The Value of Non-Traditional Credentials in the Labor Market*

Susan Athey & Emil Palikot

October 5, 2026

Abstract

Non-traditional credentials can convey workers' skills, but earning a credential does not ensure that employers see it. We test whether helping workers share existing qualifications improves employment. Two randomized experiments with over 1.7 million Coursera learners combine reminders with simpler sharing on LinkedIn. Both increase certificate views. Among 45,037 learners with linked profiles, encouragement raises reported new jobs by 0.82 percentage points (SE 0.34) and new jobs with new employers in related fields by 0.71 points (SE 0.27). Adding certificate to resumes improves a language model's assessment of job fit. Views, display, and observed resume scores capture 56 percent of the employment difference. Greater local-market exposure to the experiment is associated with shorter time to fill in vacancies consistent with faster matching. The results show a return to helping workers communicate qualifications they already hold.

*Athey: Stanford University; Palikot: Northeastern University. We thank Eric Karsten, Yixi Jiang, and Anna Zhao from Coursera for collaborating on this project. We thank Keshav Agrawal and Elena Pittarokoili for excellent research assistance. We thank Aaron Kaye, Ambre Nicolle, Lukasz Grzybowski, and participants at seminars at LMU Munich, the University of Warsaw, Boston University, and Northeastern University, as well as the European Winter and North American Summer Meetings of the Econometric Society, NEUDC, Advances with Field Experiments, and the Platform Strategy Symposium, for helpful comments. The Golub Capital Social Impact Lab at Stanford Graduate School of Business provided funding for this research. This research has been subject to review and approval by Research Compliance Office at Stanford University, protocol number IRB-59983 and registered at AEA RCT registry AEARCTR-0009438 .

1 Introduction

Workers without formal educational credentials experience substantially lower employment rates and earnings than their credentialed counterparts (Card, 1999; Goldin and Katz, 2018). A central question for policy is whether these poor outcomes reflect skill deficiencies or information frictions that prevent employers from observing workers' existing capabilities. If it is the former, these workers may need intensive training to acquire marketable skills. If it is the latter, interventions that simply make workers' skills visible to employers may have substantial benefits at minimal cost. Non-traditional credentials can help convey these skills, but earning a certificate does not ensure that employers see it. We study whether helping workers make already-earned credentials visible improves their labor-market opportunities.

We study this question through two randomized experiments with more than 1.7 million Coursera learners. All had completed a course and earned a certificate before assignment. The interventions reminded learners to add the certificate to LinkedIn and simplified sharing. They provided no additional instruction, assessment, or credential. This design separates the return to signaling a credential from the decision to acquire it.

We find that helping learners share their credential increases the probability of reporting employment. Among 45,037 learners whose LinkedIn profiles we observe, encouragement increases the probability of a new job by 0.82 percentage points (standard error 0.34), or 4.9 percent of the control mean. New jobs with new employers in fields related to the certificates (*new job in scope*) increase by 0.71 points (SE 0.27), or 7.4 percent. We estimate the value of credential sharing by instrumenting sharing with treatment assignment. We find a local average treatment effect of 7.46 percentage points (3.09) for learners induced to initiate sharing on the new job outcome and 6.43 points (SE 2.49) for the new job in scope outcome.

Both experiments increase the visibility of credentials. In the first, an email reminder and simplified sharing link raise the probability of a certificate-page view by 4.20 percentage points, from 16.6 percent. The increase is larger for learners without college degrees and those in developing countries. The second experiment targets these populations through an in-app notification. It raises certificate views by 0.66 points. Both interventions also increase views by people other than the learner. The reminders therefore change both what learners display and the attention their credentials receive.

Employment gains tend to be larger for learners with lower baseline employability. To establish

employability, we predict whether subjects in the control group report new jobs using pre-treatment characteristics (Athey et al., 2025). A ten-percentile increase in predicted employability is associated with a 0.30-point decline in the treatment effect (standard error 0.15). The effect is 2.17 points (SE 0.66) in the bottom 40 percent and -0.25 points (SE 0.57) in the remaining 60 percent.

A reminder to add a certificate might also prompt a learner to list a job they already hold. Yet restricting attention to positions dated at least four months after the cohort month still yields estimated effects of 0.54 percentage points (SE 0.26) for any job and 0.46 percentage points (SE 0.20) for a new job in scope. The employment effect could also reflect greater engagement on LinkedIn, but we find no evidence of broader profile updating. Treatment and control groups have similar profile completeness, measured by the number of characters in each profile after excluding newly reported jobs and added certificates (Kolmogorov–Smirnov $D = 0.010$, $p = 0.22$). We also find no difference between groups in the probability of adding information outside the jobs section, such as education or skills.

We examine two ways in which certificate sharing may improve employment prospects: helping workers get noticed and helping them make a stronger impression once considered. To study the assessment channel, we use an LLM-as-a-judge framework (Zheng et al., 2023): a language model, acting as a recruiter, evaluates each observed pre-treatment profile and a counterfactual version with added certificate information. This comparison isolates how the added information changes the model's assessment, holding all other profile characteristics fixed. We find that the added certificates substantially raise assessments of candidate job fit, in particular, for candidates with low baseline scores. To connect these assessments and other observed signals to employment, we combine certificate views, detected display, and the observed resume score in a surrogate index that predicts employment (Athey et al., 2026). We find that the surrogate index captures 56 percent of the unadjusted employment effect.

These findings raise a broader question: how much labor-market value is lost when workers fail to communicate qualifications they already possess? If undisclosed credentials prevent employers from identifying suitable candidates, improving disclosure could shorten vacancies and allow workers to use their existing skills sooner. Yet gains to encouraged learners may come at other workers' expense. Assessing the broader value of disclosure therefore requires evidence on whether employers find suitable workers faster, as well as on which workers obtain jobs. This motivates our examination of vacancy duration in the markets exposed to the intervention.

We examine employer-side outcomes using Revelio Labs data, which record LinkedIn profiles displaying credentials from the courses in our experiment and job advertisements for relevant positions.

First, we show that the number of profiles displaying credentials from the experiment's Coursera courses increases by 10.5 percent relative to profiles displaying other-provider credentials, after adjusting for prior trends. Second, we examine whether vacancy duration falls in markets with greater projected exposure to the encouragement. Across 779 city–field markets in twelve countries, a one-standard-deviation increase in projected exposure is associated with a 1.44-day reduction in vacancy duration (SE 0.46).[1] The negative association is also present in developing-country markets; the U.S. estimates are close to zero. This comparison assumes that more- and less-exposed markets would have followed similar paths after accounting for the included trends, common shocks, and field-specific changes associated with baseline market size. The evidence is consistent with faster matching.

Related literature. The education literature distinguishes skills from signals of ability (Spence, 1973; Lange and Topel, 2006), using credential thresholds and employer learning to separate these channels (Tyler et al., 2000; Aryal et al., 2022). We hold the acquired qualification fixed and vary assistance with communicating it.

Our experiment builds on research on how information reaches employers. Pallais (2014) creates public work histories and evaluations for inexperienced workers, while Carranza et al. (2022) disseminate skill assessments to workers and firms. Both find employment gains. In Uganda, disclosing non-cognitive skill certificates changes beliefs, matching, and earnings conditional on employment (Bassi and Nansamba, 2022). Certification combined with application assistance in Ethiopia (Abebe et al., 2021) and LinkedIn training in South Africa (Wheeler et al., 2022) also improve labor-market outcomes. We study a narrower intervention: reminders and simpler sharing for workers who already possess a verifiable qualification.

The MOOCs literature documents unequal participation and completion (Hansen and Reich, 2015) and learners' use of courses for skills, career transitions, and self-presentation (Laryea et al., 2021). Observational studies report employment and career benefits (Hadavand et al., 2018; Zhenghao et al., 2015), but completion is selected. Experiments offering access to Coursera in Costa Rica and Colombia have low completion rates and no detectable or imprecisely estimated employment effects (Novella et al., 2024; Majerowicz and Zárate, 2024). Our post-completion intervention isolates assistance with disclosure from access, learning, and selection into certification.

[1] We observe when an advertisement is no longer open, rather than whether a hire occurred. Our preferred interpretation is that the vacancy was filled, although a posting can expire or be withdrawn unfilled. Holding vacancy demand and posting practices fixed, an influx of candidates with relevant credentials should facilitate filling; it provides no direct reason for faster withdrawal of unfilled vacancies.

Research on credential value distinguishes perceptions from employer behavior. Respondents prefer MOOC certificates to no credential in hypothetical hiring choices (Rivas et al., 2020). Correspondence experiments find fewer callbacks for business degrees from for-profit online institutions than from nonselective public institutions (Deming et al., 2016), and for online than in-person degrees (Lennon, 2021). We complement these comparisons with reported employment among actual certificate earners, varying disclosure assistance rather than credential type.

Our market analysis connects to research on displacement. Crépon et al. (2013) randomize program intensity across markets; Pallais (2014) studies aggregate effects under assumptions about market trends. Our vacancy-duration comparison follows the latter approach.

The assessment exercise uses an LLM-as-a-judge framework (Zheng et al., 2023), with the measurement qualifications emphasized in economic applications (Ludwig et al., 2026; Asirvatham et al., 2026). The surrogate-index framework (Athey et al., 2026) relates observed signals to employment; neither exercise directly measures employers' hiring decisions.

Section 2 describes the shared setting, the two experiments, and measurement. Section 3 presents effects on visibility and employment and examines reporting. Section 4 studies who benefits and the information channel. Section 5 turns to local labor markets. Section 6 concludes.

2 Setting, Experiments, and Measurement

2.1 Experimental Treatments

Coursera offers online courses taught by universities and other education providers. Learners who meet a course's requirements can receive a certificate with a unique verification page. To display it on LinkedIn, they ordinarily copy information from Coursera into their professional profile. Only 16.7 percent of controls in our linked sample had the earned certificate detected on their profile at follow-up.

Both experiments reduced the effort required to share an existing credential. Each paired a reminder with a link that simplified adding certificate information to LinkedIn. The experiments differed in delivery channel and population (Table 1), but shared the same comparison: certificate earners randomly offered assistance with sharing versus controls with access to the ordinary process. The treatment combines reminders and simplification; the design does not separate their effects.

Experiment 1. Coursera randomized 855,000 certificate earners in equal proportions between June and August 2021. Treated learners received an email reminding them to share, with a link to a certificate

TABLE 1
THE TWO EXPERIMENTS

	Experiment 1	Experiment 2
Timing	June–August 2021	September 2022–March 2023
Population	Broad population of certificate earners	Earners in four fields, in developing countries or without a college degree
Randomized population	855,000	Approximately 880,000
Encouragement	Email with a simplified sharing link	In-app notification with a simplified sharing link; one reminder
Assignment	Equal treatment and control groups	Equal probability within monthly batches and strata
Shared outcome	Certificate-page views	Certificate-page views
Additional measurement	–	Delivery, clicks, detected display, and reported employment

Notes: Both interventions occur after certificate completion. Experiment 2 strata use course domain, developing-country status, and college-degree status. The strict survey-verifiable population is 880,663; the view and LinkedIn samples use different coverage and eligibility boundaries, reconciled in Appendix D.

page offering a streamlined route to LinkedIn; controls received no experimental email. Coursera later incorporated this email into the standard learner experience. Platform records measure visits to each learner's certificate page. The broad population allows us to compare responses by education and country. The larger visibility gains among learners without college degrees and in developing countries informed the population of the second experiment. Appendix H reports the first experiment's sample and balance.

Experiment 2. The second experiment covered learners completing credentials in Business, Computer Science, Data Science, or Information Technology who lived in developing countries (non-OECD) or lacked a college degree. Assignment took place with equal probability in monthly batches, stratified by domain, country group, and degree status.

Our sample has 880,663 learners. Business accounts for 37 percent, Computer Science for 25 percent, Data Science for 24 percent, and Information Technology for 14 percent. Credentials span 7,355 courses; individual courses and guided projects account for approximately 82 and 17 percent, respectively. Appendix C gives course examples and geographic coverage. The experiment was registered before randomization as AEARCTR-0009438. The registration identifies sharing and employment as primary outcomes; Appendix A provides further details.

Treated learners receive a sharing notification on a visit to the Coursera application after certificate completion. Learners who did not click could receive one reminder on a subsequent visit. The notification emphasized the achievement and invited the learner to add the certificate to LinkedIn

through a simplified process. Appendix B reproduces the wording and interface.

2.2 Assignment, Receipt, and Initiating Sharing

Platform records distinguish assignment, returning to the app, seeing a notification, and clicking to share. Among treated learners in the view sample, 89 percent returned, 48 percent saw a notification, and 4.1 percent clicked. In the linked employment sample, the rates are 96 percent, 62 percent, and 10.98 percent. We retain learners regardless of whether they returned or received a message.

Our primary estimates are intent-to-treat (ITT) effects of assignment to encouragement. We also estimate the average treatment effect on the treated (ATT), defining treatment as access to encouragement among assigned learners who launch the app. We instrument assignment times app launch with assignment and retain both arms. If assignment has no effect on nonlaunchers, this estimates the effect of encouragement among treated app-launchers (Bloom, 1984). Controls can visit the app but cannot access the experimental encouragement.

Finally, we instrument clicking to share with assignment to estimate a local average treatment effect (LATE) for learners induced to initiate sharing. This interpretation requires assignment to affect employment through that action; changes in search effort among nonclickers would violate the exclusion restriction. Controls cannot click the experimental notification but can share through ordinary channels. A click initiates sharing but the learner has to still approve it once redirected to LinkedIn.

2.3 Data and Outcomes

Coursera records provide assignment, certificate characteristics, the delivery and click indicators, and visits to certificate verification pages. Registration surveys provide country, education, and gender; course records also contain skill-assessment measures. Missing survey responses remain explicit categories. Appendix D documents linkage and coding.

Visibility. The main visibility outcome is whether the certificate's verification page receives a visit. We also distinguish LinkedIn referrals and visits classified by Coursera as coming from someone other than the learner, using login, network, and timing information. These visits measure attention to the certificate; visitors need not be recruiters, and employers can read a displayed credential without opening its verification page. The view sample contains 892,929 learners. Appendix D documents coverage.

LinkedIn profiles and sample selection. Learners could provide LinkedIn URLs at enrollment, before the intervention. We observe profiles for 45,037 learners: 22,487 control and 22,550 treated. The first month which contributes 8,356 learners was initially intended as a pilot. We include it throughout and report the main employment table without it in Appendix S. Assignment, receipt, click, and primary outcomes are complete in the linked sample.

Learners with linked profiles differ from the broader view sample. They more often completed Data Science credentials (30.2 versus 23.5 percent), less often completed guided projects (10.1 versus 16.6 percent), and more often reported education (83.3 versus 35.8 percent). They also respond more often to the sharing feature. The employment estimates therefore apply to learners with linkable profiles. Within the six-cohort administrative sample, profile coverage is similar across arms: 5.15 percent in control and 5.19 percent in treatment.

Display and employment. We identify display by finding the earned certificate's unique Coursera verification code in the scraped profile. Either arm can display the credential. Incomplete extraction can generate false negatives, so detected display measures an outcome of encouragement rather than receipt of treatment.

New Job records a new position, including moves within an employer. *New Job in Scope* additionally requires an employer change and a title in a field related to the certificates. We additionally examine jobs dated at least four months after the cohort. Appendix E gives the date rules and classifications.

Previous studies support using LinkedIn histories to measure employment. Among 154 U.S. inventors, profiles agree with survey reports on employers in 91 percent of inventor-years and on annual employer changes in 92 percent (Ge et al., 2016). At the firm level, Berry et al. (2026) validate workforce demographics against survey and company records. Public LinkedIn profiles also contain fewer false claims about responsibilities than private profiles or traditional resumes in an experiment (Guillory and Hancock, 2012).

3 Effects on Visibility and Employment

Encouragement increases both credential visibility and reported employment. We establish the visibility response in both experiments before turning to employment in the linked sample from the second experiment.

3.1 Sharing and Attention

Table 2 shows that both delivery channels increased certificate-page views. In Experiment 1, the email raised the probability of any view by 4.20 percentage points (standard error 0.10), a 25 percent increase over the control mean of 16.6 percent. The increase in views classified as coming from others was 1.62 points (0.08). Thus the response extended beyond learners revisiting their own certificates.

TABLE 2
EFFECTS ON CREDENTIAL VISIBILITY

	Control (%)	ITT (pp)	Learners
Panel A. Experiment 1: Email encouragement			
Any certificate-page view	16.6	4.20	855,000
		(0.10)	
Any view by others	13.6	1.62	855,000
		(0.08)	
Panel B. Experiment 2: In-app encouragement			
Any certificate-page view	19.5	0.66	892,929
		(0.08)	
Any view by others	14.6	0.29	892,929
		(0.07)	
Any LinkedIn-referred view	16.8	0.65	892,929
		(0.08)	
LinkedIn-referred view by others	12.7	0.27	892,929
		(0.07)	

Notes: Standard errors in parentheses. Experiment 1 uses differences in means and Welch standard errors. Experiment 2 uses reconstructed design-cell effects and administrative controls, with HC1 standard errors. September counts come from April 2023 and the later cohorts from October 2023; these are cohort-specific accumulated views, not a common-duration outcome. Visits concern certificate verification pages. Views by others exclude visits classified as the learner's own.

The visibility response is larger for learners with weaker conventional signals. In Experiment 1, any views increase by 4.9 percentage points among learners without a college degree and by 2.6 points among graduates, a difference of 2.3 points (0.2). The increase is 6.9 points in developing countries and 1.4 points elsewhere, a difference of 5.5 points (0.2). These comparisons can reflect both willingness to share and the attention a shared credential attracts (Appendix H.1).

In Experiment 2, encouragement increases any view by 0.66 percentage points (0.08), or 3.4 percent of the control mean. LinkedIn-referred views by others increase by 0.27 points (0.07). Among learners with linked profiles, 10.98 percent of the treated click to share, and detected display rises from 16.7 to 19.4 percent; the adjusted display effect is 2.45 points (0.35). Both delivery channels make existing credentials more visible. At the time of the Experiment 2, e-mail from Experiment 1 was a part of the standard app experience; thus, the in-app notifications in Experiment 2 are in addition to it. The

magnitudes of the two treatment effects also reflect differences in population, receipt, and follow-up, so they do not isolate the relative effectiveness of email and in-app messages.

3.2 New Jobs

Encouragement increases reported new jobs by 0.82 percentage points (0.34), from a control mean of 16.8 percent (Table 3). New jobs with new employers in fields related to the certificates increase by 0.71 points (0.27), or 7.4 percent of the 9.6 percent control mean. Excluding the September pilot gives an any-job ITT of 0.92 percentage points (0.38); Appendix S reports the full table.

TABLE 3
EFFECTS ON REPORTED EMPLOYMENT

	Control (%) (1)	ITT (pp) (2)	ATT (pp) (3)	LATE (pp) (4)
Panel A. Recorded new-position outcomes				
New Job	16.8	0.82	0.85	7.46
		(0.34)	(0.35)	(3.09)
New Job in Scope	9.6	0.71	0.73	6.43
		(0.27)	(0.28)	(2.49)
Panel B. Four-month restriction				
New Job	8.7	0.54	0.57	4.96
		(0.26)	(0.27)	(2.40)
New Job in Scope	4.9	0.46	0.47	4.15
		(0.20)	(0.21)	(1.87)

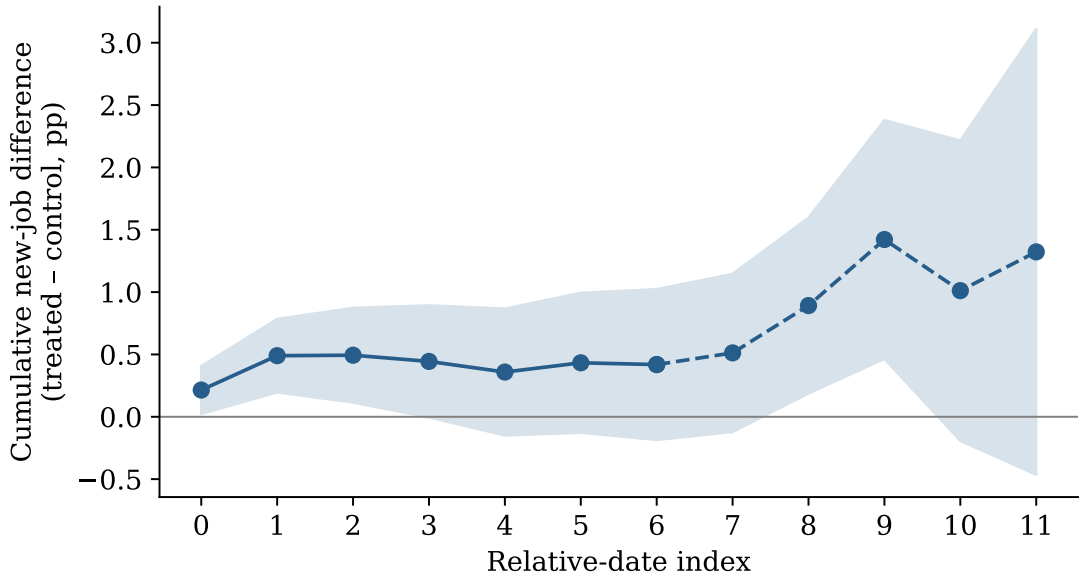
Notes: All 45,037 unique matched learners, including September. Standard errors in parentheses. ITT is assignment; ATT instruments assignment times app launch with assignment; LATE instruments clicking to share with assignment. Their adjusted first stages are 96.26 and 10.98 percentage points. Control means are strata-weighted. All columns include the full set of pre-treatment controls. Appendix S reports the same table without the pilot cohort.

The ATT among treated app-launchers is 0.85 percentage points (0.35) for any new job and 0.73 points (0.28) for a related-field job with a new employer. These effects are close to the ITT because 96 percent of assigned learners launch the app. Click LATEs are 7.46 points (3.09) and 6.43 points (2.49), respectively, using an adjusted first stage of 10.98 points. These are effects for learners induced to initiate sharing.

The estimates are similar with fewer controls and with no strata weighing. Unadjusted effects without strata weights are 0.97 (0.36) percentage points for any job and 0.83 (0.28) for an in-scope job. Adjusting only for randomization strata gives 0.98 (0.36) and 0.83 (0.28) Appendix F).

Figure 1 plots cumulative new-job differences over the available follow-up. The line is solid through relative-date index six and dashed thereafter. Fewer recruitment batches contribute to the later part of the dashed curve: all six remain through index seven, followed by five, three, two, and one at indices

Figure 1: Difference in Cumulative New-Job Rates



Notes: Treated-minus-control cumulative New Job differences, in percentage points, with pointwise 95 percent intervals using two-sample binomial standard errors. The solid line ends at index six; the dashed line extends the series using batches with available follow-up. All 45,037 learners enter indices zero through seven; the samples at indices eight through eleven contain 39,531, 23,622, 15,614, and 7,602 learners. We retain complete calendar months through April 2023 for the September batch and August 2023 for later batches, preceding their documented collection periods. The relative-date index is job-month index minus cohort plus one; zero includes the preceding month. Appendix E describes the date conventions.

eight through eleven. Changes in this part of the curve can therefore reflect changes in the mix of batches as well as employment. Appendix J reports cohort-specific estimates.

The new-job gain is concentrated in starts with new employers, including entries from nonemployment: the effect is 0.69 percentage points (0.31), compared with −0.04 points (0.16) for same-employer moves. The overall new-job effect is positive both for learners with a recorded job in the month before treatment and for those without one, but the difference between these groups is imprecise (Appendix K).

3.3 Employment or Profile Reporting?

A reminder to add a certificate may also prompt a learner to report a job obtained before treatment. This explanation predicts changes concentrated in jobs already held or accompanied by updates elsewhere in the profile. We examine later job starts and unrelated profile fields.

The employment difference persists for jobs dated well after the cohort month. Requiring a start at least four months later gives effects of 0.54 percentage points (0.26) for any job and 0.46 points (0.20) for a related-field job with a new employer. The corresponding ATT estimates are 0.57 and 0.47

points. Positions with these later start dates are harder to explain as newly reported jobs already held at treatment.

The employment effect could also reflect increased engagement with LinkedIn. Learners prompted to display a certificate might update other parts of their profiles, increasing the chance that they report a job they already held. We first examine whether treatment changes overall profile completeness. For each of the 45,037 linked learners, we count characters in recorded education and pre-cohort employment fields, excluding certificates, courses, and new-job entries. These exclusions prevent the certificate and employment outcomes themselves from mechanically changing the measure. A two-sample Kolmogorov–Smirnov test compares the full distributions of these character counts across treatment and control. It does not reject equality ($D = 0.00990$, $p = 0.220$; Appendix G.1).

We then examine specific profile entries that might change with greater LinkedIn engagement: education histories, listed courses, licenses or certifications, and earlier employment records. We find no statistically detectable response in these outcomes. The effect on the number of education entries is -0.0001 (0.011), and the effect on adding an education entry dated 2023 or later is -0.19 percentage points (0.13). Counts of courses and licenses or certifications also show no detectable changes (Appendix N.1). Greater engagement might also prompt learners to fill gaps in earlier employment histories by adding positions or end dates. The effect on recording an end date for a pre-cohort position is -0.08 percentage points (0.31), and the effect on listing at least two earlier positions is 0.09 points (0.29). None of the eight employment-history checks rejects equality after Holm adjustment for multiple testing (Appendix N.2).

Together, these results provide suggestive evidence that LinkedIn engagement does not differ between treatment and control beyond reporting new jobs and displaying the added certificate. The checks measure recorded profile content; they do not capture visits, messages, or job-search activity and therefore cannot establish that all forms of engagement are unchanged.

4 Who Benefits, and What Does the Credential Convey?

The experiment raises both credential visibility and reported employment. We next examine how the two are connected. A credential's value may depend on how much it adds to the information employers already have about a worker. We first compare employment and sharing responses across learners with different predicted employment prospects. We then measure what a credential adds to an otherwise identical profile and ask how much of the employment difference is reflected in the signals

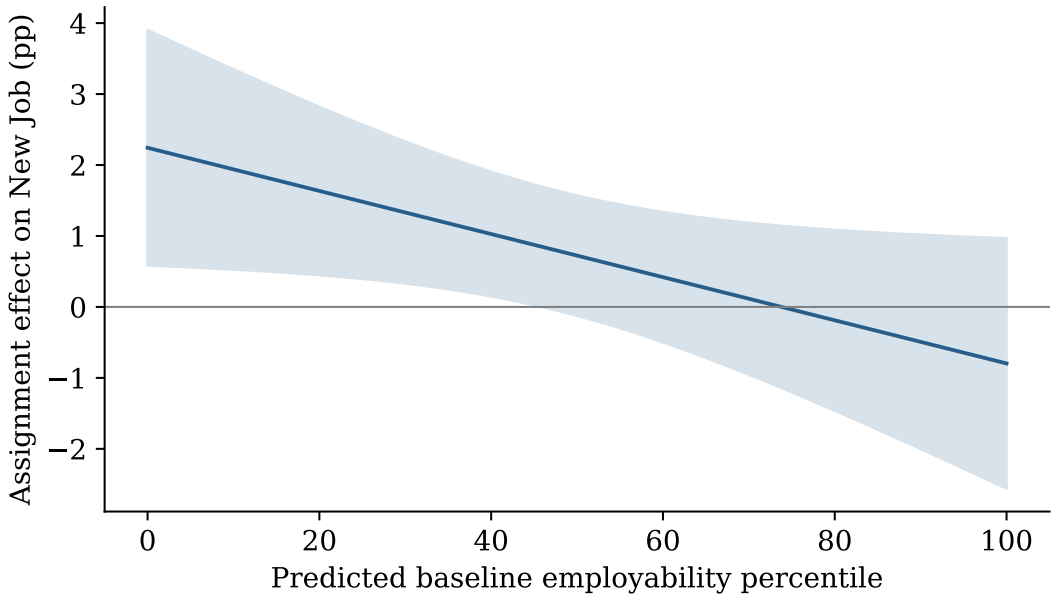
learners actually display.

4.1 Employment Gains by Baseline Employability

Employment gains are larger for learners with lower predicted employability. We rank learners by their predicted probability of employment without encouragement, using pre-treatment education, gender, country group, course characteristics, cohort, and skill measures (Athey et al., 2025). A gradient boosting model trained on 7,495 controls supplies predictions for a separate evaluation sample of 30,027 learners. Each evaluation learner's own employment outcome therefore plays no role in determining their rank. Appendix M describes the model and sample split.

Among learners in the bottom 40 percent of predicted employability, encouragement raises the new-job rate by 2.17 percentage points (standard error 0.66), relative to a control rate of 14.18 percent. This is an increase of about 15 percent. In the upper 60 percent, the estimated effect is -0.25 points (0.57), relative to a control rate of 18.25 percent. The difference in effects is 2.43 points (0.87). The continuous comparison gives the same pattern: a ten-percentile increase in employability rank is associated with a 0.30-percentage-point reduction in the treatment effect (0.15). Figure 2 plots this gradient. Both comparisons adjust for cohort.

Figure 2: Employment Effects along Predicted Employability



Notes: Fitted assignment effect from a linear probability model of New Job on assignment, employability percentile, their interaction, and cohort indicators. Shading is a 95 percent confidence interval. Sample is 30,027 test set learners.

The larger employment gain could arise because learners with lower predicted employability

respond more to the encouragement, or because sharing a credential has a larger return for them. The first explanation predicts larger increases in sharing and visibility in this group. Table 4 examines this prediction using notification receipt, clicking to share, detected certificate display, and certificate-page views in the same evaluation sample.

The sharing and visibility responses do not show the employment gradient. Clicking to share rises by 10.58 percentage points in the lower-employability group and 11.23 points in the upper group; any certificate-page view rises by 1.92 and 1.83 points, respectively. All four outcomes increase in both groups, and none of the four differences between groups is statistically significant at five percent. The estimates remain imprecise enough to permit moderate differences in take-up. They nevertheless provide no evidence that the group with the larger employment gain responds more strongly on these measured margins.

TABLE 4
TAKE-UP, VISIBILITY, AND EMPLOYMENT BY EMPLOYABILITY

	Bottom 40%	Upper 60%	Difference
Notification receipt	62.15	61.13	1.03
	(0.62)	(0.51)	(0.80)
Click to share	10.58	11.23	−0.65
	(0.40)	(0.33)	(0.52)
Detected display	2.22	3.18	−0.96
	(0.69)	(0.58)	(0.90)
Any certificate view	1.92	1.83	0.09
	(0.90)	(0.74)	(1.16)
New Job	2.17	−0.25	2.43
	(0.66)	(0.57)	(0.87)

Notes: ITT effects in percentage points on the same 30,027 evaluation learners and with the same grouping and cohort adjustment as the employment comparison in Section 4.1. HC1 standard errors in parentheses. Difference is bottom 40 percent minus upper 60 percent.

These results are consistent with a larger return to visible qualifications among learners with weaker employment prospects. They do not measure what employers learn. We next examine how an evaluator responds to the credential by varying its presence within otherwise identical profiles.

4.2 Credentials Improve Profile Perception

Sharing a credential may help a worker obtain a job by changing how the worker's qualifications are assessed once the profile is seen. The field experiment identifies the effect of encouragement on employment, but does not isolate this assessment channel. A comparison across workers cannot by itself separate the information conveyed by a displayed credential from differences in their skills

and experience. We therefore compare two versions of the same resume, varying only the certificate information. A language model, acting as a recruiter, lets us apply common assessment instructions across these pairs and measure how the added information changes assessed fit for a relevant vacancy.[2]

We construct a *baseline* profile for each learner by removing identifiable later additions from the LinkedIn profile collected at the end of the study.[3] An *enhanced* version adds a Coursera credential reflecting the learner's course domain and holds all other information fixed. Comparing the two versions measures the judge's response to the added credential. A third, *observed-profile* version retains the credentials learners actually displayed and supplies the score used in the surrogate analysis below.[4]

For all 45,037 learners, Claude Haiku 4.5 rates candidate–job fit from 0 to 100 against the same junior job description within each course domain, following an LLM-as-a-judge framework (Zheng et al., 2023). Each learner has two independent evaluations of the baseline profile. We use the second evaluation to group profiles and the primary evaluation to calculate the gain: the enhanced score minus the primary baseline score. An unusually low draw used to form a group therefore does not also inflate the gain by entering its calculation. This separates the grouping from the regression to the mean that would arise from reusing the same evaluation. The groups here reflect assessed profile fit, whereas the preceding subsection ranks learners by predicted employment.

Adding a credential raises the score by 15.97 points (standard error 0.09) across all 22,487 controls. The gain is largest where the independent baseline evaluation indicates the weakest fit: 24.88 points (0.16) in the first quartile, compared with 5.19 points (0.09) in the top quartile. The difference is 19.69 points (0.18; Appendix Table A13). Figure 3 shows a declining gain across ten-point bins of the independent baseline score, from 23.51 points for scores below 10 to 0.66 points for scores of 90 or above.

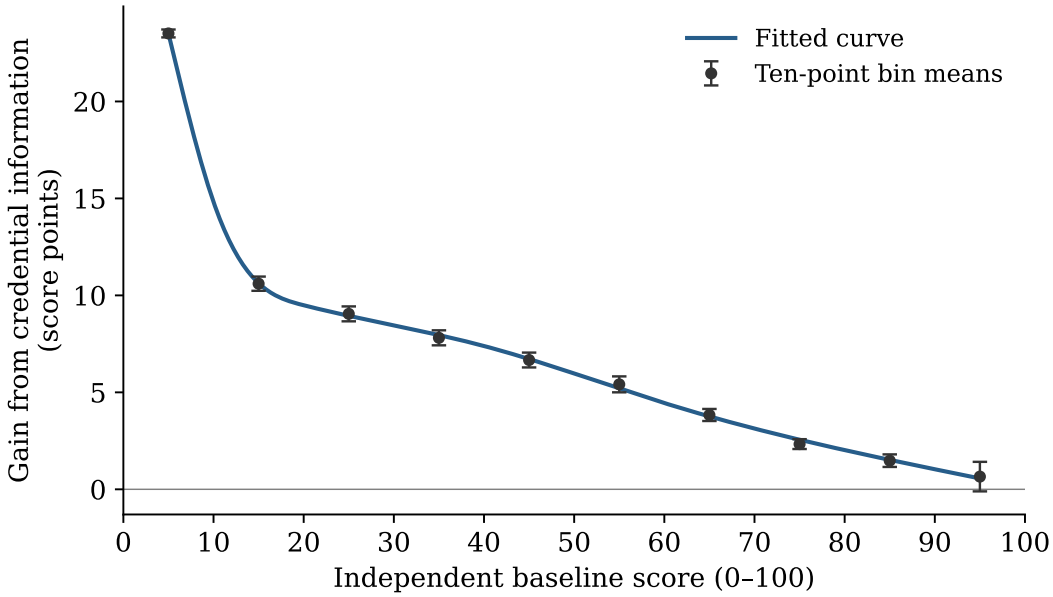
The largest assessment gains occur when the available profile says little else about the learner. Profiles that list no education, experience, certifications, or courses make up 97 percent of the first

[2]This approach is related to Athey et al. (2022), who use generative models to create versions of borrower profile images that differ in one feature at a time and test participants' responses in randomized survey experiments. The shared principle is to vary displayed information while holding the rest of the profile fixed. In our exercise, the variation is the addition of credential information and the evaluator is an LLM, so the outcome is a model assessment rather than an observed employer decision.

[3]Bloom et al. (2026) find that 19.7 percent of established U.S. LinkedIn users in their sample revised a past job's title or description during 2020–2026, with edits concentrated around job transitions. Such revisions can remain after identifiable later additions are removed, so our reconstructed baseline need not reproduce the profile visible before treatment. The paired comparison holds the remaining text fixed but cannot recover its earlier wording.

[4]For the observed-profile score, we exclude jobs with recorded start dates on or after September 1, 2022. Appendix Q details the construction rules.

Figure 3: Additional Credential Information and Profile Assessments



Notes: Claude Haiku 4.5 evaluations of all 22,487 control profiles, including September. The horizontal axis uses the independent baseline evaluation; gains subtract the primary baseline score from the enhanced-profile score. Points show mean gains in ten-point bins of the independent score, plotted at bin midpoints. The bins hold between 29 profiles (scores of 90 to 95, the highest score) and 12,034 profiles (scores below 10). Bars are pointwise learner-level 95 percent intervals conditional on realized scores. The curve is a natural cubic regression spline fitted to the bin means with equal weight per bin and interior knots at scores 10, 20, 40, and 60, shown between the first and last bin midpoints. Intervals exclude uncertainty about the model and prompt.

quartile and 88 percent of the second. The 0–100 scale also leaves less room for improvement at high scores.

As a robustness check, we repeat the profile comparison using Llama 3.1 8B. The findings are robust to this alternative LLM: adding a credential improves assessed fit, with larger gains in the lowest than in the highest quartile of the independent Haiku baseline ranking. On the same 18,326 controls, the mean gain is 8.28 points under Llama and 15.95 under Haiku. Appendix R reports the comparison and correlations between the models' scores.

The paired comparison adds a credential to every profile, whether or not the learner actually shared it. To connect this potential information gain to observed employment, we combine the gain with actual certificate display. Let s_i^0 and s_i^1 denote learner i's primary baseline and enhanced Haiku scores, and let D_i equal one when the earned certificate is detected on LinkedIn. We construct

$$Q_i = D_i\left(s_i^1 - s_i^0\right).$$

For a learner with a ten-point hypothetical gain, Q_i equals ten if the certificate is displayed and zero

otherwise. Negative gains are retained. The measure combines how much the credential could add with whether the learner makes it visible.

In the pooled sample, ten additional points of Q_i are associated with a 9.00-percentage-point higher new-job probability (robust standard error 0.48), controlling for treatment assignment. Separate regressions give similar estimates for treatment learners, 8.80 points (0.66), and controls, 9.23 points (0.69). The association therefore also appears among learners who were not assigned encouragement. Controlling additionally for the primary baseline score reduces the control estimate to 7.20 points (0.66), with a 95 percent confidence interval of 5.90–8.51 points. Appendix Table A15 reports both specifications for the pooled and separate samples.

This association links the assessment exercise to employment, but its coefficient combines differences in sharing with differences in hypothetical gains. Sharing is chosen. The coefficient may therefore reflect job search, skills, or profile updating as well as information conveyed to employers. Holding profile text fixed isolates the judge's response to adding a credential; it does not identify the employment return to that score change.

4.3 How Much of the Employment Effect Do the Observed Signals Capture?

The preceding exercise measures the information a credential could add. We now return to the randomized comparison and ask how much of the employment difference is reflected in what learners actually display and the attention their credentials receive. Following Athey et al. (2026), we combine these observed signals into a *surrogate index*: a prediction of employment based on signals and covariates. Comparing the index across treatment and control translates the differences in observed signals into a predicted employment difference.

The signals are any certificate-page view, detection of the earned credential on LinkedIn, and the Haiku score of the observed profile. This last measure assesses the profile with its displayed credentials after removing jobs with dated starts on or after September 1, 2022. Assignment raises this score by 0.42 points (standard error 0.22). This contrast measures the change in actual profile assessments across experimental groups. The paired gain instead measures what adding a credential to every baseline profile could do.

We estimate the relationship between these signals and employment among controls, then use the same relationship to predict employment in both groups. The index also includes demographic, course, skill, and recorded employment-history covariates. Five-fold cross-fitting ensures that each learner's

prediction uses a model fitted without that learner's outcome. If the observed signals fully capture the employment effect, treatment assignment should carry no further information about employment once we hold the signals and covariates fixed.

This requirement is *conditional-mean surrogacy*. Let Y denote employment, Z treatment assignment, S the signals, and X the covariates. The condition is $E[Y \mid S, X, Z = 1] = E[Y \mid S, X, Z = 0]$. It would fail, for example, if encouragement changed job search in ways the signals do not capture. We also require overlap: signal–covariate combinations observed among treated learners must be represented among controls. Under these conditions and a correctly specified prediction model, the difference in mean index values recovers the employment effect. Random assignment alone does not guarantee these conditions.

Table 5 compares an index using views alone with one adding detected display and the observed-profile score. Both use the same 45,037 learners and covariates. The views-only index predicts a 0.27-percentage-point employment difference, or 28 percent of the unadjusted 0.97-point difference between treatment and control. Adding display and the observed score raises the predicted difference to 0.54 points, or 56 percent.

TABLE 5
EMPLOYMENT CONTRAST CAPTURED BY VIEWS, DISPLAY, AND RESUME ASSESSMENT

Index inputs	Index contrast (pp)	Share of contrast (%)	Residual (pp)
Views only	0.27	28	0.70
			(0.34)
Views, display, and observed score	0.54	56	0.44
			(0.33)

Notes: Each index predicts employment using a control-trained linear regression, cross-fitted over the same five folds, with the listed signals and the same covariates. Views denotes any certificate-page view. Index contrast is its treated-minus-control mean; the share divides by the 0.97-point employment contrast on the same 45,037 learners. Residual is the assignment coefficient in a regression of employment on assignment and the index, with an HC1 standard error conditional on the fitted index. The share is descriptive, not an identified causal mediation fraction.

Because we observe employment, we can also check whether assignment still predicts it after controlling for the fitted index. With views alone, the remaining assignment coefficient is 0.70 percentage points (standard error 0.34; $p = 0.039$). With all three signals it falls to 0.44 points (0.33; $p = 0.182$). Views alone therefore leave a detectable employment difference. Adding display and the profile score reduces that difference, but its imprecision prevents a firm conclusion about whether the three signals fully capture the effect. A zero residual is only a necessary implication under the fitted model: failure to reject it does not establish surrogacy, and rejection can also reflect model misspecification.

The observed signals thus predict part of the experimental employment difference, with more of the difference captured when display and profile assessment accompany views. The comparison does not separate the employment contributions of attention, display, and assessment. Sharing is chosen, and several signals are measured alongside employment, so the 56 percent share is descriptive rather than a causal mediation estimate. The employability and index results also use the main New Job measure and inherit the observation-window and reporting qualifications in Section 3.

The evidence supports the interpretation that making existing qualifications visible can help workers with weaker employment prospects. It leaves room for changes in search behavior and does not trace each step from employer attention to hiring. Individual gains also leave open whether employers fill vacancies sooner. The next section examines this question in the markets where learners seek work.

5 Local Labor Markets and Vacancy Duration

Our credential-sharing treatments could help employers fill job openings that would otherwise remain vacant. However, they could also help participating learners get jobs that other workers would otherwise have obtained. Crépon et al. (2013) distinguish between these possibilities by randomizing treatment intensity across markets. Our experiments do not provide this variation. Instead, we compare outcomes across markets with different exposure to the intervention, adjusting for common shocks and market trends (see Pallais (2014) for a similar approach).

We focus on Experiment 2 and proceed in three steps. First, we define local labor markets as metropolitan area–occupation pairs. Second, using professional profiles and job advertisements from Revelio Labs, we document an increase in the number of workers with Coursera credentials in these markets following our experiment. Third, we compare how quickly employers fill vacancies across markets with different exposure to our intervention. This analysis shifts the focus from individual employment gains to whether greater credential visibility is associated with faster matching.

In the vacancy analysis, we are interested in the effect of higher versus lower experimental exposure on how quickly vacancies fill in a local labor market. Let $Y_{mt}(d)$ denote mean advertisement duration in market m and posting month t under exposure d. For two exposure levels, $d_{\text{high}} > d_{\text{low}}$, define

$$\theta(d_{\text{high}}, d_{\text{low}}) = \mathbb{E}_{m,\, t \in \mathcal{E}} \left[Y_{mt}(d_{\text{high}}) - Y_{mt}(d_{\text{low}}) \right],$$

where $\mathcal{E}$ denotes the experiment period. This comparison measures the effect of moving from lower to higher exposure, giving equal weight to each market and each month of the experiment. The

market-level outcome captures effects on both the experimental subjects and other workers competing for the same vacancies. A value of $\theta(d_{\text{high}}, d_{\text{low}}) = 0$ would be consistent with participants gaining at other workers' expense. A negative value would be consistent with greater credential visibility helping employers fill vacancies sooner than they otherwise would.[5]

5.1 Markets, Outcomes, and Exposure

We focus on the twelve countries with the most experimental participants, together accounting for about 75% of the sample.[6] In each country, we define local labor markets as metropolitan areas crossed with four fields: Business, Computer Science, Data Science, and Information Technology.[7] We retain junior vacancies (Revelio seniority categories 1 and 2) and classify them into these fields based on Revelio's occupation categories. Appendix T lists the occupations included in each field and country-level sample counts.

We retain markets with at least 100 advertisements and ten distinct employers in each comparison period, and at least one advertisement in every month from October 2021 through March 2023. These restrictions focus the analysis on larger markets. The resulting sample contains 779 markets in 217 metropolitan areas: 426 markets in 107 U.S. metropolitan areas and 353 markets in 110 metropolitan areas in developing countries. Appendix Table A19 lists the metropolitan areas by country. In these markets, we observe 3,667,217 advertisements in the experiment period.

Each vacancy in the Revelio data has a posting date and a closure date. We measure advertisement duration as the number of days between these dates, capping individual durations at 180 days and assigning a duration of zero to advertisements removed on the day they were posted. Our outcome of interest is the average duration of advertisements posted in a given market and month. This measure captures how long advertisements remain active. Closure may reflect hiring, withdrawal, or expiration, so it does not establish that the vacancy was filled. [8]

To measure treatment exposure in each local labor market, we use control-group learners in the LinkedIn-matched sample to estimate how participants are distributed across markets. We locate these learners using the metropolitan area and occupation of their latest pre-experiment job. We then

[5] Throughout the analysis, we use distinct time intervals for the experiment period and the pre-experiment period to establish baselines, conduct placebo test, and check for pre-trends. Appendix Y summarizes all the dates.

[6] These countries are the United States, Brazil, Egypt, India, Indonesia, Nigeria, Pakistan, the Philippines, Saudi Arabia, Singapore, the United Arab Emirates, and Vietnam.

[7] Defining local labor markets jointly by geography and occupation or occupational group follows Azar et al. (2020), Marinescu et al. (2021), and Berger et al. (2024).

[8] There is no direct reason to expect the treatment to induce employers to withdraw unfilled vacancies sooner. Other studies also use the duration of online job advertisements to study employer recruitment, including Bassier et al. (2025) and Davis and Samaniego de la Parra (2024).

assume that the same distribution holds for the full experimental population and allocate participants to markets in these proportions and express exposure per 100 pre-experiment advertisements (using advertisements posted from October 2021 through March 2022, the same calendar months one year earlier, totaling 4,748,253 advertisements across the 779 sample markets). Appendix T provides construction details.

Figure 4 shows the distribution of projected experimental exposure across local labor markets. Of the 779 markets, 368 (47.2 percent) have zero projected exposure. Weighting markets equally, the median is 0.68 expected encouraged learners per 100 pre-experiment advertisements. The distribution nevertheless has a long upper tail: exposure reaches 16.79 at the 75th percentile, 60.15 at the 90th percentile, and 111.78 at the 95th percentile. Panel A shows the full distribution; Panel B uses a logarithmic scale to show variation among the 411 markets with positive exposure.

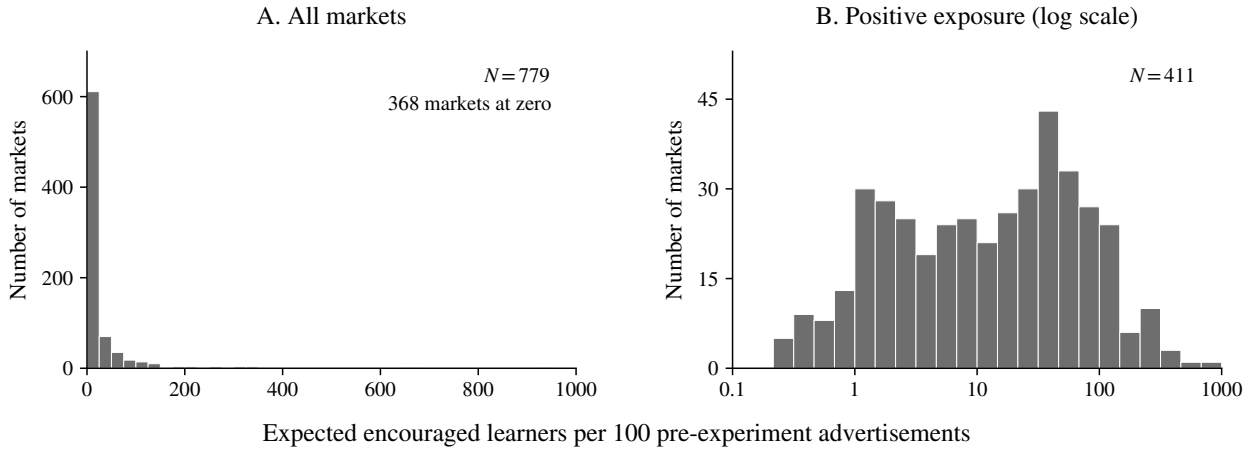


FIGURE 4
PROJECTED ENCOURAGEMENT EXPOSURE ACROSS LOCAL LABOR MARKETS

Notes: Projected exposure across 779 local labor markets. Exposure is the expected number of encouraged learners per 100 advertisements in October 2021–March 2022. Panel A includes zero exposure. Panel B excludes the 368 zero-exposure markets; the axis labels report exposure in levels.

Exposure depends on enrollment concentration, not on randomized differences in market treatment intensity. Prior jobs may imperfectly locate subsequent search, and the linked controls may not represent all learners in their strata. Expected encouragement therefore measures projected exposure, not additional active job seekers.[9]

[9]Online posting coverage also limits the interpretation of exposure. Dalton et al. (2026) link U.S. Burning Glass postings to representative vacancy data and find that small and mid-sized establishments are less likely to post openings online than large establishments. If similar undercoverage reduces Revelio's posting denominator relative to all openings, we overstate expected encouragement. Our observed vacancy outcomes may therefore disproportionately reflect larger establishments recruiting online within each city–field market.

5.2 Credential Visibility in the Broader Market

We now examine whether credentials from the experiment's courses become more prevalent in Revelio dataset following the launch of the experiment. We identify Coursera credentials using issuer names and verification URLs, then match their titles to the experiment's course catalog. Our benchmark is the prevalence we would expect if the pre-experiment trend continued.

We use monthly data from September 2021 through April 2022 to estimate a linear trend in the log ratio of experiment-course Coursera credentials to credentials from other providers. We project this trend forward to June 2022–February 2023 and compare the observed ratios with their predicted values.[10] Appendix U provides construction details.

Table 6 shows the the number of displayed credentials from the experiment's courses, relative to other-provider credentials, are 10.5 percent above the projected trend across the twelve countries. Giving countries equal weight yields an increase of 12.6 percent. Restricting profiles to the 779 markets used in the vacancy analysis yields an increase of 8.4 percent, with increases of 16.9 percent in U.S. markets and 5.0 percent in developing-country markets. Appendix Figure **??** plots monthly counts of experiment-course Coursera and non-Coursera certificates alongside their fitted pre-experiment trends and Appendix U compares U.S. workers without college degrees, the group targeted by the experiment, with college graduates, who fall outside its U.S. eligibility criteria.

TABLE 6
GROWTH IN DISPLAYED CREDENTIALS FROM EXPERIMENT COURSES

	Trend-adjusted growth (%)	
Profile sample	Pooled	Equal country weights
All profiles in twelve countries	10.50	12.59
All study markets (779)	8.38	9.79
U.S. study markets (426)	16.93	16.93
Developing-country study markets (353)	4.96	9.17

Notes: Experiment-course Coursera credentials relative to credentials from other providers. The first row uses all profiles in the twelve countries covered by the vacancy data. Market rows restrict profiles to the markets in Table 7. Pooled estimates assing the same weight to each credentials. Equal-country estimates re-weigh by country.

[10] Dates refer to when credentials were granted, not when they were added to profiles. We omit January in both periods because reported grant dates bunch in that month. We include June and July 2022 in the comparison period because some credentials granted in those months were included in the experiment. Matching course titles identifies credentials from the experiment's courses, but does not establish that their holders participated in the experiment.

5.3 Comparison and Identifying Conditions

We estimate

$$\begin{aligned} Y_{mt} = \alpha_m + \lambda_{ct} + \mu_{kft} + \gamma_{ft}\ell_m \\ + b_m\tau_t + q_m\tau_t + \beta d_m P_t + \varepsilon_{mt}, \end{aligned} \tag{1}$$

where Y_{mt} is mean advertisement duration in market m and posting month t. Each market belongs to a metropolitan area c, country k, and job type f. The variable d_m measures exposure to the experimental encouragement, and P_t indicates the experiment period, October 2022–March 2023. The comparison period is October 2021–March 2022. Appendix V reports April–July 2022 placebo estimates and separate tests of equality with the experiment coefficients; the pooled equal-weight quadratic placebo coefficient is imprecise and not statistically distinguishable from zero.

The fixed effects capture persistent differences across markets (α_m), monthly shocks shared across fields within a metropolitan area (λ_{ct}), and monthly shocks shared across metropolitan areas within a country's job type (μ_{kft}). We also adjust for market size measured by the number of pre-experiment advertisements (ℓ_m). Our baseline specification includes market-specific linear trends in calendar time, τ_t, and weights markets equally. We also report specifications with quadratic trends, with market weights proportional to the number of pre-experiment advertisements, and with employer fixed effects.

The coefficient β measures the additional change in advertisement duration associated with a one-standard-deviation increase in exposure, after accounting for the controls and trends. A negative coefficient indicates shorter durations in more-exposed markets.

A causal interpretation requires that, absent treatment, more- and less-exposed markets would have followed the paths captured by the controls and market-specific trends. Other changes affecting vacancy duration must therefore be unrelated to exposure after these adjustments. We also assume that encouragement has no effect before the intervention and that each market's exposure captures any relevant spillovers from other markets, while allowing spillovers among workers within a market. Interpreting β as the average causal effect additionally requires an effect that is linear in exposure.

These assumptions would fail, for example, if more-exposed markets experienced hiring booms that shortened vacancy durations beyond what the controls and trends predict. Workers might also acquire credentials in anticipation of faster hiring. Appendix W reports a small, imprecise pre-experiment association between vacancy closure and displayed experiment-course credentials; it does not rule out this concern. Changes in Revelio's coverage or employers' withdrawal practices that are systematically

related to exposure could also confound the comparison. Appendix X tests conditional pre-period deviations over longer histories. The pooled joint tests do not reject, but imprecision and departures in the U.S. series limit their interpretation.

5.4 Exposure and Vacancy Duration Across Country Groups

Table 7 reports estimates of equation (1) with equal market or posting weights, linear or quadratic trends, and with and without employer fixed effects. Panel A pools all twelve countries, Panel B includes only U.S. markets, and Panel C includes only developing-country markets. In the pooled sample, greater exposure is associated with shorter vacancy duration in all specifications. The baseline specification, which weights markets equally includes linear trends, and no employer fixed effects, gives a 1.44-day reduction per standard deviation of exposure (standard error 0.46).[11] Weighting markets by pre-experiment posting counts gives a 2.31-day reduction (0.92).

TABLE 7
ENCOURAGEMENT EXPOSURE AND ADVERTISEMENT DURATION

	Equal market weights				Posting weights			
	Linear trend		Quadratic trend		Linear trend		Quadratic trend	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. All countries (779 markets)								
	−1.444***	−1.385**	−1.444***	−1.369**	−2.306**	−1.351**	−2.306**	−1.345**
	(0.437)	(0.477)	(0.461)	(0.497)	(0.874)	(0.447)	(0.922)	(0.450)
Panel B. United States (426 markets)								
	0.060	−0.173	0.060	−0.168	−0.166	−0.183	−0.166	−0.182
	(0.241)	(0.127)	(0.254)	(0.125)	(0.340)	(0.189)	(0.358)	(0.189)
Panel C. Developing countries (353 markets)								
	−2.604**	−2.295**	−2.604**	−2.283**	−3.520**	−1.912**	−3.520**	−1.900**
	(0.827)	(0.927)	(0.872)	(0.977)	(1.292)	(0.695)	(1.362)	(0.698)
Employer FE	No	Yes	No	Yes	No	Yes	No	Yes

Notes: Exposure coefficients for advertisement duration in days. The sample covers October–March of 2021–22 and 2022–23. Exposure is projected encouraged learners per 100 pre-experiment advertisements, standardized within each panel using posting-weighted standard deviations. Panels A–B use the twelve-country destination map; Panel C restricts origins and destinations to eleven non-OECD countries. All specifications include market, city-by-month, and country–field-by-month effects, field-specific log baseline posting count interacted with month, and the indicated market-specific trends. Posting weights use pre-experiment counts. Column (3) is the main specification. Parentheses report standard errors clustered by country in Panels A and C (12 and 11 clusters) and metropolitan area in Panel B (107 clusters). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel B shows little association between exposure and vacancy duration in the United States. The

[11] In the pooled sample, the posting-weighted standard deviation is 20.48 projected encouraged learners per 100 pre-experiment advertisements, approximately one additional encouraged learner per five advertisements. The 1.44-day estimate therefore corresponds to an exposure difference of this size. Across the 779 markets, weighting markets equally, median exposure is 0.68 and the 90th percentile is 60.15. Adding one standard deviation to the median gives 21.16, approximately the 77th percentile.

main estimate in column (3) is 0.06 days (standard error 0.25), and the corresponding posting-weighted estimate is −0.17 days (0.36). Adding employer fixed effects yields estimates of −0.17 days (0.13) and −0.18 days (0.19), respectively. In developing countries, all eight coefficients are negative and statistically significant at the 5 percent level. Panel C's main estimate implies a 2.60-day reduction (0.87), increasing to 3.52 days (1.36) with posting weights. Employer fixed effects reduce these magnitudes to 2.28 days (0.98) and 1.90 days (0.70), respectively.

Although employer fixed effects improve precision, the U.S. estimates remain consistent with moderate reductions or increases in vacancy duration for a common change in exposure. Furthermore, we cannot reject equality of the U.S. and developing-country responses in any of the eight specifications; two-sided p-values range from 0.47 to 0.90.

The U.S. results are consistent with the modest average effects of job-search assistance in high-income countries discussed by Carranza and McKenzie (2024), though their imprecision limits this interpretation. The same review highlights more promising recent evidence on search and signaling interventions in developing countries. Our developing-country results are consistent with this evidence and with the view that information frictions limit matching in these settings. Bandiera et al. (2022) emphasize the potential for credible skill signals to improve matching in African labor markets, while Carranza et al. (2022), Bassi and Nansamba (2022), and Abebe et al. (2021) document gains from skill certification and communication.

Under the market-trends assumption, the pooled and developing-country results are consistent with faster matching. They extend the evidence beyond treated learners' profiles, but vacancy durations cannot establish whether the private gains translate into higher total employment.

These market-level comparisons are observational. Their causal interpretation depends on the controls and market-specific trends accounting for exposure-related differences in how vacancy duration would otherwise have evolved, and on the assumed common linear response to exposure being correctly specified.

Two Placement Benchmarks. We provide two benchmarks to help interpret what these results suggest in terms of job search outcomes for the treated populations. Benchmark 1 assumes that all experimental participants projected to a market find a job on the first day of search, up to the number of available vacancies. Benchmark 2 assumes 94.6 days of job search on average (the value for the control group in the LinkedIn matched sample). We use a normal distribution centered at this value, with a 30-day

standard deviation and negative times excluded. In both benchmarks, we draw vacancies at random and replace their duration. These scenarios hold posting volumes fixed.

Figure 5 illustrates four outcome curves: the observed outcomes, fitted zero-exposure comparison, and both benchmarks in four markets varying with the exposure. Blue curves show observed experiment-period outcomes; orange dashed curves show fitted distributions with exposure set to zero. Green and purple curves show Benchmarks 1 and 2. A higher curve means that a larger share of advertisements has closed by that day. We can notice that in markets such as Belo Horizonte (top left) or Philadelphia (top right), even the most optimistic job search assumption for the experimental subjects does not change the market outcomes, because of low exposure. In contrast, in higher-exposure markets such as Singapore or Karachi, our benchmarks give a substantial change in the average outcomes; much larger than the estimated effects.

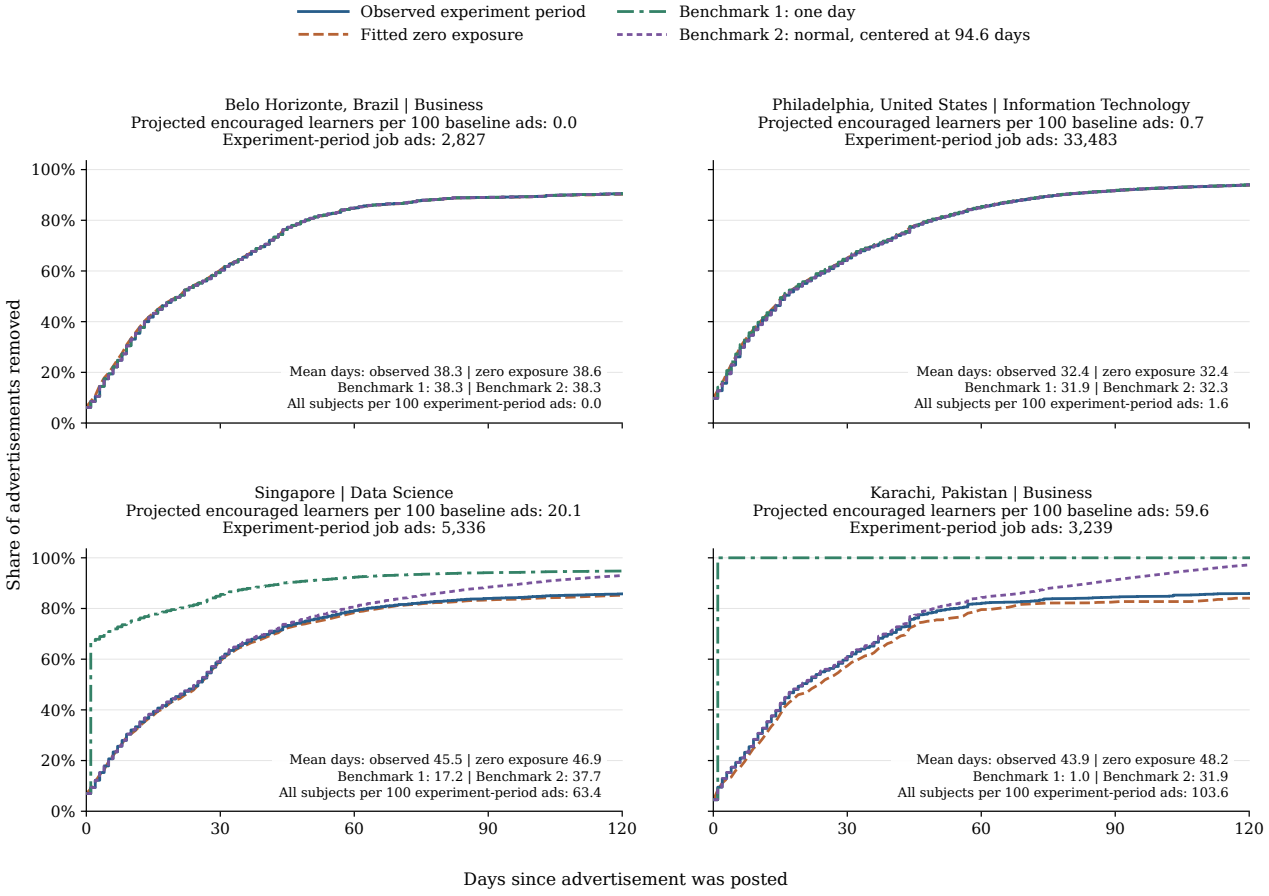


Figure 5: Vacancy-Removal Distributions and Two Placement Benchmarks

Notes: Exposure is projected encouraged learners per 100 baseline job advertisements (October 2021–March 2022). Blue curves use October 2022–March 2023 postings; orange dashed curves set exposure to zero in threshold regressions using Table 7's baseline specification. Green and purple curves apply Benchmarks 1 and 2.

Across all 779 markets, Benchmark 1 reduces mean vacancy duration from 37.5 to 27.9 days, a decline of 9.6 days (25.7 percent). Benchmark 2 reduces it to 35.2 days, a decline of 2.3 days (6.0 percent). The distribution includes 368 markets with zero projected subjects, for which neither benchmark changes duration. In 90 markets, subjects outnumber postings and excess subjects remain unallocated.

Figure 6 shows how these scenarios shift the distribution of mean vacancy duration across markets. It overlays densities under the fitted no-intervention counterfactual, observed outcomes, and the two benchmarks, with corresponding averages of 39.0, 37.5, 27.9, and 35.2 days. Months and markets are weighted equally. The four-city figure pools postings within each market; the density figure instead summarizes the distribution of market means across the full sample.

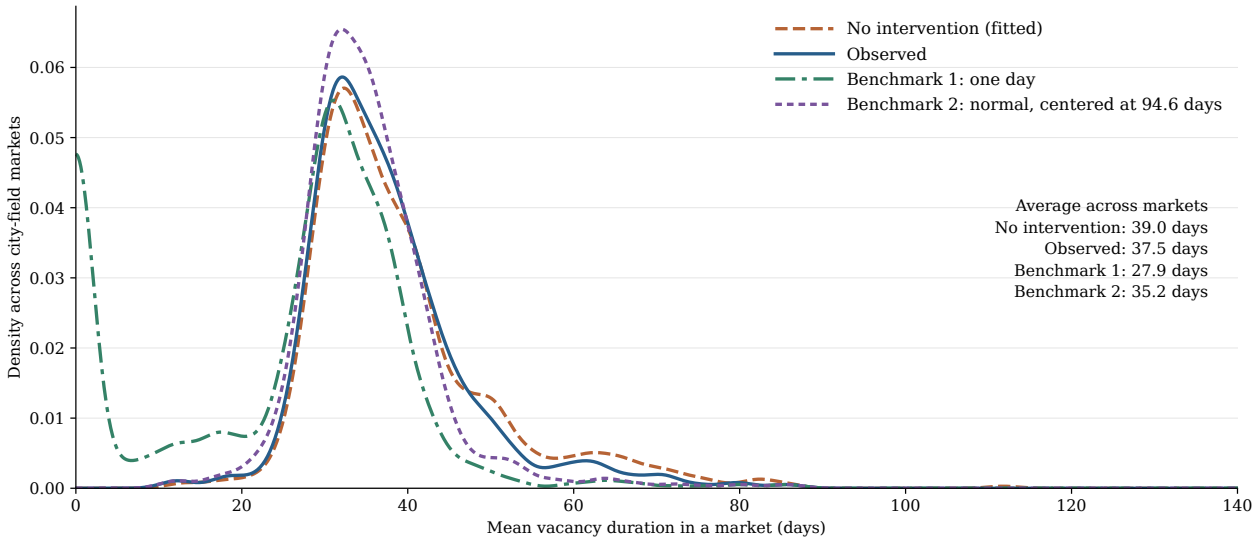


Figure 6: Market Vacancy Duration Under Four Scenarios

Notes: The four curves show kernel densities of mean vacancy duration across all 779 city–field markets. Within each market, durations are capped at 180 days and averaged equally across the six experiment-period months; markets receive equal weight. The orange curve uses fitted predictions from Table 7, Panel A, column (2), with encouragement exposure set to zero, retaining the fitted effects, size interactions, and trends. The blue curve uses observed outcomes. Green and purple apply Benchmark 1 (one day) and Benchmark 2 (normal timing centered at 94.6 days, with a 30-day standard deviation and negative times excluded). Benchmark 2 uses expected durations over the timing distribution.

6 Conclusion

Helping workers communicate qualifications they already hold can improve their labor-market opportunities. In two experiments, reminders and simpler sharing make Coursera credentials more visible. In the experiment linked to employment histories, encouragement also increases reported new jobs. Because the intervention takes place after course completion, the result identifies a return to assistance

with disclosure, separate from acquiring the qualification.

The evidence points to information as one channel. Adding a credential improves a model's assessment of an otherwise identical resume, and views, display, and observed resume scores capture part of the employment difference. The reporting checks do not rule out changes in what learners record. Greater market exposure is associated with shorter vacancy duration, consistent with faster matching, though this comparison cannot establish net job creation.

Training providers ordinarily focus on helping learners acquire skills and complete credentials. Our results identify another task: helping certificate holders communicate what they have earned. A provider can reduce the effort required to display a verifiable qualification without offering additional instruction or assessment. Whether these private gains translate into additional employment depends on employers' responses and effects on other workers. The findings show that a qualification's economic value depends in part on whether its holder makes it visible.

References

Abadie, A., Angrist, J., Frandsen, B., and Pischke, J.-S. (2025). Harvesting differences-in-differences and event-study evidence. NBER Working Paper 34550, National Bureau of Economic Research.

Abebe, G., Caria, A. S., Fafchamps, M., Falco, P., Franklin, S., and Quinn, S. (2021). Anonymity or distance? job search and labour market exclusion in a growing african city. *The Review of Economic Studies*, 88(3):1279–1310.

Aryal, G., Bhuller, M., and Lange, F. (2022). Signaling and employer learning with instruments. *American Economic Review*, 112(5):1669–1702.

Asirvatham, H., Mokski, E., and Shleifer, A. (2026). GPT as a measurement tool. Working Paper 34834, National Bureau of Economic Research.

Athey, S., Chetty, R., Imbens, G. W., and Kang, H. (2026). The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. *Review of Economic Studies*, 93(4):2284–2312.

Athey, S., Karlan, D., Palikot, E., and Yuan, Y. (2022). Smiles in profiles: Improving fairness and efficiency using estimates of user preferences in online marketplaces. Working Paper 30633, National Bureau of Economic Research.

Athey, S., Keleher, N., and Spiess, J. (2025). Machine learning who to nudge: causal vs predictive targeting in a field experiment on student financial aid renewal. *Journal of Econometrics*, 249:105945.

Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2020). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, 66:101886.

Bandiera, O., Elsayed, A., Smurra, A., and Zipfel, C. (2022). Young adults and labor markets in africa. *Journal of Economic Perspectives*, 36(1):81–100.

Bassi, V. and Nansamba, A. (2022). Screening and signalling non-cognitive skills: experimental evidence from uganda. *The Economic Journal*, 132(642):471–511.

Bassier, I., Manning, A., and Petrongolo, B. (2025). Vacancy duration and wages. *The Review of Economics and Statistics*, pages 1–28.

Berger, D., Herkenhoff, K., Kostøl, A. R., and Mongey, S. (2024). An anatomy of monopsony: Search frictions, amenities, and bargaining in concentrated markets. *NBER Macroeconomics Annual*, 38(1):1–47.

Berry, A., Maloney, M., and Neumark, D. (2026). The missing link? Using LinkedIn data to measure race, ethnic, and gender differences in employment outcomes at individual companies. In Akee, R., Katz, L. F., and Loewenstein, M. A., editors, *Race, Ethnicity, and Economic Statistics for the 21st Century*, chapter 8, pages 273–323. University of Chicago Press.

Bloom, H. S. (1984). Accounting for no-shows in experimental evaluation designs. *Evaluation Review*, 8(2):225–246.

Bloom, N., Moore, G., Simon, L. K., and Wilkie-Rogers, C. (2026). Time travel on professional profiles. Working Paper 35546, National Bureau of Economic Research.

Card, D. (1999). The causal effect of education on earnings. *Handbook of labor economics*, 3:1801–1863.

Carranza, E., Garlick, R., Orkin, K., and Rankin, N. (2022). Job search and hiring with limited information about workseekers' skills. *American Economic Review*, 112(11):3547–3583.

Carranza, E. and McKenzie, D. (2024). Job training and job search assistance policies in developing countries. *Journal of Economic Perspectives*, 38(1):221–244.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *Quarterly Journal of Economics*, 128(2):531–580.

Dalton, M., Kahn, L. B., and Mueller, A. I. (2026). Do online job postings capture job vacancies? Working Paper 35697, National Bureau of Economic Research.

Davis, S. J. and Samaniego de la Parra, B. (2024). Application flows. Working Paper 32320, National Bureau of Economic Research. Revised December 2025.

Deming, D. J., Yuchtman, N., Abulafi, A., Goldin, C., and Katz, L. F. (2016). The value of postsecondary credentials in the labor market: An experimental study. *American Economic Review*, 106(3):778–806.

Ge, C., Huang, K.-W., and Png, I. P. L. (2016). Engineer/scientist careers: Patents, online profiles, and misclassification bias. *Strategic Management Journal*, 37(1):232–253.

Goldin, C. and Katz, L. F. (2018). The race between education and technology. In *Inequality in the 21st Century*, pages 49–54. Routledge.

Guillory, J. and Hancock, J. T. (2012). The effect of LinkedIn on deception in resumes. *Cyberpsychology, Behavior, and Social Networking*, 15(3):135–140.

Hadavand, A., Gooding, I., and Leek, J. T. (2018). Can MOOC programs improve student employment prospects? SSRN Working Paper 3260695.

Hansen, J. D. and Reich, J. (2015). Democratizing education? examining access and usage patterns in massive open online courses. *Science*, 350(6265):1245–1248.

Imbens, G. W. and Angrist, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.

Lange, F. and Topel, R. (2006). The social value of education and human capital. In Hanushek, E. A. and Welch, F., editors, *Handbook of the Economics of Education*, volume 1, chapter 8, pages 459–509. Elsevier, Amsterdam.

Laryea, K., Paepcke, A., Mirzaei, K., and Stevens, M. L. (2021). Ambiguous credentials: How learners use and make sense of massively open online courses. *The Journal of Higher Education*, 92(4):596–622.

Lennon, C. (2021). How do online degrees affect labor market prospects? evidence from a correspondence audit study. *ILR Review*, 74(4):920–947.

Ludwig, J., Mullainathan, S., and Rambachan, A. (2026). Large language models: An applied econometric framework. *Annual Review of Economics*, 18:283–316.

Majerowicz, S. and Zárate, R. A. (2024). Massive open online courses and labor market outcomes: Evidence from colombia. Working paper, Universidad de los Andes and University of Toronto. August 20, 2024.

Marinescu, I., Ouss, I., and Pape, L.-D. (2021). Wages, hires, and labor market concentration. *Journal of Economic Behavior & Organization*, 184:506–605.

Novella, R., Rosas-Shady, D., and Freund, R. (2024). Is online job training for all? experimental evidence on the effects of a coursera program in costa rica. *Journal of Development Economics*, 169:103285.

Pallais, A. (2014). Inefficient hiring in entry-level labor markets. *American Economic Review*, 104(11):3565–3599.

Rivas, M. J., Baker, R. B., and Evans, B. J. (2020). Do moocs make you more marketable? an experimental analysis of the value of moocs relative to traditional credentials and experience. *AERA Open*, 6(4):2332858420973577.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374.

Tyler, J. H., Murnane, R. J., and Willett, J. B. (2000). Estimating the labor market signaling value of the ged. *The Quarterly Journal of Economics*, 115(2):431–468.

Wheeler, L., Garlick, R., Johnson, E., Shaw, P., and Gargano, M. (2022). LinkedIn(to) job opportunities: Experimental evidence from job readiness training. *American Economic Journal: Applied Economics*, 14(2):101–125.

Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., and Stoica, I. (2023). Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, Datasets and Benchmarks Track*.

Zhenghao, C., Alcorn, B., Christensen, G., Eriksson, N., Koller, D., and Emanuel, E. J. (2015). Who's benefiting from moocs, and why. *Harvard Business Review*.

Appendix

A Pre-Registration and Analysis Plan

Experiment 2 was registered at the AEA RCT Registry on August 3, 2022, before the first randomization batch, as AEARCTR-0009438, "The value of online credentials in the labor market." No formal analysis plan document was attached to the registration, so no estimators or covariate sets were pre-specified.

What was pre-registered. The registration describes the intervention as an in-app message encouraging learners to add their credential to LinkedIn together with a simplified sharing process; the population as eligible learners who complete a course and obtain a credential, randomized within batches, stratified on observed characteristics, at the individual level; and a planned sample of 640,000 learners (320,000 treatment, 320,000 control). Primary outcomes are whether the learner shared the credential and whether the learner got a new job, promotion, internship, or pay increase, measured through surveys and LinkedIn data collection. Secondary outcomes cover usage of the credential: the number of clicks on the credential, the source of the clicks, the number of activity logs, and the section of the profile in which the credential was shared.

Registered outcomes and realized coverage. The intent-to-treat effects of assignment on credential sharing and on new employment (Table 3, ITT column) estimate treatment effects on pre-registered primary outcomes, measured through LinkedIn data collection as anticipated in the registration; the certificate-view analyses of Section 3.1 correspond to the pre-registered secondary outcomes on credential clicks and their sources. The realized population departs from the plan in scale: recruitment ran for seven monthly batches rather than the anticipated four months, yielding 880,663 learners satisfying strict survey-verifiable eligibility rather than the planned 640,000, and the employment analysis is conducted in the LinkedIn Matched Sample of 45,037 unique learners (Appendix D). The *New Job* outcome includes changes of position within the same employer. Appendix K separates these moves from other transitions in an exploratory decomposition; the position histories do not establish whether a same-employer move is a promotion. Internships and pay increases are not analyzed separately.

Analyses beyond the registration. The registration did not specify estimators or covariate sets. The linear models, receipt and click adjustments, occupational refinement, employability grouping, scoring

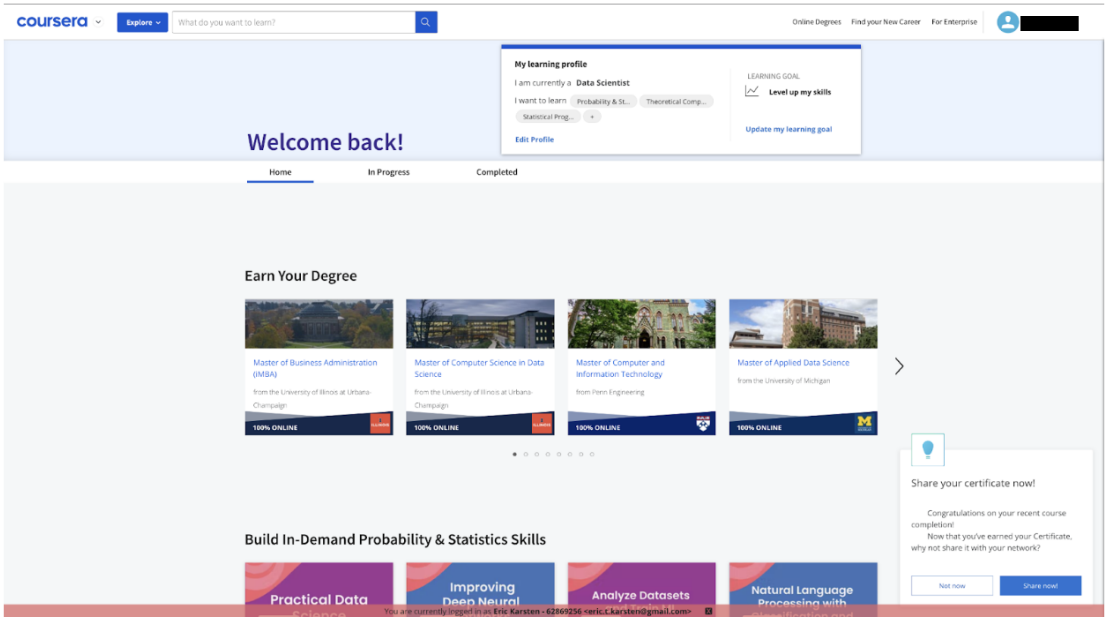
exercise, surrogate indices, and market extension were developed subsequently. The absence of a pre-analysis plan does not change the randomized comparison, but the registration should not be read as prespecifying these analytical choices. Experiment 1 was a prior platform experiment and was not registered under this record.

B Notification Content and Interface

The first message invited the learner to add the certificate to LinkedIn in two clicks and highlighted the fraction completing the course. A second message, for learners who did not click, emphasized credential visibility and described itself as the final reminder. The following wording reproduces the intervention; its claim about profile views is promotional treatment content, not an estimate from this study.

> Do you want to boost your career? Only [XYZ]% of learners manage to complete [course name] on Coursera and get a certificate. Let everyone know you did it! Add the certificate to your LinkedIn profile in just two clicks.
>
> Looking to boost your career? LinkedIn profiles with credentials receive 6x more views! Don't waste your hard-earned certificate! Add the certificate to your LinkedIn profile in just two clicks. PS. This is your last reminder.

Figure A1: Sharing Notification in the Coursera Application



Notes: The experimental notification appears at the bottom right. Controls retain ordinary certificate-sharing options.

C Details of the Experimental Sample Construction

C.1 List of Developed Countries

We classify countries by OECD membership: all learners with an account registered in any of the following OECD member countries are labeled as 'developed country': Austria, Australia, Belgium, Canada, Chile, Colombia, Costa Rica, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Israel, Italy, Japan, Korea, Republic of, Latvia, Lithuania, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Slovak Republic, Slovenia, Spain, Sweden, Switzerland, Turkey, United Kingdom, United States.

We classify learners outside the listed OECD countries as developing. This is a coding choice rather than an income classification: OECD membership includes several middle-income economies and excludes some high-income nonmembers.

C.2 Sample of Courses Included in the Experiment

Table A1 shows a sample of the 7,355 courses from which the learners included in the experiment graduated.

TABLE A1
SAMPLE OF COURSES INCLUDED IN THE EXPERIMENT

Credential type	Certificate name
Information Technology	
Course	Fundamentos do Suporte Técnico
Course	Fundamentals of Red Hat Enterprise Linux
Course	Introduction to Cloud Computing
Course	AWS Cloud Practitioner Essentials
Course	Crash Course on Python
Course	Technical Support Fundamentals
Course	Introduction to Web Development with HTML, CSS, JavaScript
Course	AWS Cloud Technical Essentials
Course	Explore Core Data Concepts in Microsoft Azure
Business	
Guided Project	Develop a Company Website with Wix
Course	Capital-investissement et capital-risque
Course	Teamwork Skills: Communicating Effectively in Groups
Course	Excel Skills for Business: Advanced
Course	Foundations of Project Management
Course	Assess for Success: Marketing Analytics and Measurement
Course	Bookkeeping Basics
Guided Project	Create a Project Charter with Google Docs
Course	Foundations of Digital Marketing and E-commerce
Course	Developing Innovative Ideas for Product Leaders
Guided Project	Introduction to Microsoft Excel
Course	Construction Project Management
Course	Esports Teams and Professional Players
Guided Project	Designing and Formatting a Presentation in PowerPoint
Data Science	
Course	SQL for Data Science
Course	Neural Networks and Deep Learning
Specialization	Introduction to Data Science
Guided Project	Introduction to Business Analysis Using Spreadsheets: Basics
Professional Certificate	Google Data Analytics
Professional Certificate	IBM Data Analyst
Course	Foundations: Data, Data, Everywhere
Course	Introduction to Genomic Technologies
Computer Science	
Course	Introduction to Front-End Development
Course	Python Programming: A Concise Introduction
Guided Project	Get Started with Figma
Course	Foundations of User Experience (UX) Design
Course	Programming for Everybody (Getting Started with Python)
Course	JavaScript Basics
Course	Blockchain: Foundations and Use Cases
Course	HTML, CSS, and Javascript for Web Developers
Course	Responsive Website Basics: Code with HTML, CSS, and JavaScript
Guided Project	Build a mobile app with Google Sheets on Glide and no coding

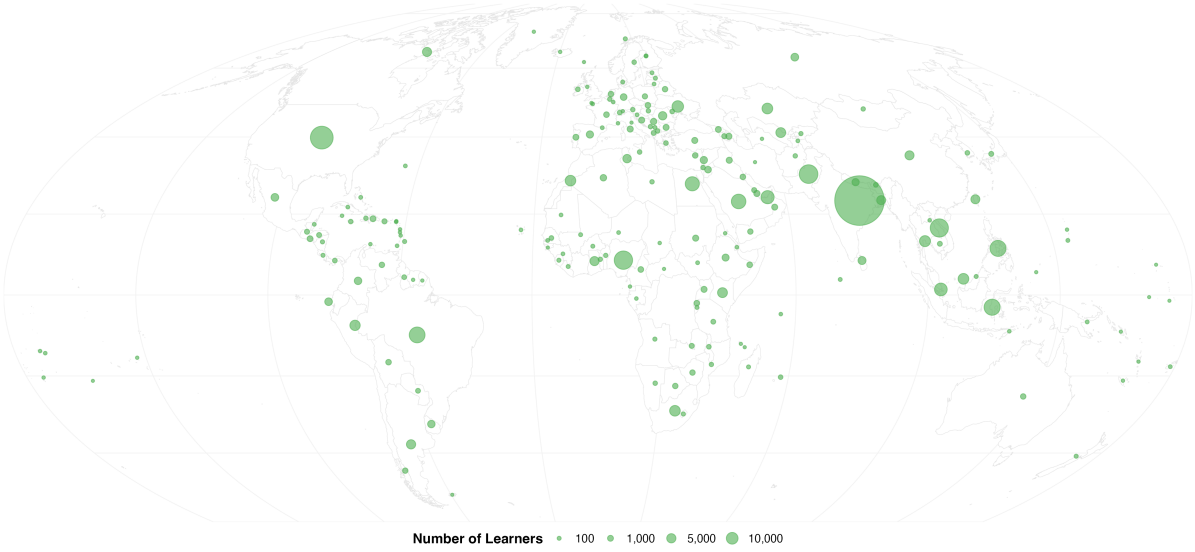
Notes: The table lists a sample of the courses from which learners included in the experiment graduated, with each course's primary domain and credential type.

C.3 Geographical Distribution of Experimental Subjects

Figure A2 displays the global distribution of learners who participated in the experiment. Each circle represents the number of participants from a given country, with the size of the circle proportional

to the number of observations. The sample spans over one hundred countries, with particularly high concentrations of learners from India, the United States, Brazil, and several European countries. Participation from African, East Asian, and Latin American countries is also substantial, illustrating the broad international reach of the experiment.

Figure A2: Geographical Distribution of Experimental Subjects



Notes: The figure shows the geographical distribution of experimental subjects. The size of each circle corresponds to the number of subjects in a country.

C.4 LinkedIn Feature Engineering

- **Current Enrollment in Educational Program**: This binary feature is set to *TRUE* if the end date of the participant's educational program is later than 2022. Only the year of the start and end dates of the educational programs is available in our dataset.
- **Level of Education**: This categorical feature classifies the participant's highest level of education based on keywords found in the title of their degree. The classifications are as follows:
  - *Master's Degree*: Identified through keywords such as 'master', 'msc', 'maestría', or 'ma'.
  - *Bachelor's Degree*: Identified through keywords like 'bsc', 'bs', 'bachillerato', or 'bachelor'.
  - *Doctor's Degree*:Identified through keywords such as 'doctor', 'doutorado', or 'docteur'.
  - *Degree*: This is marked as *TRUE* if any of the above conditions are satisfied.

- **University Ranking of the Latest Academic Degree**: Using the national rankings provided by a public Kaggle dataset (https://www.kaggle.com/datasets/mylesoneill/world-university-rankings), we assigned rankings based on the latest available data, predominantly from 2015. Rankings may vary annually, and institutions may hold different positions in different years.

- **Years Since Latest Academic Degree**: This feature calculates the difference in years between 2023 and the year of the participant's most recent academic degree. If the result is negative, indicating that the participant has not yet graduated, the value is set to 0.

Recorded employment outcomes. The employment analysis uses the new-position indicator in the linked analysis file. New Job in Scope additionally requires a new employer and a job title in a field related to the certificates. Same-employer moves need not be promotions. Appendix E documents the date rules. The decomposition of Appendix K preserves the original new-position outcome and classifies its events using the position histories.

D Data Construction and Matching

Coursera records supply randomized assignment, the delivery funnel, certificate identifiers and grant dates, registration-survey characteristics, skill scores, and certificate-page views. We append the September pilot export to the later-cohort export and link learners by Coursera's hashed identifier. Skill records also match the certificate-grant date. LinkedIn histories supply reported employment and profile fields; verification-code matches identify display of the experimental credential. Every analysis uses one record per learner, and assignment agrees across the linked sources.

Coverage and collection. The September–March administrative sample in the two relevant arms contains 899,715 learners. Of these, 880,663 satisfy strict survey-verifiable eligibility: a known non-OECD registration country or education below an associate degree. The platform boundary is broader, including OECD learners with an associate degree. Missing view counts are not treated as zero.

The visibility sample combines 765,351 learners from the later export with 127,578 September learners with observed counts, for 892,929 unique learners. The earlier September file contains 127,968 learners in the relevant arms, of whom 390 lack counts. September's cumulative counts come from April 2023; the later export is from October 2023. The resulting outcomes measure accumulated visibility over each cohort's available window. Cohort adjustment does not turn them into equal-duration outcomes.

The retained profile extract contains 45,098 valid identifiers; 45,052 link to the administrative sample. Fifteen linked profiles lack a usable employment-analysis record: fourteen lack the required employment/covariate record and one has an analysis row without the employment outcome. The final employment sample is 45,037, comprising 36,681 later-cohort learners and 8,356 from September. September profiles were collected in May–June 2023; later cohorts in September 2023. Exact learner-level acquisition dates are unavailable.

Coding. Missing administrative categories remain explicit. Professional experience is converted to years; negative and missing values retain the source convention of zero. Job-history controls are retrospective and may reflect later reporting. The design-only employment specification excludes them. The broad occupational outcome uses the fixed title lookup described below. Detected display searches profile text for the experimental certificate's verification code; incomplete extraction can generate false negatives.

TABLE A2
EXPERIMENT 2 ANALYSIS SAMPLES

Sample	Control	Treated	Total
Visibility, including pilot	447,760	445,169	892,929
Employment, including pilot	22,487	22,550	45,037
September employment cohort	4,161	4,195	8,356
Employment without pilot	18,326	18,355	36,681

Notes: Unique learners. Different collection windows and eligibility boundaries distinguish the view and employment samples. All primary Experiment 2 results include September; the final row is used only for the pilot-exclusion comparison.

E Job-Title Classification and Employment-Outcome Checks

Classification. New Job in Scope requires New Job, a change in recorded employer, and a title in an occupational field related to the certificates. The fixed lookup combines the existing crosswalk with a review of 3,404 previously unmatched normalized titles using GPT-6 Astra. Review uses original title wording without treatment assignment. The scope covers computing, engineering, professional business and management, and the specified design and digital-marketing occupations. Lookup normalizes capitalization and whitespace.

We apply this same frozen lookup to September. Its added profiles contain 586 distinct titles not covered by the lookup, affecting 599 learners; these remain unresolved. Across the complete sample, 1,922 titles are unresolved and 24,674 profiles lack a current title. Unresolved or missing scope is

coded zero while retaining the learner in the denominator. Thus broad occupational scope is measured conservatively; we do not claim that every September title has received a new review.

Timing. The main New Job indicator retains the source date rule, including some starts before or in the cohort month. The relative-date index equals job-month index minus cohort plus one; the source retains nonnegative values. The later-start rows in Table 3 require an index of at least five, corresponding to a job-month index at least four months after the cohort. Cohorts derive from certificate-grant dates rather than verified individual assignment timestamps.

The expanded sample contains 7,771 positive New Job outcomes: 493 at relative-date index zero, 683 at index one, and 6,595 at later indices. The in-scope outcome is nested within New Job at every restriction.

F Sample and Estimator Comparisons

We reconstruct design cells from cohort, credential domain, registration-country group, and degree status. Known OECD, non-OECD, and unknown countries form separate groups. Bachelor's, master's, doctoral, and professional degrees form the college group; lower degrees and missing education form separate groups. Original allocation-stratum identifiers are not retained, so these are controls based on the documented design rather than verified allocation strata.

The additional controls are registration education, gender, credential type, developing-country status, professional experience, and prior technical and business employment. The last three are retrospectively constructed. Survey missingness remains explicit. Skill scores do not enter the main employment specification; their missingness does not cause complete-case deletion. All tables use HC1 inference unless stated otherwise.

ATT and LATE. Let Z_i denote assignment and A_i an app visit after experiment launch. The ATT instruments $D_i^{\text{app}} = Z_i A_i$ with assignment, retaining the full matched sample. Under no effect of assignment on nonlaunchers, this one-sided adjustment identifies the effect of encouragement among treated app-launchers (Bloom, 1984). The LATE instead instruments clicking to share with assignment, under exclusion and monotonicity (Imbens and Angrist, 1994). A recorded first homepage visit, notification receipt, or click establishes app access; access need not establish that every learner saw the message.

The adjusted first stages are 96.26 percentage points for app access, 10.98 for clicking, and 61.57 for

notification receipt. Outcome and first-stage regressions use the same learners and controls; two-stage-least-squares standard errors include first-stage uncertainty. Receipt is a separate, narrower margin than the ATT reported in Table 3.

TABLE A3
EMPLOYMENT EFFECTS ACROSS SPECIFICATIONS

	Unadjusted (1)	Design cells (2)	Additional controls (3)	Strict eligibility (4)
New Job	0.97	0.98	0.82	0.86
	(0.36)	(0.36)	(0.34)	(0.35)
New Job in Scope	0.83	0.83	0.71	0.75
	(0.28)	(0.28)	(0.27)	(0.28)
New Job, four months	0.61	0.62	0.54	0.52
	(0.27)	(0.27)	(0.26)	(0.27)
In Scope, four months	0.51	0.51	0.46	0.45
	(0.21)	(0.21)	(0.20)	(0.21)

Notes: ITT effects in percentage points with HC1 standard errors. First three columns use 45,037 learners; strict eligibility uses 43,372. The last two columns include the additional controls from Table 3. Outcomes and date restrictions are identical across columns.

G Sample Composition and Profile Content

Table A4 compares the 892,929 learners in the certificate-view sample with the 45,037 learners in the linked employment sample; both include September. The linked sample has a larger share of Data Science credentials (30.23 versus 23.51 percent) and a smaller share of guided projects (10.05 versus 16.61 percent). Developing-country registrations remain the majority in both samples, although their share is lower among linked learners (85.62 versus 89.84 percent).

The largest differences concern survey reporting. Education is reported for 83.26 percent of linked learners, compared with 35.82 percent of the view sample; gender is reported for 89.69 versus 46.95 percent. These gaps indicate that the linked sample contains learners with more complete recorded survey information. They do not establish differences in underlying education or gender composition, since reporting itself differs across samples.

Linked learners also engage more often with the sharing intervention. Among those assigned to treatment, 96.25 percent launch the app, 61.58 percent have recorded notification receipt, and 10.98 percent click to share. The corresponding rates in the view sample are 89.32, 47.53, and 4.05 percent. These post-assignment measures describe the intervention's reach and take-up in each sample. Together, the comparisons show why the employment estimates should be interpreted for learners with linkable

profiles: this group differs in credential mix, survey reporting, and responsiveness to encouragement. Extending its employment effects to the broader population would require assumptions about how treatment effects vary with sample selection.

TABLE A4
SELECTED SAMPLE CHARACTERISTICS

Characteristic (%)	View sample	LinkedIn sample
Data Science credential	23.51	30.23
Guided Project credential	16.61	10.05
Developing-country registration	89.84	85.62
Education reported	35.82	83.26
Gender reported	46.95	89.69
Treated: app launch	89.32	96.25
Treated: notification receipt	47.53	61.58
Treated: click to share	4.05	10.98
Learners	892,929	45,037

Notes: All samples include September. App, receipt, and click rates condition on treatment assignment; other shares use both arms.

G.1 Profile Content and Reporting

We summarize profile completeness by counting characters in recorded education and pre-cohort employment fields, excluding certificate and course fields, identifiers, file metadata, and derived classification flags. We remove entire employment entries whose start dates fall in or after the learner's treatment-cohort month, as well as the latest entry whenever it defines the new-job outcome. We also exclude entries with missing or unparsed start dates and year-only dates that cannot establish that the position began before the cohort month. This check includes September and retains all 45,037 unique learners in the expanded matched sample (22,487 control and 22,550 treated), with empty fields contributing zero characters. Mean length is 197.46 characters in control and 200.29 in treatment. A two-sided Kolmogorov–Smirnov test gives $D = 0.00990$ ($p = 0.220$, using the asymptotic reference distribution); a test of equality of means gives $p = 0.290$. These tests do not detect a difference in the retained profile content. The measure does not capture visits, messages, or search activity. Separate checks of unrelated profile entries and pre-cohort employment histories appear in Appendices N.1 and N.2.

H Experiment 1 (Mail Spin): Summary Statistics and Covariate Balance

Summary Statistics Table A5 reports summary statistics for the main variables used in the analysis, separately for the treatment and control groups. Each entry shows the number of observations, mean,

standard deviation, minimum, and maximum. The unit of observation is an individual learner.

The key variables capture prior online learning activity and background characteristics. *Prior Course Completions* and *Prior Specialization Completions* measure the number of Coursera courses and specializations completed by each learner before the intervention. *Career Goal* is an indicator equal to one if the learner explicitly stated a career-related objective on the platform. *No College Degree* identifies learners without a university degree, *Developing Country* flags those residing in developing economies, and *Tech Domain* indicates whether the course belongs to a technology-related field. *First Skill Score* measures baseline proficiency in the first assessed skill dimension (measured at course completion), and *Course in English* equals one if the course language is English.

In addition to these numeric and binary covariates, the dataset contains several categorical variables that provide context: the learner's *country of origin*, the *name of the course*, the *skills taught* in the course (up to three skill tags), the *primary domain* and *subdomain* classification of each course, as well as metadata such as *course length* and *difficulty level*.

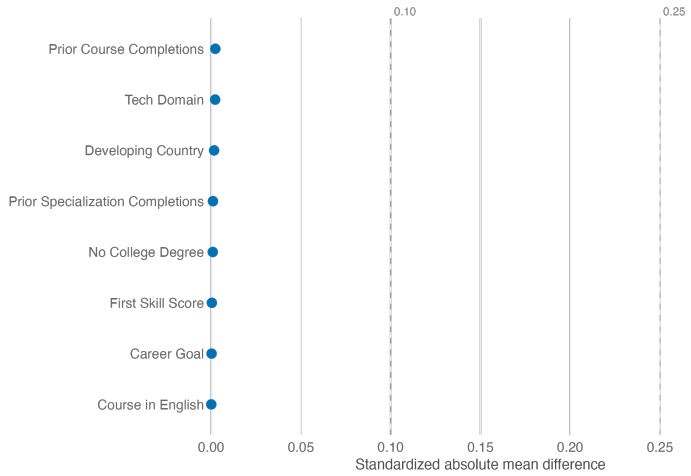
TABLE A5
SUMMARY STATISTICS BY TREATMENT STATUS, EXPERIMENT 1

Variable	Learners	Mean	SD	Min	Max
Panel A. Treatment					
Prior course completions	427,337	3.36	6.05	0.00	877.00
Prior specialization completions	427,337	0.23	0.79	0.00	105.00
Career Goal	427,337	0.12	0.32	0.00	1.00
No College Degree	427,337	0.67	0.47	0.00	1.00
Developing Country	427,337	0.50	0.50	0.00	1.00
Tech Domain	427,337	0.50	0.50	0.00	1.00
First Skill Score	242,739	2.90	1.96	−2.60	13.99
Course in English	427,337	0.96	0.19	0.00	1.00
Panel B. Control					
Prior course completions	427,439	3.35	6.10	0.00	848.00
Prior specialization completions	427,439	0.23	0.83	0.00	136.00
Career Goal	427,439	0.12	0.32	0.00	1.00
No College Degree	427,439	0.68	0.47	0.00	1.00
Developing Country	427,439	0.50	0.50	0.00	1.00
Tech Domain	427,439	0.50	0.50	0.00	1.00
First Skill Score	242,538	2.91	1.97	−2.64	13.71
Course in English	427,439	0.96	0.19	0.00	1.00

Notes: Baseline characteristics of Experiment 1 learners, by assignment. Columns report the number of learners, mean, standard deviation, minimum, and maximum. Prior completions count Coursera courses and specializations completed before the intervention. Career Goal indicates a stated career objective; No College Degree, absence of a university degree; Developing Country, residence in a developing economy (Appendix C.1); Tech Domain, a technology-related course; and Course in English, an English-language course. First Skill Score measures proficiency in the first assessed skill dimension at course completion.

Covariate Balance in Experiment 1 Figure A3 plots standardized absolute mean differences (SMDs) for the main covariates between the treatment and control groups. Each point represents the difference in means for a given baseline variable, standardized by the pooled standard deviation. Dashed vertical lines at 0.10 and 0.25 mark thresholds for small and moderate imbalance. Across all variables, standardized differences are close to zero and well below conventional thresholds.

Figure A3: Difference in Covariate Values Between Treatment and Control



Notes: Each point shows the standardized absolute mean difference between the treatment and control groups of Experiment 1 for a baseline covariate, standardized by the pooled standard deviation. Dashed vertical lines at 0.10 and 0.25 mark conventional thresholds for small and moderate imbalance.

H.1 Heterogeneous Visibility Effects in Experiment 1

Table A6 reports the average treatment effect and conditional average treatment effects by subgroup for Experiment 1 (the email campaign). The outcome is whether a certificate received a view, and we use a difference-in-means estimator. The visibility gains are substantially larger for learners from developing countries (a 6.9 percentage point increase, 39% from baseline, versus 1.4 percentage points for learners from developed countries) and for learners without college degrees (4.9 versus 2.6 percentage points). New learners and learners in non-tech domains also see larger impacts. These patterns motivated the population we target in Experiment 2.

TABLE A6
AVERAGE AND HETEROGENEOUS TREATMENT EFFECTS IN EXPERIMENT 1

	Characteristic = 0		Characteristic = 1		
	Baseline (1)	ATE (2)	Baseline (3)	CATE (4)	Difference (5)
Panel A. Overall					
Overall	0.166	0.042	–	–	–
	(0.000)	(0.001)			
Panel B. Heterogeneity by subgroup					
No College Degree	0.212	0.026	0.144	0.049	0.023
	(0.000)	(0.002)	(0.000)	(0.001)	(0.002)
Developing Country	0.157	0.014	0.175	0.069	0.055
	(0.000)	(0.001)	(0.000)	(0.001)	(0.002)
New User	0.131	0.011	0.180	0.054	0.043
	(0.000)	(0.001)	(0.000)	(0.001)	(0.002)
Tech Domain	0.135	0.059	0.198	0.024	−0.035
	(0.000)	(0.001)	(0.000)	(0.001)	(0.002)

Notes: The sample is the population of Experiment 1 (the email campaign). The outcome is Certificate Views, a binary variable equal to 1 if the certificate received a view and 0 otherwise. Baseline columns report the mean outcome in the control group. ATE and CATE columns report the difference in means (treated − control) with Welch standard errors in parentheses. Panel A reports the overall ATE. Panel B reports conditional ATEs for learners with the row characteristic equal to 0 (columns (1)–(2)) and equal to 1 (columns (3)–(4)): No College Degree, Developing Country, New User (below the sample mean of prior course completions), and Tech Domain. Column (5) reports the difference in CATEs between characteristic values 1 and 0, with its standard error in parentheses.

I Visibility Effects by Treatment Margin

Table A7 reports receipt and click IV effects in the same unique-learner view sample and with the same controls as the assignment estimates.[12] These first stages differ from those in the matched employment sample.

[12] In the full certificate-view sample, 28 control learners have a recorded receipt flag. We therefore use the treated-minus-control difference in receipt as the first stage and report receipt-IV estimates without imposing one-sided receipt.

TABLE A7
VISIBILITY EFFECTS BY TREATMENT MARGIN

Outcome	ITT	Receipt IV	Click IV
Any certificate-page view	0.66	1.38	16.16
	(0.08)	(0.17)	(2.04)
Any view by others	0.29	0.62	7.24
	(0.07)	(0.16)	(1.83)
Any LinkedIn-referred view	0.65	1.36	15.90
	(0.08)	(0.16)	(1.92)
LinkedIn-referred view by others	0.27	0.57	6.73
	(0.07)	(0.15)	(1.72)

Notes: Effects in percentage points, HC1 standard errors in parentheses, on 892,929 learners including September. Each first stage uses the outcome regression sample and administrative controls. Receipt and click IV effects concern different treatment margins. A small number of controls have recorded receipt flags; these flags are retained in the receipt first stage.

J Results by Recruitment Cohort

Table A8 reports detected display and employment for each of the six matched cohorts. September was a pilot and had earlier profile collection. The cohort rows use the available follow-up for each batch.

TABLE A8
ASSIGNMENT EFFECTS BY RECRUITMENT COHORT

Cohort	Learners	Detected display	New Job	New Job in Scope
September	8,356	1.02	0.69	0.68
		(0.85)	(0.79)	(0.59)
October	7,602	2.26	1.32	0.86
		(0.90)	(0.91)	(0.73)
November	8,012	3.05	0.40	0.73
		(0.83)	(0.85)	(0.69)
December	8,008	3.04	2.57	2.05
		(0.86)	(0.86)	(0.69)
January	7,553	4.57	1.40	0.79
		(0.91)	(0.88)	(0.70)
February	5,506	1.68	−1.12	−0.52
		(1.02)	(0.94)	(0.75)

Notes: Unadjusted differences in percentage points, HC1 standard errors in parentheses. All assigned matched learners remain in each denominator. September was the pilot. Profiles for September were collected earlier than those for October–February; these outcomes use available cohort-specific follow-up.

K Employment Transitions

Table A9 separates the type of new job from employment before treatment. Panel A first reports the overall New Job effect, then distinguishes same-employer moves from starts with a new employer. The latter includes entry from nonemployment and study as well as moves between employers. Same-

employer moves are a proxy for promotions: the histories identify employer continuity but do not establish an increase in rank or pay. We retain the established employer-name and date-matching rules. Unresolved employer or history records, and entries describing nonwork, remain in the aggregate outcome but are not assigned to either employer category.

The overall effect is 0.82 percentage points (0.34). The estimate for new-employer starts is 0.69 points (0.31), while the same-employer estimate is -0.04 points (0.16). Thus the identifiable gains occur primarily in starts with new employers. The remaining 491 records contribute 0.16 points (0.10); the two displayed employer categories consequently do not exhaust the overall effect.

Panel B compares the overall New Job effect across employment-status groups defined before treatment. A learner has a recorded prior job if any dated work position spans the calendar month preceding the certificate-grant cohort: its start is no later than that month and its end is in or after that month, or explicitly recorded as ongoing. We exclude titles explicitly describing unemployment or study while retaining paid student jobs and internships. Missing and year-only dates do not establish employment in a specific month. Of the 29,290 learners with no qualifying prior job, 539 have potentially relevant work entries with incomplete dates. This group may therefore include workers with incomplete histories as well as unemployed learners and those outside the labor force.

The effect is 1.27 percentage points (0.74) among the 15,747 learners with a recorded prior job and 0.54 points (0.29) among the 29,290 without one. The difference is 0.73 points (0.80; $p = 0.361$), so the estimates do not establish a difference across groups. These are within-group treatment effects from separate regressions with the main control set. Employment status is reconstructed retrospectively from dated profiles; interpreting this contrast as baseline heterogeneity requires encouragement not to change the recording of prior employment.

TABLE A9
NEW JOBS BY EMPLOYER TRANSITION AND PRIOR EMPLOYMENT

Outcome or group	Learners	Control (%)	ITT (pp)
Panel A. Type of new job			
All new jobs	45,037	16.77	0.82
			(0.34)
Same-employer move (promotion proxy)	45,037	3.09	−0.04
			(0.16)
New employer, including employment entry	45,037	12.67	0.69
			(0.31)
Panel B. Employment before treatment			
Recorded job in preceding month	15,747	33.35	1.27
			(0.74)
No recorded job in preceding month	29,290	8.03	0.54
			(0.29)
Difference between groups			0.73
			(0.80)

Notes: All results include September. HC1 standard errors are in parentheses; control variables follow Table 3. Panel A retains all learners in each regression. Same-employer moves use overlapping or adjacent recorded jobs at a matched employer and can include lateral moves. New-employer starts include moves from another employer, returns after a recorded gap, and first recorded jobs, including student entries. The 491 unclassified new-job records remain in All new jobs; their control mean is 1.01 percent and ITT is 0.16 percentage points (0.10). Panel B estimates separate regressions by whether a dated work record spans the calendar month before the certificate-grant cohort. No recorded prior job does not establish unemployment. The difference is the first group minus the second ($p = 0.361$); its standard error allows separate nuisance coefficients in the two groups.

L Sharing Responses by Employability

Table 4 compares the bottom 40 percent with the upper 60 percent in the same held-out evaluation sample as the employment comparison. Receipt, clicking, detected display, and views are assignment outcomes. None of their between-group differences rejects equality at five percent. Separate subgroup IV estimates would combine differences in take-up and outcomes and do not identify the reason for the employment gradient.

M Construction of Predicted Employability

The prediction model uses registration education, gender, primary course domain, credential type, developing-country status, cohort, and skill-score levels, normalization, and missingness. It excludes retrospective employment-history predictors. After sorting learners by identifier, we allocate one-third to training within assignment-by-cohort cells using seed 3303. The model fits only the 7,495 training controls; the 30,027 evaluation learners supply no training outcomes.

The Bernoulli gradient boosting model has 100 trees, interaction depth five, learning rate 0.01, bag fraction 0.5, and minimum node size ten. Hyperparameters are fixed; three-fold training validation

records predictive performance. The final fit uses seed 20307. Predicted ranks break ties using a permutation with seed 72831. The bottom group contains 12,011 learners, and the upper group 18,016.

For the continuous comparison, let r_i be the predicted rank scaled to $[0, 1]$. We estimate

$$Y_i = \alpha + \tau Z_i + \gamma r_i + \delta Z_i r_i + \lambda_{c(i)} + \varepsilon_i. \tag{2}$$

The gradient per ten-percentile increase is 0.1δ. The group comparison estimates separate regressions in the bottom 40 percent and upper 60 percent, including cohort indicators centered within group and their interactions with assignment. HC1 standard errors condition on the fitted predictor. Group-effect differences use the sum of the independent group variances.

N Placebo Tests

This appendix examines effects on unrelated profile content and recorded work histories as checks on the employment results of Section 3.2. Failure to reject a reporting response does not establish its absence.

N.1 General LinkedIn Profile Activity

Table A10 re-estimates the reference outcomes and unrelated profile components with the same sample and controls. The employment and Coursera-content effects are positive. The four unrelated-entry checks show no statistically detectable increase; these nulls do not establish that every form of reporting or search activity is unchanged.

N.2 Reporting of Earlier Employment

The position histories permit checks closer to the employment outcome itself. Table A11 compares recorded pre-cohort positions, end dates, and undated entries. None of the eight placebos rejects equal means after Holm adjustment. For example, the any-earlier-end-date effect is -0.08 percentage points (0.31), and the effect on listing at least two earlier positions is 0.09 points (0.29). The checks retain all matched learners and the original history classifications. They do not resolve uncertainty in the cohort anchors or observe when entries were added.

TABLE A10
ASSIGNMENT EFFECTS ON OTHER PROFILE COMPONENTS

	Control	Effect	p
Any Coursera item (%)	29.582	0.996 (0.399)	0.013
New Job (%)	16.770	0.819 (0.339)	0.016
New Job in Scope (%)	9.579	0.705 (0.273)	0.010
New education entry (%)	2.095	−0.186 (0.131)	0.158
Number of licenses/certifications	4.050	0.002 (0.075)	0.977
Number of courses	0.542	0.038 (0.028)	0.172
Number of education entries	0.954	−0.000 (0.011)	0.991

Notes: All 45,037 learners; controls match Table 3. Parentheses report HC1 standard errors. Binary effects are percentage points, counts are levels. The education-entry indicator requires a recorded start year of 2023 or later. These checks do not measure all LinkedIn activity.

TABLE A11
ASSIGNMENT EFFECTS ON RECORDED POSITION HISTORIES

	Control	Effect	p
Any earlier position with an end date	0.368	−0.0008 (0.0031)	0.805
Ended before cohort	0.347	−0.0006 (0.0031)	0.846
Number with an end date	1.464	−0.0040 (0.0170)	0.815
Number ended before cohort	1.293	−0.0051 (0.0157)	0.747
Two or more earlier positions	0.341	0.0009 (0.0029)	0.757
Number of earlier positions	1.761	0.0076 (0.0180)	0.674
Any undated position	0.044	−0.0004 (0.0018)	0.843
Number of undated positions	0.080	0.0006 (0.0044)	0.885

Notes: All outcomes in levels, including binary indicators. All 45,037 learners; controls match Table 3; HC1 standard errors in parentheses. These eight outcomes form the history-placebo family, with no rejection after Holm adjustment.

O Certificate Views and Reported Employment

Table A12 describes the association between certificate views and recorded employment. The regressions include an intercept and use unique learners; the comparison is descriptive, not a causal effect of receiving views.

TABLE A12
ASSOCIATION OF CERTIFICATE VIEWS WITH NEW JOBS

Dependent variable	New Job coefficient (pp)
Any certificate-page view	10.03
	(0.62)
Any view by others	10.10
	(0.61)
Any LinkedIn-referred view	10.44
	(0.62)
LinkedIn-referred view by others	10.19
	(0.60)

Notes: Each regression uses the view indicator as its dependent variable and New Job as its regressor, with an intercept, on 45,037 learners. HC1 standard errors in parentheses. These are descriptive associations, not effects of receiving views.

P Sample Job Descriptions

The evaluation uses four fixed job descriptions: Junior Business Analyst, Junior Software Engineer, Junior Data Scientist, and Junior IT Support Specialist. Each learner is scored against the job assigned to the primary domain of the learner's credential. The retained scoring code loads these four texts; it does not identify the model version originally used to draft them. We therefore document the evaluation texts rather than attribute their generation to a particular model. An earlier pilot with multiple jobs per domain does not enter the reported scores.

The following descriptions show the evaluation context. They are hypothetical jobs, not vacancies to which the learners are observed to apply. The relatively demanding degree and skill requirements help determine the level of the resulting scores and limit their interpretation as a general measure of employability.

Position: Junior Business Analyst

Overview: We are looking for an analytical and detail-oriented Junior Business Analyst to support our business operations. You will assist in analyzing business processes, gathering requirements, and helping drive data-informed decision making.

Required Qualifications:

- Bachelor's degree in Business Administration, Economics, Finance, or related field
- Strong analytical and critical thinking skills
- Proficiency in Microsoft Office Suite (Excel, PowerPoint, Word)
- Excellent written and verbal communication skills
- Ability to work with cross-functional teams
- Basic understanding of business processes and operations
- Strong attention to detail

Preferred Qualifications:

- Internship experience in business analysis, consulting, or related field
- Familiarity with data visualization tools (Tableau, Power BI)
- Basic SQL knowledge for data querying
- Experience with process mapping and documentation
- Understanding of project management principles
- Exposure to CRM or ERP systems
- Knowledge of business intelligence concepts

Key Skills: Excel, PowerPoint, Data Analysis, Business Analysis, Requirements Gathering, Process Mapping, Communication, Presentation Skills, Problem Solving, SQL, Tableau, Power BI, Stakeholder Management, Documentation, Critical Thinking, Project Management, Teamwork

Position: Junior Software Engineer

Overview: We are seeking a motivated Junior Software Engineer to join our development team. You will contribute to the design, development, and maintenance of software applications while learning from experienced engineers.

Required Qualifications:

- Bachelor's degree in Computer Science, Software Engineering, or related field
- Strong foundation in programming fundamentals and data structures
- Knowledge of at least one programming language (Python, Java, C++, or JavaScript)
- Understanding of algorithms and computational complexity
- Familiarity with version control systems (Git)
- Problem-solving skills and logical thinking
- Ability to work collaboratively in a team environment

Preferred Qualifications:

- Internship or project experience in software development
- Familiarity with object-oriented programming principles
- Understanding of software development lifecycle
- Experience with debugging and testing
- Exposure to Agile/Scrum methodologies
- Contributions to open source projects or personal coding projects
- Knowledge of web technologies (HTML, CSS, JavaScript)

Key Skills: Python, Java, C++, JavaScript, Git, Data Structures, Algorithms, Object-Oriented Programming, Debugging, Testing, Problem Solving, Teamwork, Communication, Agile, Software Development, Version Control, Code Review

Position: Junior Data Scientist

Overview: We are seeking a Junior Data Scientist to join our analytics team. You will work on data-driven projects, perform statistical analysis, build predictive models, and help transform data into actionable insights.

Required Qualifications:

- Bachelor's or Master's degree in Data Science, Statistics, Mathematics, Computer Science, or related field
- Strong foundation in statistics and probability
- Programming skills in Python or R
- Knowledge of data manipulation and analysis libraries (pandas, NumPy)
- Understanding of machine learning fundamentals
- Ability to work with large datasets
- Strong analytical and problem-solving skills

Preferred Qualifications:

- Experience with machine learning libraries (scikit-learn, TensorFlow, PyTorch)
- Familiarity with data visualization tools (Matplotlib, Seaborn, Plotly)
- Knowledge of SQL and database querying
- Experience with Jupyter notebooks
- Understanding of A/B testing and experimental design
- Exposure to big data technologies (Spark, Hadoop)
- Portfolio of data science projects or Kaggle competitions

Key Skills: Python, R, Statistics, Machine Learning, Data Analysis, pandas, NumPy, scikit-learn, SQL, Data Visualization, Matplotlib, Statistical Modeling, Predictive Analytics, Data Cleaning, Jupyter, Problem Solving, Communication, TensorFlow, A/B Testing, Experimentation

Position: Junior IT Support Specialist

Overview: We are looking for a customer-focused Junior IT Support Specialist to provide technical assistance to our organization. You will troubleshoot hardware and software issues, support end-users, and help maintain our IT infrastructure.

Required Qualifications:

- Bachelor's degree in Information Technology, Computer Science, or related field (or equivalent experience)
- Strong understanding of computer hardware, software, and networking fundamentals
- Familiarity with Windows and/or macOS operating systems
- Basic knowledge of networking concepts (TCP/IP, DNS, DHCP)
- Excellent customer service and communication skills
- Ability to troubleshoot technical issues methodically
- Strong attention to detail and documentation skills

Preferred Qualifications:

- CompTIA A+, Network+, or similar IT certifications
- Experience with help desk ticketing systems
- Knowledge of Active Directory and user management
- Familiarity with cloud platforms (Microsoft 365, Google Workspace)
- Understanding of cybersecurity best practices
- Experience with remote desktop support tools
- Scripting knowledge (PowerShell, Bash)

Key Skills: Windows, macOS, Linux, Networking, TCP/IP, Hardware Troubleshooting, Software Installation, Help Desk, Technical Support, Customer Service, Active Directory, Microsoft 365, Google Workspace, Ticketing Systems, Documentation, Problem Solving, Communication, Cybersecurity, Remote Support, CompTIA A+

Q Profile Scoring: Construction and Validation

All 45,037 learners have valid baseline, enhanced, and observed-profile scores from Claude Haiku 4.5 (model identifier `claude-haiku-4-5-20251001`), and a second, independent evaluation of the baseline profile from the same model. We retain the previously scored learners and add the 8,356 September profiles using the same judge and four fixed job descriptions (Appendix P). Both cohorts use the same judge.

Constructed texts. The paired baseline removes certification entries dated after September 1, 2022. Its job text reports durations rather than dates. The enhanced version adds a standardized Coursera credential naming the course domain and up to three recorded skills, including an artificial identifier. The pair holds the remaining text fixed. The effect therefore concerns the added credential and does not measure an observed employer decision.

The observed-profile version retains certifications as posted and removes jobs with structured start dates on or after September 1, 2022. This observed score, rather than the hypothetical paired gain, enters the preferred surrogate index.

Prompt. The judge is asked to act as an HR recruiter, evaluate the resume against the fixed domain-matched job description, and return a score from 0 to 100 with a brief explanation in JSON. The prompt asks it to consider relevant experience and skills, education, overall fit, and presentation. Failed requests are retried; they are not assigned zero. Shared text can use cached or averaged scores, so learner-level uncertainty does not include independent model uncertainty for every row.

Replication and extension checks. We preserve every existing score when appending September. Construction checks on 2,000 earlier learners reproduce the baseline and observed text; enhanced text agrees after masking synthetic identifiers. The extension uses a reproducible synthetic-identifier seed and averages repeated evaluations of identical texts. These implementation differences are recorded in the replication archive.

Independent baseline evaluation. Every learner has a second baseline evaluation that shares no scoring request with the primary baseline or enhanced score. It comes from a replicate batch: a copy, in a new random order, of the batch in which the learner's baseline and enhanced profiles were first

TABLE A13
CREDENTIAL GAINS BY INDEPENDENTLY RANKED BASELINE QUARTILE

Controls	Learners	Ranking evaluation	Scores entering the gain: Baseline	Scores entering the gain: Enhanced	Gain (points)
All controls	22,487	19.88	21.53	37.51	15.97 (0.09)
First quartile	5,622	2.63	4.54	29.42	24.88 (0.16)
Second quartile	5,622	4.92	5.22	28.72	23.51 (0.14)
Third quartile	5,621	17.68	22.67	32.99	10.32 (0.12)
Top quartile	5,622	54.27	53.69	58.89	5.19 (0.09)
First minus top quartile					19.69 (0.18)

Notes: Claude Haiku 4.5 scores on a 0–100 scale for all 22,487 controls, including September. Quartiles rank learners by an independent baseline evaluation (column Ranking evaluation) that shares no scoring request with the Baseline and Enhanced columns. Gain is the enhanced score minus the primary baseline score; the ranking evaluation does not enter the gain. Learners are sorted by ranking score and then by anonymous learner identifier and divided into four groups of equal size by position, so the identifier orders tied scores. The ranking scores at the three group boundaries are 3, 5, and 32, and learners tied at these values fall on both sides. Profiles that list no education, experience, certifications, or courses share one ranking evaluation within a course domain and scoring run; they make up 97, 88, 2, and 0 percent of the four quartiles, so the division between the first two quartiles largely reflects the order of ties. Parentheses report standard errors that reflect learner variation conditional on realized scores and group assignments; they exclude model-evaluation uncertainty and the dependence created by scores shared across identical texts. The last row is the difference in mean gains between the first and top quartiles. Figure 3 uses the same evaluations.

scored, evaluated in a separate request with the same model, prompt, and job description. The original scoring replicated a subset of batches, which gives this evaluation for 13,342 learners, including 6,705 controls. For 8,572 further learners we reuse an existing replicate evaluation of the identical baseline text and job description; 8,550 of them have the profile that lists no education, experience, certifications, or courses. The remaining 23,123 learners, including 11,467 controls, were evaluated in 2,344 new replicate batches. When a profile text is shared by many learners, the primary score averages several evaluations; none of them is the independent evaluation. The two evaluations of a profile can return the same number. Independence refers to the scoring requests, not to the values.

The primary and independent baseline scores have a correlation of 0.931 across all 45,037 learners and 0.868 among the 22,794 learners whose profile text is longer than that blank profile. For these learners, a second evaluation of the enhanced profile has a correlation of 0.828 with the primary enhanced score, and the gain computed from the second evaluations has a correlation of 0.490 with the primary gain. Individual gains are therefore measured less reliably than score levels. Repeat evaluations of the observed profile exist only for 100 September profiles (correlation 0.828). The new

evaluations score somewhat lower than the earlier ones: on 120 existing replicate batches that we scored again, the difference is −2.25 points (standard error 0.56, clustered by batch). Adding this difference back to the new evaluations before ranking changes the first-minus-top difference in Table A13 from 19.69 to 19.97 points. These diagnostics quantify repeated evaluation under this model; they do not validate scores as actual hiring probabilities.

Table A14 reports relationships with recorded employment. These regressions describe associations rather than causal effects.

TABLE A14
PROFILE SCORES AND RECORDED EMPLOYMENT AMONG CONTROLS

Score	Coefficient (pp)	R^2
Baseline score	6.00	0.133
	(0.12)	
Hypothetical paired gain	−9.25	0.101
	(0.17)	
Observed-profile score	5.31	0.113
	(0.12)	

Notes: Separate regressions of New Job on each score and an intercept, using all 22,487 controls including September. Coefficients are percentage points per ten score points, with HC1 standard errors in parentheses.

Table A15 reports the employment regressions using $Q_i = D_i(s_i^1 - s_i^0)$, the hypothetical gain multiplied by actual sharing. Unlike the unweighted gain in Table A14, this measure assigns zero to learners whose earned certificate is not detected on LinkedIn. Panel A reports the positive pooled and separate-arm associations discussed in Section 4. Adding the baseline score in Panel B reduces the estimates to 6.99 percentage points (standard error 0.46) pooled, 6.81 (0.62) in treatment, and 7.20 (0.66) in control.

TABLE A15
CERTIFICATE SCORE GAINS, ACTUAL SHARING, AND EMPLOYMENT

	Pooled	Treatment	Control
Panel A. Gain interacted with sharing			
Gain × actual sharing	9.00	8.80	9.23
	(0.48)	(0.66)	(0.69)
R^2	0.012	0.012	0.013
Panel B. Controlling for the baseline score			
Gain × actual sharing	6.99	6.81	7.20
	(0.46)	(0.62)	(0.66)
Baseline score	5.94	5.99	5.90
	(0.09)	(0.12)	(0.12)
R^2	0.140	0.140	0.141
Treatment assignment included	Yes	No	No
Learners	45,037	22,550	22,487

Notes: Linear probability models with New Job as the dependent variable. The regressor is $Q_i = D_i(s_i^1 - s_i^0)$, where s_i^0 is the baseline Haiku score, s_i^1 is the enhanced score after adding a standardized Coursera credential, and D_i indicates detection of the earned certificate on LinkedIn using the certificate-code matching measure. It is zero for non-sharers; negative gains are retained. Each panel contains three separate regressions with an intercept. Panel A includes Q_i; Panel B adds the baseline score. Pooled regressions also include randomized treatment assignment. No other covariates are included. Coefficients are percentage points per ten score points; parentheses report HC1 robust standard errors conditional on the realized scores. Samples include September.

R Comparison of Llama and Haiku Assessments

We assess whether the scoring results depend on the choice of judge by comparing Llama 3.1 8B and Claude Haiku 4.5 on the same 18,326 control learners. Both models evaluate the baseline and enhanced profiles against the same domain-specific junior job descriptions. The comparison covers the October 2022–February 2023 cohorts, for which both sets of scores are available; it excludes September.

The models' assessments are positively correlated: the Pearson correlation is 0.75 for baseline scores and 0.56 for enhanced scores, with corresponding rank correlations of 0.73 and 0.49. Both judges find that adding a credential improves assessed fit, although the average gain is larger under Haiku: 15.95 points compared with 8.28 under Llama.

Table A16 compares gains in the first and top quartiles of the independent Haiku baseline evaluation (Table A13), using the same learners for both judges. In the first quartile, the mean score increases by 10.48 points under Llama and 24.79 points under Haiku. In the top quartile, the increases are 4.44 and 5.27 points, respectively. Both models therefore show larger gains in the first quartile, but differ substantially in the size of those gains: the first-minus-top difference is 6.04 points (standard error 0.27) under Llama and 19.52 points (0.20) under Haiku. Agreement about individual learners' gains is weaker than agreement about score levels, with a correlation of 0.17 across all matched controls.

TABLE A16
LLAMA AND HAIKU CREDENTIAL GAINS BY INDEPENDENTLY RANKED BASELINE QUARTILE

Model	Baseline	Enhanced	Gain (points)
Panel A. All matched controls ($N = 18,326$)			
Llama	34.84	43.12	8.28
			(0.10)
Haiku	21.50	37.44	15.95
			(0.10)
Panel B. First quartile ($N = 4,537$)			
Llama	21.07	31.55	10.48
			(0.17)
Haiku	4.53	29.32	24.79
			(0.17)
Panel C. Top quartile ($N = 4,543$)			
Llama	57.52	61.96	4.44
			(0.21)
Haiku	53.58	58.85	5.27
			(0.10)

Notes: Llama 3.1 8B and Claude Haiku 4.5 assessments on a 0–100 scale. Both models score the same baseline and enhanced profiles against the same domain-specific junior job descriptions. Panel A includes the 18,326 controls scored by both models. Panels B and C use the first and top quartiles of all 22,487 controls ranked by the independent Haiku baseline evaluation, as defined in Table A13, restricted to learners scored by both models. The groups are therefore the same for both models, and the ranking evaluation enters neither model's gain. Llama uses its first baseline evaluation; Haiku uses the primary scores of the main analysis. The evaluations differ in instructions and batching. Standard errors in parentheses condition on realized scores and group assignments and exclude model-evaluation uncertainty.

S Employment Results Excluding the September Pilot

September 2022 was the pilot month of Experiment 2. The main analyses include this cohort. For completeness, Table A17 repeats Table 3 for the 36,681 learners in the October 2022–February 2023 cohorts, using the same outcome definitions, controls, and ITT, ATT, and LATE estimators. This comparison also excludes the cohort whose profiles and view counts were collected earlier. The any-job ITT is 0.92 percentage points (0.38), compared with 0.82 (0.34) including September; the in-scope ITT is 0.77 (0.31), compared with 0.71 (0.27).

TABLE A17
REPORTED EMPLOYMENT EXCLUDING THE SEPTEMBER PILOT

	Control (%) (1)	ITT (pp) (2)	ATT (pp) (3)	LATE (pp) (4)
Panel A. Recorded new-position outcomes				
New Job	17.2	0.92	0.95	8.36
		(0.38)	(0.40)	(3.47)
New Job in Scope	10.0	0.77	0.80	7.01
		(0.31)	(0.32)	(2.82)
Panel B. Four-month restriction				
New Job	8.9	0.54	0.57	4.96
		(0.30)	(0.31)	(2.70)
New Job in Scope	5.2	0.40	0.42	3.65
		(0.23)	(0.24)	(2.12)

Notes: Effects and standard errors are percentage points. All 36,681 unique matched learners enter each regression. ITT is the effect of assignment. ATT is the effect of encouragement among treated app-launchers, estimated by instrumenting assignment × app launch with assignment. LATE instruments clicking to share with assignment. The first stages are 95.95 and 10.95 points, respectively; HC1 robust standard errors include first-stage uncertainty. Controls are fixed effects for reconstructed cohort–domain–country-group–degree cells, education, gender, credential type, developing-country status, professional experience, and prior technical and business employment. Missing administrative categories are explicit. Section 2.2 states the assumptions; Appendix F defines app access. New Job in Scope requires a new position with a new employer and a title in a field related to the certificates; unresolved titles are coded zero. Panel B restricts events to relative-date indices of at least five; all learners remain in the denominator. Appendix E documents title classification and the date rules.

T Sample Construction and the Exposure Map

Table A18 reports eligible advertisements across the 12 countries. The unrestricted sample contains 1,018 metropolitan-area–field cells in 256 metropolitan areas across twelve countries. Requiring at least 100 advertisements and ten employers in each comparison period, and posting activity in all 18 months, retains 779 cells in 217 metropolitan areas. We apply these rules separately by field within each city. Of the retained cells, 411 have positive projected exposure; 368 meet the coverage rules but receive no destination weight from the control-based map. The rules retain persistent activity and multiple employers, not necessarily stable employer composition. Results need not extend to markets with intermittent coverage. Table A19 summarizes all the metropolitan areas in our sample.

Occupation crosswalk. The four fields contain the following occupation categories: Business—Marketing Coordinator, Product Marketing, Project Coordinator, and Project Management; Computer Science—Software Engineer, Application Developer, and QA Engineer; Data Science—Data Analyst, Data Engineer, and Data Scientist; Information Technology—Cybersecurity Engineer, Infrastructure Engineer, IT Operations, SAP Analyst, Security Analyst, Security Continuity Analyst, Systems Analyst,

TABLE A18
CITY–FIELD SAMPLE BY COUNTRY

			Advertisements	
	Cities	Markets	Pre-experiment	Experiment
Brazil	37	131	189,064	185,583
Egypt	2	8	12,297	15,390
India	49	140	535,845	527,723
Indonesia	5	14	24,341	29,386
Nigeria	2	7	6,067	6,119
Pakistan	5	14	30,452	29,578
Philippines	2	8	62,939	64,855
Saudi Arabia	3	11	5,723	7,354
Singapore	1	4	89,505	64,820
United Arab Emirates	2	8	24,048	25,779
United States	107	426	3,754,275	2,695,675
Vietnam	2	8	13,697	14,955
Total	217	779	4,748,253	3,667,217

Notes: Pre-experiment period: October 2021–March 2022. Experiment period: October 2022–March 2023. Each city–field market has at least 100 advertisements and ten employers per period and an advertisement in every month of October 2021–March 2023. All eligible postings receive weight one, including those from the full India and Brazil extracts. These counts describe Revelio coverage rather than all vacancies in a country.

Systems Engineer, and Technical Support Engineer.

Projection and scale. Of the 892,929 learners in Experiment 2, 667,746 registered in the twelve vacancy-data countries. Registration country and course field define 48 source strata. We estimate destination shares from 4,406 control learners linked to Revelio profiles with an identifiable prior job: the latest job that began by May 2022 and was not known to have ended before June. Its metropolitan area and occupation determine the destination market.

The strata contain seven to 677 located controls each. Registration and prior-job countries must both be covered, but destinations may cross country borders. We estimate shares before imposing outcome-support restrictions. Of 333,873 expected treated learners, 322,114.8 project into supported markets. Median exposure is 0.68 expected treated learners per 100 pre-experiment advertisements; the 90th percentile is 60.15. We standardize exposure by its pre-experiment-posting-weighted SD within each regression sample.

Panels A and B of Table 7 use the twelve-country map; Panel B restricts outcomes to U.S. destinations. Among located controls, 368 have U.S. destination jobs, and 138 of the 426 supported U.S. markets receive positive projected exposure. Panel C restricts both registration and destination countries to the eleven developing countries under the experiment's non-OECD classification. It uses 608,522 learners in 44 source strata and a map based on 4,034 controls. Coverage rules retain 353 developing-

TABLE A19
CITIES IN THE LOCAL LABOR MARKET SAMPLE

Country	Cities
Brazil	Aracaju; Bauru; Belem; Belo Horizonte; Blumenau; Brasilia; Campinas; Campo Grande; Caxias do Sul; Cuiaba; Curitiba; Florianopolis; Fortaleza; Goiania; Joao Pessoa; Joinville; Juiz de Fora; Londrina; Maceio; Manaus; Maringa; Natal; Piracicaba; Porto Alegre; Recife; Ribeirao Preto; Rio de Janeiro; Salvador; Santos; Sao Jose do Rio Preto; Sao Jose Dos Campos; Sao Luis; Sao Paulo; Sorocaba; Teresina; Uberlandia; Vila Velha.
Egypt	Alexandria; Cairo.
India	Agra; Ahmedabad; Amritsar; Aurangabad; Bengaluru; Bhopal; Bhubaneshwar; Chandigarh; Chennai; Cochin; Coimbatore; Dehra Dun; Delhi; Erode; Guwahati; Hyderabad (India); Indore; Jaipur; Jalandhar; Jammu; Jodhpur; Kanpur; Kolkata; Kota; Kozhikode; Lucknow; Ludhiana; Madurai; Mangalore; Meerut; Mumbai; Mysore; Nagpur; Nashik; Navi Mumbai; Neyveli; Patna; Ponnani; Pune; Raipur; Rajkot; Ranchi; Surat; Thiruvananthapuram; Tiruchirappalli; Udaipur; Vadodara; Vijayawada; Visakhapatnam.
Indonesia	Bandung; Jakarta; Semarang; Surabaya; Yogyakarta.
Nigeria	Abuja; Lagos.
Pakistan	Faisalabad; Islamabad; Karachi; Lahore; Multan.
Philippines	Cebu City; Manila.
Saudi Arabia	Dammam; Jeddah; Riyadh.
Singapore	Singapore.
United Arab Emirates	Abu Dhabi; Dubai.
United States	Albany; Albuquerque; Anaheim; Anchorage; Atlanta; Austin; Bakersfield; Baltimore; Baton Rouge; Birmingham (Alabama); Boise; Boston; Buffalo; Charleston; Charlotte; Chattanooga; Chicago; Cincinnati; Cleveland; Colorado Springs; Columbia; Columbus; Corpus Christi; Dallas; Dayton; Denver; Des Moines; Detroit; El Paso; Eugene; Fayetteville; Fort Collins; Fort Lauderdale; Fort Wayne; Fort Worth; Fresno; Gainesville; Grand Rapids; Greensboro; Greenville; Hartford; Honolulu; Houston; Huntsville; Indianapolis; Jackson; Jacksonville; Kansas City; Knoxville; Las Vegas; Lexington; Lincoln; Little Rock; Long Beach; Los Angeles; Louisville; Lubbock; Madison; Memphis; Miami; Milwaukee; Minneapolis; Mobile; Montgomery; Nashville; New Orleans; New York City; Oklahoma City; Omaha; Orlando; Philadelphia; Phoenix; Pittsburgh; Portland; Providence; Raleigh; Reno; Richmond; Riverside; Rochester; Sacramento; Salt Lake City; San Antonio; San Diego; San Francisco; San Jose (California); Savannah; Seattle; Sioux Falls; Spokane; Springfield; St Louis; Stockton; Syracuse; Tacoma; Tallahassee; Tampa; Toledo; Tucson; Tulsa; Virginia Beach; Washington; West Palm Beach; Wichita; Wilmington; Winston Salem; Worcester.
Vietnam	Hanoi; Ho Chi Minh City.

Notes: The 217 metropolitan areas containing the 779 city–field markets in Table 7. Each listed area has at least one field that meets the posting-coverage requirements. A metropolitan area can contribute up to four markets: Business, Computer Science, Data Science, and Information Technology. Cities use Revelio metropolitan-area labels, with the common suffix omitted. The sample includes the full India and Brazil posting extracts.

country markets, including 131 in Brazil and 140 in India. The map and exposure scale are fixed across specifications within each panel.

U Credential Visibility in Revelio

This appendix documents the displayed-credential data and trend adjustment in Section 5.2, then compares U.S. education groups.

Data and timing. The extract contains 105.6 million certificate records from approximately 29.0 million users in the twelve countries covered by the vacancy analysis. Coursera issuer names and URLs identify the provider; normalized titles identify experiment courses. We deduplicate within user, title, and grant date. We do not match to identify a randomized participant.

For the market rows of Table 6, we assign profiles to the 779 study markets using Appendix T's occupation crosswalk. We select each holder's latest job beginning by May 2022 and exclude profiles whose selected job was known to have ended before June. We check location after selecting the job, so an older job cannot replace an unlocatable latest job. Job country, metropolitan area, and field fix market membership across grant cohorts. The position extract covers relevant occupations with starts from 2018 onward; profiles lacking assignment information are excluded. The sample contains 1,020,993 holders with a credential in the fifteen included grant months and covers all 779 markets: 426 U.S. and 353 developing-country markets. Numerators and denominators use the same profiles. Equal-country estimates use profile country for the unrestricted sample and job-market country for the market samples; the weighting schemes coincide for the U.S. sample.

The baseline is September 2021–April 2022 and the post window is June 2022–February 2023. We omit January to limit date bunching and exclude May from both means. These grant cohorts begin earlier than the September–March experimental sample used for vacancy exposure. Grant dates identify neither profile additions nor encouragement delivery.

Pooled credential growth with trend adjustment. Across the twelve countries, mean monthly experiment-course listings rise from 36,257 to 53,095. The mean monthly ratio of Coursera certificates to other-provider certificates rises by 33.4%. To adjust for the pre-existing differential trend, we pool counts across countries in each grant month and fit a linear trend to the log ratio of experiment-course Coursera listings to other-provider listings over the seven baseline months. We extend this trend into the eight post-period months and average the differences between the observed and projected log

ratios. Exponentiating this average and expressing the result as a percentage gives a trend-adjusted increase of 10.5%. Table 6 repeats this calculation for each geographic sample. For the equal-country estimates, we apply the same procedure to the mean within-country log ratio in each month, giving each country equal weight.

U.S. eligibility comparison. In the United States, the experiment targets learners without a college degree. Figure A4 compares experiment-course Coursera credentials relative to other-provider credentials within education groups. The unadjusted ratio rises by 70.1% below bachelor's and 29.6% among graduates. Monthly experiment-course listings rise from 937 to 1,674 and from 10,367 to 14,134, respectively. Counting distinct holders or using an online-learning-provider benchmark preserves this ordering. The comparison is consistent with the targeting rule, though profile education need not equal education reported to Coursera.

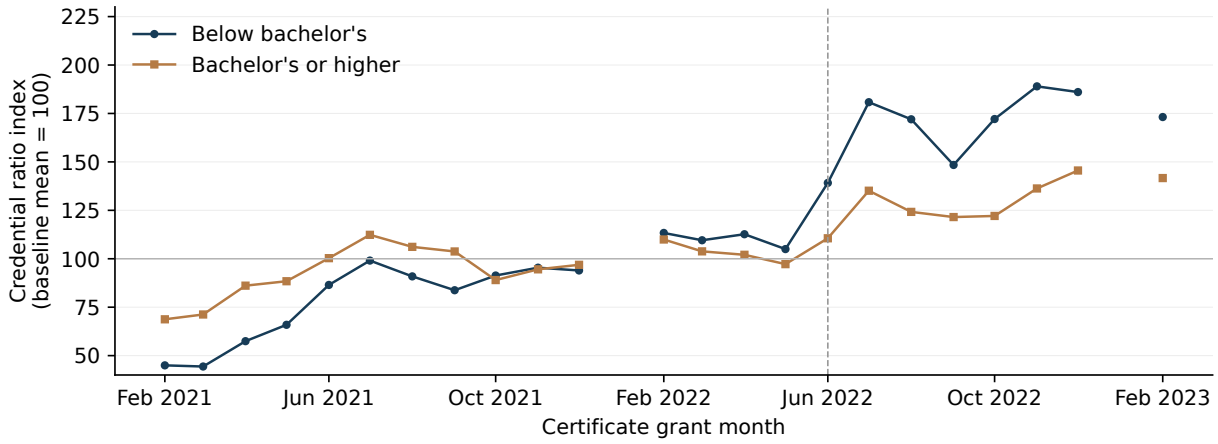


Figure A4: Credential visibility by education in the United States

Notes: Experiment-course Coursera certificates divided by other-provider certificates within education group, indexed to the September 2021–April 2022 mean of that monthly ratio. These descriptive series do not remove prior trends. January is omitted. The vertical line marks the June 2022 grant cohort, not the date of encouragement or profile sharing. High School and Associate define below bachelor's; Bachelor, Master, MBA, and Doctor define bachelor's or higher. Unknown education is excluded.

These descriptive contrasts cannot separate encouragement-induced display from changes in credential acquisition, profile reporting, or data coverage. The vacancy comparison continues to require the market-level identifying assumptions in Section 5.3.

V April–July Placebo Comparison

In this section, we treat the April–July 2022 period as a placebo. We replace the experiment indicator in equation (1) with an April–July indicator and omit later posting months. The outcome remains mean advertisement duration capped at 180 days. Each specification retains its market sample, exposure map, scale, and baseline posting weights. We refit all fixed effects, field-specific log baseline size-by-month controls, and the indicated market trends over these ten months: 779 markets and 7,790 market-months pooled.

Table A20 presents the estimates. Across pooled specifications, placebos range from -0.57 to -0.22 days, versus -2.31 to -1.44 during the experiment. The developing-country quadratic placebo is -0.56 days (0.65), compared with -2.60 (0.87) during the experiment. The U.S. placebo is 0.37 days (0.28). Only the posting-weighted pooled linear-trend regression rejects zero at 5 percent conventionally ($p = 0.042$). Nonrejection establishes neither equivalence to zero nor parallel trends. The placebo also changes the recruiting season and refits trends rather than testing an out-of-sample forecast.

TABLE A20
EXPOSURE AND ADVERTISEMENT DURATION: APRIL–JULY PLACEBO

	(1)	(2)	(3)	(4)
A. All countries				
April–July placebo	−0.459	−0.220	−0.565	−0.418
	(0.270)	(0.358)	(0.246)	(0.279)
95% interval	[−1.054, 0.136]	[−1.009, 0.568]	[−1.107, −0.023]	[−1.032, 0.196]
p, placebo equals zero	0.118	0.551	0.042	0.162
Markets	779	779	779	779
Placebo market-months	7,790	7,790	7,790	7,790
B. United States				
April–July placebo	0.261	0.366	−0.117	−0.129
	(0.266)	(0.284)	(0.168)	(0.195)
95% interval	[−0.266, 0.789]	[−0.198, 0.929]	[−0.449, 0.216]	[−0.516, 0.257]
p, placebo equals zero	0.328	0.201	0.488	0.508
Markets	426	426	426	426
Placebo market-months	4,260	4,260	4,260	4,260
C. Developing countries				
April–July placebo	−0.943	−0.559	−0.788	−0.527
	(0.463)	(0.650)	(0.355)	(0.432)
95% interval	[−1.973, 0.088]	[−2.007, 0.888]	[−1.578, 0.003]	[−1.489, 0.435]
p, placebo equals zero	0.069	0.409	0.051	0.251
Markets	353	353	353	353
Placebo market-months	3,530	3,530	3,530	3,530
Market weights	Equal	Equal	Posting	Posting
Market trend	Linear	Quadratic	Linear	Quadratic

Notes: Entries are duration differences in days per one-standard-deviation higher projected exposure. The placebo compares April–July 2022 with October 2021–March 2022. Outcome: market-month mean advertisement duration, capped at 180 days. Regressions include market, city-by-month, and country–field-by-month effects, field-specific log baseline posting count interacted with month, and the indicated market trends, all refitted over these ten months. The market sample, exposure map, scale, and baseline posting weights are fixed within panels. Exposure SDs are 20.48, 1.29, and 38.09 in Panels A, B, and C. Parentheses contain CR1 standard errors; brackets give 95 percent confidence intervals. Clustering is by country in A and C (12 and 11 clusters) and metropolitan area in B (107 clusters), with t_{G-1} critical values. Column (2) uses equal weights and quadratic trends. Columns (1)–(4) match the specifications without employer fixed effects in Table 7, columns (1), (3), (5), and (7).

W Credential Presence and Vacancy Closure Before the Experiment

Workers who expect employers to hire more workers or fill vacancies faster may enroll in courses or begin displaying certificates they earned earlier. To check this possibility we examine whether displayed credentials covary with vacancy closure before the experiment.

The sample covers 779 study markets and includes only displayed Coursera credentials whose normalized titles match the experiment's course catalog. We restrict both credential grants and vacancy postings to months before June 2022, the earliest post-period grant cohort in the visibility comparison. Following that comparison, we omit January and May, leaving six months: October–December 2021 and February–April 2022. The resulting sample contains 4,674 market-months, 30,386 holder-months,

and 44,211 certificates. Each worker counts once per grant month, regardless of how many qualifying credentials they received. Market-months without qualifying credentials remain in the sample with zero counts.

We examine whether market-months with faster vacancy closure also have more workers with displayed credentials granted that month. We regress monthly holder counts on the closure speed of vacancies posted in the same market-month, capping individual advertisement durations at 180 days. Closure speed measures how far mean duration falls below the pre-experiment mean, in standard-deviation units; we calculate both the reference mean and the standard deviation of 9.72 days with equal weights across the retained market-months. A positive coefficient therefore indicates more credential holders where vacancies close faster. The comparison adjusts for market effects and quadratic trends, city-by-month and country–field-by-month effects, and field-specific interactions between log baseline posting size and month.

Table A21 reports the results. We find that one-SD faster closure is associated with 0.043 additional holders per market-month (standard error 0.121), against a mean of 6.50. The certificate-count estimate is 0.185 (0.307). Both log-count estimates are negative and imprecise.

TABLE A21
EXPERIMENT-COURSE CREDENTIALS AND VACANCY CLOSURE BEFORE THE EXPERIMENT

	Counts		$\log(1+\text{count})$	
	Workers	Certificates	Workers	Certificates
	(1)	(2)	(3)	(4)
One SD faster closure	0.043	0.185	−0.029	−0.026
	(0.121)	(0.307)	(0.044)	(0.044)
95% interval	[−0.223, 0.309]	[−0.490, 0.861]	[−0.126, 0.068]	[−0.124, 0.071]
Clustered p-value	0.730	0.558	0.525	0.566
Markets	779	779	779	779
Grant months	6	6	6	6
Market-months	4,674	4,674	4,674	4,674

Notes: Coefficients on one-SD faster mean vacancy closure, with individual advertisement durations capped at 180 days. The SD is 9.72 days, calculated with equal weights using only the 4,674 retained pre-experiment market-months. Log outcomes are the natural logarithm of one plus each count. All columns use equal market weights, market effects and quadratic trends, city-by-month and country–field-by-month effects, and field-specific log baseline posting size interacted with month, fitted only on the retained months.

The estimates provide little evidence of a positive relationship, but the intervals permit meaningful associations. Six months provide limited variation after quadratic trends. These comparisons also cannot establish whether workers acquire or display credentials in anticipation of hiring: grant dates do not identify display dates, and capped advertisement durations can extend 180 days beyond posting.

X Conditional Pre-trend Tests

In this section, we explore whether there are pre-experimental differences in outcome trends between markets with high and low treatment exposure. We find that extending the pooled history to November 2020 retains 747 of 779 markets and does not jointly reject the conditional pre-period restrictions. Estimates are imprecise, and the longer U.S. series shows departures before multiple-testing adjustment.

Historical coverage. In the main vacancy analysis, we compare the experiment period, October 2022–March 2023, with the pre-experiment period, October 2021–March 2022. To examine earlier outcome trends, we use the preceding months in our existing Revelio extract, which begins in June 2020, and supplement it with records from June 2008 through May 2020. We process these additional records using the same occupation and market definitions, junior-vacancy eligibility criteria, and duration restrictions as in the main analysis, then aggregate them to the market-month level. This extension adds 1.07 million eligible postings. Only 5,585 date from 2008–2015, providing insufficient coverage for a continuous panel across our markets, so we focus on later years. Table A22 reports how many of the original markets have at least one eligible posting in every month from each proposed start date through January 2022.

TABLE A22
HISTORICAL COVERAGE OF THE CURRENT MARKET SAMPLE

First posting month	Months	Pooled	United States	Developing
Jan 2016	73	35	35	0
Jan 2017	61	139	139	0
Jan 2020	25	438	407	31
Jun 2020	20	475	424	51
Nov 2020	15	747	426	321
Jan 2021	13	769	426	343
Aug 2021	6	779	426	353

Notes: Numbers of original markets with at least one eligible advertisement in every month from the indicated start through January 2022. The original samples contain 779 pooled, 426 U.S., and 353 developing-country markets. Missing market-month durations are not set to zero. Eligibility preserves the existing occupation and metropolitan-area definitions, seniority categories 1 and 2, and duration screen of 0–184 days, followed by a cap at 180 days.

The sample changes as we extend the historical window because fewer markets have postings in every month. In the analysis that follows, we focus on two sets of windows: one beginning in August 2021 that includes all 779 pooled markets, and longer windows that retain at least 95 percent of the markets in each sample. These longer windows begin in November 2020 for the pooled sample,

January 2021 for developing countries, and January 2020 for the United States, retaining 747, 343, and 407 markets, respectively. We choose these dates based on data coverage. Going further back substantially reduces the sample: a window beginning in January 2017 includes only 139 U.S. markets and no developing-country markets.

X.1 Timing, Specification, and Inference

We test whether pre-experiment outcomes differ systematically between markets with high and low projected exposure after accounting for the fixed effects, baseline-size controls, and market-specific trends in the main specification. The model already allows markets to differ in their outcome levels and in their linear and quadratic trends. We therefore ask whether exposure predicts any additional monthly differences during the pre-period.

Using the notation of equation (1), we estimate

$$Y_{mt} = \alpha_m + \lambda_{c(m),t} + \mu_{k(m),f(m),t} + \gamma_{f(m),t}\ell_m + b_m\tau_t + q_m\tau_t^2 + d_m\delta_t + u_{mt}.$$

The specification includes market effects α_m, city-by-month effects $\lambda_{c(m),t}$, and country–field-by-month effects $\mu_{k(m),f(m),t}$. The term $\gamma_{f(m),t}\ell_m$ allows the association between baseline market size and duration to vary by field and month. Each market also has its own linear and quadratic trends, $b_m\tau_t + q_m\tau_t^2$. The coefficient δ_t measures the remaining exposure-related difference in month t.

To separate these monthly differences from the components already captured by market effects and trends, we impose

$$\sum_{t\in\mathcal{T}} \delta_t = 0, \qquad \sum_{t\in\mathcal{T}} \tau_t\delta_t = 0, \qquad \sum_{t\in\mathcal{T}} \tau_t^2\delta_t = 0,$$

where $\mathcal{T}$ contains the T months in the test window. We jointly test whether all δ_t equal zero, giving $T-3$ independent restrictions. A rejection indicates that the controls and trends do not fully account for exposure-related differences in pre-experiment outcomes.

We cluster standard errors by country in the pooled and developing-country samples and by metropolitan area in the United States. We jointly test whether all pre-period exposure coefficients are zero using the bootstrap sup-t procedure in Abadie et al. (2025). The test compares the largest absolute monthly t-statistic with its distribution under the joint null, obtained from a wild-cluster bootstrap. We reestimate the full model in each replication, using all 4,096 possible cluster-level sign patterns in the pooled sample, all 2,048 in the developing-country sample, and 4,999 draws in the United States.

This procedure accounts for testing multiple months and for dependence across their estimates. This approach performs well when there are many monthly coefficients but few clusters.

X.2 Joint Tests and Monthly Deviations

Table A23 presents the results. The joint tests do not reject the pre-period restrictions in the pooled sample: the p-value is 0.305 for the window retaining all markets and 0.405 for the extension beginning in November 2020. The developing-country extension gives a p-value of 0.555. For the U.S. extension, the joint test rejects at the five percent level ($p = 0.038$).

TABLE A23
CONDITIONAL PRE-TREND TESTS ACROSS HISTORICAL WINDOWS

Sample	Posting window	Markets	Months	Joint p-value
A. All markets in each original sample				
Pooled	Aug 2021–Jan 2022	779	6	0.305
Developing countries	Aug 2021–Jan 2022	353	6	0.523
United States	Oct 2020–Jan 2022	426	16	0.446
B. Earliest start retaining at least 95 percent of markets				
Pooled	Nov 2020–Jan 2022	747	15	0.405
Developing countries	Jan 2021–Jan 2022	343	13	0.555
United States	Jan 2020–Jan 2022	407	25	0.038
C. Posting cohorts after Experiment 1				
Pooled	Sep 2021–Jan 2022	779	5	0.145
Developing countries	Sep 2021–Jan 2022	353	5	0.283
United States	Sep 2021–Jan 2022	426	5	0.098
D. Later posting cohorts: follow-up may overlap the pilot				
Pooled	Aug 2021–Jul 2022	779	12	0.423
Developing countries	Aug 2021–Jul 2022	353	12	0.701
United States	Oct 2020–Jul 2022	426	22	0.393
Pooled	Nov 2020–Jul 2022	747	21	0.316
Developing countries	Jan 2021–Jul 2022	343	19	0.393
United States	Jan 2020–Jul 2022	407	31	0.061
E. Additional U.S. history				
United States	Jan 2017–Jan 2022	139	61	0.207
United States	Jan 2020–Nov 2020	407	11	0.413

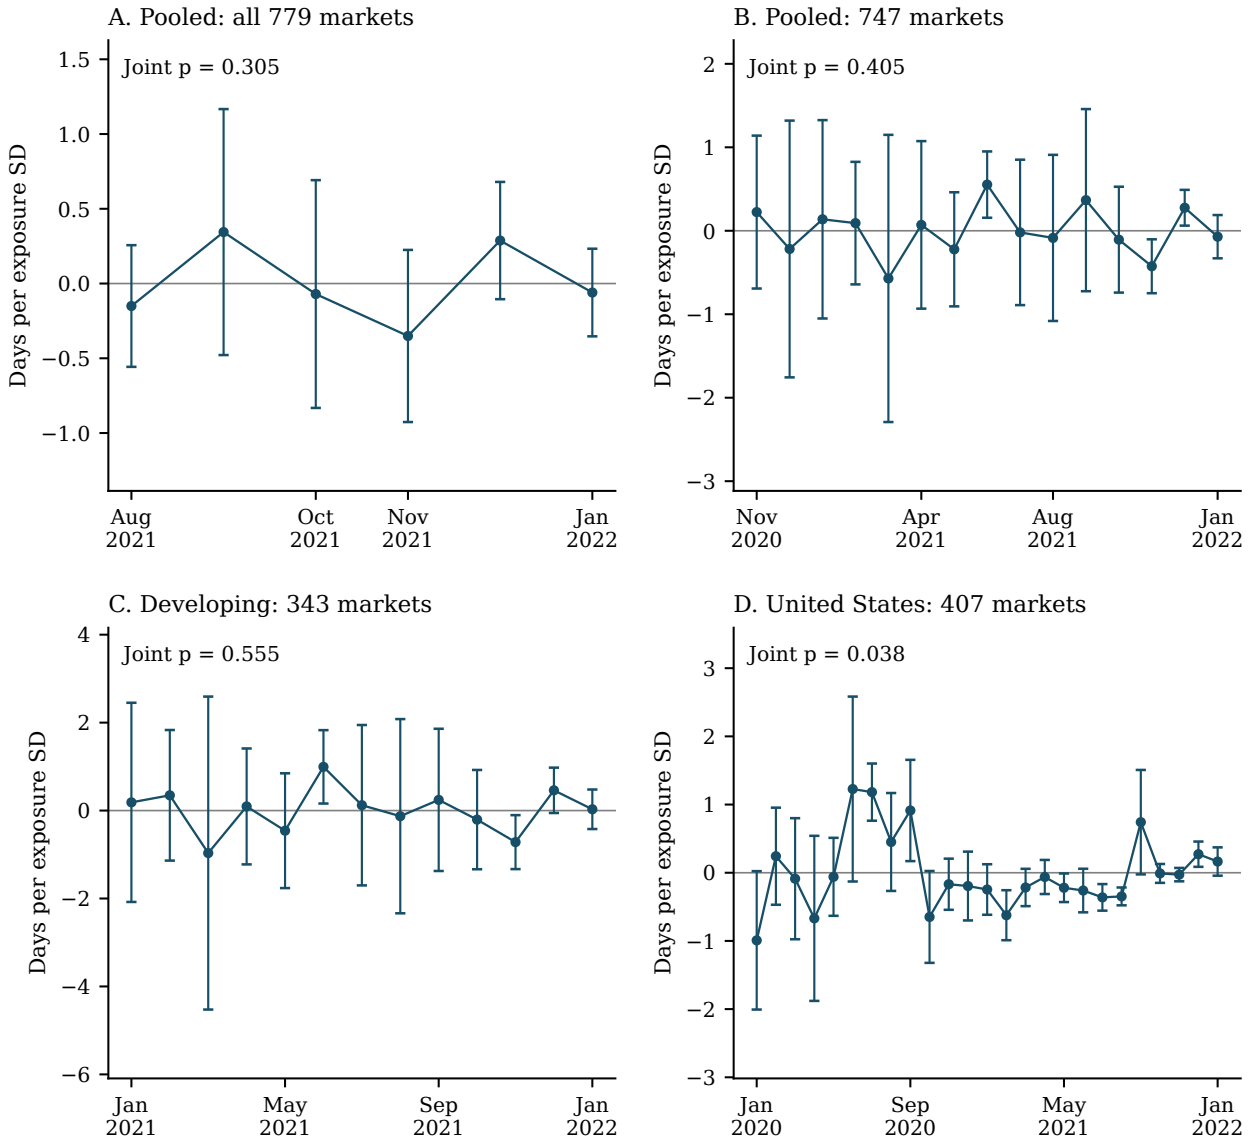
Notes: All specifications include market effects and quadratic market trends, city-by-month effects, country–field-by-month effects, and field-specific log baseline posting size interacted with month. Markets receive equal weights, and exposure retains its original scale. "Joint p-value" reports the bootstrap sup-t test of the joint null that all identified monthly exposure coefficients are zero, conditional on these controls and trends. Standard errors are clustered by country in the pooled and developing-country samples, with twelve and eleven clusters, respectively, and by metropolitan area in the United States, with 107 clusters except in the 2017-start sample, which has 67. In Panels A–C, duration follow-up ends before the August 2022 pilot, although the longer windows can include Experiment 1. Panel D extends each earlier window to include posting cohorts whose follow-up may overlap the pilot, holding its market sample fixed. In the second row of Panel E, follow-up ends before Experiment 1.

Figure A5 shows the monthly coefficients and pointwise confidence intervals. These intervals describe uncertainty for each month separately; the sup-t test assesses the monthly coefficients jointly.

In the pooled and developing-country samples, the coefficients fluctuate around zero without a clear pattern. The U.S. estimates are more precise and show a more pronounced pattern: the coefficients turn negative after May 2021 and become positive around January 2022. These movements indicate exposure-related differences in pre-period outcomes remaining after adjustment for the controls and market-specific trends.

Restricting the window to September 2021–January 2022 gives joint p-values of 0.145, 0.283, and 0.098 for the pooled, developing-country, and U.S. samples, respectively. This short window leaves only two independent monthly contrasts after accounting for quadratic trends. The later windows in Panel D do not reject jointly at the five percent level, although duration follow-up may overlap the pilot. Extending the U.S. history back to January 2017 also yields no joint rejection ($p = 0.207$), but this comparison includes only 139 markets and provides no corresponding evidence for developing countries.

Figure A5: Exposure-related Deviations from Pre-period Quadratic Paths



Notes: Monthly mean-duration deviations per original-sample SD of projected exposure, after all controls in equation (1). Bars are pointwise 95% confidence intervals.Panel A uses August 2021–January 2022; Panels B–D use the earliest starts retaining at least 95% of each original sample, all ending January 2022. Exposure SDs are 20.48 in pooled, 38.09 in developing-country, and 1.29 in U.S. markets, in expected encouraged learners per 100 baseline advertisements. Vertical scales differ. The long windows include Experiment 1 and changes in posting sources.

Y Time Periods for Vacancy Analysis

TABLE A24
TIME INTERVALS USED IN SECTION 5

Component	Interval or cutoff	Role in the analysis
A. Calendar windows and market assignment		
Vacancy baseline	Oct. 2021–Mar. 2022 (6 months)	Baseline outcomes, exposure denominator, posting weights, and log baseline size; also the baseline for the April–July placebo.
Vacancy experiment period	Oct. 2022–Mar. 2023 (6 months)	Main comparison, fitted zero-exposure outcomes, and both placement benchmarks. Matches the baseline's recruiting season.
Posting coverage	Oct. 2021–Mar. 2023 (18 months)	Require activity in every month. Apr.–Sep. 2022 enter this screen but not the main duration regression.
Experimental population for exposure	Sep. 2022–Mar. 2023 (7 cohorts)	Experiment 2 learners, including the September pilot, projected through the control-based destination map.
Pre-experiment job used for location	Job began by May 2022; not known to have ended before June 2022	Latest recorded pre-June job locates controls and credential holders, subject to these cutoffs. Assignment stays fixed across cohorts.
Credential visibility: baseline	Sep. 2021–Apr. 2022, excluding Jan. 2022 (7 months)	Fit the pre-period trend in the log ratio of experiment-course Coursera to other-provider credentials.
Credential visibility: post period	June 2022–Feb. 2023, excluding Jan. 2023 (8 months)	Average deviations from the baseline trend. May 2022 enters neither window; January omissions limit grant-date bunching.

TABLE A24
TIME INTERVALS USED IN SECTION 5 (CONTINUED)

Component	Interval or cutoff	Role in the analysis
B. Placebos and historical diagnostics referenced in Section 5		
April–July placebo	Apr.–July 2022 versus Oct. 2021–Mar. 2022 (10 months total)	Exclude August because preliminary piloting had begun. Earlier postings can remain open into piloting (Appendix V).
Pre-experiment credential–closure association	Oct.–Dec. 2021 and Feb.–Apr. 2022 (6 months)	Grant and posting cohorts before June 2022, retaining January and May omissions. No later month enters estimation or standardization (Appendix W).
Historical tests: all original markets	Aug. 2021–Jan. 2022, pooled and developing; Oct. 2020–Jan. 2022, U.S.	Earliest balanced windows retaining every market in each sample. Ending in January keeps 180-day follow-up before the August 2022 pilot.
Historical tests: at least 95% of markets	Nov. 2020–Jan. 2022, pooled; Jan. 2021–Jan. 2022, developing; Jan. 2020–Jan. 2022, U.S.	Retain 747, 343, and 407 markets. These longer histories can include Experiment 1 (June–Aug. 2021).
Historical timing sensitivities	Sep. 2021–Jan. 2022; alternatively, extend the preceding historical windows through July 2022	The first excludes Experiment 1 posting cohorts but cannot exclude persistent effects. July-ending windows can include durations extending into piloting.
Additional U.S. history	Jan. 2017–Jan. 2022; Jan.–Nov. 2020	The first uses 139 continuously covered markets; the second keeps 180-day follow-up before Experiment 1.
Historical six-month placebo	Aug. 2021–Jan. 2022 versus earlier months in each 95%-coverage window	Earlier comparison months: Nov. 2020–July 2021 pooled; Jan.–July 2021 developing; Jan. 2020–July 2021 U.S.

Notes: Calendar endpoints are inclusive, subject to the stated omissions. Credential windows refer to grant dates, not dates of profile addition or encouragement; the June 2022 cutoff does not date intervention delivery. Vacancy windows refer to posting dates, not removal dates. Panel B summarizes supporting analyses; detailed historical windows appear in Table A23.