More Women in Tech?
Evidence from a field experiment addressing social identity

Lucía Del Carpio
INSEAD

Maria Guadalupe
INSEAD

PRELIMINARY AND INCOMPLETE

Abstract

This paper investigates whether social identity considerations and norms may be driving occupational choices by women. We implement a randomized field experiment to analyze how the self-selection of women into the technology sector changes when we randomly vary the recruitment message to potential applicants to a 5-month software coding program offered only to low income women in Peru and Mexico. In addition to a control message with generic information, in a treatment message we correct misperceptions about expected returns for women and their ability to pursue a career in technology. This de-biasing message doubles the probability of applying (from 7% to 15%). We then analyze the stark differential self-selection patterns for the treatment and the control groups to infer the potential barriers that may explain occupational segregation. We find evidence that both expectations about monetary returns in the sector and a perceived "identity" cost (as reflected by an IAT test and survey measures) of a career in technology operate as barriers. We interpret our results in the light of a Roy model where women are endowed with returns to skill in the technology and services sector, and bear an identity cost of working in technology (*à la* Akerlof Kranton, 2000). Through a follow up experiment in Mexico DF we are able to rule out alternative explanations for our results and point to what dimensions of the initial treatment matter most. Our results suggest social identity can explain persistent occupational segregation in this setting and point towards policy interventions that may alleviate it.

1. Introduction

In spite of significant progress in the role of women in society in the last 50 years, an important gender wage gap persists today. Scholars have shown that a large share of that gap can be explained by the different industry and occupational choices men and women make. However, the reasons behind those stark differential choices are still unclear (Blau and Kahn, 2017). In this paper we propose and study "social identity" as a key driver of women's occupational choices, and in particular, its predominant feature:

persistent occupational gender segregation (see e.g. Bertrand, 2011; Goldin, 2014; Bertrand and Duflo, 2016).

Starting with at least Roy (1951) economists have tried to explain how people self-select into certain occupations/industries. We argue that women are likely to select a career not just as a function of the marginal returns to their skills in that sector (as in the classic Roy, 1951, model), but also of their beliefs on expected success given existing gender norms (Akerlof Kranton, 2000), and possibly to other preferences such as the non-monetary benefit of working in an environment with few other women, or other attributes of the industries (Goldin 2014).[1] The fact that social identity matters has long been recognized and shown to be relevant empirically by social psychologists who have designed and tested strategies to reduce bias and stereotypical thinking (Spencer and Steele, 1999; see survey by Paluk and Green 2009). But much of this evidence is in the lab and looks at very short-term outcomes. In economics, there is very little empirical evidence, but Bertrand Kamenica and Pan (2015) show that gender identity norms can explain a number of important patterns in marriage.

The goal of this paper is to bring together, and into the field, the economics of self-selection and psychology de-biasing literatures to investigate how important are identity considerations in the occupational choices women make. We focus on the technology sector for two key reasons: first, because it is a sector with high growth potential and good employment prospects; second, because it is predominantly male.

Our framework builds on the Roy (1951)/Borjas (1987) model of self-selection and introduces identity considerations following Akerlof and Kranton (2000). Women decide whether to enter the technology industry (rather than go to the services sector) as a function of their endowment of "technology" skills, "services" skills and what we will refer to an identity cost of entering the technology sector that it is a male dominated sector.[2] As in the standard Roy model (without identity) the self-selection will depend on the correlation between the two types of skills and identity relative to the dispersion of skills and identity. Depending on these correlations and dispersions,

[1] Goldin (2014) shows that women, fearing discrimination, select occupations that focus on objective performance measures and industries that allow for more flexibility.

[2] "Identity" can be thought of is a non-monetary/psychological cost of working in an industry where the social norm is very different from one's social category. In our case this is the identity cost of a low income Latin American woman, relative to the norm of the technology/software development industry, which is typically male and upper middle class.

we may observe positive or negative self-selection into the technology sector both along the skills and the identity dimensions: i.e. we may end up with a sample that is on average more or less skilled, and more or less "biased".

Against this framework, we run two field experiments that aimed to de-bias women against the perception that women cannot be successful in the technology sector. This can occur by raising expected monetary return to skills and/or by reducing the perceived identity costs. In both experiments, we randomly varied the recruitment message to potential applicants to a 5 month "coding" bootcamp and leadership training program, offered only to women from low-income backgrounds by a non-for-profit organization in Latin America.[3] We ran the first field experiment in Lima (Peru) where female coders represent only 7% of the occupation. In addition to sending a control group message with generic information about the program (its goals, content and requirements), in a treatment message, we added a paragraph aiming to correct misperceptions about women's abilities to pursue a career in technology, provide role models and highlight the fact that the program is offered solely to women. We argue the message increased the expected returns of a woman pursing a career in technology both by reducing the identity cost and increasing expected monetary returns (we argue it is impossible to separate the two in a message). Subsequently, applicants to the program were invited to attend a set of tests and interviews that will determine who will be selected to the training. In those interviews we were able to collect a host of characteristics on the applicants, in particular those implied by the framework as being important and reflected in the self-selection from the message: their expected monetary returns of pursuing a career in technology and of their outside option (a services job), their cognitive skills, and two measures of implicit gender bias (an implicit association test (IAT) we created specifically to measure how much they identify gender (male/female) and occupational choice (technology/services) as well as a survey based measure of identification with traditional female role). We also collected an array of other demographic characteristics, and games aimed to elicit preferences and aspirations, which allow us to rule out alternative mechanisms for our findings.

[3] The goal of the organization is to identify high potential women, that because of their background may not have the option, knowledge or tools to enter the growing technology sector, where it is hard to find the kind of basic coding skills offered in the training.

We find that application rates doubled from 7% to 15% when women receive the de-biasing message. This increased significantly the applicant pool to the training program. We then analyze the self-selection patterns in the two groups to assess what are the barriers that are being loosened by the message. We essentially estimate the equilibrium self-selection following an exogenous shock to the perceived returns to a career in technology. Our results suggest that there is negative self-selection in both average technology skills, average services skills, as well as in cognitive skills. This implies that we are in a world of comparative (not absolute) advantage in technology vs. services skills, i.e. that the correlation between the two types of skills is positive. However, the correlation is smaller than the relative dispersion in those skills in the economy. So, in the margin "worse" women apply.

We also find negative self selection on identity costs: on average, women with higher identity cost as measured by the IAT and the traditional gender role survey measure apply following our de-biasing message, the marginal woman is "more biased". In the light of our model, this result suggests, first, that identity matters and that identity costs vary across women. Otherwise we would not find a significant change in average identity. Second, in the light of the model it implies that the correlation between identity costs and skills is not too large (relative to the dispersion of the two variables).

Overall, however, what firms and organizations care about is the right tail of the skills distribution: do we have more qualified women to chose from now? We find that the overall increase in applicants also raises the numbers of high-cognitive ability applicants: the de-biasing message significantly increases cognitive and tech specific abilities of the top group of applicants (those that would have been selected for training). Why did higher cognitive skill women apply even if on average selection is negative? Besides the obvious answer of noise in the distribution of skills or the effect of the experiment, another reason within our framework would be that given the distributions of skill and identity, there are some high skill women that are also high identity costs women that did not apply before treatment that are induced to apply when expected returns to skill increase or the expected identity cost falls. We also measured a number of other characteristics and preferences of applicants, which allow us to rule out certain possible mechanisms of the effects we find.

In a second experiment in Mexico City we aimed to disentangle what was the information in the first message that the women in Lima responded to. This allows us to directly test whether it is beliefs about the returns, the non-monetary component to being in an environment with fewer women and/or being presented with a role model which mattered most in our first message. It also allows us to rule out that it is any kind of information provided about women that makes a difference. Now the control treatment was the complete message and in each of three treatments we took out one feature of the initial message (returns, network of women and role model) at a time. We found that women respond mostly to the presence of a role model, and also to hearing about the high expected returns for women in the technology sector. In contrast, the information that they would have a network of other women upon graduating made no significant difference to application rates.

A specificity of our setting is that the training is offered only to women, and all applicants know that. This has the advantage that we can design a message that is specifically targeted to women without being concerned about negative externalities on men by providing, for example, a female role model. It therefore allows us to investigate mechanisms that would be harder to investigate as clearly in the presence of men. This comes at the cost that we do not know how men would respond in a setting where they also see the de-biasing message, and that we cannot say anything about the role of monetary returns or identity for men or other social categories or what kind of message would work as an encouragement to men.

This paper contributes to the literature on how women self-select to different industries (Goldin, 2014). Field experimental evidence on this topic is limited. For example, Flory, Leibbrant and List (2015) show that that women shy away from competition. But none of these focuses on the explicit de-biasing or correction of misperceptions.

We also relate to the literature on socio-cognitive de-biasing under stereotype threat in social psychology (Steele and Aronson, 1995). It is by now well established in this literature that disadvantaged groups under-perform under stereotype threat and the literature has devised successful de-biasing strategies (Good, Aronson, and Inzlicht, 2003; Kawakami et al., 2017; Forbes and Schmader, 2010). While this literature focuses on the effect of de-biasing on performance we focus on its effect on self-selection (we

cannot assess the effect of de-biasing on performance itself, but it is unlikely to be very big in our setting given the context of the test and surveys).

We also contribute evidence to a very limited literature on the performance effects of restricting the pool of applicants through expected discrimination or bias. As Bertrand and Duflo (2015) state "the empirical evidence (even non-randomized) on any such consequence of discrimination is thin at best". Ahern and Dittmar (2012) and Matsa and Miller (2013) find negative consequences on profitability and stock prices of the Norway 2006 law mandating a gender quota in corporate board seats and find negative consequences on profitability and stock prices. We identify improvements not only in the number of applicants, but also in the type of applicants and the number of top applicants available to select from, even though the average quality of candidates falls.

Finally, our paper is related to the literature showing how the way a position is advertised can change the applicant pool. Ashraf, Bandiera and Lee (2014) study how career incentives affect who selects into public health jobs and, through selection, their performance while in service. They find that making career incentives salient attracts more qualified applicants with stronger career ambitions without displacing pro-social preferences. Marinescu and Wolthoff (2013) show that providing information of higher wages attracts more educated and experienced applicants. And Dal Bó et al. (2013) explore two randomized wage offers for civil servant positions, finding that higher wages attract abler applicants as measured by their IQ, personality, and proclivity toward public sector work. In contrast to these papers we find negative self-selection on average, which highlights the fact that an informational treatment that aims at selection needs to take into account the returns of the outside option, and the correlations between returns.

Whether women (or men) self-select out of certain industries for "identity" reasons is an important question, not just because if "identity" matters it would be a barrier blocking the efficient allocation of (human) resources and hence aggregate welfare, but also because it speaks to the secular debate about nature versus nurture. Do women select out from certain industries because they are genetically different or because society is configured in a way that "biases" and conditions their choices? For example, the infamous Google engineer fired in 2017 after writing a memo to the company

seemed to think that women are intrinsically different in ways that disqualified them for a career in technology. This paper sheds light on that question.

2. Framework: Self-Selection into an industry

This section develops a simple theoretical framework to illustrate how changing the information provided on a career/an industry, affects which applicants self-select into that career. We start from a standard Roy model (Roy, 1951; Borjas 1987) adapted to our setting and add identity as a potential driver of the decision to enter an industry in addition to the relative return to skills in the two industries, as in the classic model.

Women decide between applying or not applying to the training program, i.e., whether to attempt a career in the technology sector. Each woman is endowed with a given level of skills that are useful in the technology sector *T* and skills that are useful in the services sector *S*. Assume for now that identity does not matter: Total returns in Services and in Tech are given by $W_0 = P_0 S$ and $W_1 = P_1 T$, respectively, where P_0 and P_1 are the returns to skill (e.g. wage per unit of skill) in each sector. If we log linearize and assume log normality: $\ln W_0 = w_0 + s$ and $\ln W_1 = p_1 + t$ where *lnS*= *s*~N(0, σ_s^2) and *lnT*= *t*~N(0, σ_t^2).

The probability that a woman applies to the technology sector is :

$$\Pr(Apply) = \Pr\left(p_1 + t > p_0 + s\right) = \Pr[\frac{v}{\sigma_v} > \frac{p_0 - p_1}{\sigma_v}] = 1 - \Phi[\frac{p_0 - p_1}{\sigma_v}]$$

Where $v = t - s$ and Φ is the CDF of a standard normal.

Pr (Apply) is increasing in p_1 and decreasing in p_0, such that as expected returns in technology increase, more women will apply to Tech.

This allows us to study how the selection of women (the average expected level of *t*) that apply will change with a change in returns to technology skill. One can show that

$E(T \mid Apply) = \rho_{tv}\sigma_t\lambda(\frac{p_0 - p_1}{\sigma_v})$ where $\rho_{tv} = \sigma_{tv}/(\sigma_v\sigma_t)$ is the coefficient of correlation between *t* and *v*, and $\lambda(z)$ is the inverse mills ratio, with $\lambda' > 0$. Therefore:

$$\frac{dE(T \mid Apply)}{dp_1} = \frac{\sigma_t^2 - \sigma_{st}}{\sigma_t\sigma_v}\frac{d\lambda(z)}{dp_1}.$$

Given $\frac{d\lambda(z)}{dp_1} < 0$ and $\sigma_t\sigma_v > 0$ the sign of the selection will depend on the sign of $\sigma_t^2 - \sigma_{st}$.

In particular, if $\rho > \frac{\sigma_t}{\sigma_s} \Rightarrow \frac{dE(T \mid Apply)}{dp_1} > 0$ and selection is positive, and $\rho < \frac{\sigma_t}{\sigma_s} \Rightarrow \frac{dE(T \mid Apply)}{dp_1} < 0$ selection is negative and the average Tech skills of applicants decreases in the expected returns to Tech skills.

Now we introduce the concept of identity to the basic framework. We follow Akerlof and Kranton (2000) who introduce "identity" as a potential driver of behavior in economic models. In their representation, which builds on a large body of social psychology literature, "identity" defines who people are, and in particular for our purposes, what social category they belong to. Then, society attaches norms of behavior to different social categories. For example, in a traditional society women stay at home to tend to the household and children, while men are the breadwinners and work outside the house. There are norms of behavior attached to these categories. These norms can inflict a cost to individuals that do not conform to the norm. For example, in our case, being a woman in the technology sector can represent and non-monetary (psychological) cost for women whose identity does not fall within the norm. Or a stay-at-home father may suffer an identity cost if the norm is that men are breadwinners working in the market so that being at home and not having a "career" can feel particularly costly to the individual. Note that a key difference between identity and standard preferences is that the cost is imposed by the social norm prevalent in society: in a different society the same person with the same preferences would have a different cost, for example in a society where women are the breadwinners, women would the

ones bearing the psychological identity cost of staying at home, with their underlying preferences unchanged.

Assume first that identity in our setting is a fixed cost, identical for all women, which is a "psychological"/non-monetary cost of being a woman in the technology industry given by the fact that the social norm in the Tech industry is a "male" norm. If the identity cost *i* is fixed for all women such that $\ln W_1 = p_1 + t - i$, then a reduction in *i* will have the same effect of an increase in p_1 in the previous case and increase applications leading to positive or negative selection in *t* depending on how strongly *t* and *v* are correlated relative dispersion of *t* and *v*. Conceptually, what we call "i" is a non-monetary discount that all women experience in the technology sector, that reduces their total expected returns by a constant magnitude. We will call this "identity costs" but broadly it is a more non-monetary wedge.

Result 1: Reducing identity costs (the non-monetary wedge) and/or increasing expected returns in technology increases application rates, when there are no identity costs, or identity costs are constant.

Result 2: if *i* constant for all women, a decrease in *i* will lead to more applications and a selection in *t, s*, but no selection on *i* (it is constant for all).

Now, let's consider the case where just as services and technology skills are distributed in the population, so are identity costs, with some women experiencing higher identity costs than others so that $\ln W_1 = p_1 + t - i$ with $i \sim N(0, \sigma_i^2)$.

$$\Pr(Apply) = \Pr[t - s - i > p_0 - p_1]$$

$$\Pr(Apply) = \Pr[D - i > p_0 - p_1] = 1 - \Phi[\frac{p_0 - p_1}{\sigma_h}]$$

$$D \sim N(0, \sigma_D^2), D = t - s, h = t - s - i$$

We can again examine selection into an industry as returns to skill increase (or isomorphically, the perceived mismatch between one's identity and the industry identity decreases). Note that when we interpret i as identity, we do not take a stand on whether it is one's identity or the perceived social norm that is changing (which is

something that psychologists would be interested in), we only require that the difference between the two goes down.

In this setting we will expect that the average skill differential of applicants $\frac{dE(D \mid Apply)}{dp_1} > 0$ is higher if $\rho_{Di} > \frac{\sigma_D}{\sigma_i}$. Conversely selection in D will be negative if $\rho_{Di} < \frac{\sigma_D}{\sigma_i}$. This implies that an increase in p_1 now will have a positive or negative effect on average skills depending on the correlation between skills and identity.

Result 3: Increasing expected returns can lead to positive or negative self-selection of in *t*, depending on the correlation between *t, s* and *i* in the underlying population relative to their dispersion. Similarly, it can lead to positive or negative self-selection in *s,* the outside option.

Further, we can see how average identity costs in applicants will change with an increase in expected returns:

$$E(i \mid Apply) = \rho_{ih}\sigma_i\lambda(z)$$
$$\lambda(z) = \phi(z) / \Theta(-z),$$
$$\rho_{Di} > \frac{\sigma_i}{\sigma_D} \Rightarrow \frac{dE(i \mid Apply)}{dp_1} < 0$$
$$\rho_{Di} < \frac{\sigma_i}{\sigma_D} \Rightarrow \frac{dE(i \mid Apply)}{dp_1} > 0$$

Result 4: Increasing expected returns when identity costs are distributed in the population, can lead to positive or negative self-selection in identity cost, depending on the correlation between *t, s* and *i* in the underlying population relative to their dispersion.

Note that once we introduce identity/the psychological wedge, and even in the case of negative average selection on *t*, the expected increase in p_1 through lower perceived identity costs may lead to some very high quality women applying that also have high identity costs. In this setting it is possible that even though on average selection on T is negative, it may be that some women who are high T but also have high *i* apply after the increase in p_1.

Result 5: Once we introduce a second dimension that matters, such as identity, and even in the case of negative self-selection on skills on average, we may also be able to attract more high skilled women that had also high identity costs.

As we will see, our experiment raises expected returns for women in the technology sector, so we interpret it as increasing p_1 which has both the effect of increasing expected monetary returns for women but also of reducing the discount due to identity cost. The key variables to track in this model are expected returns in tech, expected returns in the outside option, identity costs and the underlying cognitive skills.

3. Context

Our study is conducted in Lima (Peru) and Mexico City in conjunction with a non-profit organization that empowers women youth from low-income backgrounds in Peru, Mexico and Chile with education and employment in the tech sector.[4] The program recruits young women (18-30 years old) who lack access to higher education, takes them through an immersive five-month digital coding "bootcamp" and connects them, upon graduation, with local tech companies in search for coders. In what follows, we describe the key aspects of the program.

Recruitment. Calls for applications are launched twice a year. The training provider runs targeted advertising campaigns in social media while receiving publicity in various local media. Interested candidates are asked to apply online and directed to a registration website which provides detailed information about the program and the eligibility criteria.

Evaluation and selection of top candidates. Applicants must attend two examination sessions as part of the selection process and they are assessed and selected to the program based on their results in these examinations. In the first session, candidates take cognitive abilities tests as well as a simulation measuring specific coding abilities. In a second stage, interpersonal skills and traits like motivation, perseverance and commitment are evaluated through a personal interview and group dynamics.

Training. Admitted participants begin an intensive five-month training program in web development in which students achieve an intermediate level of the most common front-end web development languages and tools (HTML5, CCS3, JavaScript, Bootstrap, Sass and Github). They also receive English reading lessons given that web languages

[4] www.laboratoria.la

and tools are written in English. Technical skill development is also complemented with personal growth and mentorship activities with professional psychologists that build the students' self-esteem, communication ability, conflict-resolution capacity and adaptability.

Placement in the Job Market. Upon training completion, the organization places students in the job market. For this, the organization has built a local network of partner companies committed to hiring their graduates. These companies are also involved in the design of program's curricula as a way to ensure that participants develop skills in high demand. In addition, the organization's sustainability is based on an Impact Sourcing model in which they, as an organization, offer web development services to companies and hire recent graduates to deliver these services. On average, and combining both sources, around 2/3 of the program's trainees find a job in the tech sector upon graduation.[5]

Cost of the program. As part of their social design, the organization charges trainees a sum of around US$15 per month of training (below the actual cost of training). If trainees end up with a job in the tech sector, then they are asked to repay the full cost of the program (around US$3,000) by contributing between 10% to 15% of their monthly salary up to the total program cost.

As of 2016, the training provider was interested in increasing application rates and assessing how to attract a better pool of applicants. They felt that despite the attractiveness of the program (over 60% of their graduates in their first two cohorts found a job in the tech sector upon graduation), sector growth potential and the low risk and cost of the program, total numbers of registered applicants were relatively low.

After completing two cohorts of trainees in Lima, the organization was launching a new operation in Arequipa in the first semester of 2016, and developing training sites in Mexico City and Santiago de Chile. We tested our intervention design in a pilot in Arequipa, where the organization was not known. We then launched our first large scale experiment in Lima, their largest operation, in their call for applications for the class starting training in the second semester of 2016. We launched the second experiment in Mexico City for the class starting training in the first semester of 2017.

4. Interventions and Research Design

[5] We are currently also evaluating the impact of the program itself.

The evidence we provide in what follows comes from two experiments and the follow up surveys of applicants to the program. In the first experiment (Lima, summer 2016) we tested the effect of a "de-biasing message" with three types of information on application rates and on the characteristics of women that self-select into the program. In the second experiment (Mexico City, winter 2016) we were able to separate out the three components of the initial message to assess which was/were responsible for the increase in response rates.

The experiments aim to first, assess whether a de-biasing message is effective in increasing application rates to the training program and second, evaluate what type of selection is induced by the de-biasing. In the context of our framework, and against the background of the Roy/Borjas model, we infer from the changes in observed self-selection what are the types of barriers that women were faced with, limiting their decision to apply for training, and in particular whether "identity" plays a role.

4.1 The first experiment: Lima summer 2016

As mentioned, to apply to the training program, one has to go the organization's registration webpage. In the application page, the organization provides detailed information about the program as well as the eligibility criteria. At the end of this page, interested applicants can find the application form.

The information provided on the program that all potential applicants saw (the control) includes the following categories:

Web Development: "You will learn to make web pages and applications with the latest languages and tools. You will learn to code in HTML, CSS, Java Script and others. In 5 months you will be able to build webpages like this one (that was done by one of our graduates)".

Personal growth: "Our objective is to prepare you for work, not only to give you a diploma. That is why we complement your technical training with personal training. With creativity workshops and mentorships, we will strengthen your abilities: we will work on your self-esteem, emotional intelligence, leadership and professional abilities."

A career in the tech sector. "Our basic training lasts 5 months, but that is just the beginning. If you succeed in this course, you will start a career as coder having access to more income. Through specializations, we offer you a program of continuous formation for the next 2 years."

The only difference between our control and treatment messages is that the treatment message included two additional paragraphs aiming to "de-bias" perceptions and beliefs on the prospects of women in the technology sector. Conceptually this message included three different additional pieces of information: (1) the fact women can be successful in the sector (2) the fact that the organization gives access to a network of women in the sector and (3) a role model: the story of a recent graduate. This first experiment therefore "bundles" three different pieces of information. Our attempt to separate those out after seeing the results of this experiment is what gave rise to the Mexico City experiment a few months later where we explicitly varied these three components.

In practice this is the exact text of the de-biasing message in Lima:

"A program solely for women. The tech sector is in need for more women bringing diversity and innovation. That is why our program is solely for women. Our experience tells us that women can have a lot of success in this sector, adding up a special perspective and sensibility. We have already trained over 100 young women that are working with success in the digital sector. They all are part of our family of coders. Women youth like you, with a lot of potential."

This text was followed by the story and picture of one of the organization's recent graduates who was successfully working in the tech sector: "Get to know the story of Arabela". The actual control and treatment messages (in Spanish) can be seen in Figure 1.

3.1.1 Data Collection on Selection Days

After applying, women attended a two-day selection process where we were able to collect information on a number of relevant characteristics that try to capture the variables in the model. In particular we collected data about the following:

A) Expected returns: In a survey, we asked them what they would expect to earn after three years of experience as a web developer, and also what they would expect to earn after three years of experience as a sales person, which is a common outside option for these women. In the context of our model, this gives us a (self-reported) measure of $P_0 S$ and $P_1 T$ for those who applied. Note that it is unusual to have a measure of the outside option for those who apply, albeit subjective (in most applications of the Roy Model one

observes returns only on the selected sample –e.g migrants, or women in the workforce-, not the "expected" outside option).

B) Cognitive Skills: The first stage in the training provider's selection process comprises three cognitive tests: two exams measuring math and logic skills, and a coding simulation exercise measuring tech capabilities. A test called "Code Academy" is a coding simulation that tested how quickly test takers are to understand basic coding and put it into place. This was taken from codeacademy.com. A second test "Prueba Laboratoria" is a test the training provider developed with psychologists to test cognitive skills. We also use an equally weighted average of the two (cognitive score). Both tests are very good predictors of the probability of success in the training, in particular the Code Academy test, so we interpret these as capturing the underlying cognitive skills that are useful in technology.

C) Gender Identity: In order to measure the implicit perceptions of women and their association of women to success in technology, we used two variables. 1) The first is based on an implicit association test (IAT). The IAT measures the strength of association between different categories, and hence the strength of the stereotype. IATs have been created to study different implicit associations/biases/prejudices (race and intelligence, gender and career etc). We created a new IAT to see how much (how little!) people associate women and technology. It asks participants to associate male or female words (Man, Father, Masculine, Husband, Son vs/ Feminine, Daughter, Wife, Woman, Mother) to technology or services words (Programming, Computing, Web development, IT, Code, Technology vs/ Cooking, Hairdressing, Sewing, Hostelry, Tourism, Services, Secretariat). The test measures how much faster the applicant is to associate male to technology and female to services than the opposite combination. 2) The second variable is based on answer to survey questions. We asked participants: if you think about yourself 10 years from now, will you be: married? With children? In charge of household duties?. Three possible answers, (No, Maybe, Yes) were available to them. We coded these as 1 2 and 3 and took the average answer. The higher the score the more the woman sees herself in a "traditional" role.

D) Other variables: The training company also collected other information on applicants as part of the selection process. In the context of our work, we asked them to implement tests to estimate risk and time preferences, with the idea that the self-selection may have operated on women with different preferences. The time preference variable

elicited from applicants the minimum monetary amount (in Peruvian Soles) the applicant required to have 3 months into the future be indifferent between receiving 50 Soles today and that amount. The risk preference variable is the minimum required as certain instead of a lottery with 50% chances of winning 150 soles or 50% change of winning nothing.

4.2 The second experiment: Mexico City winter 2016

In the first experiment, the treatment included several pieces of information bundled into the message. Given the very strong response we observed from the treatment, we wanted to assess what piece(s) of information women were actually responding to. We then ran a second experiment in Mexico City, which is a larger market and where the organization was less known so that information is more salient (this was only the second cohort of trainees in Mexico, but the organization was gaining a lot of press and notoriety in Peru during the fall of 2016). Furthermore, given the success of the first experiment, the organization really wanted to use our "de-biasing" message, and was concerned about jeopardizing applications if the old control was used. So, in the second experiment, the control group is the full de-biasing message and we take out one piece of information at a time. In addition to all the basic information, the control now includes explicit messages about (1) the fact that women can be successful in the sector ("returns") (2) the fact that the organization gives access to a network of women in the sector ("women network") and (3) a role model: the story of a recent graduate ("role model"). We implement three treatments that take one piece of information out at a time.

4.3 Randomization

We randomized the messages directly at the training provider's registration website by unique user visiting the website. To randomize the information provided in the registration page we used the Visual Web Optimizer (VWO) software.[6] To boost traffic, we launched three targeted ad campaigns in Facebook. Traffic results (total and by treatment message) are shown in Table 1. Our advertising campaigns launched in

[6] The only caveat to randomization with this strategy is that if the same user logged in multiple times from different computers, she may have seen different messages. We are only able to register the application of the last page she saw.

social media -as well as program publicity obtained through various local media- led to a total traffic to the program information and registration website of 5,387 unique users. Through our randomization, roughly half of these users saw each recruitment message.

5. Impact of the de-biasing intervention: Results from the first experiment (Lima 2016)

In this section, we report four sets of results from our first experiment. In section 5.1, we evaluate the effect of receiving the de-biasing message on the size of the pool of applicants (application rates) as well as rates of attendance to the examination by type of recruitment message. In section 4.2, we examine differences in the cognitive abilities tested and used by Laboratoria to select the candidates for training. In section 4.3, we report results for a series of tests that we designed to measure the extent of gender bias and other non-cognitive abilities of applicants. Finally, in section 4.4 we report differences in socioeconomic background and aspirations from a baseline survey we implemented to all applicants.

5.1 Application rates and attendance to selection examinations

The experiment is designed to raise expected returns in technology p_1 . Column 1 in Table 2 reports the results on differential application rates by recruitment message: essentially, our de-biasing message doubled application rates--15% of those who were exposed to treatment, or 414, applied to the program, versus only 7%, or 191, in the control group, and this difference is highly significant. We had piloted the de-biasing message in Arequipa a few months earlier on a smaller target population, with a slightly different control message, and we also found a significant doubling of application rates there (See Appendix).

This result means that the simple message had an impact on women's willingness to enter the technology training. The magnitude of the effect is quite striking, but in order to understand the mechanisms driving this change in behavior we need to do more. In particular, since this is a "bundled" treatment (many things changed at the same time between the treatment and the control). For example, the treatment contains a photograph of Arabela and the control does not. Is a picture the driver? Our pilot in Arequipa did not contain any images (only text) and we obtained similar

magnitudes of the treatment there. Could it be the exact wording? As we will see later, the wording is different in our Mexico experiment and was slightly different in the Arequipa experiment, and we obtain similar results, so this suggests it is about the information provided in the treatment message, not the precise wording or the presence of a picture. Could it be that the treatment offers just more information and with more information candidates are more likely to apply? As we will see in the Mexico experiment, it is not just more information but specific types of information that women respond to. Of course, de-biasing someone is typically associated to providing new information, but the key is to understand what "priors" is that additional information affecting. So, next we turn to analyze the change in self-selection with the treatment message to infer what is the relevant information that is changing these women's priors, and to what extent identity is one of the dimensions we are affecting.

As discussed, all registered applicants have to attend two days of examinations to be evaluated for admission into the program. From the day of registration to the examination dates there could be up to a month difference. Traditionally, attendance to examinations has ranged between 30 to 35% of all registered applicants. In column 2 of Table 2 we report attendance rates to the examination dates by treatment group. We observe that, despite the much larger numbers of applicants coming from the treatment message, there is no significant difference in the ratio of applicants coming to the examinations between the two groups. So this rules out that the results we describe in what follows on selection is driven by the fact that treatment affected attendance to the exams.

It is important to highlight that differences in application rates highly influence the distribution of candidates attending the selection process. Of the total 202 candidates attending, 66% had been exposed to the treatment message.

5.2 Self-Selection Patterns: Expected returns and Cognitive Skills

Table 3 shows the differential selection following the treatment on the logarithm of expected returns in technology (column1), in sales (column 2) and the difference between the two (column 3). We regress the variables of interest on the treatment variable. Note that here we are only estimating the differential selection in treatment and control, and not a causal effect of treatment on the outcome variables. We are looking at how the equilibrium selection changes following the exogenous shock

(treatment). We discuss later on why we thing treatment effects of de-biasing on performance are minimal relative to the effect on selection.

The results in Table 3 suggest negative selection in both technology and services/sales skills. The effect is clear and highly significant in column 2 where the women that apply to Laboratoria under treatment have an outside option that is 23% lower than those in the control. In terms of our model, given P0 is unchanged with the experiment this is suggests average S falls. For technology skills, we see a negative effect (-0.115) that is not significant. But this is likely driven by the fact that if average T decreases (negative selection) as p_1 increases (the experiment message), the net effect of the two is ambiguous. P_1T . They fall, although not significantly.

In order to measure skills directly (not confounded by the returns that change with the experiment). We analyze the change in selection of cognitive skills following the de-biasing message. This is shown in Table 4. We find that average cognitive skills measure by both the "Code Academy" and "Prueba Laboratoria" tests are 0.26 to 0.28 of a standard deviation lower lower in the treatment group. There is clear negative selection in cognitive skills.

These selection patterns suggest that we are in a world of comparative advantage, where skills in technology and sales (the outside option) is positive but not very high (otherwise we'd have positive selection).

5.3. Self-Selection Patterns: Identity

We turn next to analyze self-selection patterns on our measures of gender identity in Table 5. We find that following the women that apply following the de-biasing message are on average more "biased" as measured by the IAT we developed on the association of women and technology as well as on the survey measure for "Traditional Role". The magnitude of the negative self-selection on identity is large: 0.29 of a standard deviation more biased for the IAT and 0.39 of a standard deviation higher association with a traditional role. Figures 3 and 4 show the raw distribution of the identity variables and reflects this pattern.

This suggests that the correlation between identity costs and the difference between technology and services skills is positive but not very high. Therefore the marginal women that apply are "more biased" following the treatment message.

We interpret our results as reflecting mostly "Selection" and argue that it is unlikely that the de-biasing message has a significant causal effect on some of the outcomes we measure (like social identity and cognitive skills).[7] This is because (1) up to a month passes between application and the days of the test, so any treatment effect is unlikely to persist into the selection days; (2) when applicants arrive to the training provider for the tests they have received much more information on Laboratoria and the future of its graduates, where we think that the gap in information between the two groups is much smaller once they take the test; and finally, (3) because our prior is that if anything the de-biasing, to the extent that it reduces stereotype threat (Steele and Aaronson 1999) would help them do better in tests and have lower biases, and this would bias our estimates positively. Given we still find negative selection on all dimensions we think the treatment effect of the message on performance is dwarfed by the selection effects we identify.

5.4. Selection at the Top

The results so far suggest that the average women applying is of worse technology/cognitive skills and has a higher average bias against women in technology and more traditional view of the role of women in society. This allows us to understand, in the light of the Roy model, the underlying correlation between these dimensions in our populations, as well as the type of comparative advantage in place in this economy. However, from the point of view of the training firm, one might be worried that it is not allowing them to do what they were aiming for: attracting more high quality candidates.

Fortunately, these mean effects obscure what is happening along the distribution. In fact, the training provider is interested in attracting a higher number of "right tail" candidates to select from. As overall numbers increase, do the number of highly qualified women increase in spite of the fall in the mean quality? In the bottom panel of Table 4, we compare the cognitive skills of the top 50 performers in each experimental group (50 is the size of the population to be admitted into the program). We find that those treated report significantly higher average cognitive scores and ad-hoc tech capabilities (0.37 standard deviation higher score in the Code Academy simulation and 0.36 higher average score).

[7] The only exception is expected returns in tech, were treatment is likely to raise these beliefs. In this case, we have both a treatment effect on p1 and selection on tech skills.

These results suggest that the treatment affects differentially candidates by level of cognitive ability: it increases the number of applicants at all levels of cognitive ability, but it particularly does so at the bottom of the distribution. Figure 2 shows the frequency of applicants in treatment and control that reflects this pattern.

What about social identity at the top? Panel B in Table 5 shows the difference in the average IAT, and traditional role variables for the top 50 candidates ranked by cognitive score. We have large standard errors and none of the variables is significant but, if anything, the results suggest that the average applicant, with the higher cognitive scores is more biased/has a larger identity cost in the treatment than in the control group.

These selection patterns at the top are consistent with some women applying under treatment who are high skill but also have a high identity costs, suggesting that identity not only matters on average, but also that it is one of the dimensions precluding high cognitive skill women from attempting a career in the Tech sector.

5.4 Interest in Technology, time and risk preferences

During the training provider's examination period, we were able to measure other non-cognitive traits for all applicants like time and risk preferences, and we asked women about their prior interest in Technology. [8]

Table 7 shows that there are no significant differences between those treated and non-treated in terms of prior interest in technology, time and risk preferences. This allows us to rule out that the selection is operating because the treatment affects those dimensions.

4.5 Trading-off attributes

So, overall we find that the de-biasing message doubles application rates. It attracts women at all levels of cognitive ability but disproportionately so at lower levels of ability. It also attracts women who are more biased on average in relation to the role of women in technology (as measured by the IAT and the traditional role survey variables). This self-selection pattern is consistent with a model where women apply if they perceive that their (monetary and non-monetary/psychological) returns to being in an industry is higher than a given threshold. The treatment effectively increases the

[8] Using Survey (Falk, cite)

perceived returns relative to the threshold so it should attract disproportionately more women that were just below the threshold. Women above the threshold in the control would still apply under treatment, but it is women who were below the threshold that apply under the treatment. This is exactly what we find.

However, the message increases application rates also for high cognitive ability women, which is the pool that the organization is interested in attracting. This is consistent with a multiple index model where the application threshold is determined not just by one single attribute (the standard monetary expected returns) but also by non-monetary ones. To the extent that cognitive ability is related to expected returns, we next investigate how the treatment affects expected returns at each level of ability.

6. Identifying the drivers of the bias: Results from the second experiment (Mexico D.F. 2016):

The results from the Lima experiment show that application rates doubled when women were exposed to the de-biasing message (in the pilot we ran in Arequipa application rates also doubled). However, given this was a bundled message we do not know what is the piece of information that triggered the increased application rate. In order to see that, we collaborated again in the winter of 2016 with the organization to implement the second experiment in Mexico City.

In this follow-up experiment we decomposed each prior element of treatment. To address concerns by the training provider of not maximizing the number of applicants (they had seen how applications rates doubled with our prior treatment), we considered a control group with all previous treatment components, and eliminated, one by one, each of its component so that the four experimental groups resulted as follows:

- Control: all components (success/returns, network, role model)
- T1: network and role model (eliminate success/returns)
- T2: success/returns and role model (eliminate network)
- T3: success/returns and network (eliminate role model)

Again, we randomized at the trainer providers' registration website URL by unique user and we launched three targeted advertising campaigns in Facebook to attract more traffic. Results are provided in Table 8.

Conversion rates in the control group attain 8.9%. We can then see how both treatments that eliminate the role model and the "women can be successful" component significantly reduce the probability of applying for training: the treatment that eliminates the role model reduces the conversion rate by 2 percentage points or 23%, while the treatment that eliminates the "women can be successful" component, reduces the conversion rate by 18%. We conclude thus that the key components of treatment are the role model and addressing the fact women can be successful in the sector. This second experiment also allows us to address external validity: we found similar results to the treatment in the Arequipa pilot, Lima and Mexico DF experiments, that is in different time periods and different countries, suggesting that the informational content of our experiment really is able to alter behavior and self-selection into the industry.

7. Conclusion

We experimentally varied the information provided to potential applicants to a 5-month digital coding bootcamp offered solely to women. In addition to a control message with generic information, in a first experiment we corrected misperceptions about women's ability to pursue a career in technology, provided role models, and highlighted the fact that the program facilitated the development a network of friends and contacts in the Tech sector.

Treatment exposure doubled the probability of applying to training (from 7% to 15%). On average, however, the group exposed to treatment reported an average cognitive score which is 17% below the control group. We also find that among the population that would have been selected for training (top performers in examinations), cognitive and tech specific abilities are 22% and 23% higher than those that are treated. Our empowerment message thus appears to be increasing the interest of women in pursuing a career in the tech sector at all levels of ability, but proportionally more for those with lower ability.

In a follow up experiment, we decomposed the three components of treatment: addressing the probability of success for women, the provision of a role model and the development of a network of friends and contacts. We find that the key components are the provision of a role model and the de-biasing about the success of women in the Tech sector.

References

Ahern, Kenneth R. and Amy K. Dittmar. 2012. "The Changing Of The Boards: The Impact On Firm Valuation Of Mandated Female Board Representation". *The Quarterly Journal Of Economics* 127 (1): 137-197.

Akerlof, George A., and Rachel E. Kranton. (2000) "Economics and identity." *The Quarterly Journal of Economics* 115.3: 715-753.

Ashraf, Nava, Oriana Bandiera, and Scott S. Lee. 2014. "Do-gooders and go-getters: career incentives, selection, and performance in public service delivery." *STICERD-Economic Organisation and Public Policy Discussion Papers Series* 54.

Borjas, G. J. (1987). Self-Selection and the Earnings of Immigrants. *The American Economic Review*, 531-553

Dal Bó, Ernesto, Frederico Finan, and Martín 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.

Flory, J. A., A. Leibbrandt, and J. A. List. 2014. "Do Competitive Workplaces Deter Female Workers? A Large-Scale Natural Field Experiment On Job Entry Decisions". *The Review Of Economic Studies* 82 (1): 122-155.

Forbes, Chad E. and Toni Schmader. 2010. "Retraining Attitudes And Stereotypes To Affect Motivation And Cognitive Capacity Under Stereotype Threat.". *Journal Of Personality And Social Psychology* 99 (5): 740-754.

Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104 (4) :1091-1119.

Good, Catherine, Joshua Aronson, and Michael Inzlicht. 2003. "Improving Adolescents' Standardized Test Performance: An Intervention To Reduce The Effects Of Stereotype Threat". *Journal Of Applied Developmental Psychology* 24 (6): 645-662.

Kawakami, K., D.M. Amodio, and K. Hugenberg. 2017. "Intergroup Perception And Cognition: An Integrative Framework For Understanding The Causes And Consequences Of Social Categorization". In *Advances In Experimental Social Psychology Volume 55*, 2-323.

Marinescu, Ioana and Ronald Wolthoff. 2017. ""Opening The Black Box Of The Matching Function: The Power Of Words"". *NBER Working Paper*.

Matsa, David A and Amalia R Miller. 2013. "A Female Style In Corporate Leadership? Evidence From Quotas". *American Economic Journal: Applied Economics* 5 (3): 136-169.

Paluck, Elizabeth Levy and Donald P. Green. 2009. "Prejudice Reduction: What Works? A Review And Assessment Of Research And Practice". *Annual Review Of Psychology* 60 (1): 339-367.

Roy, A. D. 1951. "Some Thoughts On The Distribution Of Earnings". *Oxford Economic Papers* 3 (2): 135-146.

Spencer, S. J., Steele, C. M., & Quinn, D. M. 1999. “Stereotype threat and women's math performance.” *Journal of experimental social psychology*, *35*(1), 4-28.

Steele, Claude M. and Joshua Aronson. 1995. "Stereotype Threat And The Intellectual Test Performance Of African Americans." *Journal Of Personality And Social Psychology* 69 (5): 797-811.

Tables

Table 1: Traffic to site

	Traffic to "Postula URL"	
	Traffic	Conversions
Total	5387	605
De-biasing message	2763	414
Control	2624	191

Table 2: Effect of de-biasing message on application rates and exam attendance

	(1)	(2)
	Application rate	Attendance
Treated	0.077***	-0.022
	(-0.01)	(-0.04)
Mean of the dependent variable in control	0.07	0.35
Observations	5387	608

Standard errors in parentheses
* p<0.10 ** p<0.05 *** p<0.01"

Table 3: Expected Returns

	(1) Log Webdev income	(2) Log Salesperson income	(3) Log salary dif.
Treated	-0.115	-0.231***	0.111
	(0.081)	(0.084)	(0.068)
Mean of the dependent variable in control	7.969***	7.534***	0.441***
	(0.066)	(0.068)	(0.055)
Observations	197	196	196
Adjusted R-squared	0.005	0.033	0.009

Standard errors in parentheses
"* p<0.10 ** p<0.05 ** p<0.01"

Table 4: Cognitive abilities

	Panel A: All Observations		
	(1) Code Academy (std)	(2) Prueba Lab (std)	(3) Cog. Score (std)
Treated	-0.268*	-0.278*	-0.316**
	(0.149)	(0.159)	(0.158)
Mean of the dependent variable in control	0.178	0.182	0.207
	(0.121)	(0.128)	(0.128)
Observations	200	174	174
Adjusted R-squared	0.011	0.012	0.017

	Panel B: Top 50 Candidates by Cognitive Score		
	(1) Code Academy (std)	(2) Prueba Lab (std)	(3) Cog. Score (std)
Treated	0.373**	-0.163	0.349**
	(0.159)	(0.190)	(0.155)
Mean of the dependent variable in control	0.552**	0.418**	0.486**
	(0.112)	(0.134)	(0.109)
Observations	100	100	100
Adjusted R-squared	0.044	-0.003	0.040

Standard errors in parentheses

"* p<0.10 ** p<0.05 ** p<0.01"

Table 5: Social Identity

Panel A: All Observations

	(1) IAT Gender/Career (std)	(2) IAT Gender/Tech (std)	(3) Traditional Role (std)
Treated	-0.125	-0.290*	0.380**
	(0.159)	(0.157)	(0.148)
Mean of the dependent variable in control	0.080	0.190	-0.252**
	(0.127)	(0.127)	(0.120)
Observations	171	178	199
Adjusted R-squared	-0.002	0.013	0.028

Panel B: Top 50 Candidates by Cognitive Score

	(1) IAT Gender/Career (std)	(2) IAT Gender/Tech (std)	(3) Traditional Role (std)
Treated	-0.262	-0.128	0.215
	(0.206)	(0.187)	(0.189)
Mean of the dependent variable in control	0.150	0.100	-0.318**
	(0.144)	(0.134)	(0.134)
Observations	92	95	100
Adjusted R-squared	0.007	-0.006	0.003

Standard errors in parentheses
"* p<0.10 ** p<0.05 ***p<0.01"

Table 6: Pairwise Correlations between variables

	(1) Log Webdev income	(2) Log Salesperson income	(3) Log salary dif.	(4) Cog. Score (std)	(5) IAT Gender/Tech (std)	(6) Traditional Role (std)
Log Webdev income	1					
Log Salesperson income	0.671***	1				
	0.00					
Log salary dif.	0.363***	-0.446***	1			
	0.00	0.00				
Cog. Score (std)	0.254***	0.235***	0.013	1		
	0.00	0.002	0.87			
IAT Gender/Tech (std)	-0.0051	0.0173	-0.281	0.0403	1	
	0.946	0.819	0.711	0.621		
Traditional Role (std)	0.081	0.017	0.077	-0.132*	-0.807	1
	0.258	0.81	0.286	0.085	0.285	

P-Values in parentheses
* p<0.10 ** p<0.05 ***p<0.01

Table 7: Other Preferences

	(1) Wanted to study technology prior to application	(2) Risk Preferences (std)	(3) Time Preferences (std)
Treated	-0.016	0.196	0.173
	(0.079)	(0.162)	(0.162)
Mean of the dependent variable in control	0.516**	-0.128	-0.113
	(0.064)	(0.131)	(0.131)
Observations	182	168	168
Adjusted R-squared	-0.005	0.003	0.001

Standard errors in parentheses
* p<0.10 ** p<0.05 *** p<0.01
Note: Time preference is the minimum required to have in 3 months instead of 50 soles today. Risk preference is

the minimum required as certain instead of a lottery with 50% chances of winning 150 soles or same chance of winning nothing

Table 8: Follow-up experiment in Mexico, Treatment Decomposition

	Dependent Variable:
	Application Rate
T1: Network and Role Model	-.016*
	(0.01)
T2: Success and Role Model	-.001
	(0.01)
T3:Network and Success	-.020**
	(0.01)
Control group	0.087***
	(0.007)
Observations	6183

Standard errors in parentheses
* p<0.10 ** p<0.05 *** p<0.01

FIGURES

Figure 1A: Application Message in Lima 2016

The Treatment message added the elements that are circled in Red to the Control

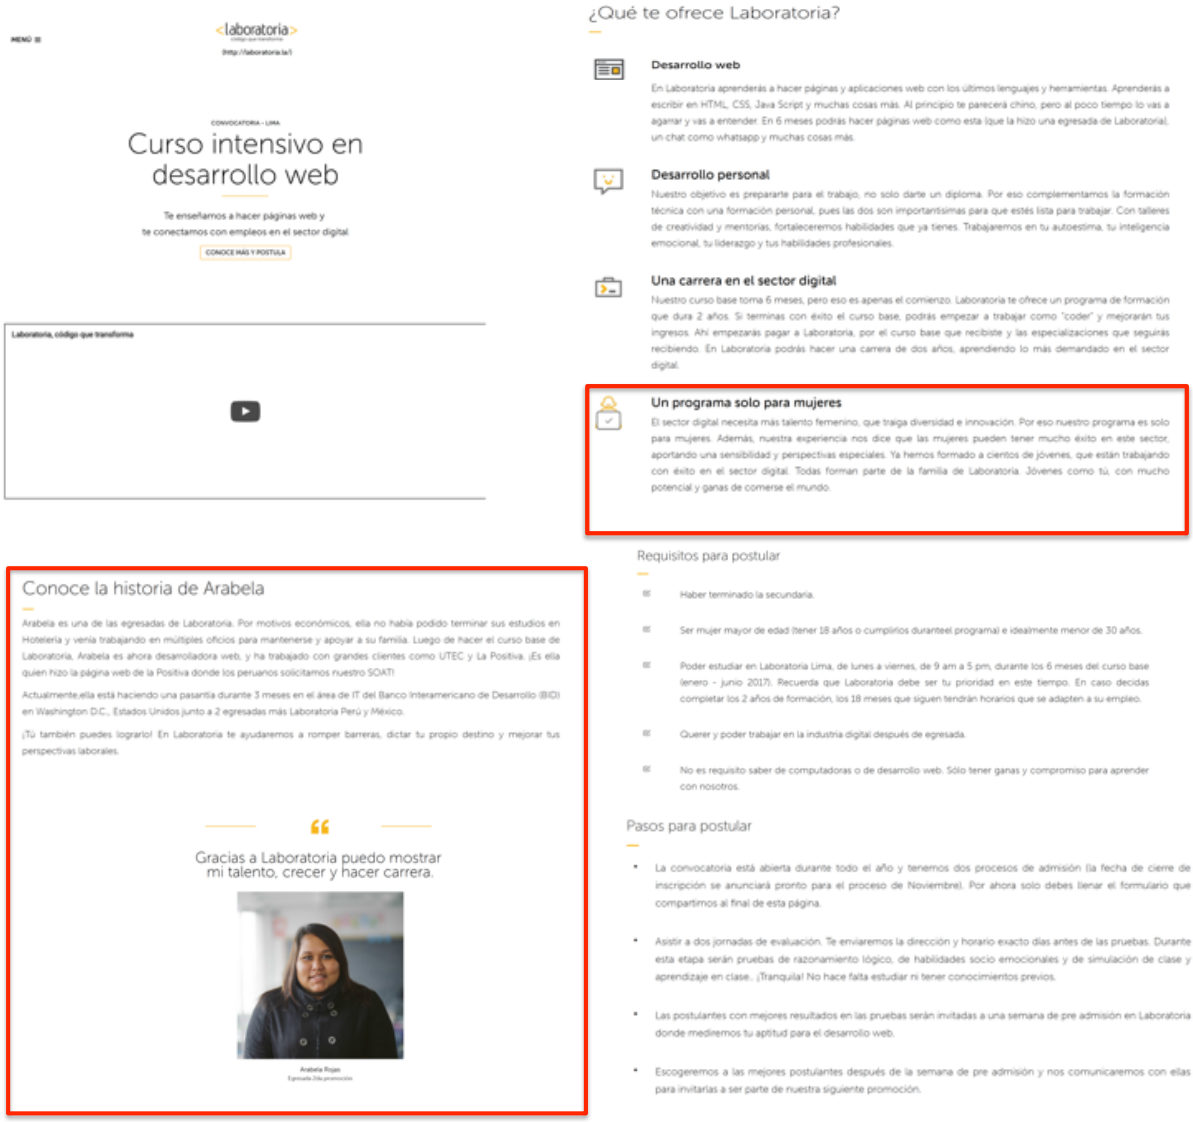


Figure 1B: Application Message (continued)

Postula

Nombres: *

Apellidos: *

Edad: *

Correo Electrónico: *

Documento de Identidad (DNI): *

Teléfono *

¿Cómo te enteraste de Laboratoria? *

- Facebook
- Radio
- Televisión
- Charla en mi comunidad
- Diarios o medios impresos
- Familia o amigo me avisó
- Otros

Si seleccionó otros medios

¿Cuál es tu motivación para estudiar en Laboratoria?: *

¡Recibe novedades de Laboratoria!*

- Acepto

Figure 2: Distribution of Cognitive Scores in Control (0) and Treatment (1)

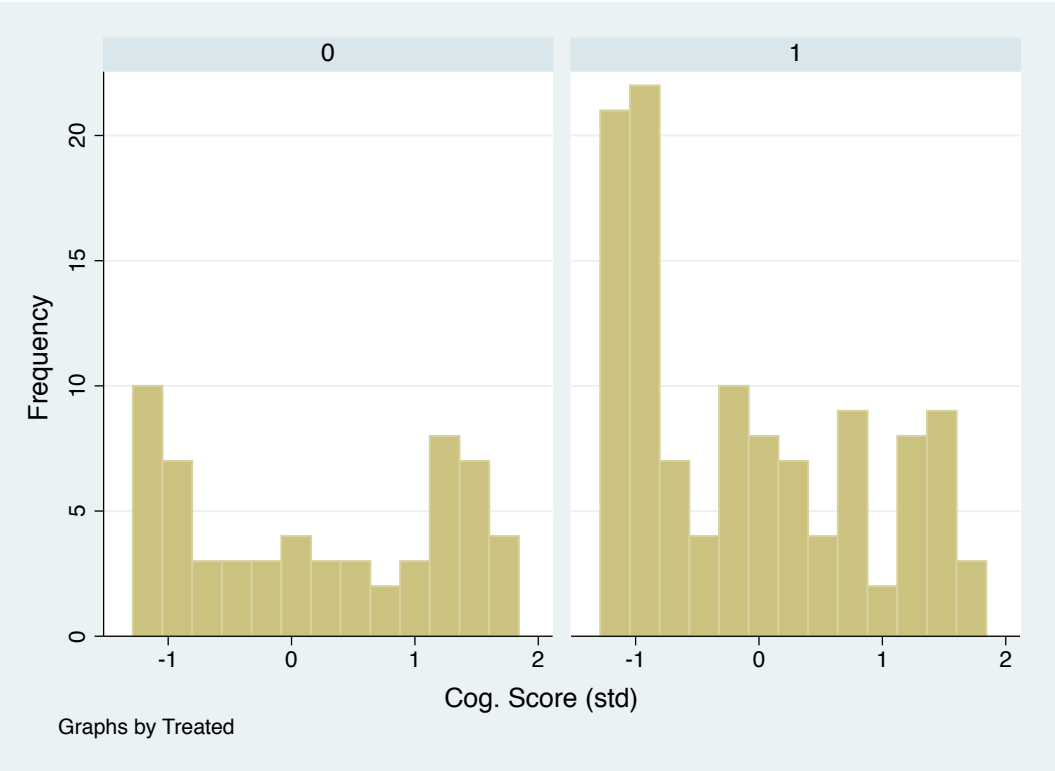


Figure 3: Distribution of Traditional Role survey variable in Control (0) and Treatment (1)

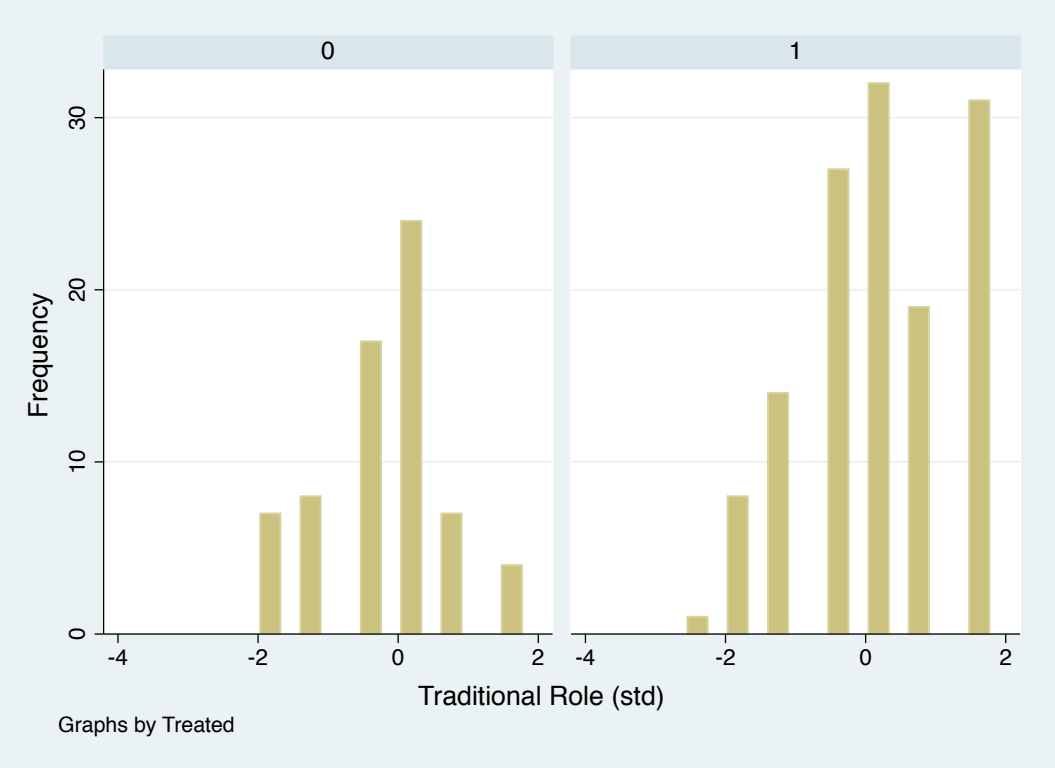


Figure 4: Distribution of IAT Technology/Services in Control (0) and Treatment (1)

