

# Geographical Discrimination in the Gig Economy

Hernan Galperin  
Annenberg School for Communication  
University of Southern California  
3502 Watt Way, Los Angeles CA 90089-0281  
hernan.galperin@usc.edu

Catrihel Greppi  
Universidad Nacional de La Plata  
La Plata, Argentina (1900)  
catrihelgreppi@gmail.com

## ABSTRACT

Online labor platforms were expected to flatten labor markets by reducing the importance of worker location and, as a result, promote employment and wage growth in developing countries based on cost differentials. We test these propositions using transactional data from Nubelo, a large online labor platform for Spanish-speaking employers/freelance workers. The results suggest that information-related frictions long observed in traditional labor markets may be exacerbated in online platforms, resulting in worker discrimination based on country of origin. We show that, after controlling for observable workers' characteristics and their job bids, foreign job-seekers are 42 percent less likely to win contracts from Spanish employers, which represent about two-thirds of all employers in the platform. We attribute this result to the activation of stereotypes that orient employers' hiring decisions, in the absence of verifiable information about the quality of individual workers. We draw implications for platform design and the governance of online labor contracts.

## *1. Introduction*

Scholars have long noted that labor markets are shaped by geography, creating special relationships between employers and workers in close proximity. However, in recent years a new narrative has emerged that challenges the role of geography in explaining labor market outcomes. This narrative asserts that the global diffusion of high-capacity communication networks and the digitization of work is increasingly making location irrelevant, creating a globally-contested market for labor (Blinder, 2006; Farrell et al., 2005). This is the narrative of the “flat world” popularized, among others, by Thomas Friedman’s best-seller ‘The World is Flat: A Brief History of the Twenty-First Century’ (Friedman, 2005).

According to this narrative, digital work has the potential to drive wage and employment growth in emerging economies through several mechanisms (Raja et al., 2013; Rossotto et al., 2012). Driven by salary differentials, employers in advanced economies are likely to gravitate towards countries with large pools of lower-cost white-collar workers. Further, online labor platforms significantly reduce problems of spatial mismatch between workers’ abilities and opportunities in local job markets. In turn, this increases the returns to skills acquisition, thus unleashing a virtuous cycle of human capital development and job growth in developing regions.

A number of stylized facts support this narrative. First, the majority of employers in online labor platforms are based in high-income countries, while the majority of workers are based in middle and low-income countries (Agrawal et al., 2013; Lehdonvirta and Ernkvist, 2011). This simple fact suggests that workers in developing regions may be able to earn higher (hourly) wages relative to opportunities in local labor markets. Second, digital work

dramatically expands the number and range of labor opportunities, facilitating access to employers in higher-wage countries and increasing the likelihood that individual skills will be matched with available jobs. Third, online labor platforms allow employers to break down large processes into so-called microtasks, enabling individuals or small labor cooperatives to compete directly with offshoring firms that intermediate between employers and workers (Fish and Srinivasan, 2011).

This study seeks to empirically test the narrative of a flattened global market for digital labor. In particular, we test the hypothesis that cost differentials will induce employers to offshore contracts to workers in low-wage countries. Our empirical strategy is based on the examination of internal data from Nubelo, one of the largest Spanish-language online labor platforms.<sup>1</sup> We obtained records for all transactions in Nubelo over a 44-month period between March 2012 and December 2015. The dataset includes basic demographic characteristics for employers and workers, and extensive platform-specific information about contracted jobs.

Our results suggest that information-related frictions long observed in traditional labor markets remain pervasive (and may possibly be exacerbated) in online labor platforms, resulting in a significant penalty for job seekers from developing countries. This penalty works in two ways. First, we show that, after controlling for observable individual characteristics and job bids, foreign workers are 42 percent less likely to win contracts from employers in Spain, the highest-wage country in our sample and where the majority of employers are based. Second, we show that Spanish workers are able to command a

<sup>1</sup> In December 2016 Nubelo was acquired by Freelancer.com.

significant wage premium, of about 16 percent, over similarly-qualified foreign workers. The combination of a hiring and a wage penalty helps explain why the attrition rate is much higher for non-Spanish workers, who are significantly less likely to remain active in Nubelo within 12 months of joining the platform.

We offer two complementary hypotheses for these results. The first relates to the nature of the contracts outsourced through online labor platforms. Most of the job opportunities available in Nubelo require a degree of co-production between buyer and seller. This differs from other platforms, in particular Mechanical Turk, where demand is typically for very small, low-skill tasks that require minimal communication between employer and worker (Irani, 2013). Because employers in Nubelo anticipate higher communication costs when working with foreign contractors, the balance tilts in favor of domestic workers.

However, the countries served by Nubelo (which mainly encompass Spain and Latin America) are relatively homogeneous in terms of culture, language and (to a lesser extent) time zone. Therefore, our second and most relevant hypothesis relates to information asymmetries and uncertainty about worker quality. In online labor platforms, employers evaluate candidates on the basis of very limited verifiable information. Further, employers are unable to observe latent worker characteristics which may be inferred in traditional hiring contexts (e.g., through personal job interviews).

Under these circumstances, stereotypes that relate country of origin to worker quality provide cognitive shortcuts that orient hiring choices. This is similar to the spatial signaling mechanism whereby employers discriminate against job seekers from US inner cities or living

in public housing (Neckerman and Kirschenman, 1991). We provide empirical evidence for this mechanism by showing that Spanish employers adjust preferences based on the amount of platform-verified information about workers. In particular, the hiring advantage for domestic workers falls as more information about individual workers is available, and as employers acquire experience in contracting foreign workers. This suggests employers discriminate due to information uncertainty rather than distaste for hiring workers from other countries.

This study contributes to the emerging literature on the dynamics and socioeconomic impact of digital labor. Our contribution to this literature is threefold. First, we corroborate previous findings about the continued salience of geographical location in labor markets in a setting where language and other cultural factors are by and large irrelevant. Second, we provide a novel measure to quantify wage differentials between foreign and domestic workers, and explore how wage penalties change when information uncertainty is reduced. Third, we propose a path-dependent mechanism that suggests why, despite very low entry-costs, workers from developing countries are less likely to remain active in digital labor markets in the long run.

## *2. Discrimination in labor markets and the emergence of digital work*

Online platforms that facilitate matching for contingent work (often called the “gig economy”) have been the subject of significant scholarly attention in recent years. From a macro perspective, several scholars have examined how market design features embedded in these platforms are exacerbating power imbalances between employers and workers (e.g.,

Kingsley et al., 2015). A related literature has examined how contingent work, despite being framed in a discourse of entrepreneurship and family-work balance, may be eroding workers' rights, trapping job seekers from disadvantaged groups in precarious work arrangements (Huws, 2015; Mann and Graham 2016).

The starting point for this body of work is that the Internet is rapidly changing how labor markets operate. Following Autor (2001), we identify three dimensions of such change. First, search costs are significantly reduced, potentially improving matching between employers and job seekers. Second, the digitization of tasks results in more workers capable of performing their tasks remotely. Third, online platforms make geographical proximity between employers and workers less relevant, potentially flattening the labor market.

This study examines the hypothesis of a flattening labor market in which geographical proximity between employers and job seekers is increasingly irrelevant. Before the diffusion of the Internet and the emergence of online platforms, spatial distance benefitted workers in close proximity to the most productive firms. This resulted in better employment opportunities and higher wages for workers in advanced economies, since workers in other countries were by and large prevented from competing for these jobs due to barriers to migration as well as high search and communication costs. Online labor platforms are hypothesized to create a more level playing field in which workers compete for contracts regardless of place of residence, nationality, race, gender or other characteristics unrelated to individual productivity (Agrawal et al., 2015).

In particular, the digitization of work has the potential to significantly reduce spatial mismatch between employers and job seekers. The spatial mismatch hypothesis posits that frictions in the housing market and under-investments in transportation result in inferior labor market outcomes for racial minorities and other disadvantaged groups due to their greater geographical distance to high-wage jobs (Kain, 1992). The key insight is that, under spatial mismatch, employers need not discriminate against racial minorities or other groups, since differential outcomes result naturally from the pool of job applicants. As geographical proximity becomes irrelevant, non-discriminating employers will hire exclusively on the basis of productivity-related characteristics, leading to better outcomes for previously disadvantaged groups.

The argument above rests on two key assumptions that deserve further examination. First, that tasks are easily codifiable and written into binding job contracts that can be monitored and enforced remotely. Second, that online employers can accurately observe the productivity-related characteristics of job seekers. Decades of scholarship on information asymmetry and transaction costs have shown that these assumptions by and large do not hold true in offline labor markets (Ashenfelter and Card, 1999; Leamer, 2007). Building on these findings, scholars have turned attention to how they operate in the context of online labor platforms.

In general, the findings suggest that information asymmetries and communication costs are far from irrelevant for explaining outcomes in online labor markets. For example, Gefen and Carmel (2008) find that most contracts in an online programming marketplace are awarded to domestic contractors. When jobs are offshored, employers prefer workers from countries with

minimal cultural distance (rather than lower costs), such as US employers hiring programmers in Canada and Australia. Hong and Pavlou (2014) find that differences in language, time zone, cultural values and levels of economic development negatively affect hiring probabilities in a global platform for IT contracts. Similar results are reported by Lehdonvirta et al. (2014), who also find that the hiring penalty for foreign job applicants increases when tasks require knowledge of formal institutions (e.g., legal work) or regular interaction with employers.

Other studies attempt to identify remedial mechanisms that mitigate these observed frictions. For example, Stanton and Thomas (2014) show that being affiliated with an outsourcing firm increases hiring and wages among inexperienced workers, helping them overcome the first-job barrier. The advantage dissipates over time and jobs as more information is available about the quality of individual workers. Mill (2011) finds that feedback from previous contracts significantly reduces the effect of geographical location on hiring. Similarly, Agrawal et al. (2013) find that the benefit of platform-verified information is disproportionately large for job applicants from less-developed countries, which suggests that employers have more difficulty evaluating quality among foreign workers. Along these lines, Ghani et al. (2014) find that ethnic Indian employers based outside India are more likely to hire Indian workers. They attribute the advantage to the familiarity of these employers with information regarding workers' qualifications rather than ethnicity-based preferences.

In general, these results are consistent with theories of statistical discrimination whereby employers, faced with uncertainty about worker productivity, attribute values to individual

job seekers based on perceived group averages (Aigner and Cain, 1977; Phelps, 1972). At its core, statistical discrimination is a theory of social stereotyping. When hiring workers, employers seek information that helps predict future productivity. If this information is too noisy or simply unavailable, stereotypes provide cognitive shortcuts that help orient hiring choices. Geographical stereotyping has long been observed in the context of traditional labor markets (Fernandez and Su, 2004). For example, Neckerman and Kirschenman (1991) found that employers in the Chicago area used the home address of job candidates as a primary screening mechanism. In a study of hiring behavior in the New York area, Newman (1999) finds a similar pattern of employer discrimination driven by job applicants' residence (more specifically associated with public housing) rather than their socioeconomic background.

Along these lines, there are several reasons why geographical stereotyping may play an even larger role in online labor platforms. Online employers are by and large unable to screen job applicants in person. Rather, they rely on two types of information generally available on a job applicant's online profile. First, platform-generated information such as the number of previously obtained jobs and a reputation score. Second, non-verifiable information voluntarily provided by applicants such as career experience and technical skills not validated by platform-administered tests.

Under some circumstances, the inability to screen candidates in person may reduce hiring biases – most famously when symphony orchestras started implementing “blind” auditions, as shown by Goldin and Rouse (2000). However, less information may also trigger stereotyping when other credible signals of worker quality are unavailable. As shown below, the amount of platform-validated information about job applicants on Nubelo (as in most online labor

platforms) is very limited. In this context, as Pallais (2014) shows, even very small differences in the amount of information available can have a significant effect on future hiring and earnings.

In addition, given the relatively small value and short-term nature of a typical digital labor contract, it is unlikely that employers will incur in high screening costs (Horton and Chilton, 2010). As a result, faced with several dozen applicants for each contract, limited verifiable information, and a short-term contract window, employers are likely to activate cognitive shortcuts in hiring decisions. Prior beliefs about the average productivity of workers based on available signals (such as country of origin) are likely to become highly salient in such contexts.

### *3. Data and descriptive results*

Nubelo matches employers who post contracts for short-term jobs with workers who bid for these jobs. Job postings typically describe the task required, the job category, the expected date of delivery, and the location of the employer. Employers select workers on the basis of the proposed bid as well as other characteristics that are visible on job applicant's online profiles. These include name, country of residence, previous work experience in the platform, and a summary feedback score from previous jobs completed. In addition, job applicants can voluntarily include other information such as a CV, a brief description of skills and work experience outside Nubelo, portfolio samples, and a personal picture.

Our dataset includes records for all transactions in Nubelo for a 44-month period between March 2012 and December 2015. They include information on all jobs posted by employers and on all bids placed by workers, both winning and unsuccessful. Unlike other platforms, Nubelo actively discourages employer-worker interaction prior to hiring. Therefore, all the information visible to employers is available in our dataset, reducing concerns about omitted variables in our estimation models. Our units of observation are the bids made by job seekers. As a result, the dataset is restricted to active contractors, by which we refer to those who have submitted at least one bid during the 44-month study period. The full dataset includes 81,497 bids made by 18,356 job seekers for a total of 5,262 jobs posted by 2,517 employers. We note appropriately when partial data subsets are used.

Nubelo targets Spanish-speaking employers and freelance workers. While 63 countries are represented in our dataset, Spain and a few large countries in Latin America account for the majority of job seekers (Table 1). On the other hand, labor demand is largely concentrated in Spain, which accounts for about two-thirds of all employers.

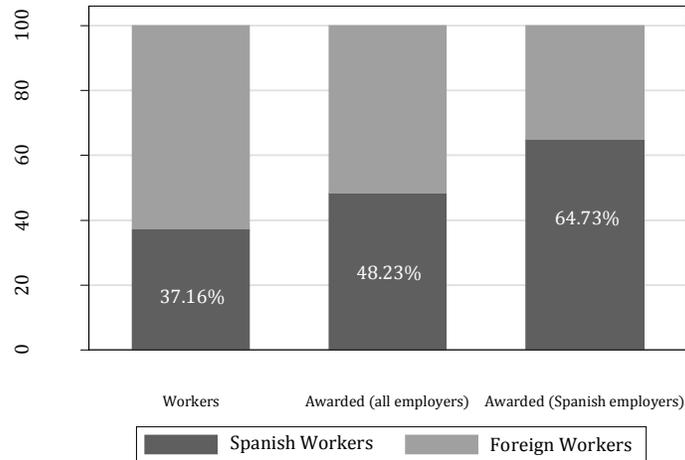
TABLE 1  
WORKERS, EMPLOYERS AND JOBS POSTED BY COUNTRY

<i>Country</i>	<i>Workers</i>	<i>%</i>	<i>Employers</i>	<i>%</i>	<i>Jobs posted</i>	<i>%</i>
Spain	6,820	37.16	1,639	65.12	3,528	67.05
Argentina	4,045	22.04	390	15.49	689	13.09
Colombia	2,144	11.68	139	5.52	222	4.22
Mexico	1,326	7.22	175	6.95	419	7.96
Venezuela	1,268	6.91	8	0.32	13	0.25
Chile	648	3.53	49	1.95	147	2.79
Peru	399	2.17	9	0.36	17	0.32
Uruguay	239	1.30	10	0.40	16	0.30
Ecuador	175	0.95	9	0.36	50	0.95
Dominican Republic	144	0.78	6	0.24	6	0.11
Others	1,148	6.25	83	3.30	155	2.95
<i>Total</i>	<i>18,356</i>	<i>100</i>	<i>2,517</i>	<i>100</i>	<i>5,262</i>	<i>100</i>

*Source:* Authors calculations based on Nubelo data

Descriptive results suggest that employers tend to favor job applicants based in Spain. As shown in Figure 1, Spanish workers win a larger-than-expected share of all jobs posted. This difference is magnified when the sample is restricted to Spanish employers: job applicants based in Spain comprise 37 percent of all workers but obtain 65 percent of the contracts originated in Spain. On the other hand, when not hiring domestically, Spanish employers hire equally from all other countries in our sample. In other words, the share of contracts awarded to workers across Latin America is proportional to their share of workers in the sample.

FIGURE 1  
WORKERS AND CONTRACTS AWARDED BY COUNTRY OF ORIGIN



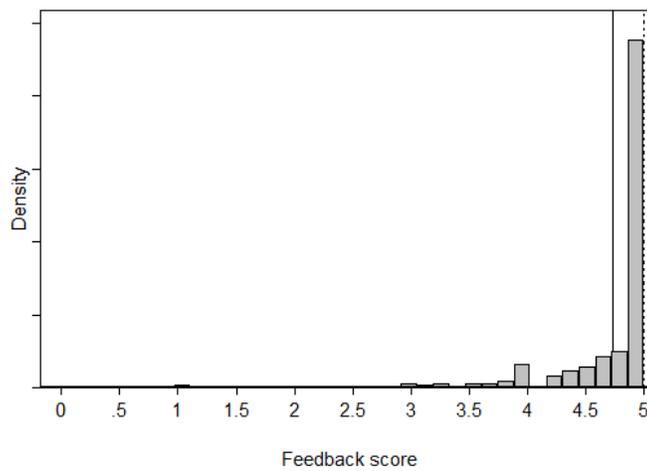
Source: Author calculations based on Nubelo data.

Nubelo supports outsourcing in a broad range of job categories. Four categories account for the vast majority of transactions: 1) software development, 2) graphic design and multimedia, 3) writing and translation, and 4) IT services. Demand is thus concentrated in relatively high-skill jobs, particularly when compared to micro-task platforms such as Mechanical Turk, where lower-skill tasks (such as image identification and data entry) are most common. As expected, the market is tighter in the job categories that require more technical skills, such as software development and IT services. By contrast, competition (as per bids to projects ratio) is particularly intense for contracts in multimedia and graphic design. The intensity of competition and contract prices vary widely across but also within job categories.

Employers find two types of information in a job applicant online profile. First, information about previous jobs obtained through the platform, along with the average feedback score

received by the worker in these contracts. This information is generated automatically by the platform, and therefore cannot be manipulated by either party. Figure 2 shows the distribution of feedback scores, which is highly skewed towards the maximum of 5 ( $\bar{x}=4.73$ ,  $SD=0.57$ ). This distribution is consistent with previous studies which find feedback scores in online marketplaces to be highly inflated.<sup>2</sup>

FIGURE 2  
DISTRIBUTION OF WORKERS' FEEDBACK SCORES (5-POINT SCALE)



*Source:* Author calculations based on Nubelo data. Solid line represents the mean value, while dotted line represents the median value.

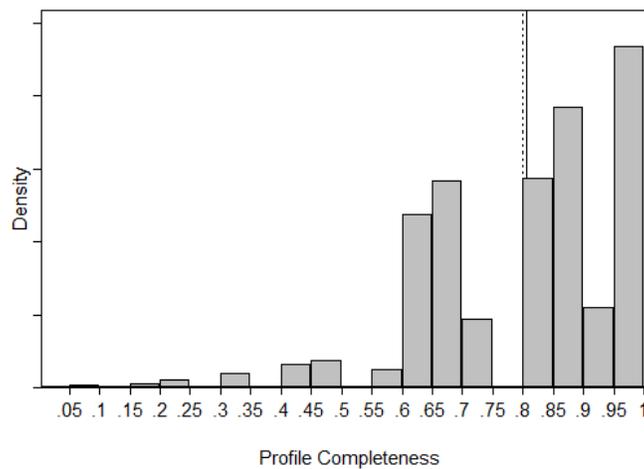
It is important to recall that feedback scores are conditional on having obtained at least one job contract in Nubelo. Most active workers (i.e., those who have submitted at least one bid during the study period) have never won a contract. At the same time, a small number of successful workers concentrate much of the job volume. This results in a superstar-type

<sup>2</sup> For example, Pallais (2014) finds that 83% of data entry workers in oDesk received a rating of at least 4, while 64% received a maximum rating of 5. Similarly, Stanton and Thomas (2014) find that about 60% of workers in oDesk received a feedback score of 5 in their first job.

distribution which is self-reinforcing, given that, as shown below, both work experience and feedback scores are significant predictors of hiring.

The second type of information available to employers is information voluntarily disclosed by workers. Nubelo encourages job seekers to complete an online profile with details about previous work experience, skills and training, a sample portfolio, and a personal picture. The platform computes the degree to which workers have completed their online profiles by assigning a certain percentage value to different data categories.<sup>3</sup> As shown in Figure 3, the average worker profile is 80 percent complete. We use this threshold to examine the effect of voluntarily disclosing more (or less) information in our probability models in Section 4.

FIGURE 3  
DISTRIBUTION OF ONLINE PROFILE COMPLETENESS

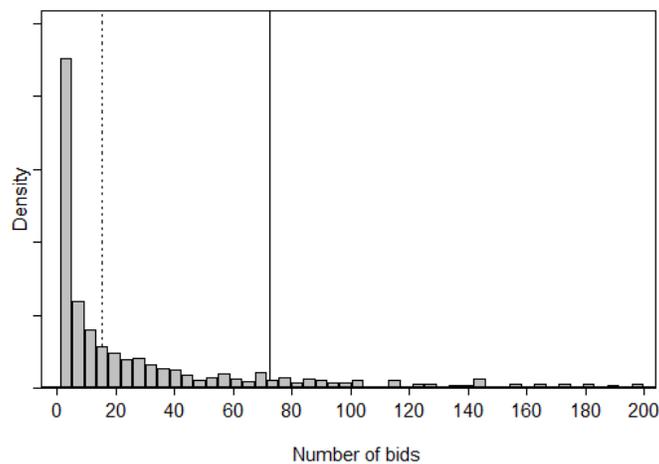


*Source:* Author calculations based on Nubelo data. Solid line represents the mean value, while dotted line represents the median value.

<sup>3</sup> For example, a personal picture adds 10 percent to the completeness of the profile, a description of previous work experience adds 5 percent, a description of skills adds 10 percent, and so forth.

Figure 4 shows the distribution of worker activity in Nubelo, measured by the number of bids submitted per job seeker during our study period. While the sample average is 72 bids, the distribution is highly skewed to the left, with a median value of just 15 bids per worker.

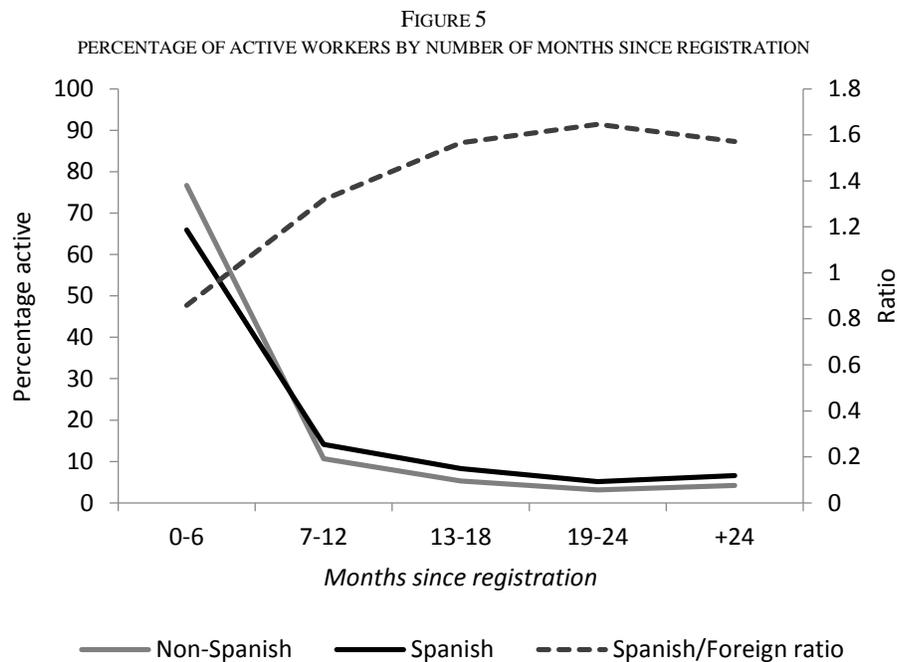
FIGURE 4  
DISTRIBUTION OF BID ACTIVITY (TRUNCATED AT 200 BIDS)



*Source:* Author calculations based on Nubelo data. Solid line represents the mean value, while dotted line represents the median value.

Finally, the evidence reveals high attrition rates, with most job seekers dropping out (or remaining inactive) within the first three months of joining the platform. In fact, about 40% drop out within the first month. This is partly to be expected given the small probability of obtaining a contract for workers without validated work experience and feedback (as discussed in Section 4). However, there are significant differences in attrition rates for Spanish and non-Spanish workers.

Workers are most active during their first 3 months after joining Nubelo, as shown in Figure 5. After 3 months, activity levels drop sharply for both groups. However, Spanish workers are more likely to remain active beyond the first trimester, with the domestic to foreign ratio stabilizing at about 1.5:1. We attribute this difference in attrition rates to the cumulative effect of hiring and wage penalties against foreign workers, as described in the next section.



Source: Author calculations based on Nubelo data.

#### 4. Method and results

Descriptive results suggest a hiring bias in favor of Spanish workers, particularly among Spanish employers. In order to formally test this proposition, we build a linear model that

estimates the probability of a worker being hired, conditional on nationality and covariates that capture bid amount, bid timing (in hours after the job is posted), other worker characteristics, and country reputation. The vector of worker characteristics includes the number of previous jobs in the platform, a dummy variable for having completed the online profile at or above the sample average of 80 percent, a dummy variable for having positive feedback from previous jobs (i.e., 4 points or more on a 5-point scale), and a dummy variable that indicates whether the job seeker has previously worked with the employer.

Country reputation is measured by the number of times the employer has previously contracted a worker from the same country as the job applicant. Given high feedback scores, we hypothesize that the more previous hires from a certain country, the more likely the employer will be to hire a worker from that country. Finally, given the variance in the intensity of competition and the value of contracts across jobs, the model includes a jobs fixed-effects term which captures both observed and unobservable differences across jobs.<sup>4</sup>

We restrict the sample to job postings from Spanish employers, for a number of reasons. First, as noted, contracts originating in Spain represent the large majority of job opportunities in Nubelo. Second, our main interest lies in labor offshoring from high-income to lower-income countries. With a GNI per capita of \$28,520 in 2015 (in current US Dollars), Spain's average income is about twice that of Argentina, the second largest employer. Further, during our study period Argentine employers were prevented from hiring outside Argentina (due to

<sup>4</sup> More formally, the estimated model is:  $\text{Hiring}_{ij} = \alpha_{ij} + \gamma \text{Foreign}_{ij} + \delta \log \text{Price}_{ij} + \lambda \log \text{Delay}_{ij} + \eta \text{CountryRep}_{ij} + \beta Z + \sigma_j + \varepsilon_{ij}$  where Hiring is the probability of worker's bid  $i$  being selected for job posting  $j$ , Foreign is a dummy (yes=1) that identifies non-Spanish workers, Price denotes bid amount (in log), Delay is the difference (in hours) between the job posting and the bid submission (in log), CountryRep denotes whether the employer has previously hired from the same country of the worker submitting bid  $i$  at the time of job posting  $j$ ,  $Z$  is a vector of worker characteristics that vary over time,  $\sigma$  controls for job fixed effects, and  $\varepsilon$  is an error term.

government-imposed limits to international payments), which eliminates variance in our main variable of interest. Mexico, the third largest employer with about 8 percent of jobs posted, has a GNI per capita of less than a third of Spain's.

In addition, given our interest in comparing outcomes for workers depending on country of residence, we further restrict the sample to job postings that received at least one bid from a Spanish job applicant and one from a non-Spanish (i.e., foreign) applicant. Finally, filtering for job postings that did not result in a positive match (i.e., where the employer did not hire), our restricted sample comprises 46,799 bids for 2,500 jobs.

#### *4.1. Hiring penalty*

Table 2 corroborates that, after controlling for bid amount and delay, country reputation, previous contracts between applicant and employer, and observable worker quality characteristics, non-Spanish job applicants are less likely to be hired by Spanish employers. Following the full model in column 7, foreignness reduces the winning odds by 2.2 percentage points. Relative to the average winning odds of 5.3 percent in the full sample, this represents a hiring penalty of about 42 percent.

TABLE 2  
HIRING PROBABILITIES (OLS WITH FIXED EFFECTS)

Dependent Variable: Hiring Probability							
	1		3	4	5	6	7
Foreign worker	-0.030*** [0.002]	-0.030*** [0.002]	-0.029*** [0.002]	-0.029*** [0.002]	-0.028*** [0.002]	-0.022*** [0.002]	-0.022*** [0.002]
<b>Controls:</b>							
Bid Amount	✓	✓	✓	✓	✓	✓	✓
Bid Delay		✓	✓	✓	✓	✓	✓
Country Reputation			✓	✓	✓	✓	✓
Work Experience				✓	✓	✓	✓
Profile					✓	✓	✓
Feedback						✓	✓
Worked w/Employer							✓
Constant	0.152*** [0.005]	0.168*** [0.005]	0.167*** [0.005]	0.148*** [0.005]	0.128*** [0.006]	0.0789*** [0.006]	0.0743*** [0.005]
N	46,799	46,799	46,799	46,799	46,799	46,799	46,799
R <sup>2</sup>	0.011	0.013	0.013	0.022	0.025	0.050	0.161
Jobs	2,500	2,500	2,500	2,500	2,500	2,500	2,500
Mean DV	0.053	0.053	0.053	0.053	0.053	0.053	0.053

Standard errors in brackets.

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Source: Author calculations based on Nubelo data.

It is interesting to note how the hiring penalty for non-Spanish applicants changes as different covariates are introduced. Model 1 represents the base estimate which only includes bid amount and nationality. In this model, the magnitude of the effect for *Foreign* is about 3 percentage points, which represents a 58 percent penalty relative to the sample mean. In models 2 to 7 the control variables are sequentially introduced. The penalty remains essentially unchanged until model 6, when individual reputation (i.e., feedback from previous jobs) is introduced. This strongly suggests that Spanish employers attribute quality to workers based on nationality in the absence of credible signals for individual workers. Once this information is available the magnitude of the foreign worker penalty drops by about a third.

#### 4.2. *Wage premium*

Descriptive statistics also suggest that Spanish employers are willing to pay a wage premium for hiring domestically. In order to quantify this wage premium, we build a linear model that estimates bid amount (in log) conditional on nationality and a vector of worker characteristics.<sup>5</sup> We then restrict the sample to projects that resulted in a Spanish applicant being hired. Hence, our coefficient of interest ( $\gamma$ ) indicates the marginal change in the amount of bids submitted by foreign workers to the price of contracts ultimately obtained by Spanish workers. In other words, it quantifies the premium that the employer was willing to pay to hire domestically, relative to alternative bids by similarly-qualified (per observable characteristics) foreign workers.

The results in Table 3 indicate that, when hiring locally, Spanish employers rejected alternative bids by non-Spanish job seekers that were, on average, 14 percent lower (model 6). This translates into a wage premium for Spanish workers of about 16 percent when calculated as a premium over alternative bids.

<sup>5</sup> More formally, the estimated model is:  $\log(\text{Price})_{ij} = \alpha_{ij} + \gamma \text{Foreign}_{ij} + \lambda \log \text{Delay}_{ij} + \beta Z + \sigma_j + \varepsilon_{ij}$  where Price is the bid by worker  $i$  for job posting  $j$ , Foreign is a dummy (yes=1) that identifies non-Spanish workers, Delay is the difference (in hours) between the job posting and the bid submission (in log),  $Z$  is a vector of worker characteristics that vary over time,  $\sigma$  controls for job fixed effects, and  $\varepsilon$  is an error term.

TABLE 3  
WAGE PREMIUM (OLS WITH FIXED EFFECTS)

Dependent Variable: Bid Amount (log)						
	1	2	3	4	5	6
Foreign worker	-0.120*** [0.012]	-0.122*** [0.012]	-0.124*** [0.012]	-0.125*** [0.012]	-0.140*** [0.012]	-0.141*** [0.012]
<u>Controls:</u>						
Bid Delay		✓	✓	✓	✓	✓
Work Experience			✓	✓	✓	✓
Profile				✓	✓	✓
Feedback					✓	✓
Worked w/Employer						✓
Constant	4.738*** [0.008]	4.615*** [0.014]	4.644*** [0.015]	4.660*** [0.018]	4.746*** [0.019]	4.747*** [0.019]
N	31,516	31,516	31,516	31,516	31,516	31,516
R <sup>2</sup>	0.003	0.007	0.008	0.008	0.014	0.014
Number of Jobs	1,626	1,626	1,626	1,626	1,626	1,626
Mean DV	281	281	281	281	281	281

Standard errors in brackets.

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Source: Author calculations based on Nubelo data.

Next we examine if the wage premium varies when more information about job applicants is available (Table 4). Our hypothesis is that the wage premium will be higher among job applicants without previous contracts, with lower-than-average feedback, and with less information on their online profiles. In order to test these hypotheses, we replicate the full model (model 6 in Table 3) for different subsamples of job postings.

In model 1, the sample is restricted to jobs that only received bids from workers without previous job experience. These results are compared to model 2, in which the sample is restricted to jobs that only received bids from workers with previous contracts in Nubelo. As in Table 3, we restrict our sample to jobs for which a Spanish worker was hired, but which also received (unsuccessful) bids from non-Spanish applicants. As expected, the wage penalty is larger among job seekers without previous contracts (model 1) with respect to job seekers

with validated work experience (model 2). In other words, the less validated information about individual workers is available, the higher the premium employers are willing to pay to hire domestically.

Models 3 and 4 are based on a similar exercise. In this case the comparison is between workers with below-average feedback scores (model 3) against workers with above-average scores (model 4). The results indicate that having poor feedback from previous contracts disproportionately affects foreign job applicants. As shown, the wage penalty in model 3 is about three times larger with respect to the sample of workers with positive feedback in model 4. It is likely that Spanish employers interpret poor feedback as a confirmation of their prior beliefs about lower average quality among foreign workers.

Finally, we replicate this exercise for jobs that only received bids from workers with below-average information in their online profile (model 5) against the case of jobs where all bidders had average or above-average information in their profiles (model 6). As shown, the foreign wage penalty is almost identical in both cases. This suggests that employers discount this type of information regardless of worker nationality.

TABLE 4  
WAGE PREMIUM (OLS WITH FIXED EFFECTS)

Dependent Variable: Log of Bid Amount						
	Previous contracts		Feedback score		Profile completeness	
	No (1)	Yes (2)	< 4 (3)	4+ (4)	< 80% (5)	80% + (6)
Foreign Penalty	-0.119*** [0.029]	-0.079*** [0.020]	-0.164** [0.066]	-0.052*** [0.017]	-0.122** [0.048]	-0.135*** [0.016]
<u>Controls:</u>						
Bid Delay	✓	✓	✓	✓	✓	✓
Profile	✓	✓	✓	✓		
Feedback	✓	✓			✓	✓
Work Experience			✓	✓	✓	✓
Worked w/employer	✓	✓	✓	✓	✓	✓
Constant	4.916*** [0.041]	4.421*** [0.049]	4.856*** [0.093]	4.482*** [0.026]	4.768*** [0.064]	4.737*** [0.022]
Observations	7,463	8,831	1,790	13,089	2,877	17,124
R <sup>2</sup>	0.007	0.008	0.006	0.004	0.010	0.017
Number of jobs	518	1,108	144	1,482	321	1,306
Mean of DV	377.5	214.0	527.3	224.6	392.1	236.6

Standard errors in brackets

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

Source: Author calculations based on Nubelo data.

## 5. Conclusions

Online platforms offer a valuable laboratory to examine how discrimination operates in a variety of social realms. In some cases, the research interest lies in understanding how the algorithms that determine which information is presented to which platform participants affects individual behavior and beliefs (Sandvig et al., 2014). This study is situated in a related line of work that examines how biases emerge from choices by participants in multi-sided platforms, and how alternative market design choices promote or mitigate discriminatory behavior. Examples include scholarship about bias in short-term property

rentals (Edelman and Luca, forthcoming), in peer-to-peer transportation (Ge et al., 2016) and in crowdfunding platforms (Pope and Sydnor, 2011).

The focus of this study is geographical discrimination in a contract labor platform that matches employers with job seekers. We find that discrimination stems from information frictions that trigger cognitive shortcuts among employers. These cognitive shortcuts are, by definition, stereotypes that associate country of origin with expected worker productivity. We corroborate this mechanism by showing that increasing the availability of validated information about individual workers tends to deactivate geographical stereotypes, reducing the hiring and wage penalties faced by foreign job applicants.

The findings have a number of theoretical implications. The first is that the reduction in search and communication costs brought about by online platforms does not result in perfectly competitive labor markets. Rather, competition continues to be highly imperfect, with outcomes shaped by the geographical proximity between employers and job applicants. In the online environment, however, geography operates less in terms of physical distance (as it does in the spatial mismatch hypothesis) and more as a signaling mechanism that orients hiring choices. This is why the “flat world” metaphor is inadequate to describe the dynamics of online labor. The metaphor correctly describes recent changes in the underlying information infrastructure of labor markets, but mistakenly identifies the source of frictions that determines employment and wage outcomes.

A flattened labor market assumes impersonal exchanges in which race, gender, nationality and other personal characteristics become irrelevant. The results of this study indicate that

this is not how digital labor platforms operate. In fact, they suggest that stereotypes may play an even larger role in determining hiring and wage outcomes, particularly when verifiable information about individual job applicants is limited and employers lack other mechanisms for screening workers.

Moreover, research has shown that most tasks cannot be easily routinized or codified, and the ones that can (such as image recognition and data entry) are increasingly being automated (Kokkodis et al, 2015). The limits to the commodification of work suggest that exchanges in online labor markets will continue to depend on human relationships and, as a result, be affected by communication costs and information frictions. In other words, the location and identities of the transacting parties will continue to matter.

A number of platform design and policy implications emerge from our findings. First, platform operators may discourage the display of information unrelated to productivity on worker's profiles, while at the same time implement mechanisms to validate skills and previous work experience. This would not only favor workers in developing countries (who, as shown, are penalized disproportionately when lacking verified experience) but also improve employer-worker matching. Horton (forthcoming) estimates that about half of the job contracts posted in digital labor platforms are never filled.

Second, platforms can further develop mechanisms that help employers find and screen job applicants. As mentioned, Nubelo discouraged employer-applicant interaction before hiring, yet other platforms in fact promote personal interviews and direct contract negotiation between parties. Following our results, these mechanisms are likely to deactivate stereotypes

in hiring choices, thus favoring foreign job seekers, and more generally promoting diversity in job categories currently associated with specific demographic groups.

Finally, employers can be nudged to hire more diversely by altering the order in which alternative candidates are presented. For example, Nubelo attempted to lower the first-job barrier (and thus reduce worker attrition) by favoring job seekers without previous contracts in the algorithm that determined the display of potential matches to employers. A similar nudge could be applied to favor foreign job seekers, or candidates from specific countries underrepresented in certain job categories. In turn, individual hiring would help build country reputation, which as this and other studies show is an important predictor of hiring (see Leung, 2012).

Finally, the question of governance for online labor contracts must be addressed to protect workers from the vulnerabilities associated with digital work. These vulnerabilities are multifaceted, ranging from lack of enforcement of minimum-wage legislation to the boundaries of employment relationships (De Stefano, 2016). This and other questions related to the protection of workers' rights only become more pressing when hiring takes place across borders. In particular, clear jurisdictional lines must be established to enable the enforcement of existing anti-discrimination laws in hiring and compensation in the context of online labor contracting.

## FUNDING

This work was supported by the International Development Research Centre (IDRC-Canada), Project #107601-001.

## REFERENCES

- Agrawal A, Horton J, Lacetera N and Lyons J (2015) Digitization and the contract labor market: A research agenda. In: Goldfarb A, Greenstein S and Tucker C (eds) *Economic analysis of the digital economy*. Chicago: University of Chicago Press, pp. 219-256.
- Agrawal, A, Lacetera N and Lyons J (2013). *Does Information Help or Hinder Job Applicants from Less Developed Countries in Online Markets?* NBER Working Paper 18720.
- Aigner DJ and Cain GC (1977) Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review* 30: 175-187.
- Ashenfelter O and Card D (1999) *Handbook of labor economics*. Amsterdam: Elsevier.
- Autor D (2001) Wiring the labor market. *Journal of Economic Perspectives* 15: 25-40.
- Blinder A (2006) Offshoring: The next industrial revolution? *Foreign Affairs* 85(2): 113-128.
- De Stefano V (2016) Introduction: Crowdsourcing, the gig-economy and the law. *Comparative Labor Law & Policy Journal* 37(3): 471-503.
- Edelman B, Luca M and Svirsky D (forthcoming) Racial discrimination in the sharing economy: Evidence from a field experiment. *American Economic Journal: Applied Economics*.
- Farrell D, Laboissière M, Pascal R, de Segundo C, Rosenfeld J, Stürze S and Umezawa F (2005) *The emerging global labor market*. McKinsey Global Institute, June.
- Fernandez, RM and Su C (2004) Space in the study of labor markets. *Annual Review of Sociology* 30: 545-569.
- Fish A and Srinivasan R (2011) Digital labor is the new killer app. *New Media & Society* 14(1): 137-152.

- Friedman T (2005) *The world is flat: A brief history of the twenty-first century*. New York: Farrar, Straus and Giroux.
- Ge Y, Knittel CR, MacKenzie D and Zoepf S (2016) *Racial and Gender Discrimination in Transportation Network Companies*. NBER Working Paper No. 22776.
- Gefen D and Carmel E (2008) Is the world really flat? A look at offshoring in an online programming marketplace. *MIS Quarterly* 32(2): 367-384.
- Ghani E, Kerr WR and Stanton C (2014) Diasporas and outsourcing: Evidence from oDesk and India. *Management Science* 60(7): 1677–1697.
- Goldin C and Rouse C (2000) Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review* 90(4): 715-741.
- Hong Y and Pavlou P (2014) *Is the world truly 'flat'? Empirical evidence from online labor Markets*. Available at SSRN: [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2371748](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2371748).
- Horton JJ (forthcoming) The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics*.
- Horton JJ and Chilton L (2010) *The labor economics of paid crowdsourcing*. In: Proceedings of the 11th ACM conference on electronic commerce, Boston, 7-11 June 2010, pp. 209-218.
- Huws U (2015) When Adam blogs: Cultural work and the gender division of labor in Utopia. *Sociological Review* 63(1): 158-173.
- Irani L (2013) The cultural work of microwork. *New Media & Society* 17(5): 720-739.
- Kain, JF (1992) The spatial mismatch hypothesis: three decades later. *Housing Policy Debate* 2: 371-460.
- Kingsley SC, Gray ML and Suri S (2015) Accounting for Market Frictions and Power Asymmetries in Online Labor Markets. *Policy and Internet* 7(4): 383–400.

- Kokkodis M, Papadimitriou P, and Ipeirotis P (2015) *Hiring behavior models for online labor markets*. Proceedings of the 8<sup>th</sup> ACM International Conference on Web Search and Data Mining, February 2-6, Shanghai, China.
- Lehdonvirta V and Ernkvist M (2011). *Knowledge Map of the Virtual Economy*. Washington DC: World Bank.
- Lehdonvirta V, Barnard H, Graham M, and Hjorth I (2014) *Online labour markets – leveling the playing field for international service markets?* In: Internet, Politics, and Policy conference, Oxford, UK, 25-26 September 2014.
- Leamer EE (2007) A flat world, a level playing field, a small world after all, or none of the above? A review of Thomas L. Friedman's *The World is Flat*. *Journal of Economic Literature* 45: 83-126.
- Leung M (2012) *Job categories and geographic identity: A category stereotype explanation for occupational agglomeration*. UC Berkeley Institute for Research on Labor and Employment Working Paper.
- Mann L and Graham M (2016) The domestic turn: Business process outsourcing and the growing automation of Kenyan organizations. *The Journal of Development Studies* 52(4): 530-548.
- Mill R (2011) *Hiring and learning in online global labor markets*. Available at SSRN: [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1957962](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1957962).
- Neckerman, KM and Kirschenman J (1991) Hiring strategies, racial bias, and innercity workers: an investigation of employers' hiring decisions. *Social Problems* 38: 801-15.
- Newman KS (1999) *No shame in my game: the working poor in the inner city*. New York: Knopf.

- Pallais A (2014) Inefficient hiring in entry-level labor markets. *American Economic Review* 104: 3565-3599.
- Phelps E (1972) The statistical theory of racism and sexism. *The American Economic Review* 62(4): 659-661.
- Pope DG and Sydnor JR (2011) What's in a Picture? Evidence of Discrimination from Prosper.com. *Journal of Human Resources* 46(1): 53-92.
- Raja S, Imaizumi S, Kelly T and Paradi-Guilford C (2013). *Connecting to work. How information and communication technologies could help expand employment opportunities*. Washington DC: World Bank.
- Rossotto C, Kuek SC and Paradi-Guilford C. (2012). *New frontiers and opportunities in work*. Washington DC: World Bank.
- Sandvig C, Hamilton K, Karahalios K., and Langbort C 2014. *Auditing algorithms: Research methods for detecting discrimination on Internet platforms*. In: Data and Discrimination pre-conference of the 64th annual meeting of the International Communication Association, Seattle, WA, 22 May 2014.
- Stanton C and Thomas C (2014) *Landing the first job: The value of intermediaries in online hiring*. Available at SSRN: <http://ssrn.com/abstract=1862109>.