NET Institute*

www.NETinst.org

Working Paper #11-17

October 2011

Hiring and Learning in Online Global Labor Markets

Roy Mill
Stanford University

* The Networks, Electronic Commerce, and Telecommunications ("NET") Institute, http://www.NETinst.org, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, "virtual networks" comprised of computers that share the same technical standard or operating system, and on network issues in general.

Hiring and Learning in Online Global Labor Markets

Roy Mill*

September 30, 2011

Abstract

This paper uses data from freelancer.com – an online platform that allows employers and freelancers to search for, and match with, each other – to study the effect of freelancers' country of origin on their likelihood to be hired. Having to rely on a relatively small number of characteristics, employers use the freelancer's country of origin to infer the expected service's quality. This setting also allows me to document how employers' experience in past hires affects their behavior in current hires. I find that freelancers from developing countries are less likely to be hired when they have no individual reputation, and as individual reputation becomes better this country effect disappears. I show that following a good match with a freelancer, employers are more likely to hire freelancers from the good match's country. These these findings are consistent with statistical – rather than purely taste-based – discrimination.

Keywords: International outsourcing; Online labor market; Information acquisition; Quality reputations; Country-of-origin effect; Statistical discrimination

JEL Classifications: D83, F15, F23, J23, J71, L24, O15

*Department of Economics, Stanford University. I would like to thank Ran Abramitzky for invaluable guidance. I would also like to thank freelancer.com for granting access to the data and providing helpful information on the platform. Financial support from the NET Institute (www.NETinst.org) 2011 Summer Research Grant as well as the Shultz Graduate Student Fellowship in Economic Policy through a grant to the Stanford Institute for Economic and Policy Research (SIEPR) is gratefully acknowledged.

1 Introduction

I study how employers and freelancers from all over the world interact in an online website that facilitates searching and matching. The platform I study – freelancer.com – allows individuals to take advantage of the large wage gaps across countries, generating a single global labor market defined by occupation and free of location considerations.

Traditionally, the search for prospective employees was geographically concentrated near the employer. The rise of the Internet allowed online marketplaces such as freelancer.com to evolve and facilitate search for employees anywhere on the globe. Employers in such online platforms inevitably face the problem of inferring prospective employees' (freelancers' henceforth) productivity without meeting them in person. One salient feature from which employers infer the expected quality is the freelancers' country. If employers associate freelancers from particular countries with higher expected productivity, they will be more likely to choose freelancers from such countries. Therefore, while a programmer in India may earn only one-tenth of what a programmer in the U.S does, American employers may still be reluctant to hire freelancers from a developing country.

I was granted access to a copy of the database that is used for the operation of freelancer.com. The database contains both public and non-public data recorded from user activity. Specifically, it contains information on employers, freelancers, projects, bids, reviews, etc.[1]

This paper is divided into two parts. Since this is among the first papers on online labor markets, the first part of the paper describes some of the main patterns that characterize user activity on it. I start by showing that the most popular types of jobs are those that are easily tradable across the internet and do not require major involvement by the employer. Then I show that services flow mainly from developing countries to developed countries, but that developed countries are still home to a significant share of freelancing. Next, I show that some markets are

[1] The copy I got was did not contain sensitive individual details such as contact and payment information, which are also unnecessary for this research.

thicker than others in terms of the number of bids per project. Thick markets are relatively easier for freelancers to enter than thin markets. Lastly, employers do not view freelancers as perfect substitutes of labor services: the winning quote is higher than the lowest quote in 40-60% of the projects. This pattern is especially noticeable in types of jobs that are more sensitive to quality.

In the second part of the paper, I analyze how employers use the freelancers' country as a signal for their quality. One could view employers as importers who buy services from exporting freelancers. For international trade in goods, many studies show how consumers frequently ascribe a higher quality to products imported from some countries relative to others and are willing to pay more for them.[2] In the context of labor markets, this relates to the large literature on discrimination where, instead of gender or race, country of origin defines groups employers discriminate between.

I document differential hiring of freelancers based on their country of origin in two ways. First, I look at how country and personal reputation interact and show that employers put more weight on the country of origin when no individual reputation is available. As an example, American employers are 17 percentage points less likely to hire a Bangladeshi freelancer than a Canadian freelancer. However, the gap disappears for Bangladeshi and Canadian freelancers with 100 good reviews.

Second, I delve into how experience in the market changes employers' perception of country quality. I follow employers and record the number of good and bad experiences they had with freelancers from each country prior to the current job they seek to fill. I show that employers are more likely to hire freelancers from countries they have better experience with. Additionally, as employers gain more experience, the relationship between the marginal experience and future projects' hiring decisions becomes weaker, since each new observation on a specific freelancer's quality adds relatively less information on the whole distribution of quality in their country.

[2]For a relatively recent review of the empirical literature on country-of-origin effects on demand, see Dinnie (2004).

Specifically, the first good experience that an employer had with a certain country increases the likelihood of another freelancer from the same country to be hired in a future project by this same employer by 0.2 percentage points. After 10 good experiences, though, the next good experience increases the probability by only 0.1 percentage point.

All these findings are consistent with statistical – rather than purely taste-based – discrimination. Statistical discrimination is related to how employers predict quality of service from freelancers' countries (see Phelps (1972); Arrow (1973)) while taste-based discrimination identifies employers preferences as the source of differential treatment (see Becker (1957)).

Previous empirical work has documented discrimination and tried to identify the type of discrimination. Bertrand and Mullainathan (2004) is an audit study that uses the interaction of other signals with race to see if they change the relative importance of race, Altonji and Pierret (2001) uses wage dynamics from longitudinal data, and List (2004) uses a series of experiments complementing each other. However, no study followed employers' repeated hiring decisions following their experience with previous hires. This paper contributes to the literature of statistical discrimination by showing the relationship between specific paths of experience in the market and subsequent hiring behavior.

Three characteristics of this unique dataset allow me to document how behavior dynamically evolves according to an employer's experience: First, complete project history allows me to follow employers from the first project until the most recent one. Second, project outcome is partially observable allowing me to classify past projects as successes or failures. Finally, many employers receive bids from freelancers hailing from countries the employers probably never had contact with. Thus, every project's outcome should have a heavier weight in the determination of beliefs over productivity.

Even though online labor markets are relatively new markets, they have already been studied in several contexts related to job search and matching. Pallais (2010) conducts an experiment on oDesk, a competing platform, where she hires freelancers randomly and measures the contribution of freelancer experience (and reputation) to the propensity to be hired later. Stanton and Thomas (2011) examine how freelancers can increase their likelihood to be hired if they join a network of freelancers that has a reputation based on the network members' experience. In the context of online international trade in goods, Hu and Wang (2010) show how consumers are willing to pay different amounts for identical products sold by retailers from different countries on eBay.

The implications of my findings to international trade in services and growth of developing countries is important to note. Wage gaps between rich and poor countries persist partly because labor cannot freely move between them due to immigration restrictions imposed by rich countries.[3] Online platforms allow developing countries to export labor *services*, but the ability to penetrate foreign markets depends on the perception of the quality of these services in the importing economies. Chisik (2003) shows how initial differences in country-level reputation can become self-fulfilling through product specialization, due to an endogenously created reputational comparative advantage that can eventually determine the whole industrial structure of a country.

The paper is constructed as follows: Section 2 constitutes the first part that describes main patterns of trade on freelancer.com. Section 3 is the second part of the analysis. It describes the empirical strategy, dataset construction and reports preliminary results. Finally, section 4 concludes.

2 Employers and Freelancers Matching through Freelancer.com

Trade on freelancer.com is organized around projects. A project is posted to the website by an employer looking for freelancers. Freelancers then bid on the project, quoting price for the service and specifying an estimate of the number of

[3]See Clemens (2011) for a survey of several studies on the gains of lifting international migration barriers.

days to completion. Projects on freelancer.com are almost always traded in full: the whole service needed is described upfront and bids are made for the completion of the project, rather than for an hourly or weekly wage.[4] The website allows employers to pay freelancers through an escrow service, thereby guaranteeing a higher level of security.[5] Figures 7 and 8 in Appendix A show a list of recent projects and a list of bids on a project randomly selected.

This section provides descriptive statistics of overall activity on the website. I break down activity from different angles, to describe the setting in which employers and freelancers operate and to show in what ways this platform is home to multiple global labor markets. Overall, since the launch of the website in 2004 – 945,688 projects were posted and matches accounted for the transfers of more than $60 million dollars in exchange for services provided. [6]

2.1 What Kinds of Services are Traded?

To reduce search costs, projects are tagged with job-types, or skills. The website offers 385 job types grouped into nine categories. A project can be classified with up to five job-types. Figure 1 shows the evolution of the nine categories over time since the inception of the platform and the stock of projects at the time the database was copied. Table 1 lists the top five job-types within each of the top five categories.

The distribution of projects by categories and jobs shows that the more popular types of services are indeed more tradable. The extent to which a task is tradable varies. Some tasks like childcare are simply non-tradable: they require the physical presence of the service provider near the seller. But even those that are tradable require different degrees of freelancer's tacit knowledge or employer involvement, and different abilities to assess freelancers' quality before hiring them.

[4]A job can be defined as "full time" and then bids can be made for hourly or weekly wages, but it is quite rare on the platform. More recently, contests were added as an additional form of outsourcing termed crowd-sourcing. However, no contests were available the time span for which I have the data.

[5]The fee for paying through the website is 10% of the total value of the transaction. Fee for full-time job is a fixed fee charged only in the first payment.

[6]Numbers are true as of February 14, 2011 at 4:00pm EST.

For example, data entry tasks can be simple and straightforward, which makes it relatively easy to screen suitable freelancers, and have them independently complete the task. Conversely, designing a new machine requires extensive knowledge in engineering that is also harder to assess in advance, and also requires more extensive communications between the service provider and the buyer.

Figure 1: The distribution of projects by category: flows and stocks.

(a) Number of new projects in a month, February 2004-January 2011.

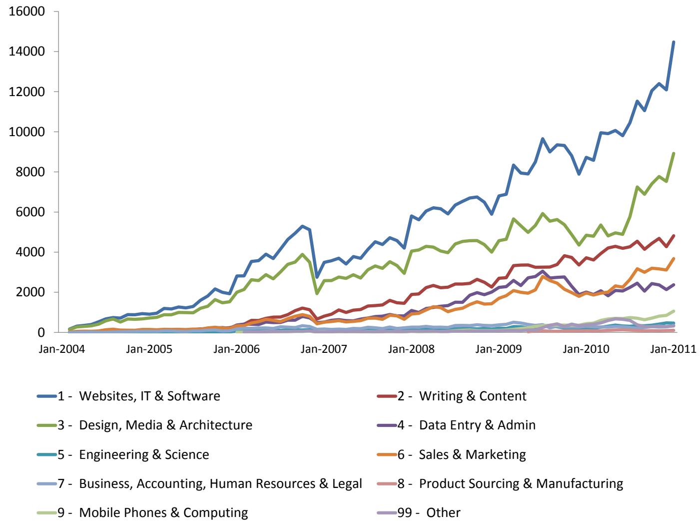

(b) Total projects within a category, February 14, 2011.

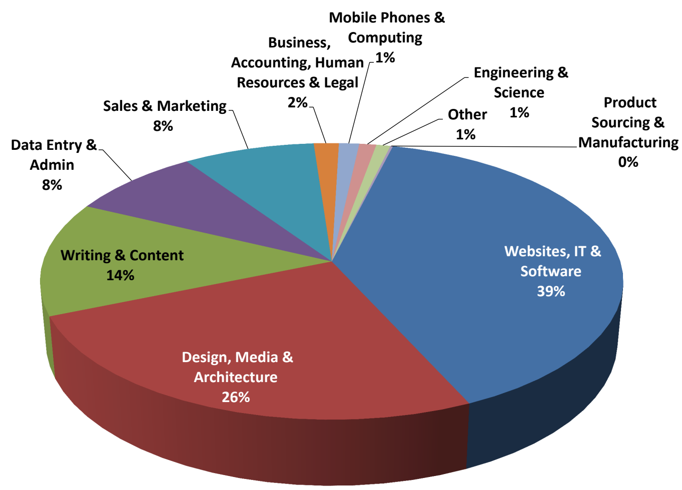

Note:
1. Since projects can be classified to more than one job type and job types can be in different categories, the same project may appear in more than one category. The numbers here are therefore given on the project-category level.
2. Data spans from the launch of the website through January 2011 inclusive for Figure 1a and through February 14, 2011 for 1b.

Table 1: Top 5 jobs in each of the top-5 categories

Websites, IT & Software		Writing & Content		Design, Media & Architecture		Data Entry & Admin		Sales & Marketing	
Job	Projects	Job	Projects	Job	Projects	Job	Projects	Job	Projects
PHP	198,864	Copywriting	81,515	Website Design	164,329	Data Entry	70,622	Internet Marketing	67,711
SEO	74,711	Articles	39,623	Graphic Design	98,738	Data Processing	42,573	Marketing	17,718
Javascript	41,161	Research	25,322	Flash	60,338	Excel	10,863	Advertising	13,015
Link Building	39,975	Proofreading	20,225	Logo Design	43,941	Virtual Assistant	4,854	Telemarketing	11,263
.NET	38,415	Article Rewriting	19,768	Photoshop	26,215	Web Search	3,861	Bulk Marketing	9,720

Indeed, the most popular job types – website construction, writing, design, data entry and bulk marketing – all fit relatively well into the story of posting a well-defined project, finding the most cost-beneficial provider and letting them complete the job wherever they are with relatively low involvement required from the service buyer.

Job types define labor markets. They are roughly defined along the lines of skills or occupations that freelancers possess. Moreover, freelancers are required to choose a set of job types they have the skills to work in.[7] Labor market thickness – the number of applicants per job – varies across job types. Table 2 lists the thickest and thinnest labor markets. Some markets exhibit a very high level of competition between freelancers. A project flagged with "Excel" has 42 freelancers bidding on average. By contrast, projects requiring iPhone related work get 10 freelancers bidding on average.

Table 2: Thickest and thinnest job types among markets with at least 3,000 projects

10 thickest job types			
Job	**Category**	**Projects**	**Bids / projects**
Excel	Data Entry & Admin	10,863	**41.88**
Web Search	Data Entry & Admin	3,861	**33.42**
Logo Design	Design, Media & Architecture	43,941	**28.70**
Data Entry	Data Entry & Admin	70,622	**28.59**
Photography	Design, Media & Architecture	4,839	**28.24**
Virtual Assistant	Data Entry & Admin	4,854	**27.31**
Illustrator	Design, Media & Architecture	6,737	**26.76**
Photoshop Design	Design, Media & Architecture	4,585	**26.71**
Illustration	Design, Media & Architecture	4,373	**25.78**
Photoshop	Design, Media & Architecture	26,215	**25.52**

10 thinnest job types			
Job	**Category**	**Projects**	**Bids / projects**
Android	Mobile Phones & Computing	3,131	**10.14**
iPhone	Mobile Phones & Computing	8,361	**9.99**
Twitter	Websites, IT & Software	4,089	**9.98**
Python	Websites, IT & Software	3,411	**9.54**
Software Architecture	Websites, IT & Software	6,695	**9.52**
Bulk Marketing	Sales & Marketing	9,720	**9.20**
Facebook	Websites, IT & Software	12,286	**9.09**
Mobile Phone	Mobile Phones & Computing	6,433	**8.94**
Sales	Sales & Marketing	9,595	**8.74**
Leads	Sales & Marketing	8,247	**8.35**

A very salient difference between thick and thin markets is the amount of investment required to enter the market. Most thin markets require either high and specific human capital (programming and mobile) or marketing infrastructure (e.g mailing lists and customer databases). In contrast, thick markets are mainly data entry and design services which most of them do not require great investments and are therefore more open for entry. Of course, thickness depends on demand factors as well – iPhone jobs are in higher demand than Web Search jobs – but from looking at the numbers of projects in both the thick and thin jobs, it does

[7]Technically this set of job types restrict freelancers to bid only on projects flagged with one of these job types. Additionally, they are also subscribed to emails announcing new jobs posted within each job type they assign themselves to.

not look like there's a systematic difference in demand.

2.2 Where are Services Demanded and Supplied?

Employers and freelancers from all over the world log on to the website. Upon signing up to the website they are required to fill in their country of residence and optionally additional geographic and other personal details.[8] Figure 2 plots the main city-level flows of services traded through the website. Aggregate at the country level, Figure 3 shows amounts earned and sent over the platform (surface on map), as well as total net exports (map color).

Figure 2: City-level flows of services

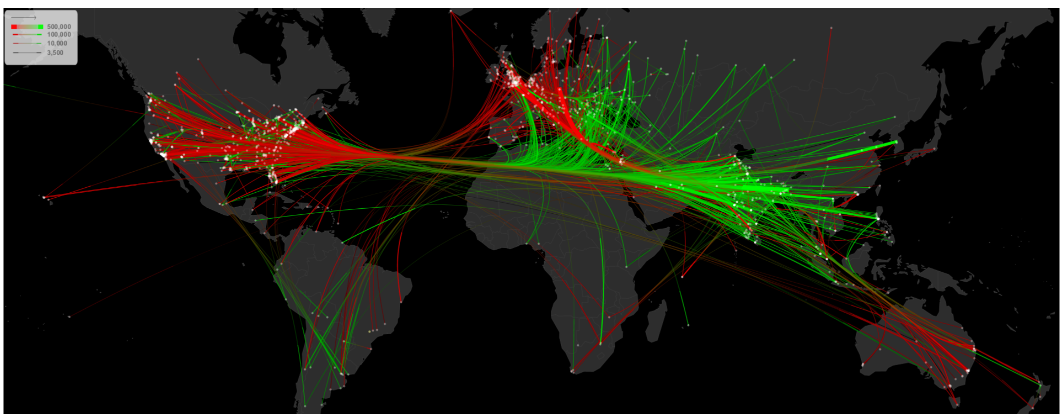

Note: the green part of every curved line stands for the origin of the service while the red part of the line stands for the destination. All city-pairs with overall project volume of at least $3500 are given a line in width proportional to the volume of trade.

[8]It is possible to report a country-of-residence falsely, but freelancer is undertaking several verification procedures to fight fraudulent abuse of the system, especially when a payment is involved.

Figure 3: Geographical distribution of net exports and money paid and earned

(a) Country size inflated by total USD ***paid*** (b) Country size inflated by total USD ***earned***

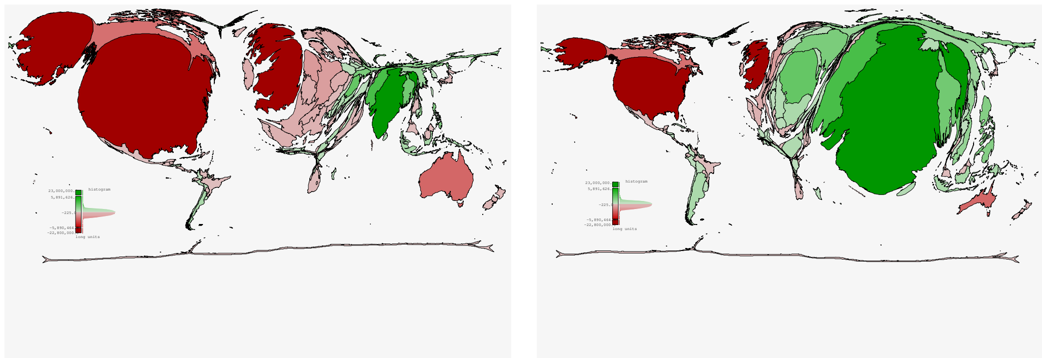

Note: green color indicates positive net exports while red color indicates negative net exports.

Looking at the maps, a few patterns are worth mentioning: First, services flow mainly from developing to developed countries. Net exporters on the website – countries in green – are mainly located in Asia, Eastern Europe, Africa and Latin America. Net importers are mainly USA, UK, Australia, New Zealand and Western Europe. Second, countries with a larger proportion of English speakers are more active, both in terms of posting projects – UK, USA, Canada, Australia and New Zealand inflated disproportionately more than other non-English speaking OECD members – and in terms of being selected to do the job – India, Egypt and the Philippines disproportionately more active than comparable countries such as China, Morocco and Thailand. Lastly, we see that not all projects are awarded to freelancers from developing countries: some developed countries, mainly the United States and Britain, are also significant *suppliers* of services (Figure 3b). Similarly on the other side of the market, employers from India, Pakistan and Bangladesh post a non-negligible number of project as well.

2.3 Quotes, Prices and Wages

Employers purchasing services would like to minimize the cost of the service, but are aware of differences in quality among providers. Thus, employers do not always select the bid with the lowest quote. In fact, the last column in Table 3 shows that only between 40 and 60 percent of the winning bids quoted the lowest price among all bids for the project. The rest of the bids won despite having another bid with a lower quote. Additionally, Table 3 reports average quote ranges for projects in selected job types.

Projects can be quite different in scope and magnitude, even within job types. However, one can see that tasks in the Mobile Phones & Computing category have relatively high quote ranges. Compared to Virtual Assistant jobs that start at almost \$70 and winning bids quote \$105 per project on average, iPhone and Android jobs start at around \$400 and winners on average earn around \$600 per project. Moreover, since quality is relatively more important, only 40% of the projects that match choose the bid with the lowest quote.

Employers in developed countries will find the prices on freelancer.com much lower than those that local companies quote. A programmer is making on average \$36/hr and a data entry keyer is making \$13.6/hr on average.[9] In India, a programmer makes between \$1 and \$6/hr while a data entry operator is making between \$0.6 and \$2.5/hr.[10]

On the website, it is not straightforward to infer wages from total revenues unless we know for sure that the freelancer is self-employed and has no employees. But assuming all freelancers are working alone, taking the number of days a freelancer who won a project promises to complete it as a rough estimate of how long it actually takes, the column before last of Table 3 shows the implied daily wage of a freelancer who won a project on average. Assuming a work day of eight hours, we see that programmers make roughly \$8/hr, which is between the Indian and the American wage. Engineering freelancers earn \$10/hr, probably due to higher

[9]Source: Occupational Employment and Wages, May 2010.

[10]Source: "Salary Snapshot for Software Engineer / Developer / Programmer Jobs" and "Salary Snapshot for Data Entry Operator Jobs", payscale.com

human capital required for these jobs.[11]

Table 3: Average quotes ranges and winning bids' premia over the lowest bid for selected job types

		All projects					Projects that were awarded					
				Average					Average			
Category	Job type	Projects	Bids per Project	Minimum quote	Average quote	Maximum quote	Projects	Bids per Project	Minimum quote	Winner's quote	Winner's quote per day	Share of projects awarded to a minimum quote
Websites, IT & Software	PHP	198,864	13.9	$ 291.64	$ 545.72	$ 1,562.49	95,870	12.2	$ 193.75	$ 292.52	$ 61.84	47%
	Javascript	41,161	13.6	$ 316.14	$ 619.29	$ 1,882.73	19,079	12.4	$ 195.67	$ 302.36	$ 60.61	46%
	Link Building	39,975	13.3	$ 189.69	$ 364.91	$ 1,078.80	20,244	12.4	$ 126.50	$ 207.52	$ 27.43	46%
	Facebook	12,286	9.1	$ 205.05	$ 532.77	$ 1,833.68	5,909	7.7	$ 122.48	$ 222.34	$ 39.65	56%
	SQL	10,549	17.1	$ 391.37	$ 717.99	$ 2,702.12	5,282	15.0	$ 247.15	$ 393.40	$ 67.60	40%
	eCommerce	10,381	18.5	$ 404.84	$ 804.17	$ 3,171.42	4,810	16.6	$ 249.68	$ 385.85	$ 59.46	38%
	Software Architecture	6,695	9.5	$ 498.11	$ 922.03	$ 2,948.85	2,734	8.7	$ 247.16	$ 350.08	$ 60.06	49%
Design, Media & Architecture	Website Design	164,329	19.6	$ 285.18	$ 555.50	$ 1,743.62	79,786	17.7	$ 183.33	$ 297.00	$ 60.56	39%
	Graphic Design	98,738	22.3	$ 211.20	$ 410.03	$ 1,363.70	54,627	20.5	$ 127.47	$ 203.61	$ 50.34	38%
	Flash	60,338	17.0	$ 310.13	$ 598.81	$ 1,881.03	29,168	15.1	$ 196.87	$ 302.39	$ 69.71	42%
	Logo Design	43,941	28.7	$ 134.20	$ 273.92	$ 989.39	26,830	26.3	$ 86.67	$ 140.04	$ 42.58	35%
	Photoshop	26,215	25.5	$ 131.02	$ 260.34	$ 1,099.72	16,604	22.7	$ 84.13	$ 130.00	$ 46.90	42%
Writing & Content	Copywriting	81,515	14.7	$ 119.42	$ 214.00	$ 786.76	51,319	12.6	$ 80.93	$ 116.73	$ 29.11	53%
	Articles	39,623	14.3	$ 94.16	$ 186.86	$ 1,011.33	23,398	12.7	$ 62.72	$ 100.05	$ 31.17	56%
	Research	25,322	15.6	$ 236.46	$ 432.14	$ 1,769.68	12,275	14.4	$ 113.82	$ 172.55	$ 44.18	52%
Sales & Marketing	Internet Marketing	67,711	13.1	$ 265.52	$ 496.17	$ 1,782.59	27,462	12.2	$ 136.43	$ 283.31	$ 64.13	50%
	Bulk Marketing	9,720	9.2	$ 259.16	$ 476.78	$ 2,237.69	3,303	8.8	$ 129.01	$ 400.98	$ 73.07	60%
	Sales	9,595	8.7	$ 624.72	$ 1,069.93	$ 3,780.03	2,054	8.7	$ 266.90	$ 395.00	$ 117.98	56%
Data Entry & Admin	Data Entry	70,622	28.6	$ 172.61	$ 313.45	$ 1,738.64	36,812	21.1	$ 91.63	$ 158.84	$ 86.78	62%
	Data Processing	42,573	25.3	$ 193.63	$ 381.77	$ 1,821.17	20,869	19.2	$ 106.43	$ 184.66	$ 118.63	59%
	Excel	10,863	41.9	$ 232.65	$ 424.32	$ 2,347.90	6,000	29.6	$ 76.36	$ 125.78	$ 49.30	54%
	Virtual Assistant	4,854	27.3	$ 164.95	$ 333.30	$ 2,485.57	2,146	22.7	$ 68.80	$ 105.81	$ 41.32	61%
Mobile Phones & Computing	iPhone	8,361	10.0	$ 660.91	$ 1,265.11	$ 3,239.84	2,715	9.2	$ 420.84	$ 638.87	$ 70.72	40%
	Mobile Phone	6,433	8.9	$ 654.65	$ 1,209.70	$ 3,300.68	1,948	8.3	$ 408.01	$ 588.22	$ 68.56	41%
	Android	3,131	10.1	$ 691.15	$ 1,271.70	$ 3,018.08	1,019	9.3	$ 384.13	$ 601.65	$ 62.44	41%
Business, Accounting, HR & Legal	Project Management	11,486	16.2	$ 509.96	$ 1,058.58	$ 4,114.59	3,539	14.3	$ 263.32	$ 481.88	$ 100.72	48%
Engineering & Science	Engineering	4,145	13.4	$ 520.50	$ 1,221.80	$ 5,408.88	1,488	12.3	$ 261.70	$ 394.08	$ 81.05	43%

2.4 Repeated Activity and Trade

Having seen labor markets form and facilitate matches across international borders it is interesting to look at the distribution of activity across individuals. As in many other markets and activities, a relatively small number of agents are responsible for a large share of total activity. This statistical skewness for employers and freelancers is shown in Figure 4.

Specifically, 64% of employers in the system posted only one project.[12] But

[11] One unexpected result is the high implied-wage in Data Entry and Data Processing jobs. A possible explanation is that such jobs can easily be divided between more than one person who can work in parallel, unlike programming or design jobs. If this is the case, then the revenue-per-day measure is more likely to represent the sum of two or more wages.

[12] Employers in the system are all users who have posted at least one project.

their projects account for only 17% of all projects on the website. On the other side of the distribution, 7% of employers posted 10 projects or more, which account for 56% of the projects. These repeatedly posting employers provide the variation in experience required in the analysis of learning in the second part of this paper.

Freelancers show an even more skewed distribution of placing bids. The top 1% of freelancers in terms of placing bids – those who bid for at least 251 projects – are responsible for 46% of the bids. The reason for this difference in the distribution may lie in the different cost of bidding on, relative to posting, a project.

Figure 4: Distribution of activity across employers and freelancers

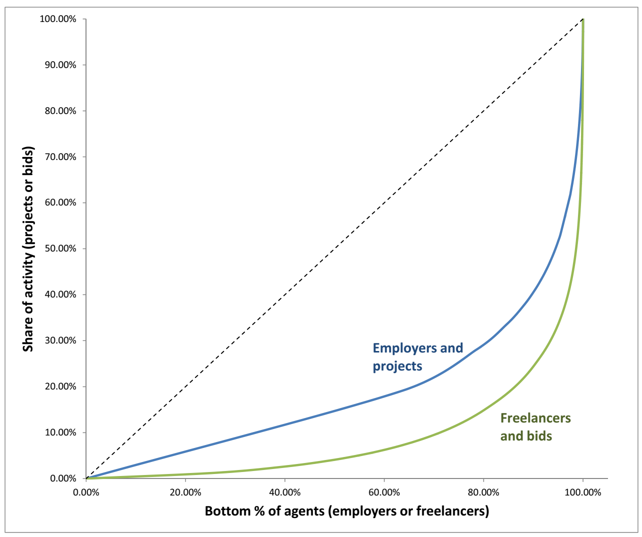

Note: Similar to a Lorenz curve, the horizontal axis has the bottom p percent of the population while the vertical axis has the share of activity created by this bottom p percent. The 45º line represents the curve had each agent created the same number of activities.

Table 4 looks at repeated employer activity by job-type. In table 4a the average number of projects an employer posts is reported. Some job-types have employers posting more projects on average than others. Writing and content is one category where employers repeatedly post. Since projects in this category sometimes need to be done repeatedly and some of them are also relatively smaller, the average employer posts between four to five projects. This is not true for job types in the bottom panel of the table that tackle more basic challenges of a business, such as designing templates or adding eCommerce capabilities to a website. These skills are needed less frequently by the same employer.

Table 4: Repeated posting and hiring of the same freelancer

(a) Repeated posting of projects within a job-type

Job name	Category	Average number of projects posted by employer
Top 10 jobs in repeated project posting		
Articles	Writing & Content	5.04
Article Rewriting	Writing & Content	4.53
Copywriting	Writing & Content	4.11
Ghostwriting	Writing & Content	3.87
PHP	Websites, IT & Software	3.22
Data Entry	Data Entry & Admin	3.18
OSCommerce	Websites, IT & Software	3.07
Proofreading	Writing & Content	2.99
SEO	Websites, IT & Software	2.86
Bottom 10 jobs in repeated project posting		
Training	Other	1.48
Publishing	Writing & Content	1.48
JSP	Websites, IT & Software	1.46
Web Security	Websites, IT & Software	1.46
Photo Editing	Design, Media & Architecture	1.44
Photography	Design, Media & Architecture	1.44
Paypal API	Websites, IT & Software	1.38
Format & Layout	Design, Media & Architecture	1.38
User Interface / IA	Websites, IT & Software	1.34
Corporate Identity	Design, Media & Architecture	1.32

Notes:
Job ranking restricted to jobs with at least 1000 matched projects.
Employers restricted to those who placed their first bid before 2010, to take out the effect of new employers entering disproportionately some job-types more than others

(b) Employers hiring freelancers previously hired by them

Job name	Category	% projects given to a freelancer previously hired by employer
Top 10 jobs in rehiring		
.NET	Websites, IT & Software	12.2%
C# Programming	Websites, IT & Software	12.1%
C++ Programming	Websites, IT & Software	10.5%
Facebook	Websites, IT & Software	10.0%
PHP	Websites, IT & Software	9.8%
Joomla	Websites, IT & Software	9.7%
JSP	Websites, IT & Software	9.7%
Android	Mobile Phones & Computing	9.6%
Drupal	Websites, IT & Software	9.3%
Bottom 10 jobs in rehiring		
Technical Writing	Writing & Content	3.4%
Virtual Assistant	Data Entry & Admin	3.4%
Forum Posting	Writing & Content	3.3%
Reviews	Writing & Content	3.3%
Blog	Writing & Content	3.1%
Publishing	Writing & Content	3.0%
Ghostwriting	Writing & Content	2.9%
eBooks	Writing & Content	2.9%
Editing	Writing & Content	2.8%
Telemarketing	Sales & Marketing	2.4%

Notes:
Job types ranking restricted to jobs with at least 1,000 projects awarded.
Every employer's first project was excluded from the % rehiring denominator since it can't be a rehire by definition.

Lastly, what is the share of projects that is given by employers to freelancers who have already done business with them before? If employers are satisfied and need an additional project to be done, we expect them to hire the same freelancer if the latter is available. On average, roughly 8% of all projects (excluding the first project of every employer) are won by freelancers who were already awarded another project by the employer in the past. The share of projects that rehire previously-hired freelancers differs by job type.

The top-10 jobs in rehiring, as Table 4b shows, are mainly more complicated development projects. These are projects where quality is very important and harder to detect before the project is delivered. Therefore employers will tend to work with freelancers they have already worked with in the past and are satisfied with.

The bottom jobs in terms of rehiring are mainly writing jobs. Writing skills can easily be detected by a short trial and therefore employers are more likely to experiment with new freelancers instead of being locked in by previous freelancers even if they were good. At the same time, some of them may inherently require different freelancers to do different projects the same employer offers (e.g forum posting, reviews and blog posts).

3 Measuring Statistical Discrimination and Learning: Empirical Strategy and Data

In this section I describe two avenues for studying the inference employers make on freelancers based on their country of origin. The first explores how information on country and personal reputation interact. The second focuses on the dynamics of how employers learn from previous experience.

To test the hypotheses in this section I restrict attention to projects flagged with "PHP" job type posted in 2010. PHP is a programming language that allows web servers to send customized pages to the website's visitors. This is a very common programming language on the web. Even prominent websites like facebook.com and freelancer.com itself use PHP. By focusing on a single market – the largest on freelancer.com – I avoid combining different labor markets with different characteristics such as how sensitive the service is to quality or how hard it is to detect it *ex ante.* It also makes the analysis easier computationally. Each observation in my dataset is a bid on a project. An observation contains project, bid, freelancer, and employer details as they appear in the original database. An indicator of whether the bid was selected by the employer will be our dependent variable of interest.

3.1 Individual Reputation vs. Country Reputation

Freelancers on the website differ on a whole range of characteristics. In addition to the freelancer's country, other details are visible on a user profile. Moreover, after each project the employer and freelancer rate each other on a 0-to-10 scale. If both employer and freelancer rated each other, these ratings become public. Perhaps the most prominent feature a freelancer and an employer may have to credibly signal their quality is the average, and number of, ratings.

How does this individual reputation mechanism interact with the country-of-origin signal? If discrimination would have been taste-based, individual reputation should not change the likelihood to be hired, since the disutility from employing a freelancer in a certain country is independent of productivity. Unlike statistical discrimination, taste-based discrimination is not the outcome of lack of information. If there is at least some component of statistical discrimination then when more credible signals of individual ability are available the country should matter less.

Since reviews do not become public if one of the sides does not review the other, those who expect bad reviews usually do not rate the other side and make the review "expired." Thus, the distribution of *public* average ratings is concentrated around the 10-out-of-10, as Figure 5 shows.[13] Even though average ratings is almost always very close to 10, the number of reviews signal both quality and experience in the market. Thus, I restrict attention to the number of reviews.

[13]Note, however, that the website has unpublished reputation available and I can observe them. This will become important for the exercise in 3.2.

Figure 5: The Distribution of Average and Number of Ratings

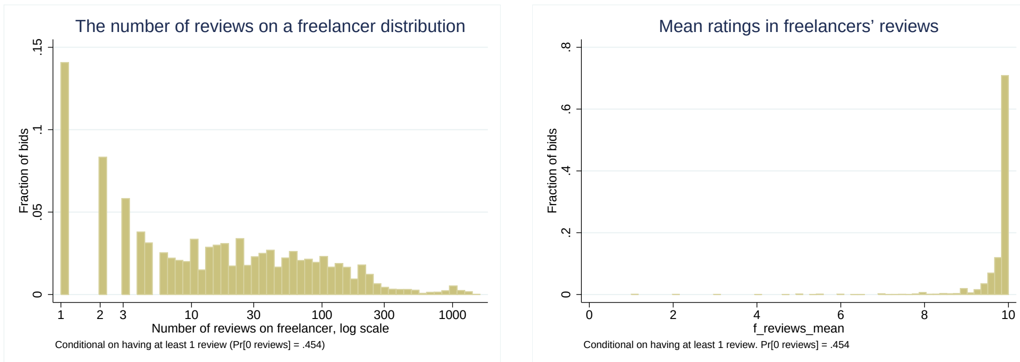

Notes:

1. Distributions plotted for all freelancers bidding on PHP projects in 2010, weighted by the number of bids.

To complete the exercise described below, I supplement the dataset with additional variables used for the analysis. First, I reconstruct freelancer reputation to have the number of publicly observable ratings *at the time the bid was made.*[14] Second, I reconstruct the number of previous matches made between the employer and the freelancer prior to posting the project.

I estimate the employer's demand for freelancers as a function of the number of reviews, price, and previous matches between the employer and freelancer. To capture differential hiring probability by country I let the constant and the regression coefficient on the number of reviews be country-specific, as equation (1) shows:[15]

$$won_{bp} = \sum_{c \in C} \gamma_c c_{bp} + \sum_{c \in C} \beta_c c_{bp} \times \ln(1 + reviews_{bp}) + \delta \ln price_{bp} + \lambda pastProj_{bp} \qquad (1)$$
$$+ \sum_{p \in P} \phi_p + \varepsilon_{bp}$$

where c_{bp} is a dummy that equals 1 if the freelancer is from country c, ϕ_p are project fixed effects, γ_c are the country fixed effects, and β_c measures, for country

[14]I am able to do so thanks to the fact that each review completed for the agent is saved individually with the date of submission.

[15]I choose a linear probability model for computational ease. For now I am interested in the signs of the coefficients rather than the accurate point estimate.

c, the return to each additional 1% reviews.

One must also take into account how $\ln price_{bp}$ is determined. A freelancer sets price so that it will maximize the expected revenue from bidding and provide ex-post positive profits. Conditional on a freelancer's country of origin and publicly available individual reputation, the freelancer will trade off a high quote for lower probability to win. When a new review becomes available the freelancer will likely increase the price quoted. Additionally, price may be correlated with other characteristics that I am not capturing but may also be correlated with winning the bid. I therefore need a "supply" shifter: a factor that determines pricing but does not affect the likelihood to be selected other than through the price.

I use fluctuations in the exchange rates to the U.S dollar as an instrument. Quotes on the website are virtually always given in U.S dollars. The exchange rate between the dollar and the local currency should affect the offered price because freelancers' costs and alternative income sources are mostly denominated in their local currency. Thus, when the local currency devalues – when each dollar is now worth more in the local currency – the freelancer can bid lower than otherwise and still get the same amount in local terms. The exchange rate is excluded from the demand equation if I restrict attention to American employers. Since their local currency is the U.S dollar their decision is unaffected by the cost side of the freelancer in ways that are not operating through the quoted price. Thus, the the supply equation will be:

$$\begin{aligned}\ln price_{bp} = &\sum_{c\in C}\alpha_c c_{bp} + \sum_{c\in C}\theta_c c_{bp} \times \ln(1 + reviews_{bp}) + \varphi pastProj_{bp} \\ &+ \xi exRate_{bp} + \sum_{p\in P}\phi_p + \nu\end{aligned} \tag{2}$$

where $exRate_{bp}$ is an exchange rate index a moving average of the exchange rate between the freelancer's country's currency and the U.S. dollar, normalized to the average exchange rate in December 2009. Figure 9 in Appendix A shows how this index evolves over 2010 for selected currencies.

If discrimination is taking place, then first we need to see that at least for unrated freelancers, those who live in some countries are significantly more likely to be hired than those from other countries. In other words, and if we assume discrimination is favoring freelancers from developed countries, I expect $\gamma_{developing} < \gamma_{developed}$. If discrimination is at least in part statistical, then more information on the ability of an individual should decrease the gap in the likelihood to be hired. Therefore the coefficients for individual reputation should switch sides: $\beta_{developing} > \beta_{developed}$, so that as individual reputation is better the gap in the prediction of the likelihood to be hired will decrease.

Results

Figure 6 summarizes the findings: Bangladeshi and Vietnamese freelancers with no reviews are much less likely to be hired relative to Canadian or Australian freelancers. However, as freelancers earn good ratings the difference slowly decreases until it disappears. The regressions are reported in Table 5. Column (1) shows that better individual reputation as well as having already worked for the employer in the past are positively correlated with being selected. Prices are negatively associated with the probability to be hired. This coefficient captures a bias caused by freelancers endogenously choosing price so as to not decrease the probability to be hired by too much. Once we use the IV approach, in column (4), the coefficient increases dramatically.

Looking at the coefficients on country dummies (γ_c) in column (4) we can see a clear ranking of countries: all countries have a negative coefficient meaning that unrated American freelancers are preferred by American employers to all other freelancers. Still, all three developed countries that are not the U.S lie above the rest of the countries. As for the slope of individual reputation, for most countries it is not statistically significant. However, the point estimates are mostly in line with an opposite ordering (relative to the country dummies). Moreover, Bangladesh is the country with the highest slope, that is also statistically and economically higher than the American slope, and the lowest coefficient on the dummy.

Figure 6: Country vs. individual reputation: a graphical representation

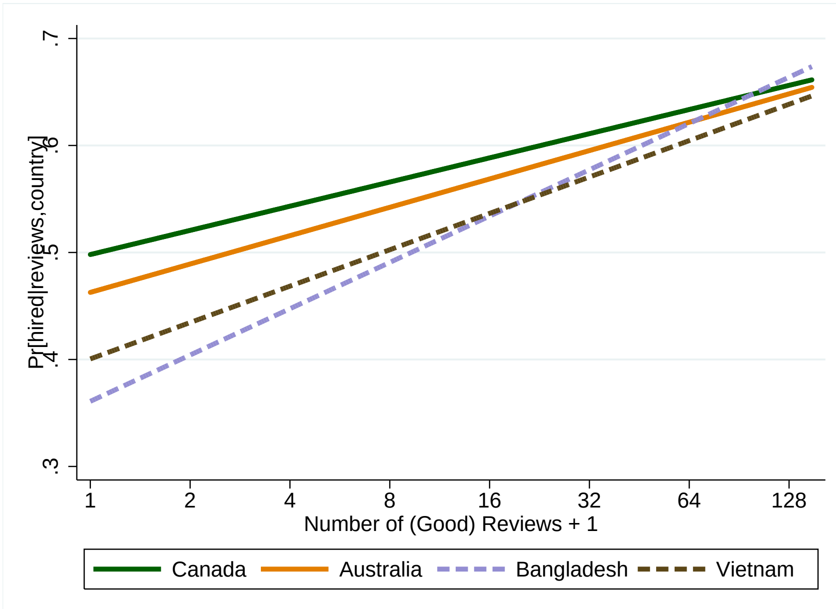

Notes:
This figure plots the effects from Table 5 for selected countries.

3.2 Employers Learning from Experience

The previous section detected evidence for statistical discrimination. I would like to show how employers form their opinions, or beliefs, about the average freelancer from country c.

Employers beliefs about the distribution of skills within a group they can identify are at the core of the theory of statistical discrimination. But where do these beliefs originate from? Following examples from Ellsberg (1961), consider two urns – one for each country – with "good"- and "bad"-marked balls. Each ball represents a freelancer that can either be of a "good" or a "bad" type, where "good" types yield higher revenues to the employer. The employer needs to decide which urn to take a ball from. The decision will depend on her belief of the proportion of "good" and "bad" balls within each urn.

However, if the employer is repeatedly taking out balls from an urn, she can infer the proportion of "good" balls in the urn. If the employer is American, she may have a good idea of the proportion in the American urn from prior experiences with American freelancers, but having no prior experience with freelancers from Bangladesh, her prior beliefs on the proportion of "good" types is not clear.[16] Therefore, each good or bad experience with the Bangladeshi freelancer is expected to have a significant effect on the employer's choice next time she faces a Bangladeshi freelancer. This intuition is more formally laid out as a simple dynamic problem in Appendix B.

Three features of employers on freelancer.com make them good candidates for the task of documenting how market experience affects subsequent observable behavior through learning: First, employers are observed hiring multiple times. This means that if they learn from their experience in one project, they should behave differently in subsequent projects and observing all their projects allows to detect such changes in behavior. After taking a "bad" ball from the Bangladeshi urn, is the employer more likely to go for the other urn or try another ball from the same urn?

Second, employers' experience gained from a project, in the form of service satisfaction, is at least partially observable. The reviews freelancers and employers fill on each other, project cancellations, filings of disputes on the website, and timely payments are all signs of either good or bad outcomes. In other words, I observe whether the ball the employer takes out of the urn is "good" or "bad."

Third, most employers are likely to have no prior experience with freelancers from developing countries, which makes every project – at least in the beginning – a significant learning experience. If employers already had a long history of trading with individuals from the freelancer's country before, the specific project that I as a researcher observe its outcome would be only one event in a long list of experiences with that country and therefore shouldn't have changed employers' beliefs about agents from that country by much.

To implement this idea I assign each freelancer the prior experience that the

[16]Actually, she may have a good idea on the general population of Americans, but not on those who self-select to this platform. With the unfamiliar country I can assume that she does not know much about either the general population's or the self-selected population's proportion of "good" types.

employer had with freelancers from this freelancer's country before posting the project. Experience is captured in two variables: number of good outcomes and bad outcomes. I use employers' ratings of freelancers, projects cancellations and disputes initiated by employers to classify past projects of each employer to "good" and "bad". For the classification of project outcome I take both publicly observable and publicly-unavailable reviews. Unlike public ratings, many non-public reviews are bad.

A good outcome is defined as a perfect review (=10) as rated by the employer. A bad outcome is defined as one of the following three: bad reviews given by the employer ($\leq$ 5), a project for which the employer initiated a dispute-resolution request, or a project that the employer canceled.[17]

I estimate the following regression:

$$\begin{aligned} won_{bp} = \alpha + \beta \ln\left(1 + reviews_{bp}\right) + \delta \ln price_{bp} + \lambda pastProj_{bp} \\ + \eta_{good} goodExp^{c}_{bp} + \eta_{good} badExp^{c}_{bp} + \eta_{good2}\left(goodExp^{c}_{bp}\right)^{2} + \eta_{bad2}\left(badExp^{c}_{bp}\right)^{2} \\ + \eta_{goodbad} goodExp^{c}_{bp} \times badExp^{c}_{bp} + \sum_{p \in P} \phi_p + \varepsilon_{bp} \end{aligned} \tag{3}$$

The partial derivatives implied by this functional form are, for good experiences:

$$\frac{\partial won}{\partial goodExp} = \eta_{good} + 2\eta_{good2} goodExp + \eta_{goodbad} badExp \tag{4}$$

and symmetrically for bad experiences. Thus, following the first implication of the model, I expect that $\eta_{bad} < 0 < \eta_{good}$, since in absence of prior experience, if the first experience with country c is good (bad) then we expect the employer to be more (less) likely to hire another freelancer from c the next project. We also expect η_{good2}, the coefficient on the quadratic term, to be negative since if we have more and more experiences that are good, the marginal good experience will update our beliefs upwards less and less. Our beliefs have already formed. The sign on $\eta_{goodbad}$ should depend on the number of good and bad experiences.

[17]Note that according to this classification, some projects will be not categorized as either good or bad. I do not count them in my measures of good and bad experience as I cannot classify them.

Table 6 reports the regressions in question. Even though the OLS specifications in columns (1)-(3) yields coefficients all in line with the hypotheses, the preferred specifications are the IV specifications: (7)-(9). Indeed, if we just look at a linear specification in column (7) then every 10 good experiences contribute 0.6 percentage points to the likelihood of hiring another freelancer from the same country, while every 10 bad experiences decrease this likelihood by 0.3 percentage points. These are not very large numbers economically.

Once we try the full quadratic specification we get that the first good experience with a freelancer from country c increases the probability of hiring another one from c in the next project by 0.26 percentage points. The effect of a first bad experience, however, is not statistically significant, and the point estimate is even positive in contrast to what we would expect.

4 Conclusion

Employers' search for employees has always been a careful process in which employers try to infer the expected productivity of prospects. Traditionally, the search for prospective employees was geographically concentrated near the employer. The rise of the Internet allowed online marketplaces such as freelancer.com to evolve and facilitate search for freelancers anywhere on the globe. Even the smallest employer can now exploit wage gaps between rich and poor countries to reduce its labor costs significantly. Online labor markets make matching between employers and employees less costly and less locally biased. Wherever they are, employers and freelancers on the website are as far as a click of a mouse from each other. If some labor services can be provided remotely, and if an employee in a developing country is a good-enough substitute for a developed-country employee, then economic theory predicts that employers from developed countries will hire more freelancers from developing countries. In the long run, at least for these services, the wage gap between rich and poor countries should decrease and labor markets should be less geographically segregated.

It is nevertheless unclear to what extent labor markets can integrate globally through such online platforms. Employers can find much cheaper labor services in new economies, but they are also entering new and unfamiliar domains where employees are harder to screen for productivity, quality, and trustworthiness, given their different observable characteristics, or lack thereof. Indeed, this paper showed a country-of-origin effect for freelancers on the website. It also showed how individual reputation may alleviate the country-of-origin effect. Lastly, it showed that some learning by employers is taking place: they respond mainly to good outcomes by hiring more freelancers from the same country of freelancers they had successful matches with.

Different implications for policy arise if statistical discrimination stems from employers not knowing the distribution of skills in the freelancer's country, or if employers know the distribution and it is inferior. If the main problem is ambiguity – employers not knowing the distribution of skills – more accurate and credible information should be disseminated to employers. If, on the other hand, it is objective differences in the distribution of productivity or quality, then additional screening by the platform that will guarantee higher average quality may help the ones who get clearance build up reputation and provide positive externalities to all potential service exporters in the way Chisik (2003) describes.

References

Agresti, Alan, and David Hitchcock. 2005. "Bayesian inference for categorical data analysis." *Statistical Methods & Applications*, 14: 297–330. 10.1007/s10260-005-0121-y.

Altonji, Joseph G., and Charles R. Pierret. 2001. "Employer Learning and Statistical Discrimination." *The Quarterly Journal of Economics*, 116(1): pp. 313–350.

Arrow, Kenneth J. 1973. "Discrimination in labor markets." , ed. Orley C. Ashenfelter and Albert Rees, Chapter The theory of discrimination, 3–33. Princeton University Press.

Becker, Gary S. 1957. *The Economics of Discrimination.* University of Chicago Press.

Bertrand, M., and S. Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *The American Economic Review*, 94(4): 991–1013.

Chisik, Richard. 2003. "Export industry policy and reputational comparative advantage." *Journal of International Economics*, 59(2): 423 – 451.

Clemens, Michael A. 2011. "Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?" *Journal of Economic Perspectives*, 25(3): 83–106.

Dinnie, Keith. 2004. "Country-of-Origin 1965-2004: A Literature Review." *Journal of Customer Behaviour*, 3(2): 165–213.

Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms." *The Quarterly Journal of Economics*, 75(4): 643–669.

Hu, Ye, and Xin Wang. 2010. "Country-of-Origin Premiums for Retailers in International Trades: Evidence from eBay's International Markets." *Journal of Retailing*, 86(2): 200 – 207. <ce:title>Special Issue: Modeling Retail Phenomena</ce:title>.

List, John A. 2004. "The Nature and Extent of Discrimination in the Marketplace: Evidence From the Field*." *Quarterly Journal of Economics*, 119(1): 49–89.

Pallais, Amanda. 2010. "Inefficient Hiring in Entry-Level Labor Markets."

Phelps, Edmund S. 1972. "The statistical theory of racism and sexism." *The American Economic Review*, 62(4): 659–661.

Stanton, Christopher, and Catherine Thomas. 2011. "Landing the First Job: The Value of Intermediaries in Online Hiring." *SSRN eLibrary*.

A Additional Figures

Figure 7: Browsing newest projects on freelancer.com: a range of jobs awaiting to be filled.

5890 Freelance Jobs Found Show 50

Project Name	Bids	Avg (USD)	Job Type	Started	Left
New Website Construction We need a new website built. We will provide the basic design, all text and all images. The site will be approx. 15 pages all using the same basic template. CMS system required. Shopping cart required. This project must be 100% complete by ...	29	$1645	PHP, Website Design	Today	4d 15h
Wordpress Project Development	1	$950	CSS, MySQL, PHP, Script Install, Wordpress	Today	15h 54m
iPhone app - Add 2-3 features and solve bugs	3	$177	iPhone, Objective C	Today	15h 54m
Html Expert need, Urget	2	$30	CSS, HTML, Javascript, PHP, Script Install	Today	15h 53m
Simple After Effects Changes	11	$73	After Effects	Today	4d 15h
30 articles needed related to furniture TRIAL	2	$30	Copywriting	Today	4d 15h
Send 5 million emails for us TRIAL	1	$200	Internet Marketing, Link Building, SEO	Today	59d 15h
Content Writing TRIAL	2	$63	Product Descriptions, Publishing, Report Writing, Speech Writing, Technical Writing	Today	59d 15h
iphone art sales application design TRIAL	1	$800	iPhone, Mobile Phone	Today	16d 15h
Script File For PTC Site TRIAL	2	$750	Graphic Design, HTML, PHP, Website Design	Today	4d 15h
Bulk Mail Solution Needed TRIAL	1	$120	Bulk Marketing	Today	2d 15h
Facebook Friends need TRIAL	4	$150	Facebook	Today	6d 15h
Data Entry "PdF file to Word Doc	43	$763	Data Entry, Data Processing	Today	4d 15h

Note: Accessed in November, 2010

Figure 8: Bids on a random project: a range of freelancers to choose from.

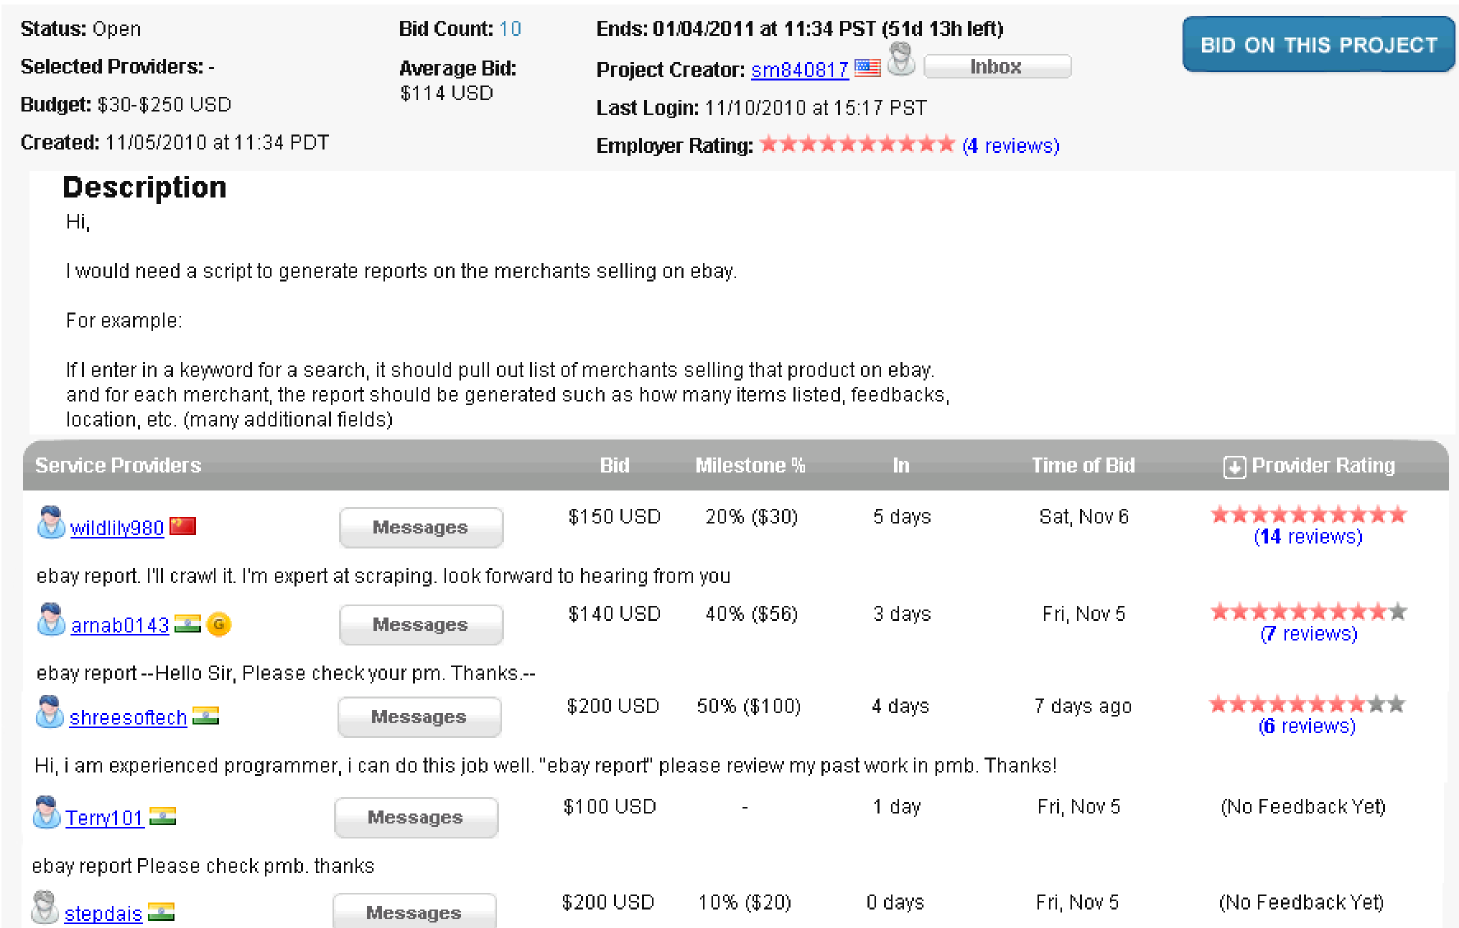

Note: Accessed in November, 2010

Figure 9: Selected exchange rate fluctuations in 2010

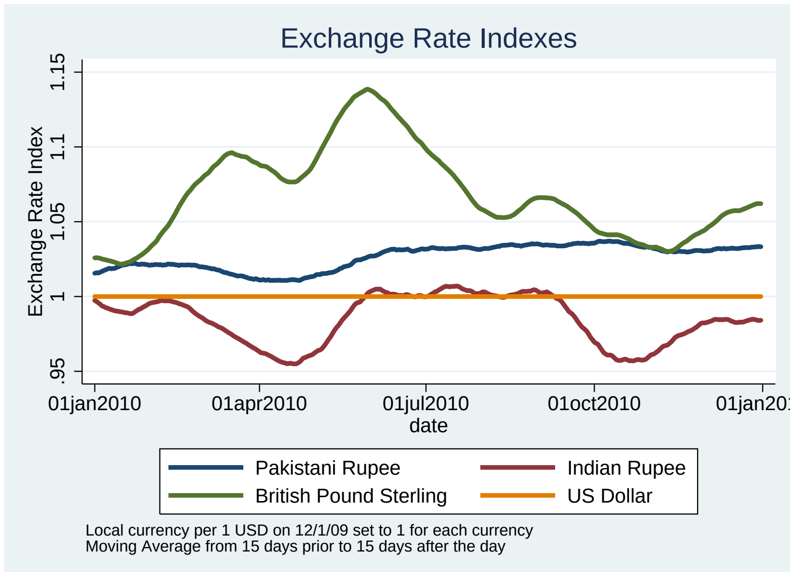

B Modeling the Employer's Decision

Since learning and statistical discrimination are concepts that have been used in more than one context and setup in economics I am going to lay out a simple model of how employers experiment with freelancers of unfamiliar countries and learn about employees from these countries as they gain experience with them in the market.

Suppose there is one risk-neutral employer serially posting projects: hiring a freelancer, observing the outcome of the project and then posts another project. On each project two new freelancers place a bid: one from a known country (k) and one from an unknown country (u). There are only these two countries in the market. The freelancers quote two constant wages: w_k and w_u for the known- and unknown-country freelancers respectively. Assume that $p_k - w_k > 0$ so that the employer will always hire a freelancer.[18] Freelancers can either be "good" types or "bad" types. Good and bad types yield revenue q normalized to 1 and

[18]Only the employer's behavior is modeled here. Freelancer behavior is not modeled here. The analysis will be equivalent if it wasn't an employer but a consumer deciding between buying an apple from a known brand or an apple from an unknown brand.

0 respectively. Thus, each project's *ex ante* payoff to the employer from hiring freelancer i is:

$$\pi_i = \mathbb{E}q_i - w_i \tag{5}$$

Assume that the proportion of good types in the known country is known and denoted by p_k. In contrast, the proportion of good types in the unknown market is unknown and is subjectively distributed according to $p_u \sim F_u$. This means that the employer is uncertain about both candidates, but the unknown-country candidate introduces a second layer of uncertainty: ambiguity regarding the distribution of the good-vs.-bad types in the country. Thus, the *ex ante* payoff for the known- and unknown-countries' freelancers is:

$$\pi_k = p_k - w_k \quad , \quad \pi_u = \mathbb{E}\left[p_u | \, \Omega_t\right] - w_u \tag{6}$$

where Ω_t is the information the employer at the time she makes the decision for project t.

Experience, embodied in Ω_t is the only difference between projects. After each project the employer learns about the hired freelancer's type, and therefore Ω_t will be the number of good and bad types in the unknown country.[19],[20] Thus, we can rewrite (6) in more tangible terms:

$$\pi_k = p_k - w_k \quad , \quad \pi_u = \mathbb{E}\left[p_u | \, n_g, n_b\right] - w_u \tag{7}$$

where n_g and n_b are the number of good and bad freelancers from the unknown country.

Experience makes the employer's decision more informed. Therefore a forward looking employer thinks about her information in future projects when deciding whether to hire a freelancer from an unknown country presently. We can write

[19]For the sake of simplicity I assume that for the known country the employer does not keep track of good and bad experience. This is equivalent to having a long record of trading with freelancers from the known country, thereby making the projects I observe negligible in the formation of beliefs on known-country freelancers.

[20]I also assume that the proportion of good types does not change with time, and therefore the order of good and bad experiences does not matter for the inference of p_u.

the dynamic problem as a Bellman equation:

$$\mathbb{V}(n_g, n_b) = \max\{p_k - w_k \;\;,\;\; \mathbb{E}[p_u | n_g, n_b] - w_u\} + \beta \mathbb{E}\mathbb{V}\left(n_g^{'}, n_b^{'}\right) \tag{8}$$

where the laws of motion for $n_g^{'}$ and $n_b^{'}$ are:

$$\begin{aligned} n_g^{'} &= n_g + \mathbb{I}\{c_i = u\} \times \mathbb{I}\{q_i = 1\} \\ n_b^{'} &= n_b + \mathbb{I}\{c_i = u\} \times \mathbb{I}\{q_i = 0\} \end{aligned} \tag{9}$$

Lemma 1. *If the employer hires a known-country freelancer at time t then she will hire a known-country freelancers for all $\tau > t$ as well. The value of hiring locally – present and future payoffs aggregated – is then $\frac{p_k - w_k}{1-\beta}$.*

Proof. note that if the employer hires a known-country freelancer, $\mathbb{I}\{c_i = u\} = 0$, then $\mathbb{V}\left(n_g^{'}, n_b^{'}\right) = \mathbb{V}(n_g, n_b)$. Thus, the problem will be identical the next period because no further experience with the unknown-country freelancers was gained. Therefore in each period after t the employer solves the same problem and gets the same policy: hire a known-country freelancer. Since this is the case, each period's payoff from that period on is equal to $p_k - w_k$ and therefore $\sum_{t=0}^{\infty} \beta^t (p_k - w_k) = \frac{p_k - w_k}{1-\beta}$ □

Thus, I can simplify (8) and get:[21]

$$\mathbb{V}(n_g, n_b) = \max\left\{\frac{p_k - w_k}{1-\beta} \;\;,\;\; \widehat{p_u} - w_u + \beta\left[\widehat{p_u}\mathbb{V}(n_g + 1, n_b) + (1 - \widehat{p_u})\mathbb{V}(n_g, n_b + 1)\right]\right\} \tag{10}$$

where $\widehat{p_u} = \widehat{p_u}(n_g, n_b) \equiv \mathbb{E}(p_u | n_g, n_b)$. This equation embodies the tradeoff that the employer faces: experimenting with unknown-country freelancers may yield better

[21]Note that

$$\begin{aligned} \mathbb{E}\mathbb{V}\left(n_g^{'}, n_b^{'}\right) &= \mathbb{E}_{F_u}\left[\mathbb{E}\left[\mathbb{V}\left(n_g^{'}, n_b^{'}\right) \middle| p_u\right]\right] = \mathbb{E}\left[p_u \mathbb{V}(n_g + 1, n_b) + (1 - p_u)\mathbb{V}(n_g, n_b + 1)\right] \\ &= \mathbb{V}(n_g + 1, n_b)\mathbb{E}[p_u] + \mathbb{V}(n_g, n_b + 1)\mathbb{E}[1 - p_u] \\ &= \mathbb{V}(n_g + 1, n_b)\mathbb{E}[p_u] + \mathbb{V}(n_g, n_b + 1)[1 - \mathbb{E}[p_u]] \\ &= \mathbb{V}(n_g + 1, n_b)\widehat{p_u} + \mathbb{V}(n_g, n_b + 1)(1 - \widehat{p_u}) \end{aligned}$$

continuation value if the experience is good. The probability that this experience will, indeed, be good, depends on the current best guess of what the proportion of good employees in the unknown country is.

What will determine the dynamics is how the expected value of the subjective distribution of the proportion will change with good and bad experiences. In other words, the function $\widehat{p_u}\left(n_g, n_b\right)$. I assume that employers follow Bayesian rules in updating their subjective distributions. The final nail will be to choose the prior subjective distribution.

Agresti and Hitchcock (2005) surveys possible priors proposed for binomial distributions. One plausible prior could be $U\left(0,1\right)$, assigning equal probability for each possible p_u. When that is the case, the posterior distribution's expected value is:

$$\widehat{p_u}\left(n_g, n_b\right) = \frac{n_g + 1}{n_g + n_b + 2}$$

Even if the prior is not $U\left(0,1\right)$, assume that $\frac{\partial \widehat{p_u}}{\partial n_b} < 0 < \frac{\partial \widehat{p_u}}{\partial n_g}$. This means that every good experience increases the expected quality of the unknown country while every bad experience decreases it.

First testable prediction: Employers are more (less) likely to hire freelancers from the unknown country after a good (bad) experience than before it.

Second testable prediction: As the employer earns a higher "stock" of experiences, the change in hiring likelihood induced by each additional experience is smaller.

Both employer and freelancer Canadian

Table 5: Country vs. individual reputation: Country-of-origin effects for rated and unrated freelancers

		All		U.S employers ONLY	
		OLS	OLS	OLS	IV
		(1)	(2)	(3)	(4)
	Quote in $	-0.076 ***	-0.075 ***	-0.076 ***	-0.566 ***
		(0.001)	(0.001)	(0.002)	(0.115)
	ln(number of reviews + 1)	0.020 ***	0.021 ***	0.024 ***	0.045 ***
		(0.000)	(0.001)	(0.002)	(0.005)
	Have matched in the past once	0.020 **	0.020 **	0.026 ***	0.021 ***
		(0.009)	(0.009)	(0.006)	(0.001)
Freelancers from:					
United Kingdom	from the UK?		0.004	-0.007	-0.036 ***
			(0.004)	(0.007)	(0.012)
	UK X ln(reviews + 1)		0.010 ***	0.008 *	0.003
			(0.003)	(0.005)	(0.005)
Canada	from Canada?		0.005	0.011	-0.013
			(0.006)	(0.010)	(0.014)
	Canada X ln(reviews + 1)		-0.002	-0.007	-0.012
			(0.004)	(0.007)	(0.008)
Australia	from Australia?		0.008	-0.020	-0.054 ***
			(0.008)	(0.013)	(0.019)
	Australia X ln(reviews + 1)		0.035 ***	0.027 **	-0.006
			(0.007)	(0.011)	(0.012)
India	from India?		-0.037 ***	-0.044 ***	-0.080 ***
			(0.002)	(0.003)	(0.010)
	India X ln(reviews + 1)		-0.003 **	-0.004 *	-0.002
			(0.001)	(0.002)	(0.002)
Pakistan	from Pakistan?		-0.036 ***	-0.042 ***	-0.100 ***
			(0.003)	(0.004)	(0.015)
	Pakistan X ln(reviews + 1)		-0.001	-0.002	0.003
			(0.001)	(0.002)	(0.003)
Bangladesh	from Bangladesh?		-0.044 ***	-0.051 ***	-0.180 ***
			(0.003)	(0.005)	(0.031)
	Bangladesh X ln(reviews + 1)		0.001	0.001	0.018 ***
			(0.002)	(0.003)	(0.005)
Vietnam	from Vietnam?		-0.013 **	-0.028 ***	-0.127 ***
			(0.005)	(0.008)	(0.025)
	Vietnam X ln(reviews + 1)		-0.007 ***	-0.009 ***	0.004
			(0.002)	(0.003)	(0.005)
All other non-US countries	is other?		-0.020 ***	-0.034 ***	-0.101 ***
			(0.002)	(0.004)	(0.016)
	other X ln(reviews + 1)		0.009 ***	0.010 ***	0.006 **
			(0.001)	(0.002)	(0.003)
	Observations	336,556	336,556	137,619	137,398
	Project Fixed Effects	26,447	26,447	11,218	11,215

Notes:

1. The omitted category is U.S-based freelancers. The coefficient of $\ln(reviews + 1)$ in columns (2)-(4), therefore, pertains to U.S freelancers as well.
2. "All other non-US countries" represents a grouping of all countries that are neither listed in the table nor the U.S and treated as one country.
3. Quote in $ instrumented in column (4) using a moving average of exchange rate 15 days before through 15 days after the day of the bid.
4. Sample restricted to PHP projects in 2010. Experience and other historical variables constructed for all skills and time on the website.
5. Standard errors clustered at the project level. *** p<0.01, ** p<0.05, * p<0.1

Table 6: Learning from experience: Hiring as a function of the employer's experience with the freelancer's country

	All			U.S employers only					
	OLS	OLS	OLS	OLS	OLS	OLS	IV	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Quote in $	-0.076 ***	-0.076 ***	-0.076 ***	-0.076 ***	-0.076 ***	-0.076 ***	-0.458 ***	-0.494 ***	-0.502 ***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.040)	(0.044)	(0.044)
Prior experience with freelancer's country									
Good / 10	0.003 ***	0.017 ***	0.020 ***	0.004 **	0.016 ***	0.017 ***	0.006 ***	0.025 ***	0.026 ***
	(0.001)	(0.005)	(0.005)	(0.002)	(0.003)	(0.003)	(0.001)	(0.003)	(0.003)
Bad / 10	-0.007 ***	-0.025 ***	-0.023 ***	-0.006 ***	-0.014 ***	-0.011 **	-0.003 **	-0.000	0.005
	(0.001)	(0.004)	(0.004)	(0.002)	(0.004)	(0.004)	(0.001)	(0.004)	(0.004)
$(\text{Good} / 10)^2$		-0.0004 **	-0.0004 **		-0.0006 ***	-0.0004 **		-0.0010 ***	-0.0007 ***
		0.0002	0.0002		0.0001	0.0002		0.0001	0.0002
$(\text{Bad} / 10)^2$		0.0011 ***	0.0014 ***		0.0004	0.0006 **		-0.0005 *	-0.0002
		0.0002	0.0003		0.0003	0.0003		0.0002	0.0003
(Good / 10) X (Bad / 10)			-0.0021 ***			-0.0017 **			-0.0027 ***
			0.0006			0.0008			0.0007
ln(number of reviews + 1)	0.020 ***	0.020 ***	0.020 ***	0.022 ***	0.022 ***	0.022 ***	0.039 ***	0.041 ***	0.041 ***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.002)	(0.002)	(0.002)
Number of previous	0.020 **	0.021 **	0.020 **	0.025 ***	0.025 ***	0.025 ***	0.022 ***	0.021 ***	0.021 ***
	(0.009)	(0.009)	(0.009)	(0.006)	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
Observations	336,556	336,556	336,556	137,619	137,619	137,619	137,398	137,398	137,398
Project Fixed Effects	26,447	26,447	26,447	11,218	11,218	11,218	11,215	11,215	11,215

Notes:
1. Good and bad experience are counted in units of 10 experiences to ease reporting of the coefficients. To get the effect of a single good or bad experience, one should divide the coefficients by 10 (quadratic and cross-term coefficients by 100).
2. Quote in $ instrumented in columns (7)-(9) using a moving average of exchange rate 15 days before through 15 days after the day of the bid.
3. Sample restricted to PHP projects in 2010. Experience and other historical variables constructed for all skills and time on the website.

4. Standard errors clustered at the project level. *** p<0.01, ** p<0.05, * p<0.1