The Making of a Public Sector Worker:
The Causal Effects of Temporary Work Assignments to Poor Areas*

Patrick Agte
Stockholm School of Economics

Mariel Bedoya
Columbia University
(Job Market Paper)

November 1, 2026

Click here for the latest version

Abstract

Can a temporary policy intervention permanently change where people choose to work? We study this question using Peru's mandatory rural service program, which assigns health professionals to posts across the country for one year. We exploit exogenous variation from a lottery system determining psychology graduates' service assignments. Using administrative data and an original survey, we find that psychologists assigned to poorer locations are significantly more likely to subsequently work in the public sector (17%) and poorest districts (54%) up to three years after the program. We then present evidence that the experience changes worker preferences and beliefs. Psychologists assigned to poorer areas score 0.11 standard deviations higher on a prosociality index, and report greater willingness to work serving the poor: both geographically and by type of employer (Ministry of Health facilities). We rule out inertia, differential skill formation, and network effects as alternative explanations. These findings demonstrate the potential of preference endogeneity, not just financial incentives, to attract workers to underserved areas.

*Contact information: patrick.agte@hhs.se; mrb2285@tc.columbia.edu. We are grateful to Alex Eble, Cristian Pop-Eleches, and Sebastián Otero for their continued encouragement, advice, and patience. We thank Janet Currie, Thomas Fujiwara, Rohini Pande, Carolina Concha-Arriagada, Arielle Bernhardt, Mateus Dias, Seema Jayachandran, Christopher Neilson, Natalia Rigol, Wojciech Kopczuk, Jack Willis, Cristóbal Otero, Diana Moreira, and seminar participants at Princeton and Columbia for helpful advice and comments. We want to thank IPA and GRADE for their data collection support and especially our team members for their excellent work on this project: Zaragoza Limachi, Jorge Luis Tipisman, and Dianela Espinoza. This project was made possible in collaboration with the SERUMS National Coordination Unit of the Peruvian Ministry of Health. Funding was provided by the Princeton Research Program in Development Economics, the Princeton Industrial Relations Section, the CDEP Student Research Grant Program, and the J-PAL Governance Initiative. All remaining errors are our own.

1 Introduction

Ensuring adequate provision of public services in underserved areas remains a critical policy challenge. Despite higher social returns to working in these areas, workers systematically sort away from them. This pattern has particularly severe consequences in sectors with positive externalities such as healthcare, where the projected global shortage is expected to reach 12 million workers by 2030 (Boniol et al., 2022). This crisis is particularly acute in rural areas and developing countries: the International Labour Organization estimates that rural shortages are roughly twice as large as urban ones across low- and middle-income countries (Scheil-Adlung, 2015), while high-income countries hold nearly half the global health workforce despite accounting for only a fifth of the world's population (Boniol et al., 2022). Existing policy responses have focused primarily on financial incentives to attract workers to underserved areas, yet there is limited evidence on whether the work experience itself can durably shift workers' preferences and beliefs about such postings, and through this channel, their long-run career choices.

Compulsory temporary service programs have been used in over 70 countries around the world (Frehywot et al., 2010) as an alternative policy tool to reduce staffing needs and retain health workforce in remote and underserved locations. These programs often require newly graduated health workers to spend a fixed period working in rural and poor areas in order to access professional privileges such as earning a certification needed for practice, receiving educational financial support, accessing public sector employment, or specialist training. Peru's *Servicio Rural y Urbano Marginal de Salud* (SERUMS), in operation since 1981, has been one such program: it places recently graduated health workers in facilities in the country's most vulnerable and underserved areas for one year in exchange for subsequent access to public sector employment and specialist training. It uses a lottery to assign participants to posts. The nature of the experience raises the possibility that the assignment itself, rather than the compensation attached to it, could serve as a lever for changing preferences and updating beliefs of workers exposed to these communities, thus expanding the pool of workers willing to staff these areas. Despite their proliferation, rigorous evidence on the effects of these programs on labor market outcomes remains scarce.

This paper uses this setting to estimate whether temporary work assignments to undesirable areas can induce workers to pursue public sector careers and serve low-income communities. For identification, we exploit random variation in the geographic assignment of psychologists, the only profession still subject to random assignment under SERUMS. We examine effects on subsequent labor market choices, and on the job-related preferences and beliefs that may explain them. Our empirical strategy takes advantage of exogenous variation in psychologists' job choice sets generated by the program. Participants select their 12-month service location in a randomly determined order: those drawn first chose from the full set of available posts, while those drawn later select from a progressively smaller set of remaining options. As a result, individuals with worse lottery draws face more limited choices and are more likely to take positions in poorer or otherwise less desirable areas. This random ordering generates plausibly exogenous variation in exposure to disadvantaged locations, which we exploit to estimate the reduced-form impact of program assignments on subsequent job choices and to explore the mechanisms that may drive these effects.

Our data come from two main sources: administrative records from the Ministry of Health in Peru on health worker placements and an online survey we conducted among psychologists who had completed the rural service program over the last three years. In the survey, we collected information on employment outcomes, the participant's reported rank order during the assignment progress, and potential mechanisms that could mediate results (e.g., perceived demand-side explanations and changes in supply-side preferences and beliefs after the rural service program). We use these to create a prosociality index based on stated job preferences, redistributive preferences, and (past and hypothetical) donation decisions. We supplement the surveys with administrative data on public sector employees.

We confirm that psychologists with a worse lottery draw completed the rural service program in more vulnerable, rural, and poorer locations. We then show that this experience affects their subsequent career outcomes. Compared to psychologists who were given first choice over locations, respondents in the bottom lottery-rank tercile are 8.9 percentage points, or 17%, more likely to work for the public sector after their service concludes. These effects are concentrated among workers in facilities that provide healthcare services to poor households. In addition to changes in type of employer, we also observe a change in

respondents' later career location. Psychologists in the bottom lottery-rank tercile are 5.8 percentage points, or 54%, more likely to subsequently work in the poorest districts in the country. Consistent with this result, we find similar effects for a subsample of psychologists for which we observe true lottery rankings through government video recordings and later employment outcomes in administrative data.

We then test a series of explanations for what drives these career effects. One possibility is that the experience changes the jobs for which these psychologists believe they are competitive: those who complete the program in poorer areas might worry that they received worse training, making them less attractive to private-sector employers. A second possibility is that service changes psychologists' professional networks: workers assigned to poorer areas might perceive fewer opportunities to develop networks that help them find a job. A third possibility is that assignment changes psychologists' perceived availability of job options across sectors.

We find little support for these channels. We first rule out inertia: less than 5% of respondents continue to work at the facility or district where they completed the rural service program. We then find no evidence that service in less desirable areas changes a psychologist's perceived hireability, professional networks, or perceived availability of job options across sectors. First, assigned rank tercile does not affect participants' perceptions of how the rural service affected their skill formation as valued by public- and private-sector employers. Second, we do not find that respondents in the bottom lottery-rank tercile give different answers about whether networks helped them get their current job. This null is consistent with network formation being a relatively minor feature of the program overall: only 20% of respondents report that someone they met through the program helped them get their current job. Finally, we show that lottery rank does not affect perceived availability of job options in the public and private sectors.

In contrast, we find several pieces of evidence that service shapes a worker's willingness to work in rural areas and changes their beliefs and attitudes regarding others. Psychologists in the bottom lottery-rank tercile are more likely to say that the rural service program increased their willingness to work for public facilities that help the poor, and are more likely to accept a job in a poor area. They also score 0.11 standard deviations higher on a prosociality index

than psychologists in the top tercile. We further find that psychologists in the bottom tercile change their beliefs about how effective the public sector is at achieving its goals, and about the underlying reasons for poverty. Effects are larger among respondents born in Lima — the capital and Peru's most prosperous area — consistent with a first-time-exposure channel, though we treat this as suggestive given the number of subgroups examined.

This paper contributes to multiple bodies of research. First, our findings supplement existing work on policies supporting spatial market integration to address imbalances in the geographic distribution of public service delivery. Previous work has focused on the role of financial incentives (Glewwe et al., 2010; Muralidharan and Sundararaman, 2011; Duflo et al., 2011; Ashraf et al., 2014; Costa et al., 2024). Although it has been standard policy to attract workers to undesirable locations with higher wages, research has shown this approach may be insufficient or costly (Costa et al., 2024). Even more so, evidence suggests they may crowd out another feature critical for effective service delivery: having motivated, prosocial applicants who are willing to serve these communities and most likely to stay (Bobba et al., 2021; Dal Bó et al., 2013; Ashraf et al., 2014; Deserranno, 2019). More importantly, these policies treat workers' preferences and beliefs as given, despite recent evidence suggesting that mandatory assignment to an unfamiliar region can reduce familiarity bias and expand workers' willingness to live outside their origin communities long-term (Okunogbe, 2024). Our paper is, to the best of our knowledge, the first to show that governments can engineer and expand the pool of high-skilled workers willing to work in the public sector and poorer areas of the public sector through temporary service programs. This provides support for the expansion of temporary service programs, a common feature of many developing country policy toolkits. In related work, Dobbie and Fryer Jr (2015) show that participation in the volunteer service program "Teach for America" makes participants more likely to work in the education sector after the program. Because they identify off the extensive margin of participation, their design cannot separate the effect of assignment itself from other program components, such as training content. We address this by exploiting within-program variation in assignment intensity in a public policy implemented at scale, across a broader range of employment outcomes.

Second, our paper speaks to existing work on the long-run effects of early career expe-

riences (Oreopoulos et al., 2012; Angrist and Chen, 2011; Fadlon et al., 2022; Dobbie and Fryer Jr, 2015; Rasul, 2026). Our paper is particularly related to Fadlon et al. (2022), who find that female physicians in Denmark who received unfavorable internship positions are more likely to sort into less desirable local labor markets in the long run, partly due to worse professional networks. We interpret this divergence as arising from two key differences in our setting. First, our service program is a mandatory, uniform point of entry into the public sector labor market, while Fadlon's internships mark entry into a specialty career ladder where subsequent advancement depends heavily on professional networks. Second, "unfavorable" placement in our context means direct exposure to rural poverty in a developing-country setting, rather than a less desirable position within an otherwise similar career track; this distinction is particularly relevant given our focus on prosocial preferences and beliefs about poverty as channels. Consistent with the first distinction, we find no evidence that lottery rank affected either network formation or the role networks played in obtaining respondents' current job, reinforcing that the returns to such networks are far less relevant to labor market sorting in our context. In contrast, we find a variety of evidence suggesting that changes in supply-side preferences and beliefs, linked to prosociality in particular, are a key driver of the career effects we estimate. Our paper thus contributes to this literature by identifying a channel, prosocial motivation activated at public sector labor-market entry, through which assignment shapes long-run sorting, distinct from the network-based mechanism previously documented.

Finally, our paper contributes to the literature on the endogeneity and malleability of beliefs and preferences by showing that contact with disadvantaged communities can shift high-stakes behavior, not merely attitudes. Existing work in this literature largely documents effects on attitudes and beliefs — social justice, inequality, prosociality, and social behaviors — that rarely translate into consequential behavioral change (Lowe, 2021; Rao, 2019; Londoño-Vélez, 2022; Mo and Conn, 2018; Mousa, 2020; Kremer et al., 2019; Kosse et al., 2020; Dobbie and Fryer Jr, 2015; Rivera, 2025). A notable exception is Okunogbe (2024), who study a similar mandatory national service program in Nigeria intending to promote national integration and find that exposure to other ethnic regions increases migration, interethnic romantic relationships, and national pride. We extend this evidence on

hard behavioral outcomes to a labor market setting: workers' choice of employer and sector — arguably a higher-stakes, more policy-relevant behavior than migration or relationship formation. While the source of variation in their paper is ethnicity, our paper's exposure considers a bundle of characteristics that make an area undesirable, including poverty, remoteness, and other vulnerability factors such as altitude, temperature, and transportation access. Our paper thus weaves together the insight of contact theory by showing (1) that its logic generalizes beyond ethnicity to other forms of interpersonal difference, and (2) that it can shift hard, policy-relevant behavior — public-sector labor supply — not just attitudes or beliefs.

The rest of the paper is organized as follows. Section 2 describes the institutional setting and the characteristics of the service program. In Section 3, we describe the data and survey components. Section 4 presents the study design. Section 5 reports our main results on program placement and career choices. Section 6 introduces the conceptual framework guiding the mechanisms analysis results in Section 7. Section 8 presents heterogeneity of effects based on participant characteristics and lottery types, and Section 9 contains conclusions.

2 Background

This section provides background on the healthcare sector in Peru and describes the institutional features of the rural service program.

2.1 The Health Workforce in Peru

Similar to many other countries, Peru faces a geographical imbalance of health workers. The health worker density is almost twice as high in urban than in rural areas (PAHO, 2017). Appendix Figure A1 shows the relationship between the number of psychologists in the public sector per 100,000 people and a district-level poverty index.[1] The index is created by FONCODES, a government agency, based on poverty indicators in the 2007 census.[2] The

[1]Districts are the third-level administrative unit in Peru and have a median population of 4,366 people according to the 2007 census.

[2]The index is based on the following indicators: the share of homes without drinking water, the share of homes without proper sanitation, the share of homes without electricity, the share of households with at least one malnourished member, the illiteracy rate of women aged 15 and over, and the population share of children aged 0-12 years.

index ranges from 0 to 1 and indicates the percentile rank of the district relative to the rest of the country. The government uses this ranking to classify districts into poverty quintiles such that each quintile has roughly the same population.[3] We use the poverty index and its quintiles throughout the paper as our preferred measure of location classifications. The blue line in Appendix Figure A1 shows that the richest districts have around 22 psychologists per 100,000 people, whereas the poorest districts only have a psychologist density of around 22 in the richest districts of around 13. At the same time, poorer districts show lower rates of mental healthcare utilization. Conditional on having a mental health condition, the share of people who seek care in the poorest districts is about half the share in the richest districts (orange line). Combined with lower psychologist density, this suggests poor districts face a compounding disadvantage: fewer providers, and lower care-seeking even among those who need it.

We find a similar pattern when looking at a proxy of unfilled vacancies at public facilities. In the poorest-quintile districts, 51% of public primary facilities mandated to provide mental healthcare services lack a psychologist, compared to only 17% in the richest-quintile districts. This is consistent with unmet demand for psychologists in remote health centers,[4] despite the existence of bonuses for all employers of primary care facilities.[5]

2.2 The Mandatory Rural Service Program

To address the shortage of healthcare professionals in rural areas, the government launched the *Servicio Rural y Urbano Marginal de Salud* (SERUMS) program in 1981. The law requires that healthcare professionals work for one year in a rural or semi-urban area in order to be eligible to work in the public sector.[6] The assignment of applicants to healthcare facilities differs across professions. For professions with a nationwide exam, such as physi-

[3]Since poorer districts tend to have a smaller population, the poorest quintile consists of more districts.

[4]Evidence from a Regional Health Management Office in Lima (DIRIS) illustrates this pattern: in 2023, a call for 36 psychologists across 27 health centers received 152 eligible applicants but resulted in 5 health centers with unfilled positions at the end of the process, all located in rural and semi-urban areas. We do not have comparable administrative records for other regions, so this case is illustrative rather than representative.

[5]The delivery of these bonuses is subject to a set of requirements, including completion of specific training courses and a second specialization aligned with primary care.

[6]Some professions, including physicians, are also required to complete SERUMS to apply for residency programs.

cians and nurses, facilities are assigned based on merit ranking. In contrast, professions without a nationwide exam are assigned facilities through a randomized lottery. We focus on psychologists, the largest profession for which assignment was determined by a lottery.[7]

Psychologists are eligible to apply to the program after the completion of a 5-year bachelor's degree. According to anecdotal evidence, many psychologists still apply to the program even if they do not intend to immediately work in the public sector to maximize future career options. A back-of-the-envelope calculation suggests that roughly 70% of psychology graduates apply.[8]

The tasks of psychologists during the rural service program are the same nationwide and range from the diagnosis and treatment of mental health conditions, such as depression, anxiety, schizophrenia, and alcoholism, to providing support to victims of domestic violence.

While all positions are located in rural or semi-urban areas, there is substantial variation in the poverty level of different positions in the program. Appendix Figure A2 shows the geographical distribution of the poverty level of each location. Positions are available across the country and only in the top three poverty quintiles. All positions are paid the same (nominal) salary except for approximately 20% of positions that are located in emergency or border areas and receive around 30% higher salaries. [9]

[7]Until 2024, psychology and nutrition were the only professions assigned by lottery. These professions were not assigned by merit because the associations responsible for developing and administering the exams lacked the resources to implement a standardized nationwide exam. In 2025, the Ministry of Health assumed responsibility for exam creation and administration; since then, all professions have been assigned based on merit.

[8]The calculation is based on the following statistics: (i) university records document that there were 2,537 psychology graduates in 2015, (ii) there were 1,723 SERUMS positions for 3,505 unique psychology applicants across both lottery rounds in 2019, (iii) the average respondent in our survey applied to SERUMS three times. If we assume that the numbers are roughly constant across years and that psychologists stop applying after five times, we would predict that $\sum_{i=0}^{4} 2537 * 0.51^i = 4998$ psychologists should apply in a given year. Comparing this to the actual number of applicants in 2019 leads to an application rate of 70%.

[9]The government offers bonus points, which can be used for preferential treatment in future applications to government jobs, to health workers applying to jobs in districts with different poverty quintiles (see Appendix Table A1 for details). However, since 87% of positions are located in the two poorest quintiles, the variation in bonus points is relatively small. Most of the psychologists in our pilot activities were not aware that the bonus points applied to them, and only 26% of our survey respondents say that the bonus points were important for them when they chose their SERUMS location. Bonus points are very important for other professions like physicians and nurses since they are also relevant for residency program applications.

The assignment process. The program is heavily oversubscribed for psychologists: in the first round of 2021, 3,886 applicants competed for 336 paid positions.[10] Overall, only around 10% of applicants to paid positions are selected for employment via the program annually. Assignments occur twice per year. Approximately one month before each lottery, the government publishes a list of available vacancies. This list includes only limited information: the geographic location of the facility, the managing institution or provider, facility size, the poverty quintile of the location, and whether the position is in an emergency or border area. Among these characteristics, the poverty quintile is the most salient, as applicants commonly use it as a proxy for remoteness and safety. Applicants may also supplement this information through personal networks and social media groups.

Figure 1 describes the application and assignment process. Conditional on applicants meeting the eligibility requirements, each applicant has to apply to one of twelve regional lottery sites based on the location of their university. During registration, applicants choose the provider or institution type, within their assigned lottery site, with whom they would like to work.[11] Applicants within each lottery site-provider cell are then assigned positions via lottery. There are 12 regions that serve as sites and up to four participating providers—the Ministry of Health (MINSA); the local government; the social contribution-based system; and either the police (PNP), army, marines, or national airforce.[12] In the lottery, software is used to draw random numbers of participants. When a participant's number is called, they are asked to select a facility in which to work from the set of remaining vacancies within the region-provider cell.[13] All lottery choices are completed before lottery numbers are assigned, and the random draw takes place. Each lottery is conducted separately by region–provider combination within a given cycle, and applicants cannot enroll in more than one lottery nor

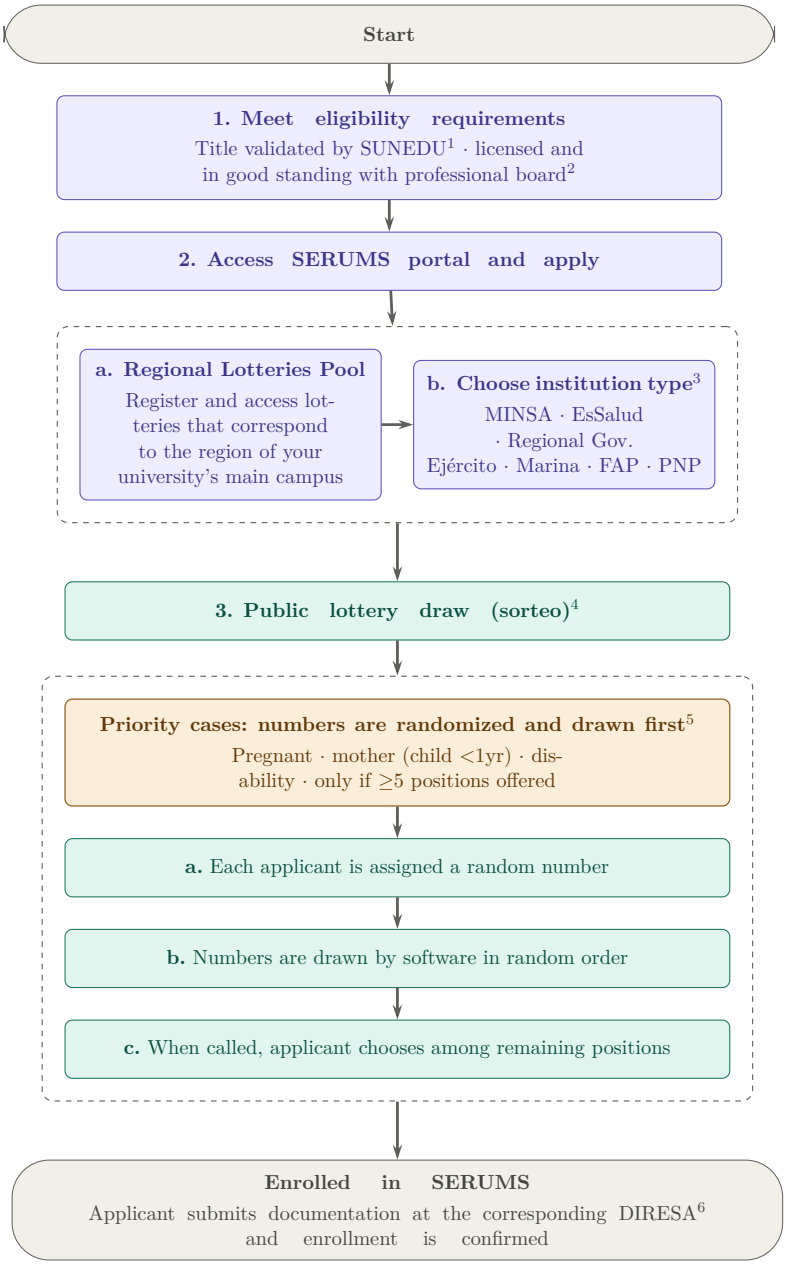
[10]The program also offers unpaid positions. We focus on paid positions given that unpaid positions are usually offered in richer locations and assigned separately. Paid positions are only available in districts classified into the three poorest poverty quintiles.

[11]The application process does not require that applicants state their ordered preferences beforehand, although they may do so independently.

[12]This yields up to 48 separate lotteries on each lottery cycle. However, in practice 87% of SERUMS positions for psychologists are located in Ministry of Health facilities. Regional lottery sites typically offer positions within a single region, whereas the central lottery offers positions across eleven regions. Fifty-two percent of applicants are assigned through the central lottery site.

[13]Exceptions include pregnant women, women with infants, and applicants with disabilities, who are allowed to choose first. These cases are removed from our analysis.

Figure 1: SERUMS Assignment Process for Psychologists



[1] *SUNEDU*: Superintendencia Nacional de Educación Superior Universitaria — the national body that validates university degrees in Peru.
[2] "Licensed and in good standing" translates *colegiado* (registered member of the professional board) and *habilitado* (license currently active and not suspended).
[3] Each regional lottery always includes positions from three providers — MINSA, EsSalud, and the Regional Government — plus positions from one of four possible Armed Forces or Police institutions (Army, Marines, Air Force, or Police). With 26 regions and up to 4 providers, this implies up to 104 lotteries taking place simultaneously during each process.
[4] The *sorteo* applies to psychologists and nutritionists. Physicians, nurses, and other health professionals are assigned by merit order (GPA + national examination score).
[5] Priority cases restricted to quintil 3 placements (quintil 2 if unavailable). Priority cases are dropped from our main sample.
[6] *DIRESA*: Dirección Regional de Salud (Regional Health Directorate). Replaced by *DIRIS* (Dirección de Redes Integradas de Salud) in Lima Metropolitan Area.

find out how many participants enroll in a given lottery before the draw takes place.

Figure A3 illustrates this process using a hypothetical lottery for Amazonas region facilities provided by the Ministry of Health. At the start of each lottery, all applicants face the same set of available facilities within the chosen region and provider. Numbers are drawn, and applicants are called sequentially in random order to select from the remaining positions. As the lottery proceeds, the set of available facilities shrinks, so applicants drawn later face more limited options.

The randomization is conducted digitally by the central government using the website `random.org`. When called, applicants may withdraw or may be absent, in which case the lottery proceeds to the next number. If a participant misses their opportunity to choose because of either withdrawal or absence, they are permanently removed from that period's lottery, though they are still permitted to either enter a separate lottery for unpaid positions held the following week or wait six months for the next lottery cycle. In practice, active withdrawal from paid positions is rare. More commonly, applicants are absent when called (10%) or make procedural errors (5%), such as failing to bring the required documentation or attempting to select a facility that has already been assigned for the third time. As discussed later, we find no significant differences in acceptance rates by lottery position. The process ends once all facilities have been assigned.

Career choices after service in the SERUMS program. After completing their service, psychologists can make job choices without further restrictions. Less than 5% of respondents continue to stay at the facility in which they completed SERUMS. Psychologists have various career options. According to our survey, 68% work in the health sector after SERUMS. Another 18% work in the education sector, while the rest work in a variety of areas, including public administration, marketing, and human resource management. Within the healthcare sector, most psychologists either work for the Ministry of Health, the social contribution-based system, or private facilities. The facilities of the Ministry of Health provide healthcare to beneficiaries of the public social protection scheme, *Seguro Integral de Salud* (SIS), consisting of poor households and informal workers. ESSALUD is a public contribution-based social security system for salaried workers with a separate network of facilities. Private facilities consist of individual practices or larger hospitals for which patients

either have to pay out-of-pocket or receive coverage through private insurance schemes. As shown in Figure A5, the large majority of poor households receive healthcare from Ministry of Health facilities. Ministry of Health facilities are usually also the only formal healthcare providers in poorer districts (Appendix Figure A4). For psychologists, the salaries across the different institutions are similar. Instead, the main difference between the institutions is that vacancies at the Ministry of Health tend to be located in more remote areas. In our sample, 58% of private sector jobs are located in the two richest district categories, whereas only 24% of public sector jobs are located in the same areas.

Multiple studies have shown that health workers overall have strong preferences to live in cities (Miranda et al., 2012). Skewed perceptions about life in rural areas are also likely to contribute to location decisions, since health workers tend to come from urban and more affluent socio-economic backgrounds. Relative to the general population, psychologists in our sample come from more urban areas (78% vs 70%), have more educated parents (67% vs 47% have mothers who completed secondary school), and are less likely to speak an indigenous language as their mother tongue (7% vs 21%). Those who still decide to work in poor areas tend to be more prosocial.

3 Data

Our primary data source is an online survey that we conducted between 2022 and 2024. Our main sample consists of 1,560 psychologists who have completed a paid SERUMS assignment at a Ministry of Health facility between November 2019 and November 2023. For this sample, we obtained the name, email address, phone number, university, SERUMS placement, and SERUMS lottery details from the Ministry of Health.

All respondents received an initial email with the link to the online survey. They were informed that the survey was conducted by an independent research team to analyze potential improvements to the SERUMS program. After the initial email, we used a call center to make follow-up calls and remind respondents to fill out the survey. Respondents could also request to be sent a WhatsApp message with the survey link. In addition to the follow-up calls, we sent two reminder emails over the course of 1.5 weeks. In the final stage of the data collection, we conducted an abbreviated phone survey with respondents who had not filled

out the online survey by that time.[14]

Fifty-five percent of the main sample either completed the online or the phone survey, which is high compared to similar studies (Dobbie and Fryer Jr, 2015; Mo and Conn, 2018; Wu et al., 2022). In our preferred specification, we always pool phone survey responses, partial online survey responses, and completed online survey responses. We show in the appendix that the results also hold when we restrict the sample to respondents who completed the online survey.

Lottery Rank: Unfortunately, the administrative data we access do not contain information on lottery rank. Consequently, our primary measure of an applicant's position in the lottery is derived from self-reported rank during the lottery assignment in survey responses. We ask respondents about the number of positions that were available, their position in the lottery, and the number of places left when it was their turn to choose. If the answers were internally inconsistent or the reported total number of positions was inconsistent with administrative data, we tried to call back the respondents to clarify their answers.

Ex-ante, there is no reason for a candidate to strategically misremember their position in the lottery: unlike qualifications or performance measures, one's rank at the moment of choosing carries no information about the quality of the candidate, and thus no incentive to misreport it favorably. For the same reason, self-reported rank is unlikely to be systematically distorted by social desirability bias. The absence of a strategic incentive to misreport does not rule out non-strategic measurement error, however. To assess this, we obtained Zoom video recordings of 46 lottery draws conducted between 2020 and 2022, covering a total of 387 applicants, and compared actual rank from these recordings to self-reported rank for the subsample also present in our survey data. We use the actual rank from the lottery recordings whenever possible in our main measure of lottery rank and also to conduct internal validity checks using this subsample. Reassuringly, we find that the self-reported and observed lottery rank for the video recordings subsample is highly correlated (the Pearson correlation coefficient is equal to 0.7), and 76% of respondents say that they are certain or very certain about their answers regarding the lottery assignment. The average difference between self-

[14]We further surveyed a subset of respondents who had applied to SERUMS in May 2021 but did not receive a position. We use the data from this survey for descriptive statistics in the paper.

reported and actual rank is not statistically distinguishable from zero; however, self-reports are compressed toward the center of the true rank distribution, which is likely to attenuate estimates that rely on the self-reported measure.[15]

Employment: Our employment module asked whether the respondent worked for pay in the past week and, if they worked, in which sector, institution, and district their current job is located. The survey also asked about secondary jobs at other institutions as well as the respondent's monthly salary for each job (based on six salary categories). We supplement the employment outcomes with administrative data on the universe of Ministry of Health employees in Peru.

Job preferences and beliefs about job attributes: We asked respondents to rate six additional job attributes from 1 to 7 according to their importance. These measures were salary, work-life balance, intellectual satisfaction, compatibility with a partner, work environment, and local infrastructure. We further asked respondents to compare positions at the Ministry of Health and in the private sector according to their salary, their opportunities to help the poor, and their work-life balance on a 5-point scale. Similarly, we also ask them to compare jobs at the Ministry of Health in the poorest and richest districts according to their local infrastructure, their opportunities to help the poor, their safety levels, and their compatibility with a partner. We further collected the respondent's opinions on whether they agree that the SERUMS experience increased their willingness to work for the Ministry of Health and to work in rural areas.

Measures of Prosociality: We collect measures of prosociality across three dimensions. First, we ask respondents about how important it is for them to have a job that provides opportunities to help the poor on a scale from 1 to 7. Second, we ask respondents whether they think that the government should increase aid to the poor. Third, we collect four different measures of donation behavior. We ask respondents whether they donated time or money to an organization that supports those in need in the past 30 days. We also adopt a measure from the Global Preference Survey that asks how much the respondents would

[15]Figure A6 compares self-reported rank to actual rank observed in Zoom recordings of a subsample of lottery draws. Out of the 357 applicants observed in the Zoom recordings, 177 are also present in survey responses. Regressing self-reported on actual rank yields a slope of 0.77 ($p < 0.001$), significantly different from one, and the gap between self-reported and true rank widens systematically for applicants with worse true ranks, suggesting mean-reverting measurement error.

donate to a good cause if they were to receive 300 soles (≈ 158 USD PPP) unexpectedly (Falk et al., 2018). Previous research has shown that this is correlated with incentivized decisions in a dictator game (Falk et al., 2016). In addition, we conducted a dictator game in the 2022 survey and a spectator game in the 2023 survey.

Skills and network formation: We obtained information on how much the respondents believe that the SERUMS experience improved their theoretical knowledge, clinical skills, language abilities, and professional networks. We ask whether these changes in skills directly helped the respondents to get their current jobs. For networks, we included the following question: "did any member of your professional network that you met through SERUMS help you find or get hired at your current job?". We also obtained information on the perceived availability of jobs by asking whether respondents agree that they could get a job at the Ministry of Health, at ESSALUD, and in the private sector if they wanted.

4 Empirical Strategy

Our empirical strategy exploits the randomized order in which psychologists applying to region-provider cells, via the SERUMS program, are able to select from the set of available work facilities. Within each lottery, applicants are randomly ordered and choose sequentially from a common set of available facilities. Applicants drawn early in the lottery can choose from the full set of positions, whereas applicants drawn later face progressively more restricted choice sets, with the final applicant who receives a position choosing only whether or not to work in the last remaining facility. This mechanism generates exogenous variation in applicants' choice sets as a function of their position in the lottery (i.e., lottery rank).

We restrict the analysis to Ministry of Health lotteries with at least six available facilities to ensure sufficient variation in lottery position. We further exclude individuals classified as special cases who were allowed to choose first and therefore did not participate in the randomized assignment mechanism.

We estimate the following specification for psychologist i participating in lottery j:

$$Y_{i(j)} = \alpha + \gamma f(PercentileRank_{i(j)}) + \delta_j + X'_{i(j)}\beta + \varepsilon_{i(j)}. \quad (1)$$

The variable $PercentileRank_{i(j)}$ captures applicant i's position in the lottery, expressed as a percentile within lottery j. The function $f(\cdot)$ is specified flexibly. We present results using both a linear specification and indicators for terciles of the percentile rank distribution, which do not impose a linear relationship. The outcome variable $Y_{i(j)}$ denotes the post-SERUMS outcome of interest for individual i. The term δ_j represents lottery fixed effects, defined at the level of site $\times$ institution type $\times$ lottery year, which absorb all common features of the choice environment faced by applicants within a lottery. The vector $X_{i(j)}$ includes predetermined individual characteristics: age, gender, an indicator for having attended a public university, an indicator for Spanish as a native language, and an indicator for whether the survey was administered by phone. Robust standard errors are reported.

Our identification relies on three main assumptions. First, *random assignment of lottery position*: within each SERUMS lottery, applicants are randomly ordered, implying that percentile rank is independent of applicants' potential outcomes and baseline characteristics. We provide empirical support for this assumption by showing that predetermined characteristics are smooth across the distribution of percentile ranks. Columns 1–2 of Table 1 show that true lottery rank is not significantly correlated with gender or university type. We also do not observe systematic differences in acceptance rates across lottery ranks (columns 3-6).

Second, *no manipulation of rank*: applicants cannot influence their position in the lottery, and the ordering is determined mechanically by the assignment procedure. The Zoom recordings support this assumption. Figure A9 displays the use of the website `random.org` to draw the random sequence of numbers that defines applicant's order of choice. In addition, before the lotteries were recorded on Zoom, regulations required the presence of a notary to certify the transparency of the process.

Third, *monotonicity of choice sets*: earlier lottery positions weakly expand the set of available facilities relative to later positions, as applicants choose sequentially without replacement, and they only have one time to choose a single facility. Under these assumptions, differences in outcomes across lottery positions can be interpreted as causal effects of facing more favorable choice sets.

Outcomes and lottery rank are measured using survey data, raising the possibility of selective survey response. Selection into the survey would bias our estimates if response

Table 1: Internal Validity Checks

	Balance		Assignment Decision			
	Female	Public University	Accepted	Rejected	Mistake	Absent
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Linear Specification						
Percentile Rank	0.037	-0.019	-0.083	0.028	0.015	0.040
	(0.070)	(0.044)	(0.063)	(0.034)	(0.032)	(0.049)
Panel B: Discrete Specification						
Middle Tercile	0.050	0.070**	0.018	-0.023	-0.003	0.008
	(0.047)	(0.034)	(0.042)	(0.020)	(0.020)	(0.033)
Bottom Tercile	0.015	-0.010	-0.047	0.021	0.013	0.014
	(0.048)	(0.029)	(0.046)	(0.026)	(0.024)	(0.035)
pval: Middle Tercile = Bottom Tercile	0.485	0.028	0.149	0.073	0.480	0.867
Outcome Mean	0.787	0.065	0.810	0.049	0.031	0.110
Observations	418	418	444	444	444	444

Notes: This table tests for balance in baseline sociodemographic characteristics (female, enrolled at a public university) and application outcomes (acceptance, rejection, mistakes – application errors when it was time to choose, and absenteeism during the draw) across lottery rank, using the full sample of applicants observed in Zoom lottery recordings (N = 444). Each column reports coefficients from a separate regression of the column variable on lottery rank, controlling for lottery stratification fixed effects. Panel A reports coefficients on continuous percentile rank; Panel B reports coefficients on lottery-rank tercile dummies, with the top tercile as the omitted category. Robust standard errors appear in parentheses. The row "pval: Middle Tercile = Bottom Tercile" reports the p-value from a test of equality between the middle- and bottom-tercile coefficients. * p<0.10, ** p<0.05, *** p<0.01.

probabilities vary systematically with lottery position and are correlated with outcomes and individual characteristics. We address this concern in several ways. First, lottery position is determined prior to both facility choice and survey participation, making it unlikely that rank directly affects survey response except through post-SERUMS outcomes. Second, we show that observable baseline characteristics conditional on survey completion, are balanced across percentile ranks (Table 2), as well as survey initiation and completion (Table A2). Third, results are robust to restricting the sample to respondents who completed the survey. Together, these findings mitigate concerns that selective survey participation drives our results.

Table 2: Balance Check Conditional on Survey Participation

	Female	Public University	Age at SERUMS Start	Born in Lima	Native Language is Spanish
	(1)	(2)	(3)	(4)	(5)
Panel A: Linear Specification					
Percentile Rank	-0.038	-0.047	-0.623	-0.061	-0.025
	(0.054)	(0.036)	(0.650)	(0.049)	(0.028)
Panel B: Discrete Specification					
Middle Tercile	-0.025	0.002	-0.436	-0.027	-0.056**
	(0.037)	(0.025)	(0.465)	(0.034)	(0.023)
Bottom Tercile	-0.030	-0.021	-0.491	-0.035	-0.015
	(0.039)	(0.025)	(0.476)	(0.035)	(0.021)
pval: Middle Tercile = Bottom Tercile	0.895	0.414	0.915	0.832	0.120
Outcome Mean	0.778	0.200	29.873	0.198	0.958
Observations	793	793	768	709	710

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a dummy for whether the survey was completed by phone. This table tests for differential survey response by lottery rank among applicants who participated in the survey, regressing each baseline characteristic (columns 1–5) on lottery rank. Panel A reports coefficients on continuous percentile rank; Panel B reports coefficients on lottery-rank tercile, with the top tercile as the omitted category. All regressions include lottery stratification fixed effects and a dummy for whether the survey was completed by phone. The row "pval: Middle Tercile = Bottom Tercile" reports the p-value from a test of equality between the middle- and bottom-tercile coefficients. Sample sizes vary across columns due to item non-response on individual survey questions. Robust standard errors appear in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

5 Main Results

This section first characterizes participants' choice sets and program placements as a function of their position in the lottery. We then examine how receiving a later lottery draw affects labor market outcomes after the program ends.

5.1 Exogenous Choice Sets and Revealed Preferences

As described earlier, participating psychologists choose their postings in a lottery: applicants are drawn in random order, and each one picks a facility from those still available. To understand what a later draw means for an applicant, this section uses the exact draw positions observed in the Zoom recordings to reconstruct the set of facilities available at each

applicant's turn, and it analyzes the characteristics of that set and of the facility chosen from it. The analysis uses 332 regular applicants in 29 lotteries with three or more applicants.[16]

Throughout, the Poverty Index refers to the government's 2006 deprivation index, which ranges from 0 to 1 and is used to classify vulnerability at the district level. The Vulnerability Index is a multidimensional score created by the Ministry of Health in 2023 to classify vulnerability at the facility level. We use it as a z-score, standardized over all public facilities in the country.[17] Higher values of either index indicate poorer or more vulnerable postings. As in the reduced-form results, draw position is grouped in terciles of (pick − 1) divided by the number of spots filled; the top tercile contains the first picks.

The analysis proceeds in three steps. First, we compare the set of facilities left at an applicant's turn with the starting choice set. Second, we separate the effect of the set left from the effect of the applicant's own choice within it. Third, we estimate a conditional logit model of the facility chosen on the attributes of the facilities left, to reveal preferences over facility attributes. We express each estimate as the extra distance applicants would travel to avoid that attribute.

The results show three things. First, the set of facilities left at an applicant's turn is less desirable than the starting choice set: earlier applicants take the less poor and less vulnerable facilities. Second, the effect of the set outweighs the effect of choice: later applicants end up in poorer, more vulnerable postings because the set they face is less desirable, not because they choose less desirable postings within it. Third, within the set they face, applicants avoid more vulnerable postings: a facility one standard deviation more vulnerable must be about 50 km closer to be as attractive.

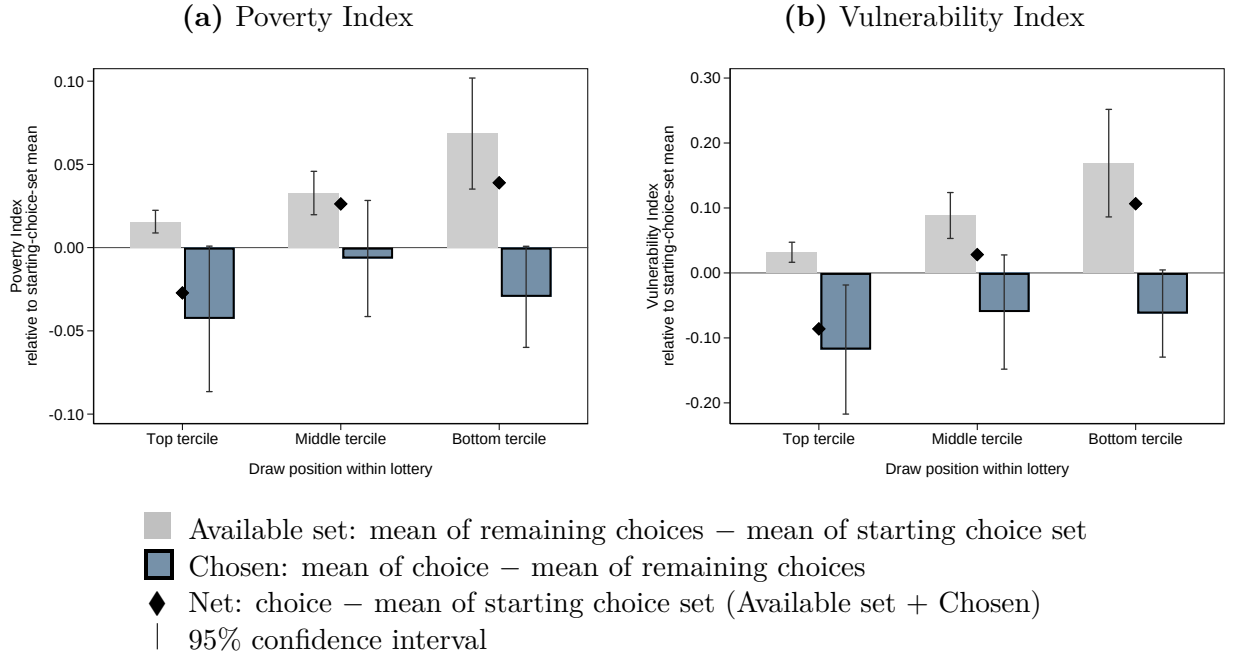
Figure 2 presents the first two findings by decomposing the gap between the chosen facility and the mean of the starting choice set into two parts: the gap between the mean of the facilities still available at the applicant's turn and the starting-choice-set mean (Available set), and the gap between the chosen facility and the mean of the facilities still available

[16]There are also 32 priority applicants (pregnant applicants and applicants with disabilities), who pick first. Because they take vacancies, they are kept when reconstructing the remaining sets, but they are not counted in the position results. The 17 lotteries with fewer than three applicants (11 of them with a single applicant) are excluded from the position analysis.

[17]The raw vulnerability score has mean 48.3 points and standard deviation 19.0 points across 9,125 facilities, and it is the sum of eleven components, including monetary poverty, altitude, internet access, the basic health team and transport conditions.

(Chosen). Their sum is the Net. In the top tercile, the facilities still available are close to the starting choice set: their mean is only 0.02 above the starting-choice-set mean for the Poverty Index and 0.03 standard deviations above it for the Vulnerability Index. Figure A10 shows the same pattern for the median and most desirable facility left in the set. The applicant's choice dominates, so the Net is negative: applicants end up in less poor and less vulnerable postings than the mean of the starting choice set (-0.03 for the Poverty Index and -0.09 standard deviations for the Vulnerability Index). As the draw order progresses, the available set becomes less desirable, and its confidence intervals exclude zero in every tercile. The Chosen bars are negative in every tercile. The choice effect is largest in the top tercile, where there is the most to choose from (-0.04 for the Poverty Index and -0.12 standard deviations for the Vulnerability Index), and smaller later on. In the bottom tercile, the available set adds 0.07 to the Poverty Index and 0.17 standard deviations to the Vulnerability Index, while the applicant's choice removes 0.03 and 0.06, leaving a net of +0.04 and +0.11 standard deviations, respectively. The set drives the gap; applicants' own choices partly offset it.

Figure 2: Available set and chosen facility, by draw position



Notes: Means across applicants in each tercile of draw position, relative to the mean of the applicant's starting choice set. Capped lines are 95% confidence intervals. Higher values indicate poorer or more vulnerable postings.

To measure how applicants trade off facility attributes, we estimate a conditional logit model. Psychologist i in lottery j obtains utility

$$U_{i(j)k} = z_k'\theta + \theta_d\, d_{i(j)k} + \eta_{i(j)k}, \qquad k \in R_{i(j)}, \tag{2}$$

from facility k in $R_{i(j)}$, the set of facilities left at the applicant's turn, where z_k are the facility's attributes, $d_{i(j)k}$ is its distance from the applicant's origin in 100 km, and $\eta_{i(j)k}$ is independent and identically distributed type I extreme value. The applicant picks the facility with the highest utility, so the probability that applicant i chooses facility k is

$$\Pr\left(c_{i(j)} = k \mid R_{i(j)}\right) = \frac{\exp(z_k'\theta + \theta_d\, d_{i(j)k})}{\sum_{\ell \in R_{i(j)}} \exp(z_\ell'\theta + \theta_d\, d_{i(j)\ell})}. \tag{3}$$

The model fits the lottery for two reasons. First, the serial dictatorship is strategy-proof: applicants gain nothing by picking anything other than their preferred facility among those left, so the facility chosen reveals their ranking of the set. Second, the set $R_{i(j)}$ depends on the applicant's draw position, which is random, and not on the applicant's own preferences, so the choice sets are exogenous. We report each attribute as a distance equivalent, $100 \times \theta_z/\theta_d$: the number of kilometers closer a facility must be to offset a one-unit increase in attribute z.

Table 3 presents the results. The facility attributes in the model are the Poverty Index, the Vulnerability Index, travel time from the main regional health office (DIRESA), the share of rural population in the district, and the facility tier (I-1 to II-1). We present four specifications: the Poverty Index alone, the Vulnerability Index alone, both indices together, and all attributes together. All four include distance. The sample is 296 of the 332 applicants: 26 have only one facility left at their turn, and 10 chose a facility with a missing attribute. We report randomization p-values, which come from 1,000 draws in which each applicant's choice is replaced by a random pick from the set left at their turn. The results show that applicants avoid more vulnerable facilities holding distance fixed (randomization $p < 0.01$): a facility one standard deviation more vulnerable must be 52 km closer to be as attractive. The estimate is stable when the Poverty Index is added (48 km, $p = 0.06$)

and when all attributes are included (48 km, $p = 0.11$). The Poverty Index alone has a negative coefficient holding distance fixed (8 km per 0.1, $p = 0.07$), but it goes to zero once the Vulnerability Index is included ($p = 0.78$). This is consistent with the two indices being correlated within choice sets (0.59). Part of the poverty selection in Figure 2 is distance: poorer facilities are also farther from the applicant's origin ($p = 0.02$ without distance vs $p = 0.07$ with it). Finally, travel time from the regional office is the strongest attribute after distance (40 km per point, $p < 0.01$), and we find no selection on facility tier or rural share.

Altogether, results from this analysis show that the draw is an exogenous shift in exposure to more vulnerable facilities, working through a less desirable choice set. The reduced-form effects can therefore be read as the effects of being placed in a more vulnerable facility. Within any set, applicants avoid farther and more vulnerable postings; the Vulnerability Index absorbs the effect of the Poverty Index, and the two cannot be separated in this sample.

Table 3: Choice model: chosen facility and attributes of the facilities left at the applicant's turn

	(1)	(2)	(3)	(4)
Poverty Index	−0.53		−0.10	−0.18
	(0.31) [0.07]		(0.37) [0.78]	(0.52) [0.74]
Vulnerability Index		−0.32	−0.30	−0.25
		(0.13) [<0.01]	(0.15) [0.06]	(0.16) [0.11]
Travel time from the DIRESA (points)				−0.21
				(0.05) [<0.01]
Rural population share				0.25
				(0.41) [0.57]
Facility tier (1 = I-1 to 5 = II-1)				0.07
				(0.10) [0.49]
Distance from origin (100 km)	−0.64	−0.62	−0.62	−0.53
	(0.13) [<0.01]	(0.13) [<0.01]	(0.13) [<0.01]	(0.13) [<0.01]
Distance equivalent (km closer to offset)				
0.1 higher Poverty Index	8		2	3
	[0, 20]		[−9, 13]	[−14, 28]
1 SD higher Vulnerability Index		52	48	48
		[15, 110]	[5, 109]	[−7, 128]
1 point more travel time				40
				[19, 77]
0.1 higher rural share				−5
				[−23, 9]
1 tier higher				−13
				[−58, 23]
Applicants	296	296	296	296
Applicant × facility pairs	2,813	2,813	2,813	2,813

Notes: Conditional logit of the facility chosen on the attributes of the facilities left at the applicant's turn. Robust standard errors in parentheses; randomization p-values in brackets (1,000 draws replacing each choice with a random pick from the same set). Distance equivalent: km closer a facility must be to offset the stated change, $100 \times (\theta_z/\theta_d) \times \Delta z$, with 95% bootstrap confidence intervals in brackets (applicants resampled within lottery). Excludes 26 applicants with one facility left and 10 whose chosen facility lacks an attribute.

5.2 Placement During the Program

As described in Section 4, the SERUMS assignment process follows a one-sided lottery in which psychologists are randomly ordered and select sequentially from a common set of available health centers within the lottery they self-selected into. Applicants drawn earlier face a larger set of available positions, while those drawn later choose from increasingly restricted choice sets. Because gross salaries are fixed across the majority of positions, differences in placement characteristics by lottery rank reflect variation in non-wage attributes of the

available positions.

We examine whether lottery rank predicts the characteristics of program placements. Table 4 shows that respondents who selected later in the lottery chose positions in smaller facilities with fewer staff (column 1), located in poorer locations (column 2), with higher shares of rural population (column 3), and with higher vulnerability. The vulnerability index is an index created by the Peruvian government that measures community-level vulnerability across different dimensions, including altitude, temperature, internet access, physician availability, monetary poverty, usage of native languages, and risks associated with transportation, such as type, travel time, and road risks.

Table 4: Effect of Lottery Rank on Program Placement

	Village Clinic	Poverty Index	% Rural Population	Vulnerability Index
	(1)	(2)	(3)	(4)
Panel A: Linear Specification				
Percentile Rank	0.049*	0.089***	0.057*	0.207***
	(0.026)	(0.028)	(0.030)	(0.073)
Panel B: Discrete Specification				
Middle Tercile	0.013	0.044**	0.021	0.068
	(0.017)	(0.020)	(0.021)	(0.052)
Bottom Tercile	0.042**	0.063***	0.039*	0.136***
	(0.018)	(0.021)	(0.021)	(0.053)
pval: Middle Tercile = Bottom Tercile	0.166	0.376	0.443	0.247
Outcome Mean	0.035	0.547	0.554	-0.045
Observations	792	792	792	791

Notes: Robust standard errors in parentheses. All regressions include lottery stratification fixed effects and controls for gender, public university, and a dummy for wether the survey was done over the phone. The Vulnerability Index is a Peruvian government index (2023) measuring community-level vulnerability across altitude, temperature, internet access, physician availability, monetary poverty, native language, transportation, travel time, and road risks.

Panel A shows the linear specification results, while Panel B the tercile dummies specification. Relative to psychologists whose lottery rank falls in the top tercile (Panel B), those in the bottom tercile are 4.2 percentage points more likely to choose small village clinics and that rank 6.3 percentile points higher in the government poverty index, have a rural popu-

lation share 3.9 percentage points higher, and score 0.13 standard deviations higher in the vulnerability index. This is consistent with findings from the linear specification reported in Panel A. Altogether, these results show that more undesirable positions – defined as smaller, poorer, riskier, and more remote positions– are left for participants choosing last. This, in turn, could be generating meaningful differences in job environments that have the potential of shaping subsequent labor market outcomes.

5.3 Labor Market Outcomes

We next examine the effects on subsequent career outcomes (Table 5). While later lottery draws do not seem to affect labor market participation, we find that respondents in the bottom rank tercile are 8.9 percentage points more likely to work for the public sector than respondents in the upper rank tercile (column 2). These effects are driven by an increase of 7.7 percentage points in Ministry of Health employment (column 4), the main government institution that provides healthcare services to the poor. Column 3 shows that the gains in public sector employment come at the expense of private sector employment (including private sector work at secondary jobs). Later lottery draws reduce the likelihood of working in the private sector by 8.3 percentage points, almost the same magnitude of the effect reported in column 2. On average, participants' labor market outcomes are retrieved one year after the program ended.

Table 5: Effect of Lottery Rank on Post-Service Employment Sector

	Any Work	Public Sector	Private Sector	Ministry of Health
	(1)	(2)	(3)	(4)
Panel A: Linear Specification				
Percentile Rank	0.023	0.136**	-0.122**	0.124**
	(0.055)	(0.061)	(0.054)	(0.060)
Panel B: Discrete Specification				
Middle Tercile	-0.063	0.011	-0.080**	0.044
	(0.041)	(0.047)	(0.039)	(0.045)
Bottom Tercile	0.005	0.089**	-0.083**	0.077*
	(0.039)	(0.044)	(0.040)	(0.043)
pval: Middle Tercile = Bottom Tercile	0.131	0.115	0.940	0.508
Outcome Mean	0.774	0.537	0.254	0.324
Observations	710	710	710	710

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, and a dummy for whether the respondent's university was public, and a dummy variable that indicates whether the survey was done over the phone.

Table 6 documents our estimates of the relationship between lottery ranking and type of employer and job attributes, such as location (columns 1 and 2), profession (column 3), and nominal and real net salaries (columns 4 and 5). To calculate wages in real terms, we use the spatial deflator produced by the Peruvian National Institute of Statistics to account for differentials in purchasing power between rural and urban districts.

Column 2 shows that inertia is not driving employment effects. On average, less than 5% of psychologists stay in the same district in which they did the SERUMS program. This also does not differ by lottery rank, ruling out this mechanism as a potential driver. There isn't evidence that other job attributes are changing, most participants keep working as psychologists after the program independently of where they did the program. Despite no evidence of nominal and real wages differing by lottery rank, large point estimates provide suggestive evidence of increased nominal and real wages for participants in the bottom lottery tercile, consistent with a statistically significant increased likelihood of working in the poorest districts for this group. Column 1 in Panel B shows that the share of respondents who work in the poorest districts is 5.8 percentage points higher, or 53%, for psychologists in the

bottom rank tercile ($p = 0.066$).

Table 6: Effect of Lottery Rank on Post-service Job Attributes

	Works in Poorest Districts (1)	Same District as SERUMS (2)	Works in Health (3)	Nominal Salary (PEN) (4)	Real Salary (PEN) (5)
Panel A: Linear Specification					
Percentile Rank	0.085*	0.020	0.016	165.321	175.668
	(0.044)	(0.025)	(0.065)	(184.342)	(224.941)
Panel B: Discrete Specification					
Middle Tercile	0.041	0.013	-0.005	-223.450	-266.271
	(0.033)	(0.021)	(0.048)	(140.142)	(169.957)
Bottom Tercile	0.058*	0.012	0.007	121.089	127.413
	(0.032)	(0.019)	(0.047)	(134.092)	(164.672)
pval: Middle Tercile = Bottom Tercile	0.622	0.987	0.814	0.020	0.030
Outcome Mean	0.108	0.043	0.512	1885.827	2277.816
Observations	689	689	703	620	620

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, and a dummy for whether the respondent's university was public, and a dummy variable that indicates whether the survey was done over the phone. Outcomes are unconditional on employment status.

A potential concern is that differential selection into the survey could explain our results. While we find no differences in attrition rates by lottery rank, it is still possible that different types of respondents choose to participate in the survey. We address this concern by replicating column 4 of Table 3 with the subsample of psychologists for which we observe the lottery ranking in the video rankings. We then combine this information with administrative data on the universe of Ministry of Health employees, up to three years after the program ended. Reassuringly, we find similar increases in Ministry of Health employment for the bottom-rank tercile group in this subsample (column 1, Table 7). Columns 2 and 3 suggests participants in the bottom lottery tercile are also 12.9 percentage points more likely to work in the bottom poverty quintiles and worked on average 2.4 months more for the Ministry of Health. Figure 3 also shows this effect is stable one, two, and three years, respectively, after the program ended for all participants in our sample. Taken together, these results

Figure 3: Effect Over time on Ministry of Health Employment



Notes: Reduced form estimates. The figure plots the coefficient on an indicator for being drawn in the bottom tercile of the lottery draw, relative to the top tercile (omitted category). The outcome is an indicator for working at the Ministry of Health, measured one, two, and three years after the program. All regressions include lottery stratification fixed effects and controls for gender, public university, and phone survey mode. The sample consists of psychologists for whom we observed the lottery ranking in the official video recordings and did the program, and the outcome comes from administrative data on the university of Ministry of Health employees. Whiskers denote 90 and 95% confidence intervals .

document a substantial change in the employment status of respondents up to 33 months after the end of the program, based on their rank in the SERUMS lottery.

6 Conceptual Framework

How can assignment to more undesirable areas affect later labor market outcomes? Labor market outcomes are the product of two-sided matches. On the demand side, the private and public sectors compete for newly graduated psychologists. While private health centers set wages to attract the worker types that maximize profits, the public sector's objective function is more complex, and its wages are often assumed to be exogenously determined (Costa et al., 2024). Because these wages are set administratively rather than to clear the market—and do not fully compensate for the disamenities of remote, poor postings—desirable posts are over-subscribed and must be rationed by a non-price mechanism; the SERUMS lottery is that mechanism, and the common preference ranking it implies—all else equal, participants

Table 7: Effect of Lottery Rank on Ministry of Health Employment

	Works for Ministry of Health	Works for Ministry of Health (Bottom Poverty Quintiles)	Total Months Worked at Ministry of Health
	(1)	(2)	(3)
Panel A: Linear Specification			
Percentile Rank	0.144*	0.167**	3.115*
	(0.080)	(0.085)	(1.672)
Panel B: Discrete Specification			
Medium Rank Tercile	0.082	0.106*	1.129
	(0.057)	(0.059)	(1.132)
Bottom Rank Tercile	0.133**	0.129**	2.416**
	(0.056)	(0.060)	(1.208)
Outcome Mean	0.577	0.369	6.615
Observations	357	357	357

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects, and a vector of controls that consists of gender and a dummy for whether the respondent's university was public. The sample consists of psychologists for whom we observed the lottery ranking in the official video recordings and did the program, and the outcome comes from administrative data on the university of Ministry of Health employees.

prefer less poor posts—is what allows lottery order to map into the poverty of the assigned area.

From the supply side, workers choose a job given their preferences and beliefs over the job attributes of a given facility; these include—but are not limited to—wages, amenities, and characteristics of the location such as local infrastructure, and, especially for public service, the desire to contribute to community welfare (prosociality). Formally, a worker i who has completed assignment chooses a subsequent job j to maximize $U_{ij} = \beta_i' \, \mathbb{E}_i[\mathbf{x}_j] + \gamma_i \, \mathbb{E}_i[w_j] + \varepsilon_{ij}$, where $\mathbf{x}_j$ collects pecuniary and non-pecuniary attributes, w_j is the community-welfare content of the job, γ_i is the weight the worker places on that content (prosociality), and $\mathbb{E}_i[\cdot]$ denotes the worker's beliefs. Exposure during the program can therefore act through two distinct channels: by revising beliefs $\mathbb{E}_i[\cdot]$ over attributes and welfare content, or by shifting the preference parameters themselves (β_i, γ_i). Previous literature has shown that monetary

incentives can attract and sort talent (Bobba et al., 2021; Dal Bó et al., 2013), but may crowd out the prosocial and motivated applicants who are most likely to stay (Ashraf et al., 2014; Deserranno, 2019). This paper approaches that tension from the opposite margin: rather than asking how compensation selects prosocial workers at recruitment—where higher pay can crowd them out—we ask whether the assignment experience itself can engineer prosocial preferences or correct biased beliefs among workers already in the pipeline, offering a non-selection lever for staffing underserved areas.

The SERUMS program has the potential to affect both beliefs and preferences over job attributes and community welfare through two levels of exposure. First, because of the randomized assignment, all participants are exposed to communities and peers different from their own, as each health center receives participants from across the country. Given that the sample includes only participants in the paid version of the program—which targets facilities in rural and urban-marginal areas located in the three poorest poverty quintiles—every participant experiences some degree of exposure; the level effect of program participation itself is therefore not identified absent a comparison group that did not participate. Second, and the focus of this paper, participants are exposed to different *degrees* of community vulnerability: relative to participants drawn earlier in the lottery, those drawn later were more likely to be assigned to more vulnerable rural and poor areas, as explained in section 5.2. The lottery thus identifies the effect of the *intensity* of exposure among participants, not the effect of exposure versus none, and the estimate is best interpreted as a conservative, attenuated measure of the full effect of program participation.

Contact theory predicts that exposure to an out-group—when it occurs under conditions of equal status and intergroup cooperation toward a common goal, backed by institutional support—can reduce prejudice and increase tolerance between groups (Brown and Hewstone, 2005). In the terms above, such contact could revise beliefs about the served community or shift the prosociality weight γ_i itself. Consistent with this theory, a growing literature in economics suggests that exposure to diverse social environments can shape individuals' attitudes and beliefs about social justice, inequality, and prosociality, as well as their social behavior (Lowe, 2021; Rao, 2019; Londoño-Vélez, 2022; Mo and Conn, 2018; Mousa, 2020; Kremer et al., 2019; Kosse et al., 2020; Dobbie and Fryer Jr, 2015; Okunogbe, 2023; Rivera,

2025). It could also be, however, that because participants serve from a position of authority rather than equal status—given the nature of the doctor–patient relationship—little updating of beliefs or attitudes occurs.

Alternatively, a location's working and living conditions may lead workers to value these attributes differently or to overcome familiarity bias. For instance, despite fewer amenities, participants may find the slower pace of work in rural areas more appealing and come to value work-life balance more highly; or confronting lower-quality infrastructure and fewer peers may lead them to update their preferences over the importance of infrastructure, the work environment, or intellectual stimulation at work. More generally, the program may help participants overcome a fear of the unknown and update biased beliefs about working in vulnerable areas, which in turn may raise their likelihood of choosing to work in the public sector and in poorer areas later in their careers. Because a single experience can move beliefs and preferences at the same time—observing the scale of unmet need both updates beliefs about how many people require a psychologist's care and may raise the weight γ_i placed on serving them—our outcomes cannot fully separate the two channels. Attribution-type measures (for example, whether respondents view poverty as the result of external circumstances rather than individual effort) come closer to isolating belief change, whereas hypothetical-donation and distributive-preference measures conflate beliefs and preferences.

On the demand side, work experience in poorer versus less poor areas may lead employers to infer different productivities. Employers may assume that participants placed in more vulnerable facilities—which are smaller and dedicated exclusively to primary care—had less access to training and theoretical knowledge than those placed in larger facilities, and may therefore perceive them as less employable. Alternatively, participants in smaller facilities may have had to perform more procedures on their own and could thus be perceived as having acquired skills that make them more productive. Two features qualify this demand-side channel. First, because placement is randomized, facility type is uninformative about a worker's underlying ability, so a fully informed employer should not treat it as a signal of type; any statistical-discrimination effect therefore requires that employers be unaware of, or discount, the randomness of assignment. Second, and separately, the assignment may genuinely alter the skills a worker acquires—so facility type can affect productivity directly

even though it conveys no information about innate ability. Ex ante, the demand side's response to experience in more vulnerable locations remains unclear.

The reduced-form effect on employment estimated above is, by construction, a composite. The framework makes clear that an effect of assignment on subsequent public-sector and poor-area employment could arise on the supply side—through updated beliefs or shifted preferences over job attributes and community welfare—or on the demand side, through how employers interpret experience in more vulnerable facilities. Distinguishing these channels matters because they carry different welfare and policy implications: a supply-side effect working through prosociality would suggest that mandatory rotation can build a workforce willing to serve underserved populations, whereas a demand-side effect working through employer inference would instead reveal how the labor market prices that experience. The analysis that follows brings this distinction to the data, drawing on measures of beliefs and prosocial preferences to probe the supply side and on perceived employability to probe the employer response.

7 Mechanisms

Following the hypotheses presented in section 6, we now test the potential mechanisms underlying these effects. We first examine if demand-side factors could be shaping post-program job matches. Next, we examine effects on supply-side preferences and beliefs. Finally, we examine effects on measures of prosociality.

Perceptions of Demand — It is possible that the employment effects we document are driven by demand-side mechanisms: employers may statistically discriminate against workers placed in more vulnerable areas, or they may rationally update their beliefs based on genuine skill differences acquired during service. To shed light on these mechanisms, we examine whether placement location shapes workers' perceptions of their own human capital formation. We ask respondents two sets of questions. First, we ask whether the program improved their theoretical knowledge, practical skills, ability to communicate in a native language, or professional network, with responses ranging from "didn't improve" to "improved a lot." Second, for respondents currently employed, we ask whether each of these same dimensions — theoretical knowledge, practical skills, native language communication,

and professional network — helped them obtain their current job, with responses ranging from "fully agree" to "fully disagree." This follow-up question allows us to assess whether any particular type of skill acquired during service was especially transferable to their current job. We aggregate both sets of items into two separate standardized indexes — the "Skill Acquisition Index" and the "Skill-to-Job Transfer Index" — each constructed as a weighted average of standardized Likert scores.

Serving in a more vulnerable health center did not produce differential skill gains along any of the dimensions examined (Table 8, columns 1–5), nor do respondents report higher or lower acquisition of job-transferable skills (Table 9, columns 1–5). In addition to this, we note that only 20% of all psychologists say that someone they met through SERUMS helped them to find or get hired at their current job. This emphasizes that a rural service program like SERUMS differs from residency programs for physicians that have been examined in previous work (Fadlon et al., 2022) and that play a much more pivotal role in the network formation of participants. Despite the null results, it is worth noting that the point estimates in the linear specification tend to be negative for the Skill Acquisition Index and its sub-items, with some non-monotonic heterogeneity across terciles: participants assigned to the middle tercile show positive rather than negative point estimates relative to those in the bottom tercile. For the Skill-to-Job Transfer Index, we similarly observe that participants assigned to the middle tercile report a 0.3 standard deviation increase in the perceived importance of native language communication skills for obtaining their current job.

We complement this analysis by asking respondents to rate, on a 4-point Likert scale, how employable they believe they are across the three main employer types: Ministry of Health facilities (MoH), facilities covered by the formal social security system (EsSalud), and the private sector. Any evidence of differentiated employability perceptions despite no meaningful differences in skills gained would be consistent with statistical discrimination as a demand-side mechanism driving the employment effects. Columns 1–4 of Table 10 show no differences in the perceived likelihood of obtaining a job in the public sector (MoH or EsSalud), the private sector, or any sector. Taken together, these results suggest that the employment effects we document are unlikely to be driven by demand-side mechanisms. This interpretation is further supported by qualitative interviews with public and private sector

employers, who report that placement location during the program is not a relevant factor in their hiring decisions — what matters most to them is the worker's area of specialization and any postgraduate training pursued after SERUMS.

Table 8: Perceived Skill Gains from the Program

		Program Improved...			
	Skill Acquisition Index	Knowledge	Skills	Native Language	Network
	(1)	(2)	(3)	(4)	(5)
Panel A: Linear Specification					
Percentile Rank	-0.140	-0.172	-0.061	-0.177	-0.148
	(0.116)	(0.153)	(0.159)	(0.162)	(0.168)
Panel B: Discrete Specification					
Middle Tercile	0.057	0.116	0.011	0.119	-0.018
	(0.083)	(0.111)	(0.121)	(0.119)	(0.120)
Bottom Tercile	-0.070	-0.079	-0.007	-0.094	-0.101
	(0.082)	(0.112)	(0.114)	(0.117)	(0.121)
pval: Middle Tercile = Bottom Tercile	0.167	0.103	0.886	0.098	0.517
Outcome Mean	0.009	0.008	0.018	0.002	0.007
Observations	485	485	485	485	485

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, and a dummy for whether the respondent's university was public, and a dummy variable that indicates whether the survey was done over the phone. Columns 2-5 report effects for the index components of perceived skill gains from the program in terms of theoretical knowledge, practical skills, native language, and network. Higher values reflect higher perceived skill gains in each aspect. Each column reports standardized Likert scores, where higher values indicate a higher degree of perceived skill gains in each aspect. The skill acquisition index reports an average of all values.

Beliefs and Preferences over Job Attributes — We next turn to supply-side mechanisms. We begin by examining whether assignment to a more vulnerable post shifts workers' preferences over job attributes, starting with their willingness to work in different sectors and locations. In the 2022 survey, we ask participants whether the SERUMS program increased their willingness to work at a Ministry of Health facility and in poor districts (Table 11, columns 1–2), with responses ranging from 1 ("strongly agree") to 4 ("strongly disagree"). In the 2023 survey, we elicit the same preferences using hypothetical job offer scenarios, following Okunogbe (2024) (Table 11, columns 3–4). Respondents are first asked to choose

Table 9: Effect of Lottery Rank on Skills-to-job Transfer Perceptions

		... Skills Acquired Helped Get Job			
	Skills-to-job Transfer Index	Knowledge	Practical Skills	Native Language	Network
	(1)	(2)	(3)	(4)	(5)
Panel A: Linear Specification					
Percentile Rank	0.117	-0.057	0.148	0.320	0.057
	(0.123)	(0.194)	(0.206)	(0.205)	(0.195)
Panel B: Discrete Specification					
Middle Tercile	0.176**	0.245	0.191	0.306**	-0.037
	(0.087)	(0.153)	(0.152)	(0.155)	(0.150)
Bottom Tercile	0.076	-0.007	0.109	0.211	-0.009
	(0.090)	(0.141)	(0.150)	(0.148)	(0.137)
pval: Middle Tercile = Bottom Tercile	0.288	0.108	0.610	0.576	0.867
Outcome Mean	0.001	0.002	0.005	0.008	-0.010
Observations	343	343	343	343	343

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, and a dummy for whether the respondent's university was public, and a dummy variable that indicates whether the survey was done over the phone. Columns 2-5 report effects for the index components which disaggregate by skill type. Each column reports standardized Likert scores, where higher values indicate a higher degree of agreement. The skill-to-job transfer index reports an average of all values. Sample restricted to current workers.

between two positions at equal salary (2,500 PEN monthly): a private sector health center in Lima versus a public MINSA I-3 health center in Lima, and separately, a public MoH I-3 health center in Lima versus one in Huánuco. [18] Based on their initial response, they are then presented with a follow-up scenario in which the initially unchosen option offers a higher salary (3,500 PEN), allowing us to assess the strength of their preference. The constructed outcome variables capture whether respondents prefer the public sector option — or the Huánuco posting — even when the alternative pays more, providing a revealed preference measure of non-pecuniary job attribute valuation. This design isolates sector preference and

[18]The monthly salary of 2,500 PEN (approximately USD 668 at 2023 exchange rates) was chosen to approximate the entry-level public sector wage for psychologists at the time of the survey. I-3 health centers in Peru are public healthcare facilities regulated by the MoH that represent the third tier of complexity within the First Level of Care (which ranges from I-1 to I-4), these facilities serve as a bridge between small rural health centers and urban ones, with two basic teams permanently staffed (a physician, nurse, and OBGYN).

Table 10: Perceptions of Employability

		Perceived Employability ...		
	Perceived Employability Index	at Ministry of Health	at ESSALUD	at Private Sector
	(1)	(2)	(3)	(4)
Panel A: Linear Specification				
Percentile Rank	-0.168	-0.161	-0.192	-0.150
	(0.125)	(0.160)	(0.150)	(0.157)
Panel B: Discrete Specification				
Middle Tercile	0.005	0.046	0.043	-0.073
	(0.084)	(0.112)	(0.108)	(0.107)
Bottom Tercile	-0.065	-0.020	-0.079	-0.095
	(0.089)	(0.117)	(0.107)	(0.110)
pval: Middle Tercile = Bottom Tercile	0.463	0.600	0.296	0.848
Outcome Mean	0.009	0.011	0.012	0.004
Observations	617	617	617	617

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, and a dummy for whether the respondent's university was public, and a dummy variable that indicates whether the survey was done over the phone. Columns 2-4 report effects for the index components of perceived employability at Ministry of Health Facilities, ESSALUD, and at the private sector. Each column reports standardized Likert scores, where higher values indicate a higher degree of perceived employability in each sector. The perceived employability index reports an average of all values.

location preference separately, bringing us closer to revealed preferences by removing the confound of expected wage differences across options. [19]

Evidence suggests that psychologists assigned to the bottom tercile of the lottery rank are 9.4 percentage points more likely to strongly agree that the program increased their willingness to work at a Ministry of Health facility (Panel B, column 1, $p < 0.05$), with a consistent pattern in the linear specification: a 10 percentage point increase in the lottery percentile rank increases stated willingness to work at MoH by 1.7 percentage points ($p < 0.01$). Despite no corresponding effect on stated willingness to work in poor districts (column

[19] According to the Institution of National Statistics of Peru, Huánuco is the seventh poorest region in Peru, with a monetary poverty rate of 39.2% in 2023 compared to 28.1% in Lima. The gap is more pronounced along the multidimensional poverty dimension established by the Poverty and Human Development Initiative (OPHI): 59.8% of Huánuco residents are multidimensionally poor, compared to 12.3% in Lima.

2), an interesting pattern emerges in the job vignette outcomes. A 10 percentage point increase in the lottery percentile rank increases the likelihood of preferring a rural posting outside Lima over the capital, even when the rural option pays less (column 4, $p < 0.05$). This effect is driven by respondents in the bottom tercile, who are 16.3 percentage points more likely to prefer the rural posting ($p = 0.085$), and differs significantly from the middle tercile ($p = 0.030$), suggesting a non-monotonic pattern in location preferences across lottery ranks. The divergence between columns 2 and 4 is consistent with the distinct dimensions each question captures: column 2 targets willingness to work in a poor district specifically, while column 4 captures preferences over geographic location and urban proximity more broadly. This distinction is also apparent in baseline preferences: only 23.4% of control group participants report increased willingness to work in a poor district, compared to 65.2% who prefer a non-capital posting even when it pays less, suggesting that respondents treat these as meaningfully different job attributes. The null result in column 3 should be interpreted cautiously given the smaller sample size available for the job vignette outcomes ($n = 151$).

Tables 12 and 13 present results on additional job preferences and beliefs, respectively. Table 12 asks respondents to rate the importance of salary, work-life balance, safety, work environment, local infrastructure, compatibility with their partner's job, and intellectual satisfaction when choosing a job, on a scale from 1 ("Not important at all") to 7 ("Very important"). Table 13 focuses on attributes more likely to vary across sector and poverty dimensions, asking respondents whether they expect higher salaries, better work-life balance, and greater intellectual satisfaction in public versus private sector positions (columns 1–3), and better local infrastructure, safety, and compatibility with their partner's job in poor versus rich districts (columns 4–6). Neither preferences over job attributes nor most beliefs differ systematically across lottery ranks. The only exception is beliefs about salaries in public versus private sector facilities, which display a non-monotonic pattern across lottery ranks: bottom tercile respondents report higher expected public sector salaries relative to the private sector compared to middle tercile respondents ($p = 0.012$). One potential explanation for this pattern is heterogeneity in local private sector competition across placement areas. Workers assigned to the most vulnerable posts operate in areas with limited private sector presence, making public sector wages appear relatively attractive. In contrast, workers

assigned to intermediate areas may perceive greater private practice viability due to lower psychologist density and unmet demand, raising their expectations of private sector returns relative to public sector wages — though we lack direct evidence on private sector wage variation across placement areas to test this formally.

Taken together, these results suggest some supply-side changes in job preferences following the completion of the SERUMS program, concentrated in sector and location dimensions: workers drawn later in the lottery and thus assigned to more vulnerable posts report greater willingness to work at Ministry of Health facilities and show a stronger preference for jobs in more rural areas even when they pay less. These shifts are not accompanied by changes in other job attribute preferences — including salary, work-life balance, or intellectual satisfaction — nor in most beliefs about how these attributes vary across sectors or poverty dimensions. This pattern points toward a mechanism beyond job attribute revaluation, which we explore in the following section through the lens of what we define as prosociality.

Prosociality — We next examine how SERUMS exposure shaped the prosocial beliefs, preferences, and behaviors of psychologists. We define prosociality broadly as other-regarding orientations associated with community welfare, encompassing four dimensions: social beliefs (attribution of poverty to external circumstances versus individual effort), policy preferences (support for government redistribution), normative preferences (preference for helping vulnerable populations through work), and reported behavior (past and hypothetical donation of time and money). Table 14 presents results for each dimension separately, alongside a summary index that averages standardized scores across preference and behavior dimensions.

Relative to psychologists in the top lottery tercile, respondents in the bottom lottery tercile score 0.11 standard deviations in the prosociality index. Psychologists in the bottom lottery tercile are also more likely to agree that external circumstances and not individual-level decisions are the main reasons why people are poor ($p = 0.080$). Finally, when asked to rate public and private sector jobs in terms of which one provides more opportunities to help the poor, we observe that respondents in the bottom lottery tercile have a more favorable view of the public sector, suggesting that their experience changed their view of the effectiveness of public facilities ($p = 0.073$).

A concern with the observed increase in prosociality is that the survey questions were not

Table 11: Effect of Lottery Rank on Supply-Side Preferences

	Willingness to Work		Job Vignettes	
	Program Increased Willingness to Work at MoH	Program Increased Willingness to Work in Poor Districts	Prefers MoH over Private even when Private Pays More	Prefers Rural over Urban even when Urban Pays More
	(1)	(2)	(3)	(4)
Panel A: Linear Specification				
Percentile Rank	0.170***	-0.064	0.050	0.280**
	(0.065)	(0.070)	(0.131)	(0.131)
Panel B: Discrete Specification				
Middle Tercile	0.089*	0.063	0.097	-0.097
	(0.049)	(0.055)	(0.103)	(0.115)
Bottom Tercile	0.094**	-0.058	0.005	0.163*
	(0.047)	(0.048)	(0.102)	(0.094)
pval: Middle Tercile = Bottom Tercile	0.922	0.029	0.441	0.030
Outcome Mean	0.222	0.234	0.227	0.652
Observations	552	483	151	151

Notes: Robust standard errors in parentheses. All regressions include lottery stratification fixed effects and a vector of controls comprising gender, a dummy for whether the respondent's university was public, and a dummy for whether the survey was conducted by phone. Columns 1–2 are estimated on the 2022 survey round; columns 3–4 are estimated on the 2023 survey round. All outcomes are binary. In columns 1–2, the outcome equals one if the respondent fully agrees that the SERUMS program increased their willingness to work at a Ministry of Health (MoH) facility or in poor districts, respectively, and zero otherwise. In columns 3–4, the outcome equals one if the respondent prefers the MoH or rural option even when the alternative pays 1,000 PEN more per month (3,500 PEN vs. 2,500 PEN monthly), and zero otherwise. The rural posting in column 4 refers to a public MoH I-3 health center in Huánuco in a district classified as poverty quintile 5. I-3 refers to the third tier of complexity within Peru's First Level of Care, representing mid-sized facilities that bridge small rural health posts and larger urban centers.

incentivized. To address this, we administered a dictator game in the 2022 survey in which respondents allocated a total of 19 PEN ($\approx$ 5 USD) between themselves and two NGOs — one distributing Vitamin A to children in rural Peru and another promoting art and culture in urban areas. As shown in Appendix Table A8 (columns 4–5), we find no effects along this margin. One potential explanation is crowding out: more prosocial respondents may already support other organizations, or increased public sector work and in poor areas may substitute for alternative forms of support. Consistent with this, we find suggestive evidence

Table 12: Effect of Lottery Rank on Additional Job Preferences

	... Is Very Important for Choosing a Job						
	Salary	Work-Life Balance	Safety	Work Environment	Local Infrastructure	Compatibility with Spouse	Intellectual Satisfaction
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Linear Specification							
Percentile Rank	0.065	0.187	-0.035	-0.113	-0.122	-0.019	0.183
	(0.153)	(0.151)	(0.172)	(0.190)	(0.206)	(0.262)	(0.166)
Panel B: Discrete Specification							
Middle Tercile	0.052	0.139	-0.068	-0.132	-0.297*	-0.095	0.035
	(0.110)	(0.108)	(0.124)	(0.144)	(0.156)	(0.191)	(0.114)
Bottom Tercile	0.080	0.130	0.013	-0.093	-0.063	0.047	0.193
	(0.106)	(0.106)	(0.122)	(0.135)	(0.149)	(0.184)	(0.118)
pval: Middle Tercile = Bottom Tercile	0.819	0.939	0.559	0.794	0.162	0.491	0.197
Outcome Mean	5.953	6.000	5.928	6.104	5.316	4.837	5.995
Observations	688	690	688	685	683	534	551

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and controls for gender, whether the respondent attended a public university, and whether the survey was conducted by phone. Outcomes are measured on a 7-point Likert scale ranging from 1 ("Not Important at All") to 7 ("Very Important"). Columns 6 and 7 (spouse location and intellectual challenge) were only administered in the first survey round.

that respondents in the bottom rank tercile are more likely to report that they did not donate in the dictator game because they are already supporting other organizations or helping the poor through their current work. In the 2023 survey, we administered a spectator game instead, in which respondents allocated a fixed amount ($\approx$ 1.5 USD) across three NGOs varying in activity and location. We find suggestive evidence that respondents with lower lottery ranks are more likely to donate to an NGO working on domestic violence in rural areas, though the results are noisy and inconclusive.

Despite this caveat, the broader pattern of results suggests that placement location during SERUMS shaped the beliefs and preferences of respondents afterwards, with prosociality emerging as a plausible supply-side mechanism. Qualitative evidence corroborates this interpretation. When asked about their SERUMS experience, one respondent reflected that she "learned that it is not only enough to put all your effort to be able to study [...] but the distance from schools, the lack of public services in their villages, the type of upbringing, the training received and the example to follow greatly influence many to reject studying and/or

Table 13: Effect of Lottery Rank on Beliefs about Job Attributes

	MoH vs. Private Sector			Poor vs. Rich District		
	Salary	Work-Life Balance	Intellectual Satisfaction	Local Infrastructure	Safety	Compatibility with Spouse
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Linear Specification						
Percentile Rank	0.468***	0.256	0.027	-0.151	0.120	-0.034
	(0.178)	(0.167)	(0.162)	(0.212)	(0.202)	(0.180)
Panel B: Discrete Specification						
Middle Tercile	-0.087	0.193	-0.085	0.119	0.195	-0.033
	(0.125)	(0.122)	(0.118)	(0.158)	(0.147)	(0.134)
Bottom Tercile	0.253*	0.093	0.027	-0.094	0.079	-0.038
	(0.130)	(0.119)	(0.116)	(0.153)	(0.145)	(0.129)
pval: Middle Tercile = Bottom Tercile	0.012	0.465	0.384	0.197	0.471	0.975
Outcome Mean	3.019	2.957	3.252	2.304	2.871	2.665
Observations	622	619	618	479	479	478

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and controls for gender, whether the respondent attended a public university, and whether the survey was conducted by phone. Responses are measured on a 1–5 scale, where 1 indicates the attribute is perceived as better in the private sector (columns 1–3) or in a richer district (columns 4–6), and 5 indicates the attribute is perceived as better in a Ministry of Health facility or in a poorer district. Columns 4–6 were only asked in the first round of surveys.

progressing." Another noted that the program should be expanded so more people are "able to live that great experience in order to have a better identification with the other realities of our Peru." These accounts suggest that SERUMS assignments were formative experiences that meaningfully shifted respondents' perspectives on poverty and social inequality. This is consistent with the fact that healthcare professionals in Peru tend to come from relatively privileged socioeconomic backgrounds, meaning that for many, SERUMS represents their first sustained exposure to rural poverty.

8 Heterogeneity and Robustness

In this section, we explore heterogeneity in the effects by participants' sociodemographic characteristics to assess who is most affected by worse placements resulting from later lottery draws. Later, we present a series of robustness checks.

Table 14: Effect of Lottery Rank on Prosocial Preferences and Beliefs

		Prosociality Index Components			Prosocial Beliefs	
	Prosociality Index	Preference for Helping the Poor	Support for Government Aid	Donation Index	Poverty Due to Circumstances	MoH Offers More Opportunities to Help the Poor
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Linear Specification						
Percentile Rank	0.123	0.180	0.148	-0.021	0.092	0.166**
	(0.078)	(0.136)	(0.134)	(0.091)	(0.064)	(0.067)
Panel B: Discrete Specification						
Middle Tercile	-0.028	0.043	-0.053	-0.101	0.027	-0.010
	(0.057)	(0.100)	(0.100)	(0.065)	(0.050)	(0.051)
Bottom Tercile	0.112**	0.147	0.118	0.030	0.081*	0.086*
	(0.054)	(0.097)	(0.093)	(0.064)	(0.046)	(0.048)
pval: Middle Tercile = Bottom Tercile	0.017	0.325	0.096	0.064	0.282	0.065
Outcome Mean	-0.003	-0.008	-0.004	0.007	0.571	0.645
Observations	706	691	698	674	698	621

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, a dummy for whether the respondent's university was public, and a dummy variable that indicates whether the survey was done over the phone. Outcomes in columns 1–4 are standardized using the mean and standard deviation of the bottom tercile. The Prosociality Index is the row mean of the Preference for Helping the Poor, Support for Government Aid, and Donation Index standardized components. The Donation Index is itself the row mean of three standardized donation variables (hypothetical donation, past monetary donation, and past time donation). *Poverty Due to Circumstances* is a binary variable equal to one if the respondent attributes poverty primarily to circumstances outside a person's control. *MINSA Offers More Opportunities to Help the Poor* is a binary variable equal to one if the respondent expects significantly more or more opportunities to help the poor in a MINSA post relative to the private sector.

First, we examine the possibility that exposure to more vulnerable areas during the program generates differential treatment effects across participant characteristics (place of origin, gender, age, dependence status, and university type and quality) and lottery characteristics. To test this, we construct an outcome index as a weighted average of our main results: (1) public sector work, (2) willingness to work in poor areas, (3) willingness to work for the Ministry of Health, and (4) our Prosociality Index. Figure 4 presents the results of this analysis for the linear specification. Consistent with our previous results, we find evidence that effects on our main outcomes are concentrated among respondents with less

prior contact with rural poverty, measured as being born in Lima (Appendix Table 4), and less concentrated among respondents who completed the program in their birthplace. We also find suggestive evidence of effects among respondents from private and elite universities.[20] We do not find heterogeneous effects along other sociodemographic dimensions such as gender or having dependents.

We do not find evidence that specific lottery characteristics play a role in driving these effects, contradicting the hypothesis that selective sorting is occurring between lotteries. Specifically, effects do not vary by the level of oversubscription, the number of available vacancies, or the standard deviation of the poverty index within a lottery.[21]

An interesting result is that effects appear concentrated among respondents above 30 (the sample average). One potential explanation consistent with our main mechanism is that SERUMS represents a more formative experience for older professionals, who arrive with more established career expectations and clearer prior beliefs about the health sector. For these respondents, exposure to rural poverty during the program may represent a more significant departure from their prior experience, leading to larger updates in prosocial beliefs and preferences. Younger respondents, still in the process of forming professional identities and preferences, may be less responsive to the placement experience along these dimensions.

We conduct various robustness checks for our main results (Appendix Tables A11–A14). Our findings hold if we (i) exclude controls, (ii) restrict the sample to respondents who at least partially completed the online survey, (iii) restrict the sample to respondents who fully completed the online survey, or (iv) exclude respondents who are uncertain or very uncertain about their lottery rank. The only exception is the effect on work in the poorest districts, for which the magnitude of the coefficients remains similar but the estimates become statistically insignificant when the sample is restricted.

A further concern is that the Covid-19 pandemic may affect the external validity of the findings. We find similar but noisier results when restricting the sample to respondents who completed SERUMS before the pandemic but entered the labor market during the first

[20] Top ranked universities are defined as those appearing in a set of global university rankings.

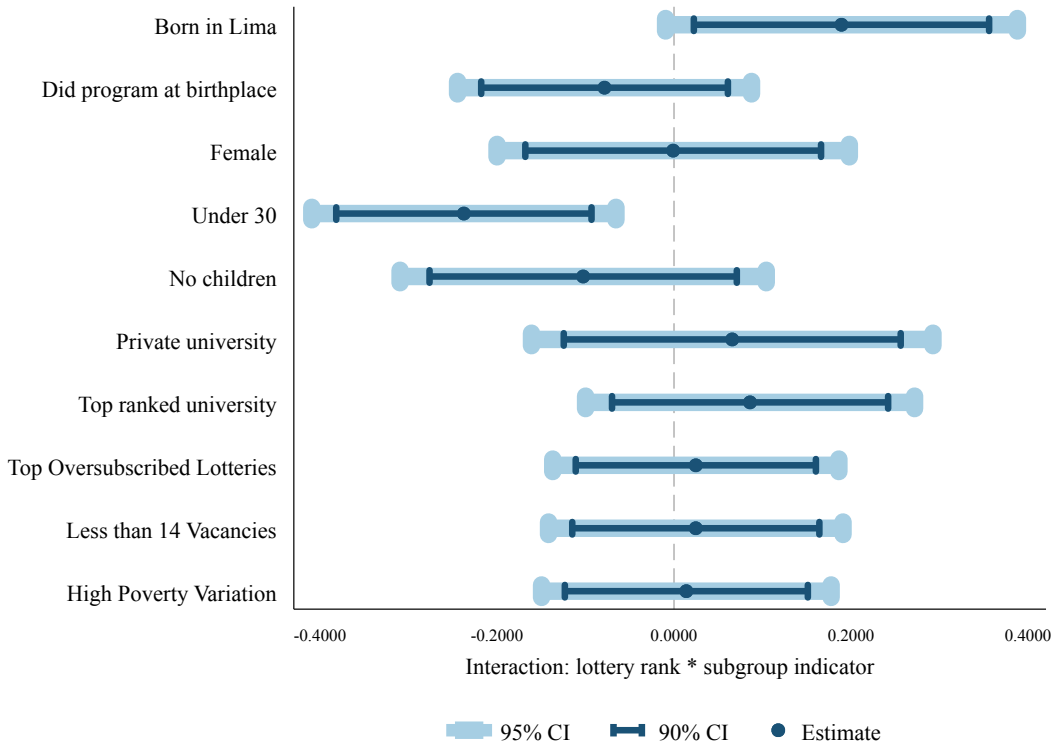
[21] Top oversubscribed lotteries are defined according to the three lottery sites appearing most oversubscribed in Figure **??**: La Libertad, Sede Central, and Lambayeque. Thresholds for a low number of vacancies and high poverty variation are defined by sample averages (14 and 0.2 SD, respectively).

wave of cases in Peru, or to respondents who completed SERUMS during the pandemic but entered the labor market after vaccines were widely available.[22] Given the reduced precision, however, we cannot rule out that the effect size differs across these subsamples.

Finally, we report p-values based on randomization inference, reassigning the lottery rank to each respondent 2,000 times (Appendix Table A10). The permutation-based p-values are somewhat larger than those from the main specification, with the Prosociality Index ($p = 0.098$) falling outside conventional significance thresholds under the discrete specification and public sector work ($p = 0.072$) and Ministry of Health employment ($p = 0.090$) falling just outside conventional significance thresholds under the percentile rank specification. Given that the corresponding OLS estimates are statistically significant at conventional levels, the larger RI p-values are consistent with the limited sample size available for this analysis

[22] An exception is the likelihood of working in the poorest districts for respondents who entered the labor market during the pandemic.

Figure 4: Heterogeneity Analysis



Notes: Robust standard errors. All regressions include lottery stratification, fixed effects, and a control for phone survey mode. The coefficient is the interaction between lottery rank (1–100) and each subgroup indicator. Outcome: row-mean of (1) Working for MoH, (2) Prosociality Index, (3) Willingness to work for MoH, (4) Willingness to work in poor areas.

9 Conclusion

This paper uses exogenous variation in choice sets within a mandatory rural service program in Peru to demonstrate that temporary work assignments can influence important policy outcomes like public sector employment, particularly in underserved areas. A suggestive mechanism analysis shows these changes are mediated by changes in worker beliefs and preferences, with prosociality playing a central role.

The paper's findings speak to two distinct literatures. First, they contribute to research on policies addressing the geographic misallocation of public service delivery, which has focused primarily on the role of financial incentives in attracting workers to underserved areas. We show that temporary service assignments are also an effective policy lever: by exposing workers to disadvantaged communities, the experience itself durably reshapes their preferences and beliefs about serving there, expanding the pool of workers willing to do so beyond what compensation alone can achieve. Second, our finding that prosociality mediates these effects contributes to a growing literature showing that contact with disadvantaged communities can increase prosocial behavior (Rao, 2019), extending it by demonstrating that such changes can affect high-stakes outcomes such as career choices.

One limitation of our identification strategy is that we cannot rule out that lottery rank itself generates psychological responses — such as disappointment or resentment from being assigned later — independent of actual placement conditions. To the extent that such responses drive our results, the estimated effects would reflect a combination of exposure to vulnerable areas and rank-based psychological reactions rather than placement experience alone. The age heterogeneity result, where effects are concentrated among older respondents, is consistent with this concern, though it does not constitute direct evidence of it. In addition to this, it should be noted that psychologists do not rely heavily on physical infrastructure, which may limit the generalizability of our findings. Physicians and nurses, by contrast, may become disillusioned with public sector work due to inadequate facilities in remote areas, potentially reversing the effects documented here. More evidence will be needed to assess the validity of our findings for other healthcare professions and service contexts.

The findings also speak to the optimal design of recruitment and posting policies for

public sector workers, Given that temporary service assignments can durably expand the pool of workers willing to serve disadvantaged communities, our results provide justification for expanding a variety of social service programs, including mandatory rural service programs like SERUMS, voluntary schemes such as Teach for America, and corporate efforts to encourage social work among employees. They further highlight the importance of accounting for the endogeneity of preferences and beliefs in economic theory and policy design more broadly. Efficient compensation schemes, for instance, might offer high lump-sum payments to encourage workers to take up postings in poor areas for at least a limited period, since preferences and beliefs can themselves be shaped by the assignment. The same logic helps resolve a potential trade-off between hiring intrinsically and extrinsically motivated candidates (Ashraf et al., 2020), since the work experience itself may strengthen a worker's intrinsic motivation over time.

While our findings speak to the role of placement experience in shaping career preferences, evaluating whether temporary service programs outperform alternative recruitment mechanisms such as financial incentives to reduce shortages, or on other dimensions such as quality of service, remains an important question for future research. Workers with greater familiarity with local communities, for instance, may be better positioned to build the trust needed to deliver effective care — a countervailing consideration our design cannot speak to.

References

Angrist, J. D. and S. H. Chen (2011). Schooling and the vietnam-era gi bill: Evidence from the draft lottery. *American Economic Journal: Applied Economics 3*(2), 96–118.

Ashraf, N., O. Bandiera, E. Davenport, and S. S. Lee (2020). Losing prosociality in the quest for talent? sorting, selection, and productivity in the delivery of public services. *American Economic Review 110*(5), 1355–94.

Ashraf, N., O. Bandiera, and B. K. Jack (2014). No margin, no mission? a field experiment on incentives for public service delivery. *Journal of public economics 120*, 1–17.

Bobba, M., T. Ederer, G. Leon-Ciliotta, C. Neilson, and M. G. Nieddu (2021). Teacher compensation and structural inequality: Evidence from centralized teacher school choice in perú. Technical report, National Bureau of Economic Research.

Boniol, M., T. Kunjumen, T. S. Nair, A. Siyam, J. Campbell, and K. Diallo (2022). The global health workforce stock and distribution in 2020 and 2030: a threat to equity and 'universal'health coverage? *BMJ global health 7*(6).

Brown, R. and M. Hewstone (2005). An integrative theory of intergroup contact. *Advances in experimental social psychology 37*(37), 255–343.

Costa, F., L. Nunes, and F. M. Sanches (2024). How to attract physicians to underserved areas? policy recommendations from a structural model. *Review of Economics and Statistics 106*(1), 36–52.

Dal Bó, E., F. Finan, and M. A. Rossi (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics 128*(3), 1169–1218.

Deserranno, E. (2019). Financial incentives as signals: experimental evidence from the recruitment of village promoters in uganda. *American Economic Journal: Applied Economics 11*(1), 277–317.

Dobbie, W. and R. G. Fryer Jr (2015). The impact of voluntary youth service on future outcomes: Evidence from teach for america. *The BE Journal of Economic Analysis & Policy 15*(3), 1031–1065.

Duflo, E., P. Dupas, and M. Kremer (2011). Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in kenya. *American economic review 101*(5), 1739–1774.

Fadlon, I., F. P. Lyngse, and T. H. Nielsen (2022). Early career, life-cycle choices, and gender. Technical report, National Bureau of Economic Research.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics 133*(4), 1645–1692.

Falk, A., A. Becker, T. J. Dohmen, D. Huffman, and U. Sunde (2016). The preference survey module: A validated instrument for measuring risk, time, and social preferences.

Frehywot, S., F. Mullan, P. W. Payne, and H. Ross (2010). Compulsory service programmes for recruiting health workers in remote and rural areas: do they work? *Bulletin of the World Health Organization 88*, 364–370.

Glewwe, P., N. Ilias, and M. Kremer (2010). Teacher incentives. *American Economic Journal: Applied Economics 2*(3), 205–227.

Kosse, F., T. Deckers, P. Pinger, H. Schildberg-Hörisch, and A. Falk (2020). The formation of prosociality: causal evidence on the role of social environment. *Journal of Political Economy 128*(2), 434–467.

Kremer, M., G. Rao, and F. Schilbach (2019). Behavioral development economics. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2, pp. 345–458. Elsevier.

Londoño-Vélez, J. (2022). The impact of diversity on perceptions of income distribution and preferences for redistribution. *Journal of Public Economics 214*, 104732.

Lowe, M. (2021). Types of contact: A field experiment on collaborative and adversarial caste integration. *American Economic Review 111*(6), 1807–44.

Miranda, J. J., F. Diez-Canseco, C. Lema, A. G. Lescano, M. Lagarde, D. Blaauw, and L. Huicho (2012). Stated preferences of doctors for choosing a job in rural areas of peru: a discrete choice experiment. *PloS one 7*(12), e50567.

Mo, C. H. and K. M. Conn (2018). When do the advantaged see the disadvantages of others? a quasi-experimental study of national service. *American Political Science Review 112*(4), 721–741.

Mousa, S. (2020). Building social cohesion between christians and muslims through soccer in post-isis iraq. *Science 369*(6505), 866–870.

Muralidharan, K. and V. Sundararaman (2011). Teacher performance pay: Experimental evidence from india. *Journal of political Economy 119*(1), 39–77.

Okunogbe, O. (2024). Does exposure to other ethnic regions promote national integration? evidence from nigeria. *American Economic Journal: Applied Economics 16*(1), 157–192.

Okunogbe, O. M. (2023). Does Exposure to other ethnic regions promote national integration?: Evidence from Nigeria. Technical report, Working Paper.

Oreopoulos, P., T. Von Wachter, and A. Heisz (2012). The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics 4*(1), 1–29.

Rao, G. (2019). Familiarity does not breed contempt: Generosity, discrimination, and diversity in delhi schools. *American Economic Review 109*(3), 774–809.

Rasul, I. (2026). Experiences and economic decision-making.

Rivera, R. (2025). Do peers matter in the police academy? *American Economic Journal: Applied Economics 17*(2), 127–164.

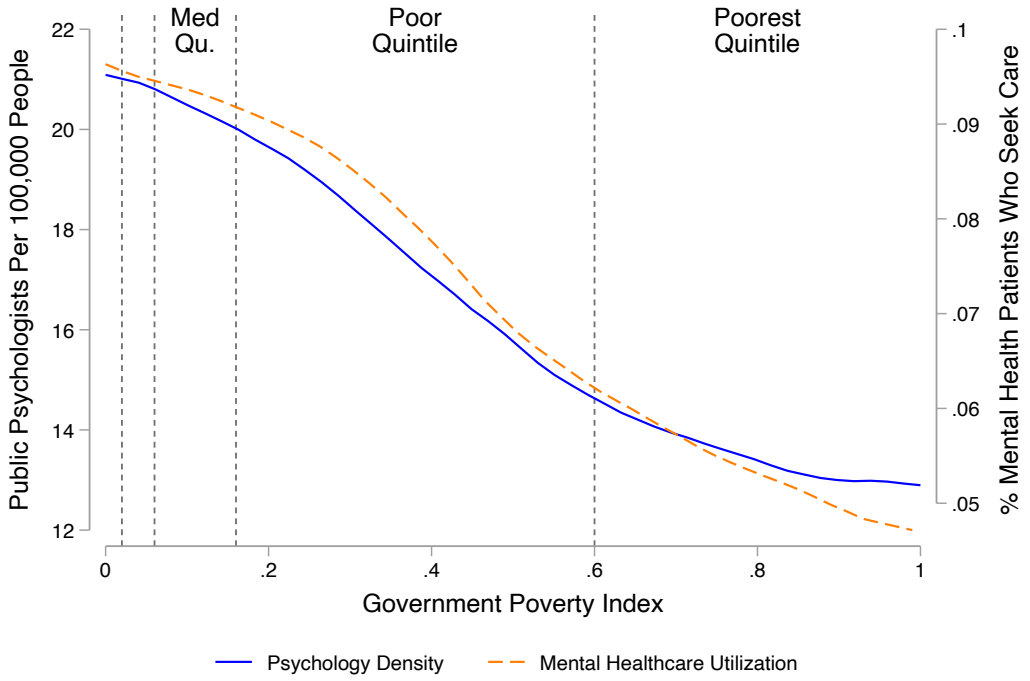
Scheil-Adlung, X. (2015). Global evidence on inequities in rural health protection : new data on rural deficits in health coverage for 174 countries. ILO Working Papers 994876213402676, International Labour Organization.

Wu, M.-J., K. Zhao, and F. Fils-Aime (2022). Response rates of online surveys in published research: A meta-analysis. *Computers in human behavior reports 7*, 100206.

A. Appendix

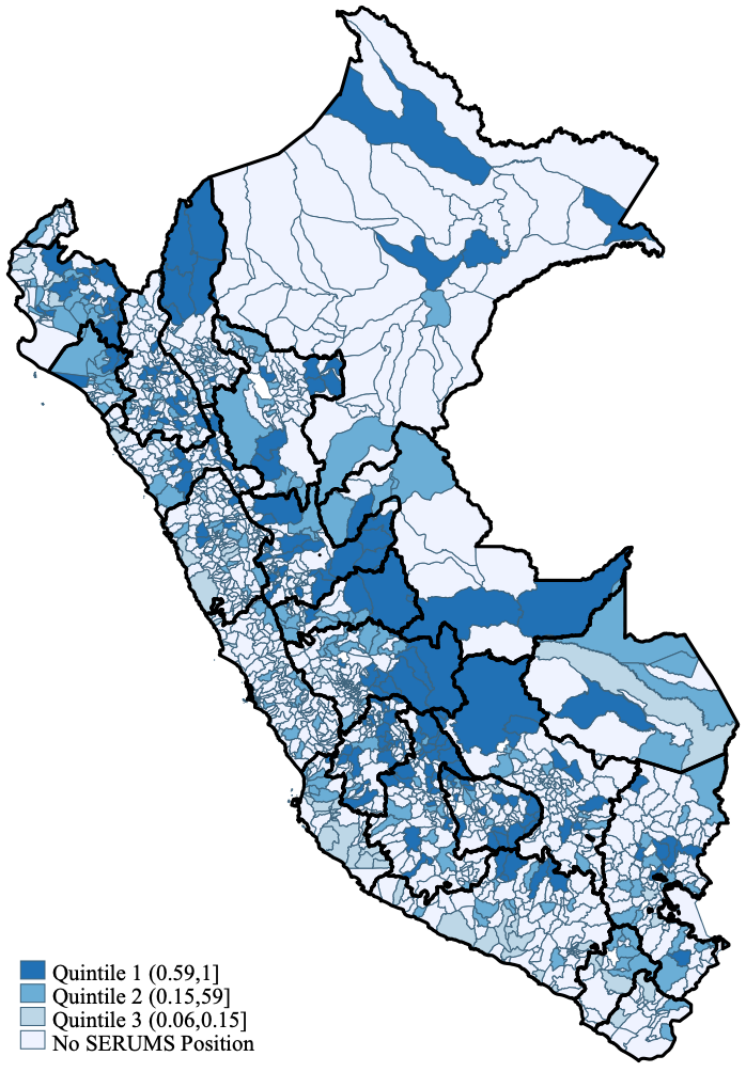
1 Appendix Figures

Figure A1: Psychologist Density and Mental Healthcare Utilization by District Poverty



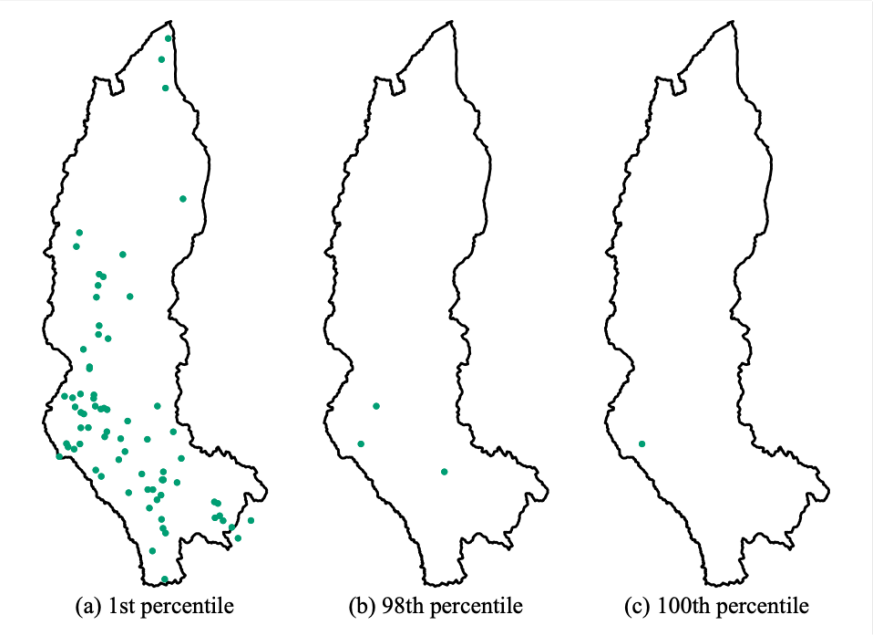
Notes: The blue line shows psychologist density (public psychologists per 100,000 people) and corresponds to the left y-axis. The orange line shows mental healthcare utilization and corresponds to the right y-axis. Psychologist density is based on administrative data from the Ministry of Health, and mental healthcare utilization is based on the National Household Survey (ENAHO). The x-axis shows the district-level government poverty index from FONCODES. The vertical lines indicate the poverty quintile cutoffs from FONCODES.

Figure A2: Geographical Distribution of SERUMS Positions by District Poverty Quintile



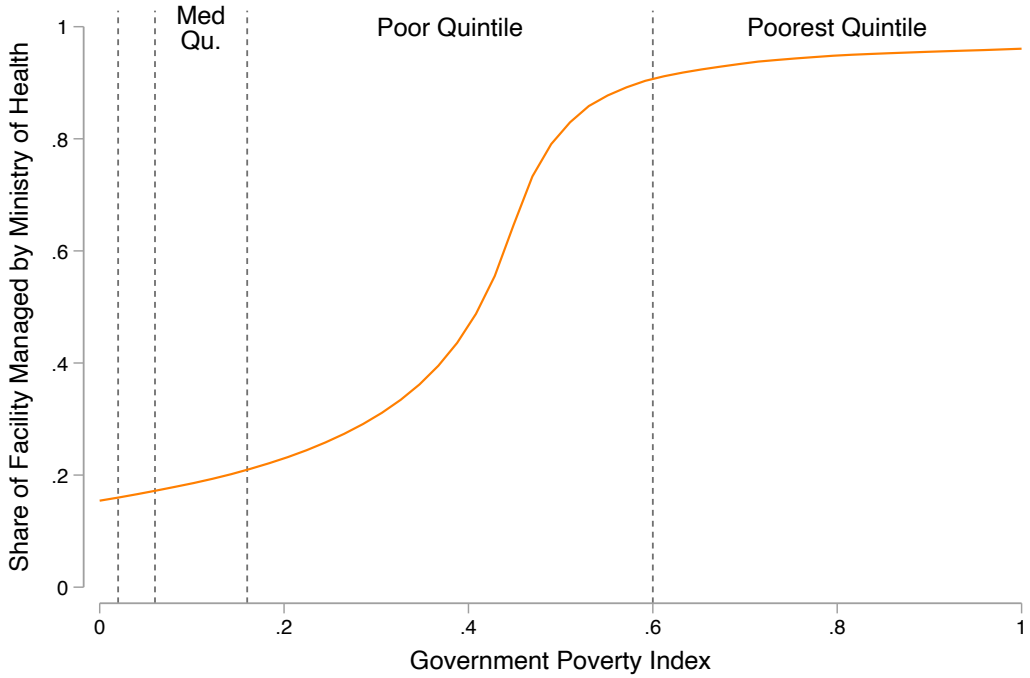
Notes: Each polygon represents a district in Peru, with dark black lines indicating regional boundaries. The figure displays the geographical distribution of SERUMS positions across the country. Districts are shaded according to poverty quintiles based on the government's poverty index. The legend reports the range of the poverty index for each quintile in parentheses. SERUMS positions were available only in districts within the top three poverty quintiles.

Figure A3: SERUMS lottery design example for the Amazonas region



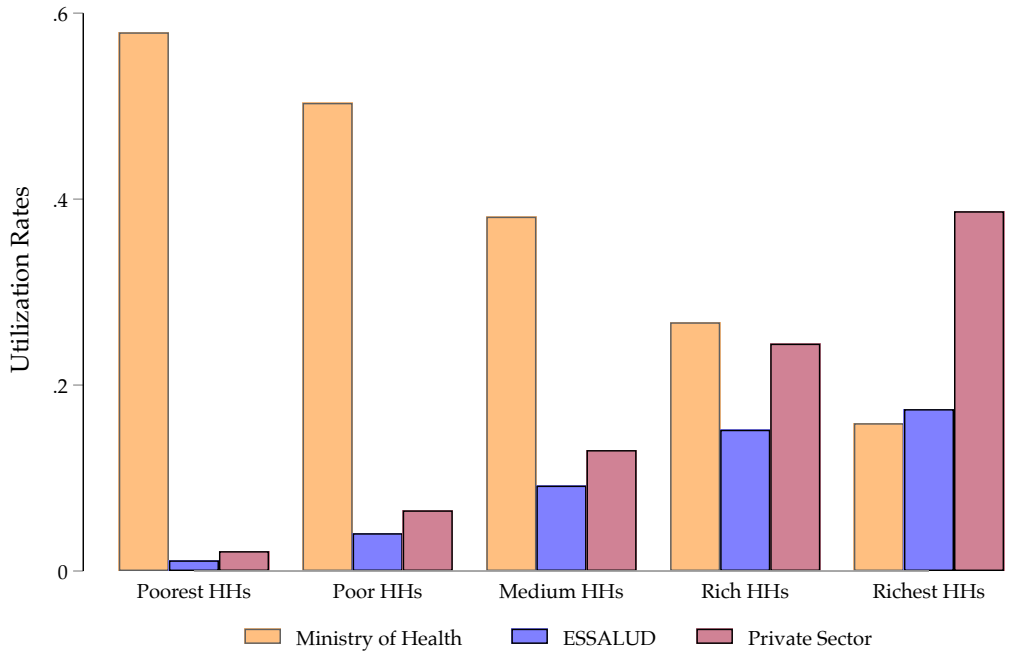
Notes: The figure presents a hypothetical lottery in the Amazonas region of Peru, illustrating the exogenous reduction in applicants' choice sets as a function of lottery rank. Panel (a) shows the full set of positions offered in lottery j by the Ministry of Health in the Amazonas region. An applicant drawn first in the lottery (i.e., at the 1st percentile of the lottery rank, excluding applicants who do not receive a position) can choose from the complete set of available positions. Panel (b) displays the positions remaining for an applicant drawn at the 98th percentile of the lottery. Panel (c) shows the positions remaining for the applicant drawn last in the lottery, for whom only one position remains available.

Figure A4: Share of Health Facilities Managed by Ministry of Health by District Poverty



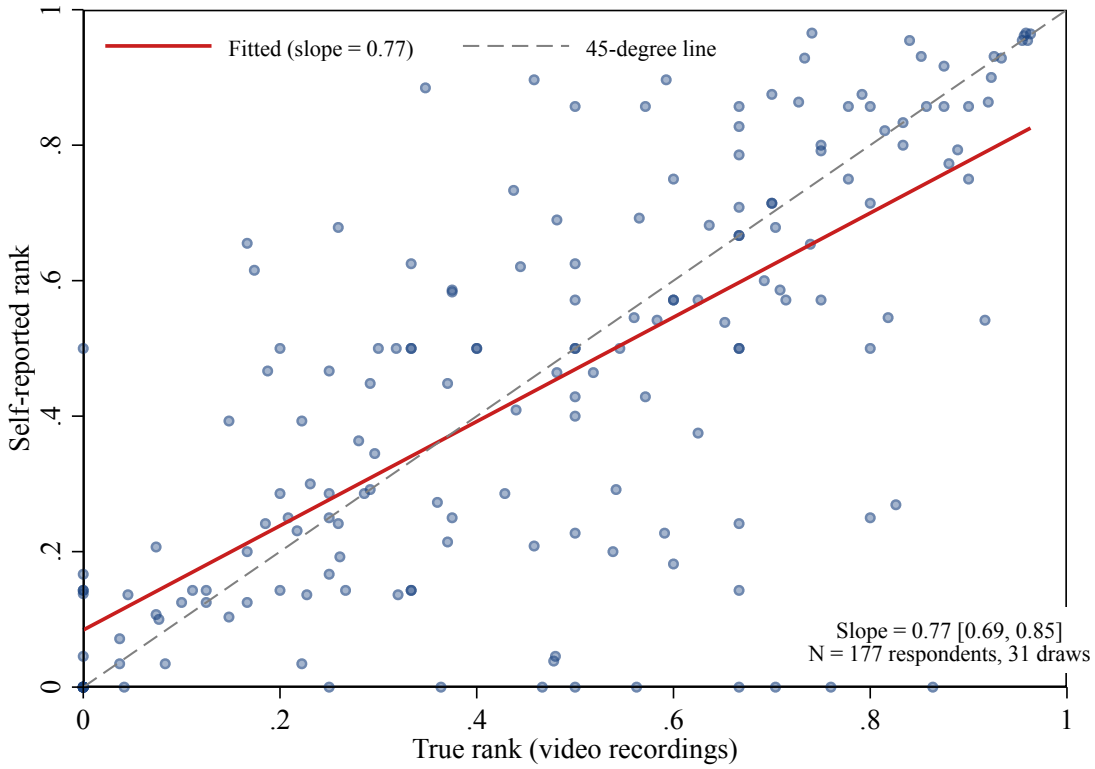
Notes: The y-axis shows the district-level share of facilities managed by the Ministry of Health according to the National Registry of Health Establishments. The x-axis shows the district-level government poverty index from FONCODES. The vertical lines indicate the poverty quintile cutoffs from FONCODES.

Figure A5: Healthcare Utilization by Household Wealth



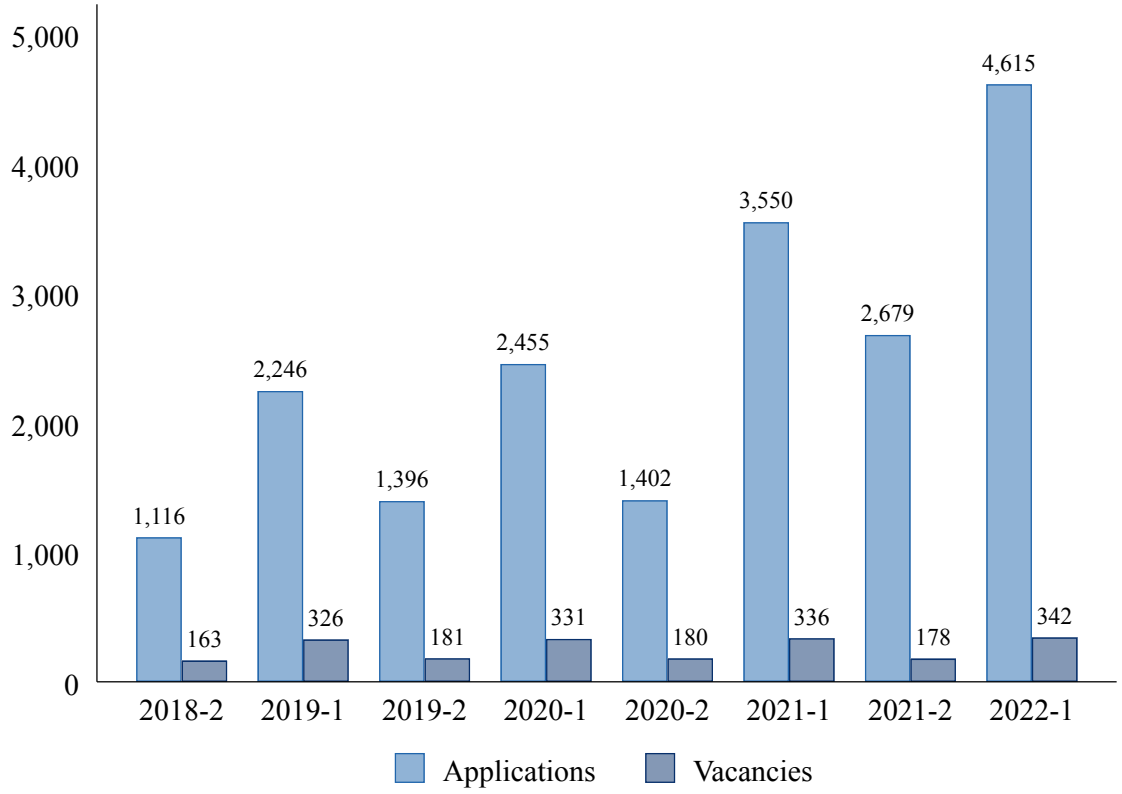
Notes: The data consists of the Peru Demographic Health Surveys from 2003 until 2017. The x-axis shows DHS wealth quintiles. The y-axis shows healthcare utilization rates for under 5-year old children who suffered from fever in the past two weeks. Omitted outcomes include visits to private pharmacies, healthcare from friends and family, and no treatment at all.

Figure A6: Self-Reported vs. True Lottery Rank



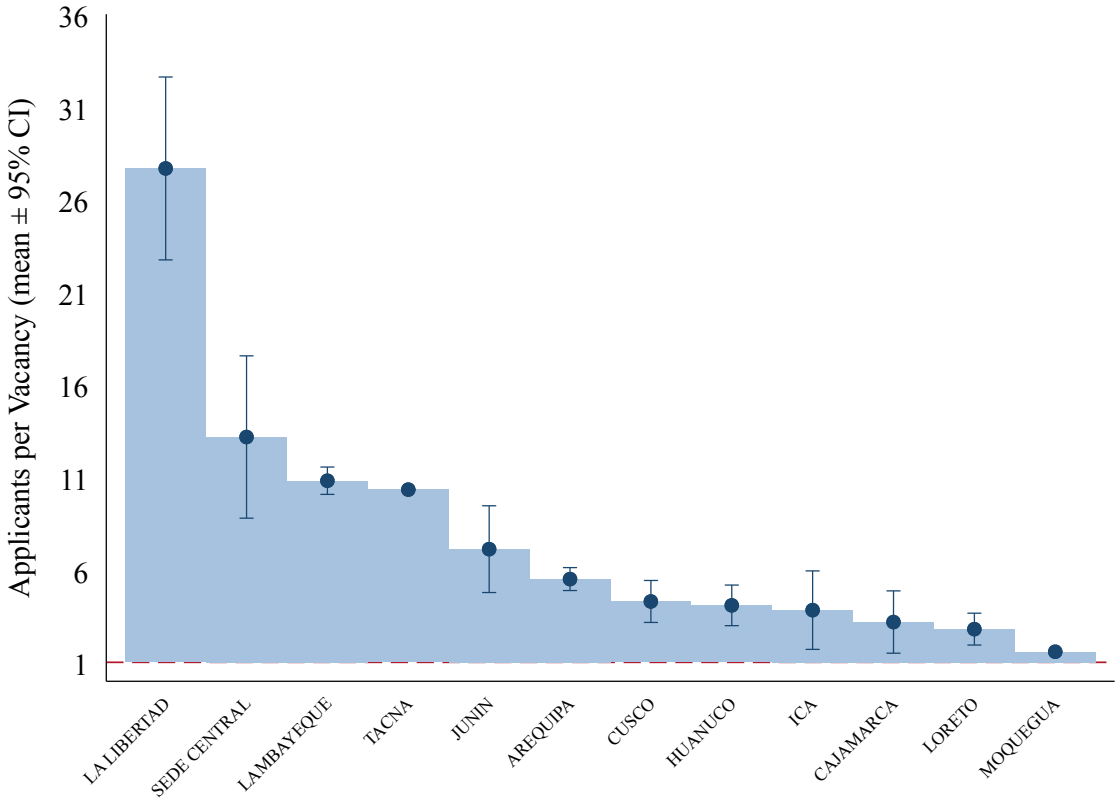
Notes: Figure: This figure plots self-reported against true lottery rank for 177 respondents in 31 lottery draws for which both Zoom recordings and survey information are available. Self-reported and true rank are highly correlated ($r = 0.72$, $p = 0.000$). The grey dashed line indicates the 45-degree line, and the red line plots the OLS fit from regressing self-reported rank on true rank, with standard errors clustered at the lottery-draw level. A slope below one is consistent with classical measurement error and implies attenuation of estimates that rely on the self-reported rank measure.

Figure A7: Number of Applications and Vacancies by Lottery Cycle



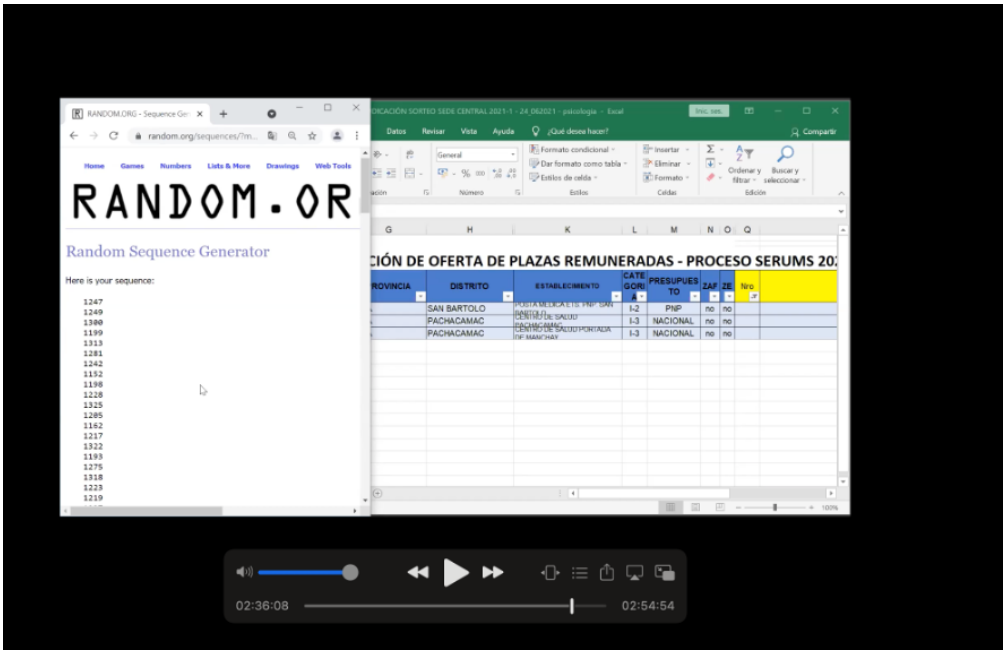
Notes: Number of applicants and vacancies per lottery cycle, 2018-2 through 2022-1, from administrative records of the Ministry of Health. Each year contains two lottery cycles. All vacancies were filled in every cycle except one in 2021-2.

Figure A8: Oversubscription Ratios by Lottery Site



Notes: Mean number of applicants per vacancy at each of the twelve lottery sites, 2018-2 to 2022-1, from administrative records of the Ministry of Health. Points show the mean across lottery cycles and bars show 95% confidence intervals. The dashed line indicates an oversubscription ratio of one.

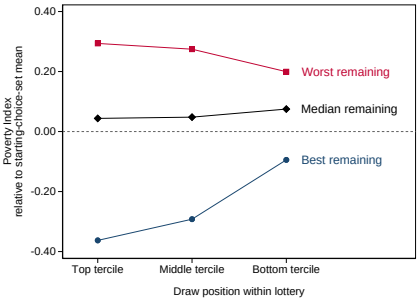
Figure A9: Screenshot of Zoom Lottery



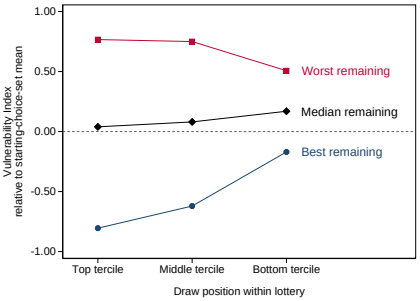
Notes: The figure presents a screenshot from an actual lottery for psychology applicants to the SERUMS program. On the left is the `random.org` platform, which is used to generate a randomized sequence of applicant numbers. On the right is the remaining choice set available at a given lottery stage.

Figure A10: Remaining facilities relative to the starting choice set, by draw position

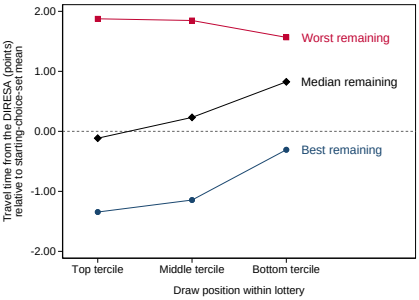
(a) Poverty Index



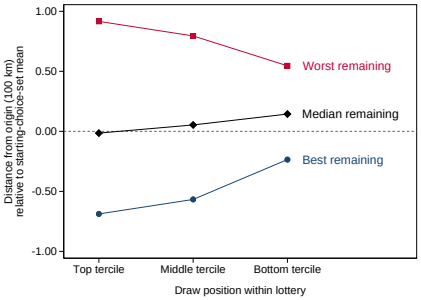
(b) Vulnerability Index



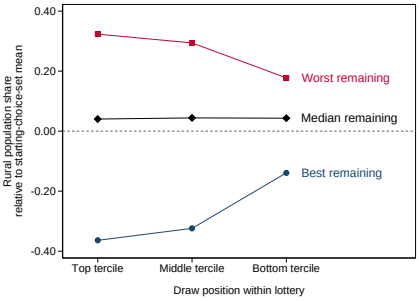
(c) Travel time from the DIRESA (points)



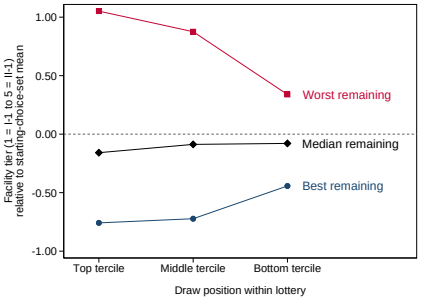
(d) Distance from origin (100 km)



(e) Rural population share

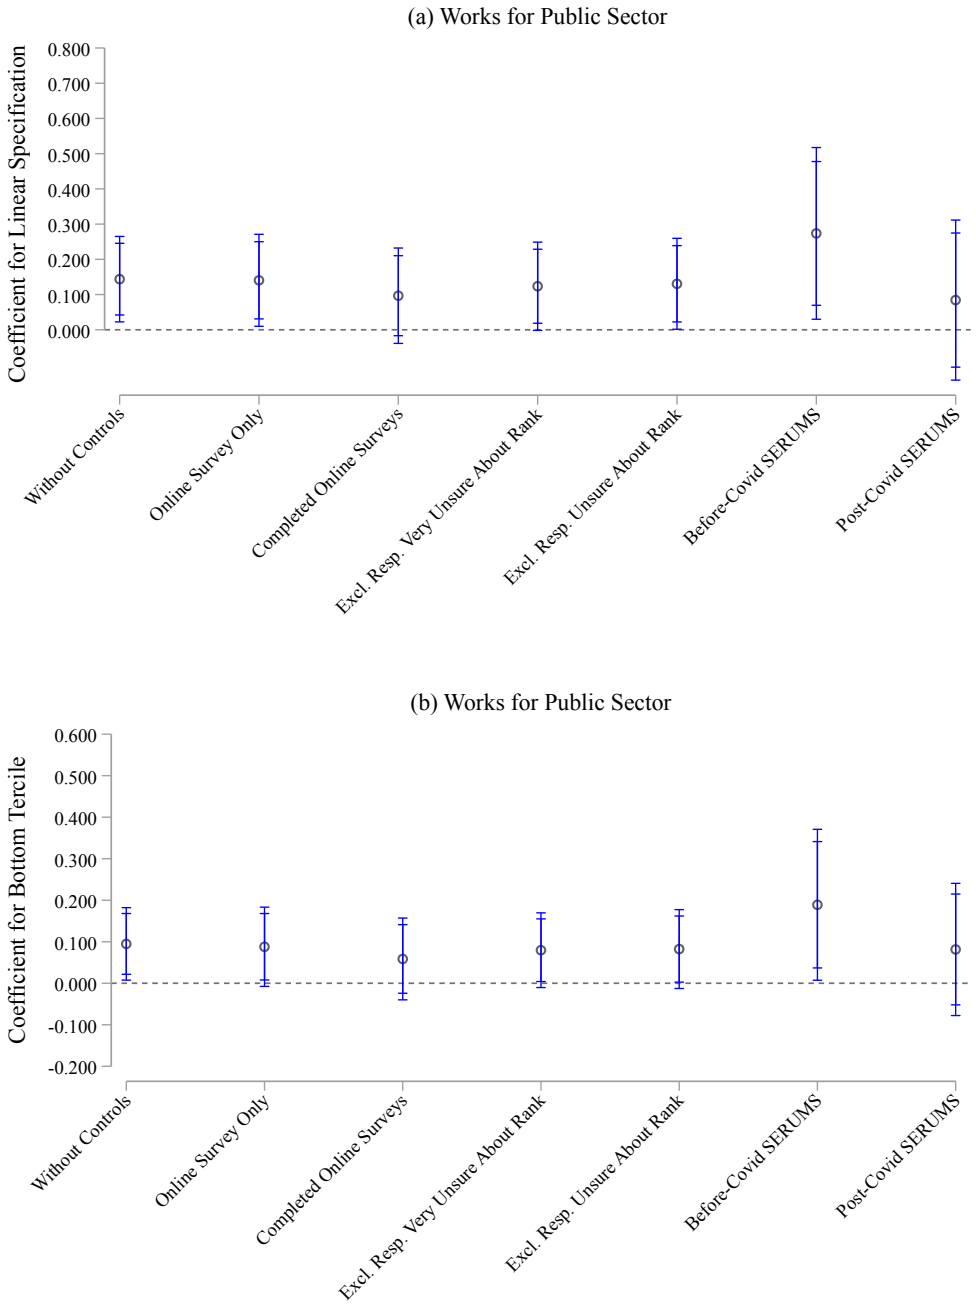


(f) Facility tier (1 = I-1 to 5 = II-1)



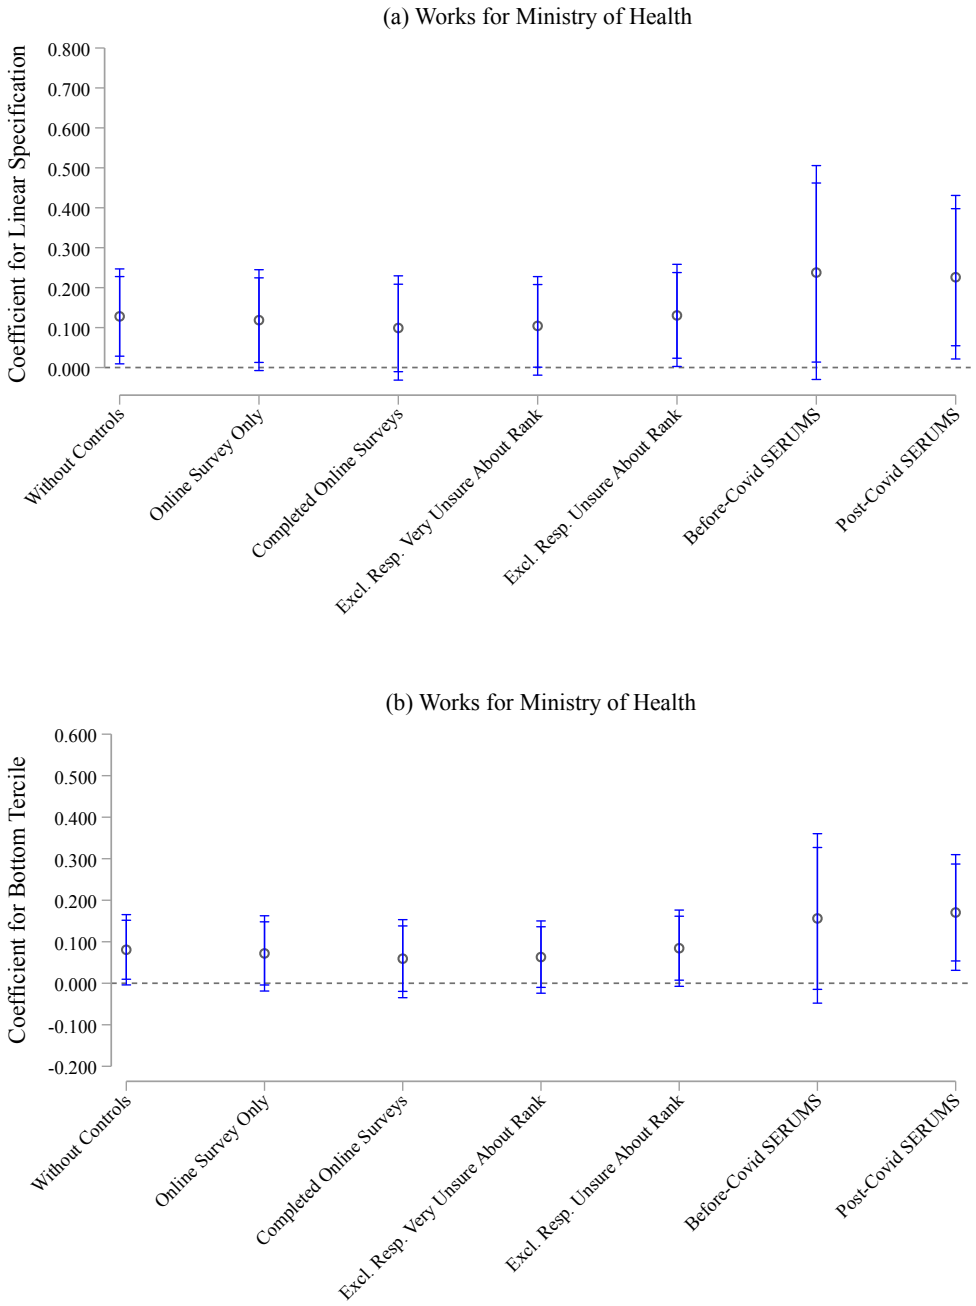
Notes: Lowest, median and highest value among the facilities left at the applicant's turn, minus the mean of the starting choice set, averaged across applicants in each tercile. Table A4 reports the levels. For facility tier, lowest and highest refer to facility size, not desirability.

Figure A11: Robustness Checks: Public Sector Employment



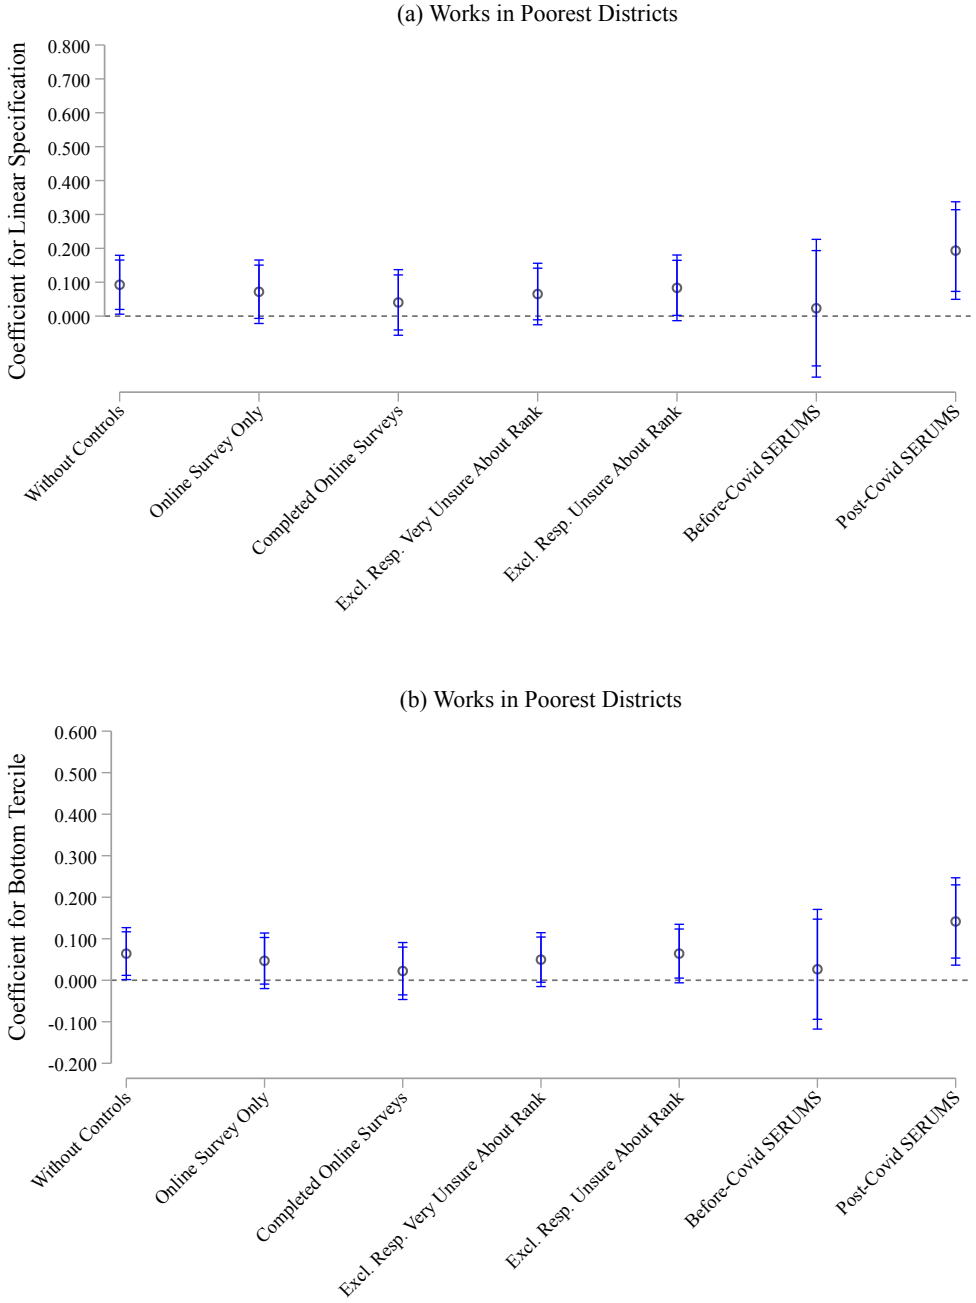
Notes: Panel (a) plots the regression coefficient on the linear lottery percentile rank variable, and Panel (b) plots the regression coefficient on the bottom tercile indicator, across seven specifications: without controls (N=710), online survey participants only (N=623), completed online surveys only (N=582), excluding respondents very unsure about their lottery rank (N=662), excluding respondents unsure about their lottery rank (N=587), SERUMS cohorts before COVID-19 (N=142), and SERUMS cohorts after COVID-19 (N=226). All specifications include strata fixed effects and control for gender, public-university attendance, and a dummy for wether the survey was completed by phone. The figures display 90 and 95 percent confidence intervals.

Figure A12: Robustness Checks: Ministry of Health Employment



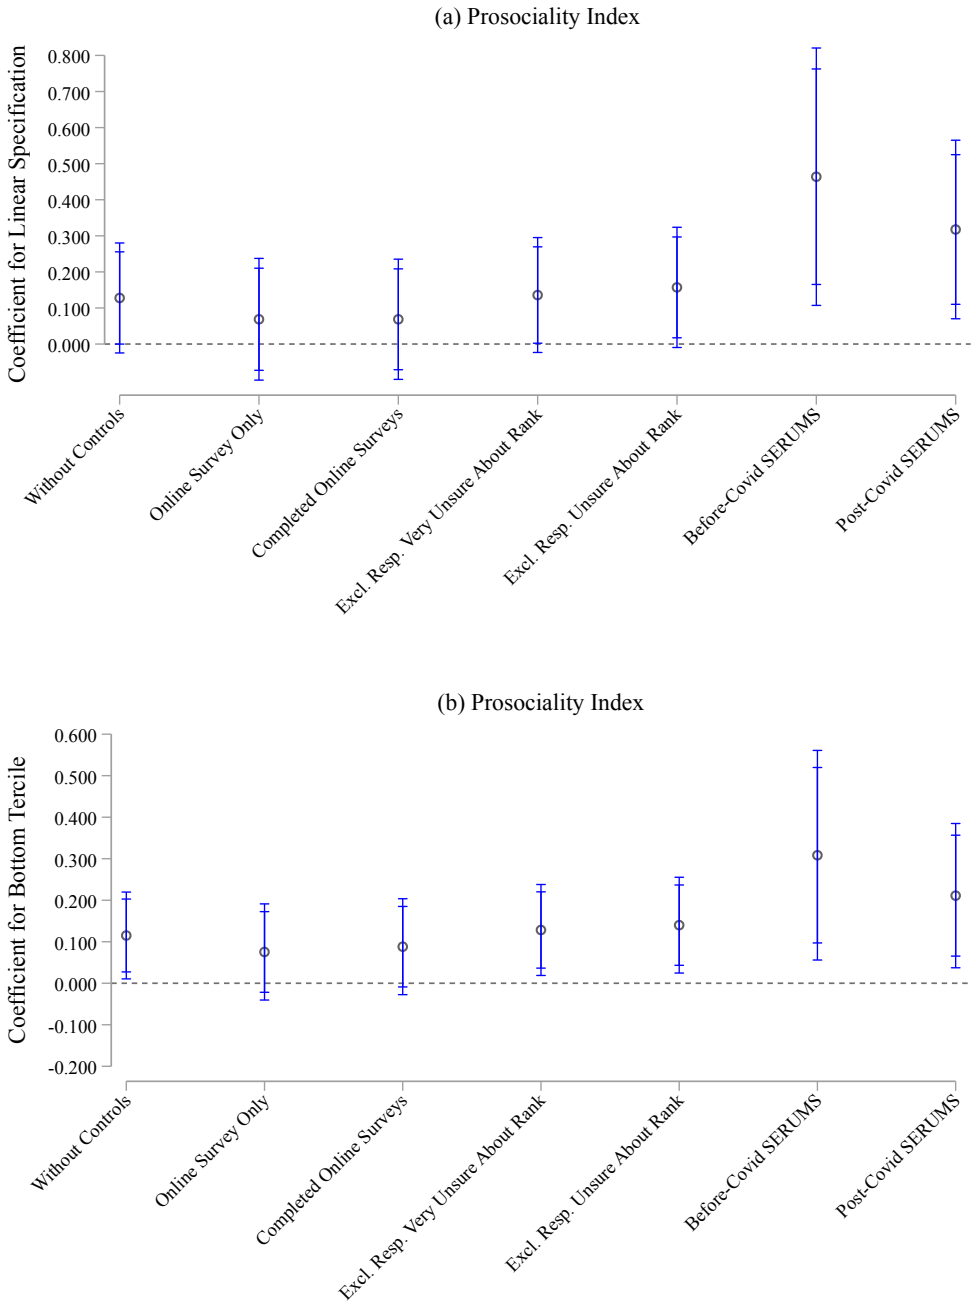
Notes: Panel (a) plots the regression coefficient on the linear lottery percentile rank variable, and Panel (b) plots the regression coefficient on the bottom tercile indicator, across seven specifications: without controls (N=710), online survey participants only (N=623), completed online surveys only (N=582), excluding respondents very unsure about their lottery rank (N=662), excluding respondents unsure about their lottery rank (N=587), SERUMS cohorts before COVID-19 (N=142), and SERUMS cohorts after COVID-19 (N=226). All specifications include strata fixed effects and control for gender, public-university attendance, and a dummy for wether the survey was completed by phone. The figures display 90 and 95 percent confidence intervals.

Figure A13: Robustness Checks: Employment in Poorest Districts



Notes: Panel (a) plots the regression coefficient on the linear lottery percentile rank variable, and Panel (b) plots the regression coefficient on the bottom tercile indicator, across seven specifications: without controls (N=689), online survey participants only (N=606), completed online surveys only (N=567), excluding respondents very unsure about their lottery rank (N=642), excluding respondents unsure about their lottery rank (N=569), SERUMS cohorts before COVID-19 (N=138), and SERUMS cohorts after COVID-19 (N=220). All specifications include strata fixed effects and control for gender, public-university attendance, and a dummy for wether the survey was completed by phone. The figures display 90 and 95 percent confidence intervals.

Figure A14: Robustness Checks: Prosociality Index



Notes: Panel (a) plots the regression coefficient on the linear lottery percentile rank variable, and Panel (b) plots the regression coefficient on the bottom tercile indicator, across seven specifications: without controls (N=706), online survey participants only (N=619), completed online surveys only (N=582), excluding respondents very unsure about their lottery rank (N=658), excluding respondents unsure about their lottery rank (N=583), SERUMS cohorts before COVID-19 (N=140), and SERUMS cohorts after COVID-19 (N=225). All specifications include strata fixed effects and control for gender, public-university attendance, and a dummy for wether the survey was completed by phone. The figures display 90 and 95 percent confidence intervals.

2 Appendix Tables

Table A1: Bonus Point Structure

	Poverty Index	Public Sector
Poverty Q1	$0.6 - 1.0$	15%
Poverty Q2	$0.15 - 0.59$	10%
Poverty Q3	$0.06 - 0.14$	5%
Poverty Q4	$0.03 - 0.05$	2%
Poverty Q5	$0 - 0.02$	0%

Notes: The second column shows the cutoff points for the poverty quintile definitions. The third column shows the bonus points that are given to psychologists for public sector application based on the poverty quintile of the SERUMS location.

Table A2: Attrition Check

	Survey Started (1)	Survey Finished (2)
Panel A: Linear Specification		
Percentile Rank	0.094	0.086
	(0.091)	(0.093)
Panel B: Discrete Specification		
Middle Tercile	-0.009	-0.057
	(0.065)	(0.067)
Bottom Tercile	0.082	0.067
	(0.065)	(0.066)
pval: Middle Tercile = Bottom Tercile	0.169	0.068
Outcome Mean	0.598	0.589
Observations	316	316

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects. Participants without valid contact information were excluded from the analysis.

Table A3: Lotteries observed on Zoom compared with the rest of the oferta

Oferta lotteries	Lotteries	Facilities (mean)	Facilities (median)	Poverty Index	Facility tier	Vulnerability score	Vulnerability Index	Difficulty grade	ZAF share	ZE share
Observed on Zoom	67	5.8	2	0.39	2.94	37.5	-0.56	3.13	0.08	0.03
Not observed	203	5.7	2	0.38	3.01	38.5	-0.51	3.14	0.13	0.05

Notes: All paid psychology lotteries in the 2020–2022 oferta (process × sede × DIRESA × institution); 67 oferta lotteries correspond to the 46 observed on Zoom. For each attribute, the table reports the median across the lottery's facilities, averaged across lotteries. ZAF (remote or border zone) and ZE (emergency zone) are the shares of the lottery's facilities in each zone, averaged across lotteries.

Table A4: Remaining facilities by draw position

Attribute	Draw position	Spots filled	Facilities left	Min	p25	Median	p75	Max
Panel A. Poverty Index, rural population share and facility tier								
Poverty Index	Top tercile	100	14.5	0.22	0.46	0.63	0.75	0.88
	Middle tercile	122	9.5	0.29	0.49	0.63	0.77	0.86
	Bottom tercile	107	3.6	0.51	0.57	0.67	0.77	0.80
Facility tier (1 = I-1 to 5 = II-1)	Top tercile	100	14.5	2.57	2.90	3.17	3.69	4.38
	Middle tercile	122	9.5	2.64	2.97	3.27	3.73	4.24
	Bottom tercile	110	3.6	2.92	3.02	3.28	3.54	3.70
Rural population share	Top tercile	100	14.5	0.16	0.35	0.56	0.72	0.84
	Middle tercile	122	9.5	0.19	0.35	0.56	0.72	0.81
	Bottom tercile	109	3.6	0.38	0.42	0.56	0.66	0.70
Panel B. Vulnerability score and distance								
Vulnerability score (points)	Top tercile	100	14.5	33.45	41.86	49.49	57.15	63.31
	Middle tercile	122	9.5	37.08	43.65	50.40	57.92	63.12
	Bottom tercile	110	3.6	46.21	48.10	52.65	57.07	59.07
Vulnerability Index	Top tercile	100	14.5	-0.78	-0.34	0.06	0.47	0.79
	Middle tercile	122	9.5	-0.59	-0.24	0.11	0.51	0.78
	Bottom tercile	110	3.6	-0.11	-0.01	0.23	0.46	0.57
Distance from origin (100 km)	Top tercile	100	14.5	2.89	3.22	3.56	3.95	4.50
	Middle tercile	122	9.5	3.28	3.54	3.90	4.34	4.65
	Bottom tercile	110	3.6	3.43	3.56	3.81	4.11	4.21
Panel C. Vulnerability score attributes (Points: higher values indicate more difficulty)								
Difficulty grade (1 to 5)	Top tercile	100	14.5	2.71	3.32	4.08	4.38	4.64
	Middle tercile	122	9.5	3.01	3.41	4.02	4.48	4.64
	Bottom tercile	110	3.6	3.66	3.76	4.10	4.48	4.54
Travel time from the DIRESA (points)	Top tercile	100	14.5	1.05	1.41	2.28	3.40	4.27
	Middle tercile	122	9.5	1.29	1.82	2.66	3.77	4.28
	Bottom tercile	110	3.6	2.15	2.53	3.28	3.83	4.02
Type of transport from the DIRESA (points)	Top tercile	100	14.5	2.91	3.23	3.33	3.44	3.93
	Middle tercile	122	9.5	3.08	3.24	3.34	3.50	3.76
	Bottom tercile	110	3.6	3.18	3.25	3.31	3.40	3.44
Number of transport routes (points)	Top tercile	100	14.5	0.00	0.00	0.01	0.06	0.20
	Middle tercile	122	9.5	0.02	0.02	0.03	0.08	0.11
	Bottom tercile	110	3.6	0.05	0.05	0.07	0.09	0.09
Transport risks (points)	Top tercile	100	14.5	0.28	0.80	1.29	1.73	2.44
	Middle tercile	122	9.5	0.52	0.86	1.31	1.84	2.44
	Bottom tercile	110	3.6	1.07	1.18	1.51	1.82	2.08
Transport frequency (points)	Top tercile	100	14.5	0.00	0.00	0.01	0.24	1.31
	Middle tercile	122	9.5	0.00	0.00	0.08	0.58	1.28
	Bottom tercile	110	3.6	0.18	0.18	0.31	0.70	1.09
District altitude (points)	Top tercile	100	14.5	1.75	3.62	5.62	6.57	7.45
	Middle tercile	122	9.5	2.15	3.64	5.54	6.25	7.34
	Bottom tercile	110	3.6	3.99	4.29	5.37	6.22	6.74
District temperature (points)	Top tercile	100	14.5	0.56	3.46	3.94	4.64	5.59
	Middle tercile	122	9.5	1.41	3.32	3.92	4.52	5.33
	Bottom tercile	110	3.6	2.84	3.40	4.00	4.44	4.48
Facility internet access (points)	Top tercile	100	14.5	0.00	0.00	0.07	1.15	3.25
	Middle tercile	122	9.5	0.00	0.00	0.10	0.86	2.83
	Bottom tercile	110	3.6	0.14	0.14	0.25	0.77	1.32

Table A4 (continued)

Attribute	Draw position	Spots filled	Facilities left	Min	p25	Median	p75	Max
Indigenous-language population (points)	Top tercile	100	14.5	2.94	3.85	4.79	5.38	5.89
	Middle tercile	122	9.5	3.61	4.50	4.99	5.47	5.71
	Bottom tercile	110	3.6	4.74	4.80	5.09	5.31	5.31
Basic health team (points)	Top tercile	100	14.5	0.00	0.00	0.00	0.00	0.75
	Middle tercile	122	9.5	0.00	0.00	0.00	0.12	0.29
	Bottom tercile	110	3.6	0.00	0.00	0.02	0.05	0.05

Notes: Mean across applicants in each tercile of the minimum, 25th percentile, median, 75th percentile and maximum of the attribute over the facilities left at their turn. Facilities left is the mean number available. Panel C items are components of the vulnerability score, in difficulty points; higher values indicate a more difficult posting. Rural share is missing for four facilities in districts created after 2006.

Table A5: Effect of Lottery Rank on Perceived Skill Gains and Skill-to-job Transfer Perceptions of Current Workers

	Program Helped to Improve					Improvements in ... Helped with Getting Job				
	Skill Acq. Index	Knowledge	Clinical Skills	Language Skills	Networks	Skills-to-Job Index	Knowledge	Clinical Skills	Language Skills	Networks
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Linear Specification										
Percentile Rank	-0.207	-0.307	-0.158	-0.140	-0.222	0.117	-0.057	0.148	0.320	0.057
	(0.151)	(0.194)	(0.206)	(0.206)	(0.211)	(0.123)	(0.194)	(0.206)	(0.205)	(0.195)
Panel B: Discrete Specification										
Middle Tercile	0.076	0.149	0.065	0.194	-0.105	0.176**	0.245	0.191	0.306**	-0.037
	(0.099)	(0.143)	(0.150)	(0.153)	(0.149)	(0.087)	(0.153)	(0.152)	(0.155)	(0.150)
Bottom Tercile	-0.096	-0.129	-0.055	-0.087	-0.114	0.076	-0.007	0.109	0.211	-0.009
	(0.107)	(0.142)	(0.148)	(0.147)	(0.153)	(0.090)	(0.141)	(0.150)	(0.148)	(0.137)
pval: Middle Tercile = Bottom Tercile	0.111	0.060	0.440	0.096	0.951	0.288	0.108	0.610	0.576	0.867
Outcome Mean	0.005	-0.020	0.035	-0.045	0.051	0.001	0.002	0.005	0.008	-0.010
Observations	343	343	343	343	343	343	343	343	343	343

Notes: Robust standard errors appear in parentheses. All regressions include lottery-stratification fixed effects and a vector of controls consisting of gender, a dummy for whether the respondent's university was public, and a dummy for whether the survey was completed by phone. The sample is restricted to respondents who were currently working for pay at the point of the survey.

Table A6: Effect of Lottery Rank on Perceived Employability

	at Ministry of Health		at ESSALUD		in Private Sector	
	Agree	Strongly Agree	Agree	Strongly Agree	Agree	Strongly Agree
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Linear Specification						
Middle Tercile	-0.001	0.067	0.014	0.035	-0.045	-0.002
	(0.050)	(0.042)	(0.053)	(0.036)	(0.046)	(0.037)
Panel B: Discrete Specification						
Middle Tercile	-0.001	0.067	0.014	0.035	-0.045	-0.002
	(0.050)	(0.042)	(0.053)	(0.036)	(0.046)	(0.037)
Bottom Tercile	-0.020	0.062	-0.007	0.010	-0.004	-0.018
	(0.050)	(0.043)	(0.051)	(0.035)	(0.043)	(0.038)
pval: Medium Tercile = Bottom Tercile	0.730	0.918	0.722	0.518	0.386	0.702
Outcome Mean	0.663	0.163	0.516	0.128	0.779	0.155
Observations	617	617	617	617	617	617

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of gender, a dummy for whether the respondent's university was public, and a dummy for wether the survey was completed by phone.

Table A7: Effect of Lottery Rank on Donation Subindex Components

	Donated Time (1)	Donated Money (2)	Hypothetical Scenario (3)
Panel A: Linear Specification			
Percentile Rank	0.139	-0.142	0.024
	(0.137)	(0.131)	(0.136)
Panel B: Discrete Specification			
Middle Tercile	-0.028	-0.207**	-0.039
	(0.104)	(0.097)	(0.099)
Bottom Tercile	0.119	-0.040	0.076
	(0.097)	(0.093)	(0.098)
pval: Medium = Bottom	0.166	0.109	0.277
Outcome Mean	0.054	0.128	-0.055
Observations	673	672	674

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and a vector of controls that consists of age, gender, a dummy for whether the respondent's university was public, a dummy for whether Spanish is the respondent's native language, and birth region fixed effects. Columns 1-3 and 5-7 were also asked as part of the phone survey and include further a dummy variable for whether the survey was done by phone. The sample in columns 1-4 consists of psychologists who completed SERUMS at the point of the survey. The sample in columns 5-7 consists of psychologists who were currently doing SERUMS or who completed the program three months prior.

Table A8: Effect of Lottery Rank on Decisions in Dictator and Spectator Games

				Donations in Dictator Game		Donations in Spectator Game		
	Donations Health	Donations Culture	Donations Total	Vitamin A NGO	Urban Culture NGO	Urban Domestic Violence NGO	Rural Domestic Violence NGO	Urban Culture NGO
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Linear Specification								
Percentile Rank	-0.841	0.515	-0.325	-1.044	0.732	-0.261	0.348	-0.086
	(0.682)	(0.552)	(1.058)	(0.875)	(0.701)	(0.414)	(0.508)	(0.397)
Panel B: Discrete Specification								
Middle Tercile	-0.425	0.263	-0.162	-0.590	0.434	-0.413	0.711*	-0.298
	(0.441)	(0.400)	(0.728)	(0.552)	(0.493)	(0.332)	(0.418)	(0.287)
Bottom Tercile	-0.560	0.159	-0.401	-0.692	0.262	-0.191	0.248	-0.057
	(0.473)	(0.392)	(0.746)	(0.603)	(0.493)	(0.301)	(0.369)	(0.296)
pval: Middle Tercile = Bottom Tercile	0.778	0.824	0.768	0.861	0.764	0.495	0.295	0.399
Outcome Mean	2.852	1.447	4.300	2.503	1.562	1.726	2.169	1.105
Observations	592	592	592	457	457	135	135	135

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and controls for gender and whether the respondent attended a public university. These questions were not included in the phone survey. Column 1 aggregates columns 4, 6, and 7; column 2 aggregates columns 5 and 8; column 3 aggregates columns 1 and 2. The **Dictator Game** (first survey round) asked respondents whether they would donate any portion of their 19 PEN (≈$5.50 USD) incentive to one of two NGOs: one distributing Vitamin A to children in rural Peru, and one promoting art and culture in urban areas. The **Spectator Game** (second survey round) informed participants that they had been awarded an additional 5 PEN (≈$1.50 USD) to allocate among three NGOs: two focused on domestic violence victims in urban and rural areas respectively, and one promoting art and culture in urban areas.

Table A9: Heterogeneity Analysis — Linear Specification

		Index Components			
	Outcome Index	Works for Ministry of Health	Prosociality Index	More Willing to Work for MoH	More Willing to Work in Poor Areas
	(1)	(2)	(3)	(4)	(5)
Panel A: Born in Lima					
Percentile Rank × Born in Lima	0.181*	-0.063	0.354*	0.021	0.251
	(0.103)	(0.152)	(0.186)	(0.148)	(0.180)
Percentile Rank	0.047	0.139**	0.045	0.164**	-0.116
	(0.047)	(0.069)	(0.091)	(0.074)	(0.079)
Born in Lima	-0.112**	-0.037	-0.192*	-0.010	-0.121
	(0.054)	(0.085)	(0.105)	(0.085)	(0.096)
Outcome Mean	0.185	0.322	-0.003	0.222	0.234
Observations	707	707	703	550	481
Panel B: Did Program at Birthplace					
Percentile Rank × Did Program at Birthplace	-0.039	0.063	-0.224	0.242*	0.009
	(0.084)	(0.124)	(0.157)	(0.129)	(0.146)
Percentile Rank	0.107*	0.100	0.228**	0.058	-0.068
	(0.058)	(0.084)	(0.104)	(0.088)	(0.100)
Did Program at Birthplace	0.038	-0.025	0.094	-0.082	-0.000
	(0.049)	(0.072)	(0.092)	(0.077)	(0.088)
Outcome Mean	0.185	0.322	-0.003	0.222	0.234
Observations	707	707	703	550	481
Panel C: Under 30 yrs					
Percentile Rank × Under 30 yrs	-0.196**	-0.022	-0.138	-0.337**	-0.265*
	(0.086)	(0.132)	(0.166)	(0.145)	(0.156)
Percentile Rank	0.221***	0.141	0.221*	0.411***	0.120
	(0.071)	(0.107)	(0.132)	(0.122)	(0.129)
Under 30 yrs	0.034	0.007	-0.008	0.062	0.072
	(0.048)	(0.072)	(0.093)	(0.074)	(0.083)
Outcome Mean	0.185	0.325	-0.005	0.222	0.234
Observations	704	704	700	552	483
Panel D: No Children					
Percentile Rank × No Children	-0.101	0.040	-0.168	-0.235	-0.045
	(0.105)	(0.161)	(0.196)	(0.154)	(0.154)
Percentile Rank	0.172*	0.167	0.261*	0.308**	-0.055
	(0.088)	(0.130)	(0.158)	(0.129)	(0.124)
No Children	0.072	-0.006	0.082	0.065	0.127
	(0.054)	(0.081)	(0.110)	(0.081)	(0.083)
Outcome Mean	0.193	0.333	-0.016	0.221	0.236
Observations	488	488	484	483	480
Panel E: Private University					
Percentile Rank × Private University	0.078	0.185	-0.138	0.114	0.155
	(0.107)	(0.166)	(0.210)	(0.170)	(0.157)
Percentile Rank	0.021	-0.027	0.235	0.077	-0.188
	(0.097)	(0.151)	(0.192)	(0.155)	(0.135)
Private University	-0.061	-0.135	0.098	-0.085	-0.012
	(0.068)	(0.097)	(0.133)	(0.093)	(0.098)
Outcome Mean	0.186	0.324	-0.003	0.222	0.234
Observations	710	710	706	552	483
Panel F: Top Ranked University					
Percentile Rank × Top Ranked University	0.098	-0.017	0.150	-0.034	0.190
	(0.095)	(0.140)	(0.163)	(0.162)	(0.177)
Percentile Rank	0.055	0.109	0.128	0.143*	-0.131
	(0.052)	(0.076)	(0.104)	(0.082)	(0.088)
Top Ranked University	-0.096*	-0.123	-0.035	-0.077	-0.094
	(0.054)	(0.079)	(0.098)	(0.082)	(0.094)
Outcome Mean	0.178	0.315	-0.021	0.217	0.237
Observations	638	638	634	489	424

Notes: Robust standard errors appear in parentheses. All regressions include lottery stratification fixed effects and controls for gender, whether the respondent attended a public university, and whether the survey was done by phone. The outcome index is the unweighted average of the four index components. Each panel reports results from a separate regression that interacts percentile rank with the indicated subgroup indicator. All subgroup indicators are binary. Columns 4 and 5 have fewer observations due to these outcomes being administered only in the first round of surveys.

Table A10: Randomization Inference

	p-value (Bottom Rank Tercile)	p-value (Percentile Rank)
Works For Public Sector	0.104	0.072
Works for Ministry of Health	0.146	0.090
Prosociality Index	0.098	0.186

Notes: The exercise randomly re-assigns the lottery rank to each respondent 2,000 times. For each iteration, we then repeat our analysis. The p-value is calculated as the share of the placebo coefficients that are larger in magnitude than the actual coefficient (in absolute terms).