculations put the response in the low single digits, and the estimate, if anything, overshoots the arithmetic as the employment estimates overshoot the tercile gap, so it too is approximate. Taken together, the domestic margins say that the adjustment has so far taken the form of waiting. The would-be BPO workers of the exposed regions appear as unemployed rather than in other sectors or out of the labor force, at least within the year that the panel can see. The full grid of estimates behind both panels is in Supplemental Appendix Table S8.

4.5 The Overseas Margin

For half a century the Philippines has adjusted to weak labor demand at home by sending workers abroad, and remittances insure households against local income shocks (Yang and Choi, 2007). The BPO industry grew as the alternative to that outlet, because the internet moved the work to the workers instead of moving the workers to the work (Baldwin and Forslid, 2023). The internet built that channel out of the tasks that can be delivered remotely (Blinder and Krueger, 2013), and the tasks that can be delivered remotely are also the tasks that a language model can perform. What generative AI removes, in every country that exports services, is therefore the channel that paid workers foreign wages at home for their actual skills. The option that remains requires physical presence, and it lies mostly in different work, in care, construction, and domestic service rather than in office tasks. Rerouting into that option is not automatic, because migration carries costs, visas are rationed, and the skill mismatch is itself a deterrent. Whether the displaced take it is an empirical question, and the answer arrives first in the Philippines, where migration frictions are lower than anywhere else. The present data allow a first reading, and I test the rerouting further in companion work as deployment records accumulate.

The first reading is a stock that has not moved and flows that have begun to stir. The stock of overseas workers in exposed regions shows no response at lags up to a year (Panel (b) of Figure 4), a pattern the migration literature predicts, since departures respond to destination demand only with time (Theoharides, 2018). Administrative deployment records see the flows that stocks cannot. In the Department of Migrant Workers' deployment counts by region of origin (Department of Migrant Workers, 2026), first-time deployments rise with adoption in exposed regions, by 1.4 percent per adoption point (SE 0.6), while rehires fall by 0.8 percent (SE 0.2) and total deployment does not move, a rotation toward new migrants (Supplemental Appendix Table S9). With sixteen regional clusters, neither flow estimate survives wild-bootstrap inference, so the reading is a margin beginning to stir rather than one that has moved.

The rerouting, if it comes, is not a neutral swap, because virtual migration paid part of the cross-border wage premium while workers stayed home in work that used their education, and its welfare and distributional consequences, for the workers who leave, for the families who stay, and for a country that loses the work but regains the remittances, are open questions that the deployment data will begin to answer. Three findings therefore hold at once: the sector's occupation mix shifts toward less exposed work, the would-be entrants queue as unemployment, and the overseas flows show the first signs of movement where it would appear first. The policy implication, which follows from the mechanism rather than from the point estimates, is in Supplemental Appendix Section S9.

5 Conclusion

Exported services were to be the development path that manufacturing once was, a way for labor-abundant countries to sell labor to the world without moving it. That path has met a technology that performs the work it runs on. The first signs of displacement are now measurable in the Philippines, the economy that depends on the path most. Even at this early stage, when the sector as a whole is still growing and no wave of layoffs has occurred, the data show a growth stall: BPO employment keeps rising but rises systematically less where exposure is highest. The stall tracks the accumulation of AI adoption among US client firms, with the short lag that hiring freezes and attrition imply, and it leaves wages and hours untouched. The pattern also reconciles the two sides of the evidence: a technology that raises what each agent produces reduces the new agents a client needs unless service demand expands in proportion, so the same adoption appears as productivity gains in deployment studies and slowed hiring in labor market data. Between 2021 and mid-2025 the stall amounts to roughly a quarter of a million foregone jobs in the most exposed provinces.

The estimates make predictions that new data can check. US adoption is still rising while the trend terms are not, so each new LFS round will separate the dose model from its alternatives. If the mechanism is foregone entry, the effects should reach older workers only slowly. The estimates here concern the directly exposed sector. The study of import competition from China followed a similar path: the first estimates showed that exposed local labor markets lost manufacturing employment, and later work traced where the affected workers went, into other sectors, out of the labor force, and onto public transfers (Autor et al., 2013). The first look at those margins finds the sector shifting its mix toward less exposed work, the workers in the unemployment rows of the exposed regions, and overseas flows that have only begun to stir. Robots ran this experiment for goods, reaching the labor markets of offshore manufacturing through reshoring (Faber, 2020). Generative AI is running it for services, and whether the displaced reroute into the physical channel will show first in deployment flows from the exposed regions, where the rotation toward first-time migrants already visible in the administrative records is what a beginning would look like. What made early detection possible here was nothing unusual. Observed usage data were crossed with sub-national labor microdata and dosed with client-country adoption, and both ingredients now exist for most countries. For economies whose comparative advantage sits in AI-exposed services, the most practical conclusion of this paper is that such measurement can and should run continuously, before the aggregate statistics move.

References

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58.

Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, 13(1):25–35.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Anthropic (2026). Anthropic economic index data releases. Hugging Face dataset repository. Accessed July 2026.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–2168.

Baldwin, R. and Forslid, R. (2023). Globotics and development: When manufacturing is jobless, and services are tradeable. *World Trade Review*, 22(3–4):302–311.

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative AI. NBER Working Paper 32966, National Bureau of Economic Research.

Blinder, A. S. and Krueger, A. B. (2013). Alternative measures of offshorability: A survey approach. *Journal of Labor Economics*, 31(S1):S97–S128.

Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L. S., Goldschlag, N., Haltiwanger, J. C., Kroff, Z., and Savage, K. (2024). Tracking firm use of AI in real time: A snapshot from the business trends and outlook survey. Working paper 32319, NBER.

Borri, N., Liu, Y., and Tsyvinski, A. (2026). AI premium. Working Paper 35451, National Bureau of Economic Research.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Brynjolfsson, E., Chandar, B., and Chen, R. (2025a). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Working paper, Stanford Digital Economy Lab. Version of November 13, 2025.

Brynjolfsson, E., Li, D., and Raymond, L. (2025b). Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942.

Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., and Tavares, M. M. (2024). Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note SDN/2024/001, International Monetary Fund.

Chang, J.-H. and Huynh, P. (2016). ASEAN in transformation: The future of jobs at risk of automation. Working Paper 9, International Labour Office, Bureau for Employers' Activities (ACT/EMP), Geneva.

Contact Center Association of the Philippines (2024). Annual member survey on industry attrition. As reported in A. M. Monzon, "Philippine contact center industry sees declining attrition rates," *Philippine Daily Inquirer*, May 28, 2024.

Cucio, M. and Hennig, T. (2025). Artificial intelligence and the philippine labor market: Mapping occupational exposure and complementarity. IMF Working Paper WP/25/43, International Monetary Fund.

Department of Migrant Workers (2026). Monthly compendium of OFW statistics: Deployed landbased workers by region of origin. Republic of the Philippines, Manila. Accessed July 2026.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308.

Errighi, L., Bodwell, C., and Khatiwada, S. (2016). Business process outsourcing in the philippines: Challenges for decent work. Ilo asia-pacific working paper series, International Labour Organization, Bangkok.

Faber, M. (2020). Robots and reshoring: Evidence from Mexican labor markets. *Journal of International Economics*, 127:103384.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

Government of India (2024). Economic survey 2023–24. Technical report, Ministry of Finance.

Hampole, M., Papanikolaou, D., Schmidt, L. D. W., and Seegmiller, B. (2025). Artificial intelligence and the labor market. NBER Working Paper 33509, National Bureau of Economic Research.

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., and Ganguli, D. (2025). Which economic tasks are performed with AI? evidence from millions of Claude conversations. arXiv preprint 2503.04761, arXiv.

Hui, X., Reshef, O., and Zhou, L. (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science*, 35(6):1977–1989.

Humlum, A. and Vestergaard, E. (2025). Still waters, rapid currents: Early labor market transformation under generative AI. NBER Working Paper 33777, National Bureau of Economic Research.

IT and Business Process Association of the Philippines (2025). Philippine IT-BPM industry caps 2024 with milestone revenue and employment. Press release, January 2025. `https://ibpap.org/news-room/21`.

Korinek, A. and Stiglitz, J. E. (2021). Artificial intelligence, globalization, and strategies for economic development. NBER Working Paper 28453, National Bureau of Economic Research.

Massenkoff, M. and McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. Economic research report, Anthropic. March 5, 2026. `https://www.anthropic.com/research/labor-market-impacts`.

NASSCOM (2024). Technology sector in india: Strategic review 2024. Technical report, National Association of Software and Service Companies.

Nayyar, G., Hallward-Driemeier, M., and Davies, E. (2021). *At Your Service? The Promise of Services-Led Development*. World Bank, Washington, DC.

Official Gazette of the Republic of the Philippines (2021). Presidential proclamations establishing information technology parks and centers, 1995–2021. Manila.

Philippine Economic Zone Authority (2021). List of operating economic zones. Pasay City.

Philippine Statistics Authority (2025). Labor force survey public use files, 2021–2025. PSA Data Archive, Quezon City. Accessed July 2026.

Rajan, R. and Lamba, R. (2024). *Breaking the Mold: India's Untraveled Path to Prosperity*. Princeton University Press, Princeton, NJ.

Rodrik, D. and Sandhu, R. (2025). Servicing development: Productive upgrading of labor-absorbing services in developing economies. *Global Policy*, 16(2):306–318.

Solon, G., Haider, S. J., and Wooldridge, J. M. (2015). What are we weighting for? *Journal of Human Resources*, 50(2):301–316.

Theoharides, C. (2018). Manila to Malaysia, Quezon to Qatar: International migration and its effects on origin-country human capital. *Journal of Human Resources*, 53(4):1022–1049.

US Census Bureau (2025). Business trends and outlook survey data releases. Washington, DC. Accessed July 2026.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Working Paper 3482150, SSRN.

Yang, D. and Choi, H. (2007). Are remittances insurance? Evidence from rainfall shocks in the Philippines. *World Bank Economic Review*, 21(2):219–248.

Supplemental Appendix

Early Signs of the Generative AI Labor Market Shock in the Global South:

Evidence from Philippine Business Process Outsourcing

Nico Ravanilla

July 21, 2026

Contents

S1 Crosswalk Construction and Coverage

The occupation weights originate on the 2018 US Standard Occupational Classification, and the Labor Force Survey codes occupations on the Philippine Standard Occupational Classification (PSOC 2012). I connect the two through ISCO-08, mapping each 2018 SOC occupation to its 2010 SOC predecessors and then to ISCO-08 unit groups using the official BLS crosswalks. Because a single SOC occupation can map to several ISCO unit groups, its OEWS May 2023 employment (US Bureau of Labor Statistics, 2023) is split equally across its mapped groups, and the exposure of an ISCO unit group is the employment-weighted mean of the exposure scores of its constituent occupations. The procedure covers 421 of the 436 ISCO-08 unit groups. Because the LFS public use files record PSOC at two digits, the analysis operates at the ISCO-08 sub-major level, where PSOC 2012 and ISCO-08 coincide. The 87 province-level areas of the LFS workplace coding are the 81 provinces of that classification, the four districts of Metro Manila, and the independent cities of Cotabato and Isabela. As the final check on coverage, the share of provincial employment matched to an exposure weight averages 100 percent across the 87 areas, with a minimum of 99.8 percent.

S2 Anthropic Economic Index: Release Harmonization

The Anthropic Economic Index has been released in successive waves whose schemas and classifiers differ, and the replication package harmonizes six usage windows spanning December 2024 to May 2026, separately for the consumer product (Claude.ai) and the first-party API. Usage records are matched to O*NET tasks in each window, and matching succeeds for essentially every task row. Within each window I renormalize usage shares to the classified total, since the share of unclassified usage varies across releases. Two caveats matter for anyone using the harmonized panel. First, there are level breaks between release families, which reflect classifier and taxonomy changes rather than changes in usage, so growth should be computed within a release family or absorbed with window effects. Second, the country-level data cover the consumer product only, while enterprise API usage, where automation-type use concentrates, is published at the global level. Neither caveat affects the main estimates in the paper, because the paper uses only the time-invariant observed-exposure weights from the labor market release, and the time variation comes from the adoption survey instead.

S3 The Adoption Dose: Construction and Alternatives

Table S3 reports the quarterly dose series and how each segment is constructed. The spine is the US Census Bureau's Business Trends and Outlook Survey, which asks firms biweekly whether they used AI in producing goods or services and which begins fielding that question in September 2023. Two adjustments extend the spine. Backward, I set the series to zero before 2022Q4 and backcast the three quarters from 2022Q4 to 2023Q2 by scaling the first BTOS reading with the growth of the Ramp AI adoption index (Ramp, 2025), cross-checked against ChatGPT web-traffic estimates and

Google Trends. Forward, the BTOS question changed in November 2025 from AI use in production to AI use in any business function, which produced a level jump from 9.9 to 17.3 percent, and I place the revised readings on the original basis using the ratio of the trend-extrapolated original series to the first revised observation, which equals 0.600. Three alternative constructions leave the conclusions of main-text Table 1 unchanged. The first sets the backcast segment to zero. The second uses the revised post-November-2025 readings without rescaling and absorbs the wording change with an indicator. The third ends the sample in 2025Q3, before the question change, which is also before the end of the LFS outcome panel used in the paper.

A fourth variant addresses the fact that BPO clients are not average US firms. The BTOS publishes adoption by sector, and the IT-BPM industry's client verticals concentrate in the sectors that adopted fastest. I therefore construct a client-weighted dose that averages sector-level BTOS adoption using the industry's global client-vertical mix from the IBPAP Roadmap 2028: banking, financial services, and insurance at 26 percent mapped to NAICS 52, technology, media, and telecom at 20 percent mapped to NAICS 51, retail at 10 percent mapped to NAICS 44, manufacturing at 11 percent mapped to NAICS 31, healthcare at 11 percent mapped to NAICS 62, energy at 6 percent split between NAICS 21 and 22, and the remaining 18 percent carrying the national series. The client-weighted dose runs about 40 percent above the national series at every point, because the client sectors adopt faster, and it correlates with the national series at 0.997, so the shape that identification relies on is unchanged. Table S1 re-estimates the four employment columns of main-text Table 1 with this dose at a one-quarter lag. The coefficients scale down in proportion to the higher dose level, as they must, and the test statistics are unchanged or stronger.

A fifth variant addresses client geography rather than client sector. North America accounts for about 70 percent of the Philippine industry's buyer base, with Europe including the UK at about 15 percent and APAC at about 15 percent, from the same IBPAP Roadmap 2028. Workers serving non-US clients received those countries' adoption paths rather than the US path, so the US-only dose overstates their treatment intensity and dilutes the estimated response toward zero. I construct a geography-weighted dose that averages three blocks by these shares. The North America block is the US series. The Europe block is the share of UK businesses using large language models for text generation from the ONS Business Insights and Conditions Survey, the only national series that measures generative AI use specifically. The APAC block is the mean post-2022 adoption increase in Japan's Communications Usage Trend Survey and Singapore's IMDA enterprise survey. Each foreign block is zero through 2022Q4, matching the splice convention of the main dose, and annual readings are interpolated to quarters. Foreign adoption tracks US adoption closely, and the geography-weighted dose correlates with the US-only series at 0.998. Table S2 re-estimates the four employment columns with this dose at a one-quarter lag. The coefficients are modestly larger than under the US-only dose, consistent with the foreign blocks rising somewhat faster late in the sample, and the test statistics are unchanged.

Table S1: Client-weighted dose: employment responses

	(1) BPO share	(2) log BPO	(3) log emp	(4) log emp, unwtd
Exposure × client-weighted adoption (lag 1)	-0.1050	-0.0255		
	(0.0190)	(0.0081)		
Exposure × client-weighted adoption (lag 1)			-0.0043	-0.0089
			(0.0017)	(0.0019)
Observations	1,479	1,312	67,290	67,832
R^2	0.977	0.953	0.960	0.902

Notes: Same specifications as main-text Table 1, columns 1 to 4, with the national adoption dose replaced by the client-weighted dose at a one-quarter lag. Columns 1 and 2 use the province panel with province and round fixed effects, weighted by pre-period employment, with standard errors clustered by province. Columns 3 and 4 use the province-by-industry panel with province-by-round and province-by-industry fixed effects, column 3 weighted by pre-period cell employment, with standard errors clustered by province-industry. Standard errors in parentheses.

Table S2: Client-geography-weighted dose: employment responses

	(1) BPO share	(2) log BPO	(3) log emp	(4) log emp, unwtd
Exposure × geography-weighted adoption (lag 1)	-0.1292 (0.0227)	-0.0321 (0.0112)		
Exposure × geography-weighted adoption (lag 1)			-0.0051 (0.0022)	-0.0091 (0.0024)
Observations	1,479	1,312	67,290	67,832
R^2	0.977	0.953	0.960	0.902

Notes: Same specifications as main-text Table 1, columns 1 to 4, with the national adoption dose replaced by the client-geography-weighted dose at a one-quarter lag. The dose weights the US series at 70 percent, a UK block at 15 percent, and a Japan-Singapore block at 15 percent, following the buyer-geography mix of the Philippine IT-BPM market in the IBPAP Roadmap 2028. Columns and inference as in Table S1. Standard errors in parentheses.

Table S3: The adoption dose series D_t, by quarter and construction segment

Quarter	D_t (% of US firms using AI)	Construction
2021Q1	0.00	assumed zero (pre-ChatGPT)
2021Q2	0.00	assumed zero (pre-ChatGPT)
2021Q3	0.00	assumed zero (pre-ChatGPT)
2021Q4	0.00	assumed zero (pre-ChatGPT)
2022Q1	0.00	assumed zero (pre-ChatGPT)
2022Q2	0.00	assumed zero (pre-ChatGPT)
2022Q3	0.00	assumed zero (pre-ChatGPT)
2022Q4	0.30	backcast: avg of Google Trends & Similarweb visit ratios to 2023Q1 x Ramp-based 2023Q1 value
2023Q1	1.78	backcast: BTOS 2023Q3 anchor x Ramp AI Index ratio
2023Q2	3.35	backcast: BTOS 2023Q3 anchor x Ramp AI Index ratio
2023Q3	3.70	BTOS original question (producing goods or services)
2023Q4	4.42	BTOS original question (producing goods or services)
2024Q1	4.74	BTOS original question (producing goods or services)
2024Q2	4.77	BTOS original question (producing goods or services)
2024Q3	5.47	BTOS original question (producing goods or services)
2024Q4	6.09	BTOS original question (producing goods or services)
2025Q1	7.33	BTOS original question (producing goods or services)
2025Q2	8.96	BTOS original question (producing goods or services)
2025Q3	9.38	BTOS original question (producing goods or services)
2025Q4	10.51	BTOS revised question x 0.600 wording-splice factor
2026Q1	11.09	BTOS revised question x 0.600 wording-splice factor
2026Q2	12.10	BTOS revised question x 0.600 wording-splice factor

Notes: D_t is the share of US firms using AI in production. The BTOS segment uses the US Census Bureau's biweekly survey aggregated to quarters. The backcast segment scales the first BTOS reading (3.7 percent, 2023Q3) by the growth of the Ramp AI adoption index. The rescaled segment places the revised BTOS question (in use from November 17, 2025) on the original basis using a factor of 0.600. The raw component series are included in the replication package.

S4 Tests of the Identifying Assumption

The estimates in the paper are identified by the curvature of the adoption path. The assumption is that, conditional on the exposure-by-linear-trend interactions, no omitted shock moves employment across exposure groups with the same curvature. This section tests the two implications of that assumption that the pre-release data can check. First, the pre-release exposure gradient should be linear in time, because any pre-existing curvature would be a shock of the prohibited kind operating before the technology existed. I test this by adding exposure interacted with squared and cubed time to the seven pre-release rounds (2021Q2 to 2022Q4) and testing the added terms jointly. Second, a dose-shaped regressor placed where no dose existed should attract nothing. I assign the realized adoption curve of the first seven post-release quarters to the seven pre-release rounds and estimate each design with its linear trend control, so the placebo coefficient is identified from the same kind of curvature as the real one.

Table S4 reports both tests next to each design's dose estimate. The two within-province designs pass both. In the industry design the curvature terms are jointly insignificant ($p = 0.79$) and the placebo dose attracts -0.002 (SE 0.005), a third of the real estimate and indistinguishable from zero, and in the municipal design the corresponding figures are $p = 0.90$ and a wrong-signed 0.042 (SE 0.047). The province-level rows do not pass. The BPO employment share shows pre-release curvature ($p = 0.001$), and the placebo dose attracts -0.138 (SE 0.037), which is comparable in size to the real estimate. The calendar explains why. The seven pre-release rounds span the reopening from pandemic restrictions, including renewed lockdowns in August 2021 and January 2022 that bound most tightly in the capital region, where exposure is highest, so the pre-release path of the exposed provinces is not linear, and a curvature-based estimator at the province level partly picks up the reopening cycle. The province-by-round fixed effects of the industry and municipal designs absorb the restriction cycle wholesale, and in those designs the pre-release period behaves as the identifying assumption requires. The province-level columns of main-text Table 1 should accordingly be read as the descriptive gradient, and the identification rests on the within-province and municipal designs.

Table S4: Tests of the identifying assumption, by design

	Dose estimate (lag 1)	Pre-release curvature (joint p-value)	Pre-release placebo dose
Province, BPO employment share	-0.1495 (0.0275)	0.001	-0.1380 (0.0369)
Province, log BPO employment	-0.0366 (0.0119)	0.088	-0.0738 (0.0433)
Industry within province, log employment	-0.0060 (0.0025)	0.792	-0.0022 (0.0051)
Municipality, log BPO employment	-0.0595 (0.0258)	0.898	0.0420 (0.0473)

Notes: The dose estimate is the coefficient on exposure interacted with US adoption at a one-quarter lag, as in main-text Table 1 (province columns weighted by pre-period employment, industry column weighted by pre-period cell employment, municipal column unweighted). The curvature column reports the joint p-value on exposure interacted with squared and cubed time, estimated on the seven pre-release rounds with the linear term included. The placebo column reports the coefficient on exposure interacted with the realized adoption curve of the first seven post-release quarters assigned to the seven pre-release rounds, with the linear trend control included, and its standard error in parentheses. Standard errors clustered by province, province-industry, or municipality to match each design.

Panel (a) of main-text Figure 3 adds the shape of the response in the design that passes the tests. It plots the within-province industry event-study coefficients net of a line fitted to the six estimated pre-release points. The detrended path sits at zero before the release, stays near zero in the first two quarters after it, when lagged adoption was still close to zero, and turns negative in nearly every quarter from mid-2023 onward, which is the path the dose model predicts rather than the step at the release date that a break-date model predicts. The quarter-by-quarter deviations are individually imprecise, because a trend extrapolated from six points carries wide uncertainty, and the pooled post-release deviation of -0.004 (SE 0.018) reflects that extrapolation penalty. The powered test of the response remains the joint dose-and-trend specification of main-text Table 1, which estimates the trend from all rounds jointly rather than extrapolating it.

A shock specific to the exposed industries themselves would survive the province-by-round fixed effects, and the candidate in this window is the return-to-office mandate for the sector's tax-zone firms. The Fiscal Incentives Review Board required registered firms to return workers onsite to keep their tax incentives, with the requirement binding from April 2023, a transfer window in late 2022 that let firms keep remote work by re-registering with the Board of Investments, and a formal relaxation to half-remote work in November 2024. Three features separate the estimates from this shock. First, its time path is a sequence of steps, not the smooth accumulation that identifies the dose coefficient, and the linear trend interaction absorbs a level shift better than it absorbs curvature that the mandate does not possess. Second, its sign reverses: the November 2024 relaxation loosened the constraint while the employment response continued to scale with adoption through mid-2025. Third, the mandate gives firms no reason to rotate their occupation mix away from the most exposed occupations within the sector, and Table S10 shows that rotation. The mandate plausibly contributed attrition to the sector in 2023, which would appear in the level of

BPO employment, but it does not track the adoption path that the estimates load on.

S5 Exposure-Robust Inference

Clustering by province leaves open the concern of Adão et al. (2019) that provinces with similar occupation mixes have correlated residuals. Following Borusyak et al. (2022), I re-estimate the long-difference form of the design at the occupation level, where the identifying assumption operates, using exposure-weighted averages of the outcome and the regressor and instrumenting with the occupation shocks. As the equivalence result requires, the occupation-level point estimate equals the province-level estimate. As Table S5 shows, the exposure-robust standard errors are smaller than the province-clustered ones, so correlated-share dependence is not what sustains the paper's precision. The inverse Herfindahl of the occupation weights implies 14.5 effective shocks across 42 occupation groups, and this concentration is inherent to the setting, because the concentration of Philippine employment in clerical and customer-service occupations is the very thing the paper measures. One caution applies when reading the table. The long-difference share outcome in columns 1 and 2 restates the rising pre-period gradient of the share outcome, so the positive coefficient there is not an adoption effect, and the table's purpose is inference rather than identification.

Table S5: Shock-level inference, long-difference form

	(1) share, province	(2) share, shock-level	(3) log emp, province	(4) log emp, shock-level
Exposure (raw index)	14.18	14.19	-0.34	-0.34
	(2.37)	(1.16)	(0.81)	(0.23)
Units	87	42	87	42

Notes: The outcome is the change in the province mean from pre-release rounds (2021Q2 to 2022Q4) to post-release rounds (2023Q1 to 2025Q2), regressed on the raw exposure index and weighted by pre-period employment. Columns 1 and 3 report the province-level regression with heteroskedasticity-robust standard errors. Columns 2 and 4 report the numerically equivalent occupation-level regression with robust standard errors, which are valid under shock-level identification. Standard errors in parentheses.

S6 Municipal Design Details

The municipal analysis links the LFS workplace codes to the Philippine Standard Geographic Code, and the crosswalk matches 99.9 percent of the 1,645 workplace codes. The treatment intensity is the count of a municipality's IT park and IT center Presidential Proclamations through 2021. The proclamation panel is assembled from the full texts of 287 IT-zone proclamations, of which 97.9 percent are matched to a municipality and 90 percent carry an exact signing date, with the year of issue used for the remainder. The estimation sample keeps municipalities observed with BPO employment in at least 12 of the 17 quarterly rounds, which yields a panel of roughly 1,600 municipality-round observations. Table S6 reports the estimates together with the clustering variants for the headline province-level specification.

Table S6: Municipal estimates and clustering variants

	(1) mun. log BPO	(2) mun. log BPO	(3) mun. unwtd	(4) prov. share	(5) prov. coarse cl.
IT zones (log) × trend	0.0228	0.0118	0.0367		
	(0.0162)	(0.0138)	(0.0148)		
IT zones (log) × AI adoption, lag 1	-0.0420		-0.0595		
	(0.0280)		(0.0258)		
IT zones (log) × kink		-0.0142			
		(0.0169)			
Exposure × trend				0.1202	0.1202
				(0.0118)	(0.0137)
Exposure × AI adoption, lag 1				-0.1495	-0.1495
				(0.0275)	(0.0333)
Observations	1,170	1,170	1,170	1,479	1,479
R^2	0.965	0.965	0.894	0.977	0.977

Notes: Columns 1 to 3 use the municipality-by-round panel of workplace BPO employment. IT-zone intensity is the log of one plus the count of IT park and IT center Presidential Proclamations through 2021. All municipal columns include municipality and province-by-round fixed effects, and standard errors are clustered by municipality. Columns 1 and 2 weight by pre-period municipal BPO employment. Columns 4 and 5 report the province-level dose specification with adoption at a one-quarter lag, clustered by province in column 4 and by ten island-group clusters in column 5. Standard errors in parentheses.

S7 Industry Exposure Ranking

The industry-level index assigns each two-digit PSIC division the exposure implied by its national pre-period occupation mix, using the same occupation weights as the provincial index. Table S7 reports the most and least exposed divisions. The ranking has face validity from both ends. The most exposed division is office administrative and business support services, which contains the call centers, and it is followed by gambling and betting, which in the Philippines includes online-gaming operations staffed by customer-service and back-office workers, and then by information services, finance, head offices, and computer programming. The least exposed divisions are forestry, fishing, and other primary activities. The full ranking for all 86 divisions is in the replication package.

Table S7: Industry exposure, most and least exposed PSIC divisions

Rank	PSIC	Division	Exposure	Pre-period employment (millions)
1	82	Office administrative and business support (incl. call centres)	0.412	8.56
2	92	Gambling and betting activities	0.368	1.65
3	63	Information service activities	0.366	0.50
4	64	Financial services, except insurance	0.347	3.36
5	70	Head offices and management consultancy	0.328	0.29
6	62	Computer programming and consultancy	0.300	0.95
7	73	Advertising and market research	0.291	0.25
8	65	Insurance and pension funding	0.278	0.50
9	79	Travel agency, tour operator, and reservation services	0.265	0.26
10	66	Auxiliary financial activities	0.260	0.48
11	69	Legal and accounting activities	0.238	1.07
12	99	Extraterritorial organizations and bodies	0.225	0.01
13	78	Employment activities	0.221	0.42
14	58	Publishing activities	0.217	0.20
15	72	Scientific research and development	0.212	0.16
		⋮		
82	39	Remediation and other waste management services	0.011	0.02
83	38	Waste collection, treatment, and disposal	0.009	0.21
84	96	Other personal services	0.008	16.78
85	03	Fishing and aquaculture	0.006	9.15
86	02	Forestry and logging	0.005	0.84

Notes: Exposure is the employment-share-weighted mean of occupation-level observed AI exposure over each division's national pre-period (2021Q2 to 2022Q4) occupation mix. Employment is the division's average weighted employment over the pre-period rounds. Division titles follow PSIC 2009, which aligns with ISIC Revision 4.

S8 Predicted-Index Construction

The capability-based comparison index in Panel B of main-text Figure 1 replaces the observed-usage occupation weights with the AI Occupational Exposure scores of Felten et al. (2021) and holds every other convention fixed. The AIOE scores are published on 2010 SOC codes, so they enter the same crosswalk chain at the 2010 SOC stage, receive the same OEWS employment weights through the successor-code mapping, aggregate to the same two-digit ISCO groups, and combine with the same pre-period provincial employment shares. The resulting provincial index correlates with the observed-usage index at $r = 0.92$. The deviations concentrate at the top of the distribution, where the observed-usage index ranks the Metro Manila districts higher than capability alone would, and the paper reads this as evidence that realized usage adds information where the BPO industry operates.

S9 Adjustment Margins

Table S8 reports the full grid of estimates behind the adjustment-margins figure in the main text. The province block applies the main-text specification to log employment in each destination sector, with the BPO divisions 62, 63, and 82 excluded throughout, at adoption lags of one to four quarters, in the full sample of 87 province-level areas and in the sample that excludes the four NCR districts. The full-sample agriculture coefficients illustrate why both samples are shown. The NCR districts carry large regression weights and hold agricultural cells of a few hundred workers, so their log employment is dominated by sampling noise, and excluding them moves the lag-1 agriculture coefficient from 0.117 to 0.011. No other coefficient changes sign or becomes distinguishable from zero across the two samples.

The region block estimates the same specification across the 17 regions, replacing provincial exposure with its employment-weighted regional average. Unemployment, participation, and overseas membership are attributes of a person's residence rather than of a workplace, and the public use files code residence at the region level, so the region is the finest geography at which these margins can be measured. Employment status comes from the LFS summary status variable and overseas membership from the household roster indicator for overseas Filipino workers, aggregated with survey weights by region and round, and the resulting series are included in the replication package. The national unemployment and participation series quoted in the main text come from the same extraction aggregated over regions.

Table S8: Adjustment margins: full grid of dose estimates

Outcome	Sample	Lag 1	Lag 2	Lag 3	Lag 4	N
Province level: log sector employment						
Agriculture	all areas	0.1170 (0.0237)	0.1069 (0.0211)	0.0336 (0.0085)	0.0367 (0.0436)	1,476
	excl. NCR	0.0111 (0.0153)	0.0317 (0.0247)	0.0318 (0.0261)	0.0292 (0.0260)	1,411
Manufacturing	all areas	-0.0028 (0.0098)	0.0002 (0.0092)	0.0014 (0.0070)	-0.0129 (0.0054)	1,479
	excl. NCR	-0.0314 (0.0274)	-0.0337 (0.0279)	-0.0150 (0.0218)	-0.0212 (0.0183)	1,411
Other industry	all areas	-0.0156 (0.0062)	-0.0136 (0.0053)	-0.0116 (0.0046)	-0.0188 (0.0044)	1,479
	excl. NCR	-0.0314 (0.0203)	-0.0263 (0.0232)	-0.0143 (0.0204)	-0.0208 (0.0178)	1,411
Trade	all areas	0.0078 (0.0032)	0.0086 (0.0034)	0.0051 (0.0035)	-0.0048 (0.0051)	1,479
	excl. NCR	0.0004 (0.0132)	-0.0021 (0.0108)	-0.0079 (0.0116)	-0.0252 (0.0116)	1,411
Transport	all areas	0.0032 (0.0049)	0.0004 (0.0046)	-0.0026 (0.0040)	-0.0072 (0.0047)	1,479
	excl. NCR	-0.0057 (0.0181)	-0.0047 (0.0143)	0.0084 (0.0158)	0.0180 (0.0138)	1,411
Accommodation and food	all areas	-0.0054 (0.0074)	0.0005 (0.0093)	0.0087 (0.0075)	0.0130 (0.0101)	1,462
	excl. NCR	-0.0211 (0.0301)	-0.0188 (0.0367)	0.0027 (0.0311)	0.0296 (0.0398)	1,394
Market services (excl. BPO)	all areas	-0.0166 (0.0086)	-0.0213 (0.0062)	-0.0176 (0.0063)	-0.0142 (0.0077)	1,475
	excl. NCR	0.0080 (0.0255)	-0.0211 (0.0280)	-0.0299 (0.0229)	-0.0250 (0.0268)	1,407
Public and social services	all areas	0.0028 (0.0037)	0.0017 (0.0028)	-0.0002 (0.0049)	-0.0028 (0.0036)	1,479
	excl. NCR	0.0204 (0.0140)	0.0052 (0.0137)	0.0021 (0.0156)	0.0031 (0.0176)	1,411
Other services	all areas	-0.0150 (0.0049)	-0.0103 (0.0058)	0.0054 (0.0072)	0.0045 (0.0054)	1,477
	excl. NCR	-0.0081 (0.0178)	0.0151 (0.0209)	0.0558 (0.0167)	0.0452 (0.0146)	1,409
Total non-BPO	all areas	0.0004 (0.0019)	0.0030 (0.0025)	-0.0014 (0.0025)	-0.0061 (0.0022)	1,479
	excl. NCR	0.0044 (0.0059)	0.0143 (0.0070)	0.0070 (0.0066)	0.0009 (0.0058)	1,411
Total employment	all areas	-0.0013 (0.0020)	0.0015 (0.0025)	-0.0024 (0.0021)	-0.0068 (0.0019)	1,479
	excl. NCR	0.0041 (0.0060)	0.0138 (0.0071)	0.0058 (0.0071)	-0.0004 (0.0062)	1,411
Region level: labor market margins						
Unemployment rate	all regions	0.2852 (0.0409)	0.0807 (0.0520)	0.1202 (0.0563)	0.1098 (0.0553)	282
Labor force participation	all regions	0.2713 (0.2201)	0.3535 (0.3310)	0.1563 (0.2179)	0.0011 (0.1264)	282
OFWs per 1,000 residents	all regions	-0.0981 (0.0860)	0.0551 (0.1206)	0.0954 (0.1156)	0.1290 (0.1525)	282

Notes: Each cell reports the coefficient on exposure interacted with US adoption at the indicated lag, from the main-text specification with exposure-by-trend interactions and unit and round fixed effects, with the standard error in parentheses. Province rows: log employment in the indicated sector group, weighted by pre-period employment, standard errors clustered by province, BPO divisions excluded. Region rows: unemployment rate and labor force participation in percentage points and overseas Filipino workers per 1,000 residents, exposure replaced by its employment-weighted regional average, weighted by pre-period employment, standard errors clustered by region. N refers to the lag-1 estimate.

The overseas flow estimates quoted in the main text come from the Department of Migrant Workers' published deployment statistics, which report deployed landbased workers by region of origin separately for new hires and rehires. I assemble the monthly releases into a region-by-window panel, where a window is a year-to-date span that repeats across years (January, January to May, January to July, and the full year), so that seasonality compares like with like. The Negros Island Region is merged back into Regions VI and VII for consistency across the 2024 regional reorganization, the unallocated residual category is dropped, and later releases revise earlier ones, so each figure comes from the latest available release. The specification interacts the employment-weighted regional exposure with the national adoption dose averaged over the window shifted back six months for deployment processing, controls for exposure interacted with a linear year trend, and includes region-by-window and window-by-year fixed effects, with standard errors clustered by region and wild-bootstrap p-values reported alongside the analytic ones in Table S9.

Table S9: Deployment flows by region of origin: dose estimates

	Coefficient	SE	Analytic p	Wild-bootstrap p
New hires	0.0142	0.0061	0.034	0.165
Total deployment	0.0019	0.0062	0.767	0.968
Rehires	-0.0084	0.0023	0.002	0.099

Notes: Outcomes are log deployed landbased overseas Filipino workers by region of origin, from the Department of Migrant Workers' monthly compendium, in a region-by-window-by-year panel (16 region units after merging the Negros Island Region into Regions VI and VII). Each cell reports the coefficient on regional exposure interacted with the national adoption dose, with exposure-by-trend interactions and region-by-window and window-by-year fixed effects, weighted by pre-period regional employment. Standard errors clustered by region. The wild-bootstrap p-values use 9,999 replications.

A related question is whether the displacement reallocates within the BPO sector itself, from exposed to less exposed work. Table S10 reports two tests. The division rows restrict the industry panel to the three BPO divisions and interact division exposure with the dose under province-by-round and province-by-division fixed effects. The occupation rows build a national occupation-by-round panel within the BPO divisions and interact occupation exposure with the dose under occupation and round fixed effects. The occupation mix rotates away from high-exposure occupations in both weightings, so part of the adjustment does run inside the sector, and the main-text estimates net this reallocation out. The division estimates disagree across weightings, which reflects small provincial cells for the programming and information divisions outside the hubs.

Table S10: Within-BPO reallocation: dose estimates

	Coefficient	SE	N
Divisions within provinces, weighted	-0.0487	0.0234	2,293
Divisions within provinces, unweighted	0.0822	0.0297	2,312
Occupations within the sector, weighted	-0.0733	0.0250	403
Occupations within the sector, unweighted	-0.0407	0.0207	405

Notes: Each cell reports the coefficient on exposure interacted with US adoption at a one-quarter lag, with exposure-by-trend interactions included. Division rows: log employment in BPO divisions 62, 63, and 82 by province and round, with province-by-round and province-by-division fixed effects, standard errors clustered by province-division, weighted by pre-period cell employment where indicated. Occupation rows: log national employment by two-digit occupation within the BPO divisions, occupation and round fixed effects, standard errors clustered by occupation, weighted by pre-period occupation employment where indicated.

A policy note follows from the mechanism rather than from the point estimates. If displacement continues along the adoption path, the unemployment queue documented in the main text will press for an outlet, and three responses are available. The first is to redirect the industry itself toward its less exposed segments, in healthcare information management, complex back-office work, and software services. The occupation shifts in Table S10 show the industry already moving this way for its incumbent workforce, but the less exposed segments are smaller than the work they would need to replace, retraining takes years, and the frontier of what the technology performs keeps moving toward them. The second is to generate employment in other domestic sectors. The skills of displaced service workers transfer to parts of commerce and logistics, but the domestic sectors that could absorb them at scale are the ones whose weakness made services-led development the strategy in the first place, and building them is a project of decades. The third is the overseas channel.

The overseas channel is the most feasible response in the near term because its institutions already exist. The Department of Migrant Workers negotiates bilateral labor agreements, licenses the recruitment industry, and sets deployment procedures, so the capacity to convert local slack into overseas employment is a policy variable that can move within a budget cycle rather than a decade. Agreements matched to the skills of displaced service workers, rather than to traditional deployment categories, would widen the channel where the pressure is arriving.

The overseas option carries its own caveat. Destination countries face the same technology, and their own office workers are exposed to the same automation, so the foreign demand that survives will concentrate in the in-person work that machines cannot yet perform, in care, construction, and hospitality rather than in the office occupations the displaced trained for. Widening the migration channel is therefore a bridge rather than a destination, and the skill content of the agreements determines how much of the displaced workers' education the bridge preserves while the domestic options mature.

References

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

Ramp (2025). Ramp AI index. Ramp Economics Lab. Accessed July 2026.

US Bureau of Labor Statistics (2023). Occupational employment and wage statistics, may 2023. Washington, DC.

Supplemental Appendix

Early Signs of the Generative AI Labor Market Shock in the Global South:

Evidence from Philippine Business Process Outsourcing

Nico Ravanilla

July 21, 2026

Contents

S1 Crosswalk Construction and Coverage

The occupation weights originate on the 2018 US Standard Occupational Classification, and the Labor Force Survey codes occupations on the Philippine Standard Occupational Classification (PSOC 2012). I connect the two through ISCO-08, mapping each 2018 SOC occupation to its 2010 SOC predecessors and then to ISCO-08 unit groups using the official BLS crosswalks. Because a single SOC occupation can map to several ISCO unit groups, its OEWS May 2023 employment (US Bureau of Labor Statistics, 2023) is split equally across its mapped groups, and the exposure of an ISCO unit group is the employment-weighted mean of the exposure scores of its constituent occupations. The procedure covers 421 of the 436 ISCO-08 unit groups. Because the LFS public use files record PSOC at two digits, the analysis operates at the ISCO-08 sub-major level, where PSOC 2012 and ISCO-08 coincide. The 87 province-level areas of the LFS workplace coding are the 81 provinces of that classification, the four districts of Metro Manila, and the independent cities of Cotabato and Isabela. As the final check on coverage, the share of provincial employment matched to an exposure weight averages 100 percent across the 87 areas, with a minimum of 99.8 percent.

S2 Anthropic Economic Index: Release Harmonization

The Anthropic Economic Index has been released in successive waves whose schemas and classifiers differ, and the replication package harmonizes six usage windows spanning December 2024 to May 2026, separately for the consumer product (Claude.ai) and the first-party API. Usage records are matched to O*NET tasks in each window, and matching succeeds for essentially every task row. Within each window I renormalize usage shares to the classified total, since the share of unclassified usage varies across releases. Two caveats matter for anyone using the harmonized panel. First, there are level breaks between release families, which reflect classifier and taxonomy changes rather than changes in usage, so growth should be computed within a release family or absorbed with window effects. Second, the country-level data cover the consumer product only, while enterprise API usage, where automation-type use concentrates, is published at the global level. Neither caveat affects the main estimates in the paper, because the paper uses only the time-invariant observed-exposure weights from the labor market release, and the time variation comes from the adoption survey instead.

S3 The Adoption Dose: Construction and Alternatives

Table S3 reports the quarterly dose series and how each segment is constructed. The spine is the US Census Bureau's Business Trends and Outlook Survey, which asks firms biweekly whether they used AI in producing goods or services and which begins fielding that question in September 2023. Two adjustments extend the spine. Backward, I set the series to zero before 2022Q4 and backcast the three quarters from 2022Q4 to 2023Q2 by scaling the first BTOS reading with the growth of the Ramp AI adoption index (Ramp, 2025), cross-checked against ChatGPT web-traffic estimates and

Google Trends. Forward, the BTOS question changed in November 2025 from AI use in production to AI use in any business function, which produced a level jump from 9.9 to 17.3 percent, and I place the revised readings on the original basis using the ratio of the trend-extrapolated original series to the first revised observation, which equals 0.600. Three alternative constructions leave the conclusions of main-text Table 1 unchanged. The first sets the backcast segment to zero. The second uses the revised post-November-2025 readings without rescaling and absorbs the wording change with an indicator. The third ends the sample in 2025Q3, before the question change, which is also before the end of the LFS outcome panel used in the paper.

A fourth variant addresses the fact that BPO clients are not average US firms. The BTOS publishes adoption by sector, and the IT-BPM industry's client verticals concentrate in the sectors that adopted fastest. I therefore construct a client-weighted dose that averages sector-level BTOS adoption using the industry's global client-vertical mix from the IBPAP Roadmap 2028: banking, financial services, and insurance at 26 percent mapped to NAICS 52, technology, media, and telecom at 20 percent mapped to NAICS 51, retail at 10 percent mapped to NAICS 44, manufacturing at 11 percent mapped to NAICS 31, healthcare at 11 percent mapped to NAICS 62, energy at 6 percent split between NAICS 21 and 22, and the remaining 18 percent carrying the national series. The client-weighted dose runs about 40 percent above the national series at every point, because the client sectors adopt faster, and it correlates with the national series at 0.997, so the shape that identification relies on is unchanged. Table S1 re-estimates the four employment columns of main-text Table 1 with this dose at a one-quarter lag. The coefficients scale down in proportion to the higher dose level, as they must, and the test statistics are unchanged or stronger.

A fifth variant addresses client geography rather than client sector. North America accounts for about 70 percent of the Philippine industry's buyer base, with Europe including the UK at about 15 percent and APAC at about 15 percent, from the same IBPAP Roadmap 2028. Workers serving non-US clients received those countries' adoption paths rather than the US path, so the US-only dose overstates their treatment intensity and dilutes the estimated response toward zero. I construct a geography-weighted dose that averages three blocks by these shares. The North America block is the US series. The Europe block is the share of UK businesses using large language models for text generation from the ONS Business Insights and Conditions Survey, the only national series that measures generative AI use specifically. The APAC block is the mean post-2022 adoption increase in Japan's Communications Usage Trend Survey and Singapore's IMDA enterprise survey. Each foreign block is zero through 2022Q4, matching the splice convention of the main dose, and annual readings are interpolated to quarters. Foreign adoption tracks US adoption closely, and the geography-weighted dose correlates with the US-only series at 0.998. Table S2 re-estimates the four employment columns with this dose at a one-quarter lag. The coefficients are modestly larger than under the US-only dose, consistent with the foreign blocks rising somewhat faster late in the sample, and the test statistics are unchanged.

Table S1: Client-weighted dose: employment responses

	(1) BPO share	(2) log BPO	(3) log emp	(4) log emp, unwtd
Exposure × client-weighted adoption (lag 1)	-0.1050	-0.0255		
	(0.0190)	(0.0081)		
Exposure × client-weighted adoption (lag 1)			-0.0043	-0.0089
			(0.0017)	(0.0019)
Observations	1,479	1,312	67,290	67,832
R^2	0.977	0.953	0.960	0.902

Notes: Same specifications as main-text Table 1, columns 1 to 4, with the national adoption dose replaced by the client-weighted dose at a one-quarter lag. Columns 1 and 2 use the province panel with province and round fixed effects, weighted by pre-period employment, with standard errors clustered by province. Columns 3 and 4 use the province-by-industry panel with province-by-round and province-by-industry fixed effects, column 3 weighted by pre-period cell employment, with standard errors clustered by province-industry. Standard errors in parentheses.

Table S2: Client-geography-weighted dose: employment responses

	(1) BPO share	(2) log BPO	(3) log emp	(4) log emp, unwtd
Exposure × geography-weighted adoption (lag 1)	-0.1292	-0.0321		
	(0.0227)	(0.0112)		
Exposure × geography-weighted adoption (lag 1)			-0.0051	-0.0091
			(0.0022)	(0.0024)
Observations	1,479	1,312	67,290	67,832
R^2	0.977	0.953	0.960	0.902

Notes: Same specifications as main-text Table 1, columns 1 to 4, with the national adoption dose replaced by the client-geography-weighted dose at a one-quarter lag. The dose weights the US series at 70 percent, a UK block at 15 percent, and a Japan-Singapore block at 15 percent, following the buyer-geography mix of the Philippine IT-BPM market in the IBPAP Roadmap 2028. Columns and inference as in Table S1. Standard errors in parentheses.

Table S3: The adoption dose series D_t, by quarter and construction segment

Quarter	D_t (% of US firms using AI)	Construction
2021Q1	0.00	assumed zero (pre-ChatGPT)
2021Q2	0.00	assumed zero (pre-ChatGPT)
2021Q3	0.00	assumed zero (pre-ChatGPT)
2021Q4	0.00	assumed zero (pre-ChatGPT)
2022Q1	0.00	assumed zero (pre-ChatGPT)
2022Q2	0.00	assumed zero (pre-ChatGPT)
2022Q3	0.00	assumed zero (pre-ChatGPT)
2022Q4	0.30	backcast: avg of Google Trends & Similarweb visit ratios to 2023Q1 x Ramp-based 2023Q1 value
2023Q1	1.78	backcast: BTOS 2023Q3 anchor x Ramp AI Index ratio
2023Q2	3.35	backcast: BTOS 2023Q3 anchor x Ramp AI Index ratio
2023Q3	3.70	BTOS original question (producing goods or services)
2023Q4	4.42	BTOS original question (producing goods or services)
2024Q1	4.74	BTOS original question (producing goods or services)
2024Q2	4.77	BTOS original question (producing goods or services)
2024Q3	5.47	BTOS original question (producing goods or services)
2024Q4	6.09	BTOS original question (producing goods or services)
2025Q1	7.33	BTOS original question (producing goods or services)
2025Q2	8.96	BTOS original question (producing goods or services)
2025Q3	9.38	BTOS original question (producing goods or services)
2025Q4	10.51	BTOS revised question x 0.600 wording-splice factor
2026Q1	11.09	BTOS revised question x 0.600 wording-splice factor
2026Q2	12.10	BTOS revised question x 0.600 wording-splice factor

Notes: D_t is the share of US firms using AI in production. The BTOS segment uses the US Census Bureau's biweekly survey aggregated to quarters. The backcast segment scales the first BTOS reading (3.7 percent, 2023Q3) by the growth of the Ramp AI adoption index. The rescaled segment places the revised BTOS question (in use from November 17, 2025) on the original basis using a factor of 0.600. The raw component series are included in the replication package.

S4 Tests of the Identifying Assumption

The estimates in the paper are identified by the curvature of the adoption path. The assumption is that, conditional on the exposure-by-linear-trend interactions, no omitted shock moves employment across exposure groups with the same curvature. This section tests the two implications of that assumption that the pre-release data can check. First, the pre-release exposure gradient should be linear in time, because any pre-existing curvature would be a shock of the prohibited kind operating before the technology existed. I test this by adding exposure interacted with squared and cubed time to the seven pre-release rounds (2021Q2 to 2022Q4) and testing the added terms jointly. Second, a dose-shaped regressor placed where no dose existed should attract nothing. I assign the realized adoption curve of the first seven post-release quarters to the seven pre-release rounds and estimate each design with its linear trend control, so the placebo coefficient is identified from the same kind of curvature as the real one.

Table S4 reports both tests next to each design's dose estimate. The two within-province designs pass both. In the industry design the curvature terms are jointly insignificant ($p = 0.79$) and the placebo dose attracts -0.002 (SE 0.005), a third of the real estimate and indistinguishable from zero, and in the municipal design the corresponding figures are $p = 0.90$ and a wrong-signed 0.042 (SE 0.047). The province-level rows do not pass. The BPO employment share shows pre-release curvature ($p = 0.001$), and the placebo dose attracts -0.138 (SE 0.037), which is comparable in size to the real estimate. The calendar explains why. The seven pre-release rounds span the reopening from pandemic restrictions, including renewed lockdowns in August 2021 and January 2022 that bound most tightly in the capital region, where exposure is highest, so the pre-release path of the exposed provinces is not linear, and a curvature-based estimator at the province level partly picks up the reopening cycle. The province-by-round fixed effects of the industry and municipal designs absorb the restriction cycle wholesale, and in those designs the pre-release period behaves as the identifying assumption requires. The province-level columns of main-text Table 1 should accordingly be read as the descriptive gradient, and the identification rests on the within-province and municipal designs.

Table S4: Tests of the identifying assumption, by design

	Dose estimate (lag 1)	Pre-release curvature (joint p-value)	Pre-release placebo dose
Province, BPO employment share	-0.1495 (0.0275)	0.001	-0.1380 (0.0369)
Province, log BPO employment	-0.0366 (0.0119)	0.088	-0.0738 (0.0433)
Industry within province, log employment	-0.0060 (0.0025)	0.792	-0.0022 (0.0051)
Municipality, log BPO employment	-0.0595 (0.0258)	0.898	0.0420 (0.0473)

Notes: The dose estimate is the coefficient on exposure interacted with US adoption at a one-quarter lag, as in main-text Table 1 (province columns weighted by pre-period employment, industry column weighted by pre-period cell employment, municipal column unweighted). The curvature column reports the joint p-value on exposure interacted with squared and cubed time, estimated on the seven pre-release rounds with the linear term included. The placebo column reports the coefficient on exposure interacted with the realized adoption curve of the first seven post-release quarters assigned to the seven pre-release rounds, with the linear trend control included, and its standard error in parentheses. Standard errors clustered by province, province-industry, or municipality to match each design.

Panel (a) of main-text Figure 3 adds the shape of the response in the design that passes the tests. It plots the within-province industry event-study coefficients net of a line fitted to the six estimated pre-release points. The detrended path sits at zero before the release, stays near zero in the first two quarters after it, when lagged adoption was still close to zero, and turns negative in nearly every quarter from mid-2023 onward, which is the path the dose model predicts rather than the step at the release date that a break-date model predicts. The quarter-by-quarter deviations are individually imprecise, because a trend extrapolated from six points carries wide uncertainty, and the pooled post-release deviation of -0.004 (SE 0.018) reflects that extrapolation penalty. The powered test of the response remains the joint dose-and-trend specification of main-text Table 1, which estimates the trend from all rounds jointly rather than extrapolating it.

A shock specific to the exposed industries themselves would survive the province-by-round fixed effects, and the candidate in this window is the return-to-office mandate for the sector's tax-zone firms. The Fiscal Incentives Review Board required registered firms to return workers onsite to keep their tax incentives, with the requirement binding from April 2023, a transfer window in late 2022 that let firms keep remote work by re-registering with the Board of Investments, and a formal relaxation to half-remote work in November 2024. Three features separate the estimates from this shock. First, its time path is a sequence of steps, not the smooth accumulation that identifies the dose coefficient, and the linear trend interaction absorbs a level shift better than it absorbs curvature that the mandate does not possess. Second, its sign reverses: the November 2024 relaxation loosened the constraint while the employment response continued to scale with adoption through mid-2025. Third, the mandate gives firms no reason to rotate their occupation mix away from the most exposed occupations within the sector, and Table S10 shows that rotation. The mandate plausibly contributed attrition to the sector in 2023, which would appear in the level of

BPO employment, but it does not track the adoption path that the estimates load on.

S5 Exposure-Robust Inference

Clustering by province leaves open the concern of Adão et al. (2019) that provinces with similar occupation mixes have correlated residuals. Following Borusyak et al. (2022), I re-estimate the long-difference form of the design at the occupation level, where the identifying assumption operates, using exposure-weighted averages of the outcome and the regressor and instrumenting with the occupation shocks. As the equivalence result requires, the occupation-level point estimate equals the province-level estimate. As Table S5 shows, the exposure-robust standard errors are smaller than the province-clustered ones, so correlated-share dependence is not what sustains the paper's precision. The inverse Herfindahl of the occupation weights implies 14.5 effective shocks across 42 occupation groups, and this concentration is inherent to the setting, because the concentration of Philippine employment in clerical and customer-service occupations is the very thing the paper measures. One caution applies when reading the table. The long-difference share outcome in columns 1 and 2 restates the rising pre-period gradient of the share outcome, so the positive coefficient there is not an adoption effect, and the table's purpose is inference rather than identification.

Table S5: Shock-level inference, long-difference form

	(1) share, province	(2) share, shock-level	(3) log emp, province	(4) log emp, shock-level
Exposure (raw index)	14.18	14.19	-0.34	-0.34
	(2.37)	(1.16)	(0.81)	(0.23)
Units	87	42	87	42

Notes: The outcome is the change in the province mean from pre-release rounds (2021Q2 to 2022Q4) to post-release rounds (2023Q1 to 2025Q2), regressed on the raw exposure index and weighted by pre-period employment. Columns 1 and 3 report the province-level regression with heteroskedasticity-robust standard errors. Columns 2 and 4 report the numerically equivalent occupation-level regression with robust standard errors, which are valid under shock-level identification. Standard errors in parentheses.

S6 Municipal Design Details

The municipal analysis links the LFS workplace codes to the Philippine Standard Geographic Code, and the crosswalk matches 99.9 percent of the 1,645 workplace codes. The treatment intensity is the count of a municipality's IT park and IT center Presidential Proclamations through 2021. The proclamation panel is assembled from the full texts of 287 IT-zone proclamations, of which 97.9 percent are matched to a municipality and 90 percent carry an exact signing date, with the year of issue used for the remainder. The estimation sample keeps municipalities observed with BPO employment in at least 12 of the 17 quarterly rounds, which yields a panel of roughly 1,600 municipality-round observations. Table S6 reports the estimates together with the clustering variants for the headline province-level specification.

Table S6: Municipal estimates and clustering variants

	(1) mun. log BPO	(2) mun. log BPO	(3) mun. unwtd	(4) prov. share	(5) prov. coarse cl.
IT zones (log) × trend	0.0228	0.0118	0.0367		
	(0.0162)	(0.0138)	(0.0148)		
IT zones (log) × AI adoption, lag 1	-0.0420		-0.0595		
	(0.0280)		(0.0258)		
IT zones (log) × kink		-0.0142			
		(0.0169)			
Exposure × trend				0.1202	0.1202
				(0.0118)	(0.0137)
Exposure × AI adoption, lag 1				-0.1495	-0.1495
				(0.0275)	(0.0333)
Observations	1,170	1,170	1,170	1,479	1,479
R^2	0.965	0.965	0.894	0.977	0.977

Notes: Columns 1 to 3 use the municipality-by-round panel of workplace BPO employment. IT-zone intensity is the log of one plus the count of IT park and IT center Presidential Proclamations through 2021. All municipal columns include municipality and province-by-round fixed effects, and standard errors are clustered by municipality. Columns 1 and 2 weight by pre-period municipal BPO employment. Columns 4 and 5 report the province-level dose specification with adoption at a one-quarter lag, clustered by province in column 4 and by ten island-group clusters in column 5. Standard errors in parentheses.

S7 Industry Exposure Ranking

The industry-level index assigns each two-digit PSIC division the exposure implied by its national pre-period occupation mix, using the same occupation weights as the provincial index. Table S7 reports the most and least exposed divisions. The ranking has face validity from both ends. The most exposed division is office administrative and business support services, which contains the call centers, and it is followed by gambling and betting, which in the Philippines includes online-gaming operations staffed by customer-service and back-office workers, and then by information services, finance, head offices, and computer programming. The least exposed divisions are forestry, fishing, and other primary activities. The full ranking for all 86 divisions is in the replication package.

Table S7: Industry exposure, most and least exposed PSIC divisions

Rank	PSIC	Division	Exposure	Pre-period employment (millions)
1	82	Office administrative and business support (incl. call centres)	0.412	8.56
2	92	Gambling and betting activities	0.368	1.65
3	63	Information service activities	0.366	0.50
4	64	Financial services, except insurance	0.347	3.36
5	70	Head offices and management consultancy	0.328	0.29
6	62	Computer programming and consultancy	0.300	0.95
7	73	Advertising and market research	0.291	0.25
8	65	Insurance and pension funding	0.278	0.50
9	79	Travel agency, tour operator, and reservation services	0.265	0.26
10	66	Auxiliary financial activities	0.260	0.48
11	69	Legal and accounting activities	0.238	1.07
12	99	Extraterritorial organizations and bodies	0.225	0.01
13	78	Employment activities	0.221	0.42
14	58	Publishing activities	0.217	0.20
15	72	Scientific research and development	0.212	0.16
		⋮		
82	39	Remediation and other waste management services	0.011	0.02
83	38	Waste collection, treatment, and disposal	0.009	0.21
84	96	Other personal services	0.008	16.78
85	03	Fishing and aquaculture	0.006	9.15
86	02	Forestry and logging	0.005	0.84

Notes: Exposure is the employment-share-weighted mean of occupation-level observed AI exposure over each division's national pre-period (2021Q2 to 2022Q4) occupation mix. Employment is the division's average weighted employment over the pre-period rounds. Division titles follow PSIC 2009, which aligns with ISIC Revision 4.

S8 Predicted-Index Construction

The capability-based comparison index in Panel B of main-text Figure 1 replaces the observed-usage occupation weights with the AI Occupational Exposure scores of Felten et al. (2021) and holds every other convention fixed. The AIOE scores are published on 2010 SOC codes, so they enter the same crosswalk chain at the 2010 SOC stage, receive the same OEWS employment weights through the successor-code mapping, aggregate to the same two-digit ISCO groups, and combine with the same pre-period provincial employment shares. The resulting provincial index correlates with the observed-usage index at $r = 0.92$. The deviations concentrate at the top of the distribution, where the observed-usage index ranks the Metro Manila districts higher than capability alone would, and the paper reads this as evidence that realized usage adds information where the BPO industry operates.

S9 Adjustment Margins

Table S8 reports the full grid of estimates behind the adjustment-margins figure in the main text. The province block applies the main-text specification to log employment in each destination sector, with the BPO divisions 62, 63, and 82 excluded throughout, at adoption lags of one to four quarters, in the full sample of 87 province-level areas and in the sample that excludes the four NCR districts. The full-sample agriculture coefficients illustrate why both samples are shown. The NCR districts carry large regression weights and hold agricultural cells of a few hundred workers, so their log employment is dominated by sampling noise, and excluding them moves the lag-1 agriculture coefficient from 0.117 to 0.011. No other coefficient changes sign or becomes distinguishable from zero across the two samples.

The region block estimates the same specification across the 17 regions, replacing provincial exposure with its employment-weighted regional average. Unemployment, participation, and overseas membership are attributes of a person's residence rather than of a workplace, and the public use files code residence at the region level, so the region is the finest geography at which these margins can be measured. Employment status comes from the LFS summary status variable and overseas membership from the household roster indicator for overseas Filipino workers, aggregated with survey weights by region and round, and the resulting series are included in the replication package. The national unemployment and participation series quoted in the main text come from the same extraction aggregated over regions.

Table S8: Adjustment margins: full grid of dose estimates

Outcome	Sample	Lag 1	Lag 2	Lag 3	Lag 4	*N*
Province level: log sector employment						
Agriculture	all areas	0.1170 (0.0237)	0.1069 (0.0211)	0.0336 (0.0085)	0.0367 (0.0436)	1,476
	excl. NCR	0.0111 (0.0153)	0.0317 (0.0247)	0.0318 (0.0261)	0.0292 (0.0260)	1,411
Manufacturing	all areas	-0.0028 (0.0098)	0.0002 (0.0092)	0.0014 (0.0070)	-0.0129 (0.0054)	1,479
	excl. NCR	-0.0314 (0.0274)	-0.0337 (0.0279)	-0.0150 (0.0218)	-0.0212 (0.0183)	1,411
Other industry	all areas	-0.0156 (0.0062)	-0.0136 (0.0053)	-0.0116 (0.0046)	-0.0188 (0.0044)	1,479
	excl. NCR	-0.0314 (0.0203)	-0.0263 (0.0232)	-0.0143 (0.0204)	-0.0208 (0.0178)	1,411
Trade	all areas	0.0078 (0.0032)	0.0086 (0.0034)	0.0051 (0.0035)	-0.0048 (0.0051)	1,479
	excl. NCR	0.0004 (0.0132)	-0.0021 (0.0108)	-0.0079 (0.0116)	-0.0252 (0.0116)	1,411
Transport	all areas	0.0032 (0.0049)	0.0004 (0.0046)	-0.0026 (0.0040)	-0.0072 (0.0047)	1,479
	excl. NCR	-0.0057 (0.0181)	-0.0047 (0.0143)	0.0084 (0.0158)	0.0180 (0.0138)	1,411
Accommodation and food	all areas	-0.0054 (0.0074)	0.0005 (0.0093)	0.0087 (0.0075)	0.0130 (0.0101)	1,462
	excl. NCR	-0.0211 (0.0301)	-0.0188 (0.0367)	0.0027 (0.0311)	0.0296 (0.0398)	1,394
Market services (excl. BPO)	all areas	-0.0166 (0.0086)	-0.0213 (0.0062)	-0.0176 (0.0063)	-0.0142 (0.0077)	1,475
	excl. NCR	0.0080 (0.0255)	-0.0211 (0.0280)	-0.0299 (0.0229)	-0.0250 (0.0268)	1,407
Public and social services	all areas	0.0028 (0.0037)	0.0017 (0.0028)	-0.0002 (0.0049)	-0.0028 (0.0036)	1,479
	excl. NCR	0.0204 (0.0140)	0.0052 (0.0137)	0.0021 (0.0156)	0.0031 (0.0176)	1,411
Other services	all areas	-0.0150 (0.0049)	-0.0103 (0.0058)	0.0054 (0.0072)	0.0045 (0.0054)	1,477
	excl. NCR	-0.0081 (0.0178)	0.0151 (0.0209)	0.0558 (0.0167)	0.0452 (0.0146)	1,409
Total non-BPO	all areas	0.0004 (0.0019)	0.0030 (0.0025)	-0.0014 (0.0025)	-0.0061 (0.0022)	1,479
	excl. NCR	0.0044 (0.0059)	0.0143 (0.0070)	0.0070 (0.0066)	0.0009 (0.0058)	1,411
Total employment	all areas	-0.0013 (0.0020)	0.0015 (0.0025)	-0.0024 (0.0021)	-0.0068 (0.0019)	1,479
	excl. NCR	0.0041 (0.0060)	0.0138 (0.0071)	0.0058 (0.0071)	-0.0004 (0.0062)	1,411
Region level: labor market margins						
Unemployment rate	all regions	0.2852 (0.0409)	0.0807 (0.0520)	0.1202 (0.0563)	0.1098 (0.0553)	282
Labor force participation	all regions	0.2713 (0.2201)	0.3535 (0.3310)	0.1563 (0.2179)	0.0011 (0.1264)	282
OFWs per 1,000 residents	all regions	-0.0981 (0.0860)	0.0551 (0.1206)	0.0954 (0.1156)	0.1290 (0.1525)	282

Notes: Each cell reports the coefficient on exposure interacted with US adoption at the indicated lag, from the main-text specification with exposure-by-trend interactions and unit and round fixed effects, with the standard error in parentheses. Province rows: log employment in the indicated sector group, weighted by pre-period employment, standard errors clustered by province, BPO divisions excluded. Region rows: unemployment rate and labor force participation in percentage points and overseas Filipino workers per 1,000 residents, exposure replaced by its employment-weighted regional average, weighted by pre-period employment, standard errors clustered by region. *N* refers to the lag-1 estimate.

The overseas flow estimates quoted in the main text come from the Department of Migrant Workers' published deployment statistics, which report deployed landbased workers by region of origin separately for new hires and rehires. I assemble the monthly releases into a region-by-window panel, where a window is a year-to-date span that repeats across years (January, January to May, January to July, and the full year), so that seasonality compares like with like. The Negros Island Region is merged back into Regions VI and VII for consistency across the 2024 regional reorganization, the unallocated residual category is dropped, and later releases revise earlier ones, so each figure comes from the latest available release. The specification interacts the employment-weighted regional exposure with the national adoption dose averaged over the window shifted back six months for deployment processing, controls for exposure interacted with a linear year trend, and includes region-by-window and window-by-year fixed effects, with standard errors clustered by region and wild-bootstrap p-values reported alongside the analytic ones in Table S9.

Table S9: Deployment flows by region of origin: dose estimates

	Coefficient	SE	Analytic p	Wild-bootstrap p
New hires	0.0142	0.0061	0.034	0.165
Total deployment	0.0019	0.0062	0.767	0.968
Rehires	-0.0084	0.0023	0.002	0.099

Notes: Outcomes are log deployed landbased overseas Filipino workers by region of origin, from the Department of Migrant Workers' monthly compendium, in a region-by-window-by-year panel (16 region units after merging the Negros Island Region into Regions VI and VII). Each cell reports the coefficient on regional exposure interacted with the national adoption dose, with exposure-by-trend interactions and region-by-window and window-by-year fixed effects, weighted by pre-period regional employment. Standard errors clustered by region. The wild-bootstrap p-values use 9,999 replications.

A related question is whether the displacement reallocates within the BPO sector itself, from exposed to less exposed work. Table S10 reports two tests. The division rows restrict the industry panel to the three BPO divisions and interact division exposure with the dose under province-by-round and province-by-division fixed effects. The occupation rows build a national occupation-by-round panel within the BPO divisions and interact occupation exposure with the dose under occupation and round fixed effects. The occupation mix rotates away from high-exposure occupations in both weightings, so part of the adjustment does run inside the sector, and the main-text estimates net this reallocation out. The division estimates disagree across weightings, which reflects small provincial cells for the programming and information divisions outside the hubs.

Table S10: Within-BPO reallocation: dose estimates

	Coefficient	SE	N
Divisions within provinces, weighted	-0.0487	0.0234	2,293
Divisions within provinces, unweighted	0.0822	0.0297	2,312
Occupations within the sector, weighted	-0.0733	0.0250	403
Occupations within the sector, unweighted	-0.0407	0.0207	405

Notes: Each cell reports the coefficient on exposure interacted with US adoption at a one-quarter lag, with exposure-by-trend interactions included. Division rows: log employment in BPO divisions 62, 63, and 82 by province and round, with province-by-round and province-by-division fixed effects, standard errors clustered by province-division, weighted by pre-period cell employment where indicated. Occupation rows: log national employment by two-digit occupation within the BPO divisions, occupation and round fixed effects, standard errors clustered by occupation, weighted by pre-period occupation employment where indicated.

A policy note follows from the mechanism rather than from the point estimates. If displacement continues along the adoption path, the unemployment queue documented in the main text will press for an outlet, and three responses are available. The first is to redirect the industry itself toward its less exposed segments, in healthcare information management, complex back-office work, and software services. The occupation shifts in Table S10 show the industry already moving this way for its incumbent workforce, but the less exposed segments are smaller than the work they would need to replace, retraining takes years, and the frontier of what the technology performs keeps moving toward them. The second is to generate employment in other domestic sectors. The skills of displaced service workers transfer to parts of commerce and logistics, but the domestic sectors that could absorb them at scale are the ones whose weakness made services-led development the strategy in the first place, and building them is a project of decades. The third is the overseas channel.

The overseas channel is the most feasible response in the near term because its institutions already exist. The Department of Migrant Workers negotiates bilateral labor agreements, licenses the recruitment industry, and sets deployment procedures, so the capacity to convert local slack into overseas employment is a policy variable that can move within a budget cycle rather than a decade. Agreements matched to the skills of displaced service workers, rather than to traditional deployment categories, would widen the channel where the pressure is arriving.

The overseas option carries its own caveat. Destination countries face the same technology, and their own office workers are exposed to the same automation, so the foreign demand that survives will concentrate in the in-person work that machines cannot yet perform, in care, construction, and hospitality rather than in the office occupations the displaced trained for. Widening the migration channel is therefore a bridge rather than a destination, and the skill content of the agreements determines how much of the displaced workers' education the bridge preserves while the domestic options mature.

References

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

Ramp (2025). Ramp AI index. Ramp Economics Lab. Accessed July 2026.

US Bureau of Labor Statistics (2023). Occupational employment and wage statistics, may 2023. Washington, DC.