AI, Automation, and Expertise

Bouke Klein Teeselink*[1,2,3] and Daniel Carey[2]

[1] *King's College London*
[2] *AI Objectives Institute*
[3] *King's Institute for Artificial Intelligence*

February 10, 2026

Abstract

Not all automation is equal: we demonstrate that labor market effects of AI depend not only on how many tasks are exposed, but on which tasks. Extending Autor and Thompson (2025), we distinguish automation that raises expertise requirements from automation that lowers them. Using hundreds of millions of job postings across 39 countries and exploiting ChatGPT's November 2022 release as a natural experiment, we find that expertise-raising automation increases advertised wages, while expertise-lowering automation does not. In addition, we find that aggregate AI exposure strongly predicts displacement: high-exposure occupations experience a 6.1% decline in job postings, with effects concentrated in countries with stricter employment protection and lower digital readiness. These results suggest that exposure-based measures alone obscure economically meaningful heterogeneity in the labor market effects of AI.

Keywords: Generative AI; Expertise; Large Language Models; Automation; Labor demand.

JEL: O33; J23; J31

*Contact: bouke.klein_teeselink@kcl.ac.uk.

1 Introduction

Two occupations can face identical automation exposure and nevertheless experience diametrically opposed labor market trajectories. For example, during the computerization era, accounting clerks and inventory clerks both had equally high shares of their tasks exposed to automation. Yet between 1980 and 2018, accounting clerks saw rising wages and decreasing employment, while inventory clerks saw falling wages and increasing employment. This puzzle, why equal exposure to automation can produce opposing labor market trajectories, lies at the heart of understanding how generative AI will reshape labor markets.

Autor and Thompson (2025) provide a framework to resolve this puzzle. The main insight is that tasks differ not only in their susceptibility to automation but also in the expertise they require. When automation removes an occupation's most expert tasks, the expertise threshold for entry falls, allowing less-skilled workers to enter and expanding labor supply. This causes wages to fall. This happened to inventory clerks, as automation displaced higher-expertise tasks like price calculations and conformance verification, leaving behind mostly generic physical tasks. When automation removes an occupation's least expert tasks, by contrast, only the specialized work remains, raising the expertise threshold, shrinking the pool of qualified workers, and pushing wages upward. This is the accounting clerk case, as computerization automated primarily routine, lower-expertise tasks like data entry and basic bookkeeping, leaving behind more specialized problem-solving and judgment tasks. Hence, their wages increased while their employment declined.

The insight that identical exposure can produce opposite outcomes depending on which tasks are automated has immediate implications for generative AI. Unlike earlier waves of computerization that primarily automated routine cognitive and manual tasks, AI can perform many tasks previously considered non-routine, including writing, analysis, and certain forms of judgment. The expertise framework predicts that the labor market consequences of this technological shift will depend on where these newly automatable tasks sit in each occupation's expertise distribution.

We provide the first large-scale empirical test of the expertise framework in the context of generative AI, while also offering a more general examination of occupation-level AI exposure effects. Using data on hundreds of millions of job postings from 39 countries and exploiting the release of ChatGPT as a natural

experiment, we examine how AI exposure and expertise change affect both labor demand (measured through job postings) and compensation (measured through advertised wages). We measure AI exposure using Eloundou et al. (2024); Gmyrek et al. (2025), though we note that alternative methodologies for measuring occupational exposure exist (Felten et al., 2023; Prytkova et al., 2024; Appel et al., 2025).

Our cross-country design allows us to test whether occupations experience divergent labor market trajectories depending on how automation reshapes their expertise requirements. It also provides external validity for the occupation-level exposure effects documented by single-country studies in the United States (Brynjolfsson et al., 2025), the United Kingdom (Klein Teeselink, 2025), and Denmark (Humlum and Vestergaard, 2025), by examining whether these patterns generalize across 39 economies with varying labor market institutions, AI adoption rates, and economic structures.

To operationalize the expertise framework, we construct a new measure of occupation-level expertise using LLM-based semantic classification. For all O*NET occupation-level task descriptions, we prompt a large language model to assess the expertise required, based on the specialized knowledge, skill, and judgment the task demands. We validate this expertise measure by demonstrating that it correlates strongly both with independent indicators of occupational expertise, including workers' education levels, years of experience, and wages, and with other expertise proxies. We then examine how expertise requirements would shift at varying levels of LLM exposure. To do so, we re-calculate the average occupation-level expertise level after removing all exposed tasks. When AI exposure is concentrated in low-expertise tasks, expertise is predicted to rise; when it targets high-expertise tasks, predicted expertise falls. It is important to note that our expertise change measure is based on predicted changes rather than realized changes, which, if anything, might attenuate the real impact that automation is likely to have in the future.

Our empirical strategy exploits the release of ChatGPT as a shock to automation potential. Using a difference-in-differences design, we compare labor market outcomes before and after November 2022 across occupations that differ in both their overall AI exposure and their predicted expertise change. This allows us to test how the level of exposure, as well as the distribution of exposure across expertise levels, shapes labor market outcomes. We first estimate the effects for 39 countries separately, and then construct meta-analyses across all countries to examine the

overall effects of exposure and expertise changes on labor market outcomes.

The results partially support the expertise framework. We find that expertise change affects wages as predicted: occupations where automation raises expertise requirements see higher advertised salaries, whereas those where automation lowers expertise do not. The pooled effect is a 0.14 percent increase in salaries for a one standard deviation increase in expertise. The effect, while modest in magnitude, is statistically significant at the aggregate level and consistent in direction at the country level, as 27 of 39 countries show positive point estimates. Contrasting our wage effects, we find no direct effect of expertise change on job postings. The pooled estimate is effectively zero. A possible explanation for the null-effect is that the reduction in the qualified candidate pool from rising expertise requirements is offset by displacement of incumbent workers, who may no longer be qualified or capable to do their job when expertise levels increase.

Turning to AI exposure, we confirm that occupations with greater LLM exposure experience substantial declines in labor demand, consistent with prior single-country studies (Brynjolfsson et al., 2025; Lichtinger and Hosseini Maasoum, 2025; Klein Teeselink, 2025; Azar et al., 2025). Averaged across 39 countries, we find that a one standard deviation increase in exposure is associated with a -6.1 percent reduction in on job postings. The direction of this effect is highly consistent: 31 of 39 countries show negative point estimates, with 14 statistically significant. The largest effects appear in Finland, Slovakia, and South Korea.

We then investigate what institutional and economic factors might explain heterogeneity in the effect of AI exposure on hiring across countries. We correlate country-level exposure effects with four characteristics: pre-LLM unemployment rates, employment protection legislation, AI adoption, and digital readiness. We find that countries with stricter employment protection exhibit larger displacement effects, consistent with the hypothesis that firms in highly exposed occupations face substantial uncertainty about whether AI will eventually automate much of their workforce's tasks, and when dismissal costs are high, the prudent response is to avoid hiring workers who may prove difficult to dismiss if their roles become redundant. Countries with greater digital readiness experience slightly smaller hiring reductions, suggesting that complementary digital infrastructure and capabilities help firms and workers adapt to AI-driven changes. These correlations are exploratory given the limited number of country observations, but they suggest that institutional context

might shape how AI exposure translates into labor market outcomes.

Our paper relates to four strands of the economics literature. First, we build on research on technological change and labor markets, particularly the skill-biased technological change hypothesis (Autor et al., 2003, 2006; Goldin and Katz, 2008; Spitz-Oener, 2006; Akerman et al., 2015). This literature documented that computerization disproportionately displaced middle-skill workers performing routine cognitive and manual tasks, while complementing high-skill workers engaged in abstract problem-solving and low-skill workers in non-routine manual occupations (Goos and Manning, 2007; Goos et al., 2014; Michaels et al., 2014; Autor and Salomons, 2018). A longstanding puzzle in this literature is why wages did not consistently decline in routine task-intensive occupations despite reduced employment (Acemoglu and Restrepo, 2022). The expertise framework provides a potential resolution: routine tasks were relatively expert for some occupations but relatively inexpert for others, so routine task automation had heterogeneous effects on wages depending on where those tasks sat in each occupation's expertise distribution.

Second, we relate to the task-based approach to automation, which provides a micro-founded structure for analyzing how technology reshapes labor demand (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2019; Restrepo, 2024). In this framework, automation operates at the extensive margin by substituting capital for labor in a widening range of tasks. This substitution generates a displacement effect, which reduces the demand for workers performing the newly automated tasks. At the same time, automation also generates a productivity effect by lowering production costs and potentially expanding aggregate output (Acemoglu, 2024). If the cost-savings are modest ("so-so automation"), the productivity effect may be insufficient to counterbalance the displacement of labor, leading to declines in real wages and employment for exposed groups, particularly if labor market frictions hinder reallocation, although recent firm-level evidence suggests AI investment can also drive product innovation and employment growth (Babina et al., 2024).

Third, and most importantly, our paper builds on the framework of Autor and Thompson (2025). Their model formalizes two features absent from canonical task models: hierarchical expertise, and occupational task bundling. They define expertise as a worker's capability to perform specific tasks within a hierarchy where more expert workers can perform all tasks at or below their level, while less expert

workers cannot perform tasks above their level. To measure this, they develop a metric rooted in the Efficient Coding Hypothesis, which identifies expert tasks by analyzing job descriptions (sourced from the DOT and O*NET) for words that possess both low frequency and low entropy, both of which are signals of specialized terminology. Analyzing four decades of U.S. labor market data (1980–2018) during the computerization era, their empirical results confirm that changes in expertise requirements predict labor market outcomes. Specifically, they find that increasing the expertise required (e.g., by automating inexpert tasks) raises wages but reduces employment by restricting the supply of qualified workers, whereas decreases in expertise requirements have the opposite effect.

Fourth, we engage with the rapidly growing literature measuring AI exposure and its labor market effects. Frey and Osborne (2017) sparked widespread concern with estimates that 47% of US employment faces high automation risk. With the advent of generative AI, researchers developed new exposure measures: Eloundou et al. (2024) estimate that 80% of US workers could have at least 10% of their tasks affected by LLMs, while Gmyrek et al. (2025) construct an alternative index linking AI capabilities to international occupation codes. Empirical studies exploiting the release of ChatGPT have generally found negative employment effects: Brynjolfsson et al. (2025) document a 16% decline in employment for young workers in high-exposure jobs; Klein Teeselink (2025) find 3.9% fewer job postings in high-exposure occupations in the UK; and Hui et al. (2024) show a 5.2% decline in earnings for affected freelancers on Upwork. Yet Humlum and Vestergaard (2025) find limited aggregate effects in Denmark, and even within studies showing displacement, there is substantial heterogeneity across occupations that exposure alone cannot explain (Hosseini Maasoum and Lichtinger, 2026). These findings raise the question of which factors determine whether AI exposure leads to displacement or not.

We make three contributions to these literatures. First, we provide the first large-scale empirical test of the expertise framework in the context of generative AI. Using data on hundreds of millions of job postings across 39 countries, we find partial support for the framework's predictions. Our findings suggest that the composition and quantity of automation may operate through different channels, with expertise change shaping compensation and exposure driving displacement. As AI capabilities expand and firms automate a larger share of exposed tasks, these early patterns are likely to intensify, making our findings informative for anticipating future labor

market shifts.

Second, we develop a new task-level expertise measure that offers three advantages over existing approaches. Unlike the measure of Autor and Thompson (2025), which identifies expert tasks through word frequency and entropy patterns in job descriptions, our semantic approach directly assesses the cognitive demands of each task. Our measure also operates at the task level rather than the occupation level, enabling precise identification of which specific tasks within an occupation require high expertise; this granularity is essential because two occupations with identical mean expertise can have very different task distributions and therefore different predictions for how automation reshapes their skill requirements. Finally, future research can combine our expertise measure with any task-level AI exposure index (current or future) to generate predictions about how expertise changes will percolate across occupations. As AI capabilities evolve, researchers can update the exposure component while retaining the expertise scores, producing new forecasts of wages and employment without having to re-estimate the entire framework.

Third, we provide the first systematic cross-country evidence on AI and labor markets. Prior studies, such as Humlum and Vestergaard (2025); Brynjolfsson et al. (2025); Lichtinger and Hosseini Maasoum (2025); Klein Teeselink (2025); Azar et al. (2025) often focus on only one country, which leaves open the possibility that other country-wide factors that may coincide with the introduction of ChatGPT such as interest rate changes are driving the observed effects (Iscenko and Curto Millet, 2026). The consistency of our findings across 39 economies, which span a wide variety of labor market institutions, AI adoption rates, and economic structures, provides important external validity to conclusions drawn from single-country studies and suggests that the reduction in hiring in AI exposed occupations is widespread.

The remainder of the paper is organized as follows. Section 2 describes our three data sources: task-level LLM exposure scores from Eloundou et al. (2024), our expertise measure constructed using GPT-4o, and job posting data spanning 39 countries. Section 3 outlines our empirical strategy. Section 4 presents our main results. Section 5 concludes with policy implications.

2 Data

Our empirical analysis combines three data sources: task-level LLM exposure scores, a task-level expertise measure we construct using LLM-based semantic classification, and job posting data from Revelio Labs covering 39 countries. This section describes each component and how we combine them to test the expertise framework.

2.1 LLM Exposure Scores

Our primary measure of occupation-level LLM exposure is based on Eloundou et al. (2024). They construct task-level LLM exposure measures for 19,265 task-level descriptions across approximately 1,000 occupations from the O*NET 27.2 database through expert human evaluation and LLM validation. We use their β measure, which classifies each task into one of three categories based on whether LLMs could reduce completion time by at least 50% while maintaining quality. A task receives a score of 1 (direct exposure) if an LLM alone can achieve this time reduction, a score of 0.5 (exposure with complementary tools) if additional software built on top of the LLM would be required, and a score of 0 (no exposure) if LLMs are not expected to meaningfully reduce task completion time even with complementary tools. We prefer the β measure over the more restrictive α measure, which captures only tasks automatable by standalone LLMs, for two reasons. First, the β measure better reflects how generative AI is actually deployed in practice: workers rarely use LLMs in isolation but instead integrate them with other digital tools in augmented workflows. Second, with recent advances in AI, β arguably provides a more accurate measure for the current state of AI abilities. Nonetheless, we also report results for the α measure as a robustness check.

For each occupation o, we calculate the importance-weighted mean exposure:

$$\text{Exposure}_o = \frac{\sum_{i \in T_o} w_i \cdot \beta_i}{\sum_{i \in T_o} w_i} \tag{1}$$

where T_o denotes the set of tasks in occupation o, w_i is the O*NET importance weight for task i (normalized to sum to one within each occupation), and $\beta_i \in \{0, 0.5, 1\}$ is the task-level exposure score. The importance weights come directly from the O*NET database, and reflect how central each task is to the occupation's

core functions. The exposure measure ranges from 0 to 1, with higher values indicating greater exposure to LLM automation.

A potential limitation of the Eloundou et al. (2024) measures is their reliance on U.S. O*NET task descriptions. If occupations involve systematically different tasks in other countries, for example due to differences in production technologies, regulatory environments, or labor market institutions, then applying U.S.-based exposure scores internationally may introduce measurement error. Such cross-country variation in task content would attenuate our estimates if uncorrelated with the true exposure, or bias them if systematically correlated.

To address this concern, we conduct robustness checks using the exposure measures developed by Gmyrek et al. (2025). Their measure is constructed at the ISCO occupation level using a methodology explicitly designed for international applicability. The 2023 version draws on the technical documentation of the ISCO-08 classification system, which describes occupations in terms that abstract from national specificities. The 2025 update expands this approach by incorporating survey data from 1,640 workers across all 1-digit ISCO groups, collecting 52,558 assessments of automation potential for 2,861 tasks drawn from the Polish occupational classification system. These worker assessments are then refined through expert consultations and used to train an AI assistant that predicts scores for the full ISCO task inventory. A drawback of using ISCO-based measures is that the classification contains substantially fewer occupations than O*NET (approximately 400 versus 1,000), which reduces statistical power. This is a particular concern given that our restrictive fixed effects structure already controls away a lot of relevant variation.

It is important to note that exposure does not necessarily imply automation. A task may be technically exposed to LLM capabilities without being automated in practice, whether due to organizational inertia, implementation costs, or worker resistance. Our exposure measures capture potential automation based on task characteristics instead of realized adoption. The empirical effects we estimate therefore reflect labor market responses to perceived automation potential rather than actual task displacement.

2.2 Task-Level Expertise Measure

The exposure measures described above capture *how much* of an occupation is exposed to AI automation, but they say nothing about *which* tasks are exposed. Two occupations with identical exposure scores may face very different labor market consequences if one loses its most demanding tasks while the other loses routine work. To distinguish between these scenarios, we develop a measure of task-level expertise that allows us to characterize how automation reshapes an occupation's skill requirements.

Autor and Thompson (2025) define expertise as a hierarchical capability where more expert workers can perform all tasks that less expert workers can perform, but not vice versa. Hence, expertise functions as a barrier to entry. High-expertise tasks are those that require specific training, certification, or accumulated know-how. Low-expertise tasks are generic activities that most adults can perform without specialized preparation.

To measure task-level expertise, we use GPT-4o to classify each of the 19,265 task descriptions in O*NET version 27.2. The full prompt is reproduced in Appendix A1. For each task, the model receives the task description along with its corresponding occupation title and returns an expertise score on a 0–100 scale. The prompt defines three tiers: low expertise (0–20) for generic tasks, medium expertise (21–60) for tasks that require some specialized skill, but that do not define the primary barrier to entry, and high expertise (61–100) for specialized tasks that determine who can enter the occupation. The prompt explicitly instructs the model to consider occupational context, as the same task description may require different expertise levels depending on where it sits within an occupation's task bundle. In addition, we explicitly incorporate the hierarchical definition of expertise from Autor and Thompson (2025)

The measure produces rankings that accord with intuition. Among the highest-expertise occupations are anesthesiologists, physicists, magistrate judges, and nuclear engineers, all of which require years of specialized training and credentialing. At the other extreme sit dishwashers, packers, baggage porters, and locker room attendants, occupations whose tasks most adults can perform without prior preparation. More interesting are the predicted expertise changes after removing β-exposed tasks. Treasurers, medical transcriptionists, and network administrators see large declines, as LLMs can perform much of their specialized cognitive work.

Patient representatives, pediatricians, and human resources specialists see increases, as automation removes their routine administrative tasks while leaving judgment-intensive responsibilities intact.

Face validity, however, is insufficient. We therefore also examine if our LLM-based ratings correlate with independent labor market indicators that reflect, but were not used to construct, those requirements. We consider workers' educational attainment, wages, and total compensation. To do so, we match O*NET occupations to Revelio Labs data drawn from LinkedIn profiles and estimate occupation-level regressions of each labor market outcome on our expertise measure. The results are in Appendix A2. We find a strong and consistent correlation between expertise and observed labor market outcomes. Across 38 countries, a one-standard-deviation increase in baseline expertise is associated with 0.43 standard deviations higher educational attainment, 0.53 standard deviations higher salary, and 0.42 standard deviations higher total compensation.[1] Every country-specific coefficient is positive and statistically significant.

A separate concern is whether our results depend on idiosyncratic features of the prompt we use. We address this in two ways. First, we re-run the classification exercise using an alternative prompt that frames expertise differently (see Appendix A2.1). Second, we compare our rankings to the independently constructed expertise measure of Hosseini Maasoum and Lichtinger (2026), who use a substantively different prompt and aggregation procedure. In both cases, we compute importance-weighted mean expertise at the occupation level and examine rank correlations between our main measure and these alternative measures. The Spearman correlation between our baseline and alternative-prompt measures is 0.914; between our measure and Hosseini Maasoum and Lichtinger's, 0.844. These correlations suggest that occupation-level expertise rankings are robust to reasonable variation in prompt design.

We combine our task-level expertise scores with the LLM exposure classifications to construct our main treatment variable: the predicted change in an occupation's expertise requirements under automation. For each occupation, we compute the importance-weighted mean expertise across all tasks, then recompute this mean after

[1]The United States is omitted from this validation analysis due to computational constraints. Given the consistency of results across all other countries, however, its inclusion would unlikely change our conclusions.

excluding tasks classified as exposed to LLMs. For this calculation, we treat any task with $\beta_i > 0$ as exposed. That is, tasks with either direct exposure ($\beta_i = 1$) or exposure requiring complementary tools ($\beta_i = 0.5$) are assumed to be automatable. The difference between these two quantities captures how automation would reshape the occupation's expertise profile if all exposed tasks were automated.

Formally, the change in expertise is:

$$\Delta\text{Expertise}_o = \frac{\sum_{i \in T_o^{\text{unexposed}}} w_i \cdot \text{Expertise}_i}{\sum_{i \in T_o^{\text{unexposed}}} w_i} - \frac{\sum_{i \in T_o} w_i \cdot \text{Expertise}_i}{\sum_{i \in T_o} w_i} \tag{2}$$

where T_o denotes all tasks in occupation o, $T_o^{\text{unexposed}} = \{i \in T_o : \beta_i = 0\}$ denotes tasks not exposed to LLMs, Expertise_i is task i's expertise score, and w_i is its importance weight. A positive value indicates that automation removes disproportionately low-expertise tasks, whereas a negative value indicates that automation removes high-expertise tasks.[2]

The correlation between β exposure and predicted expertise change is -0.39, which suggests that occupations with higher LLM exposure tend to experience expertise-lowering automation. This negative correlation reflects the fact that LLMs are particularly capable at language-intensive cognitive tasks, which often represent the high-expertise components of exposed occupations. Nevertheless, substantial variation exists within exposure levels, providing the identifying variation we exploit in our empirical analysis.

One concern is that different occupations may adopt AI at different rates, and these adoption rates could correlate with our expertise change measure. Our measures capture which tasks *could* be automated based on technical capabilities, not which tasks *have been* automated in practice. If high-expertise-change occupations systematically adopted LLMs faster or slower than low-expertise-change occupations, this could confound our estimates. We note, however, that differential adoption would primarily affect the timing and magnitude of effects rather than the qualitative pattern. The expertise framework predicts different *directions* of effects depending on which tasks are automated—expertise-raising automation should in-

[2]We focus on changes in mean expertise throughout the paper. Appendix A5 reports results using changes in minimum and maximum expertise, though these measures are extremely noisy, because the they depend on the correct expertise classification of a single task. For average expertise, most measurement error will even out. Hence, we only report the minimum and maximum measures for completeness.

crease wages while expertise-lowering automation should decrease them—and these directional predictions would be difficult to explain through adoption rates alone.

A major advantage of our approach is that task-level expertise scores can be combined with any measure of automation exposure. Because we score expertise at the task level independently of exposure classifications, the same expertise ratings can be paired with different assumptions about which tasks are automatable. Our main analyses use β exposure from Eloundou et al. (2024), but we also examine robustness to α exposure (their more conservative measure) and to the 2023 and 2025 exposure scores from Gmyrek et al. (2025).[3] This flexibility allows us to assess whether our findings depend on particular assumptions about which tasks LLMs can perform.

2.3 Job Posting Data

Our labor market data comes from Revelio Labs, which aggregates and processes publicly available employment information from job boards, company career pages, and recruitment platforms using machine learning algorithms. The dataset contains the near-universe of online job postings across 39 countries. Our analysis window spans August 2021 through October 2025, covering 15 months before and 36 months after the release of ChatGPT in November 2022.

Each posting includes the job title, employer, posting and removal dates, salary information (when available), and geographic location. We match each vacancy to O*NET SOC codes and ISCO codes that correspond directly to the occupations in Eloundou et al. (2024); Gmyrek et al. (2025). This allows us to link job-level observations to our occupation-level exposure and expertise measures.

For each country, we construct an occupation-by-month panel, retaining only occupations that have observations in all months of our sample period to ensure a balanced panel within each country. Our two main outcome variables are the total job posting volume, measured as the log of total new vacancies posted within each occupation-month cell, and the mean advertised annual salary for advertised positions within each cell.[4]

[3]For the Gmyrek et al. (2025) exposure measures, we separately apply our expertise prompt to the tasks in their database, which are defined at the international ISCO occupation level rather than the U.S.-specific O*NET classification.

[4]When salary is not explicitly listed, Revelio Labs infers compensation using a predictive model

Several features of online job posting data warrant discussion. First, job postings represent labor demand from employers who advertise positions online, which may differ from total labor demand in the economy. Occupations and industries that rely more heavily on informal recruitment channels (e.g., personal networks, walk-in hiring) will be underrepresented in our data. This selection is likely to vary across countries, with online posting more prevalent in high-income countries and formal sector employment. Our results therefore primarily apply to the segment of the labor market where employers actively advertise vacancies online. Second, a job posting reflects an employer's intention to hire rather than an actual hire, and the relationship between postings and employment may vary across occupations and over time. Third, our sample period of 2021–2025 follows the COVID-19 pandemic, which caused substantial disruption to labor markets worldwide. While our difference-in-differences design absorbs aggregate trends, we cannot fully rule out that pandemic-related adjustments differentially affected occupations by exposure or expertise change.

Despite these limitations, job posting data offer some major methodological advantages for studying labor market effects of AI. The data provide high-frequency observations at monthly intervals, allowing us to trace the dynamics of labor market responses before and after ChatGPT's release with greater temporal precision than annual surveys or census data. Job postings also capture employer demand in near real-time, reflecting hiring intentions before they materialize in employment statistics; this forward-looking quality is particularly valuable for studying a rapidly diffusing technology. Furthermore, the granular occupational coding enables direct matching to task-level exposure measures, which would be difficult with more aggregated data sources. Finally, the international coverage of our data permits the first systematic cross-country analysis of labor market effects of AI, providing external validity that single-country studies cannot offer.

2.4 Data Descriptives

Table 1 presents summary statistics for our main variables. Panel A describes the sample composition: 39 countries with an average of 473 occupations per country

trained on visa applications, self-reported data, and job postings with explicit salary information. The share of salaries that are imputed versus directly observed varies by country and occupation; measurement error from imputation likely attenuates our salary estimates toward zero.

Table 1: Summary Statistics

Variable	N	Mean	SD	P10	P50	P90
Panel A: Sample Composition						
Countries	39	—	—	—	—	—
Occupations per country	473	473.00	244.00	—	—	—
Months in sample	51	—	—	—	—	—
Occupation-month observations	939,382	—	—	—	—	—
Panel B: Treatment Variables (Occupation-Level)						
AI Exposure (β)	849	0.34	0.22	0.04	0.36	0.61
ΔExpertise	769	-2.59	7.69	-11.60	-0.42	3.38
Corr(Exposure, ΔExpertise)	—	-0.39	—	—	—	—
Panel C: Outcomes—Pre-Treatment (Aug 2021–Oct 2022)						
Job postings (count)	276,570	968	13,129	7	56	1,094
Log job postings	276,570	4.27	2.00	1.95	4.03	7.00
Advertised salary (USD)	276,564	44,464	22,495	16,649	42,361	73,089
Panel D: Outcomes—Post-Treatment (Nov 2022–June 2025)						
Job postings (count)	662,812	1,139	8,488	12	86	1,506
Log job postings	662,812	4.69	1.91	2.48	4.45	7.32
Advertised salary (USD)	662,812	45,549	22,709	16,994	43,558	74,733

Notes: This table reports summary statistics for the main variables. Panel A describes the sample composition across 39 countries. Panel B reports occupation-level treatment variables: AI exposure (β) measures the share of O*NET tasks exposed to LLM automation; ΔExpertise measures the change in mean (or minimum) expertise requirements if exposed tasks were automated. Panel C and D report outcome variables in the pre-treatment (August 2021–October 2022) and post-treatment (November 2022–June 2025) periods, respectively. All salaries are in USD.

observed over 51 months (August 2021 through October 2025), yielding approximately 939,000 occupation-month observations. Panel B reports the occupation-level treatment variables. The β exposure measure, which gives the share of tasks in an occupation exposed to LLM automation, has a mean of 0.34 (SD = 0.22), ranging from near-zero for occupations like construction laborers to over 0.60 for occupations like web developers and data entry keyers. This heterogeneity provides one of the main sources of variation for our identification strategy.

Our expertise change measure (ΔExpertise) has a mean of -2.59 (SD = 7.69). A positive value indicates that automation removes relatively low-expertise tasks, raising the occupation's expertise threshold; a negative value indicates that automation removes high-expertise tasks, lowering the threshold. The correlation between exposure and expertise change is -0.39, indicating that more exposed occupations tend toward expertise-lowering automation, though substantial variation exists within exposure levels.

Panels C and D report outcome variables in the pre-treatment (August 2021–October 2022) and post-treatment (November 2022–October 2025) periods. Mean job postings per occupation-month increased from 968 to 1,139, while median advertised salaries rose from \$42,361 to \$43,558. The substantial variation in job posting volume (standard deviation of 13,129 in the pre-period) reflects differences across occupations in hiring intensity and across countries in labor market size and online posting prevalence.

Table 2 presents example occupations in each quadrant of the exposure–expertise-change space. Low-exposure occupations with positive expertise change tend to require physical presence, and often operate in unpredictable environments. Wind turbine technicians diagnose failures hundreds of feet above ground; commercial divers navigate murky underwater conditions; surgical technologists assist in complex operations with lots of variation. Among high-exposure occupations, occupations with positive predicted expertise changes are human resources specialists and information security analysts. Their routine administrative tasks are exposed while their more judgment-intensive work cannot yet be done by AI.

When looking at occupations with negative expertise changes, we observe that welders and construction laborers see their most skilled tasks exposed while repetitive physical work persists. High-exposure occupations like marketing managers and executive secretaries similarly experience expertise-lowering automation, as their

Table 2: Illustrative Occupations by AI Exposure and Expertise Change

	Low AI Exposure ($\beta < 0.10$)	**High AI Exposure** ($\beta > 0.50$)
Expertise-Raising ($\Delta E > 0$)	Industrial Truck and Tractor Operators (β=0.05, ΔE=1.5); Hairdressers, Hairstylists, and Cosmetologists (β=0.07, ΔE=2.7); Locker Room, Coatroom, and Dressing Room Attendants (β=0.04, ΔE=0.0)	Human Resources Specialists (β=0.64, ΔE=14.2); Purchasing Managers (β=0.58, ΔE=7.3); Patient Representatives (β=0.50, ΔE=17.5)
Expertise-Lowering ($\Delta E < 0$)	Fast Food and Counter Workers (β=0.08, ΔE=-0.7); Food Preparation Workers (β=0.05, ΔE=-0.3); Welders, Cutters, Solderers, and Brazers (β=0.09, ΔE=-2.4)	First-Line Supervisors of Retail Sales Workers (β=0.50, ΔE=-23.0); Marketing Managers (β=0.50, ΔE=-2.4); Executive Secretaries and Executive Administrative Assistants (β=0.63, ΔE=-8.4)

Notes: This table presents example occupations in each quadrant of the AI exposure–expertise change space. AI exposure (β) measures the share of tasks in an occupation that are exposed to LLM automation. ΔE measures the change in mean expertise requirements if exposed tasks were automated; positive values indicate that automation removes low-expertise tasks, raising the occupation's expertise threshold. Examples selected based on data availability across countries.

most cognitively demanding tasks are precisely those within LLM capabilities.

3 Empirical Methodology

Our empirical strategy tests whether labor market responses to automation depend on *which* tasks get automated. At a given level of automation exposure, occupations may experience divergent labor market trajectories depending on whether automation displaces high-expertise or low-expertise tasks. We identify these effects using a difference-in-differences design that exploits the release of ChatGPT in November 2022 as a shock to automation potential, combined with pre-existing cross-sectional variation in occupations' LLM exposure and task expertise composition.

Our main specification examines the separate effects of AI exposure and predicted expertise change on labor market outcomes. We estimate the model separately for each country c:

$$Y_{ot}^{c} = \beta_1^{c}(\text{Post}_t \times \text{Exposure}_o) + \beta_2^{c}(\text{Post}_t \times \Delta\text{Expertise}_o) + \gamma_o^{c} + \delta_t^{c} + \varepsilon_{ot}^{c} \quad (3)$$

where Y_{ot}^{c} denotes the outcome variable (log job postings or log salary) for occupation o at month t in country c, Post_t is an indicator for months after November 2022, Exposure_o is the fraction of tasks in occupation o susceptible to LLM automation, and $\Delta\text{Expertise}_o$ captures the predicted change in mean expertise requirements after removing exposed tasks. We include occupation fixed effects γ_o^{c} to control for time-invariant occupation characteristics and 2-digit O*NET occupation codes interacted with year-by-month indicators δ_t^{c} to absorb sector-specific trends. This fixed effects structure is demanding: it compares occupations within the same broad occupational category in the same month, identifying effects solely from variation in exposure and expertise change among occupations that share similar sector-level trends. This approach controls away a substantial amount of variation, including any aggregate effects of AI on entire sectors. Consequently, if we detect effects under this specification, the true aggregate effects are likely larger. Standard errors are clustered at the occupation level. We restrict our analysis to OECD and BRICS countries and exclude countries with insufficient data coverage (China, Ethiopia, Iran, Egypt) or data quality issues (e.g. Iceland).

The coefficient β_1^{c} captures the effect of LLM exposure on labor market outcomes,

holding constant the change in expertise requirements. The coefficient β_2^c identifies whether occupations where automation changes expertise requirements experience differential labor market effects, conditional on their overall exposure level. We standardize both variables so that they have mean zero and unit standard deviation within each country. Coefficients thus represent the effect of a one country-level standard deviation increase in the respective measure.

Identification relies on two main assumptions. First, conditional on our fixed effects structure, occupations with different levels of exposure and expertise change would have followed parallel trends in labor market outcomes absent the release of ChatGPT. The coefficient β_1^c is identified by comparing the post-November 2022 trajectories of high-exposure versus low-exposure occupations within the same 2-digit sector, while β_2^c is identified by comparing occupations where automation removes high-expertise tasks against those where it removes low-expertise tasks. Second, ChatGPT's release must not represent a response to pre-existing labor market trends. The timing of the release, arguably determined by OpenAI's product development cycle rather than labor market conditions, and its broad, non-targeted nature support this assumption.

We assess the parallel trends assumption through event study specifications that allow us to examine whether high-exposure and high-expertise-change occupations were already on divergent trajectories prior to November 2022:

$$Y_{ot}^c = \sum_{k=-15}^{31} [\beta_{1k}^c \mathbf{1}_{t=k} \text{Exposure}_o + \beta_{2k}^c \mathbf{1}_{t=k} \Delta\text{Expertise}_o] + \gamma_o^c + \delta_t^c + \varepsilon_{ot}^c \tag{4}$$

where k indexes months relative to November 2022 and month $k = -1$ serves as the reference period. If occupations with different exposure levels and expertise changes were on parallel trajectories before November 2022, these coefficients should be close to zero.

As an additional analysis, we extend the main specification to include an interaction term between exposure and expertise change:

$$\begin{aligned} Y_{ot}^c = \beta_1^c(\text{Post}_t \times \text{Exposure}_o) + \beta_2^c(\text{Post}_t \times \Delta\text{Expertise}_o) \\ + \beta_3^c(\text{Post}_t \times \text{Exposure}_o \times \Delta\text{Expertise}_o) + \gamma_o^c + \delta_t^c + \varepsilon_{ot}^c \end{aligned} \tag{5}$$

The interaction coefficient β_3^c tests whether the impact of AI exposure differs depending on which tasks are automated. With standardized variables, the total effect of exposure on outcomes equals $\beta_1^c + \beta_3^c \times \Delta\text{Expertise}_o$. A positive β_3^c therefore implies that the effect of exposure is more favorable (less negative or more positive) when $\Delta\text{Expertise}_o$ is higher—that is, when automation removes low-expertise tasks rather than high-expertise tasks.

Estimating separate models for each of 39 countries yields country-specific coefficient estimates $\{\beta_1^c, \beta_2^c, \beta_3^c\}$. After estimating country-specific treatment effects, we aggregate results using a random effects meta-analysis framework that treats the true treatment effect as varying across countries. We estimate between-country heterogeneity using restricted maximum likelihood (REML) and compute the pooled estimate as a precision-weighted average of country-specific estimates (Viechtbauer, 2005). We present results using forest plots that display country-specific point estimates with 95% confidence intervals alongside the pooled random effects estimate. The heterogeneity statistic I^2 indicates the proportion of total variation attributable to between-country differences instead of sampling error. As robustness checks, we also estimate fixed effects meta-analyses that weight countries by their total volume of pre-treatment job postings and restrict the sample to countries where pre-treatment event study coefficients are jointly insignificant at the 5% level for all treatment variables. Section 4.1 explores which country-level factors may explain the heterogeneity in effects we observe.

Several potential threats to identification warrant discussion. First, the November 2022 release of ChatGPT coincided with several macroeconomic developments, including the Federal Reserve's aggressive interest rate increases, widespread layoffs in the technology sector, and ongoing post-pandemic labor market adjustments. To deal with these concerns, we include sector-by-time fixed effects, which absorb all observed and unobserved factors that affect a certain sector. In other words, we only compare occupations with different exposure levels and expertise change measures *within the same sector*. Additionally, for macroeconomic confounders to explain our findings, they would need to differentially affect not just high-exposure occupations but specifically those with particular patterns of expertise change, which is a correlation structure that seems unlikely to arise from general macroeconomic conditions.

Third, our identification treats November 2022 as an unanticipated shock, but AI

capabilities were developing well before ChatGPT's release. GPT-3 was announced in June 2020, and GitHub Copilot launched in October 2021, potentially allowing sophisticated employers to anticipate LLM capabilities. However, ChatGPT was the first LLM with a consumer-facing interface that demonstrated capabilities to a mass audience, and its release was not pre-announced; moreover, anticipation effects would most plausibly attenuate our estimates toward zero by spreading labor market adjustments over a longer period, suggesting our findings represent lower bounds on the true effects.

Fourth, both our exposure and expertise change measures are constructed from LLM-based classification of task descriptions, introducing measurement error. Classical measurement error in the treatment variables would attenuate our estimates toward zero. We partially address this concern through validation exercises showing that our expertise measure correlates strongly with independently observed occupation characteristics such as worker education and wages, but we acknowledge that some measurement error likely remains.

4 Results

The expertise framework generates two testable predictions. First, occupations where AI automates low-expertise tasks should experience wage increases, as fewer people will have the ability to execute the tasks of such a job, leading to a contraction of labor supply. Second, because of the decrease in labor supply, employment in occupations with increased expertise requirements should go down.

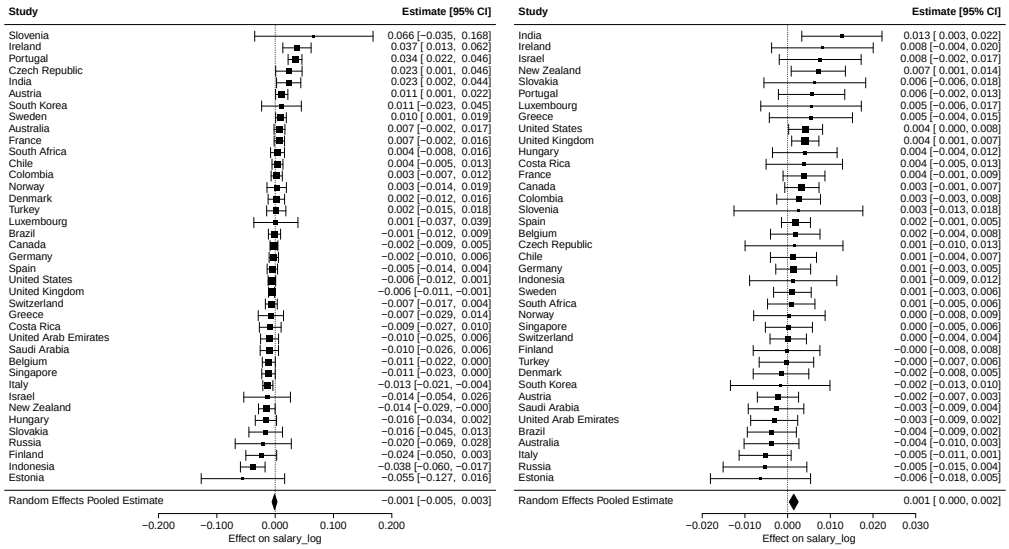
We begin with estimating the main specification (Equation (3)) to examine the effects of AI exposure and predicted expertise change on salaries and job postings. All coefficients represent the effect of a one standard deviation increase in each treatment variable.[5]

Figure 1 presents results for advertised salaries. Panel A shows that AI exposure has no significant direct effect on wages, suggesting that the quantity of tasks exposed to automation does not, by itself, predict wage changes. Expertise change, by contrast, does matter. Panel B shows that occupations where exposure is concentrated in low-expertise tasks experience an increase in advertised salaries compared

[5]standardization is done at the country-level, so there may be variation in the size of a standard deviation.

Figure 1: Meta-Analysis: Effect on Salary (Main Specification)

(a) AI Exposure (β) (b) Expertise Change



Notes: Forest plots display country-specific coefficient estimates with 95% confidence intervals. The diamond represents the pooled random effects estimate. All treatment variables are standardized to have mean zero and unit standard deviation within each country. Coefficients represent the estimated effect of a one standard deviation increase in the respective treatment variable on log average advertised salary, comparing the post-ChatGPT period (December 2022 onward) to the pre-period. Regressions include occupation fixed effects and 2-digit O*NET by year-month fixed effects. Standard errors are clustered at the occupation level.

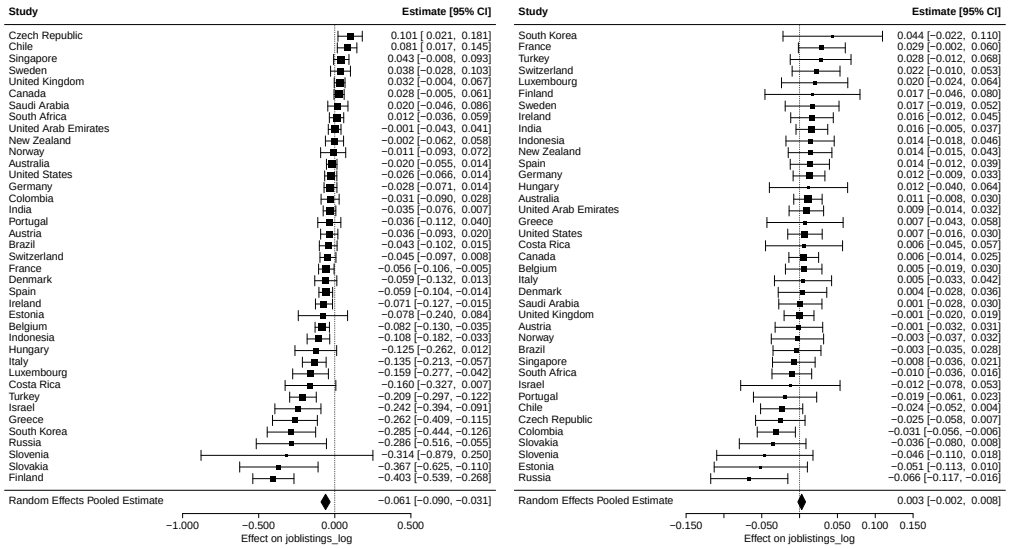
to occupations where exposure is in high-expertise tasks. Though the magnitude is modest—roughly 0.1 percent per standard deviation—the effect is consistent with theory.[6]

Figure 2 presents the results for job postings. Consistent with prior research by Brynjolfsson et al. (2025); Lichtinger and Hosseini Maasoum (2025); Klein Teeselink (2025); Azar et al. (2025), the results in Panel (a) demonstrate that occupations with higher LLM exposure experience substantial declines in hiring demand. A one standard deviation increase in exposure corresponds to a 6.1 percent decline in job postings ($p < 0.001$). The direction of this effect is consistent across countries, as 31 out of 39 countries show negative point estimates, with 14 statistically signifi-

[6]The estimated effect of expertise change becomes stronger when we estimate models that only include the expertise change variable without additionally adding the AI exposure variable. See Figure A11 in the Appendix.

Figure 2: Meta-Analysis: Effect on Job Postings (Main Specification)

(a) AI Exposure (β) (b) Expertise Change



Notes: Other details are as in Figure 1. Coefficients represent the effects on log job postings.

cant at the 5% level. The largest effects appear in Finland, Slovakia, and Slovenia. This pattern suggests that previously observed displacement from AI-exposed occupations is unlikely to be explained by national-level shocks that correlate with AI exposure. As such, this finding provides an important external validity check for previous research that focuses on one-country studies. Section 4.1 explores this heterogeneity in more detail.

Panel (b) shows the effect of predicted expertise change on job postings. We find a small, statistically insignificant effect. This null result is not consistent with the expertise framework, which predicts that increases in expertise requirements should reduce employment by shrinking the pool of qualified workers. One possible explanation is that rising expertise requirements generate both hiring and separation pressures. Workers whose skills fall below the new threshold may be displaced, forcing firms to recruit replacements even as the pool of qualified candidates shrinks. This churn may offset the reduced hiring that supply constraints and higher wages would otherwise produce.

The validity of our difference-in-differences design depends on the validity of the

parallel trends assumption: absent ChatGPT's release, occupations with different exposure levels and expertise changes would have followed similar trajectories. We assess this assumption by estimating the event study specification in Equation (4) for each country. Appendix A4 presents the full set of country-level event study plots, showing the month-by-month coefficients for both treatment variables across all four outcome-treatment combinations. The pre-treatment coefficients are generally close to zero and exhibit no systematic trends, providing little evidence of parallel trend violations. The absence of pre-trends is particularly notable given the large number of countries in our sample, as any systematic violation would likely manifest in at least some national labor markets. In Section 4.2, we show the robustness of our results to only including countries without any pre-trend violations.

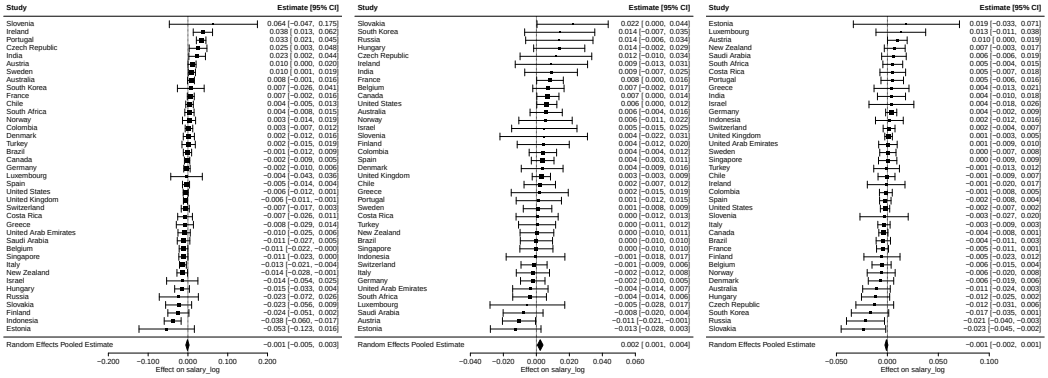
Next, we consider the interaction specification (Equation (5)). The reason this is interesting, is that exposure measures what AI could automate, but firms only reduce hiring when they can actually substitute away from labor. When expertise requirements rise, the remaining workers become harder to replace and firms cannot easily act on the automation potential because the job still requires scarce human skills. The interaction captures that the bite of exposure might depends on whether the remaining work is substitutable.

Figure 3 presents the interaction specification for salaries. Similar to our main results, we find no direct effect of exposure on salaries, and a positive and significant effect of predicted expertise change. The interaction term is indistinguishable from zero, suggesting that the wage premium from rising expertise requirements does not depend on the level of AI exposure. In other words, what seems to matter for wages is which tasks are automated, regardless of how many tasks are exposed.

Figure 4 presents the interaction model for job postings. The interaction term is positive and statistically significant, indicating that the negative effect of exposure on hiring is attenuated when automation raises expertise requirements. Occupations where AI targets low-expertise tasks experience a 4.9 percent decline in job postings, compared to 7.3 percent where AI targets high-expertise tasks. This finding suggests that exposure only fully translates into displacement when firms can substitute away from labor. When expertise requirements rise, the remaining work demands scarce human skills that AI cannot replicate, limiting firms' ability to act on the automation potential. When expertise requirements fall, the remaining tasks are more generic and substitutable, giving firms both the capability and the incentive to reduce hiring.

Figure 3: Meta-Analysis: Salary with Interaction Term

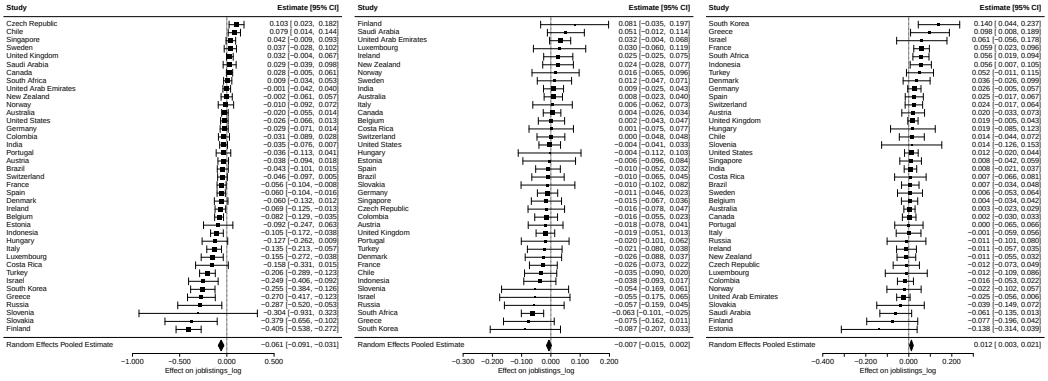
(a) AI Exposure (β) (b) Expertise Change (c) Interaction



Notes: Forest plots display coefficients from the interaction specification (Equation (5)). Panel (a) shows the main effect of AI exposure, Panel (b) shows the main effect of expertise change, and Panel (c) shows the interaction between exposure and expertise change. All other definitions are as in Figure 1.

Figure 4: Meta-Analysis: Job Postings with Interaction Term

(a) AI Exposure (β) (b) Expertise Change (c) Interaction



Notes: Other details are as in Figure 3. Coefficients represent effects on log job postings.

4.1 Cross-Country Heterogeneity

Prior evidence on the effect of AI exposure on employment predominantly comes from single-country studies (Brynjolfsson et al., 2025; Lichtinger and Hosseini Maasoum, 2025; Klein Teeselink, 2025; Humlum and Vestergaard, 2025). A unique feature of our analysis is the ability to estimate AI exposure effects across 39 countries,

which serves two purposes. First, it allows us to assess external validity: do the displacement effects documented in individual countries generalize broadly, or do they reflect idiosyncratic features of particular labor markets? Second, it enables us to examine which institutional and economic characteristics are associated with larger or smaller reductions in hiring, potentially illuminating the mechanisms through which AI affects employment.[7] As Figure 2, Panel A reveals, exposure effects on job postings vary considerably across countries, ranging from large negative effects to near-zero or even positive coefficients. We investigate what explains this variation.

We consider four potential explanations. First, labor market tightness may amplify displacement effects: in countries where firms struggled to fill vacancies prior to ChatGPT's release, AI offers an alternative to costly recruitment, making firms more willing to forgo hiring in exposed occupations. Second, stricter employment protection legislation could amplify hiring freezes: firms in highly exposed occupations face substantial uncertainty about whether AI will eventually automate much of their workforce's tasks, and when dismissal costs are high, the prudent response is to avoid hiring workers who may prove difficult to dismiss if their roles become redundant. Third, higher rates of AI adoption might accelerate displacement as firms more rapidly integrate new technologies into production processes. Fourth, greater digital readiness, which captures a country's infrastructure, skills, and institutional capacity to use digital technologies productively, could either amplify displacement through faster AI integration or attenuate it if complementary digital capabilities facilitate worker reallocation and augmentation.

To explore these potential moderators, we correlate each country's estimated exposure effect on log job postings with four country-level characteristics. For labor market tightness, we use 2022 unemployment rates from the World Bank's World Development Indicators, which compiles ILO-modeled estimates (The World Bank, 2024). For employment protection, we use the OECD's Employment Protection Legislation index, which measures the strictness of regulations governing individual and collective dismissals (Organisation for Economic Co-operation and Development, 2026). For AI adoption, we use the index contructed by Meleshenko (2023), which is a composite measure of national AI talent, infrastructure, operating en-

[7] Our estimates condition on expertise change, which differs slightly from most prior studies that examine exposure effects alone. The country-level coefficients we analyze therefore capture the effect of exposure holding expertise change constant.

vironment, research, development, and government strategy. For digital readiness, we use the IMD World Digital Competitiveness Ranking, which assesses countries' capacity to adopt and explore digital technologies for economic and social transformation (International Institute for Management Development, 2024).

Figure 5 presents the four correlation plots, with fitted lines showing precision-weighted linear regressions. Countries with stricter employment protection legislation exhibit stronger negative effects of AI exposure on hiring. This pattern is consistent with the hypothesis that high firing costs shift technological adjustment to the hiring margin: rather than dismissing workers whose tasks become automatable, firms in rigid labor markets reduce new hires, leading to larger displacement effects in job postings data. This finding echoes Bartelsman et al. (2016), who show that strict employment protection discourages firms from adopting risky new technologies because the inability to flexibly adjust the workforce makes experimentation with uncertain innovations too costly. Digital readiness correlates positively with the exposure effect, meaning that more digitally prepared countries experience smaller hiring reductions. This pattern suggests that digital infrastructure and complementary capabilities enable firms and workers to adapt more effectively to AI-driven changes, potentially attenuating displacement through smoother reallocation or more effective human-AI complementarity.

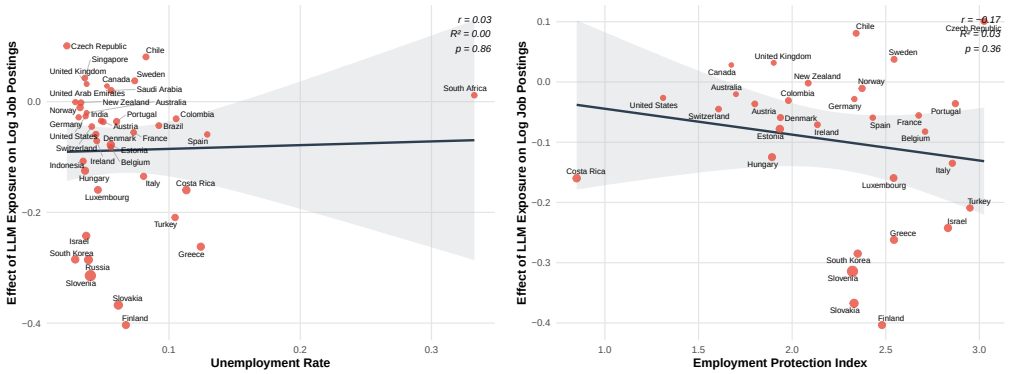
Prior labor market slack, as measured by unemployment, shows little correlation with the magnitude of displacement effects, suggesting that labor market tightness does not clearly moderate the impact of AI on hiring. The correlation with AI adoption is also weak, indicating that national adoption intensity alone does not determine the magnitude of displacement.

It is important to stress that these correlations should be interpreted cautiously. With observations limited to the number of countries in our sample, statistical power is limited, and the possibility of confounding factors cannot be ruled out. In other words, we present these results as exploratory evidence suggesting potential mechanisms rather than causal claims about the determinants of cross-country heterogeneity.

Figure 5: Correlates of Cross-Country Heterogeneity in AI Exposure Effects

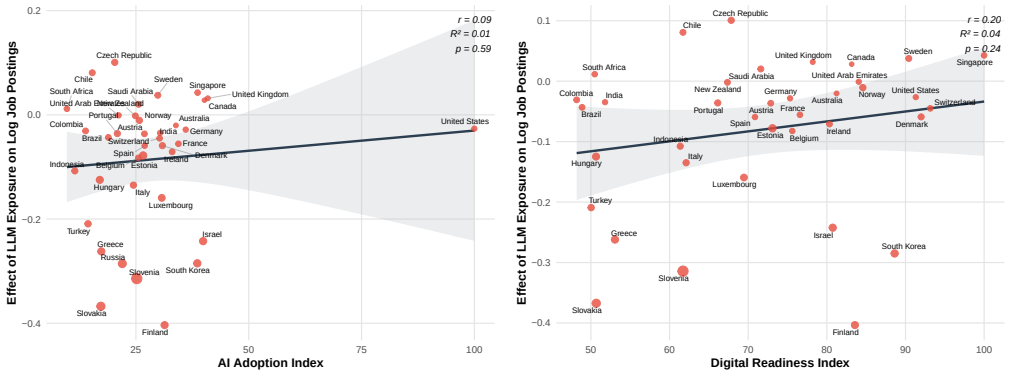
(a) Unemployment Rate (2022)

(b) Employment Protection Legislation



(c) AI Adoption Index

(d) Digital Readiness Index



Notes: Each panel plots the country-level coefficient on AI exposure (from Figure 2) against a country characteristic. The x-axis shows the effect of a one standard deviation increase in AI exposure on log job postings (negative values indicate displacement). Panel (a) uses 2022 unemployment rates from the ILO. Panel (b) uses the OECD Employment Protection Legislation index. Panels (c) and (d) use AI adoption and digital readiness indices. Point sizes are proportional to the inverse of the standard error (precision-weighted). The fitted line shows a linear regression weighted by inverse variance.

4.2 Robustness Checks

We assess sensitivity to alternative specifications and measurement choices across four dimensions.

Our main analysis uses the β exposure measure from Eloundou et al. (2024), which captures tasks that could be automated by LLMs augmented with complementary technologies such as code interpreters, web browsers, or specialized plugins. A concern is that this broader measure may overstate near-term automation poten-

tial by including tasks that require tools not yet widely deployed. We therefore re-run our analysis using the more restrictive α measure, which captures only tasks automatable by standalone LLMs without complementary tools, as measured by Eloundou et al. (2024). Because α identifies a substantially narrower set of exposed tasks (roughly 15 percent of all tasks compared to approximately 55 percent under the β measure), we expect smaller effects due to reduced treatment intensity. Consistent with this prediction, the exposure effects of job postings remain statistically significant but are attenuated in magnitude (Figure A6): the job postings effect is $\beta = -0.034$ ($p < 0.001$) versus -0.061 in the main specification. The positive expertise change effects on salary becomes insignificant under α. These results suggest that the small fraction of exposed tasks in α is insufficient to produce major effects for the expertise change channel.

A separate concern is that our O*NET-based exposure measures, which derive from the U.S. occupational classification system, may not translate well to labor markets in which occupations consist of different tasks than in the US. We address this concern by using the exposure measures from Gmyrek et al. (2025), which are constructed at the ISCO (International Standard Classification of Occupations) level specifically designed for cross-country comparability. We consider both the 2023 and 2025 versions of the exposure index.[8] For job postings, the results are highly robust (Figures A7 and A8): the 2023 version yields an exposure effect of $\beta = -0.047$ ($p < 0.001$), while the updated 2025 version yields an even larger effect of $\beta = -0.097$ ($p < 0.001$). Consistent with our main finding that exposure alone does not predict wage changes, salary effects of LLM exposure under the ISCO measures are not statistically significant ($\beta = -0.002$, $p = 0.14$ for 2023; $\beta = -0.002$, $p = 0.27$ for 2025). For the expertise change channel, the expertise change coefficients are positive but not statistically significant under either measure ($\beta = 0.001$, $p = 0.16$ for 2023; $\beta = 0.001$, $p = 0.24$ for 2025). For job postings, both the 2023 and 2025 measures yield a significant positive effect of expertise change ($\beta = 0.020$, $p < 0.001$ for 2023; $\beta = 0.026$, $p < 0.001$ for 2025).

[8]To construct predicted expertise change for the Gmyrek measures, we apply our LLM-based expertise classification to the ISCO task descriptions. We calculate pre-automation expertise as the unweighted mean of task-level expertise scores within each occupation. Post-automation expertise is calculated as a weighted mean, where each task receives weight (1 − exposure score), so that highly exposed tasks contribute less to the post-automation average. Expertise change is the difference between post- and pre-automation expertise.

Next, we consider the robustness of our results to potential violations of the parallel trends assumption. To do so, we re-estimate our study specifications for the subset of countries in which the average of the pre-trend coefficients for both the exposure measure and the expertise change are not statistically different from zero. This brings our sample to 33 countries for job postings and 27 for salary. The results are robust to this restriction (Figure A9): the job postings exposure effect remains significant ($\beta = -0.049$, $p = 0.001$), as does the positive expertise effect on salaries ($\beta = 0.001$, $p = 0.03$). Effect sizes are slightly smaller in this subsample, but the conclusions are the same.

Last, our main meta-analysis weights countries by estimation precision (inverse variance weighting), which is standard practice but may overweight small countries. The reason is that each country has roughly the same number of observations (one per occupation per month). Hence, country-level precision is not necessarily a function of country size. As an alternative, we weight countries by their pre-treatment job posting volume, which better reflects economic significance and ensures that large labor markets receive commensurate influence in the pooled estimates. Figure A10 shows the results. The conclusions are broadly similar, although the magnitude changes somewhat, as the positive effect of expertise change on salaries becomes larger, whereas the negative effect of exposure on job postings becomes smaller. As such, it appears that our findings are not driven by outsized influence of smaller countries.

5 Conclusion

This paper provides the first large-scale empirical test of the expertise framework proposed by Autor and Thompson (2025) to understand how AI is reshaping labor markets. The framework's central insight is that the labor market effects of AI depend not only on how many tasks are automated, but on which tasks. When AI automates low-expertise tasks, the remaining work requires greater specialization, raising the expertise threshold and pushing wages up as fewer workers qualify. When AI automates high-expertise tasks, the occupation becomes more accessible, expanding labor supply and pushing wages down.

Using job posting data from 39 countries and exploiting the release of ChatGPT

in November 2022 as a natural experiment, we find measured support for these predictions. Occupations where low-expertise tasks are exposed to automation experience an increase in advertised salaries compared to occupations where exposure is concentrated in high-expertise tasks. Although we do not find an effect of expertise change on job listings, this may be explained by the fact that rising expertise requirements may displace incumbent workers who fall below the new threshold, forcing firms to recruit replacements even as the qualified labor pool shrinks. For AI exposure itself, we find a significant negative effect on job postings, providing external validity for prior single-country studies documenting reduced hiring in exposed occupations (Brynjolfsson et al., 2025; Lichtinger and Hosseini Maasoum, 2025; Klein Teeselink, 2025; Azar et al., 2025). Cross-country variation in these effects correlates with institutional context, as countries with stricter employment protection and lower digital readiness experience larger hiring reductions. We additionally find evidence of an interaction between exposure and expertise change: occupations where AI automates low-expertise tasks experience displacement effects roughly one-third smaller than those where AI targets high-expertise tasks.

We acknowledge several limitations. First, our exposure and expertise measures capture potential automation based on task characteristics instead of actual automation, which likely varies across firms and countries depending on complementary investments in organizational change and infrastructure. Second, the 2.5-year post-ChatGPT window may be insufficient to observe the full labor market adjustment, particularly the supply-side constraints that arise when expertise requirements increase. Third, our expertise change measure is a predicted measure based on the assumption that all exposed tasks are automated. In reality, only a fraction of those tasks have been automated so far. Hence, our results likely underestimate the overall effect once the automation has fully played out. Fourth, we do not observe how workers reallocate effort across tasks in response to AI, and the framework's assumption that automation eliminates tasks entirely may oversimplify more complex patterns of partial automation. Finally, it could be the case that AI affects worker-level expertise, which are channel the framework does not capture.

These limitations notwithstanding, our findings carry important policy implications. In particular, we would like to stress the notion that effective intervention cannot be one-size-fits-all. Workers in occupations where AI may automate high-expertise tasks face most acute wage pressures, as the specialized knowledge that

previously justified premium wages becomes automatable. Training programs could target these workers with curricula focused on skills that remain difficult for AI to replicate, particularly those involving judgment, creativity, and interpersonal interaction. Conversely, in occupations where AI automates low-expertise tasks, workers can concentrate on higher-value expert work, and policies that encourage AI adoption in these augmenting contexts may yield productivity gains without displacing labor. Yet even within high-displacement occupations, the urgency of intervention varies. Manning and Aguirre (2025) show that AI exposure and workers' adaptive capacity, as measured by their financial buffers, transferable skills, and retraining possibilities, are positively correlated: many highly exposed occupations contain workers relatively well-equipped to navigate job transitions. At the same time, they identify 6.1 million workers in occupations that combine high exposure with low adaptive capacity, concentrated in clerical and administrative roles. Combining their insights with ours suggests that the most vulnerable workers are those facing high AI exposure that displaces high expertise tasks *and* who have limited means to adjust. Optimal policy therefore requires layering multiple dimensions of risk: exposure levels, expertise changes, and whether affected workers possess the resources to transition successfully.

References

Acemoglu, D. (2024), The simple macroeconomics of AI, Working Paper 32487, National Bureau of Economic Research.

Acemoglu, D. and Autor, D. (2011), Skills, tasks and technologies: Implications for employment and earnings, *in* 'Handbook of Labor Economics', Vol. 4, Elsevier, pp. 1043–1171.

Acemoglu, D. and Restrepo, P. (2018), 'The race between man and machine: Implications of technology for growth, factor shares, and employment', *American Economic Review* **108**(6), 1488–1542.

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. (2022), 'Tasks, automation, and the rise in US wage inequality', *Econometrica* **90**(5), 1973–2016.

Akerman, A., Gaarder, I. and Mogstad, M. (2015), 'The skill complementarity of broadband internet', *Quarterly Journal of Economics* **130**(4), 1781–1824.

Appel, R., McCrory, P., Tamkin, A., McCain, M., Neylon, T. and Stern, M. (2025), 'Anthropic economic index report: Uneven geographic and enterprise AI adoption', *arXiv preprint arXiv:2511.15080* .

Autor, D. H., Katz, L. F. and Kearney, M. S. (2006), 'The polarization of the US labor market', *American Economic Review* **96**(2), 189–194.

Autor, D. H., Levy, F. and Murnane, R. J. (2003), 'The skill content of recent technological change: An empirical exploration', *The Quarterly Journal of Economics* **118**(4), 1279–1333.

Autor, D. and Salomons, A. (2018), 'Is automation labor-displacing? productivity growth, employment, and the labor share', *Brookings Papers on Economic Activity* **2018**(1), 1–87.

Autor, D. and Thompson, N. (2025), 'Expertise', *Journal of the European Economic Association* **23**(4), 1203–1271.

Azar, J., Gine, M. and Sanz-Espín, J. (2025), 'AI is already eroding wages: Quasi-experimental evidence from occupational exposure', SSRN Scholarly Paper.

Babina, T., Fedyk, A., He, A. and Hodson, J. (2024), 'Artificial Intelligence, firm growth, and product innovation', *Journal of Financial Economics* **151**, 103745.

Bartelsman, E. J., Gautier, P. A. and De Wind, J. (2016), 'Employment protection, technology choice, and worker allocation', *International Economic Review* **57**(3), 787–826.

Brynjolfsson, E., Chandar, B. and Chen, R. (2025), 'Canaries in the coal mine? Six facts about the recent employment effects of Artificial Intelligence'. Working paper.

Eloundou, T., Manning, S., Mishkin, P. and Rock, D. (2024), 'GPTs are GPTs: An early look at the labor market impact potential of large language models', *Science* **384**(6702), 1306–1308.

Felten, E. W., Raj, M. and Seamans, R. (2023), 'How will language modelers like ChatGPT affect occupations and industries?', arXiv preprint arXiv:2303.01157.

Frey, C. B. and Osborne, M. A. (2017), 'The future of employment: How susceptible are jobs to computerisation?', *Technological Forecasting and Social Change* **114**, 254–280.

Gmyrek, P., Berg, J. and Bescond, D. (2023), Generative AI and jobs: A global analysis of potential effects on job quantity and quality, ILO Working Paper 96, International Labour Organization, Geneva.

Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rosłaniec, K. and Troszyński, M. (2025), Generative AI and jobs: A refined global index of occupational exposure, Working Paper 140, International Labour Organization (ILO), Geneva.

Goldin, C. and Katz, L. F. (2008), *The Race Between Education and Technology*, Harvard University Press, Cambridge, MA.

Goos, M. and Manning, A. (2007), 'Lousy and lovely jobs: The rising polarization of work in Britain', *The Review of Economics and Statistics* **89**(1), 118–133.

Goos, M., Manning, A. and Salomons, A. (2014), 'Explaining job polarization: Routine-biased technological change and offshoring', *American Economic Review* **104**(8), 2509–2526.

Hosseini Maasoum, S. M. and Lichtinger, G. (2026), 'Generative AI, expertise, and effective labor supply', *Available at SSRN 6059674* .

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), Large language models, small labor market effects, Working Paper 32120, National Bureau of Economic Research.

International Institute for Management Development (2024), 'Digital competitiveness ranking', Statbase. Accessed: 2026-01-23.

Iscenko, Z. and Curto Millet, F. (2026), The American worker project: Working paper, Technical report, Economic Innovation Group.

Klein Teeselink, B. (2025), Generative AI and labor market outcomes: Evidence from the United Kingdom, Working paper, SSRN.

Lichtinger, G. and Hosseini Maasoum, S. M. (2025), 'Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data'. Working paper, SSRN.

Manning, S. and Aguirre, T. (2025), How adaptable are American workers to AI-induced job displacement?, Working paper, National Bureau of Economic Research.

Meleshenko, K. (2023), 'AI global index', Kaggle. Data calculated by Tortoise Media. Accessed: 2026-01-23.

Michaels, G., Natraj, A. and Van Reenen, J. (2014), 'Has ICT polarized skill demand? evidence from eleven countries over twenty-five years', *Review of Economics and Statistics* **96**(1), 60–77.

Organisation for Economic Co-operation and Development (2026), 'Strictness of employment protection'. OECD Data Explorer. Accessed: 2026-01-23.

Prytkova, E., Petit, F., Li, D., Chaturvedi, S. and Ciarli, T. (2024), The employment impact of emerging digital technologies, Working Paper 32213, National Bureau of Economic Research.

Restrepo, P. (2024), 'Automation: Theory, evidence, and outlook', *Annual Review of Economics* **16**(1), 1–25.

Spitz-Oener, A. (2006), 'Technical change, job tasks, and rising educational demands: Looking outside the wage structure', *Journal of Labor Economics* **24**(2), 235–270.

The World Bank (2024), 'Unemployment, total (% of total labor force) (modeled ILO estimate)', World Development Indicators. Accessed: 2026-01-23.

Viechtbauer, W. (2005), 'Bias and efficiency of meta-analytic variance estimators in the random-effects model', *Journal of Educational and Behavioral Statistics* **30**(3), 261–293.

Appendix

A1 Expertise Prompt

To measure task-level expertise requirements, we prompt GPT-4o with each O*NET task description and its corresponding occupation title. The model returns an expertise score on a 0–100 scale along with a brief justification. The full prompt is as follows:

> You are scoring task expertise.
>
> DEFINITION: Expertise is defined as a specific capability or know-how required to perform a task. It is hierarchical: higher-expertise workers can perform lower-level tasks, but lower-expertise workers cannot perform tasks above their level. Expertise acts as a strict barrier to entry for the occupation.
>
> SCORING CRITERIA (0-100 scale):
>
> LOW EXPERTISE (0-20): Generic Tasks
>
> - ‘‘Generic’’ tasks that effectively all adults can perform
> - Requires NO formal training, certification, or specialized capability
> - Examples: ‘‘Box lightweight items,’’ ‘‘Greet customers,’’ ‘‘Move materials’’
>
> MEDIUM EXPERTISE (21-60): Inexpert / Supporting Tasks
>
> - Tasks bundled into an occupation that require some familiarity but are not the primary barrier to entry
> - Tasks that are ‘‘inexpert’’ relative to the occupation’s core requirements
> - Examples: ‘‘File reimbursement claims’’ (for a manager), ‘‘Count inventory’’ (for a clerk)
>
> HIGH EXPERTISE (61-100): Expert Tasks

- The specialized tasks that define the occupation's barrier to entry
- Requires specific training, certification, or scarce know-how
- If a worker cannot do this task, they cannot enter the occupation
- Examples: ``Reconcile bank statements'' (for an accountant), ``Diagnose medical conditions,'' ``Write legal briefs''

IMPORTANT CONTEXT:

- HIERARCHY: Evaluate if the task creates a ``barrier to entry.''
- BUNDLING: A task might be ``inexpert'' for one job (e.g., a manager filing receipts) but ``expert'' for another.
- GENERIC: If a task can be done by any worker without training, it is Generic (Score < 20).

INPUTS:
Occupation: [OCCUPATION TITLE]
Task: [TASK DESCRIPTION]

OUTPUT:
Return only a JSON object:
{ "expertise_score": [0-100], "barrier_to_entry": [brief description of the specific know-how required or "None" if generic] }

We apply this prompt to all 19,265 task descriptions in O*NET version 27.2. We then aggregate task-level scores to the occupation level by computing the importance-weighted mean across all tasks within each occupation, where importance weights come from O*NET's task importance ratings.

A2 Validation of the Expertise Measure

This appendix details the validation exercises summarized in Section 2.2. We conduct two sets of tests: correlations with labor market outcomes and robustness to prompt specification.

For the first test, we match O*NET occupations to Revelio Labs data aggregated from LinkedIn profiles, constructing three occupation-level outcomes: average years of education, average salary, and average total compensation. For each of the 38 countries in our sample, we estimate a cross-sectional OLS regression of the standardized outcome on standardized baseline expertise at the occupation level. We then aggregate country-specific coefficients using random-effects meta-analysis. Figures A1 to A3 present the results.

Next, we assess sensitivity to prompt wording. To do so, we first re-run the expertise classification exercise using an alternative prompt with different score thresholds and framing (see Appendix A2.1), then aggregate to occupation-level scores using the same task-level importance-weights as in our main analysis. Second, we compare our rankings to the independently constructed measure of Hosseini Maasoum and Lichtinger (2026), who use a substantively different prompt and data construction strategy. Because alternative measures need not share the same scale, we focus on rank-order comparisons. Figures A4 and A5 display the rank-rank plots.

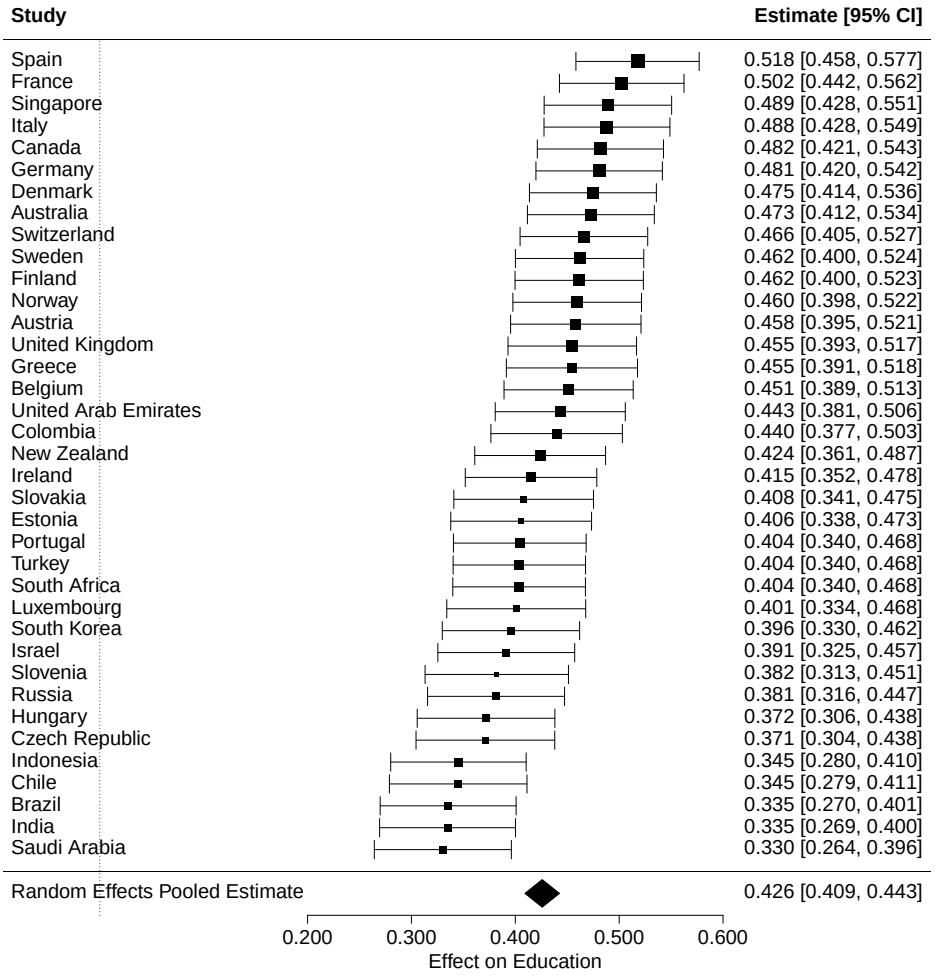


Figure A1: Validation of baseline expertise using education. Each row reports a country-specific coefficient from an OLS regression of standardized occupation-level education on standardized baseline expertise. Horizontal lines denote 95% confidence intervals. The diamond reports the pooled random-effects estimate.

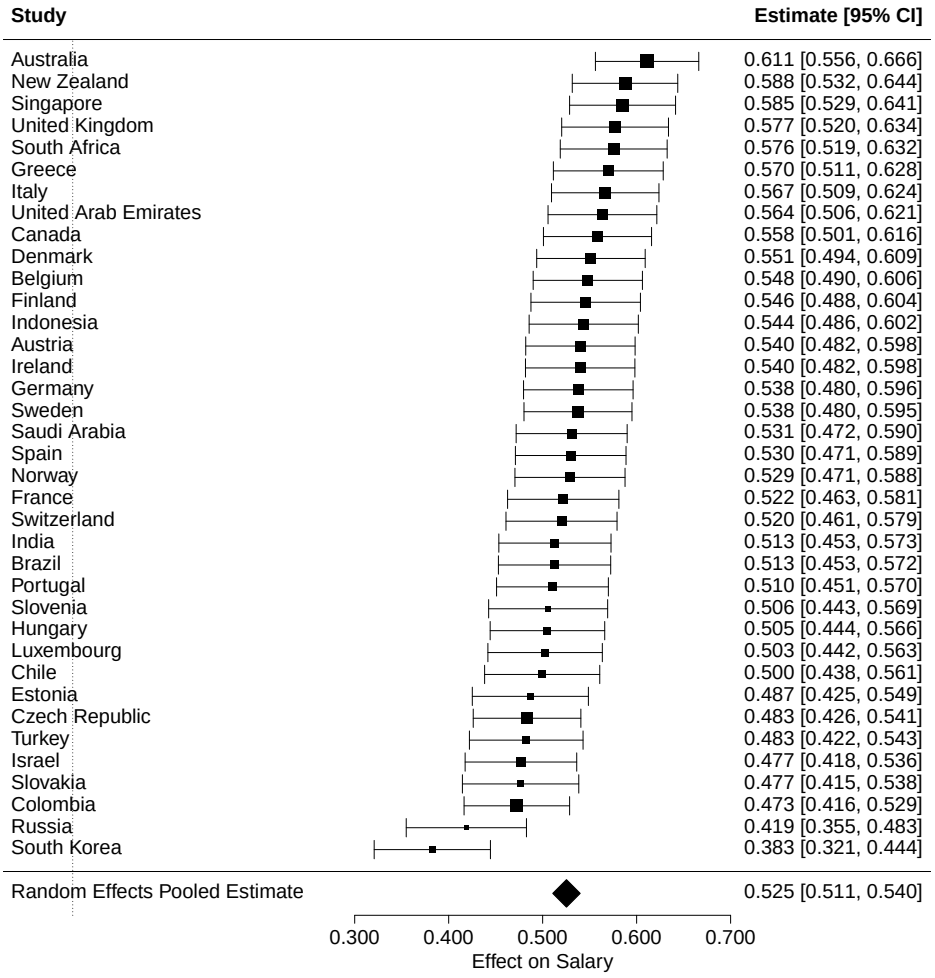


Figure A2: Validation of baseline expertise using salary. Each row reports a country-specific coefficient from an OLS regression of standardized occupation-level salary on standardized baseline expertise. Horizontal lines denote 95% confidence intervals. The diamond reports the pooled random-effects estimate.

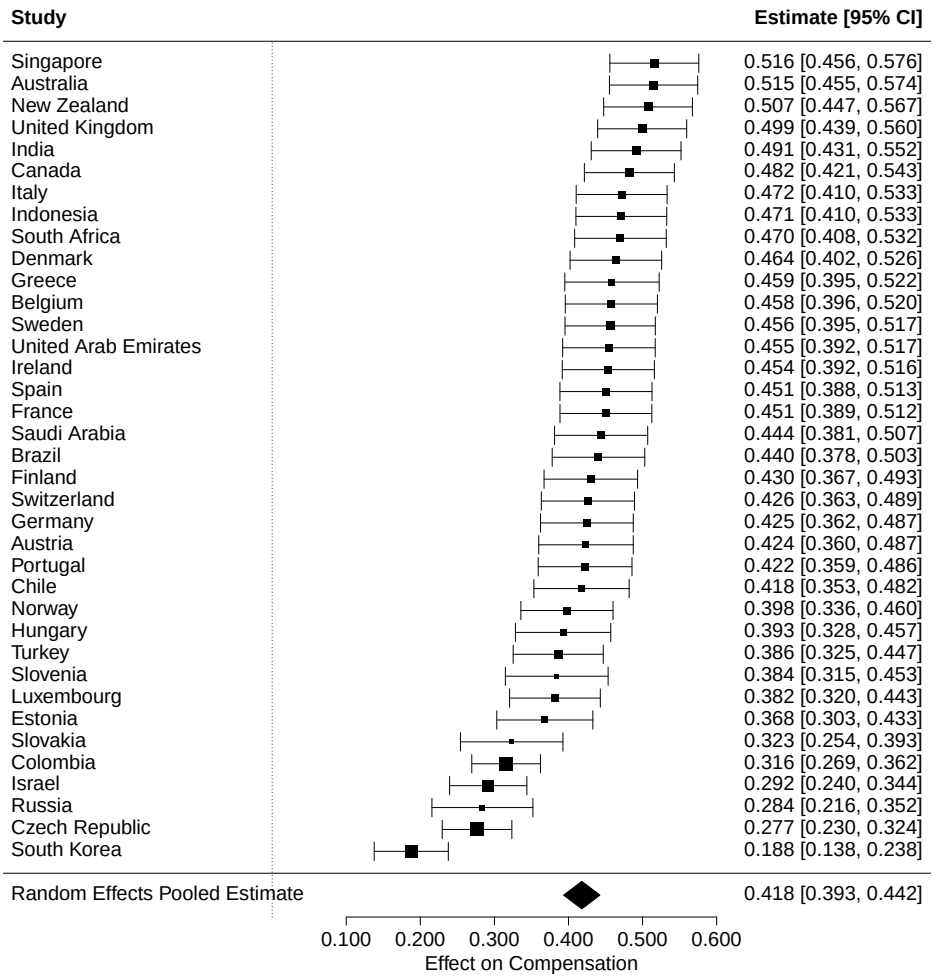


Figure A3: Validation of baseline expertise using total compensation. Each row reports a country-specific coefficient from an OLS regression of standardized occupation-level compensation on standardized baseline expertise. Horizontal lines denote 95% confidence intervals. The diamond reports the pooled random-effects estimate.

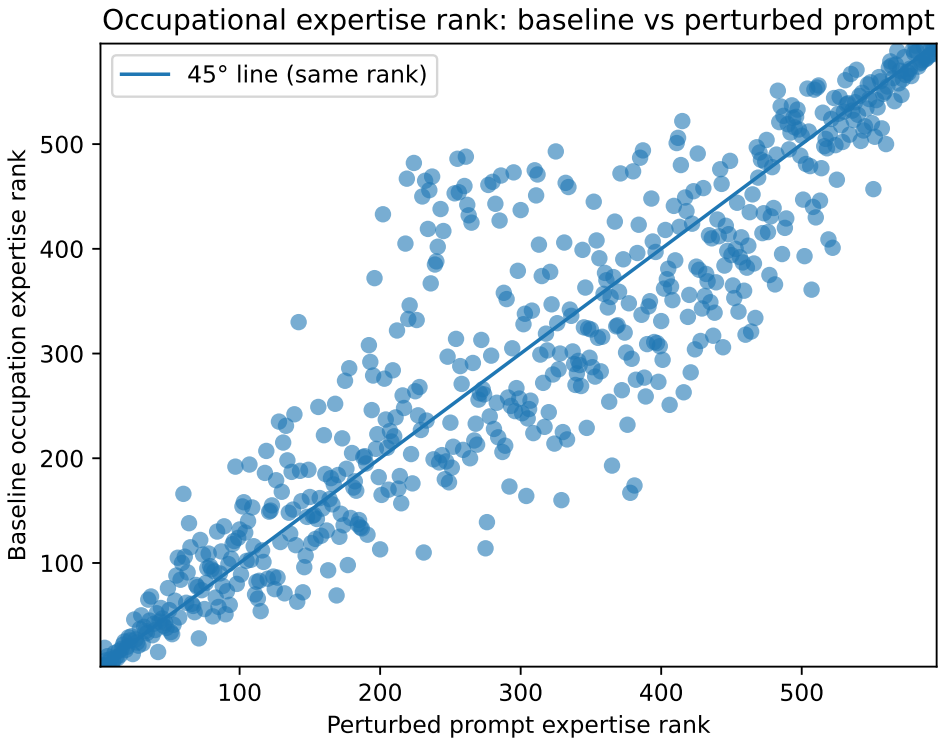


Figure A4: Correlation between rankings of occupational expertise between our baseline measure and an alternative version. The 45 degree line shows where all points would lie if the measures gave identical rankings.

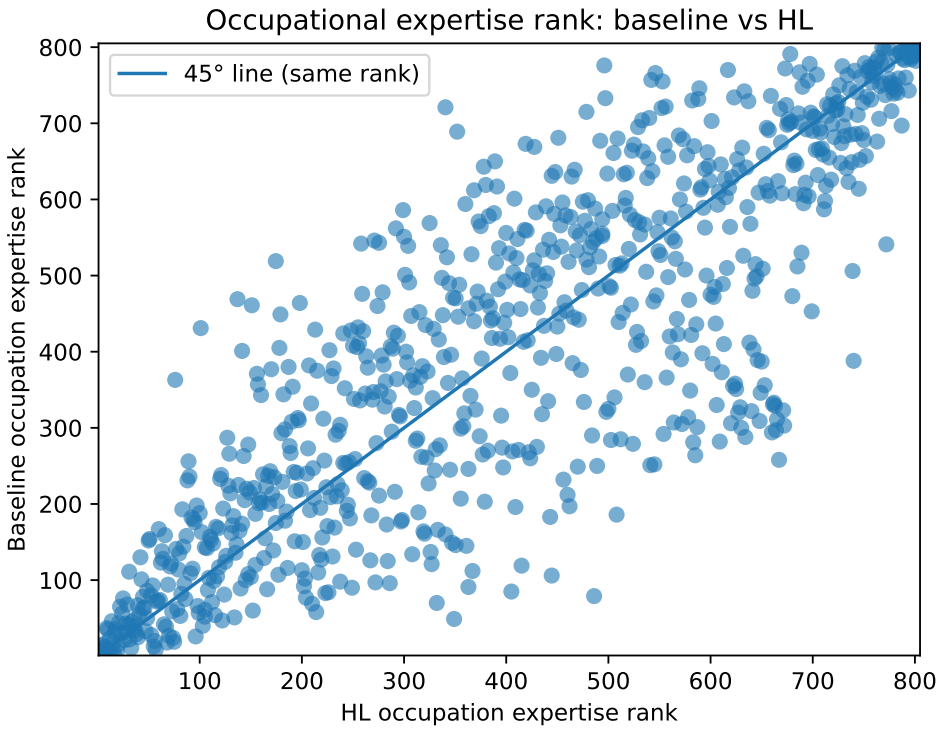


Figure A5: Correlation between rankings of occupational expertise between our baseline measure and Hosseini Maasoum and Lichtinger (2026). The 45 degree line shows where all points would lie if the measures gave identical rankings.

A2.1 Alternative Prompt

You are scoring task expertise for labor market analysis.

DEFINITION: Expertise is the level of specialized training, knowledge, or certification required to perform a task. High-expertise tasks serve as barriers to entry - workers lacking the specific capability cannot perform them.

SCORING CRITERIA (0-100 scale):

LOW EXPERTISE (0-30): Generic tasks

- Requires no formal training or certification
- Most adults could perform with minimal instruction
- Examples: ‘‘File documents alphabetically,’’ ‘‘Answer phone calls,’’ ‘‘Count inventory’’

MEDIUM EXPERTISE (31-70): Specialized tasks

- Requires occupation-specific training or experience
- Needs domain knowledge but not advanced credentials
- Examples: ‘‘Reconcile bank statements,’’ ‘‘Operate forklift,’’ ‘‘Schedule appointments using software’’

HIGH EXPERTISE (71-100): Expert tasks

- Requires formal training, certification, or advanced education
- Demands specialized knowledge that creates clear barriers to entry
- Examples: ‘‘Diagnose medical conditions,’’ ‘‘Design load-bearing structures,’’ ‘‘Argue legal motions’’

IMPORTANT CONTEXT:

- The SAME task may have different expertise levels in different occupations

- ‘‘Execute routine procedures’’ is expert for surgeons, generic for clerks
- Consider what training/knowledge would be needed to perform this task competently

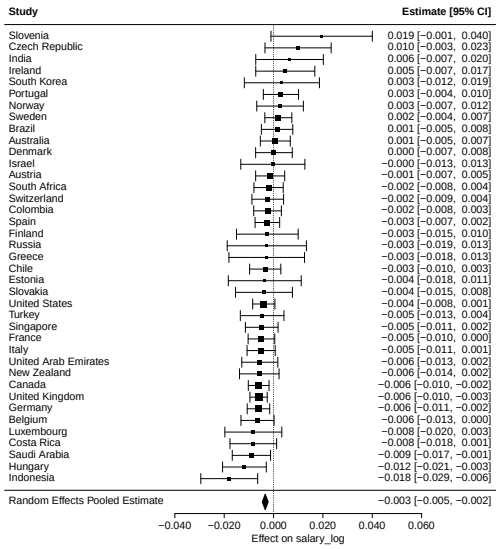
```
INPUTS:
Occupation: [OCCUPATION TITLE]
Task: [TASK DESCRIPTION]

OUTPUT:
Return only a JSON object:
{ "expertise_score": [0-100], "barrier_to_entry": [brief description
of what training/knowledge is required] }
```

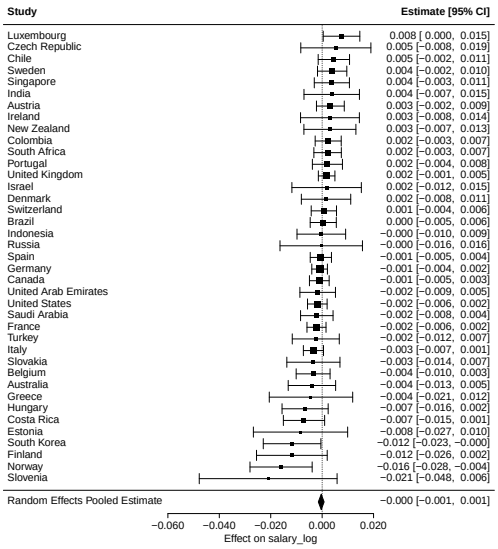
We apply this alternative prompt to all 19,265 O*NET tasks and aggregate to occupation-level scores using the same importance-weighting procedure as for our baseline measure. The resulting occupation rankings correlate 0.914 (Spearman ρ) with our baseline measure, as shown in Figure A4.

Figure A6: Robustness: Alpha Exposure Measure

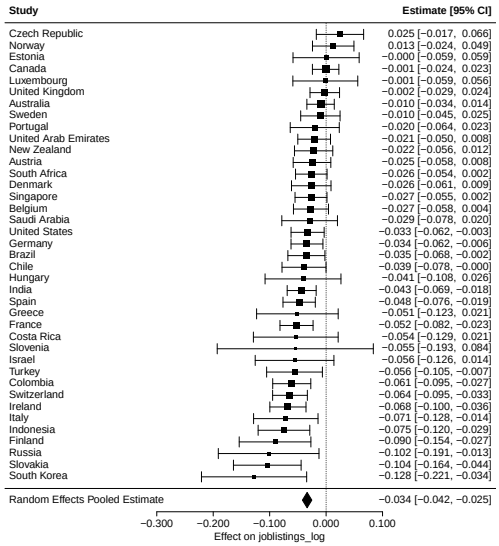
(a) Salary—Exposure



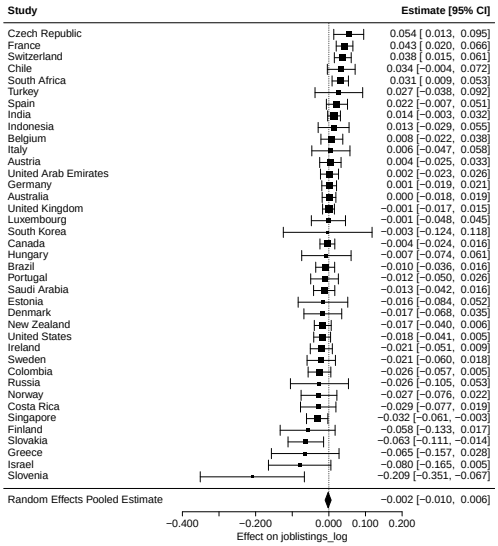
(b) Salary—Expertise



(c) Job Postings—Exposure



(d) Job Postings—Expertise

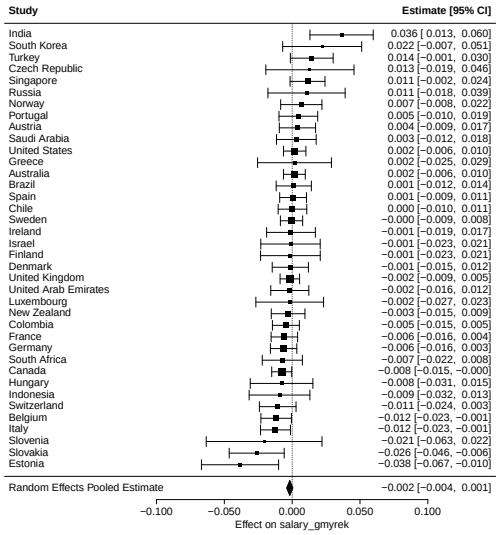


Notes: Results using the α expo
Figure 1.

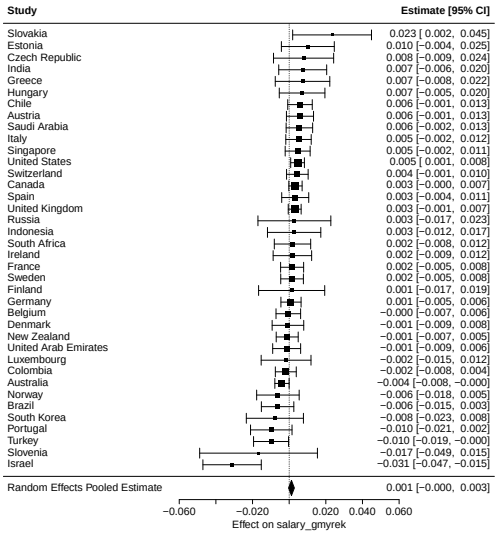
A3 Robustness Checks

Figure A7: Robustness: Gmyrek et al. (2023) Exposure Measures (ISCO Level)

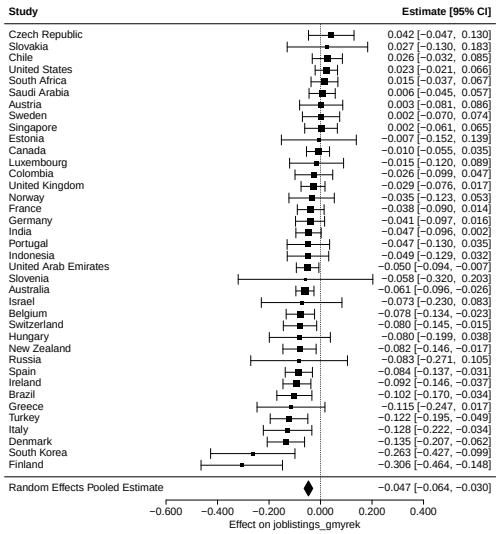
(a) Salary—Exposure



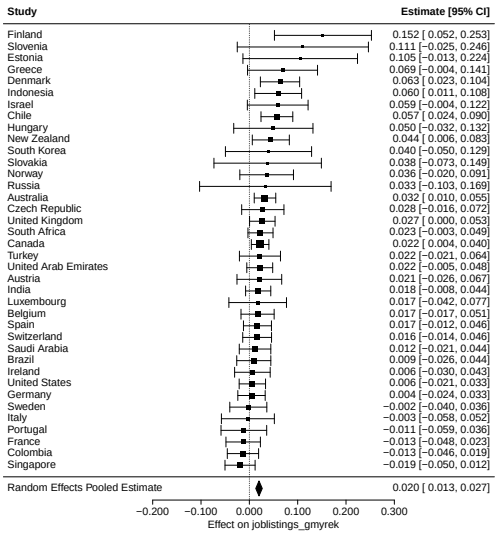
(b) Salary—Expertise



(c) Job Postings—Exposure



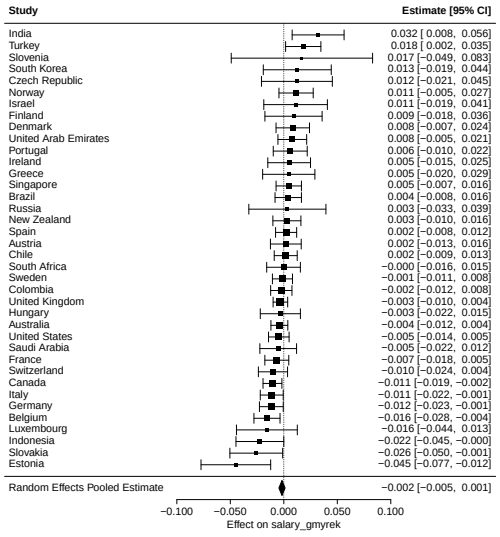
(d) Job Postings—Expertise



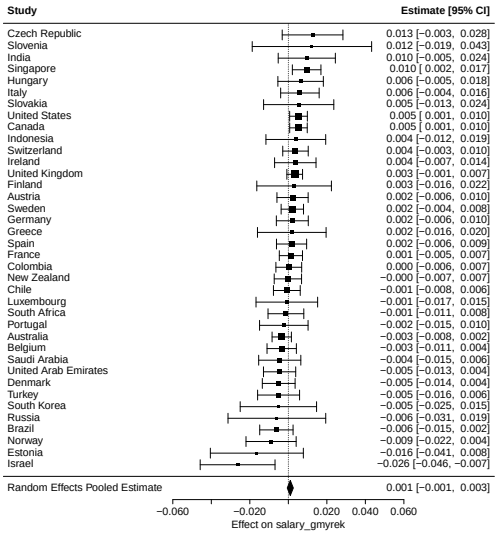
Notes: Results using exposure and expertise measures from Gmyrek et al. (2023) at the ISCO occupation level. Other details are as in Figure 1.

Figure A8: Robustness: Gmyrek et al. (2025) Exposure Measures (ISCO Level)

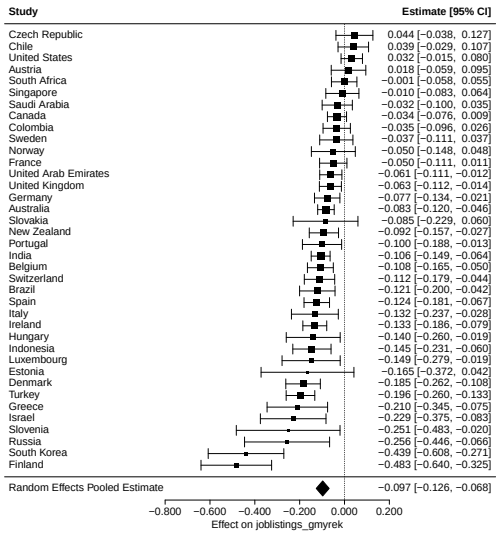
(a) Salary—Exposure



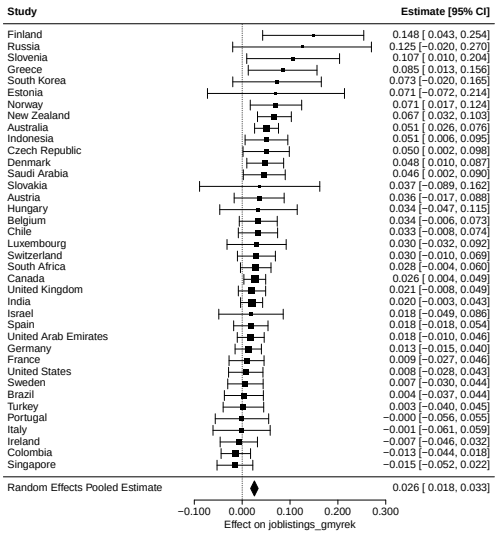
(b) Salary—Expertise



(c) Job Postings—Exposure



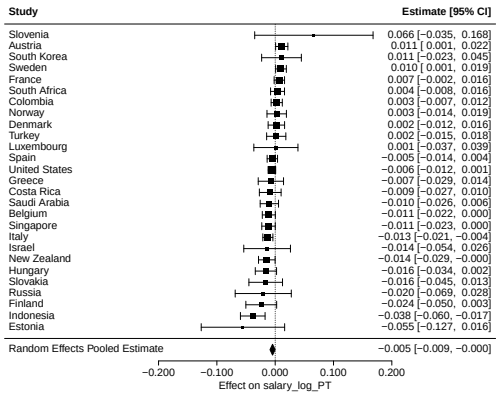
(d) Job Postings—Expertise



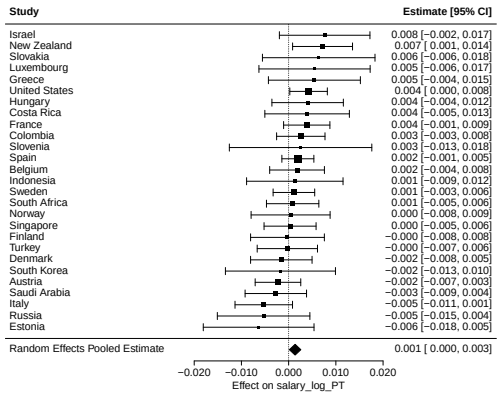
Notes: Results using exposure and expertise measures from Gmyrek et al. (2025) at the ISCO occupation level. Figure 1.

Figure A9: Robustness: Countries Passing Parallel Trends Tests

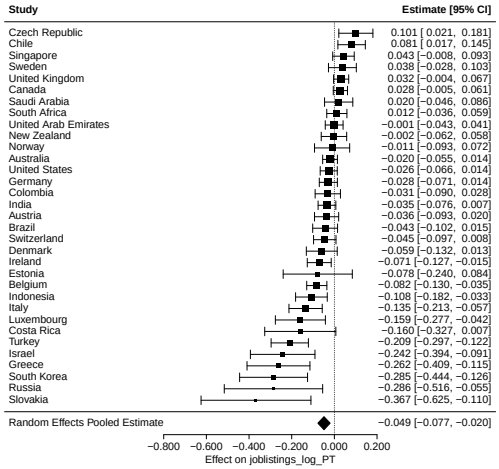
(a) Salary—Exposure



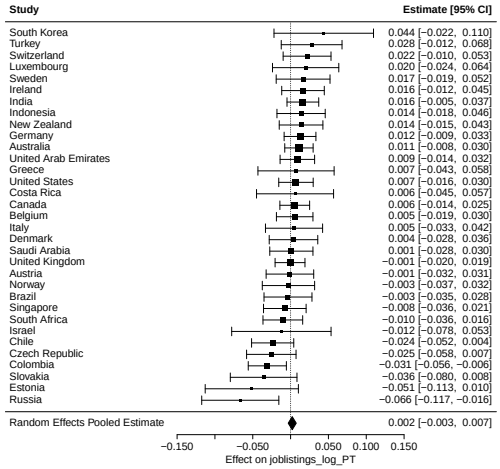
(b) Salary—Expertise



(c) Job Postings—Exposure



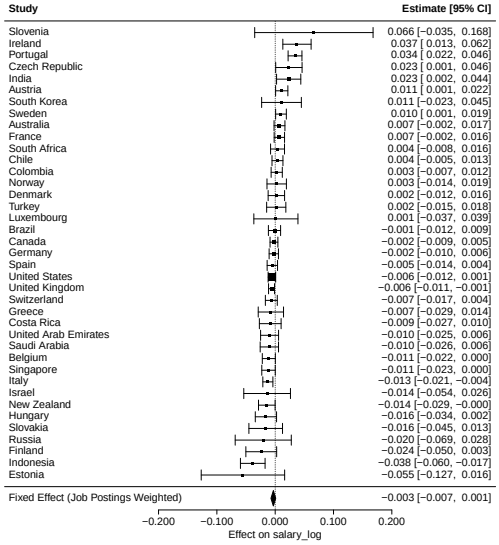
(d) Job Postings—Expertise



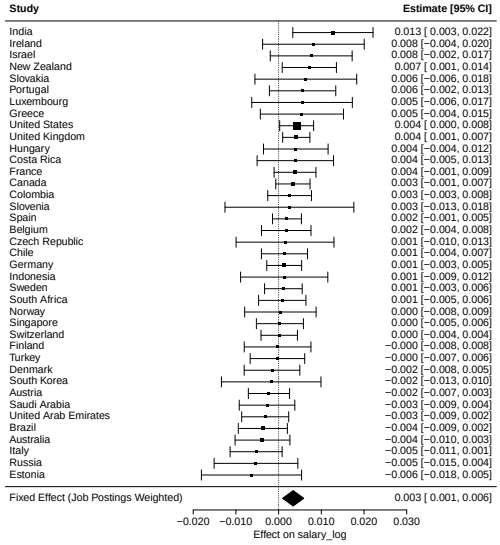
Notes: Sample restricted to countries where pre-treatment event study coefficients are on average insignificant at the 5% level. Other details are as in Figure 1.

Figure A10: Robustness: Job Posting–Weighted Meta-Analysis

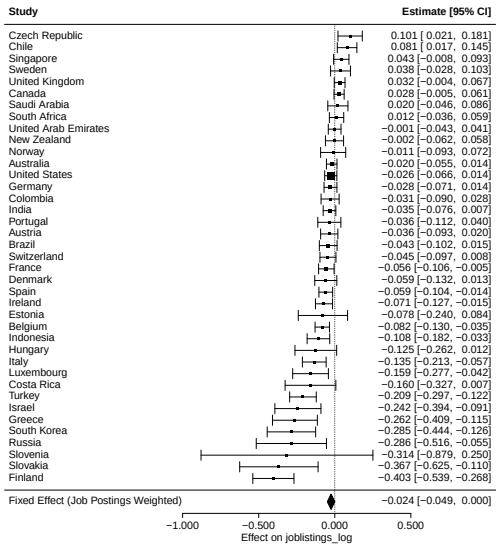
(a) Salary—Exposure



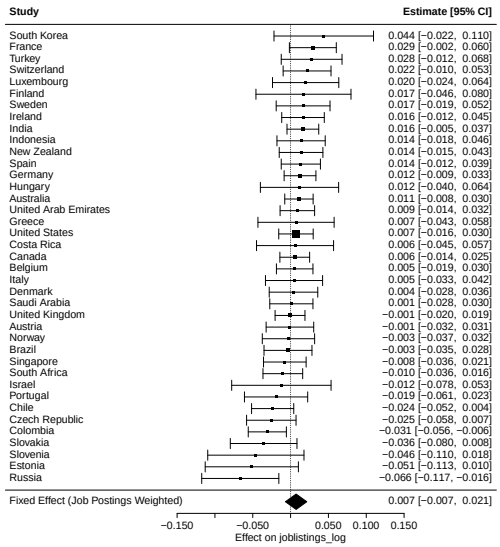
(b) Salary—Expertise



(c) Job Postings—Exposure



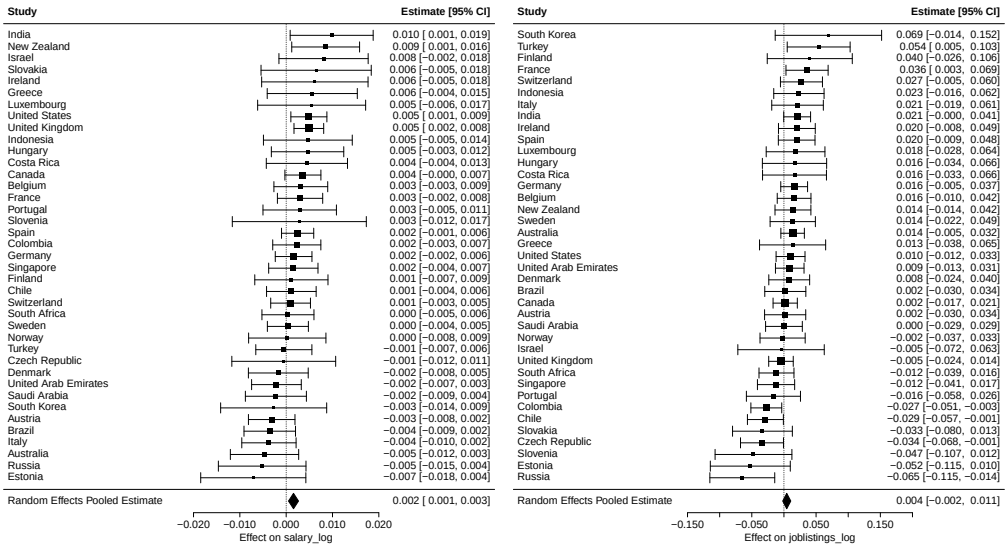
(d) Job Postings—Expertise



Notes: Fixed effects meta-analysis with countries weighted by pre-treatment job posting volume.

Figure A11: Robustness: Expertise Change Only (Without Exposure Control)

(a) Salary—Expertise (b) Job Postings—Expertise



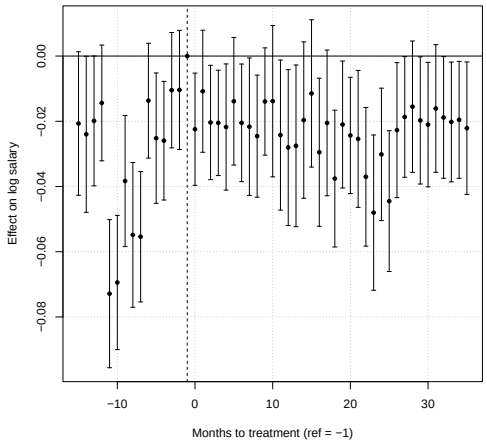
Notes: Results from models that include only expertise change without controlling for AI exposure. The expertise change effect replicates and becomes stronger compared to the main specification.

A4 Event Study Plots by Country

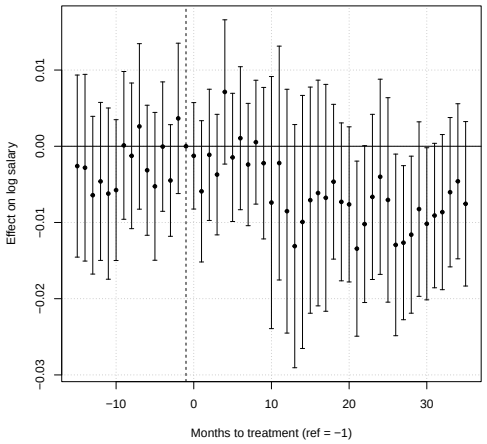
This section presents country-level event study estimates for all 39 countries in our sample. Each figure displays four panels: the effects of AI exposure and expertise change on log average salary (top row) and log job postings (bottom row). Coefficients represent the differential effect of a one standard deviation increase in the respective treatment variable relative to October 2022 (the month before ChatGPT's release). Shaded areas denote 95% confidence intervals.

Figure A12: Event Study: Australia

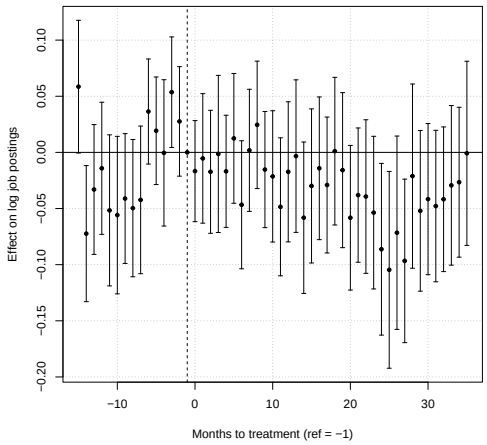
(a) Salary—Exposure



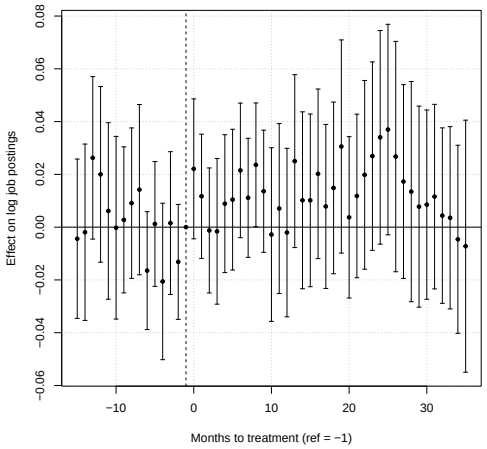
(b) Salary—Expertise Change



(c) Job Postings—Exposure



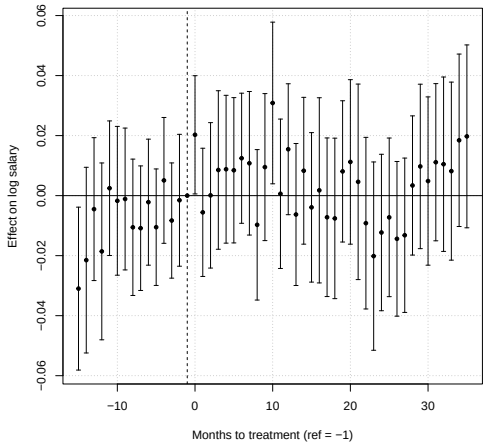
(d) Job Postings—Expertise Change



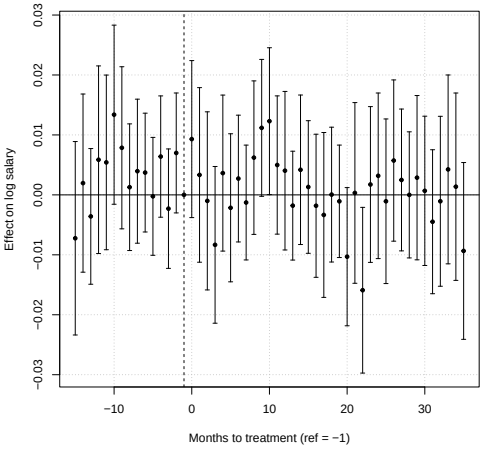
Notes: Event study estimates from the specification in Equation (4). Panels (a) and (c) plot coefficients on the interaction between AI exposure (β measure) and period indicators. Panels (b) and (d) plot coefficients on the interaction between expertise change and period indicators. All regressions include occupation fixed effects and 2-digit O*NET by year-month fixed effects. Standard errors are clustered at the occupation level. The vertical dashed line indicates November 2022, when ChatGPT was publicly released. October 2022 serves as the reference period.

Figure A13: Event Study: Austria

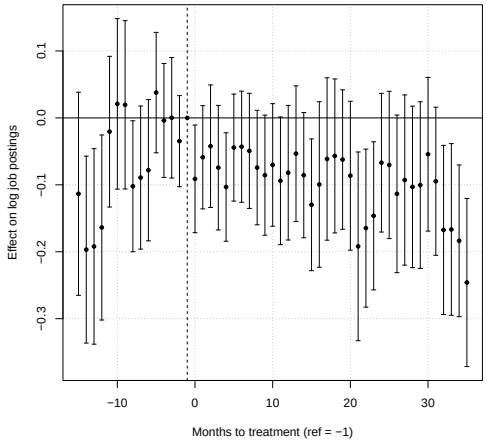
(a) Salary—Exposure



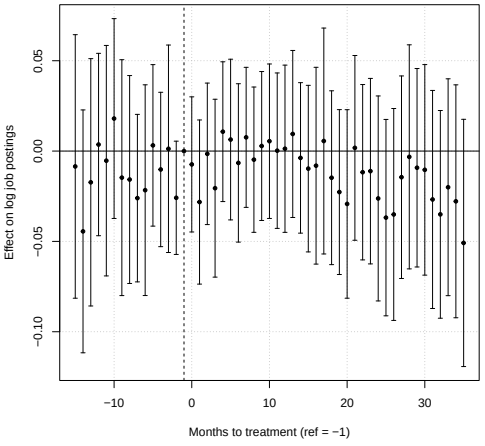
(b) Salary—Expertise Change



(c) Job Postings—Exposure



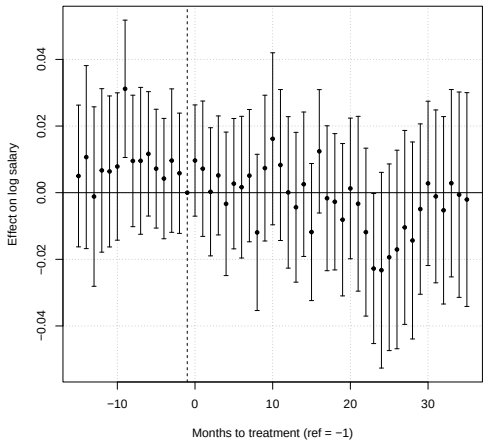
(d) Job Postings—Expertise Change



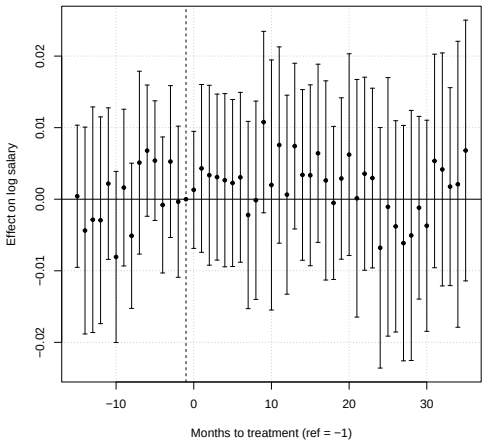
Notes: Other details are as in Figure A12.

Figure A14: Event Study: Belgium

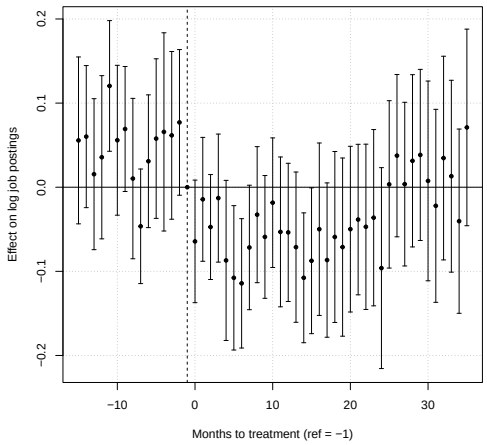
(a) Salary—Exposure



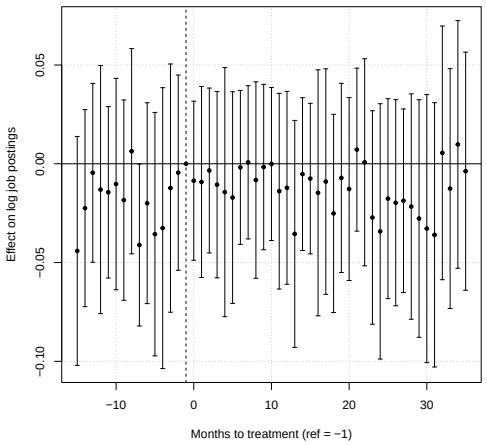
(b) Salary—Expertise Change



(c) Job Postings—Exposure



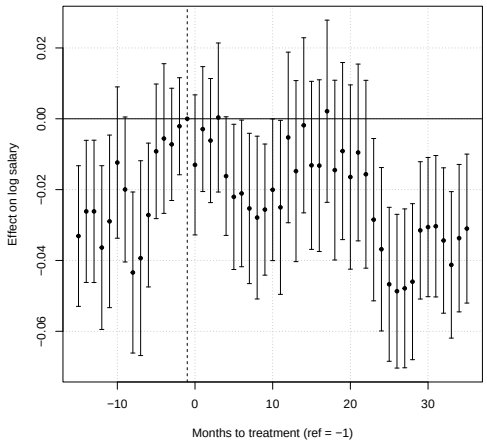
(d) Job Postings—Expertise Change



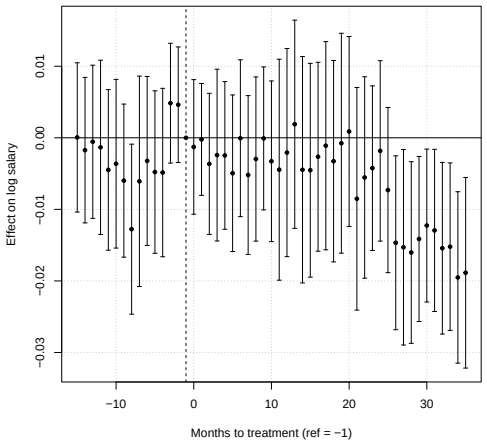
Notes: Other details are as in Figure A12.

Figure A15: Event Study: Brazil

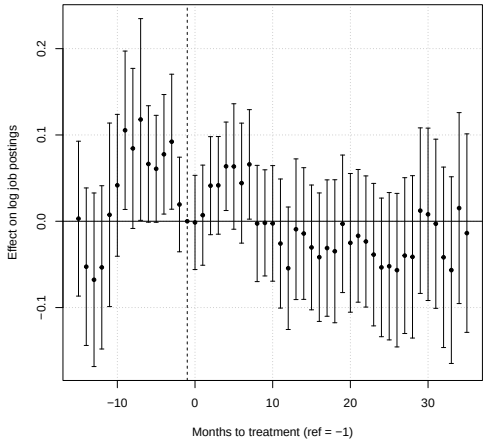
(a) Salary—Exposure



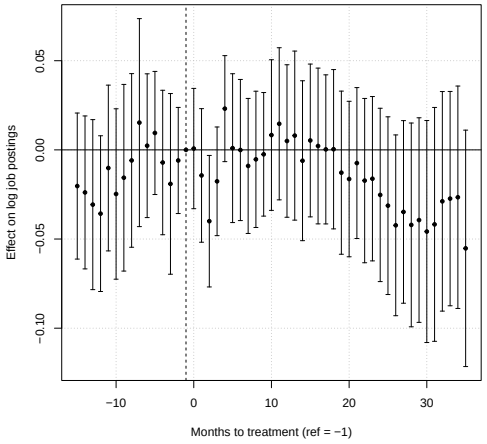
(b) Salary—Expertise Change



(c) Job Postings—Exposure



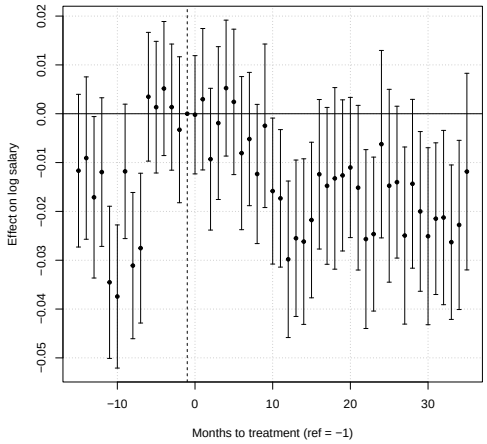
(d) Job Postings—Expertise Change



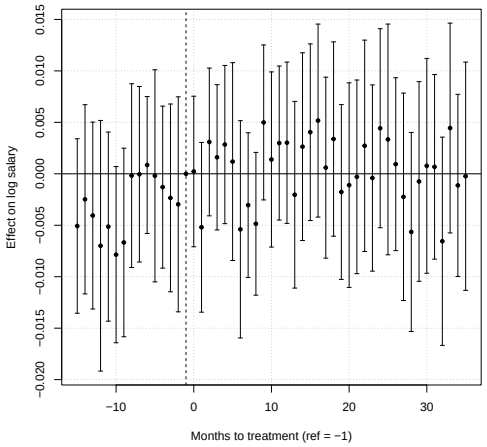
Notes: Other details are as in Figure A12.

Figure A16: Event Study: Canada

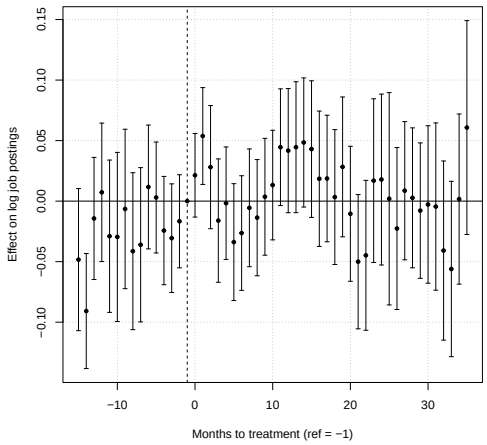
(a) Salary—Exposure



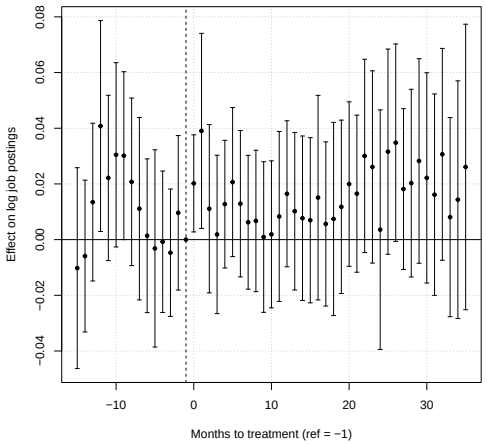
(b) Salary—Expertise Change



(c) Job Postings—Exposure



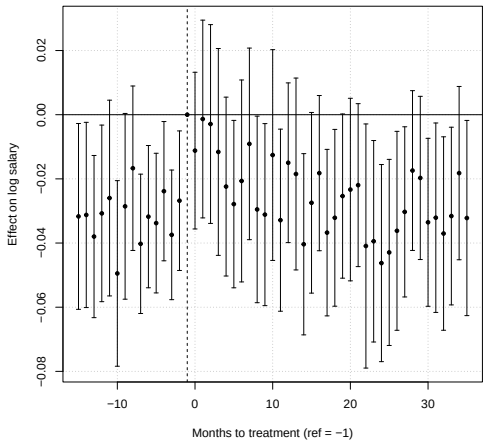
(d) Job Postings—Expertise Change



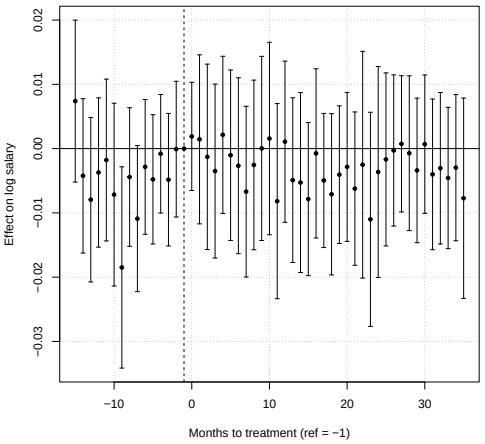
Notes: Other details are as in Figure A12.

Figure A17: Event Study: Chile

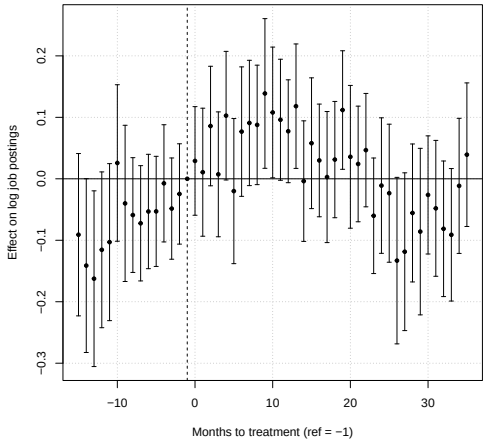
(a) Salary—Exposure



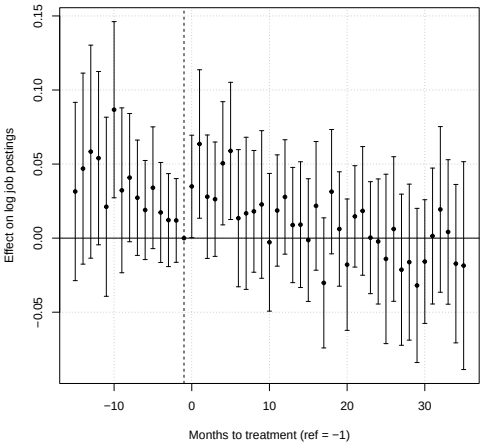
(b) Salary—Expertise Change



(c) Job Postings—Exposure



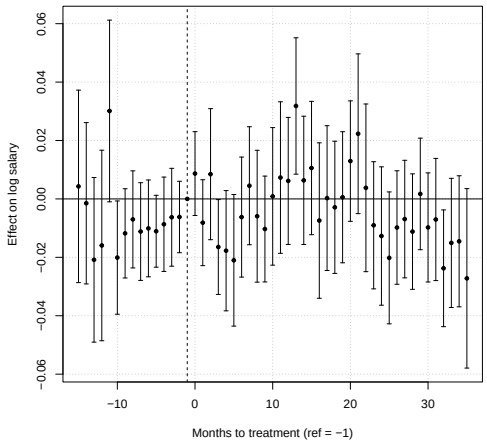
(d) Job Postings—Expertise Change



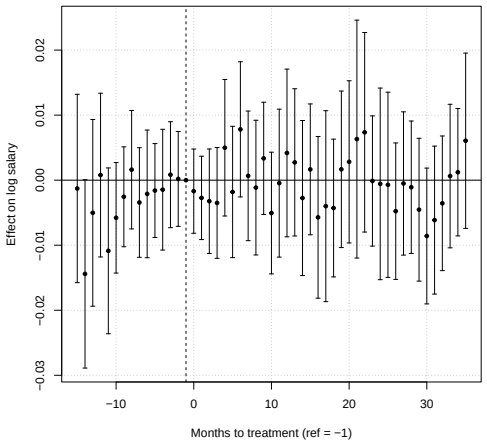
Notes: Other details are as in Figure A12.

Figure A18: Event Study: Colombia

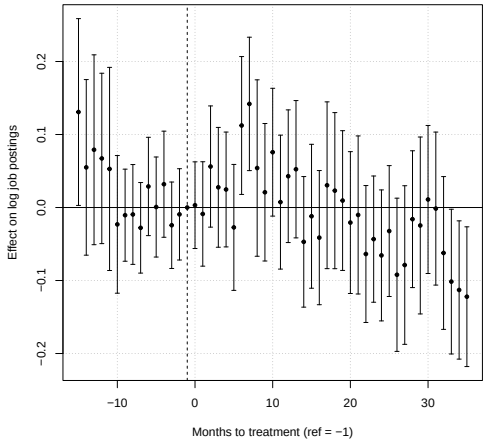
(a) Salary—Exposure



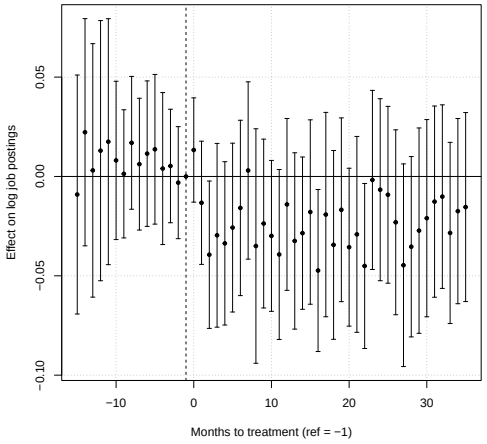
(b) Salary—Expertise Change



(c) Job Postings—Exposure



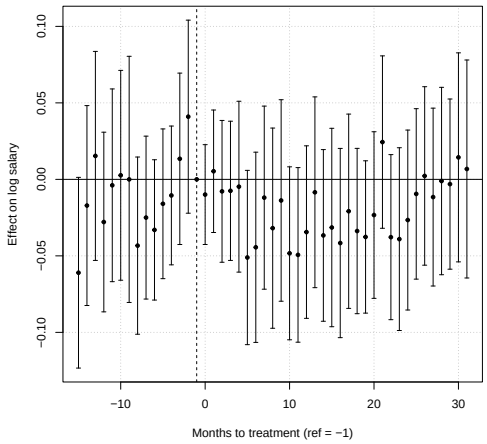
(d) Job Postings—Expertise Change



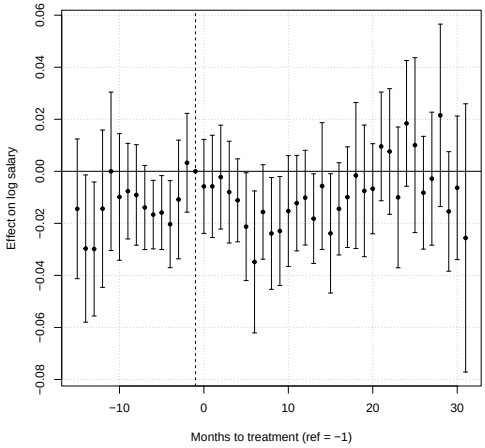
Notes: Other details are as in Figure A12.

Figure A19: Event Study: Costa Rica

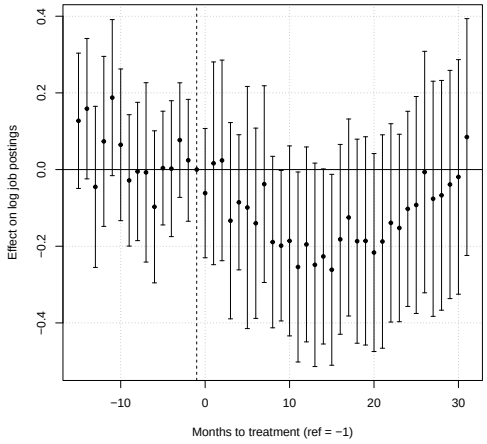
(a) Salary—Exposure



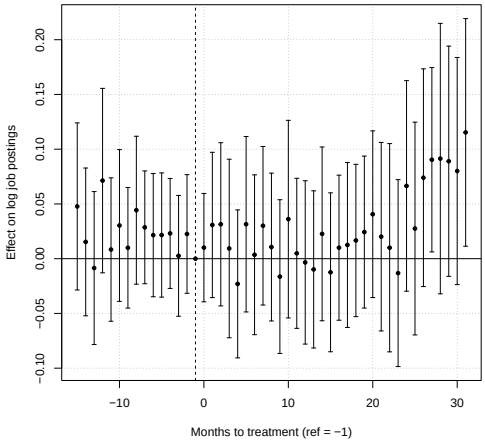
(b) Salary—Expertise Change



(c) Job Postings—Exposure



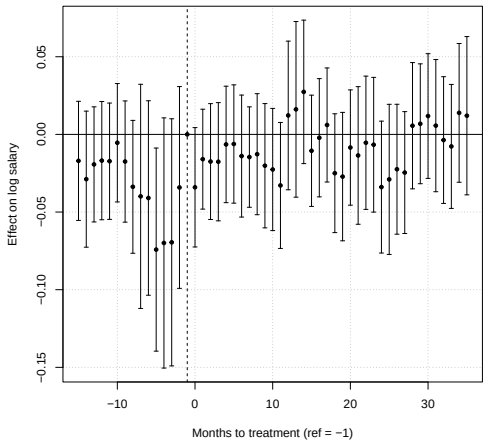
(d) Job Postings—Expertise Change



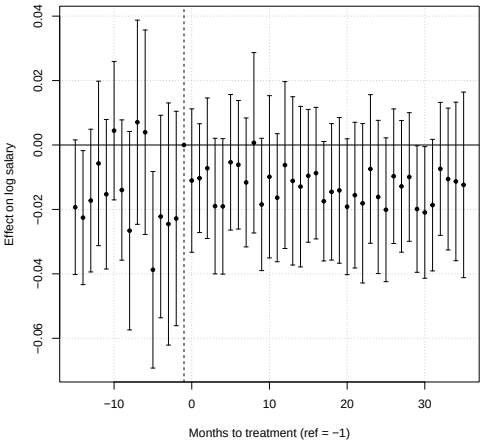
Notes: Other details are as in Figure A12.

Figure A20: Event Study: Czech Republic

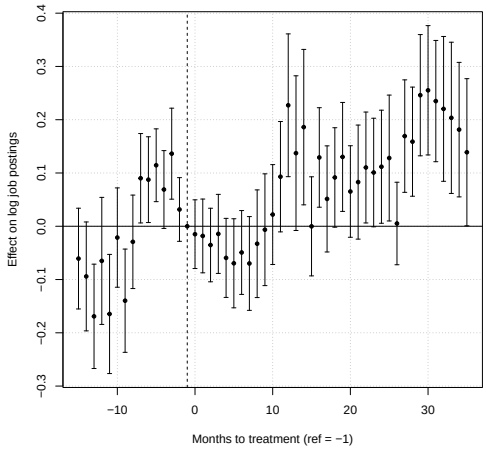
(a) Salary—Exposure



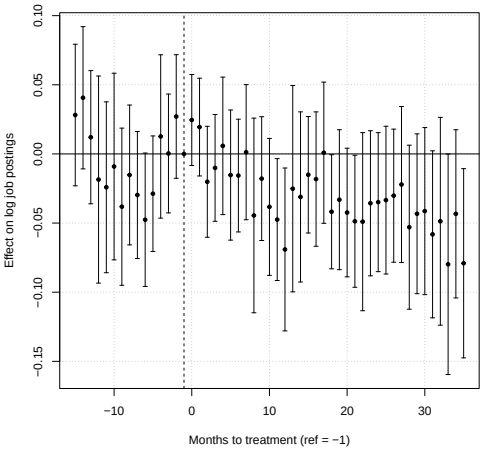
(b) Salary—Expertise Change



(c) Job Postings—Exposure



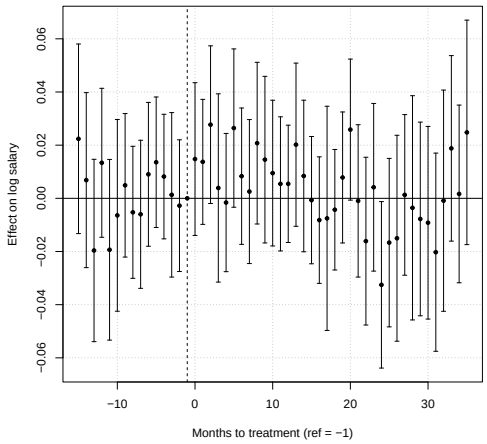
(d) Job Postings—Expertise Change



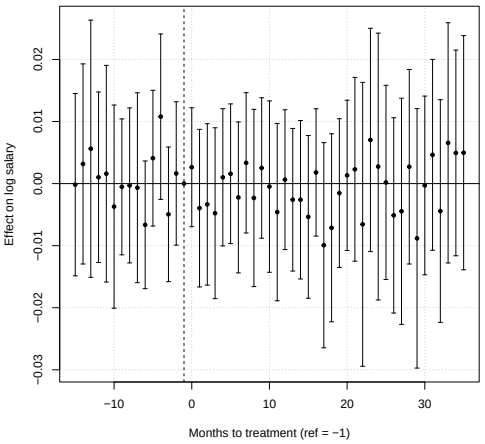
Notes: Other details are as in Figure A12.

Figure A21: Event Study: Denmark

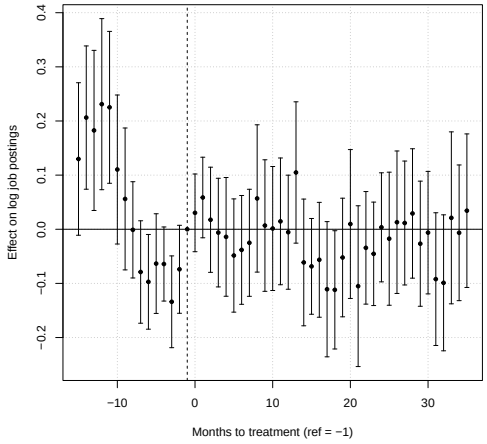
(a) Salary—Exposure



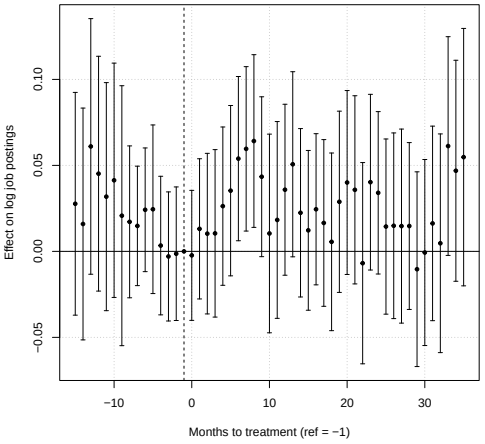
(b) Salary—Expertise Change



(c) Job Postings—Exposure



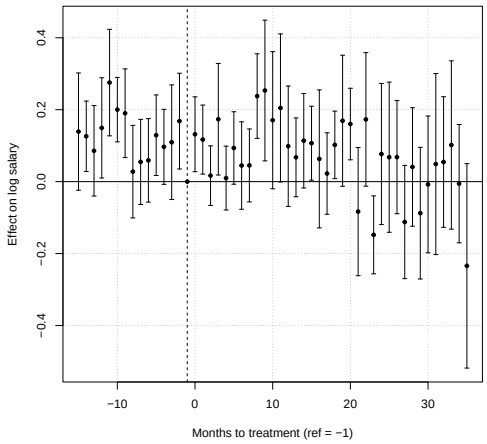
(d) Job Postings—Expertise Change



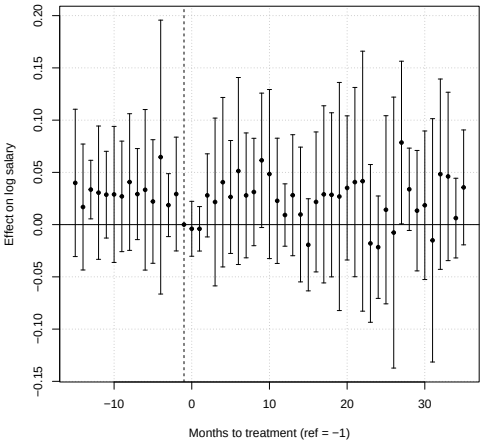
Notes: Other details are as in Figure A12.

Figure A22: Event Study: Estonia

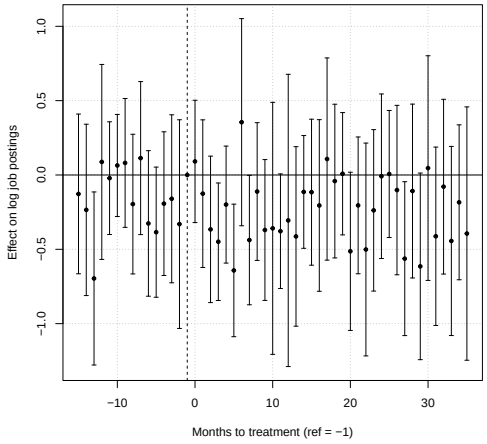
(a) Salary—Exposure



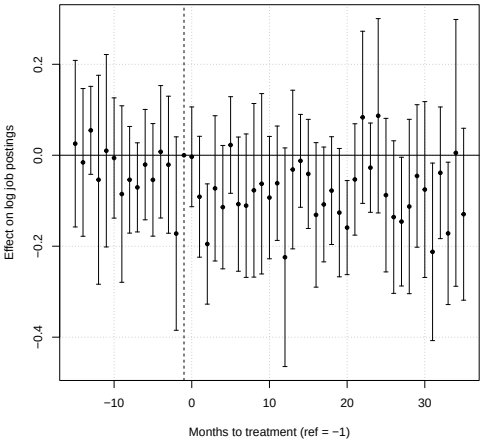
(b) Salary—Expertise Change



(c) Job Postings—Exposure



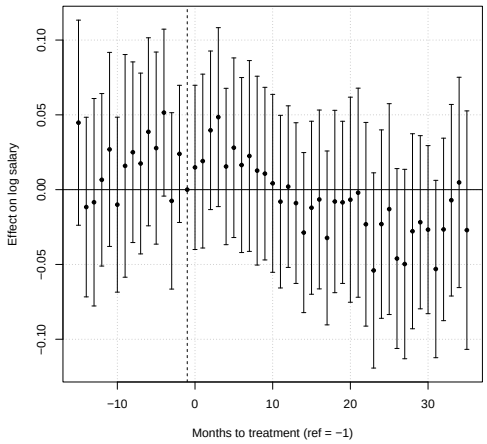
(d) Job Postings—Expertise Change



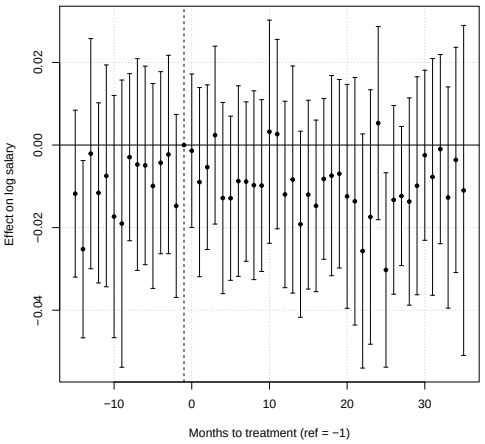
Notes: Other details are as in Figure A12.

Figure A23: Event Study: Finland

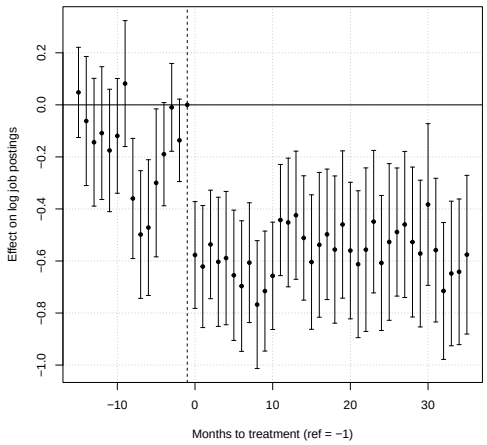
(a) Salary—Exposure



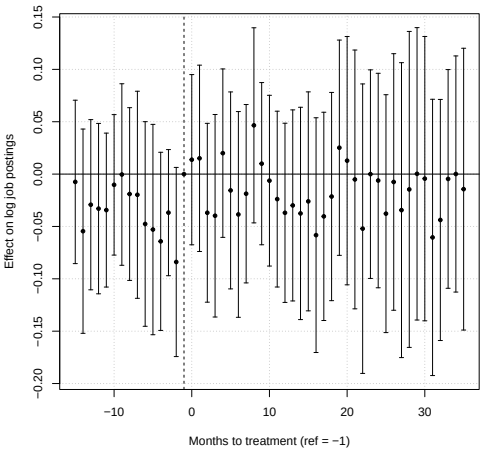
(b) Salary—Expertise Change



(c) Job Postings—Exposure



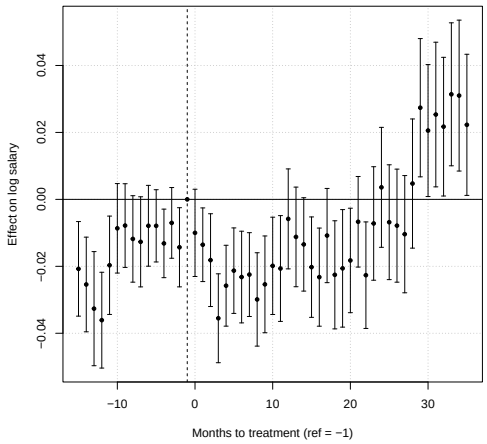
(d) Job Postings—Expertise Change



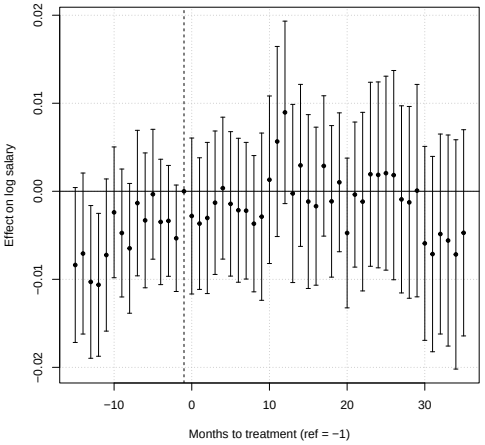
Notes: Other details are as in Figure A12.

Figure A24: Event Study: France

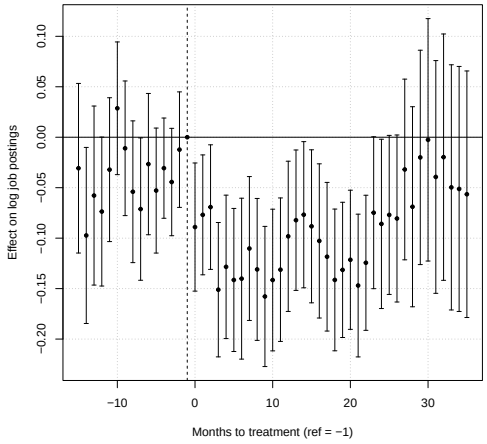
(a) Salary—Exposure



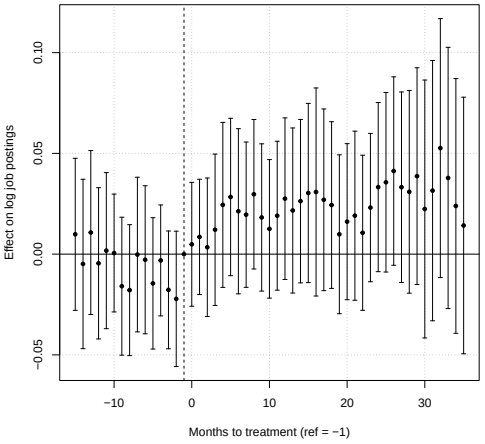
(b) Salary—Expertise Change



(c) Job Postings—Exposure



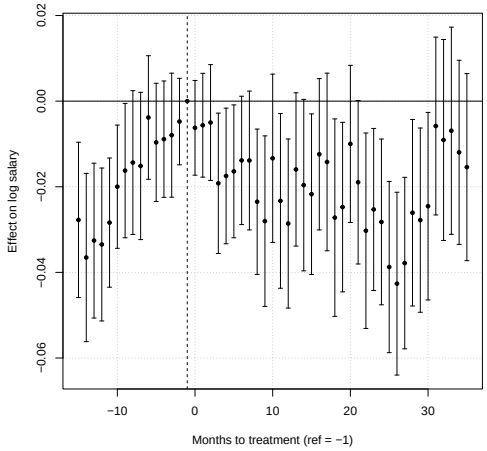
(d) Job Postings—Expertise Change



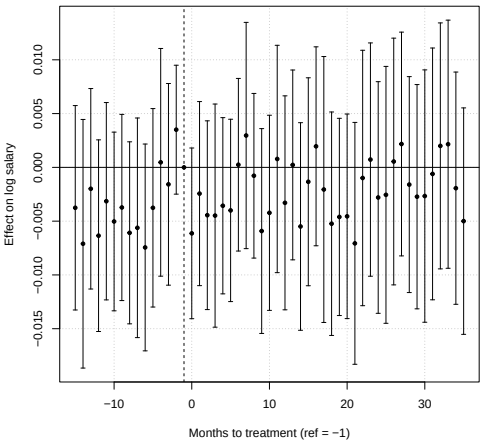
Notes: Other details are as in Figure A12.

Figure A25: Event Study: Germany

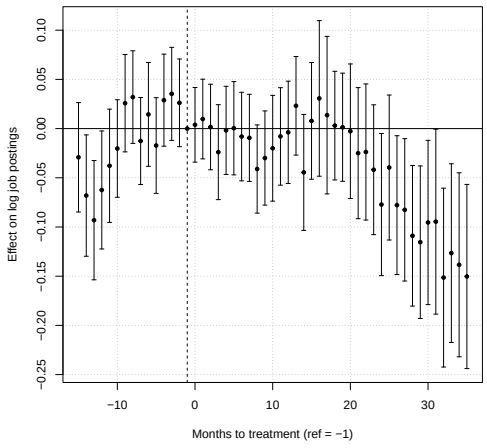
(a) Salary—Exposure



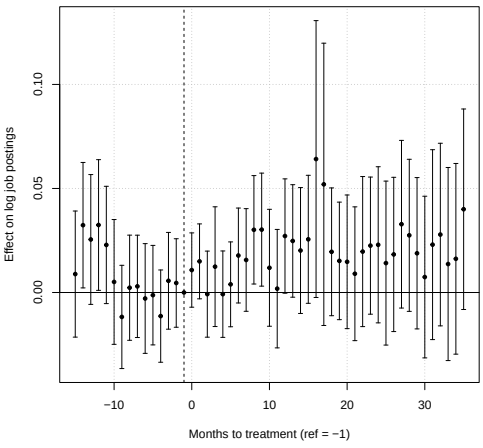
(b) Salary—Expertise Change



(c) Job Postings—Exposure



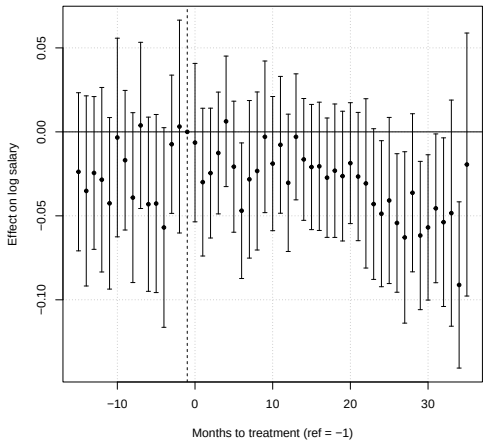
(d) Job Postings—Expertise Change



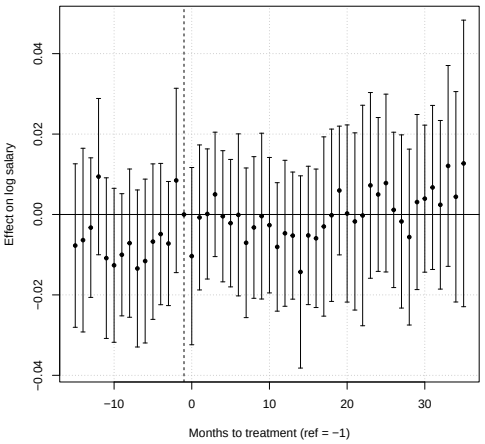
Notes: Other details are as in Figure A12.

Figure A26: Event Study: Greece

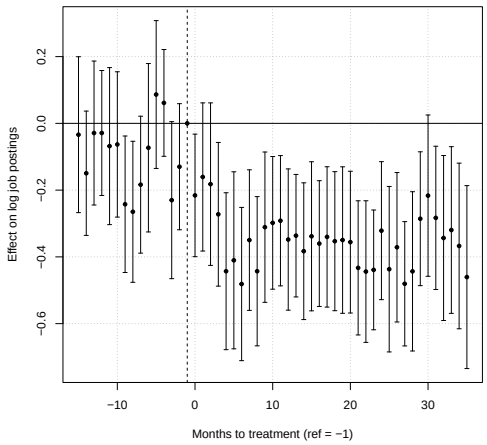
(a) Salary—Exposure



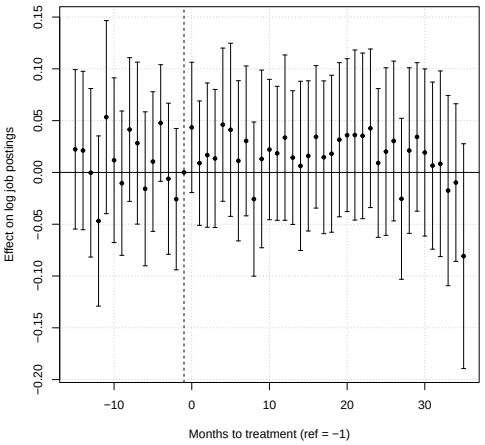
(b) Salary—Expertise Change



(c) Job Postings—Exposure



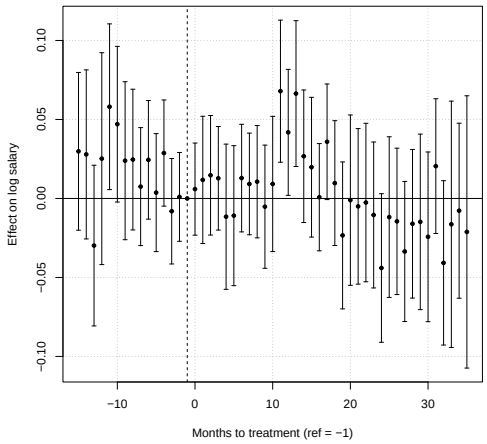
(d) Job Postings—Expertise Change



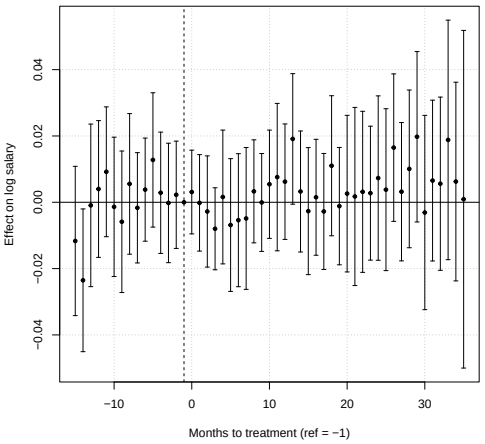
Notes: Other details are as in Figure A12.

Figure A27: Event Study: Hungary

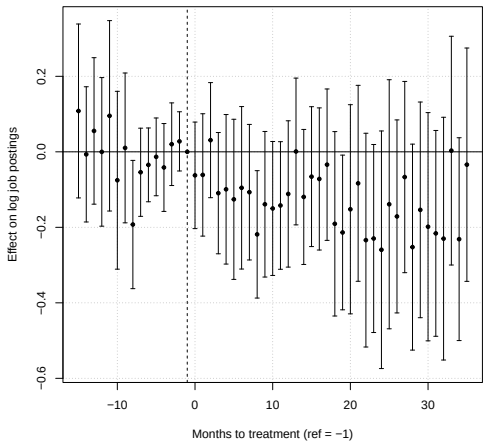
(a) Salary—Exposure



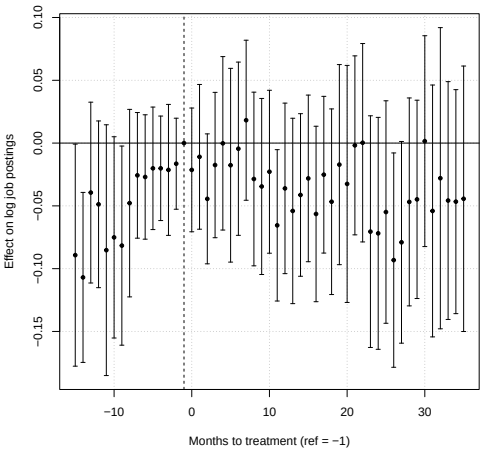
(b) Salary—Expertise Change



(c) Job Postings—Exposure



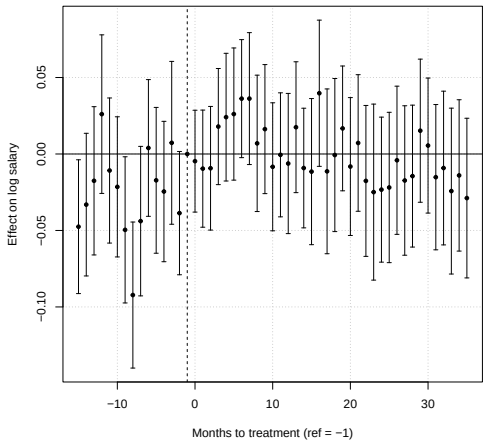
(d) Job Postings—Expertise Change



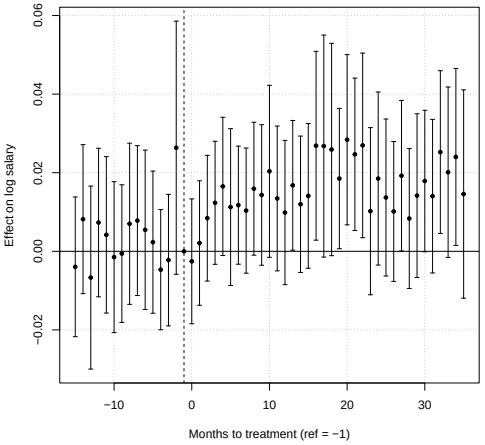
Notes: Other details are as in Figure A12.

Figure A28: Event Study: India

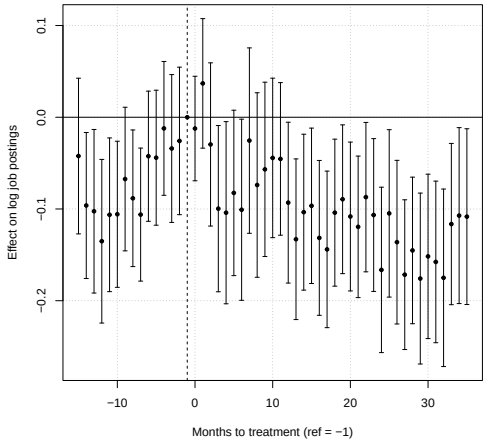
(a) Salary—Exposure



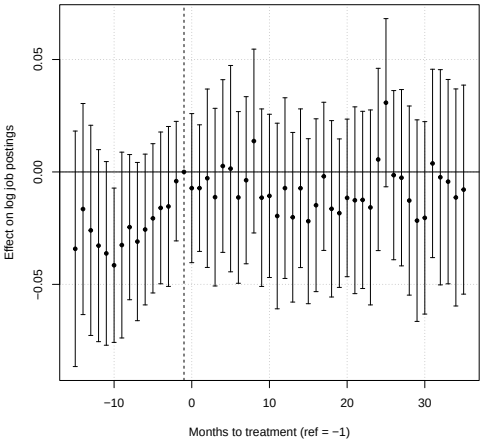
(b) Salary—Expertise Change



(c) Job Postings—Exposure



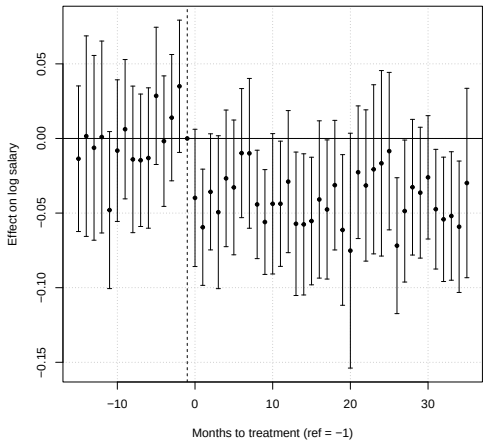
(d) Job Postings—Expertise Change



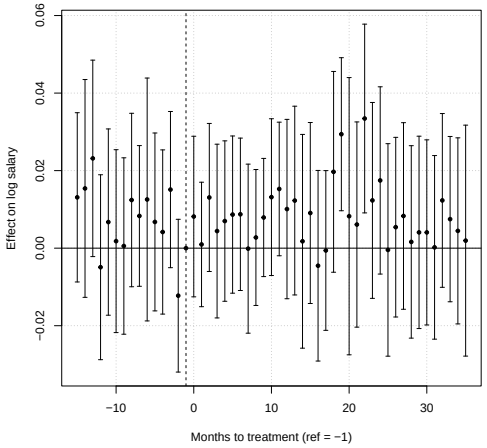
Notes: Other details are as in Figure A12.

Figure A29: Event Study: Indonesia

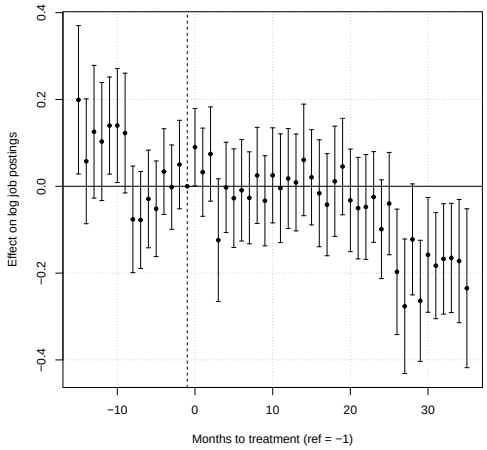
(a) Salary—Exposure



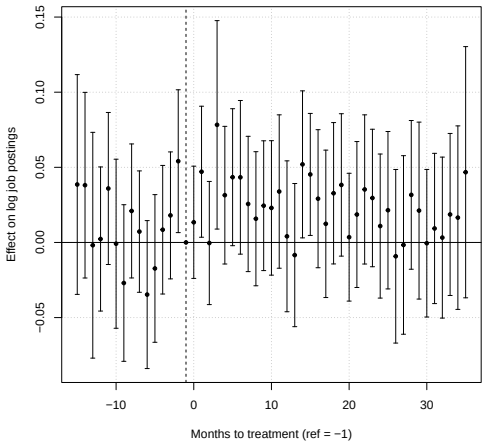
(b) Salary—Expertise Change



(c) Job Postings—Exposure



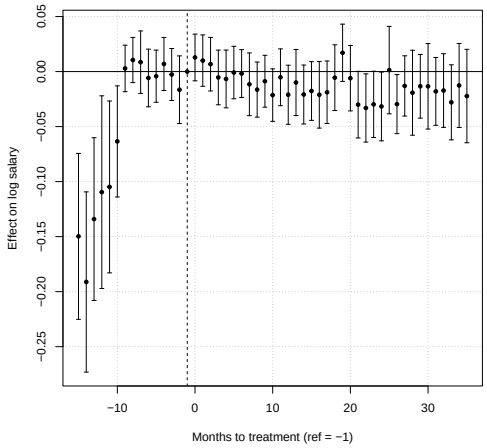
(d) Job Postings—Expertise Change



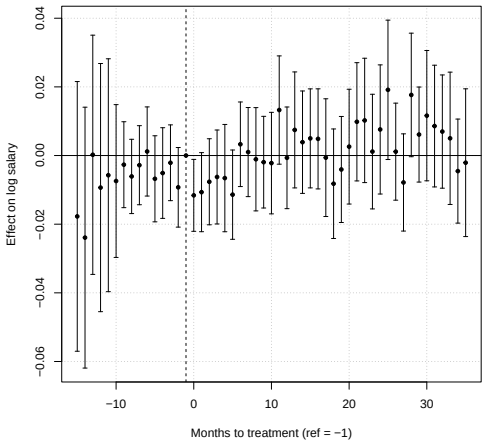
Notes: Other details are as in Figure A12.

Figure A30: Event Study: Ireland

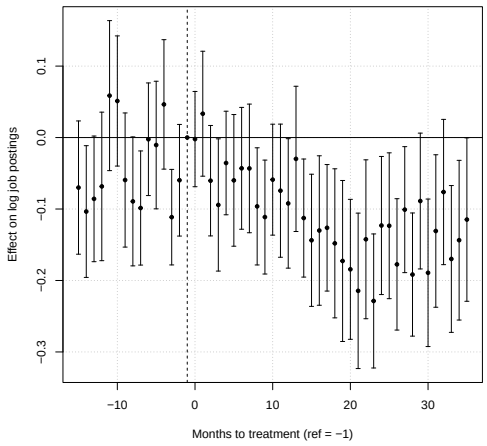
(a) Salary—Exposure



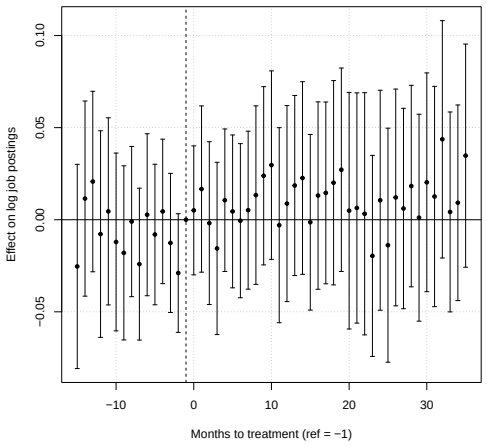
(b) Salary—Expertise Change



(c) Job Postings—Exposure



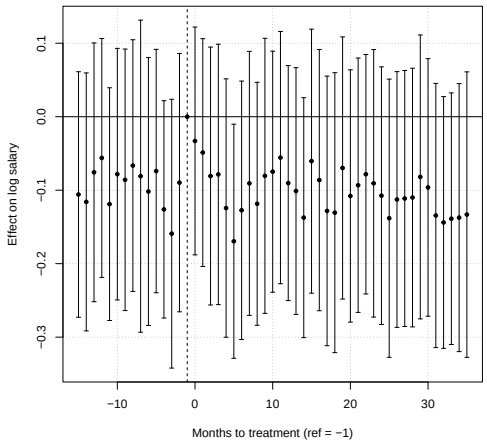
(d) Job Postings—Expertise Change



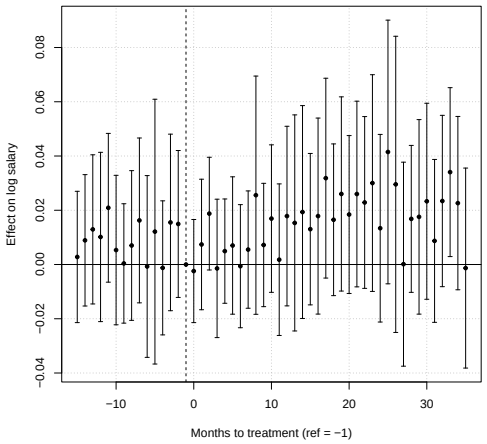
Notes: Other details are as in Figure A12.

Figure A31: Event Study: Israel

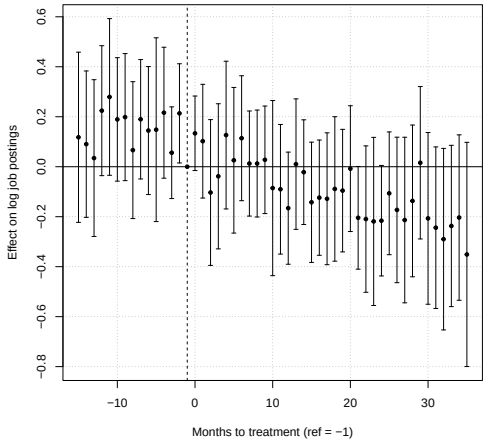
(a) Salary—Exposure



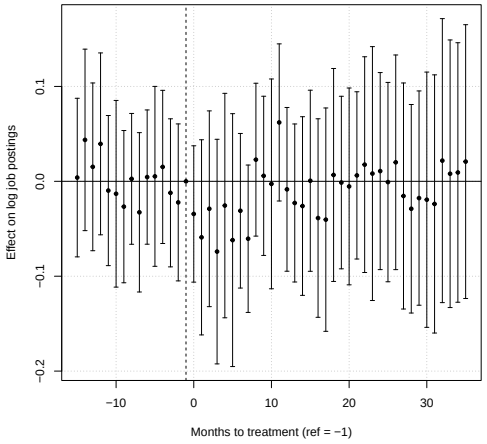
(b) Salary—Expertise Change



(c) Job Postings—Exposure



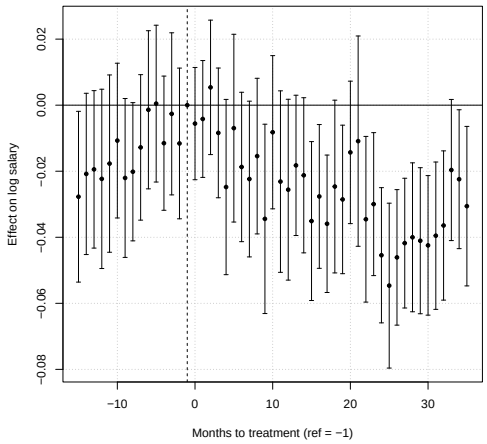
(d) Job Postings—Expertise Change



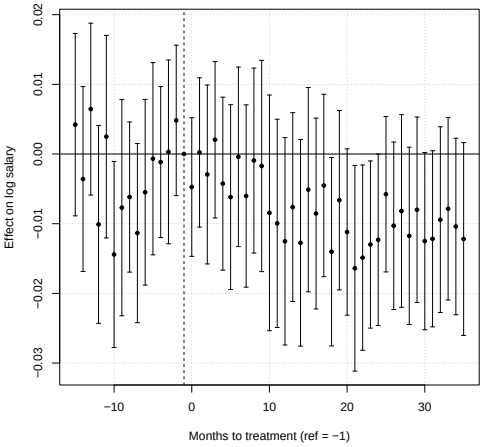
Notes: Other details are as in Figure A12.

Figure A32: Event Study: Italy

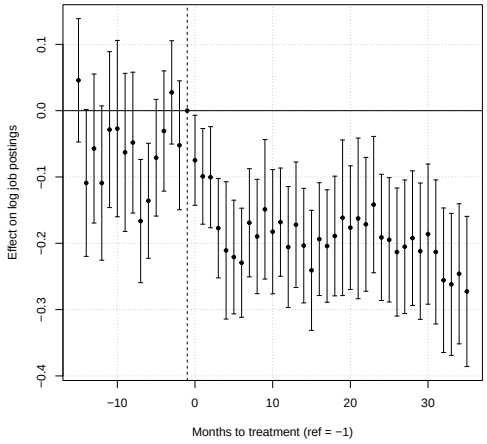
(a) Salary—Exposure



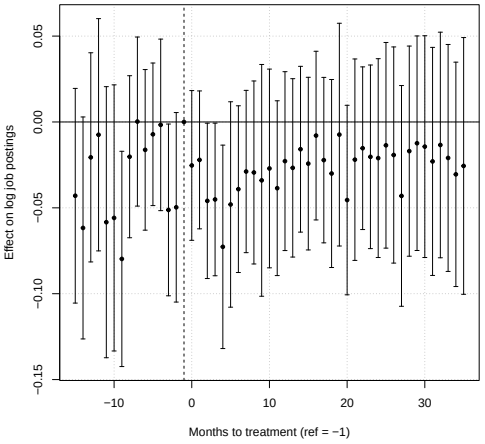
(b) Salary—Expertise Change



(c) Job Postings—Exposure



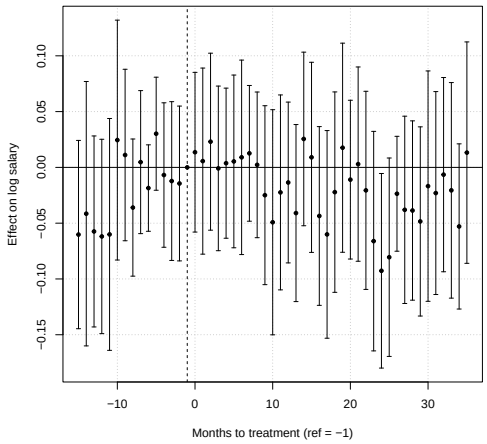
(d) Job Postings—Expertise Change



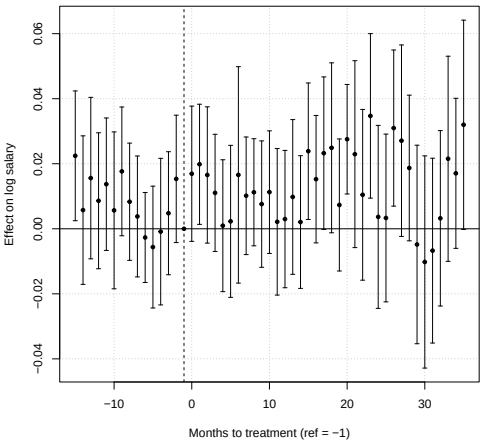
Notes: Other details are as in Figure A12.

Figure A33: Event Study: Luxembourg

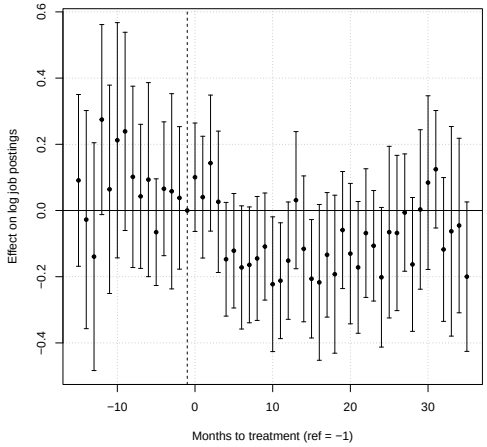
(a) Salary—Exposure



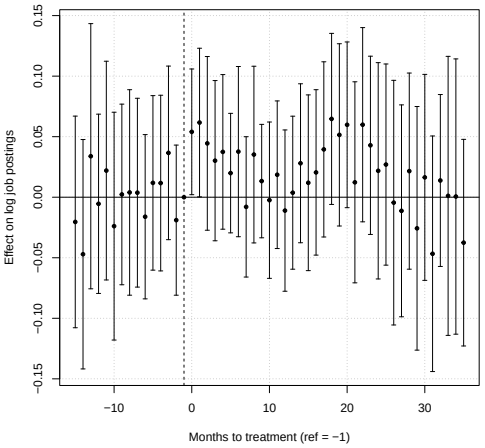
(b) Salary—Expertise Change



(c) Job Postings—Exposure



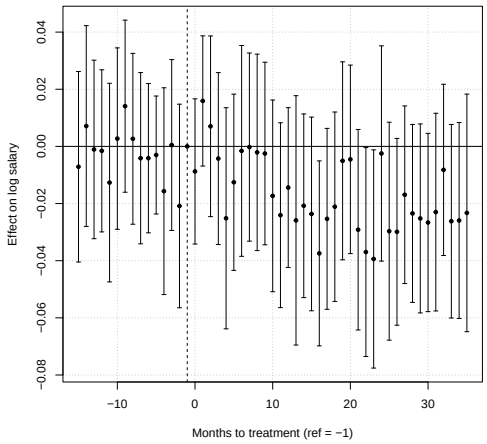
(d) Job Postings—Expertise Change



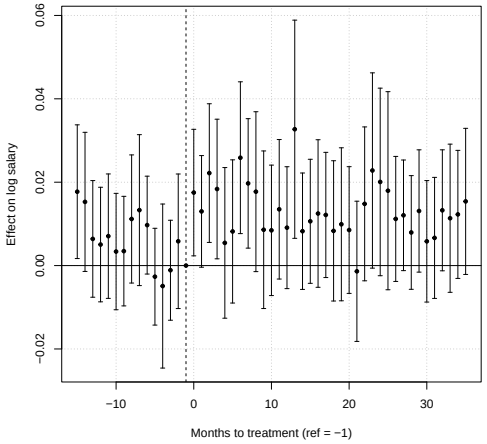
Notes: Other details are as in Figure A12.

Figure A34: Event Study: New Zealand

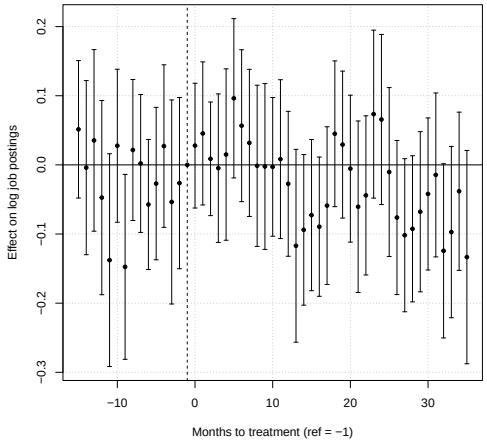
(a) Salary—Exposure



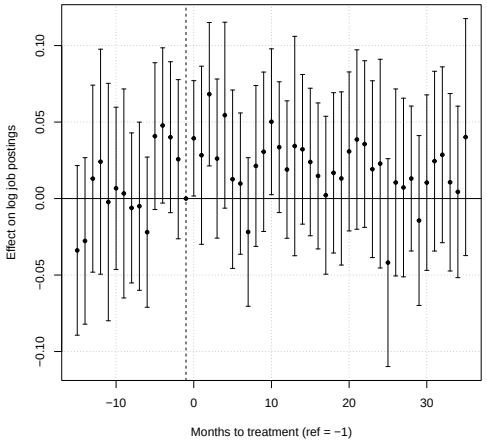
(b) Salary—Expertise Change



(c) Job Postings—Exposure



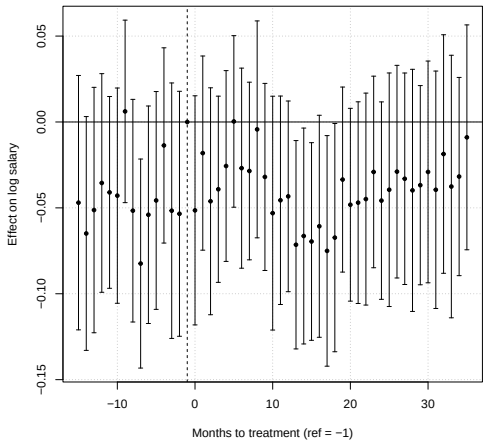
(d) Job Postings—Expertise Change



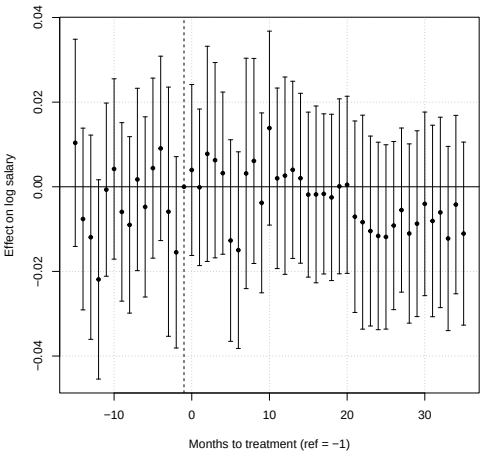
Notes: Other details are as in Figure A12.

Figure A35: Event Study: Norway

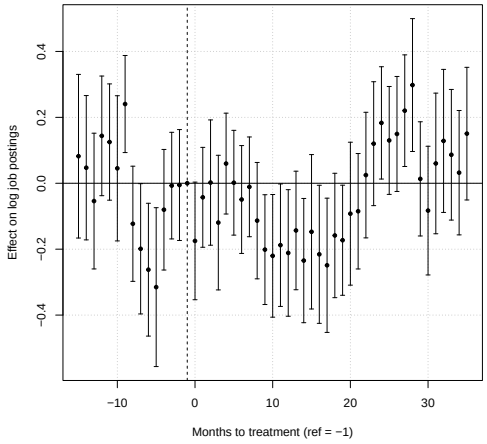
(a) Salary—Exposure



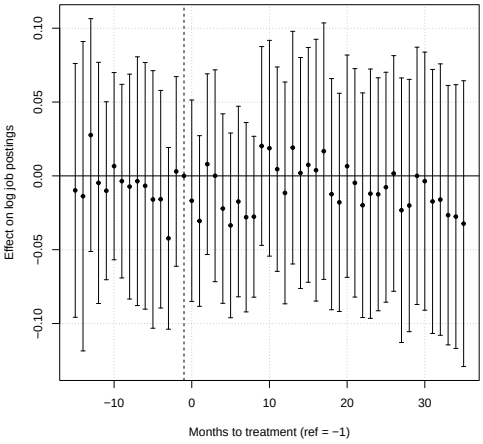
(b) Salary—Expertise Change



(c) Job Postings—Exposure



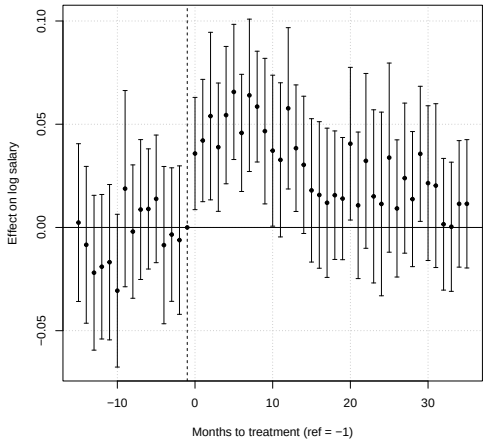
(d) Job Postings—Expertise Change



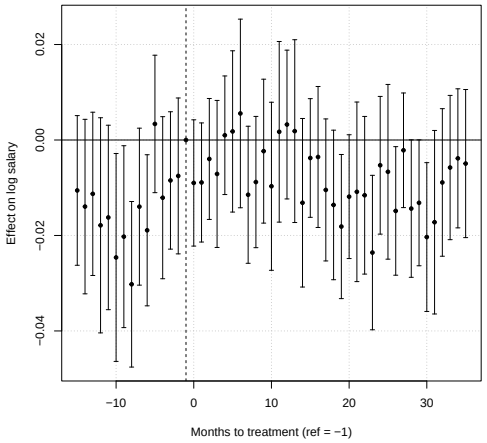
Notes: Other details are as in Figure A12.

Figure A36: Event Study: Portugal

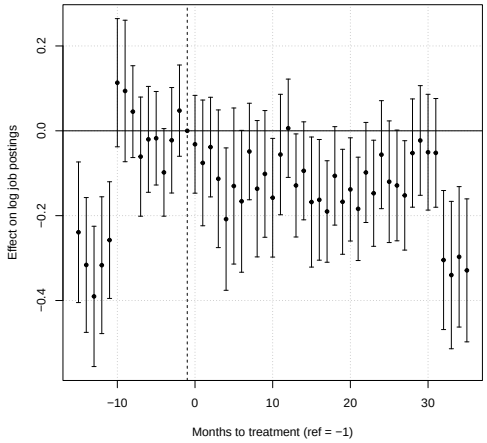
(a) Salary—Exposure



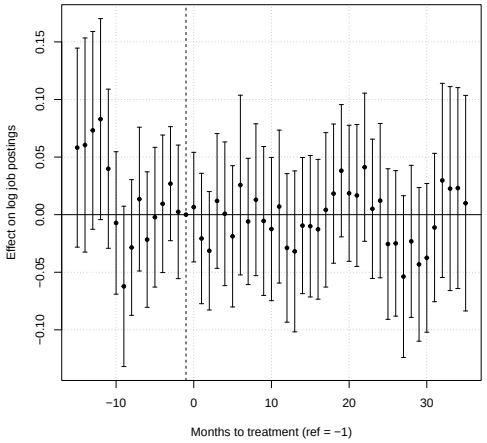
(b) Salary—Expertise Change



(c) Job Postings—Exposure



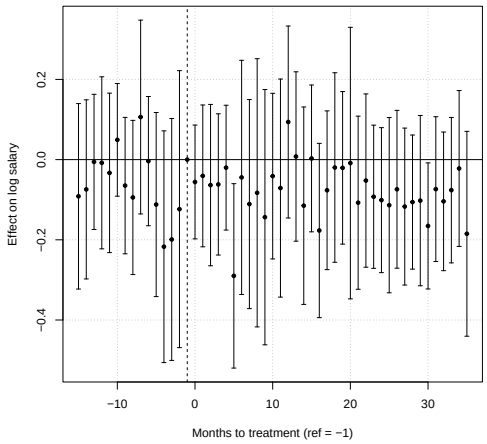
(d) Job Postings—Expertise Change



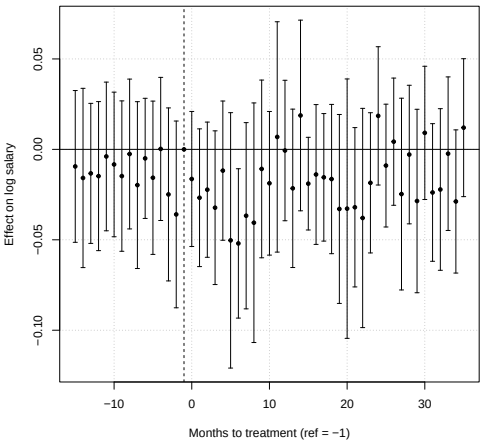
Notes: Other details are as in Figure A12.

Figure A37: Event Study: Russia

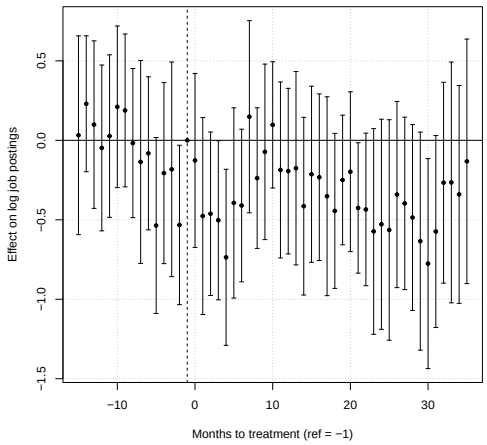
(a) Salary—Exposure



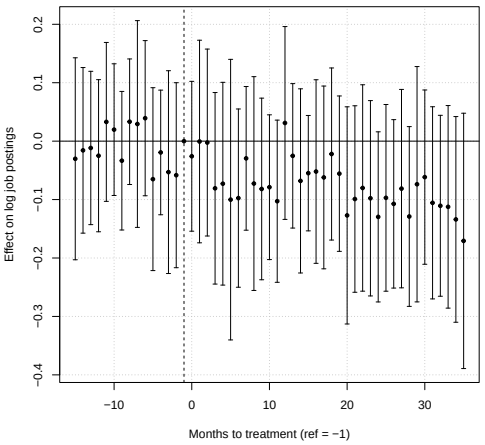
(b) Salary—Expertise Change



(c) Job Postings—Exposure



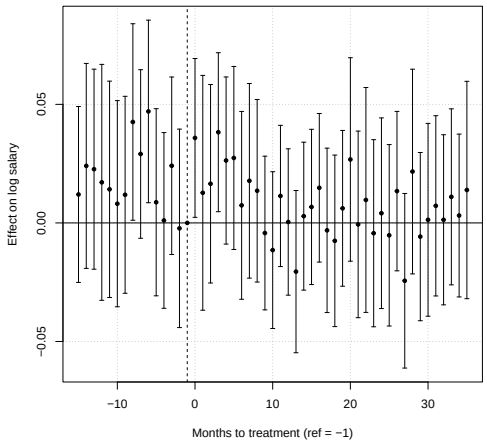
(d) Job Postings—Expertise Change



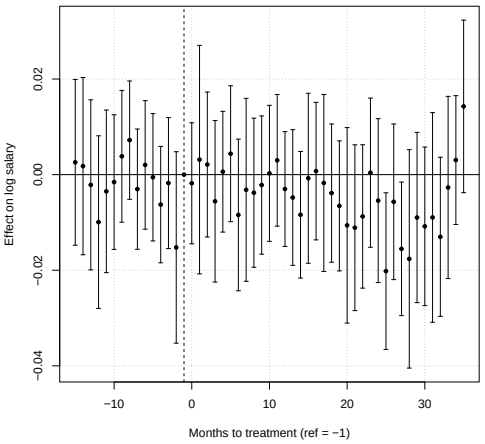
Notes: Other details are as in Figure A12.

Figure A38: Event Study: Saudi Arabia

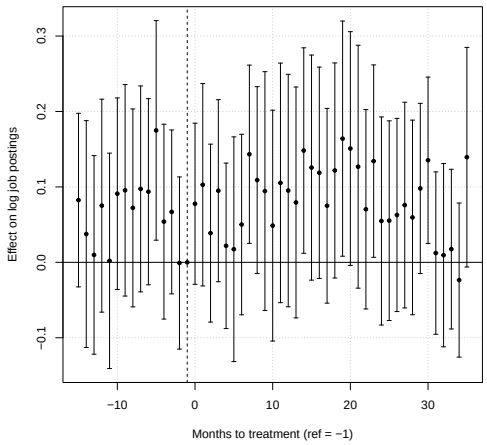
(a) Salary—Exposure



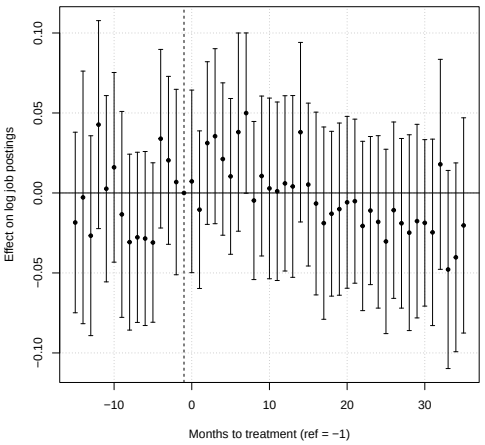
(b) Salary—Expertise Change



(c) Job Postings—Exposure



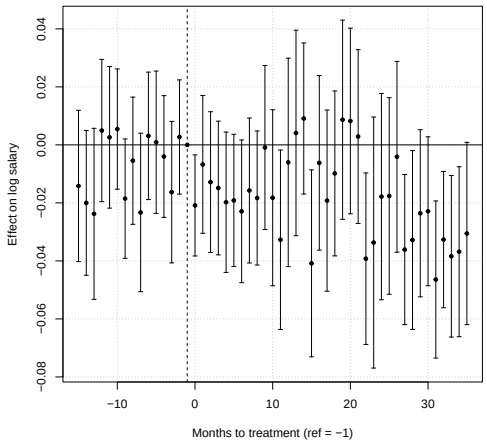
(d) Job Postings—Expertise Change



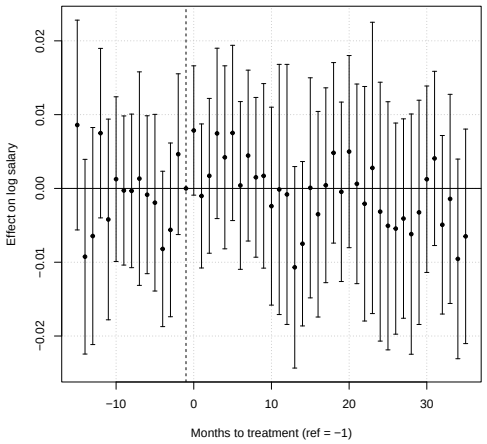
Notes: Other details are as in Figure A12.

Figure A39: Event Study: Singapore

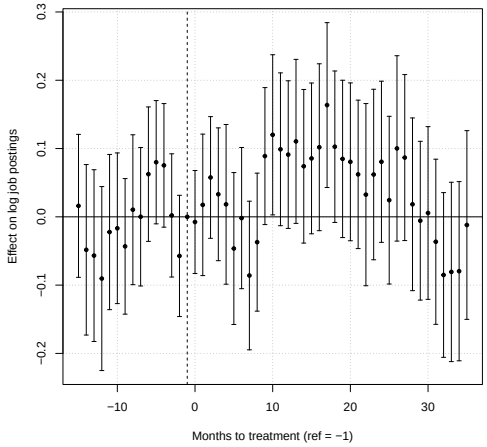
(a) Salary—Exposure



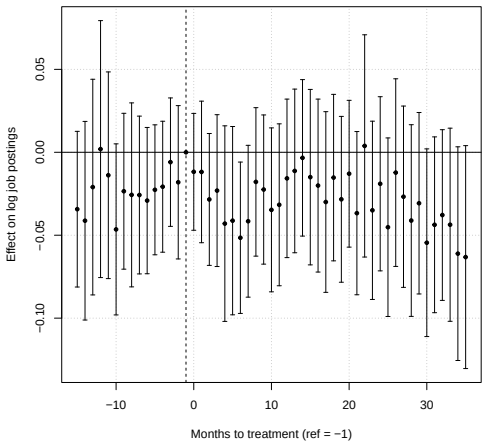
(b) Salary—Expertise Change



(c) Job Postings—Exposure



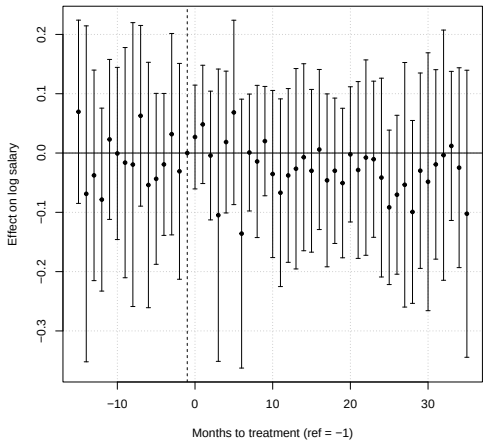
(d) Job Postings—Expertise Change



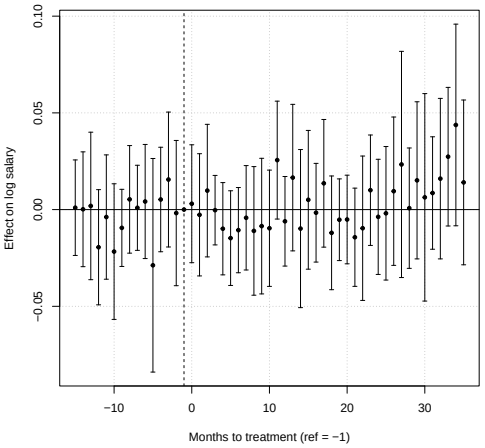
Notes: Other details are as in Figure A12.

Figure A40: Event Study: Slovakia

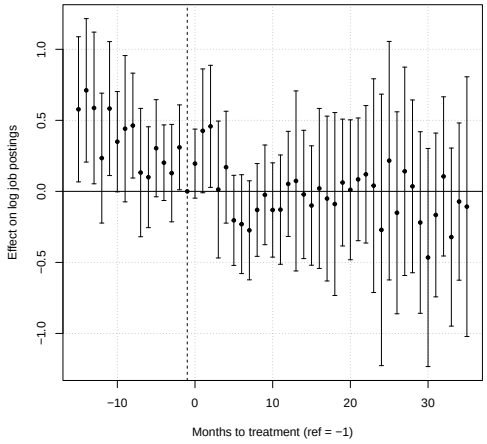
(a) Salary—Exposure



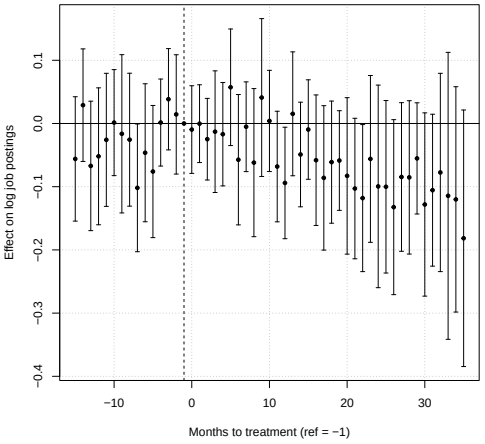
(b) Salary—Expertise Change



(c) Job Postings—Exposure



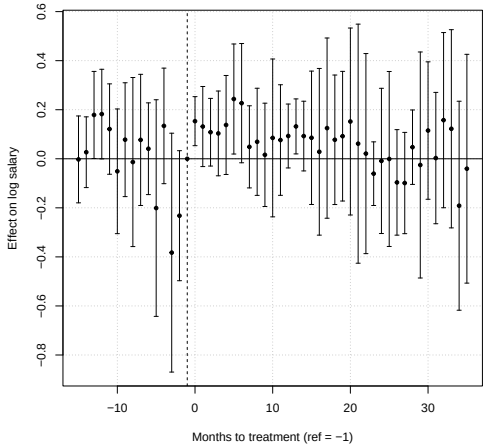
(d) Job Postings—Expertise Change



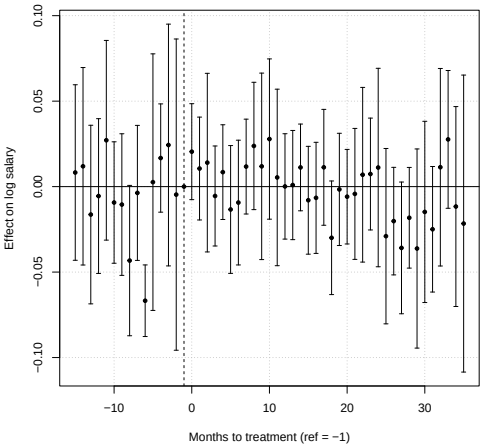
Notes: Other details are as in Figure A12.

Figure A41: Event Study: Slovenia

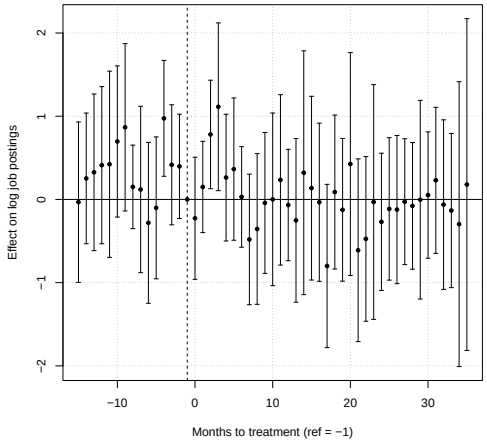
(a) Salary—Exposure



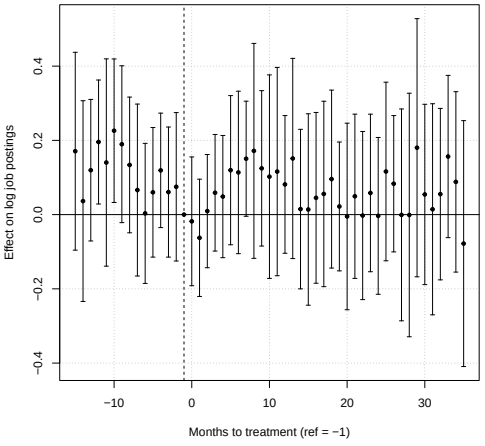
(b) Salary—Expertise Change



(c) Job Postings—Exposure



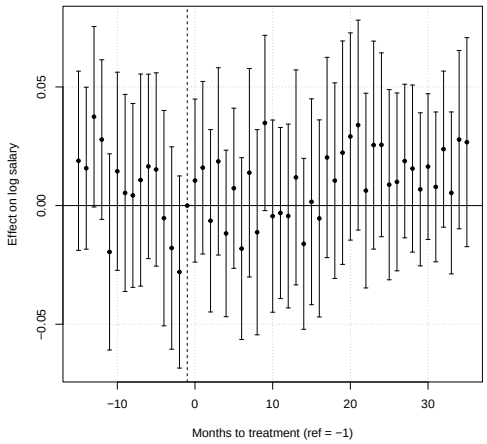
(d) Job Postings—Expertise Change



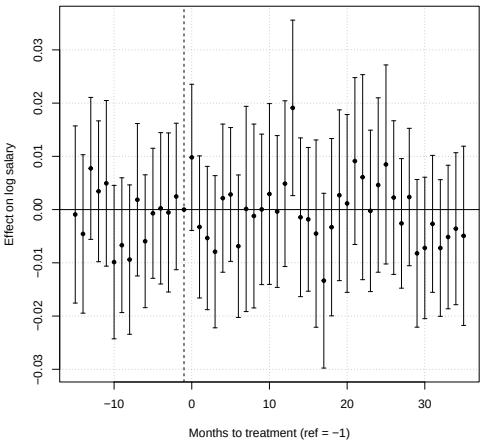
Notes: Other details are as in Figure A12.

Figure A42: Event Study: South Africa

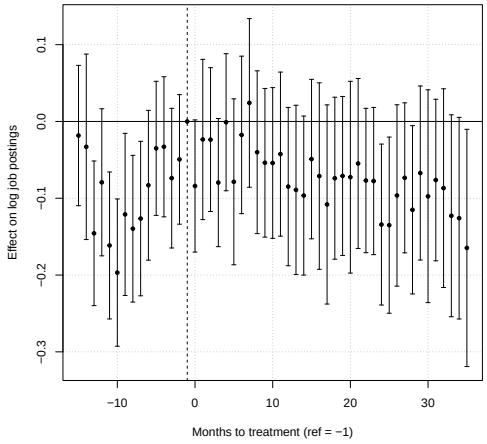
(a) Salary—Exposure



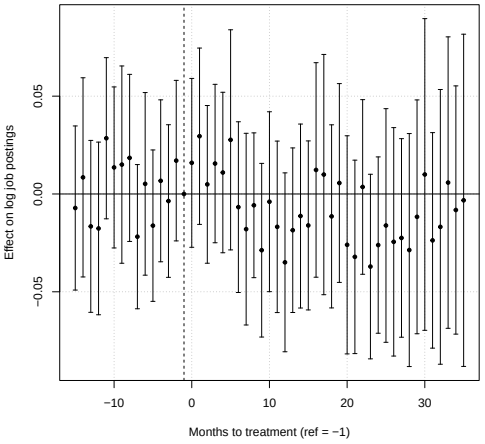
(b) Salary—Expertise Change



(c) Job Postings—Exposure



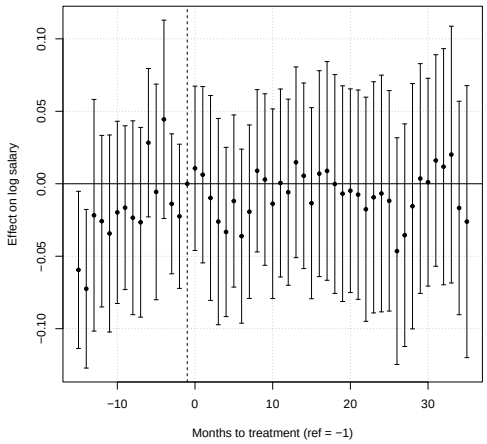
(d) Job Postings—Expertise Change



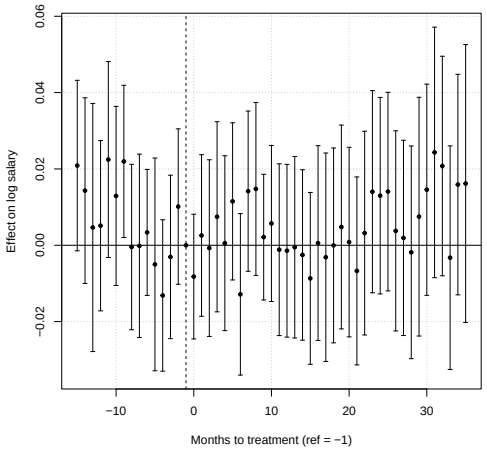
Notes: Other details are as in Figure A12.

Figure A43: Event Study: South Korea

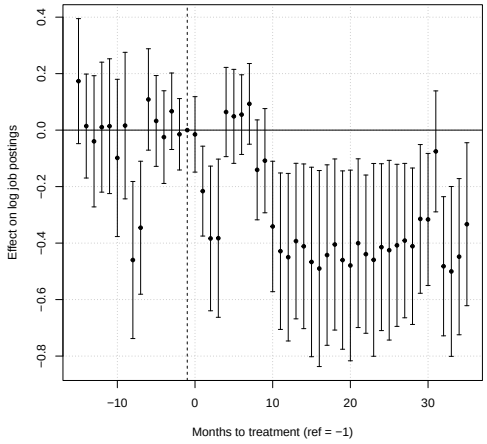
(a) Salary—Exposure



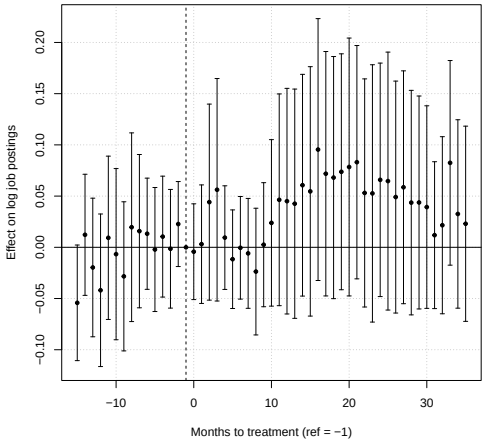
(b) Salary—Expertise Change



(c) Job Postings—Exposure



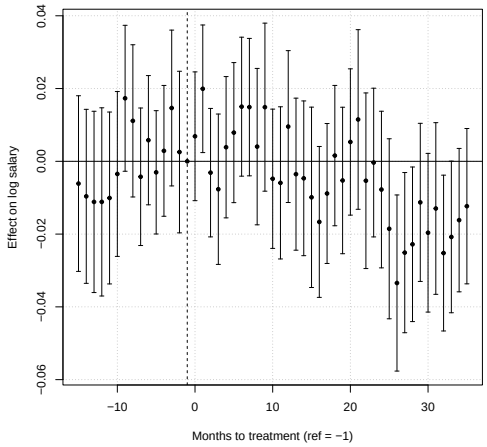
(d) Job Postings—Expertise Change



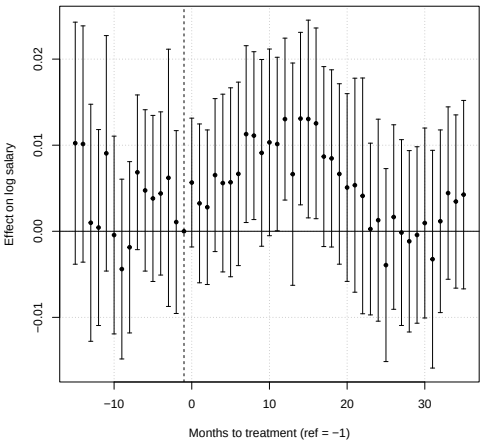
Notes: Other details are as in Figure A12.

Figure A44: Event Study: Spain

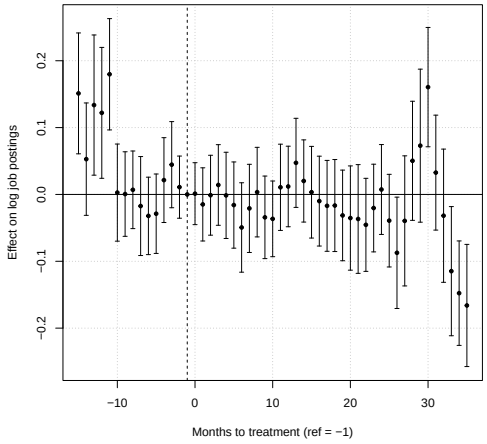
(a) Salary—Exposure



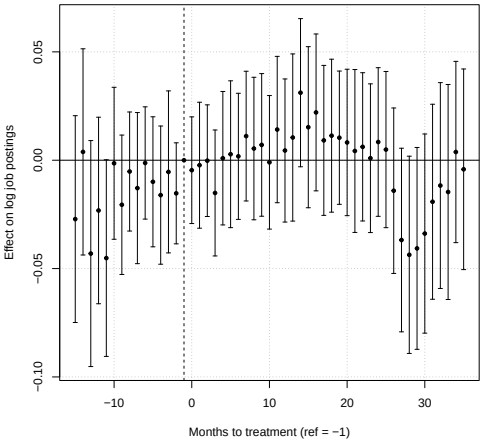
(b) Salary—Expertise Change



(c) Job Postings—Exposure



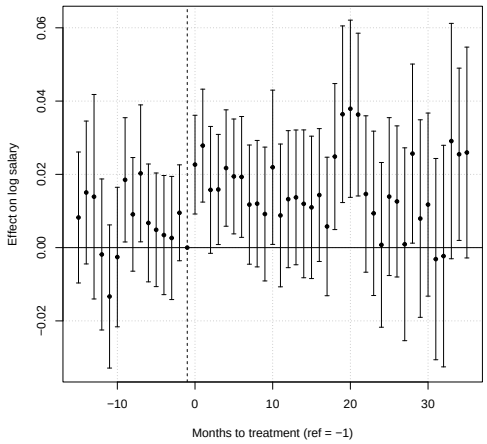
(d) Job Postings—Expertise Change



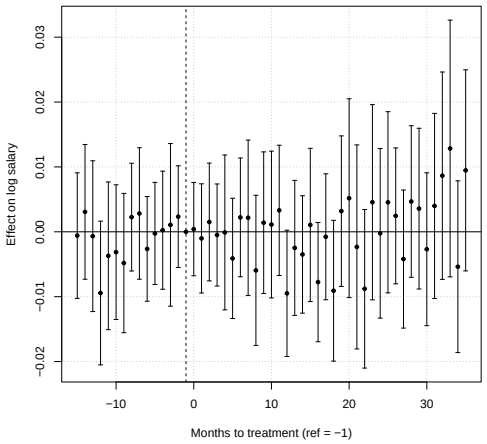
Notes: Other details are as in Figure A12.

Figure A45: Event Study: Sweden

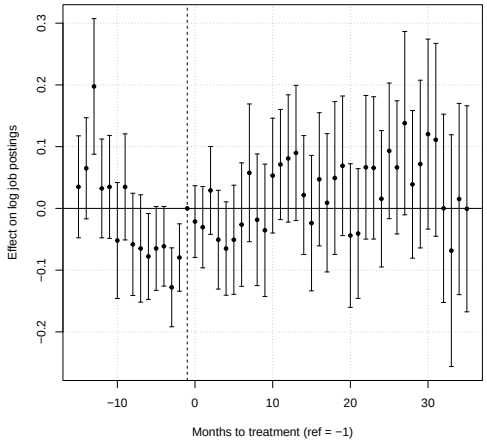
(a) Salary—Exposure



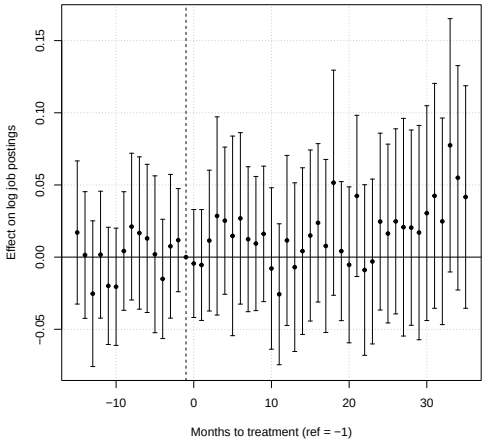
(b) Salary—Expertise Change



(c) Job Postings—Exposure



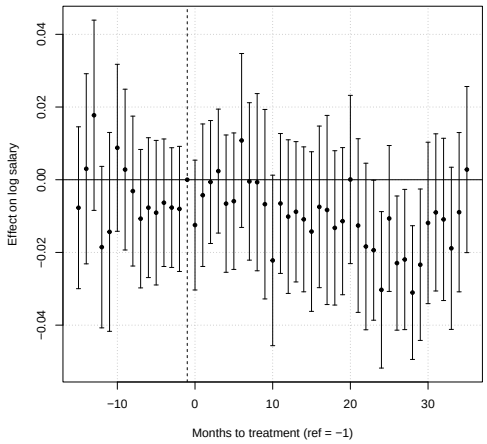
(d) Job Postings—Expertise Change



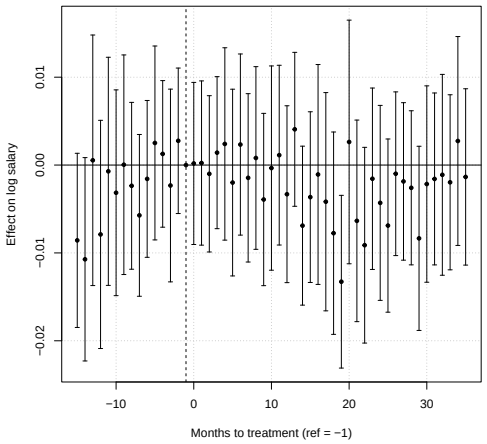
Notes: Other details are as in Figure A12.

Figure A46: Event Study: Switzerland

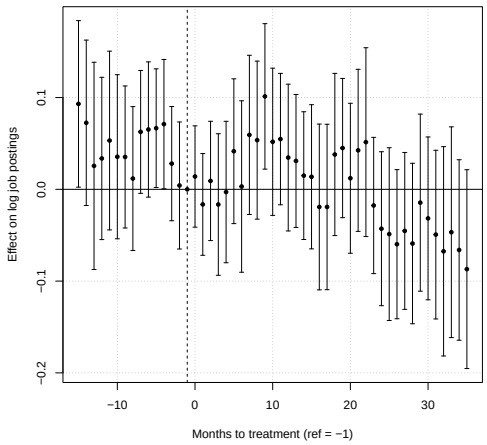
(a) Salary—Exposure



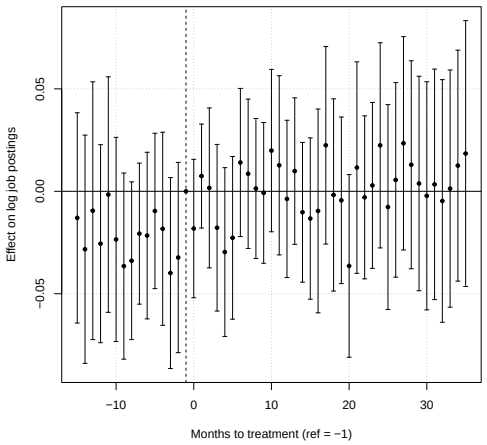
(b) Salary—Expertise Change



(c) Job Postings—Exposure



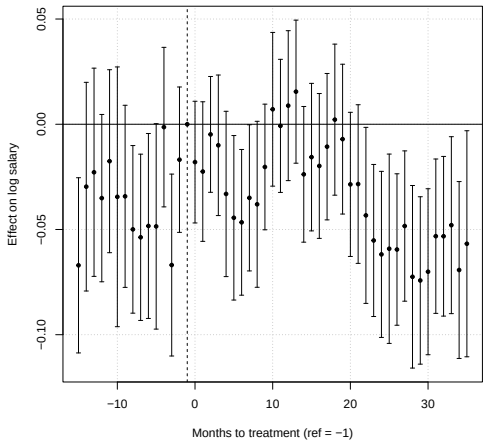
(d) Job Postings—Expertise Change



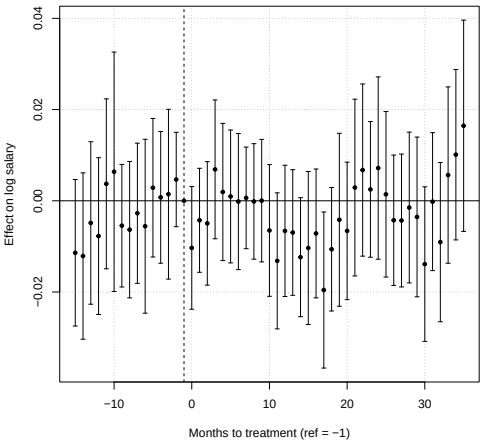
Notes: Other details are as in Figure A12.

Figure A47: Event Study: Turkey

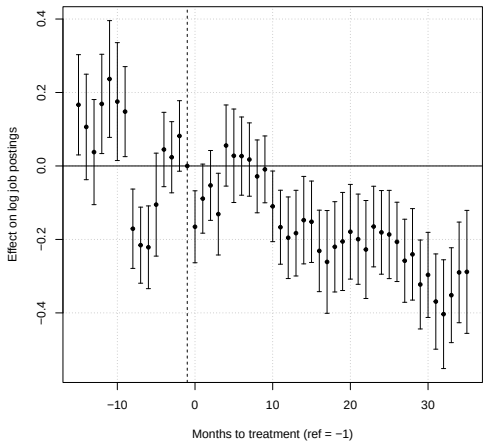
(a) Salary—Exposure



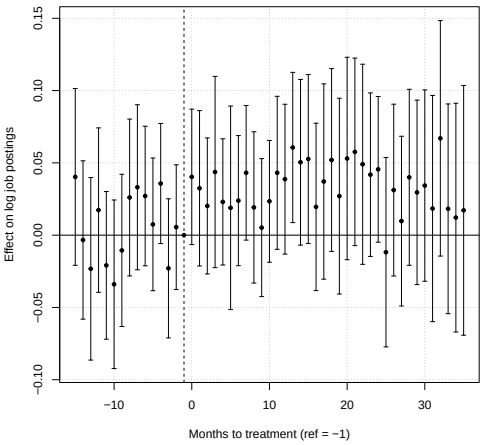
(b) Salary—Expertise Change



(c) Job Postings—Exposure



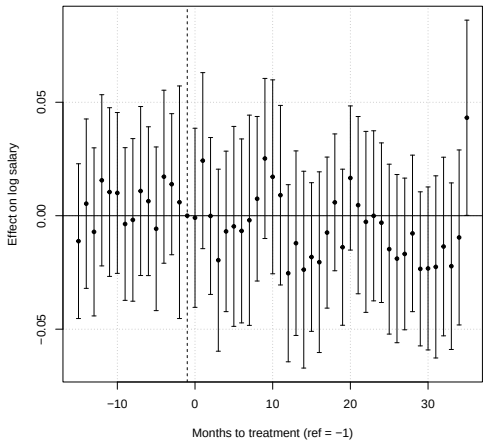
(d) Job Postings—Expertise Change



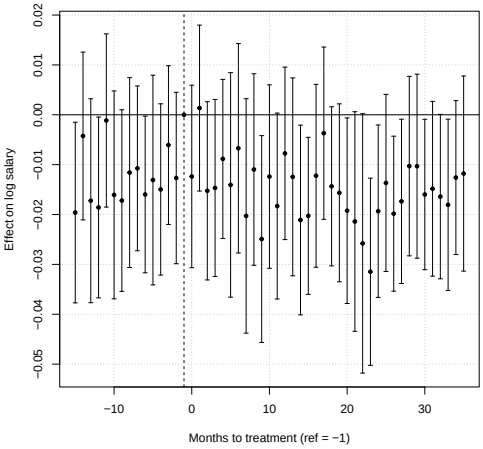
Notes: Other details are as in Figure A12.

Figure A48: Event Study: United Arab Emirates

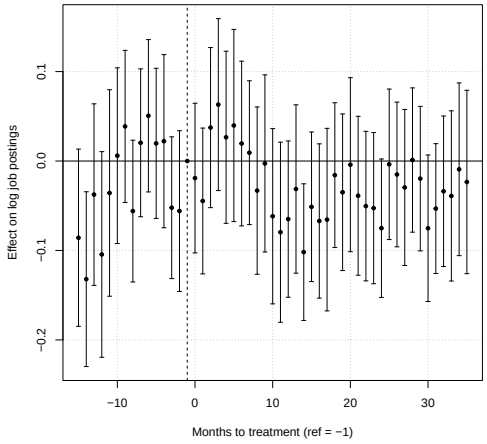
(a) Salary—Exposure



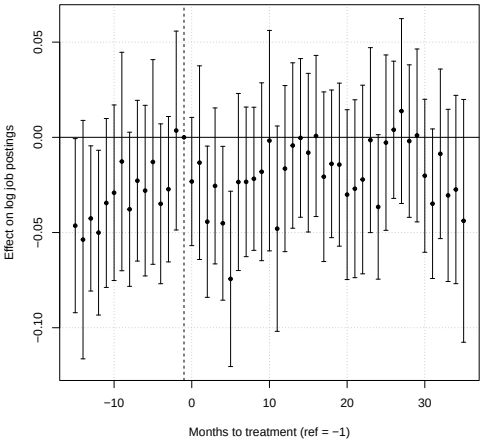
(b) Salary—Expertise Change



(c) Job Postings—Exposure



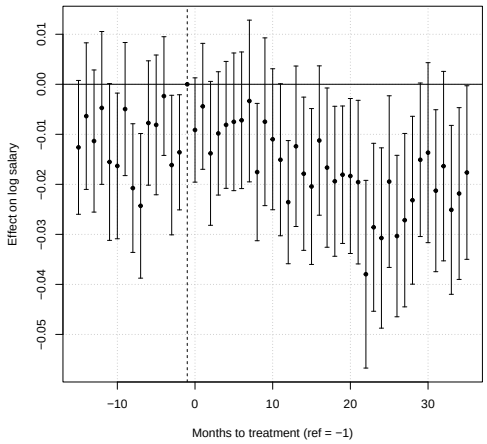
(d) Job Postings—Expertise Change



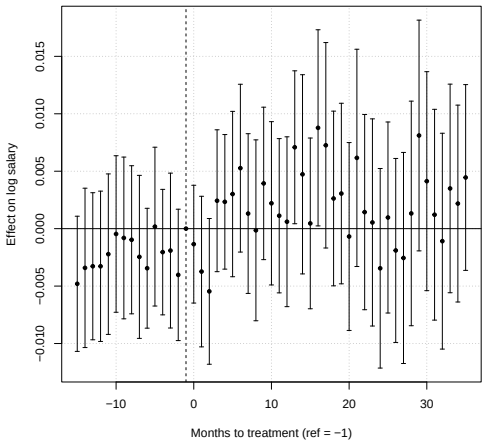
Notes: Other details are as in Figure A12.

Figure A49: Event Study: United Kingdom

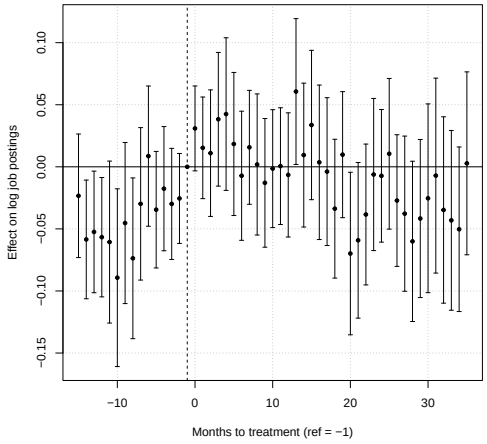
(a) Salary—Exposure



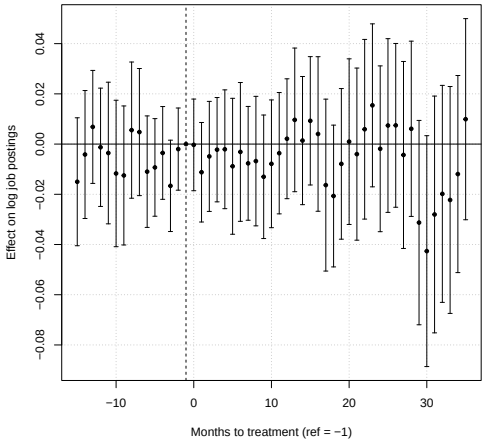
(b) Salary—Expertise Change



(c) Job Postings—Exposure



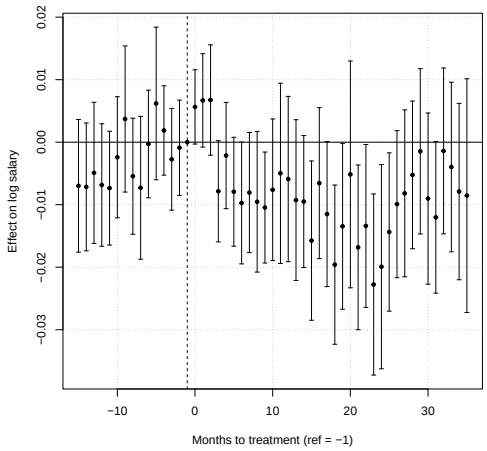
(d) Job Postings—Expertise Change



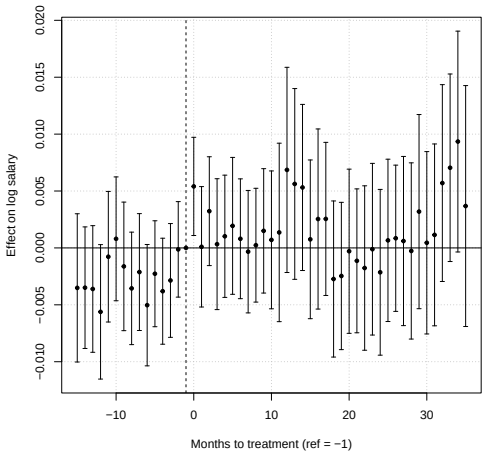
Notes: Other details are as in Figure A12.

Figure A50: Event Study: United States

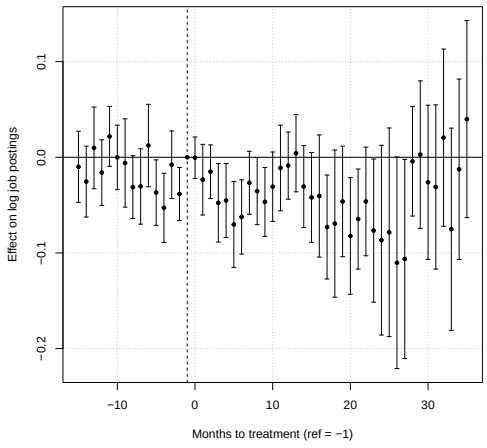
(a) Salary—Exposure



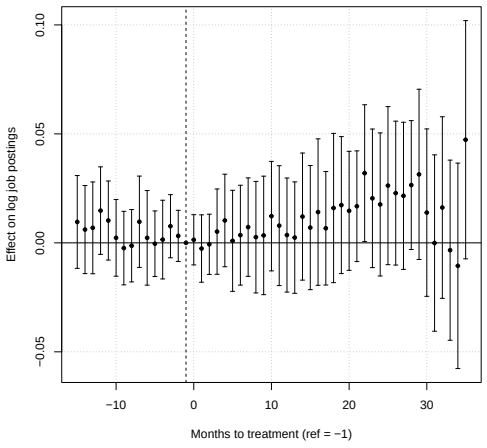
(b) Salary—Expertise Change



(c) Job Postings—Exposure



(d) Job Postings—Expertise Change



Notes: Other details are as in Figure A12.

A5 Changes in Minimum and Maximum Expertise Levels

Our main analysis uses the average change in expertise across tasks within each occupation. This choice reflects the framework's focus on the overall expertise composition of an occupation. However, alternative aggregations may capture different mechanisms. The *minimum* expertise change captures how automation affects an occupation's entry threshold—the lowest-expertise task that workers must still perform. The *maximum* expertise change captures how automation affects the highest-expertise tasks, which may drive wage determination for the most skilled workers in the occupation.

We compute these alternative measures as:

$$\Delta\text{Expertise}_o^{\min} = \min_{i \in T_o^{\text{not exposed}}} \text{Expertise}_i - \min_{i \in T_o} \text{Expertise}_i \tag{6}$$

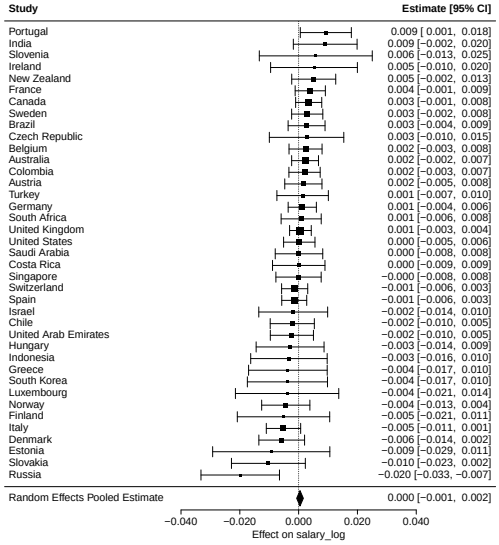
$$\Delta\text{Expertise}_o^{\max} = \max_{i \in T_o^{\text{not exposed}}} \text{Expertise}_i - \max_{i \in T_o} \text{Expertise}_i \tag{7}$$

A positive $\Delta\text{Expertise}_o^{\min}$ indicates that automation removes the lowest-expertise tasks, raising the entry barrier. A negative $\Delta\text{Expertise}_o^{\max}$ indicates that automation removes the highest-expertise tasks, lowering the ceiling of expertise required.

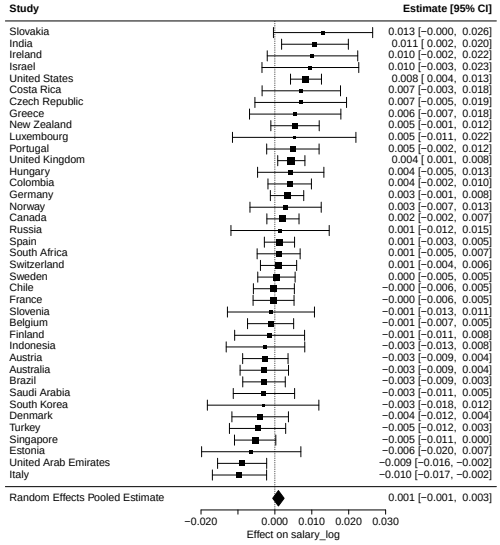
Because these measures depend on a single task rather than the average across all tasks, they are substantially noisier than our main measure. Classical measurement error would attenuate estimates toward zero. We present these results as a robustness check, with the caveat that smaller or insignificant effects may reflect attenuation bias rather than the absence of true effects.

Figure A51: Robustness: Minimum and Maximum Expertise Change

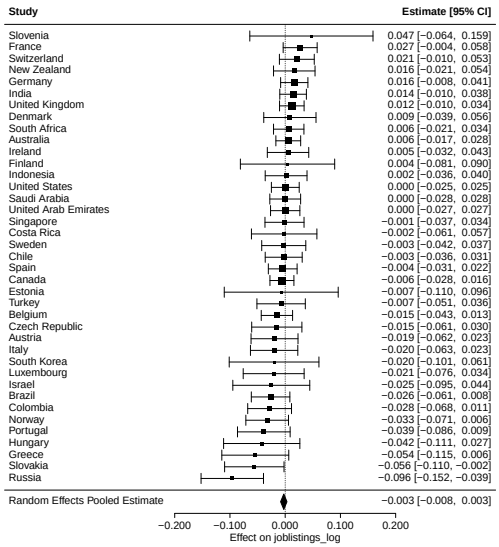
(a) Salary—Min Expertise



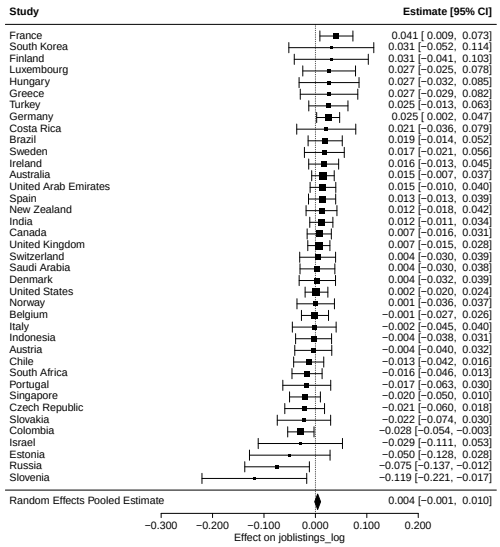
(b) Salary—Max Expertise



(c) Job Postings—Min Expertise



(d) Job Postings—Max Expertise



Notes: Left panels use *minimum* expertise change; right panels use maximum.