Journal of Development Economics

Job Seekers' Beliefs and Labor Market Demand

--Manuscript Draft--

Manuscript Number:	DEVEC-D-25-00784R1
Article Type:	Registered Report Stage 1: Proposal
Section/Category:	Experimental Papers, credit, insurance
Keywords:	Labor Market, Job Search, AI
Corresponding Author:	Zachary Kuloszewski The University of Chicago Chicago, Illinois UNITED STATES
First Author:	Erin M Kelley
Order of Authors:	Erin M Kelley
	Zachary Kuloszewski
	Gregory Lane
	Harry Moroz
	Mohit Negi
	Evelyn Vezza
Abstract:	Job seekers often face challenges in finding jobs that align with their skills and meet employer demands. To address this, we conduct a randomized controlled trial to evaluate whether artificial intelligence can assist job seekers in two key ways: 1) gaining a clearer understanding of their own skills, and 2) identifying careers that align with their skills and labor market demand. This improved understanding should enable job seekers to make more informed decisions about training programs and career paths, resulting in better labor market outcomes -- such as securing jobs more quickly, staying employed longer, and achieving greater job satisfaction.
Response to Reviewers:	Please see the attached "Response to Reviewers" Document for this re-submission for itemized responses to all reviewer comments.

Job Seekers' Beliefs and Labor Market Demand

Pre-analysis plan

Erin M. Kelley, Zachary Kuloszewski, Gregory Lane,
Harry Moroz, Mohit Negi, Evelyn Vezza*

August 2025

Abstract

Job seekers often face challenges in finding jobs that align with their skills and meet employer demands. To address this, we conduct a randomized controlled trial to evaluate whether artificial intelligence can assist job seekers in two key ways: 1) gaining a clearer understanding of their own skills, and 2) identifying careers that align with their skills and labor market demand. This improved understanding should enable job seekers to make more informed decisions about training programs and career paths, resulting in better labor market outcomes – such as securing jobs more quickly, staying employed longer, and achieving greater job satisfaction.

Keywords: Labor Market, Job Search, AI
JEL Codes: D81; D25; O12; O13; Q12; Q54

*Kelley: Harris School of Public Policy, University of Chicago. Email: erinmkelley@uchicago.edu. Kuloszewski (corresponding author): Kenneth C. Griffin Department of Economics, University of Chicago. Email: kuloszewski@uchicago.edu. Lane: Harris School of Public Policy, University of Chicago and NBER. Email: laneg@uchicago.edu. Moroz: World Bank. Email: hmoroz@worldbank.org. Negi: Harris School of Public Policy, University of Chicago. Email: mohitnegi1442000@gmail.com. Vezza: World Bank. Email: evezza@worldbank.org. We are grateful to Dave Mckenzie for helpful comments and suggestions. We thank Simon Schmid and Lina Solera from SkillLab for an outstanding collaboration. We thank Fernanda Oliveria, Uttkarsh Dandanayak, Shivant Krishnan for excellent research assistance. We thank the Becker Friedman Institute for Economics at the University of Chicago, the Agricultural Technology Adoption Initiative at J-PAL, the King Climate Action Initiative at J-PAL, and the World Bank for generously funding this project. This research has IRB approval from the University of Chicago (Protocol No. IRB25-0369), and is registered on the AEA RCT registry (Identification No. AEARCTR-015801). All remaining errors are our own.

1 Introduction

Finding a job, let alone the right job, can be difficult, which is why governments and international organizations spend millions of dollars each year to assist job seekers. A wide variety of programs have been tried, including job-training initiatives to help individuals gain new skills and experience, and labor intermediation services that equip workers with tools to improve their job search and connect them to employment opportunities. Although some of these initiatives have positively impacted employment outcomes and wages, many are expensive and difficult to implement logistically (Carranza and McKenzie, 2024a). Advancements in technology have allowed companies to harness AI to support job seekers with the information they need to navigate the labor market more effectively. This includes information that will help individuals better understand their skills, explore vocational training options, identify potential career matches, and gain insights into the labor market. While theory suggests these informational services can significantly improve job seekers' labor market outcomes, there is limited empirical evidence to support this—and even less insight into which types of information are most relevant and how they complement one another.

This research examines whether AI tools can help users 1) gain a better understanding of their skills, 2) identify the vocational training courses and careers they are best suited for; and 3) recognize which careers are in demand. To this end, we are partnering with the Secretariat of Employment and Vocational Training (SEFP) within Argentina's Ministry of Human Capital, along with the World Bank, to assess which of the three services has the greatest impact on job-seeker outcomes, and how they interact.

SEFP currently provides beneficiaries with a range of labor market services with the aim of improving skillsets tailored for quality employment in the formal sector. This includes access to vocational training programs which users can access through an online portal, where hundreds of courses are listed alphabetically. Users must browse through the list and choose the course they believe suits them best. However, discussions between the World Bank and SEFP revealed two key challenges: many users were unsure which courses matched their skillsets, and they lacked information about which courses aligned with current labor market needs. To address this, the World Bank and SEFP partnered with SkillLab, a labor market intelligence provider that uses artificial intelligence to help users build a personalized skills profile. The platform then recommends training programs and careers that best match their abilities and are in demand.

We embed a randomized controlled trial (RCT) into this initiative to evaluate how access to information 1) on skills, 2) course and career matching, and 3) labor market demand influences job seekers' ability to find employment. Specifically, SEFP program beneficiaries

are assigned to one of three versions of the SkillLab AI platform and a comparison group. The first version (SkillLab CV) enables users to input their labor market experiences, identifies their skills, and generates a tailored CV. The second version (SkillLab "Recommended Careers/Courses") includes an additional feature that recommends careers/courses tailored to the user's skill profile. The third version (SkillLab "In Demand: Careers/Courses") introduces a new feature that recommends careers and courses tailored to the user's skill profile and that align with current labor market demand. Finally, users in the comparison group will not receive access to SkillLab in 2025. All program beneficiaries— treatment and control—continue to have access to SEFP's standard portal, where vocational training courses are listed alphabetically. The only difference between the groups is that some also receive access to a version of the SkillLab platform. To measure the impact of the intervention, we will conduct two phone surveys— approximately 3 and 9 months after participants access the platform. These surveys will cover skills knowledge, job search behavior, reservation wages, employment status, and mental health.

The broad goal of SEFP and Fomentar Empleo is to address structural skill mismatches by offering skills training in key areas. A recent World Bank technical report (Vezza et al., 2025) documents rising structural obsolescence in Argentina – a mismatch between the skills employers demand and those workers possess – driven by the shift from industry to skill-intensive services and digitally mediated roles, which raises the cognitive and technological demands even in entry-level positions. We aim to complement this effort by addressing informational mismatch in the labor market – where job seekers often misidentify, understate, or inadequately signal their existing skills. We do so by guiding users through a structured recall of their prior experiences, mapping those experiences to an internationally standardized taxonomy of skills, and then matching the resulting skill-profile to in-demand occupations and training opportunities. The three treatment arms progressively add (i) articulation, (ii) algorithmic matching, and (iii) high labor-demand indicators, allowing us to isolate which type of information most effectively narrows skill gaps.

This paper makes four contributions. First, we contribute to the emerging literature on workers' awareness of their own skills. In closely related research, Kiss et al. (2023) measure job-seekers' skills and randomly provide some participants with information about how their skills compare to others. This intervention leads to significant improvements in job search and employment outcomes. We extend this research by examining how workers' beliefs about their skills interacts with information about labor market demand. By providing job seekers with information about their skills and labor market trends, we aim to determine whether they can more precisely identify and hone a set of skills that will lead to better job outcomes.

Second we contribute to an active literature on boosting the impact of vocational training

programs. As discussed in (Carranza and McKenzie, 2024b), most randomized experiments evaluating the impact of vocational training programs find relatively modest effects. These impacts are especially muted for large-scale programs implemented by governments. One policy recommendation has been to make training programs more demand-driven. Preliminary evidence suggests this approach could be effective. A qualitative review of government programs by Sanchez Puerta et al. (2015) highlights the success of demand-driven skill training initiatives across various countries. Additionally, Alfonsi et al. (2020) examine vocational training programs that offered courses aligned with pre-specified high-demand occupations, and their study is one of the few to show significant and lasting impacts. We contribute to this literature by explicitly testing how providing job-seekers with information about in-demand occupations influences their course selections and labor market outcomes.

Third, we contribute to the growing body of research examining the role of workers' beliefs about the labor market. Bandiera et al. (2021) examine impacts of a cross-randomization of vocational training and matching, and find that outcomes depend heavily on worker beliefs and how closely they align with true levels of demand. Unlike much of this literature, we focus on providing job-seekers with real-time information about in-demand occupations, in addition to information about their own skill set. This approach helps job-seekers better understand what the labor market values, enabling them to make more informed decisions about how their skill-set aligns with market demands.

Lastly, we contribute to a recent literature which attempts to provide tailored recommendations and demand-side insights to job-seekers to aid in their search. This literature aims to find information-based solutions to reduce search and matching costs at scale, such as Belot, Kircher, and Muller (2019), who show that skill-based job suggestions broadens the considered set of occupations for the recently unemployed. Bächli, Lalive, and Pellizzari (2025) show limited impacts of skill-profile and past-experience based recommendations in Switzerland, while Leduc and Tojerow (2025), who provide information about "shortage" occupations in Belgium, find this increases uptake of vocational training programs which teach general skills but does not impact employment. We build upon this literature by separately identifying the impact for skill-based and demand-based recommendations. We expand upon previous recommendation interventions by linking multiple types of information, measuring the impact of helping the job seeker triangulate skills, learning opportunities, and available vacancies. Lastly, we examine impacts of a large-scale application of these recommendations in a middle-income country context, where equilibrium access to labor information is low.

The remainder of this pre-analysis plan proceeds as follows. Section 2 provides relevant details about the research setting, including information about SEFP's existing program (Fomentar) and SkillLab. Section 3 describes our experimental design. Section 4 discusses

the data we will collect over the course of our experiment. Section 5 presents our planned analysis.

2 Context

2.1 Fomentar

The sample we work with is enrolled in two government labor market programs: Fomentar Empleo (Fomentar) and Volver al Trabajo (VAT). Fomentar is a labor market program that was launched by the Secretariat of Employment and Vocational Training (SEFP) in September 2021 to improve employment opportunities for vulnerable job seekers. VAT is a transitional program that helps former beneficiaries of the older labor market program, Potenciar Trabajo, move to Fomentar. Fomentar and VAT offer a range of services to enhance employability and job placements.[1] These services include career counseling, apprenticeships, employment mediation, and an employment subsidy that reduces employer contributions. Of particular interest to this study is their offering of vocational training programs through an online portal (Portal Empleo) that lists all courses alphabetically. However, discussions with the World Bank and SEFP revealed that users engage with the portal less than expected and often struggle to identify which courses align with their skills or meet labor market demand.

Fomentar is expected to reach 865,000 people by 2028. The program targets individuals aged 18 to 64 who are unemployed, have not had formal employment in the last three months, and are not enrolled in any other social programs. Priority is given to youth (18-24), women (25-59), older men (45-64), and ethnic/racial minorities. Table 1 presents key descriptive statistics for the combined universe of Fomentar and VAT participants (1,048,574 individuals). A majority (63%) are female, and the average participant is 32 years old. Nearly half (47%) have young children. While most participants have completed primary education (91%), substantially fewer have completed secondary education (45%). Finally, 41% have prior work experience. Figure 1 shows the geographic distribution of program participants, with a higher concentration in the northern regions of the country.

2.2 SkillLab

SEFP and the World Bank have collaborated on the Fomentar/VAT program for over five years. One key challenge they identified was helping beneficiaries recognize their skills and better target vocational training and career opportunities. To address this, they partnered

[1] To encourage participation, the Ministry provides a conditional stipend of 73 USD per month for priority groups.

with SkillLab in 2022 to give beneficiaries access to SkillLab's AI-powered career orientation services. The app includes several key features:

Skills Generation/CV The app helps job seekers identify their skills by guiding them through a series of questions based on their past jobs, education, and other experiences, such as those shown in Figure A1. For each job or educational experience a user adds, SkillLab suggests relevant skills by matching them to the ESCO taxonomy – a multilingual classification of European Skills, Competences, and Occupations. Related skills are suggested four at a time to avoid overwhelming users (users can add more if needed). SkillLab also provides clear definitions of each skill, and offers translations of the ESCO skill definitions into Latin American Spanish to ensure they are easily understood. Users then review these suggested skills, selecting the ones they have used and indicating how frequently they applied each skill on a Likert scale from "rarely" (1) to "always" (5). This feature is displayed in Figure A2, where the user is selecting skills linked in the ESCO taxonomy to a job at a local car wash. Once users have finished entering in their experiences, they have the option of generating a CV which displays standard education and work experience information, but also the identified skills associated with each experience. There is a feature that enables users to customize their CV for specific target career. In practice, it simply adjusts the order in which skills are listed within each experience in the final CV document.

Career recommender: The app recommends careers to users. SkillLab's algorithm compares a user's identified skills – based on their selected experiences and associated skills – with the skill requirements of a specific career. Using the ESCO taxonomy, each career is linked to a set of skills categorized as "essential" or "optional." A career is considered a match if the user possesses at least 50% of its essential skills. Once matched, careers appear randomly under a carousel called "Careers You Might Like".

The guiding principles for the recommendations vary by treatment arm. By default in Treatment 2, career recommendations are based solely on skill matching. However, users can indicate interest in a career by clicking the heart symbol next to it. As users "like" careers, the recommendations in the browse-able carousel adjust based on a proprietary algorithm that considers both skill matching and user preferences. The number of careers shown depends on the user's profile, with a maximum of six careers displayed in the carousel. In Treatment 3, career recommendations are provided in two browse-able carousels. The first resembles the carousel available to users in Treatment 2, with extra weight added on to careers we determine to be "in-demand," to be discussed at the end of this section. The second carousel, labeled "Careers in demand" orders in-demand careers in order of decreasing skill-match.

Figures A3 and A5 show examples for the course recommendation page in Treatment 2 and 3, respectively.

Users can also explore all available careers in the "All Career Fields" section to assess their compatibility with these careers. Careers for which there are a skill-match are also visually tagged in this extensive browsable page. By clicking on a career, the user sees a brief description of the career and skills (both essential and optional) which are attributed to that career. The user can also see courses which are tagged to this career, displayed in random order. Figure A8 shows examples of these pages.

Course recommender: The app also recommends training courses to users from the courses available through Fomentar. It links careers to courses by using a Large Language Model (LLM) to: 1) analyze the text of each course, identifying key phrases and skills and 2) generate a similarity score based on how closely the course description aligns with ESCO careers. For example: the course might refer to "conveying ideas" or "public speaking", and those will be linked to "communication", a skill that will have to be identified in a career for there to be a match. A separate Natural Language Processing algorithm tags the course with individual skills by searching course description text for shorter key words and phrases. The result is a map from each course to a set of ESCO skills and careers. SkillLab's recommendation algorithm weights recommendations between courses which, according to these mappings, are related to the careers the user has shown interest in via likes and courses for which the user has related skills.

In Treatment 2, courses are recommended using a weighting of user career interests and the user's skill-set. If the user hasn't liked any careers, the recommendation algorithm looks for courses that teach skills the user has rated below a "3" (the mid-point between "rarely" and "always" on the skill frequency scale). This is known as "boosting" the user's skills. If there are no boostable skills or stated interests, the recommender prioritizes courses that teach skills the user already has. While the algorithm doesn't specifically aim to expand the user's skill set with related skills, this naturally happens because courses typically teach groups of related skills. In Treatment 3, courses are recommended using a weighting of three distinct factors: interests, skill-alignment (identical to Treatment 2), and additionally demand data. Figures A4 and A6 gives an examples of the course recommender in T2 and T3, respectively.

Users can click on recommended courses and see why they were suggested, whether based on career interests or skill boost potential. On each course page, the user will find a course description and a list of skills associated with the course. These skills are divided into two categories: new skills and existing skills the user can "boost." The page also shows the career

that the course is linked to. Figure A7 shows an example of all of the course information available to the user about the course in-app.

In-Demand: We collaborated with SkillLab to develop a new app feature that highlights in-demand occupations. To identify these occupations at a national level, we use data from Lightcast—a company that aggregates online job postings across industries and platforms. Lightcast's database for Argentina includes 3 million job postings from 2020 to 2025. We receive regular updates of new postings, which are then mapped to 5-digit ESCO occupations. We classify an occupation as in-demand if it meets *either* of the following two criteria:

1. Occupations with high levels and non-negative growth: This includes jobs with > 5000 listing (97th) percentile and with growth rates that exceeds 0.

2. Occupations with high growth rate and important levels: This includes jobs with > 1000 listings and with growth that is above the median (37%)

Our cutoff rule aligns with established methods for identifying high-demand occupations, such as the O*NET "Bright Outlook Occupations" list. This approach balances both current job market demand and growth potential, ensuring that all recommended occupations have a sufficient level of current demand as well as positive future prospects. A visual illustration of this selection process is presented in Figure 3. We set our thresholds using data from the past two years to capture trends while smoothing out short-term fluctuations. We decided against a one-year window since labor demand in 2024 was unusually volatile due to government disruptions.

We construct in-demand measures based on nationwide job-postings data, aggregating these postings to careers at the 5-digit ESCO classification. We present in Figure 1 the spatial distribution of Fomentar Participants and in Figure 2 and the spatial distribution of LightCast job postings used to create our demand metrics, and argue that these correspond accordingly to give users in Treatment 3 nationally representative measures of labor market demand. While this may not give accurate indications of labor market demand for every locality, these signals are an improvement on ex-ante information accessible to the target population, who currently lack any reliable information on real-time labor demand. Demand indications at the 5-digit ESCO level also align more closely with how employers post vacancies, giving job-seekers precise information. We update this classification as new Lightcast data is released. SkillLab displays this information through an additional carousel on its platform that highlights in-demand occupations, ranked from the highest to lowest skill match with the user, shown in Figure A5.

User Interaction: The success of the SkillLab intervention hinges heavily on whether the users understand the skill-discovery process, provide quality inputs, can interpret and understand the information in the app, and receive recommendations which are relevant and actionable.

A first concern is that users may not provide quality inputs for the recommendation algorithms to base recommendations on. A key contribution of SkillLab is the experience-based skill elicitation which compartmentalizes the process of CV building into manageable steps for users. Skill suggestions are presented no more than 4 at a time, with descriptions in plain language. Further, the fact that skills are suggested through links to the experience limits the likelihood of users misrepresenting their skill-set. While users may manually add skills not linked to an experience through ESCO, this is rare–in our March/April 2025 pilot, only 4.9% of skills were added manually, with the vast majority (95.1%) coming from ESCO-based recommendations. Users also cannot create custom or self-defined skills. Finally, the data indicate that few users consistently rate themselves at the maximum level (5) across the skills that are presented to them, suggesting they are not systematically overstating their skills. Thus, we believe the platform effectively extracts real skill information from users.

The intervention may also be minimally effective if users are presented with recommendations which are only aligned with interests but not their existing skill-set, leading the user to pursue paths for which they are not well equipped. SkillLab's design does a great deal to mitigate this concern. Their recommendation algorithms always places positive weight on skill-matched careers, even when the user expresses interest in multiple careers for which they have no skill overlap, ensuring that there are always some skill-matched recommendations. Additionally, pilot data show that 70% of careers liked by users have some amount of skill-overlap, meaning the user already has at least one essential skill for that career. This assures us that users are expressing interest primarily in careers relevant to their skill-set, further limiting the likelihood of irrelevant recommendations. It is worth mentioning that only 19% of liked careers have over a 50% skill match, indicating that users are also employing likes for careers related to their skills but that still require further skill acquisition. We believe this imperfect skill-match with interests can enable users to get information on more aspirational careers and may inspire engagement in relevant training for skills they have some exposure to but need to further develop, leading to better outcomes ((Orkin et al., 2023) find that boosting aspirations increases labor supply in Kenya).

For the in-demand feature, demand-based recommendations may create disappointment effects which can actively suppress job-seeker up-skilling if they are not matched to in-demand careers or related courses. SkillLab carefully designed the application to minimize the risk of disappointment effects. Displayed in Figure A5, the "Careers" tab displays recommendations

in two carousels, with varying importance placed on whether a career is in-demand[2]. Thus, a user without skills relevant to any of the in-demand careers will still get suggestions that are relevant to them in the first, which mitigates any disappointment they might feel from seeing the second. Similarly in the "Learn" tab (shown in Figure A6), users without skills matching in-demand careers will still receive skill-based recommendations alongside demand-based recommendations, which we believe will mitigate the disappointment effect in this portion of the application.

Lastly, the platform would be ineffective if the set of in-demand careers were generally out of reach of the target population. The categorical distribution of occupations we label as in-demand is presented in Figure 4. We argue that this set of careers represent a set of careers which are both relevant and attainable for the target population. We see overlap in pilot data between skills matched to in-demand careers and skills already held by users, reported as off-axis data points in Figure 5. We also report the distribution of skills according to their mappings to all occupations, in-demand occupations, and users in Figure 6. We find that 5% of skills are matched to both users and in-demand occupations, indicating that in-demand careers are indeed relevant to some users. Another 18% of all skills are linked to in-demand careers but not to any users, indicating room for the recommendations to push users to acquiring *new* skills.

3 Experimental design and Implementation

3.1 Design

We design the randomized control trial to answer five specific research questions:

1. Does providing information to jobseekers about 1) their skills affect labor market outcomes?

2. Does providing information to jobseekers about 1) their skills; and 2) careers and courses that match their skills/preferences affect labor market outcomes?

3. Is information about relevant career and courses additionally beneficial for job seekers who have an understanding of their own skills?

4. Does providing information to jobseekers about 1) their skills; 2) careers and courses that match their skills/preferences; and 3) careers and courses that are in demand affect labor market outcomes?

[2]See Section C.2 for a more detailed description of the recommendation algorithm

5. Is information about careers and courses that are in demand additionally beneficial for job seekers who already have information about their own skills and relevant course/career matches?

The following schematic outlines the evaluation design. It is important to note that all program beneficiaries— treatment and control—continue to have access to SEFP's standard employment services and standard Fomentar portal, where vocational training courses are listed alphabetically. The only difference between the groups is that some also receive access to a version of the SkillLab platform.

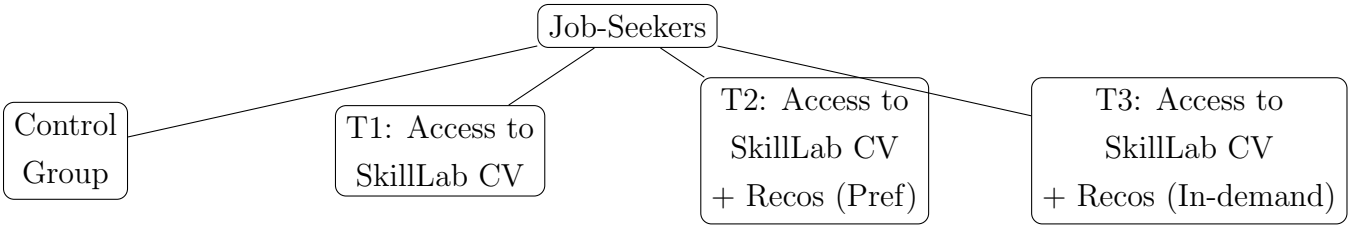

- C: Control Group
  - Will not receive access to SkillLab in 2025
- T1: Access to SkillLab CV
  - Will receive an email giving them access to SkillLab CV.
  - SkillLab CV allows the user to go through the skills documentation/discovery process by inputting their past education, jobs, and non-employment activities, creating a market-ready CV document in the process.
- T2: Access to SkillLab CV + Recommendations for Training/Careers
  - Will receive an email giving them access to SkillLab "Full".
  - These participants will be given access to the "full" SkillLab application, which includes the SkillLab CV functionality and also "Careers" and "Learn" functions that allow beneficiaries to explore career paths and training opportunities. All recommendations for trainings and career paths will be based on SkillLab's methodology detailed above.
- T3: Access to SkillLab CV + Recommendations for Training/Careers + In Demand
  - Same as T2.

– Additional recommendations for training or careers identified as "in-demand" based on the methodology outlined above.

3.2 Implementation

Implementation The research team, in conjunction with the World Bank and the SEFP, will randomize Fomentar and VAT participants into the four groups listed above. We conduct the randomization at the individual-level. We discuss implications for spillovers (which we think are limited) in the final section of the PAP.

In total, the research team will send out 500,000 invitations, representing 47% of the sum total of beneficiaries enrolled in Fomentar Empleo and Volver a Trabajo. Due to the differences in beneficiary eligibility across the two programs discussed above, we stratify randomization across enrollment in Fomentar vs VAT, by above vs. below median age, and across municipality. Where possible, the research team will work with the SEFP local offices present in each municipality (which they call employment offices) to provide additional training to their staff about the SkillLab app, should any job-seeker come to the office with questions.

Before launching the full experiment, the research team piloted two of the treatment arms in 10 randomly selected municipalities to test the implementation strategy and assess response rates. The pilot revealed that some employment offices are more responsive and engaged than others—supporting our decision to stratify by municipality, since each municipality corresponds to a specific employment office. Based on the pilot, we observed take-up rates of approximately 10%, which we use for the power calculations described below.

Timeline The full-scale will be launched starting in May 2025, pending final approvals from the Sub-secretary. The full timeline for the evaluation is as follows:

- **End of April 2025:** Sub-secretary designates 10 Employment Offices for early launch.
- **May 1-7:** Invitations for T1 (CV), T2 and T3 (SkillLab Full and Full + Demand) are sent to early participants.
- **July 2025:** Invitations for all treatments are sent to the remaining participating Employment Offices
- **August-October 2025:** Email notifications of midline surveys are sent to all control and treatment participants. Survey firm will then contact a random subset (5000) participants to conduct the survey via telephone.

- **January-March 2026:** Email notifications of endline surveys are sent to all control and treatment participants. Survey firm will then contact the same set of participants via telephone as midline.

4 Data Collection

We have three sources of data: data from SkillLab, the Ministry and our own Survey data. We detail our uses of each of these data sources in this section.

4.1 SkillLab

First, we have all of the data that SkillLab collects from its users. This includes the data that users enter about their skills, their usage (clicks and views), their preferences (which careers and skills they explicitly like), and which training courses they navigate to and click to enroll in, all with time-stamps. This allows us to establish the following:

1. Whether career interests are aligned with existing skills
2. Whether viewed/liked courses are aligned with existing skills
3. Whether career/course interests are aligned with in-demand careers

We intend to track how these different measures of alignment evolve in treatment arms with recommendations (Treatments 2 and 3). We also receive data on whether users export or download their CV and when they initiate a new session in the app.

We will also be able to observe some self-reported measures of skill alignment through an in-app survey offered to users in treatment arms 2 and 3. The survey asks users to respond on a Likert scale from 1-5, indicating how strongly they agree with the statements "I feel confident that I know which skills I need for the career(s) I'm interested in" and "I feel confident about my ability to get a job in the career(s) I'm interested in." Thess questions are asked after users have already added a minimum number of skills to their profile to be deemed "complete." Thus, we will be able to qualitatively gauge if users think they are being informed and led in a direction which will ultimately lead to more gainful employment.

4.2 Ministry of Human Capital

The Ministry of Human Capital will provide socio-demographic data on program beneficiaries at the time of registration, along with their employment status in the year prior to registering. At the end of the experiment, they will share data on whether beneficiaries

completed any training courses and whether they appear in the employment registry as formally employed, including their wage and sector of employment. This administrative data includes the location and modality of the course, as well as the firm location of any formal job. Additionally, they will provide information on whether users updated their CVs on the Portal Empleo.

In combination with the SkillLab data we collect, we can evaluate the extent to which the application empowers users to achieve "skill-laddering" where users gain skills beyond initial course/career recommendations. For example, a user completing an introductory course in a trade may realize they also require skills offered by an intermediate course in the same field, leading to more advanced skill acquisition. Course enrollments and completions are visible through Ministry data, while skills acquired and added to profiles are visible through SkillLab app data.

4.3 Survey Data

To evaluate the program's impact along other margins, the team plans to conduct midline and endline surveys, timed strategically to capture key labor market dynamics. The survey will be sent out via email to treatment and control beneficiaries in the sample. We will also hire a survey firm to conduct 5000 surveys over the phone. The midline survey will take place about three months after the treatment group gains access to the portal and one month after those who enrolled in training programs have completed them. This timing enables the team to observe any immediate changes in participants' outcomes. The endline survey will follow approximately six months after the midline, allowing the team to evaluate whether the midline effects are sustained over time or if certain outcomes, such as employment, emerge later (e.g., it may take beneficiaries up to a year to secure a job).

The surveys will measure five key outcomes: 1) skills knowledge and career aspirations, 2) job-search behavior, 3) reservation wages; 4) employment outcomes; and 5) mental health. We hypothesize that receiving targeted skills and career recommendations from the AI portal can significantly impact these outcomes. First, participants using the portal may gain a better understanding of their skills, and develop different career aspirations, which may be more aligned with interests and/or areas of high demand. Second, the portal is designed to encourage job-seekers to focus on searching for roles that match their qualifications and aspirations. Third, access to the portal is designed to help students pursue more targeted skills training, leading to better employment outcomes and greater satisfaction with their eventual jobs. Finally, improved understanding of skills, greater control over the job-search process, and enhanced employment prospects could positively influence job-seekers' mental

health. We provide details on each set of outcomes below:

Skills Knowledge and Career Aspirations

- Please state your top 3 skills.
- Please state your top 3 most desired skills.
- Think of yourself in five years, in which occupation do you *expect* to be?
- Think of yourself in five years, in which occupation *do you want to* be?

Job-Search

- How many hours did you spend searching for a job in last week?
- How many jobs did you apply for in the last month?
- Which of the following methods have you used to apply for a job in the past month? [job portal, government office, ...]
- Think about the last job you applied to or considered applying to. In which of the following occupational categories was this job?
- In the last month, have you received any direct assistance from an Oficina de Empleo in your job search? If so, how? (helped with construction of CV, helped find vacancies, helped prepare for interview, etc.)

Reservation Wage

- What is the minimum monthly wage that you would be willing to accept to take a job 30 minutes away?
- What is the probability you can get a job in 3 months? Use a 0-10 scale for probabilities
- What is the the maximum amount you expect to earn per month?
- What is the the minimum amount you expect to earn per month?
- What is the probability you think you can earn at least the average (as calculated from the maximum and the minimum)

Employment

- In the last year, how many months were you employed (please include formal and informal employment)?
- Think about your most recent job, what was your monthly wage?
- Think about your most recent job, was it in the formal or informal sector?
- Think about your most recent job, how satisfied were you with it?

Mental Health

- Over the last two weeks, how often have you been bothered by the following problems (the four item patient health questionnaire for anxiety and depression)?
  - Feeling nervous, anxious or on edge
  - Not being able to stop or control worrying
  - Feeling down, depressed or hopeless
  - Little interest or pleasure in doing things
- Please imagine a ladder with steps numbered from zero at the bottom to 10 at the top. The top of the ladder represents the best possible life for you and the bottom of the ladder represents the worst possible life for you. On which step of the ladder would you say you personally feel you stand at this time? (Cantrill ladder-present) On which step do you think you will stand about five years from now? (Cantrill ladder-future)

From these data, we will create survey measures of alignment between top pre-existing skills and (i) in-demand careers, (ii) their expected career and (iii) their desired career. These alignment variables will take a value of 1 if a user's *pre-existing* skill is tagged to a career in each career set. We will do the same with the top three *desired* skills reported in the survey. For example, a job-seeker who states Microsoft Excel as a top 3 *pre-existing* skill and states their desired career is to be an accountant, would receive a value of 1 for alignment between existing skills and desired careers. Similarly, a job-seeker who states Python as a top 3 *desired* skill and states their expected career to be a Data Analyst will also receive a 1 for this measure of alignment.

5 Hypotheses and analysis

5.1 Main analysis

Our main specification is

$$Y_{imp} = \beta_0 + \beta_1 CV_{imp} + \beta_2 SkillRecs_{imp} + \beta_3 DemandRecs_{imp} + \gamma_s + X_{imp} + \varepsilon_{imp}$$

where Y_{im} is the outcome for person i in municipality m enrolled in program p. CV_{im} is the CV treatment, $SkillRecs_{im}$ is the skills based recommendation treatment, and $DemandRecs_{imp}$ is the demand based recommendation treatment. We include strata (γ_m) fixed effects (municipality, age and program) and X_{imp} is a vector of socio-demographic controls selected via double-selection LASSO, while ε_{imp} is the error term.

We are interested in the following estimates:

1. Research Question 1 – Comparing Control to T1 ($H_a : \beta_1 \neq 0$): Identifies the bundled impact of skill-identification and skill-communication via CV building.

2. Research Question 2 – Comparing Control to T2 ($H_a : \beta_2 \neq 0$): Identifies the combined effect of skill-identification and communication, CV building and career + course recommendations.

3. Research Question 3 – Comparing T1 to T2 ($H_a : \beta_1 \neq \beta_2$): Identifies the marginal gains from career + course matching among users who already have a better understanding of their skills and enhanced ability to communicate them (through skill-identification and CV building).

4. Research Question 4 – Comparing Control to T3 ($H_a : \beta_3 \neq 0$): Identifies the combined effect of skill identification, communication, CV building, career + course recommendations and in-demand information.

5. Research Question 5 – Comparing T2 to T3 ($H_a : \beta_2 \neq \beta_3$): Identifies the marginal gains from providing in-demand information to users who already have a better understanding of their skills and have received career + course recommendations based on that.

These estimates are causal estimates assuming outcomes were ex-ante identically distributed prior to intervention and if randomization of treatment assignment was successful. We report balance in Tables 2 and 3 to support our stance that randomization was successful.

5.2 Heterogeneity

We will analyze these results by gender and age, as the Ministry of Human Capital and the World Bank expect the programs to have different impacts based on these factors. These factors may also, critically, correlate strongly with job-seeker mobility. We will observe the location and modality of Fomentar courses (in-person, online, asynchronous), the location of firms for formal jobs, and participants' home municipalities (used for stratification) through administrative data. This lets us compute home-to-course and home-to-job distances, compare distance distributions by treatment arm, and estimate treatment-effect heterogeneity by gender—revealing whether mobility constraints lead women to choose nearer or remote/self-paced courses and jobs closer to home.

5.3 Multiple hypothesis testing

Our primary analysis focuses on five key outcomes, each assessed through specific targeted questions. Our survey is short and the number of questions we ask is limited. Nevertheless, within each category (e.g. Mental Health) we can consider accounting for multiple hypotheses by computing False Discovery Rate (FDR) q-values.

5.4 Power Calculations

To ensure that our design is powered to detect reasonably-sized treatment effects, we conduct a series of power calculations for both binary and standardized continuous outcome measures. We report the results of these power calculations in Figure 7. All power calculations are based on a two-sided hypothesis test with a 5% significance level and a total sample size of 500,000 participants, with 125,000 in each treatment arm and control. In both cases of binary and continuous outcomes, we report results with take-up rates of 5, 10, and 15% to show how power varies with take-up rates both higher and lower than those observed in the pilot. In the binary outcome calculation, reported in panel A, we are powered to 80% to detect effects ranging from 2 to 4.5 percentage points, depending on resulting take-up rates. Using realized pilot take-up rates of around 10%, the MDE for a binary measure is approximately 3.5 p.p.

These effect sizes are in line with evidence presented in similar contexts and in the papers cited in (Carranza and McKenzie, 2024a) review article. Evidence from an impact evaluation of sectoral training delivered in Argentina through the same training stream as Fomentar (Plan de Formación Profesional y Continua) also finds increases of 3 percentage points (over a base of 10%) in probability of formal employment (Castillo, Ohaco, and Schleser, 2014). Thus, we believe our study to be powered according to baseline job finding rates.

In the continuous case, reported in panel B, we are powered to 80% to detect effect sizes ranging from 0.075 to 0.225 SD. Given that the upper bounds of these MDEs represent a "worst-case" relative to pilot results for take-up, and that we do not take into account any specifications which control for baseline data, we expect to be sufficiently powered for all of our specified outcomes.

5.5 Spillovers

Given our design choices of individual-level randomization, there is the possibility of spillovers occurring for some of our downstream outcomes. In what follows we outline where these spillovers might arise, and why we think they are unlikely.

At employment offices: Access to SkillLab is provided through individual, non-shareable codes. Pilot interviews revealed that participants sometimes rely on employment office staff to assist with the registration process. This raises a potential concern that information could be shared—either among job seekers or by staff—particularly around which careers are in demand, even with those not assigned to the relevant treatment group.

However, we believe the risk of such spillovers is low. Job seekers rarely visit the employment offices on the same day, reducing the chance of peer-to-peer information exchange. Additionally, staff primarily focus on helping users navigate the platform, rather than interpreting or sharing specific recommendations. Moreover, the in-demand career recommendations are at least partially personalized, as they factor in each user's prior experience, which limits their relevance to others.

To further mitigate any risk, we will provide employment offices with written implementation guidelines. These documents will emphasize that staff should not provide access to users outside of their assigned treatment group or share any recommendation information beyond what is accessible through a participant's individual platform access.

If some information-sharing does occur, our estimates of the treatment effect would likely represent a lower bound.

In enrolled courses: While SkillLab CV, full, and full + demand are separate platforms, the courses which treatment and control participants have access to through the Portal Empleo are ultimately the same. Thus, in-person courses will likely contain a mix of treatment and control participants. A concern is that those with more information (for example someone who knows a course is relevant for an in-demand career) will exert differential effort or impact the completion/performance of those in the class. While we believe this type of spillover is relatively unlikely—since classmates would need to notice that their peers were putting in more effort and attribute that change specifically to insights gained

from the app—we will test the robustness of our results by leveraging a feature of the Portal Empleo courses: variation in course modality. The Portal offers courses which are in-person, virtual live courses, and self-paced online courses. We can assess outcomes within treatment group across course modality to evaluate the extent of spillovers, as self-paced courses have little chance for these peer effects.

References

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2020. "Tackling youth unemployment: Evidence from a labor market experiment in Uganda." *Econometrica* 88 (6):2369–2414.

Bächli, Mirjam, Rafael Lalive, and Michele Pellizzari. 2025. "Helping Jobseekers with Recommendations Based on Skill Profiles or Past Experience: Evidence from a Randomized Intervention." .

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2021. "The Search for Good Jobs: Evidence from a Six-year Field Experiment in Uganda." *SSRN Electronic Journal* URL `https://www.ssrn.com/abstract=3910330`.

Belot, Michele, Philipp Kircher, and Paul Muller. 2019. "Providing advice to jobseekers at low cost: An experimental study on online advice." *The review of economic studies* 86 (4):1411–1447.

Carranza, Eliana and David McKenzie. 2024a. "Job Training and Job Search Assistance Policies in Developing Countries." *Journal of Economic Perspectives* 38 (1):221–244. URL `https://pubs.aeaweb.org/doi/10.1257/jep.38.1.221`.

———. 2024b. "Job Training and Job Search Assistance Policies in Developing Countries." *Journal of Economic Perspectives* 38 (1):221–44. URL `https://www.aeaweb.org/articles?id=10.1257/jep.38.1.221`.

Castillo, Victoria, Moira Ohaco, and Diego Schleser. 2014. "Evaluación de impacto en la inserción laboral de los beneficiarios de los cursos sectoriales de formación profesional." Tech. Rep. 6, Organización Internacional del Trabajo (OIT), Oficina de País para la Argentina, Buenos Aires. URL `https://researchrepository.ilo.org/view/pdfCoverPage?download=true&filePid=13100780200002676&instCode=41ILO_INST`.

Kiss, Andrea, Robert Garlick, Kate Orkin, and Lukas Hensel. 2023. "Jobseekers' beliefs about comparative advantage and (mis) directed search." *Available at SSRN 4593303* .

Leduc, Elisabeth and Ilan Tojerow. 2025. "Closing the Mismatch: Encouraging Jobseekers to Reskill for Shortage Occupations." .

Orkin, Kate, Rob Garlick, Mahreen Mahmud, Richard Sedlmayr, Johannes Haushofer, and Stefan Dercon. 2023. "Aspiring to a better future: Can a simple psychological intervention reduce poverty?" Tech. rep., National Bureau of Economic Research.

Sanchez Puerta, Maria Laura, David A Robalino, Victoria Strokova, Nick Lord, and Mathilde Perinet. 2015. "Skills and jobs: lessons learned and options for collaboration." .

Vezza, Evelyn, Gonzalo Zunino, Luis Laguinge, Harry Edmund Moroz, Ignacio Raul Apella, and Marla Hillary Spivack. 2025. "Understanding Labor Market Demand in Real Time in Argentina and Uruguay." Tech. rep., The World Bank.

Tables

Table 1: Summary Statistics for Fomentar/VaT Participants

Characteristic	**N = 1,048,574**[1]
Sex	
Female	63%
Male	37%
Non-binary	<0.1%
Age	32
Has Young Children	
No	53%
Yes	47%
Education	
Higher or Equal to Complete Primary	91%
Higher or Equal to Complete Secondary	45%
Ethnicity	
No Ethnic Identity Reported	95%
African	<0.1%
Afro-Argentine	4.1%
Afro-descendant	0.7%
Romani	0.3%
Has Previous Work Experience	
Not Reported	4.8%
No	54%
Yes	41%
Program	
Fomentar	63%
VAT	37%

[1] Column percentages unless otherwise noted. Age presented as mean. This table presents total eligible sample including those used in the pilot of March/April 2025 ($N = 14,227$).

Variable	(1) Control Mean/(SE)	(2) Treated Mean/(SE)	(1)-(2) Pairwise t-test Mean difference
Age	32.286	32.292	–0.007
	(0.018)	(0.011)	
Above Median Age	0.487	0.487	–0.000
	(0.001)	(0.001)	
Female	0.626	0.628	–0.002
	(0.001)	(0.001)	
Fomentar Participant	0.633	0.633	–0.000
	(0.001)	(0.001)	
Has children	0.464	0.465	–0.001
	(0.001)	(0.001)	
Primary Educ.	0.911	0.911	–0.000
	(0.001)	(0.000)	
Secondary Educ.	0.445	0.448	–0.003**
	(0.001)	(0.001)	
Imp. Phase 1	0.019	0.019	0.000
	(0.000)	(0.000)	
Imp. Phase 2	0.082	0.082	–0.000
	(0.001)	(0.000)	
Imp. Phase 3	0.899	0.899	0.000
	(0.001)	(0.000)	
N	258586	775761	1034347

Table 2: Balance Table, Control vs. Any Treatment

Note: This table reports balance tests on observable characteristics. Column (2) includes all Fomentar/VaT participants randomized to any of the treatment arms. Randomization was stratified across municipality, above/below median age, and Fomentar/VaT participation. Eligible sample are all Fomentar/VaT participants not randomized into a treatment group in the March/April pilot. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Treatment Arm Mean				*p*-value, Difference in Means t-test					
Variable	μ_C	μ_{T1}	μ_{T2}	μ_{T3}	C-T1	C-T2	C-T3	T1-T2	T1-T3	T2-T3
Age	32.286	32.271	32.309	32.297	0.014	–0.023	–0.011	–0.038	–0.026	0.012
	(0.018)	(0.018)	(0.020)	(0.018)						
Above Median Age	0.487	0.487	0.487	0.487	0.000	–0.000	–0.000	–0.000	–0.000	–0.000
	(0.001)	(0.001)	(0.001)	(0.001)						
Female	0.626	0.627	0.628	0.629	–0.000	–0.002	–0.003*	–0.002	–0.002	–0.001
	(0.001)	(0.001)	(0.001)	(0.001)						
Fomentar Participant	0.633	0.633	0.633	0.633	0.000	–0.000	–0.000	–0.000	–0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)						
Has children	0.464	0.464	0.465	0.467	0.000	–0.001	–0.003**	–0.001	–0.003**	–0.002
	(0.001)	(0.001)	(0.001)	(0.001)						
Primary Educ.	0.911	0.911	0.911	0.910	–0.001	–0.000	0.000	0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)						
Secondary Educ.	0.445	0.449	0.448	0.447	–0.003**	–0.003**	–0.001	0.000	0.002	0.001
	(0.001)	(0.001)	(0.001)	(0.001)						
Imp. Phase 1	0.019	0.019	0.019	0.019	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)						
Imp. Phase 2	0.082	0.082	0.082	0.082	–0.000	–0.000	–0.000	0.000	–0.000	–0.000
	(0.001)	(0.001)	(0.001)	(0.001)						
Imp. Phase 3	0.899	0.899	0.899	0.899	0.000	0.000	0.000	–0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)						
N	258586	258587	258587	258587	517173	517173	517173	517174	517174	517174

Table 3: Balance Table, All Treatment Arms

Note: This table reports balance tests between all arms comparing observable characteristics. Randomization was stratified across municipality, above/below median age, and Fomentar/VaT participation. Eligible sample are all Fomentar/VaT participants not randomized into a treatment group in the March/April pilot. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Figures

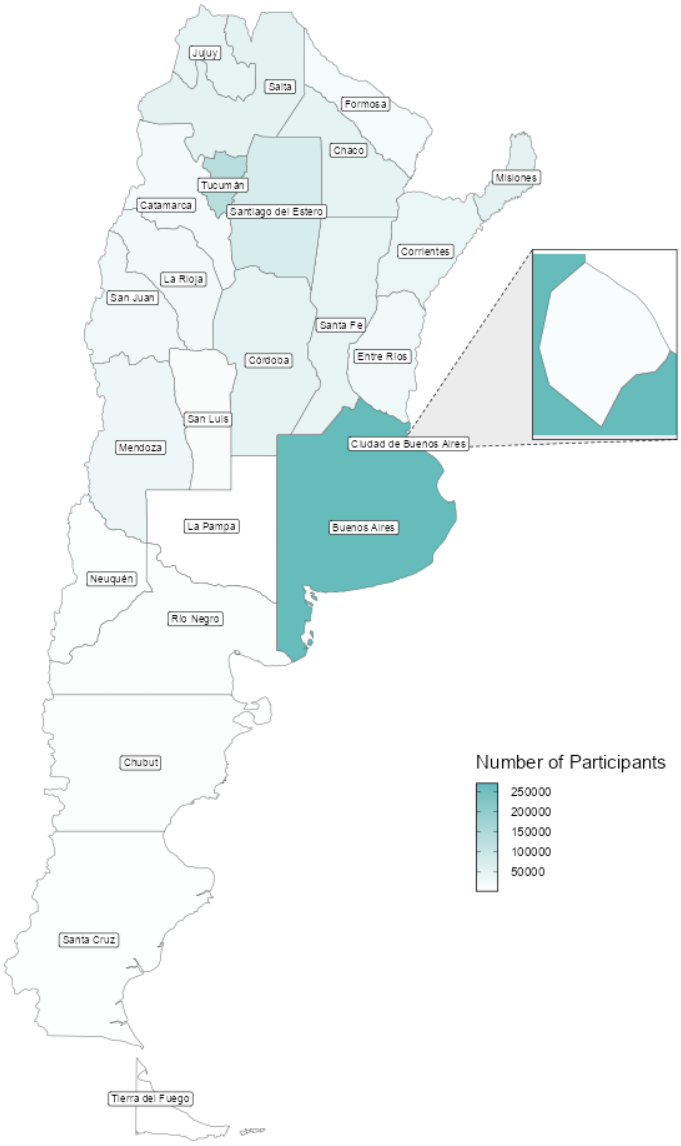

Figure 1: Eligible Participants by Province

Note: This figure reports the geographic distribution of Fomentar and VaT participants by province. Eligibility is according to SEFP administrative data as of December 2024.

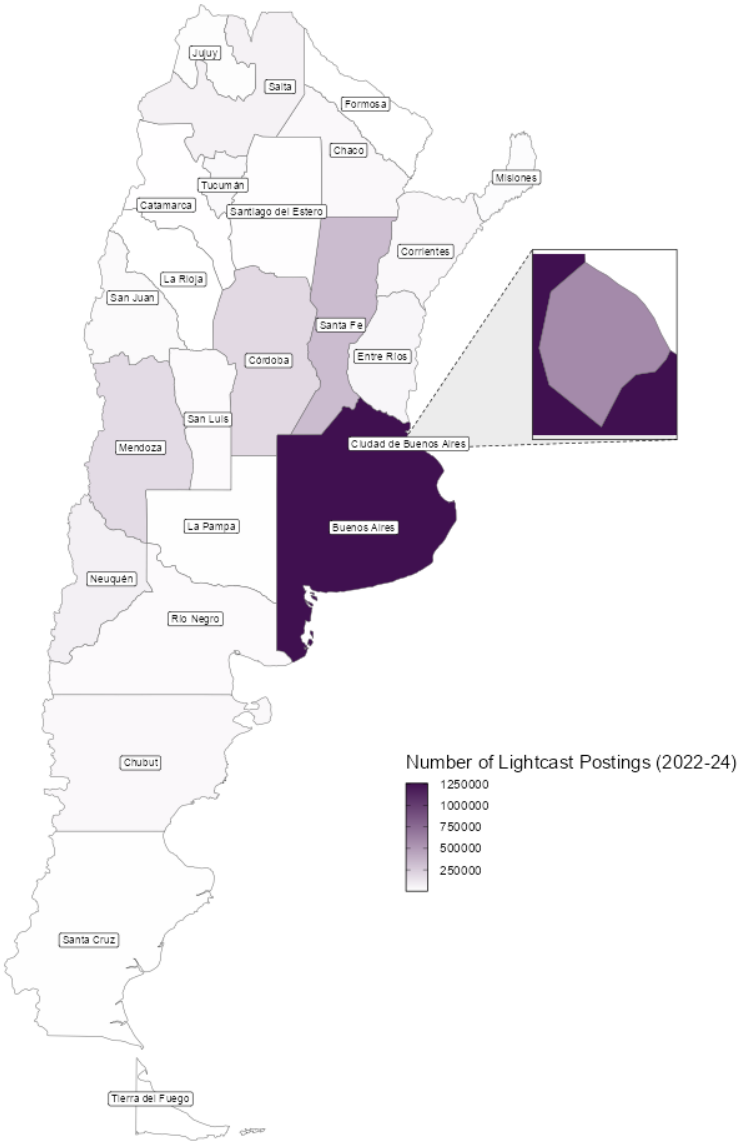

Figure 2: Geographic Distribution LightCast Job Postings

Note: This figure reports the geographic distribution of job postings in LightCast data from the time period Jan 1st 2023 to Dec 31st 2024. These postings are used to compute growth metrics and posting counts to determine "in-demand" status.

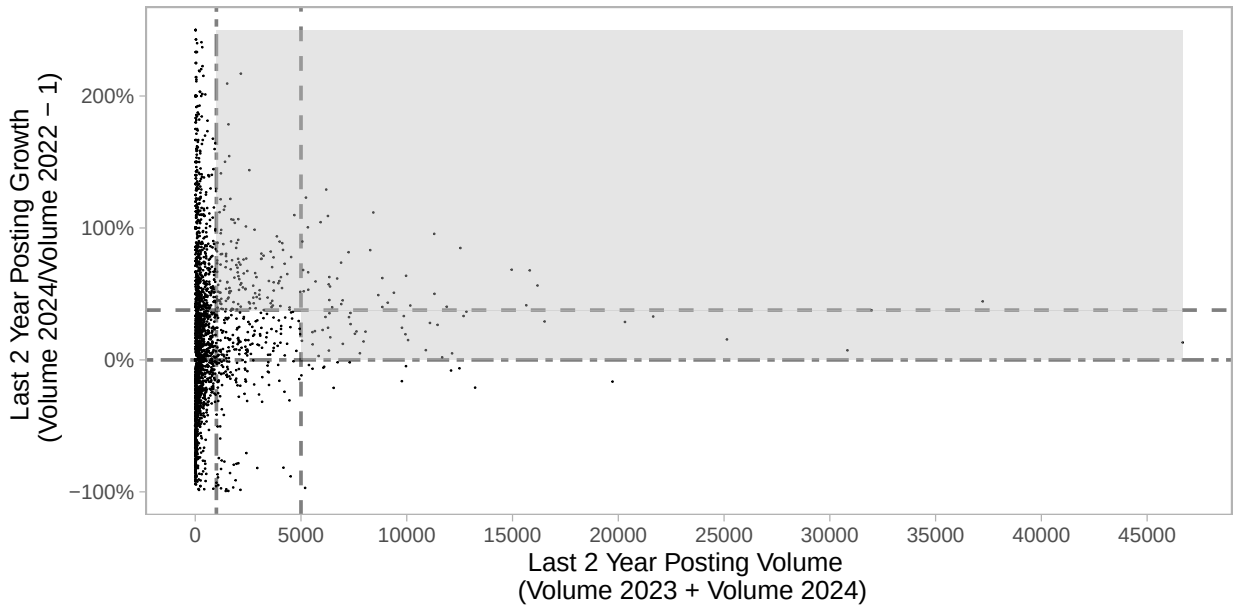

Figure 3: In-Demand Occupation Definition

Note: This figure illustrates our in-demand cutoff rule at the start of the experiment. Each point is a Level 5 ESCO occupation. The x-axis measures the sum of all postings (i.e. volume) in that occupation from Jan 1st 2023 to Dec 31st 2024. The shaded region represents the in-demand region. The y-axis measures 2 year posting growth. This is the percentage difference between the 2024 postings volume and the 2022 postings volume for that occupation.

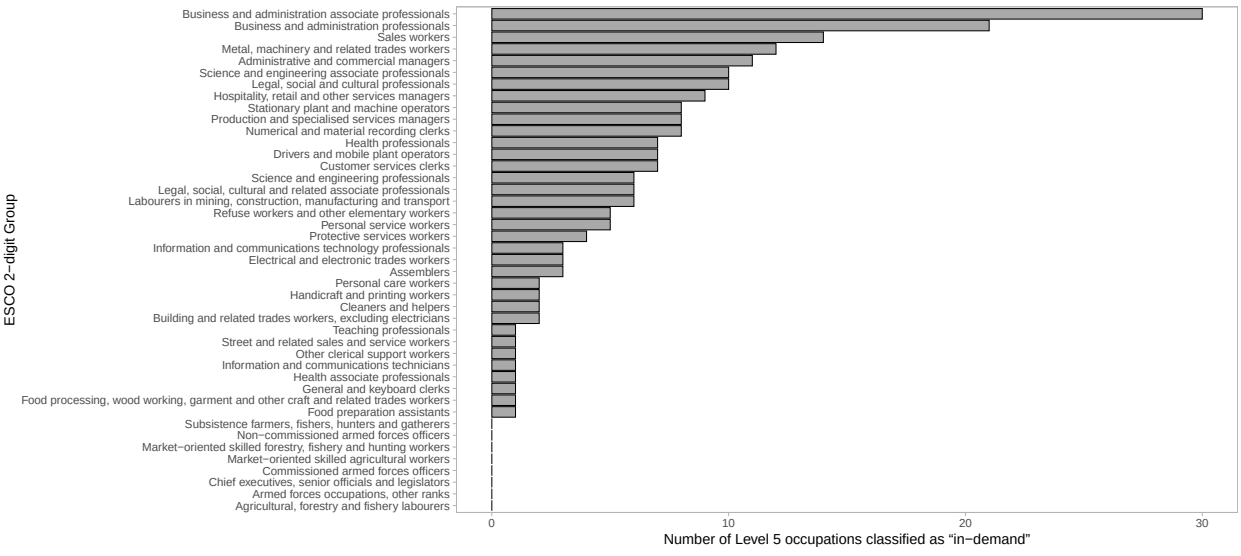

Figure 4: In-Demand Occupations by Group

Note: This figure shows the number of individual occupations at the 5-digit ESCo classification in each 2-digit ESCO occupation group determined to be in-demand based on the cutoffs defined in Figure 3.

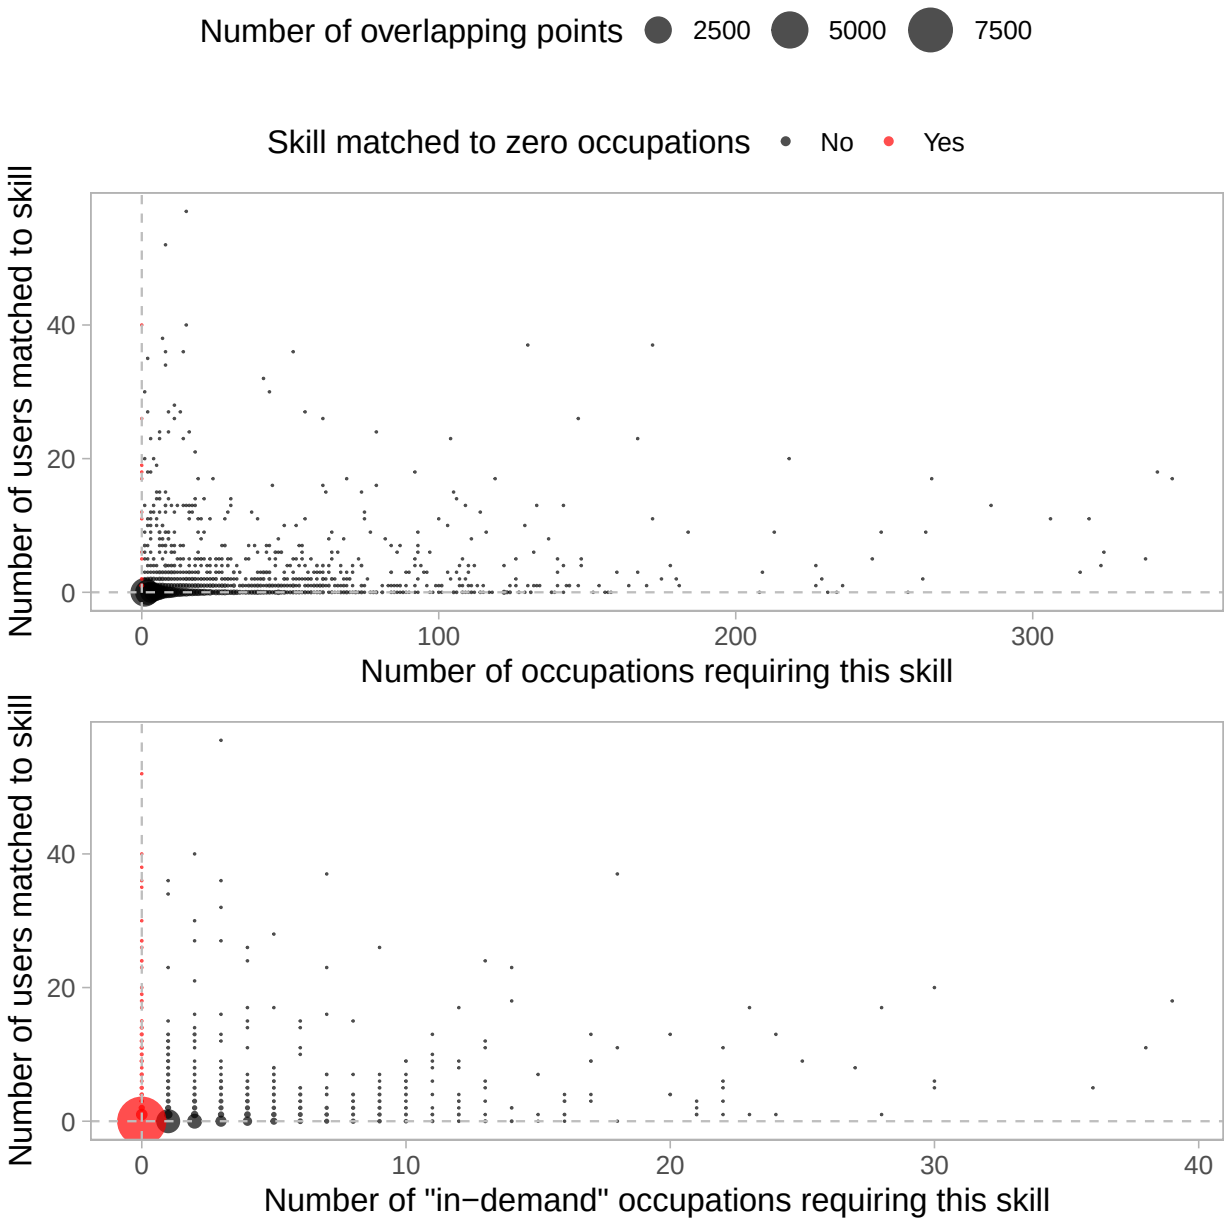

Figure 5: Overlap Between In-Demand and User Skills

Note: This figure presents two panels where each dot represents at least one skill in the ESCO taxonomy. The top panel displays the number of unique occupations matched to a skill (horizontal axis) and number of users matched to the same skill. The bottom panel reports the same exercise but subset to only include occupations which we flag as "in-demand." In each panel, dots off of both axes represent skills which are both matched to at least one user and matched to at least one occupation.

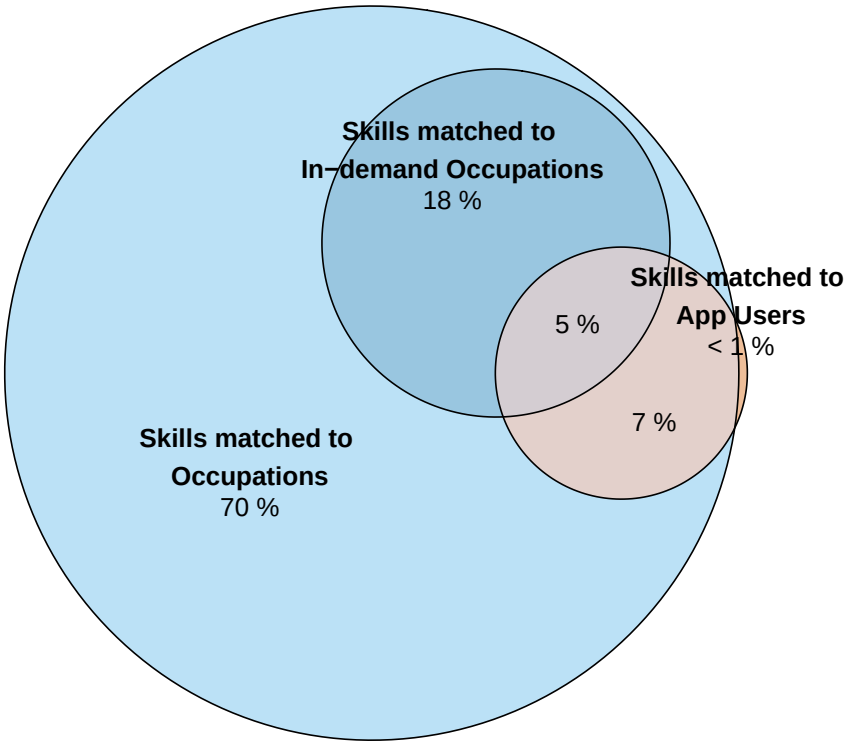

Figure 6: Skill Overlap

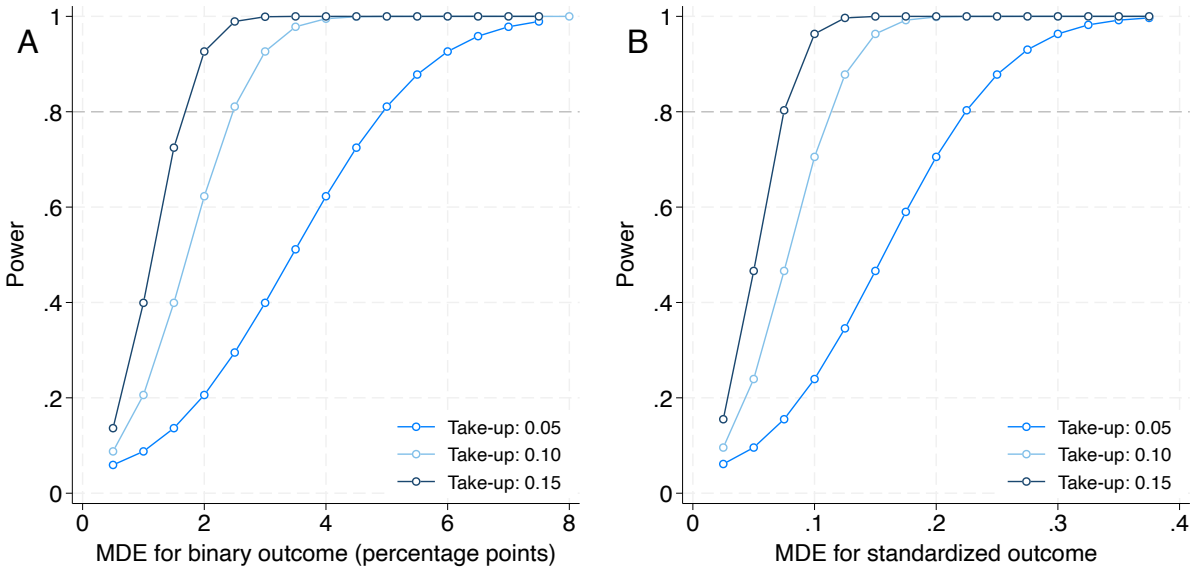

Figure 7: A representation of power calculations

Note: This figure presents the results of power calculations for binary outcomes (Panel A) and standardized outcomes (Panel B) with variation in intervention take-up rates. Panel A assumes 5% control mean in the outcome variable.

Appendix

A Portal Appearance

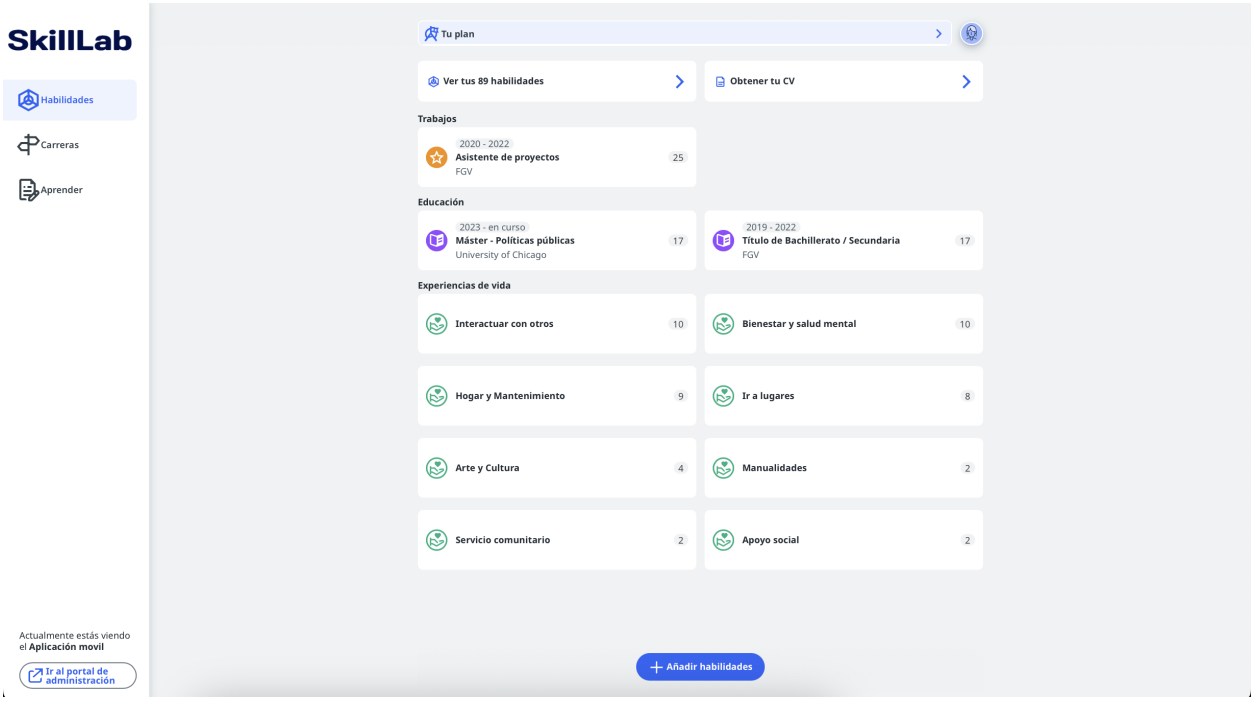

Figure A1: Skills Homepage for T1-T3

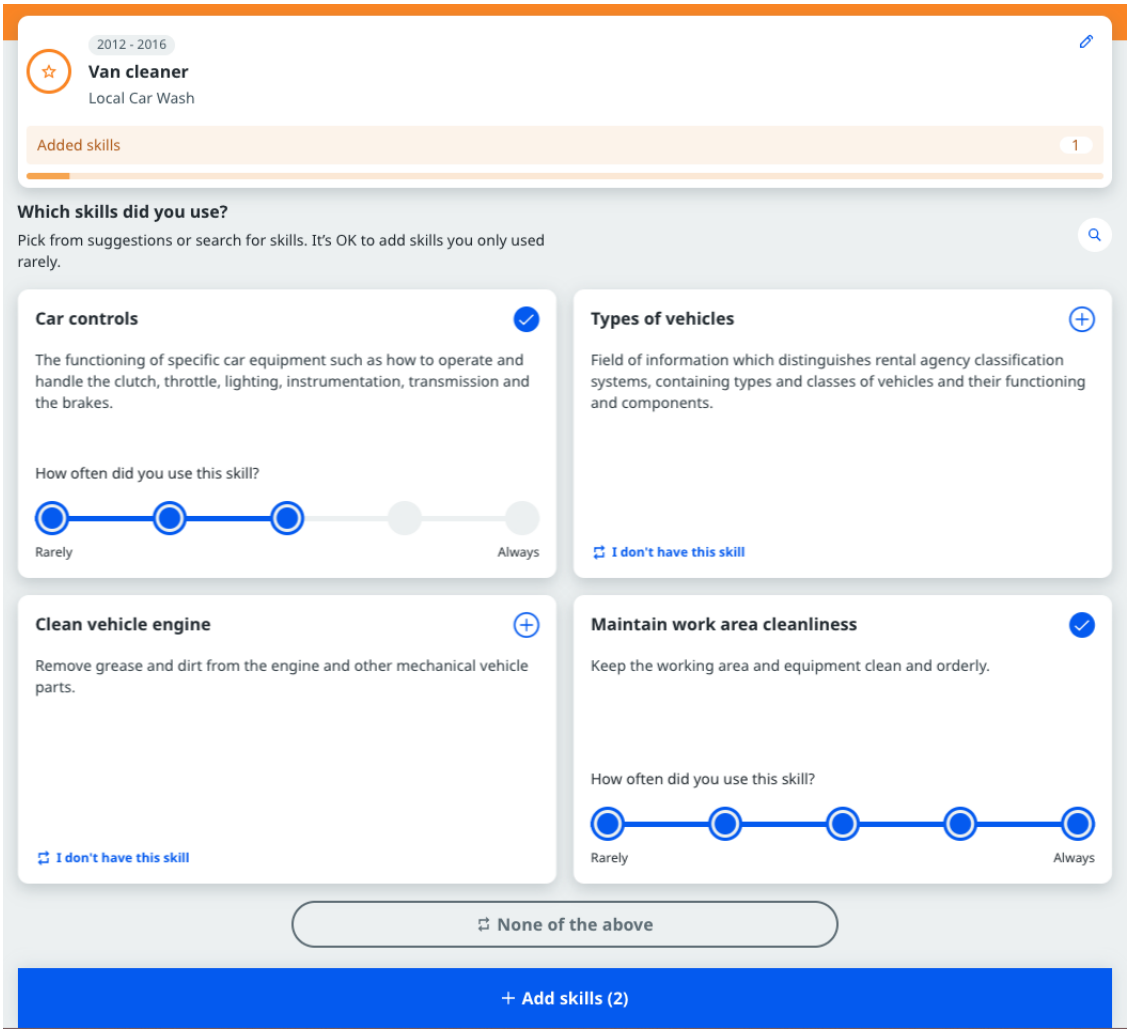

Figure A2: Example Skills Recommendation for T1-T3

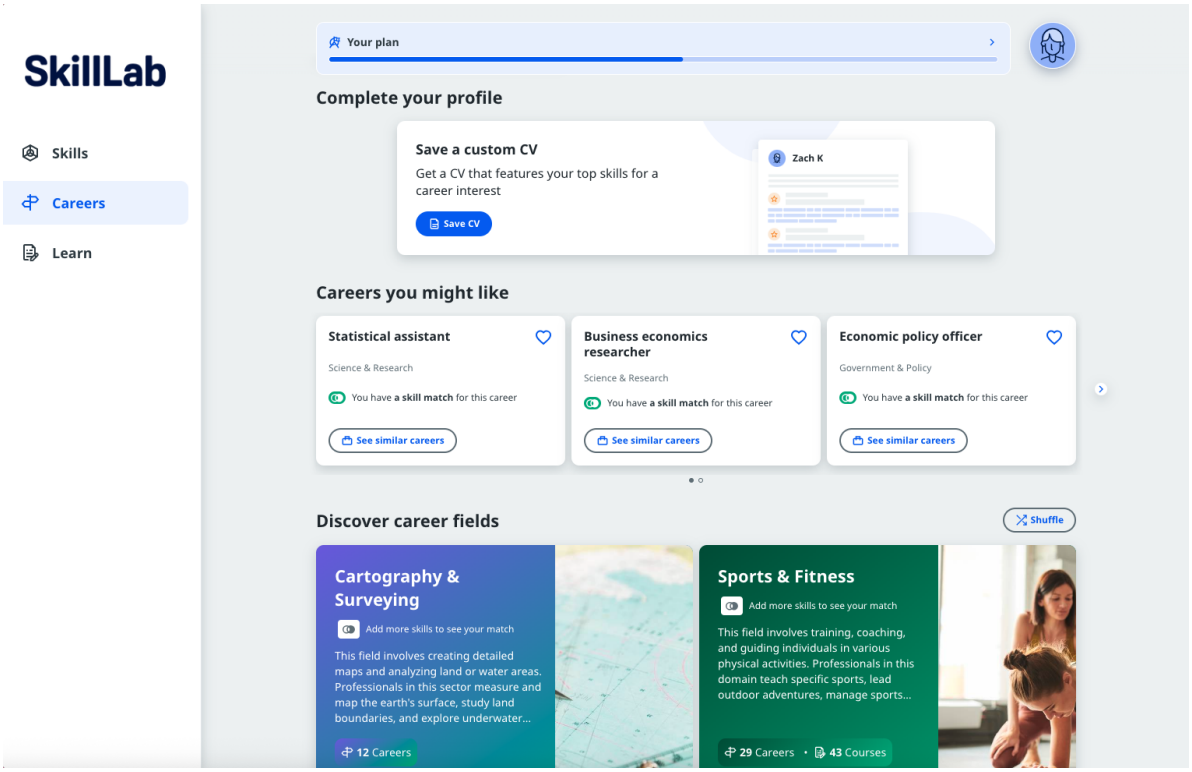

Figure A3: Careers Homepage for Treatment 2 (SkillLab Full)

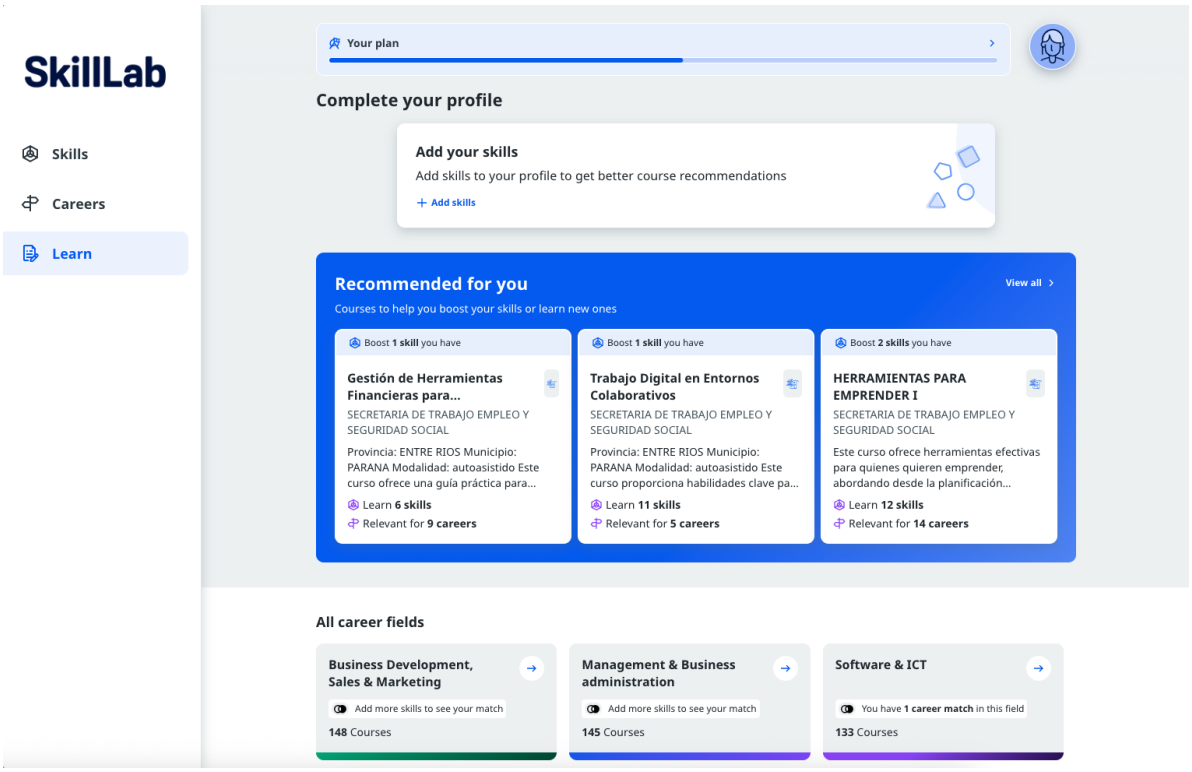

Figure A4: Courses Homepage for Treatment 2 (SkillLab Full)

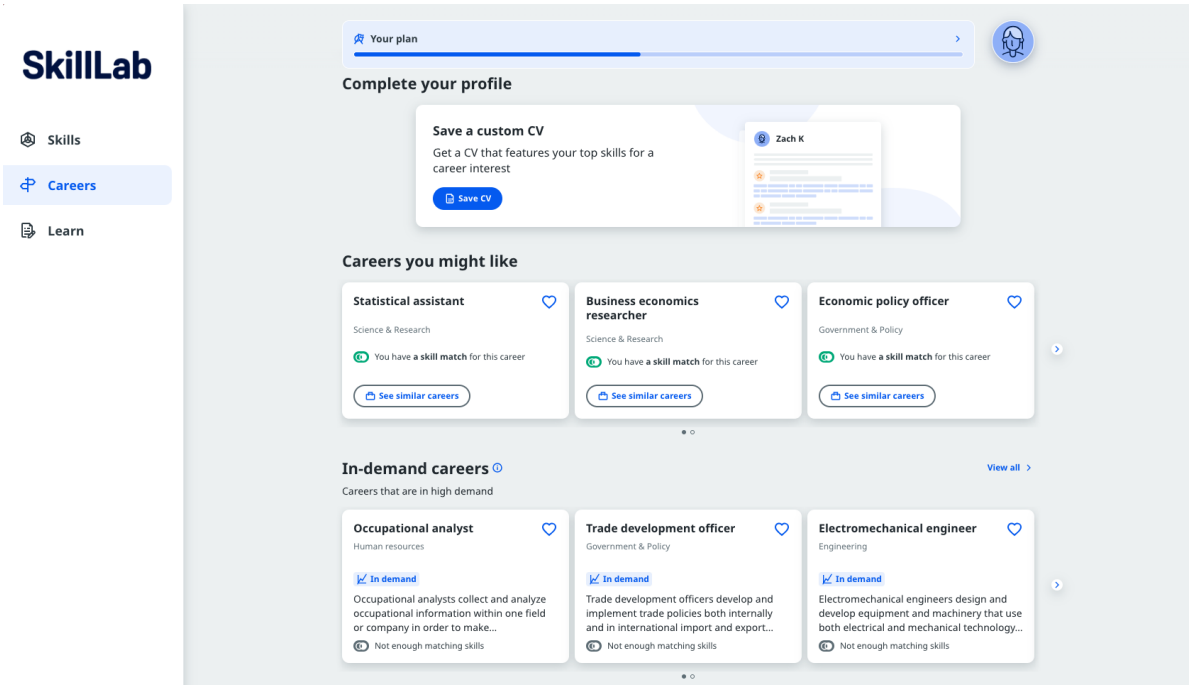

Figure A5: Careers in Treatment 3 (SkillLab Full + In-Demand)

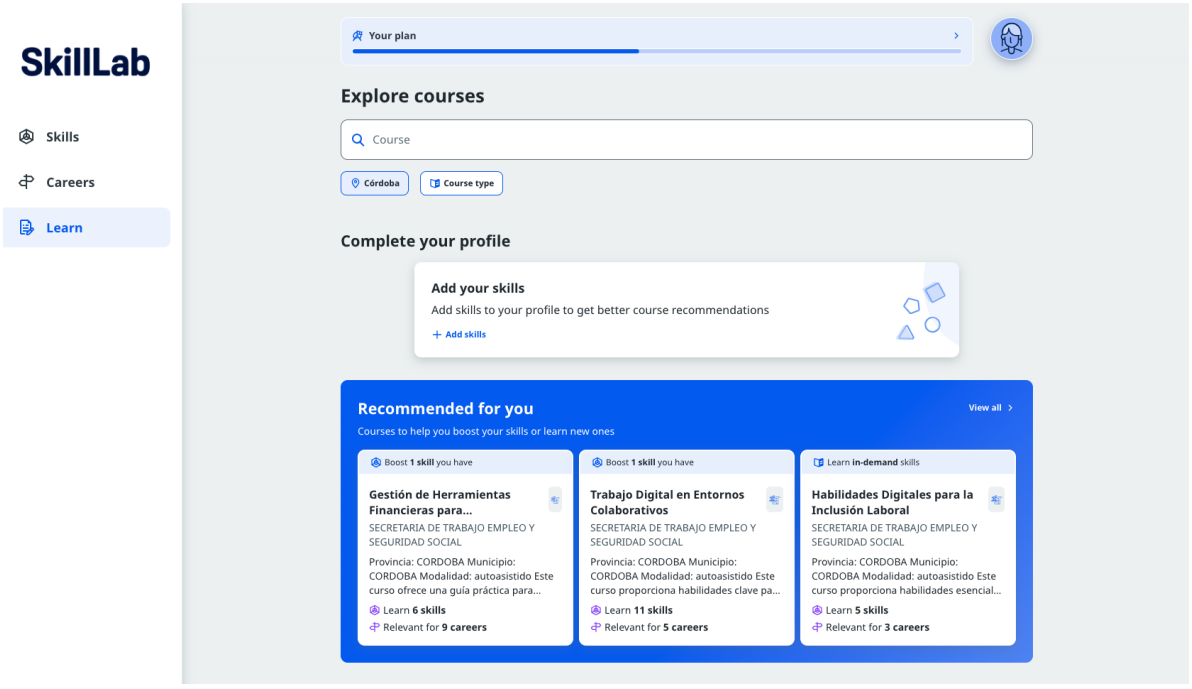

Figure A6: Courses in Treatment 3 (SkillLab Full + In-Demand)

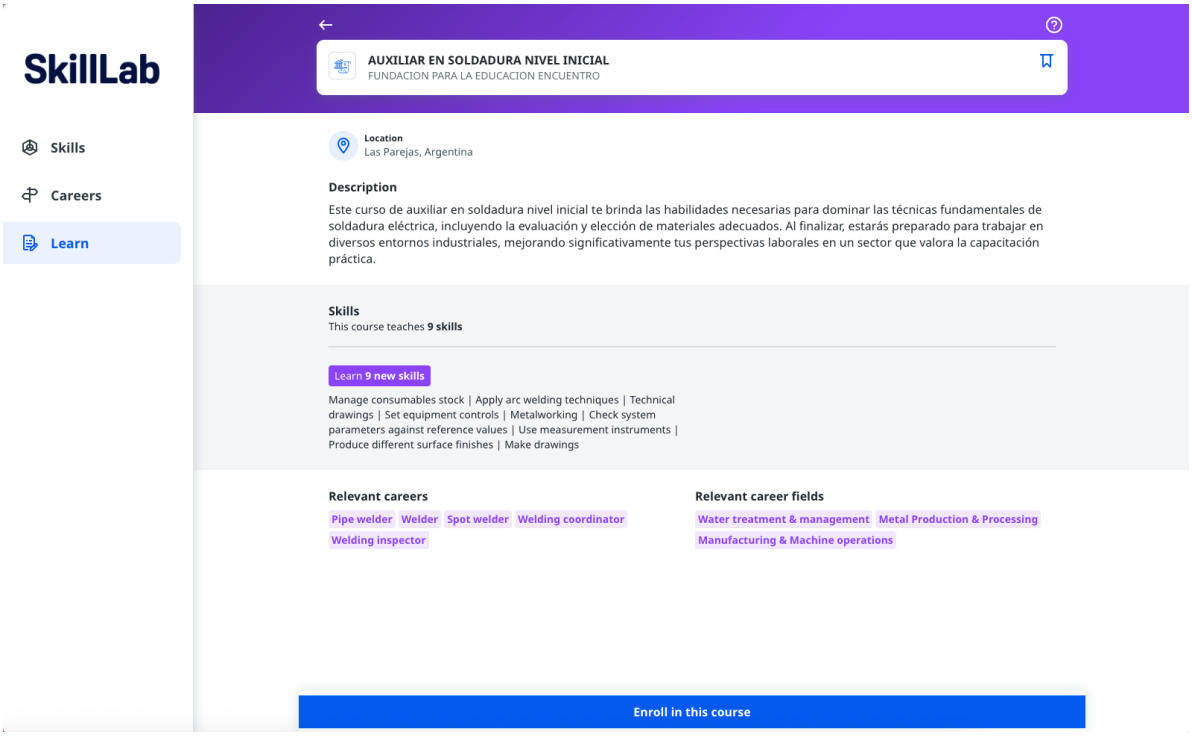

Figure A7: Course Information Page in Treatments 2 and 3

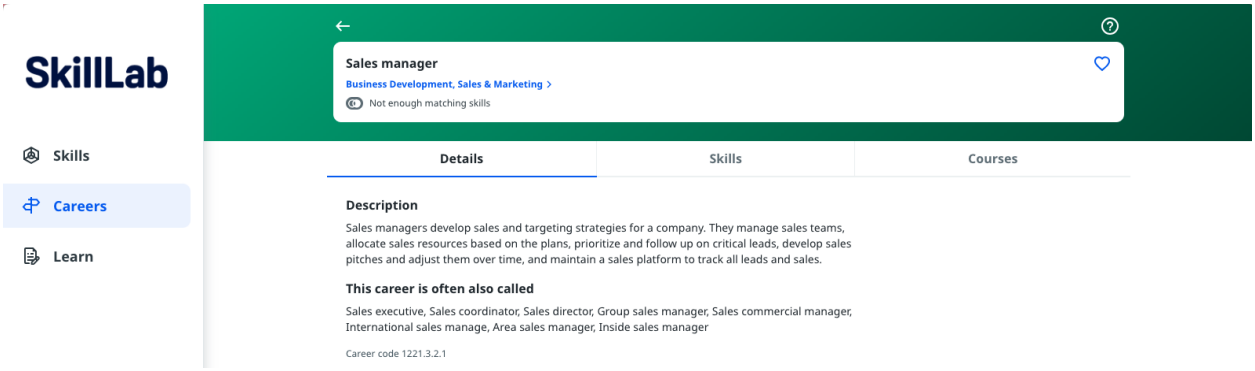

(a) Career Overview

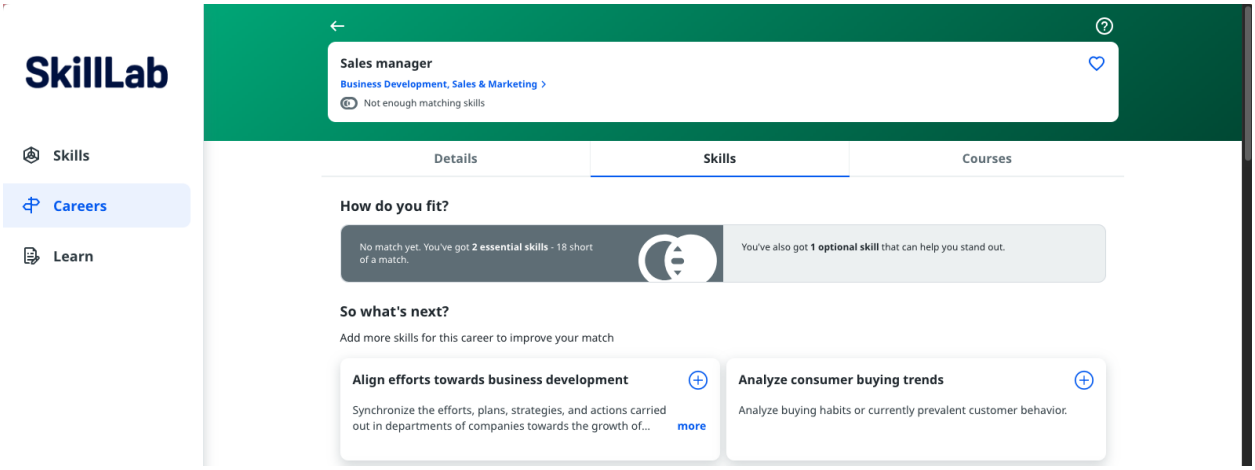

(b) Relevant Skills for Career

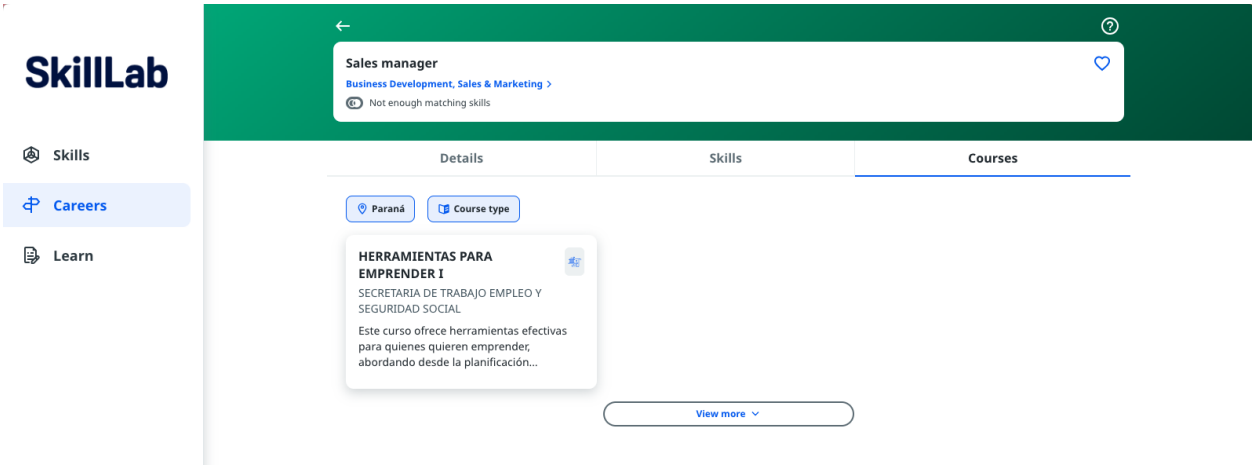

(c) Relevant Courses for Career

Figure A8: Career Information Page

B Invitation Emails

Invitation in Spanish

The following invitation will be sent to all participants randomized into T1, T2, and T3.

Subject: Descubrí tus habilidades y creá tu CV con SkillLab

Hola [nombre],

¿Te gustaría impulsar tus oportunidades laborales? ¡Te invitamos a probar SkillLab! Con esta herramienta online, podés:

- Descubrir tus habilidades rápida y fácilmente
- Crear un CV de calidad sin esfuerzo
- Obtener consejos para conseguir un empleo y mejorar tu empleabilidad

¡Hay la posibilidad también que tengas acceso a recomendaciones de cursos de formación profesional publicados en el Portal Empleo y a recomendaciones de carreras que se adapten a tus habilidades e intereses y a las tendencias del mercado laboral!

Hacé clic en el botón "COMENZAR" para dar el próximo paso hacia nuevas metas.

Para más información, hacé clic [aquí].

[Comenzar]

Invitation in English

Subject: Discover your skills and create your CV with SkillLab!

Hello [first name],

Would you like to boost your job opportunities? We invite you to try SkillLab! With this online tool, you can:

- Discover your skills quickly and easily
- Create a professional CV effortlessly
- Get tips to help you land a job

There is the possibility that you will receive access to recommendations for training courses published in the Portal Empleo and careers which match your skills, interests, and trends in the labor market!

Click on the "GET STARTED" button to take the next step towards new goals.

For more information, click [here].

[Get Started]

C The SkillLab App

C.1 Data Processes

A key contribution of the SkillLab is to help users triangulate skills, courses, and career opportunities. All user input is in the form of experiences, skills, and "likes" for interests.

Adding Skills Skills are elicited through a highly-structured manner in which users first select high-level "experiences" from education, work, or other aspects of life. When users add a relevant experience, the platform suggests four related skills to add to their profile, with the option to refresh for more suggestions, as displayed in Figure A2. Each skill appears in a box displaying its name and a brief plain-language description of 1–3 sentences. The process of matching skills to experiences follows the European Commission's well-established and rigorously vetted ESCO framework (Skills, Competences, Qualifications and Occupations). When users flag skills they applied in a given experience, they also rate how often they used each skill on a Likert scale from 1 ('rarely') to 5 ('always'). This discovery process is independent of career preferences.

Matching User Skills to Careers A key goal of SkillLab is to suggest careers for which users have a degree of *skill match*. To achieve this, SkillLab uses a taxonomy that matches each career to a set of skills. Each career has a set of associated ESCO skills, broken-down into "essential" and "optional" skills. SkillLab then determines the set of careers for which the user has a degree of skill match by comparing the skills the user has to the skills of a particular career. A career is considered a match if the user possesses at least 50% of its essential skills.

Matching ESCO Skills to Courses SEFP provided our research team with descriptions of every course made available to participants on the *Portal Empleo*. SkillLab then parses the text from these skills, and uses another proprietary model which uses Named Entity Recognition (NER) to tag each course with a set of ESCO skills likely to be taught therein using keywords or short phrases from the course descriptions. NER is a Natural Language Processing technique which searches for a pre-specified set of tags (i.e. ESCO skills) in unstructured text (Fomentar course description text). The end result is a mapping from every course offered by Fomentar to a set of ESCO skills.

Matching Courses to Careers For each course, SkillLab uses LLMs to parse the text for each course and create an embedding, which contains semantic information.[3] Similarly, embeddings are generated for each ESCO occupation in the SkillLab database. To identify semantically similar careers, the embeddings from courses and ESCO careers are compared and a similarity score is computed based on the distances between embedding vectors. Each course is then tagged with careers which have high similarity scores.

C.2 Recommendation Details

Course Recommendations In Treatment 2, courses are recommended using a weighting of user career interests and the user's skill-set. For skill-aligned recommendations, SkillLab recommends courses which, according to course-level ESCO skill tags, teach skills which the user has familiarity with but does not have great experience with (indicated by self-reported use frequency) in their skill profile ("boostable" skills). For career interests, SkillLab then also recommends courses which are tagged to the occupations liked by the user. In the absence of both boostable skills and stated interests, SkillLab prioritizes courses which teach the skills the user already has. The algorithm does not have expansion of the user's skill-set to related skills as an objective, but this often occurs in practice since courses teach bundles of related skills. These recommendations are displayed to users in a carousel which contains a clickable field with the course name and the reason for the recommendation (e.g. a course related to a "liked" career will be visibly tagged as being relevant for that career interest). The user can also browse all career fields below the recommendations. See Figure A4 for an example of the career recommendations in Treatment 2.

In Treatment 3, courses are recommended using a weighting of three distinct factors: interests, skill-alignment (identical to Treatment 2), and additionally demand data. Figure A6 in the gives an example of the course recommender in Treatment 3. In that example, the user sees three recommendations, two of which are due to a skill match, the third is teaching skills associated with at least one in-demand career.

Clicking or tapping on the card for a *course* opens a new page containing more information. At the top, the user sees location, modality, and a brief course description highlighting the aims and contents of the course. Just below, the user sees the skills tagged to the course, and the careers tagged to the course as well. At the bottom, there is a convenient link to enroll, which directly brings the user to the enrollment page within the Portal Empleo.

[3]Text embeddings are numerical vector representations of text data that capture semantic meaning, allowing machines to understand the relationships between words and phrases. This process is separate from the skill-tagging process due to the higher level of abstraction for "careers" compared to skills. LLM's are better suited to tag courses with more abstract concepts due to their enhanced ability to consider larger bodies of text.

Figure A7 shows an example of this course information page for an in-person welding course.

Career Recommendations In Treatment 2, career recommendations are initially based solely on skill match. The threshold for skill-match recommendation is 50% of *essential* skills covered by the user. Up to 6 of the highest-skill matched careers are displayed in a carousel to the user, displayed in Figure A3. However, in the event where a user has more than 6 skill-matched careers, they can optionally browse the full catalog of careers, which will be visibly identified as being skill-matched careers if appropriate. If the user expresses career interests, recommendations are updated using a proprietary weighting of skill-information and preferences.

In Treatment 3, career recommendations appear in two separate carousels in the "Careers" tab. The first carousel is nearly identical in structure to the carousel available to those in Treatment 2, weighting user skill match and interest, with some weight placed on careers which are *also* in-demand. Users in Treatment 3 will also see an additional carousel which shows 6 "in-demand" careers (defined as jobs with a high volume of postings or strong growth in the Lightcast data), ordered from highest skill match to lowest skill match. Figure in the manuscript gives an example of the first carousel under "Careers you might like", and the second carousel which more heavily weights demand data: or "Careers in demand" Examples of each of these recommendation carousels are presented in Figure A5.

Clicking or tapping on the card for a *career* opens a new screen containing three pages, one for details/overview, one for skills, and one for courses. The overview page contains a brief career description and alternative names for positions which are close approximates of the same career. The skills page shows all of the essential and optional skills for the career, with the option to add these skills directly from this page. Lastly, the courses page shows every available course tagged to the selected career. Figure A8 displays examples of each information page for the Sales Manager career.