IRLE WORKING PAPER
#113-16
September 2016

Learning to hire? Hiring as a dynamic experiential process in an online market for contract labor

Ming D. Leung

Cite as: Ming D. Leung. (2016). “Learning to hire? Hiring as a dynamic experiential process in an online market for contract labor”. IRLE Working Paper No. 113-16. http://irle.berkeley.edu/workingpapers/113-16.pdf

Learning to hire?
Hiring as a dynamic experiential process in an online market for contract labor

Ming D. Leung
Haas School of Business
University of California, Berkeley
2220 Piedmont Avenue
Berkeley, CA 94720
1-510-642-4952
mingdleung@haas.berkeley.edu

**** Working Paper ****

Learning to hire?
Hiring as a dynamic experiential learning process in an online market for contract labor

Abstract

Can employers learn to hire? This article conceptualizes hiring as a dynamic experiential learning process. Instead of examining hiring as a point in time decision, I investigate whether and how employers' past hiring experiences affect their future decisions. Drawing on evidence from a global online market for contract labor, I argue that employers revise their beliefs regarding job applicants from a particular social category following a negative hiring experience from that social category. I analyze over 16 Million applications from freelancers worldwide for over 2.2 Million jobs from 557,416 employers. I find that employers who have a negative hiring experience with a freelancer from a particular country are subsequently less likely to hire other freelancers from that country. This effect is stronger on hiring for identical subsequent jobs and weaker for other jobs. Most strikingly, evidence from the actual hiring switches following a negative experience and a simulation using data from the observed distribution of freelancers on the platform demonstrate that employers unnecessarily oversteer away from countries given the narrow distribution of observable ability among freelancers. Switches do not result in hiring from a "better" country.

Keywords

hiring, learning, labor markets, freelancing, discrimination, evaluation, country of origin

In deciding whom to hire, employers face the fundamental problem of discerning the difficult to observe ability of job applicants. Understanding how employers attempt to resolve this uncertainty through the acquisition of additional information has animated a breadth of research in this domain. For example, educational credentials are believed to signal job applicant quality (Spence 1973, Sorensen 1983). Job trials, in the form of internships, provide employers direct evidence of how well a particular employee may eventually perform in a full-time position (Baron and Kreps 1999, Sterling 2012). Past demonstrations of applicable work experiences indicate future potential to employers as well (Bills 1990, Doeringer and Piore 1971).

Since at least Phelps (1972) and Arrow (1972), employers have also been portrayed as attempting to resolve this uncertainty by relying on the visible social category of a job applicant, such as their race, gender or country of origin, to infer their potential underlying ability. If job applicant skill is correlated with their social group membership, then employers can reduce hiring uncertainty by employing individuals from those social groups who possess, on average, better skill. For example, to the extent that upper body strength is greater in men than women, employers should be more likely to hire men for a task that requires heavy lifting than women because men will be expected to perform better than women at the task (Bielby and Baron 1986). The reliance on social category cues is also believed to lead to employer 'steering' behavior, where applicants from particular social groups are encouraged to apply to jobs that are, ostensibly, more suited for them due to their race (Pager et al 2009) or gender (Fernandez and Mors 2008).

Despite the recognition that a job applicant's social category membership affects hiring outcomes because employers may infer information from them, very little work has examined whether and how employers learn about and use these vital social cues. A reasonable assumption may be that employers draw from their own past hiring experiences to inform their future hiring decisions in this regard, yet virtually no work has demonstrated whether this is the case (c.f. Pager and Karafin 2009). I suggest one reason for this is because for the most part, hiring has been studied as a *point in time decision* made by an employer in isolation of their past experiences. The few studies that have invoked a longitudinal employer perspective have generally focused on employer learning regarding a specific individual employee and not how this experience may affect employment opportunities for subsequent other job applicants (Sterling 2012, Altonji and Pierret 2001).

To remedy this, I examine hiring as a *sequence of employer decisions over time*. I adopt an experiential learning perspective (Denrell and LeMens 2007, Erev and Roth 2014) to develop an account as to whether and how employers may learn to hire, and more interestingly, whether these past experiences result in 'better' choices for the employer. One perspective, originating in economics, holds that learning from past experiences provides useful information and should result in better or more accurate outcomes. For example, Altonji and Pierret (2001) demonstrated that employee wages may change over time to reflect their underlying cognitive abilities, which was unobservable to the employer at the time of hire. They point to this outcome as evidence of how employers learn to recognize ability and reward their employees appropriately. Conversely, an alternative, sociologically minded, perspective suggests it may be very difficult to develop an accurate assessments of potential employees. Stereotypes may be very difficult to alter (Fiske 1989) and underlying ability may not vary enough among social categories to usefully inform hiring decisions (Bielby and Baron 1986, Correll and Benard 2006). Against this backdrop, I ask: How does an employer's past hiring experiences with employees from a particular social group affect their future hiring decisions of other applicants from that social group? If employers do alter their hiring behavior, do their reactions following these experiences result in improvements reflected by hiring from 'better' social groups?

From a theoretical standpoint, examining how employer beliefs regarding members of social categories over time also provides insight into whether and how stereotypical beliefs may develop or change (Fiske 1998). To date, most work in hiring has assumed that beliefs regarding the correlation between ability and particular social categories are taken for granted but have not investigated how they may come about. The longitudinal perspective I provide is congruent with more recent attempts to adjudicate between a statistical and taste-based discrimination process (Farmer and Terrell 1996, Pager and Karafin 2009), which posit different outcomes over time (Rubineau and Kang 2012). In this vein, Pager and Karafin (2009) tackle the question of racial stereotypes in hiring with a qualitative study on whether employers alter their perspectives regarding the hiring of African Americans. Their interviews reveal that employers reportedly hire African American job seekers despite holding consistently negative views of the employability of members of that social category. While a valuable finding, their results reveal employer intentions, not actual hiring decisions, thereby leaving the question as to whether perceptions may change over time and how it may result in actual changes in hiring behavior.

I examine the hiring of freelancers in an online market for contract labor, Elance.com, which provides a platform connecting skilled workers willing to work on a temporary contract basis with those seeking to employ them. This setting overcomes several of the crucial obstacles that have prevented researchers from investigating these questions in the past. First, it affords a researcher repeated observations of employers hiring over time because all transactions are recorded by the platform. Second, whether a hiring experience was unambiguously positive or negative is indicated though a feedback mechanism directly from the employer. Finally, as a researcher, I have visibility into the set of applicants who apply to every job posted on the platform as well as who was ultimately hired.

In addition to the analytical leverage this platform affords, temporary contract work (Cappelli 1999, Barley and Kunda 2004, Leung 2014) and more generally non-standard work (Pedulla 2016) is a quickly growing, non-trivial labor force category. Recent estimates classify 15.8 percent of the U.S. labor force (23.6 million individuals) in 2015 as working in alternative employment situations, up from 10.1 percent in February 2005 (Katz and Krueger 2016). Figures that include those who hold full-time or part-time employment and contract on a temporary basis estimate this figure at closer to 53 Million individuals in the US alone (Freelancer Union 2015).

I examine how employers of contract workers make hiring decisions based off of a freelancer's country of origin, a social category explicitly protected, as race and gender are, in the US (NRC 2004, Rissing and Castilla 2014). Through longitudinal observations of over 16 Million applications from freelancers worldwide for over 2.2 Million job postings from 557,416 employers, I find that employers who have a bad experience with a freelancer are less likely to subsequently hire other freelancers from that same country. This is because negative experiences are losses that contradict expectations and therefore induce the employer to reconsider their decision. I run robustness checks to ally endogeneity concerns that negative experiences are correlated with or may endogenously reduce the quantity and quality of future applicants.

Furthermore, I find the negative hiring experiences affect employer hiring decisions more so on applicants for identical subsequent jobs than for dissimilar jobs. This result suggests a statistical account of hiring while being less congruent with a taste-based perspective. Most strikingly, I find that employer learning behavior seems poorly calibrated because they overreact to very small samples of negative experiences. Given the narrow differences in distribution of 'good' and 'bad' freelancers among the different countries, employers 'oversteer' by unnecessarily avoiding countries after a single bad experience. Analyses of actual switching behavior and results of a simulation provide evidence that employers do not hire from countries with better freelancers after a negative experience. I conclude by discussing the implications of my findings for discrimination, hiring, and employer learning.

Employer hiring and learning

Employers, when evaluating job applicants, face a fundamental challenge in asserting their suitability for employment. One way that employers attempt to ameliorate this uncertainty is through gathering additional information through signals or cues (Phelps 1972, Arrow 1972). This perspective begins with the fact that employers are often unable to determine the actual underlying ability and motivations of a job applicant. Instead, an employer will draw on their beliefs as to how an applicant's salient membership in a social group (such as their country of origin) may be correlated with desirable (but unobservable to the employer) characteristics, such as their motivation, commitment, or ambition. To reduce the risk of hiring a poor performing job candidate, employers will rely on these easier to detect social cues from which they believe they can infer ability. Employers therefore express preferences for applicants from particular social groups. For example, if a job requires heavy lifting, then an employer is more likely to hire a man, instead of a woman, for that job if they believe upper body strength is, on average, greater for men than women (Bielby and Baron 1986).

While most work in this vein has generally portrayed the beliefs regarding social group stereotypes held by employers as being fixed in time (Bielby and Baron 1986), recent work has begun to theorize on whether and how employers learn about individual employees over time, moving beyond theoretical models (Farmer and Terrell 1996, Oettinger 1996). In a study that attempts to identify whether employers incorporate experiences they have with individual employees into setting their wages, Altonji and Pierret (2001) demonstrate that employee wages begin to rely less on the education level of the worker (which was a noisy cue at hire) and more on individual characteristics, such as cognitive skill. Trial employment, in the form of internships (Sterling 2012, Baron and Kreps 1999) represents another way that employers may be able to reduce the inherent risk of hiring poor performers by learning of applicant ability before making a job offer. However, while this work has emphasized employer learning regarding an individual, it does not address how an employer's past experiences may alter their future hiring decisions regarding the social category membership of the applicant.

Focusing on learning regarding members of social groups, Pager and Karafin (2009) tackle the question of racial stereotypes in hiring with a qualitative study on whether employers alter their perspectives regarding hiring African Americans. The researcher's interviews reveal that employers continue to hire African American job seekers despite holding consistently negative views of the employability of members of that social category. Employers' justified their decisions on the basis that their employees are not representative of their more general African American population. While a valuable finding, this work was unable to examine the actual hiring outcomes and merely their intentions to hire. While employer learning clearly represents an important piece of the hiring puzzle, I am aware of no work that has quantitatively linked past employee hiring experiences to their subsequent hiring behaviors. To correct this, I explore whether and how employers utilize their past experiences of a job applicant's country of origin as a cue that alters their actual subsequent hiring decisions of other job applicants from that country.

Country of origin as employment cue

People hold stereotypical beliefs regarding individuals from certain countries, which in turn, affect their perceptions of their suitability for employment. Take, for example, the contrasting perceptions that Americans hold of immigrant job seekers from Latin America and Asia. On one hand, Burns and Gimple (2000) find that Americans may generally hold negative perceptions of Latino immigrants. These negative perceptions may also influence beliefs that recent Latin American immigrants are more likely to land in low-skilled jobs (Matoo et al 2007). On the other hand, beliefs that are held regarding Asian immigrants generally characterize them as being professional and well-educated (Ho 2003, Liu 1992). Lee and Fiske (2006) find that Asians are viewed as most competent, Canadians as moderately competent, and Mexican and Latino immigrants as possessing low competence.

As direct evidence of how these perceptions play into the employability of immigrants from particular countries, Rissing and Castilla (2014) investigate the US labor certification process, where immigrant visa seekers are vetted by the Department of Labor for suitability for employment. They find, in support of their contentions, that DoL reviewers were more likely to approve applicants from Asian countries and significantly less likely to approve of applicants from Latin American countries under conditions when detailed information was less available. The authors further explore how information and social category may inform one another by exploiting the fact that some of the applications were 'audited' thereby providing more detailed information to the DoL examiner for a subset of the cases. Presented with this quasi-experiment, in the greater information instances they found no differences in the approval rates of visa seekers, which suggested that the outcomes of approvals with less information were likely to be the result of statistical, as opposed to taste-based, discrimination.

Their study supports the idea that the country of origin of a job applicant figures prominently as a cue into employment outcomes, and encourages additional investigation. With regard to a dynamic learning model, because the DoL likely does not receive feedback regarding their decisions, it remains an open question as to whether auditors would have eventually learned to be less biased in the non-audited cases. Longitudinal knowledge of which auditors handled which applications could have aided in this regard. Second, the DoL is an institutionalized central authority and represents one particular gatekeeper in the hiring process. In this case, applicants to the visa process have already been ‘hired’ by employers, with the DoL acting as an additional screen. Employer learning and hiring preferences could be usefully advanced with a more specific examination of behaviors at the actual hiring interface (Peterson and Saporta 2004). Below, I develop hypotheses and detail a setting that advances our knowledge as to whether and how employers learn from their past hiring experiences.

Learning from negative hiring experiences

I frame hiring as a sequence of events occurring repeatedly over time for an employer. Drawing from an experiential learning perspective (Denrell and LeMens 2007, Erev and Roth 2014), I focus on employer reactions to negative hiring experiences they have with an employee from a particular county. Negative hiring events are theoretically important because, "a small number of negative experiences may hold especially strong weight in shaping attitudes" (Pager and Karafin 2009: 87 from Fiske (1998)). Therefore, learning is more likely to occur following a negative event than a positive or neutral one. Because losses loom larger than gains (Kahneman and Tversky 1979) a negative event is more likely to induce employers to make changes to how they hire employees than positive ones. Negative events also present disconfirming evidence while positive events are generally considered confirming evidence which will do little to alter a decision-maker’s original beliefs (Ross and Anderson 1982).

Because inferring the ability of a job candidate is not trivial, employers will be sensitive to cues as evidence of underlying quality. Incorporating this into their evaluation, employers making an offer of employment suggest they hold at least a neutral if not a positive view of that individual’s ability, and by extension, the ability of other employees in that particular social group[1]. A negative experience will therefore be contrary to what the employer had expected, and, because it is unexpected, will act as a

[1] The assumption that hiring someone suggests the employer likely believes that the individual will be competent is a reasonable, though not necessarily required assumption. For example, an employer may hold a poor belief regarding an employee, and others from a particular social group, or no belief at all, but hired them because of a lack of options. In this case, if the employer had a negative experience, they will still likely revise their beliefs and update further downward their expectations of other members from that social group.

diagnostically useful experience. For example, to the extent I hired an individual to work for me then I would be expected to hold some belief as to their ability to successfully complete the task. However, if they fail at the task, then I will incorporate that realization into how I choose to hire subsequent employees from that social group. Therefore, a negative event should serve to alter the belief they held at the point of hire[2].

Employers, in order to improve their hiring outcomes following a negative experience, will cast about to seek cues as to how a possible hiring mistake may be ameliorated. In evaluating a freelancer, individual measures of skill, such as education and past experiences, will be harder to dispute as they are quantifiable in indicators such as education or past experiences (Bills 1988, 2003; Bidwell 2016). Employer are unlikely to alter their beliefs regarding these indicators. However, employer perceptions of the link between unobservable quality and the social group membership of individuals will be easier to re-interpret. Social category membership, such as country of origin, necessarily suggests perceptions of similarity among those individuals, as ultimately, these are what stereotypes are (Fiske 1998). Therefore, a poor interaction with an employee will be an opportunity for the employer to update their beliefs regarding other potential employees from that employee's social category[3].

Specifically, a negative experience should dampen an employer's enthusiasm for hiring individuals from that particular social category again. This is because the *a priori* beliefs an employer held regarding the underlying skills of individuals from a certain social category are contradicted with a bad experience. Therefore, a negative hiring experience will decrease an employer's expectation of the mean skill level of other individuals from that social category. Membership in that social category for a subsequent job candidate becomes a predictor of poorer ability. As a result of this updating, an employer

[2] Conversely, one could theorize that if an employer held a negative belief regarding job applicants from a particular social category, then a positive experience with them may induce the employer to update positively their beliefs in subsequent hiring decisions, as predicted by contact theory (Allport 1954). However, this is a difficult prediction to test because a negatively held belief regarding freelancers from a particular country will discourage the employer from hiring anyone from that social category in the first place. Therefore, the employer would rarely be in a situation to positively update any negatively held beliefs. More competitive labor markets may alter an employer's preferences encouraging them to hire from less preferred social groups (Berker 1957), and suggests that taste and statistical discrimination are likely intertwined. I discuss the implications of this below.

[3] Positive experiences may also serve to further improve the beliefs an employer may hold of those freelancers from a particular country under a learning story. This will be tested but not formally stated because there are other reasons why positive experiences increase the likelihood of hiring more freelancers in the future – such as country-specific investments an employer may make.

will be less likely to hire again from that social group because the employer now associates more negative information with that cue. Therefore,

> *Hypothesis 1:* An employer who has a negative past experience with an employee from a particular country will be less likely to hire any applicant from that country in the future.

An employer's evaluation of an employee as a function of their country of origin is likely more nuanced than the dichotomous decision as to whether they are skilled or unskilled. A reasonable employer may hold a belief regarding employees from a particular country being more or less suited for a specific type of job. For example, beliefs regarding gender stereotypes lead to men being perceived as excelling in jobs that require math or science skills while women may be perceptually advantaged in jobs that require writing or communication skills (Correll 2001, 2004). More germane to our setting is the finding that Asians are stereotypically associated with "tech industry immigrants" underscores the idea that associations between particular countries and occupations may be reasonable (Lee and Fiske 2006).

Employers, after having a poor experience in hiring for a certain type of job, will believe that employees from that particular country are ill-suited to work on those types of jobs. However, this experience may not apply as strongly to hiring to other types of jobs. Learning scholars refer to this as contingent sampling (Gilboa and Schmeidler 1995, Gonzalez et al. 2003). The more relevant less that employers here would learn from bad experiences is to infer that employees from a certain country are less skilled in that specific task. However, some aspects of a negative experience should be applicable across jobs because the ability to perform well in a task is a combination of task specific skills as well as more general skills, such as punctuality or conscientiousness. Given this, then following a bad experience, I expect an employer will be less likely to hire other applicants from that country for any job, but more so for identical jobs.

> *Hypothesis 2:* Negative hiring experiences will have greater effect on the hiring of employees from that country for identical jobs than for different jobs in the future.

Evidence of employers differentially altering the decision to hire applicants from a country where they had a negative experience for identical versus dissimilar jobs provides me an analytic wedge to distinguish between a statistical versus a taste-based account of hiring. Demonstrating that an employer revises their beliefs regarding the suitability of applicants from a country for a particular job is congruent with a statistical perspective of hiring while being less supportive of an animosity mechanism. Animus will persists across job types because preference or hatred towards member of a social category, as a taste-based hiring perspective would hold, is a general phenomenon and not directly related to ability in a task. By differentially altering their hiring practices following a negative experience, an employer is demonstrating the malleability of their attempts to improve their hiring practices. Circumscribing their reaction to negative experiences by applying their updating narrowly is more congruent with a statistically (and informationally driven) discrimination process.

Accurate inferences are difficult to draw from very small samples. A purely rational employer would have to hire repeatedly for the same job within one country, in order to adequately come up with an inference of the abilities of other members of that country. In the case of hiring, having only a small number of observations to on is likely to be common because hiring is costly (Sterling 2012). Unfortunately, individuals do not account for this lack of visibility into the actual distribution of abilities and are comfortable drawing conclusions from very limited experiences (Erev and Roth 2014, Hertwig and Erev 2009). For example, restaurant goers may be comfortable concluding that they do not like a certain type of cuisine after only tasting it once or only trying one dish (Denrell and LeMens 2007).

More specific to hiring, there are at least two reasons for this behavior. Because of the required expense of hiring a large enough sample from a particular social group to gather enough evidence to understand the actual underlying distribution of ability, reacting to very few (or one) negative experience may be a necessity. Second, there likely little cost to being wrong about ones sampling and learning strategy because not hiring from a social category that has been labeled as ‘bad’ does not necessarily result in any penalties. Therefore, poor hiring decisions are not likely to be corrected because it may be simply be cost-less to be inaccurate. In essence, there's not risk to choosing not to hire employees from a particular social category if there are many others available to choose from[4].

This suggests that employers who react to negative hiring events by switching away from the country where the bad experience was encountered may not necessarily result in better outcomes. A better outcome, in this case, would be to hire from a social category that provides markedly better employees. There are several reasons for this. First, invariably, inferences made from a single negative hiring experience are unlikely to convey accurate information regarding the actual underlying ability of a whole social category of individuals, so decision are made prematurely. Second, the actual differences between social groups are likely much smaller than employers may suspect (Pager and Karafin 2009), which suggests that extreme reactions may not result in markedly better outcomes. Therefore, employers who react to negative hiring experiences, especially small samples, are likely to be overemphasizing the negative experience. By altering the hiring behaviors and avoiding the country from where they had a bad experience is unlikely to yield substantially better hiring. Specifically,

> *Hypothesis 3:* Employer reactions to negative hiring experiences does not lead to hiring from a better country.

There are several scope conditions to highlight. First, the employer needs to have some level of memory of the hiring event. Because learning occurs only when an experience is noted and then reflected upon the next time a similar event is encountered, employers who are unable to capitalize on this will be unable to utilize past experiences in subsequent decisions. For example, hiring decisions within an organization may not necessarily be made by the same individuals. In these situations, institutional memory may not persist in order to be able to be learned from. Even if an employer were willing to update their beliefs regarding the social group with which they had a negative hiring experience, they may not be able to. Second, the outcome of an employment decision has to be relatively clear in terms of whether it was a success or not. This may be difficult, for example, when employees work in teams or if

[4] This argument is the corollary Becker's (1957) contention that competitive labor markets (with fewer applicants) will drive out discrimination.

their success is less skill-based and more subject to market forces. The less ambiguous the hiring outcome, the clearer the link will be between a cue and ability, in the understanding developed by an employer. Third, the social category membership under investigation should be a visible characteristic during the hiring process. For example, sexual preference may not be a visible social category by which employers could utilize to infer skill during hiring, despite the fact that employers react negatively when attention is drawn to such a cue (Tilcsik 2011).

A global market for online freelancing

I seek evidence of employer learning with data from an online market for freelancing labor, Elance.com. This website serves as a platform for those seeking to hire freelancers to connect with individuals willing to work virtually on a contract basis. Contract workers are a quickly growing labor market role (Cappelli 1999, Barley and Kunda 2004) and highly skilled individuals are consistently faced with the decision as to whether or not to consider temporary employment (Bidwell and Briscoe 2009) particularly since long-term employment with an organization is declining in prevalence (Bidwell 2013). Estimates suggest that over 53 Million individuals in the United States have worked at one point in 2014 in a temporary contract basis (Freelancer Nation Report 2015). Elance, which, recently merged with oDesk to form UpWork, is the largest and oldest online platform in this domain. With over $2 Billion worth of transactions completed since Elance's founding in 1999, today over 200,000 jobs are posted monthly on this website by employers seeking to hire temporary skilled labor. There are over 4 million registered freelancers on the platform representing 223 countries[5].

I examine how an employer's experience with hiring employees from a particular country of origin affects their subsequent decisions to hire other workers from that country. The idea that a freelancer's country of origin may be correlated with their underlying ability is one that has received

[5] Though I refer to all job seekers on this website as freelancers, 67% of them are individuals while 33% of them are companies composed of several individuals. While I do not believe this should alter my approach or findings, I account for this in the analyses. Results, available from the author, run on solely the population of individual freelancers yields substantially identical results.

attention by employers of freelancers on online platforms. Despite the measures of skill that are provided by the website, inferring underlying freelancer ability from their country of origin is an active discussion topic. While an active topic of conversation by bloggers who discuss hiring strategies on virtual labor market platforms – there are contradictory recommendations. For example, William Shaker, a self-described “thirteen year online business veteran,” reports in a blog post titled, *How To Effectively Use Elance and Odesk To Outsource Quality People*, writes that one of his tips is to, “…mostly hire from India and the Philippines – these people have amazing work ethics and show up when you ask them to.” An in-depth discussion of whether certain freelancers from particular countries possess better skills appears on the website TimeDoctor.com. Under the blog post, *How to Outsource Anything Using 6 Top Outsourcing Website*, their advice is, "In general most of our experiences with workers from India, Bangladesh and Pakistan have been negative...[they] will often create really poor excuses for not reaching milestones..." They go on to comment on other countries as well, for example, "Former Soviet Union (Russia, Belarus, Ukraine etc) – a good place to find low cost, technically superb coders,” and freelancers from the, “"Philippines – incredibly diligent workers in all areas. It’s possible to find good writers, marketers, researchers and programmers." Finally, on the website hostadvice.com, the blog post, *Elance freelancers research 2015 - Which country is the best?*, summarizes my point well by publishing the results of a survey they conducted on which countries provide the best freelancers for certain jobs. While they caution that one should, “refrain from being racist, or attribute some stereotypes on one nation or another,” they nevertheless go on to conclude that, “nationality is a clear differentiator.” Along with a host of other results from their survey, they also provide advice such as, “Don’t hire a Canadian designer to do a German designer work,” and answer the question, “Do you need a good programmer?” by suggesting you should, “hire Argentinians and stay away from Australians.”

The hiring process

The hiring process begins with an employer filling out a job posting on the website. See Figure 1 for a sample. Job postings are organized by job categories that represent the spectrum of business tasks that can be accomplished virtually. The example project listed in Figure 1 is under the category of “architecture” which is a sub-category of the more general domain of “Engineering & Manufacturing” jobs. In all, there are 162 categories of jobs on Elance, examples of which include Logo Design, Website Programming, Legal Advice, and Voiceovers (a full list appears in the Appendix). Every posted job is required to be listed under one, and only one, of these categories. These categories serve to circumscribe similar jobs and to exclude dissimilar ones. While they may vary in how similar they are to one another, jobs in one category can be considered to be more similar to other jobs in the same category and different from jobs in other categories. These categories assist both freelancers in seeking work and employers in identify previous relevant experience of freelancers on the website (Leung 2014). As illustrated in Figure 1, job postings include basic information that describes the scope of the project, this information is structured by the website (not free form text) and includes budget requirements (i.e. <$500, $500-$1000, etc), timeframe expected for completion, and specific skills the employer is seeking. The job posting also includes a text description of the job details.

From the employer perspective, all bidders for each job are summarized by their country of origin and region on the right hand side of the job posting, as demonstrated in Figure 1. This is visible to an employer when they return to the job posting to see who has applied. For example, this particular job posting attracted 7 freelancers from North America and 11 from Eastern Europe and India/Southern Asia (as well as freelancers from other parts of the globe). This further underscores how salient the freelancer’s country of origin is to potential employers.

[Insert Figure 1 about here]

Once a job posting is published on the website, there are no barriers to applying and any freelancer is able to apply to any job. To apply, a freelancer submits a job proposal. Figure 2 depicts an abbreviated list of freelancers who applied to the job in Figure 1. At this summary level, an employer can see the freelancers profile photo, the price at which they are willing to complete the described tasks, along with a summary of the freelancers’ qualifications. Notice that the country of origin for each freelancer is also prominently displayed in the form of their country’s flag as well as the name of the country.

[Insert Figure 2 about here]

An employer is able to see a specific freelancer's profile in more detail by clicking on any of the job proposals that have been submitted. Once freelancers submit job proposals, the employer can review each freelancer's bid. Information regarding the price at which the freelancer is willing to complete the task and details regarding their background of past jobs and employee ratings are visible to an employer. See Figure 3 for an example of a detailed freelancer profile. The detailed freelancer profile also prominently displays their country of origin in the form of a flag and the name of their country right below their user name and next to their photo. The country of origin for each freelancer is not ambiguous on this website. Each freelancer is associated with one and only one country. This identifier is accurate as Elance authenticates a freelancer's identity through the use of bank records and their IP address.

On the right side of the freelancer proposal, are some statistics that point to the past successes of the freelancer. The Level Score is a numerical value that ranges from 0 to 13 (in our observation window) that represents a measure of quality as calculated by Elance. This measure, the algorithm is proprietary to Elance, takes into account information such as the number of jobs a freelancers has completed, the Star Rating they have received, the number of successful bids as a ratio of unsuccessful ones, as well as private feedback ratings. The statistics on the right side also include the number of jobs the freelancer has completed, the public star ratings (on a 1-5 scale) they have received from past employers, the number of clients they have had, and their total earnings to date.

[Insert Figure 3 about here]

After applications are submitted, an employer is free to choose to hire whomever they wish. If an employer chooses to hire anyone, 90% of the time they hire one freelancer, though employers are able to hire more. Once a freelancer is hired, employer and employee exchange job details and the job is completed virtually. Payment is conducted through the website.

Upon completion of the job, an employer is urged by Elance to submit feedback regarding the freelancer's performance on their job. Feedback is solicited right after an employer pays for the job and indicates it is complete. The website displays a screen asking the employer to rate the freelancer on the recently completed job. See Figure 4 for the feedback page each employer is asked to complete. The first question, "How likely are you to recommend this freelancer to a friend or a colleague?" allows responses from 1-10. Elance uses this question to gather private information regarding whether an employer had a good or poor experience with the freelancer they recently hired[6]. It is important to note that the answer to this question is hidden from the freelancer themselves and is only reported in aggregate once a freelancer has received over 10 reviews from their past employers. I use the fact that an employer can leave private feedback of this nature is what I use to identify whether an employer has had a positive or negative experience with a particular freelancer. I will return to this measure in more detail below when I detail my analytical strategy. The star ratings that a freelancer receives in their public profile and listed next to each completed job are from the subsequent questions, which are on a 1-5 scale. Freelancers receive an average of 4.2 stars on their jobs.

[Insert Figure 4 about here]

Data and Variables

I examined all activity on the website from its inception in late 1999 through early 2013. This encompassed 16,694,447 bids from 496,077 freelancers for 2,248,605 jobs posted by 557,416 different employers. Freelancers represent 223 countries and autonomous regions (such as Hong Kong) on this platform. The top five countries, in terms of number of freelancers, are the United States, India, Pakistan, Great Britain, and Canada. These five countries accounted for 75.24% of all bids placed on the website. A list of the top 35 countries, which represents almost 95% of all bids, is reported in the Appendix. The average employer posted 4.03 jobs. As of 2012, the average number of freelancers an employer has hired was 6.99. This distribution is skewed with a standard deviation of 27.8. 25% of employers have hired six or more freelancers.

[6] Functionality on this website has been changing. In the past, the employer was asked, "Would you consider hiring this freelancer again," with choices being "Yes", "Maybe" or, "No". However, an employer may like the performance of a freelancer but have no need to hire them again in the future. The question was reworded to ask about an employer's likelihood of recommending them to a "friend or colleague."

I focus on jobs where none of the freelancers that worked with the employer in the past applied. I am therefore estimating the likelihood that an employer will hire a freelancer among a set whom they have never employed before. Eliminating previously hired freelancers allows me to better isolate the informational effect of a freelancer's country of origin from information an employer may hold of a particular freelancer, which will be confounded. If a past employed freelancer applies to subsequent a job, a poor experience with that freelancer could merely reduce the likelihood of the employer hiring *that* freelancer again and not all freelancers from that same country. Including all freelancers will therefore obscure the effect of an employer's negative experience on *other* freelancers from that country[7].

Dependent variable

The dependent variable is the likelihood of a freelancer being hired for a job. For each of the bids that an employer received for their job posting, I coded the freelancer that was hired as '1' and those that were not as '0'. On average, each job received 12 job applications. The number of applicants ranged from 1 to 973 with a standard deviation of 13.

Independent variables

I used the response left by employers to the question as to whether they would recommend the freelancer to a friend or a colleague as an indicator of their positive or negative experience with that freelancer. As described above, responses to this question were stored in the form of a "Yes," "Maybe," or "No" and, more recently, from a 1-10 scale. Of the 157,240 responses posted by employers, 137,460 (87.4%) are "Yes", 9,432 (6%) are "Maybe", and 10,348 (6.6%) are "No." I considered the "No" responses to indicate the employer had a negative experience with that freelancer. Of the 1,143,921 responses posted by employers under the 1-10 scale, 14,017 (1.23%) were recorded as 1, indicating that an employer had a negative experience with that employee. Of note is that 959,498 (83.9%) of responses were 10, suggesting that most employers had positive experiences with the freelancers they hired.

[7] Models that include jobs with freelancers who have and have not worked for the employer in the past yield identical results.

I created a time varying measure for each employer of their cumulative number of negative experiences with freelancers from a specific country. The measure incremented each time a bad experience was recorded for each country. Across all employers, this ranged from 0 to 28 with an average of 0.04. This measure was logged as there is likely to be diminishing learning value to additional information regarding the ability of freelancers from a particular country. Results are identical if the non-logged measure was used in the regression models. Results are also consistent if the bad experience measure was dichotomized to =1 if there were *any* bad experiences and zero otherwise. A Chi-squared test of model fit was significantly worst with a dichotomous variable.

In order to test my second hypothesis, I created a variable that cumulates the poor experiences that an employer had with freelancers from a particular country for a particular job category type. Recall that all the jobs on this platform are categorized into one, and only one job category. These job categories designate generally recognized tasks, a list which is available in the Appendix. Because these job categories designate different jobs, then bad experiences should have a stronger effect on hiring other freelancers from that specific country for those specific jobs than for other, ostensibly different jobs.

Control variables

I included controls at the level of the bidder, the country they are from, and the employer. For the bidding freelancer, I included an indicator as to whether they are an individual or a company, a time-varying count (logged) of the number of previous jobs they had completed before placing their bid as a measure of experience; their time-varying Level score and time-varying Star ratings as measures of quality; the amount of their bid in USD (logged)[8], and an indicator as to whether they reside in the same country as the employer. In order to account for potential time-varying levels of competition and quality from a freelancer's fellow countrymen, I included the number of other bidders from this country for this job as a measure of the competition any individual bidder may face as it pertains to their fellow freelancers; I also included the time-varying average Level score of all freelancers from that country as a measure of overall

[8] I accounted for the 0's in these variables by adding 0.01 to all measures and logging the outcomes.

country quality bidding on the job at a particular time. At the level of the employer, I included the time-varying number of freelancers they had previously hired in total and from a specific country, as clearly there could be stickiness to hiring practices. Summary statistics and correlations appear in Table 1.

[Insert Table 1 about here]

Model and Results

I test whether and how past bad experiences with freelancers alters an employer's likelihood of hiring a freelancer from that country again by regressing the likelihood an employer will hire a freelancer from a particular country on the number of negative experiences an employer has had with freelancers from that particular country. Specifically, I estimate the following:

$$\text{Hired}_{ijkl} = \beta_1 * \text{NumberNegativeExperiences}_{lj} + \beta_2 * \text{FreelancerControls}_i + \beta_3 * \text{CountryControls}_k + \beta_4 * \text{EmployerControls}_l + \gamma_k + \varepsilon_k$$

Where the likelihood of freelancer 'i', from country 'j', applying for job 'k', listed by employer 'l' to be hired is a function of the number of negative experiences that employer 'l' has had with freelancers from country 'j' before posting job 'k'. I include vectors of controls for the freelancer, the country from which they reside, and the employer. Finally, γ_k represents job fixed-effects and ε_k the error term. Because the outcome is either '1' for winning the bid and '0' otherwise, I model this as a logistic regression[9]. I include job-fixed effects, thereby allying concerns that freelancers from different countries apply to different kinds of jobs. Any general differences between different types of jobs and different types of job categories are unlikely to drive the empirical patterns. The main source of identifying variation is between-country differences in the history of an employer getting a bad experience for a given job. I expect the estimate for β_1 to be negative, suggesting that employers are less likely to hire another freelancer if they had poor past experiences with others from that country.

[9] Probit and Linear Probability Models yield substantively identical results and available from the author.

Results

Results of the logistic regression predicting which freelancer was hired are reported in Table 2. Model 1 includes only the estimates of the control variables. Individuals are overall less likely than organizations to be hired. Applicants who have more experience as well as possessing a higher level score (indicators for experience and quality) are more likely to be hired. Star ratings exert a negative effect on hiring which seems counterintuitive. However, recall that most freelancers have a very high star rating, the average being 4.4 with a standard deviation of 0.4, which provides little variation. Also, level score includes star rating variation, so most of the advantage may be included in that measure. If a freelancer is from the same county as the employer, they are also more likely to be hired. Interestingly, the effect of the amount bid is positive and significant, suggesting that more expensive bids are more likely to be chosen. One thing to note is that many freelancers bid $0 in the hopes of either gaining experience and, hopefully, positive feedback; or to negotiate more bay directly with the employer. For these reasons, the lowest bidders are not normally hired as they are generally not experienced or may be untrustworthy. At the level of the country, the greater competition from other freelancers from an applicant's identical country, the less likely that freelancer will be chosen, suggesting that employers are sensitive to within country competition. The average level score for each country positively influences the likelihood of any freelancer from that country being hired. Finally, at the level of the employer, the more freelancers they have hired in the past, the more likely they will continue to hire. Also, the more freelancers they have hired from a particular country in the past, the more likely they are to continue to do so in the future.

[Insert Table 2 about here]

I included the cumulative measure of bad experiences that an employer has had with freelancers from each particular country in Model 2. The estimated effect of this variable is negative and significant (β=-0.1112, σ= 0.0041) demonstrating that employers do learn from previous experiences with freelancers by incorporating the negative experiences into their subsequent hiring decisions. Specifically, if an employer increases from no bad experiences to 1 bad hiring experience they are, on average 29% (1-exp$^{(-2.3 \times -0.1112)}$) less likely to hire an applicant from that country again, for any type of job[10].

I test hypothesis two in Model 3 in Table 2. In support of this more nuanced description of how a poor experience with a freelancer from a particular country affects hiring, I find that an employer is significantly and dramatically less likely to hire a freelancer from that country again when it is an identical job than for any other job. The coefficient for the specific country and job category is greater than the bad experiences with a country only effect (β=-0.8040 versus β=-0.0968). More concretely, increasing the number of bad experiences an employer has from 0 to 1, decreases their likelihood of hiring another freelancer from that country for that identical job by a dramatic 535% (1-exp$^{(-2.3 \times -0.8040)}$) versus the reduction in likelihood to hire individuals from that country for ANY job of 25% (1-exp$^{(-2.3 \times -0.0968)}$).

This result is congruent with a statistical discrimination story as employers incorporate learning regarding the ability of freelancers from a *particular country* and their ability to perform a *particular task*. This negative effect employers' exhibit of avoiding freelancers from a particular country is not completely eliminated for other types of jobs and suggests some of the negative experiences are applied more generally. This is reasonable, as some abilities are not task specific, such as timeliness and conscientiousness, so likely affect ability across a range of tasks. Jobs may also vary in how similar they are to one another, suggests some particular skills (or lack of) may transfer. Conversely, because animus would be expected to be applied more broadly across all types of jobs, this seems less congruent with a taste-based interpretation. I discuss this in the conclusion.

To understand whether updating yielded better hiring outcomes as outlined in my third hypothesis, I compared the countries where a negative employment outcome was experienced to countries employers hired from after this negative experience to identify whether observed employer behavior resulted in hiring from better countries. Conceptually, consider each country as a 'barrel' that is filled

[10] As mentioned above, using a dichotomous variable indicating any negative experience yields substantively identical results, so results are not driven by outliers of employers having many bad experiences.

with freelancers. Barrels will therefore vary in the percentage of 'good' and 'bad' job applicants, with better barrels filled with a greater percentage of good freelancers. An employer will improve their hiring decisions by avoiding barrels that have a greater percentage of 'bad' freelancers and employing from barrels that have a greater percentage of 'good' freelancers[11]. Employers infer which barrels are better or worst through their experiential learning process[12]. If avoiding a barrel where an employer had a negative experience is a useful strategy, then this action should lead to an employer hiring from a better barrel, defined as one with a greater percentage of good freelancers.

I first identified good versus bad freelancers among the complete population of job seekers who have completed at least one job on the website by labeling those who received *any* negative feedback as bad freelancers and those who have not received negative feedback as good freelancers[13]. I then calculated the percentage of good freelancers in each country and assigned this value to that country. The result of the distribution of percentage good freelancers is reported in Figure 5. Most freelancers are tightly clustered at the 98% mark. There is a long tail of freelancers who hail from the worst countries, but these are very small countries and the ratios reflect that. For example, the worst countries are Liberia and Guinea and consist of 50% bad freelancers, of which there are a total of six individuals. Freelancers from the largest countries such as the US and India have a higher proportion of good freelancers and are also at or close to the average. The United States is the modal country, with exactly 98.02% good freelancers.

[Insert Figure 5 about here]

I isolate the behaviors that employers made and identify the country where they encountered a negative hiring experience to the next country they hired from. Of the total 24,365 negative hiring experiences, 16,462 of them resulted in an employer changing the country they hired from on a

[11] Of course, barrels can represent any social category, such as gender or race as well.
[12] Employers may hold beliefs as to the distribution of good and bad types in each barrel, though this doesn't alter my findings. I discuss the implications of this below.
[13] This is a conservative way to identify good and bad types. One could argue there are even less bad freelancers with a less extreme way to distinguish them. This would only serve to strengthen my results because the distribution of skill among countries will further narrow.

subsequent job[14]. A test of the difference of means, reported in Model 1 in Table 3, of the percent of good freelancers between the first county with a negative hiring experience compared with the subsequent country they hired from yield no differences (Diff=2.0e-06, σ=0.0002, p=0.992). Employers who change countries do not subsequently hire from a country with a higher percentage of 'good' freelancers. Perhaps learning occurs gradually, so the first switch may not yield better results but subsequent ones may? In subsequent models, I partition the results between first switch and subsequent switches. Models 2 reports only the first switches, that is the initial bad experience an employer experienced; in Model 3, I examine the subsequent switches occurring after the first one[15]. Results in both Models 2 and 3 yield substantively identical results, employers switching countries from where they hire do not result in hiring from countries populated by a larger percentage of 'good' freelancers, whether after the first bad experience (Diff=1.5e-05, σ=0.0003, p=0.953) or after subsequent ones (Diff=-3.4e-05, σ=0.0003, p=0.908).

[Insert Table 3 about here]

Intuitively, switching may not be desirable because the difference in skill distribution among the different countries is narrow. If ability differed dramatically, and therefore were highly correlated by country, then switching may result in hiring from a substantively 'better' country if one initially hired from a particularly poor one. On the other hand, if the distribution of skills were narrower among countries, (say identical for all) then switching is unlikely to result in a substantively 'better' country, just a different one. I demonstrate this conceptually by simulating and then compare the pre- versus post-switching outcomes initiated by a negative experience using the actual population of freelancers.

As another way to demonstrate whether updating may or may not yield better hiring outcomes on this platform, I simulate hiring behavior and compare the pre- versus post- switching outcomes initiated by a negative experience using the actual population of freelancers. I structure the simulation by first modeling a draw of the country of origin of a negative hire. Utilizing the actual distribution of good versus bad freelancers reported in Figure 5, I randomly drew (hired) from this distribution of the observed

[14] A non-switch effectively yields an identical barrel to draw from, so if they were treated as 'switches' it would effectively be zero improvement.
[15] Model 3 switches are therefore conditional on employers having more than one bad experience.

population of all bad freelancers who have ever completed a job on the website[16]. The likelihood of an employer drawing a bad freelancer from a particular country is a function of how prevalent they are in the observed population, i.e. there are two bad freelancers from Guinea, which equates to a 0.02% likelihood that an employer will hire and have a bad experience with a freelancer from there (2 bad freelancers / 9950 total bad freelancers). On the other hand, an employer has a 36% chance of hiring and having a negative experience with a freelancer from the US (3637 bad US freelancers / 9950 total bad freelancers)[17]. The employer, upon having this bad experience, removes this country from subsequent consideration and then draws (hires) from the new distribution comprised of the remaining freelancers with those from the negative experience country removed[18]. I simulated 10,000 hires and switches.

[Insert Figure 6 about here]

Figure 6 shows the new estimated mean of the distributions from which the employer is now drawing, simulated 10,000 times. The top horizontal dashed-line is the mean percentage of good freelancers from all countries in the original total population (98.0%). As suggested by Figure 6, 10,000 simulated draws will result in subsequent distributions that have either equal lower means. The modal outcome is for an employer to subsequently draw from a distribution populated by freelancers with a mean of 96.7% good types. The intuition behind this finding is that the most likely scenario is for an employer's *first* hire to be from a big country, which, in this case, is comprised of mostly good freelancers. Therefore, eliminating this country from the subsequent choice set likely eliminates a good 'barrel' from subsequent consideration. The narrow distribution of remaining freelancers coupled with fact that there are many freelancers from worst countries remaining, leads to the observed outcome.

I also estimated the standard deviations of the realized distributions. Analyses, unreported for brevity, demonstrate that there is a greater likelihood that an employer who eliminates all the freelancers from the country in which they had a negative hiring experience will result in a distribution that exhibits

[16] This is a conservative population that is likely better than the whole population if those without any experience are included. I cannot estimate whether those freelancers with no experience are good or bad because they have never been evaluated.
[17] Note, I assume all freelancers are equally active and have an equal likelihood of being hired. This may actually overstate how active the bad freelancers are, as there is likely a survival bias towards being a good freelancer.
[18] I recognize employers do not always switch countries. Analyses, not reported for brevity, regarding the sensitivity of this assumption does not substantively alter the conclusions.

greater variance than the initial one they drew from, which include all freelancers. Specifically, of the 10,000 simulated draws, 7527 resulted in a greater standard deviation from the original distribution and the remaining 2473 were equal or less. Intuitively, as there is a greater likelihood an employer will eliminate a big country from somewhere near the peak and center of the original distribution, the variance of the remaining freelancer population from which to draw from likely increases.

[Insert Figure 7 about here]

Figure 7 reports the distribution of outcomes of the percent of good types in the country the employer subsequently switches to. The modal outcome is to subsequently hire from a country that has the same mean percentage of good types, 98%, as originally encountered. Another way to view these results is to simply categorize outcomes of the new, switched to country, into categorically better, same, or worst outcomes compared to the percentage of good freelancers from the previous country of hire. A good outcome is defined as a hire from a country with a greater percentage of good types and a bad outcome being a hire from a country with a lower percentage of good types than previously realized. Results, unreported for brevity, show that of the 10,000 switches following a bad experience, the most likely outcome is to not improve the country one hires from (4198/10,000, 42%) and the second most likely outcome is to end up drawing from a worst country (3226/10,000, 32%). In sum, given the distribution of the observable good and bad freelancers and the countries they are from, switching after a negative hiring experience is more likely to result in either no improvement or a worst outcome. Employers who instead decide to continue to employ from the same country do not risk a worst outcome.

Addressing alternative explanations

Are employers learning about the employees from those countries or learning about their own ability to interact with employees from those countries? This is a global market and the differences among freelancers from different countries are not limited to their perceived ability in certain tasks, but also in other ways that affect a working relationship. For example, cultures differ in their perceptions of how ‘time’ is conceived with some countries having a much more literal and precise interpretation while others having a more flexible one (Levine 1997). Country differences also reflect other concrete differences that affect a working relationship, such as the language spoken and the time zone they are in. Perhaps an employer may believe they speak Spanish well enough to hire a freelancer from Latin America (or Portuguese in the case of Brazil), but upon hiring from there realize they do not. If these differences account for difficulties in how some freelancers and employers work together, then the switch out of these countries could be a function of an employer learning they are not able to work with employees from a certain country and not due to the perceived skill of freelancers from that country.

I included three additional control variables to account for the possibility that employers may be switching due to incompatibilities. First, utilizing Hofstede's (2001) well-known measure of cultural distance, I included the variables of Power Distance and Individualism as controls for how similar or different other countries are to the country where the negative experience was realized. I choose these two variables because they are most central to the dynamics of a working relationship. Power Distance reflects how individuals in a society accept or reject a hierarchical order and Individualism reflects ones preference for or against how much individuals are expected to be responsible for taking care of one another – outside of one's immediate family. To the extent we believe that the negative experience was precipitated by difference in ability to work with one another, controlling for other countries that are similar along these dimensions should account for this. Results, unreported for brevity, continue to demonstrate a lower likelihood of hiring from the same country despite controls for other similar countries[19], suggesting this is not what employers are reacting to.

Second, to account for the switch away from a particular country being the result of language incompatibilities, I included an indicator for those countries that have an identical official or co-official language to the one where the negative hiring experience occurred. For example, if the negative experience was from a freelancer who resided in Canada, then the US and India as well as France would receive an indicator value of 1 because both French and English is spoken in Canada. Including this

[19] Interestingly, there is a positive effect on similarity, suggesting possible cultural spillover effects, which are beyond the scope of this paper.

variable does not alter the results. I also addressed this by running the regressions on only a subset of English speaking countries[20]. This negates any language effect, as both employers and freelancers would be speaking the same official or co-official language. Results, unreported for brevity, remain the same.

Finally, I account for whether the time zone had a hand in the likelihood of switching by including an indicator for whether the subsequent job applicants share a time zone with the country where the employer previously had a negative experience. Results, unreported for brevity, continue to demonstrate that regardless of time zone compatibility, employers are less likely to hire from the country where they had a poor experience, over and above other countries in identical time zones.

Are these results endogenous to the employer's negative hiring experience? The action of leaving negative feedback for a past employee could cause structural changes in the freelancers who apply to an employer's subsequent jobs. This would suggest the results are not a function of learning, but rather changes in opportunity to hire. For example, employers could receive a lower number or lower quality of applicants from the disparaged country. If it becomes known to other freelancers from a country that an employer had a bad experience with other freelancers from that country, it may decrease the likelihood of those freelancers from applying to work with that employer in the future.

Recall that the "would work with again" ratings were private, but star ratings are public. To the extent an employer had a bad experience with a freelancer – they could decide to leave a poor public rating as well as poor private ones[21]. Functionality on the website allows other freelancers to see how an employer rated their past experiences with freelancers. If an excessively poor rating was received by a freelancer – other freelancers from that country could perceive this as a possible risk to working with this employer in the future. This may reduce their likelihood of applying to work for this employer. In short, an employer may receive fewer applicants from a maligned country or lower quality ones from more

---

[20] The top eight employing countries represent 77.8% of all jobs on the platform and all have English as an official or co-official language.

[21] Analyses demonstrate that private negative ratings result in more negative public ones. For example, those freelancers receiving a 'yes' to the question "would you hire them again" received an average public star rating of 4.8/5.0. Those freelancers who received a 'no' to the same question received an average public star rating of 2.2/5.0.

desperate freelancers. Reducing the number of applicants or their overall quality from a particular country would reduce the likelihood that anyone from that country would be employed in the future[22].

My main regressions above included controls for the number of applicants from each country as well as the overall quality of applicants from that country, which should guard against this concern. However, to further allay concerns, I investigate whether a poor review by an employer negatively affects the overall observable quantity and quality of the applicants from that country. I estimate two dependent variables. First, as a quantity measure, I estimate the number of applicants from a particular country for a particular job. Second, as a measure of quality, I estimate the average level score of all applicants from a particular country who applied to a particular job. The independent variable of interest is the number of bad experience an employer had with freelancers from that country. I modeled the count of job applicants by country with a Poisson regression and I models the level score dependent variable with an OLS[23]. Results of the within-employer fixed-effects regressions are reported in Table 4 below.

[Insert Table 4 about here]

Model 1 estimated the effect of an employer having a bad experience with a freelancer on the number of freelancers from that country who apply to subsequently listed jobs. The coefficient is positive, which suggests that employers receive a greater number of applicants from a country with which they had a poor hiring experience from. This may seem counter intuitive, but recall that an employer could leave positive public feedback yet have had a poor experience and leave negative private feedback. In these instances, merely hiring a freelancer will likely induce other freelancers from that country to apply to subsequent jobs.

---

[22] A related outcome could be that the employer explicitly discourages those from a particular country from applying to their job posting. For example, employers on the website have been known to express particular preferences for freelancers from a particular country by stating this explicitly in their job postings. However, to the extent this is occurring – it only serves to confirm my contention that employers 'learn' from their past poor experiences by discouraging freelancers from a particular country from applying. Regardless, my robustness check accounts for this behavior.

[23] Employers do not receive applicants from all countries for all jobs. Because there are so many countries, it is more likely an employer received no applications from freelancers in smaller countries for a job posting. To guard against this, I also modeled this as a logistic regression with 1 indicating whether ANY freelancer from a particular country applied and a 0 indicating no freelancer from that country applied. Conditional on the employer receiving any applicants from that country in past job postings, results, unreported for brevity, demonstrate there is still no effect of negative ratings on subsequent applicants from that country.

In Model 2 I test whether an employer who has had a bad experience with a particular freelancer reduces the subsequent quality of bidders from that country. The quality of all freelancers from a particular country is a function of the feedback that an employer leaves for freelancers from that country. To ensure my measure of overall quality from a country is not sensitive to this, I calculate the average of all bidders from a particular country, excluding the ones that the employer has hired in the past. Results in Model 2 demonstrate that there is no significant effect of having a poor experience with a freelancer and the overall quality of subsequent bidders from that country.

Countries vary by quality in ways that may not be completely observable to a researcher, but observable to an employer – therefore affecting the results. For example, poor grammar or spelling in their job applications could signal poorer freelancer skills in a particular country. The observed results could also be driven by predominantly by large countries. I address these concerns by including dummy indicators for the top 20 countries, which encompasses 92% of all bidding activity in this market, as a country fixed-effects specification. The variation we observe will be within-country and takes into account any time invariant heterogeneity among countries. Results of a within job posting logistic regression of the likelihood that a freelancer is hired are reported in Table 5. The US, being the largest, is the omitted country. Note the substantial and significant variation in the effect of the country of origin on the likelihood of a freelancer being hired. More importantly, in continued support of my original finding, the coefficient estimating the effect of previous poor experiences with a freelancer continues to exert a negative and significant effect on the likelihood that an employer will hire from that country again.

[Insert Table 5 about here]

Readers may be concerned that employers who had a negative experience differ from those that did not. For example, employers without a negative hiring experience are significantly less active than those that have. I account for this in two ways. First, I created a matched sample with the CEM procedure in Stata 12.0 and ran identical models reported in Table 2. Specifically, I matched employers by the total number of jobs they have hired for on the website, the total number of different job categories they have hired across, and also the total number of different countries – thereby more stringently controlling for the size of the employer, the breadth of their hiring needs, and the variety of freelancers they have hired in the past. Results unreported for brevity continue to yield identical results. Second, instead of utilizing the job-fixed effects specification, I estimated models with employer fixed-effects. This eliminates from consideration those employers who have not had a negative hiring experience but allows me to control for the time-invariant differences among employers. Unreported results demonstrate that the effect continues to persist in the employer-fixed effects regression. Increasingly negative hiring experiences continue to exert a significant and negative effect on the likelihood and employer will hire another freelancer from that country in the future.

If the measure of the number of negative experiences with a particular country serves a cue of perceived quality, then other indicators of freelancer quality should serve to offset the negative effect of this measure. Therefore, better measures of objective quality or experience should overcome the negative cue that a freelancer's country of origin signals because there should be employers account for these compensatory differences. Demonstration of these effects should further bolster the idea that a freelancer's country of origin is acting as a cue of difficult to observe quality. In analyses, unreported for brevity, I included the interactions with two objective measures of freelancer quality – the level score and the number of past job experiences. Coefficients of the interactions of these measures with the measure of negative experiences with the freelancer's country, as expected, were positive and significant. This suggests that possessing better objective measures of quality serves to directly offset the potentially negative cue that is conveyed by ones country of origin.

Finally, in additional support of the idea that this is an instance of learning from experience, I test the effect of the saliences of the job(s) which received negative experiences as well as how recent the negative event occurred. I included a variable of the logged, cumulative amount paid to freelancers where an employer noted a negative hiring experience. I reason the more 'expensive' the negative experience, the more salient lesson learned. I also included a variable indicating how many days past the negative hiring event the subsequent job was posted and interacted this with the negative event. Results reported in Table 6 confirm the fact that more expensive jobs exert a greater negative effect and the longer past the negative event, the less of an effect it exerts. Of note is that the main effect of the count of the number of bad experiences continues to remain negative and significant – suggesting that any bad experience is a learning experience, but that costly and recent ones have an additive effect.

[Insert Table 6 about here]

Discussion and Conclusion

I presented a cognitive perspective of hiring by framing it as an experiential learning process. My investigation revealed that employers update their beliefs regarding job applicants from particular social groups, in my case country of origin, from their previous negative hiring experiences. Most interestingly, reactions that employers adopt do not seem to result in 'better' outcomes – an example of the limits to employer's cognitive hiring abilities. To my knowledge, there is virtually no work that has examined hiring behaviors as a longitudinal experiential learning process. Furthermore, because the opportunity for discrimination is believed to be highest at the point of the hiring decision (Peterson and Saporta 2004, Fernandez and Greenberg 2013) the ability to observe actual employer behaviors at this point in the process is a valuable step forward in the hiring literature. Examining behaviors at the point of hire in an online setting such as Elance is also valuable because I isolate the decision, and therefore the learning process, to an individual employer. These hiring decisions are not the product of a committee or subject to corporate policies as organizational investigations may be reflecting. Finally, because I have access to the full pool of applicants, I am able to assuage criticisms that applicant sorting behaviors may alter the outcomes of hiring decisions (Pager and Pedulla 2015).

These finding inform several facets of hiring research. First, a stream of work has documented widespread preferential and discriminatory hiring practices by employers due to a job applicant's social category, such as disadvantaged outcomes for African American and Latino job seekers relative to white applicants (Pager et al 2009), the fact that women are disadvantaged in the labor market relative to men (Jacobs 1989, Reskin and Roos 1990, Peterson and Morgan 1995) even for jobs that do not require firm specific skills (Fernandez-Mateo 2009); or how openly gay men are less likely to receive callbacks for interviews (Tilcsik 2011); and, most recently, immigrant workers face differential likelihoods of getting green card approvals depending on their country of origin (Rissing and Castilla 2014). The longitudinal perspective I adopted here introduces a useful way to address the debate as to whether these outcomes are the result of animus or more cognitively motivated processes because the two mechanisms, under certain conditions, predict different outcomes over time (Rubineau and King 2012). For example, while I do not have information on the race or gender of a freelancer here[24], I believe that this longitudinal perspective can aid us in understanding whether impressions of applicants from these more historically relevant social categories may be altered by learning. My findings suggest stereotypical beliefs may be malleable and points to practical implications because it can inform policy remedies aimed at reducing recognized disparity in hiring by social category membership (Bielby and Baron 1986, Fernandez and Greenberg 2013). Recent work points to potential improvements in the labor market outcomes for women. Perhaps the introduction of affirmative action policies that induce experiences with less traditional hires suggests that even animus can be reduced through positive experiences (Allport 1954).

Second, this paper also highlights the limits to employer learning which also usefully contributes to whether discrimination can be simply partitioned into either statistical or taste-based processes because we need to grapple with the fact that statistical discrimination may be self-perpetuating. An implication of my findings is that employer learning is limited to the breadth of experiences they wish to undertake – in this case, they can only learn about whom they hire, not those they don't. If an employer holds a, possibly erroneous, belief that individuals from a particular social category are 'better' and they hire based on that belief, the employer may never alter their belief. In this case, employers would never 'learn' that there could be a better pool of applicants as they would never have a reason to hire and therefore experience them. This limited learning process could explain the extreme sex segregation observed by Bielby and Baron (1986). Also, because today's outcomes may affect the incentives for tomorrow's behavior (Lundberg and Startz 1983, Coate and Loury 1993) the poor statistical story of discriminatory hiring may

[24] These could be gleaned from the freelancer's profile photo, which is a highly labor intensive task that is being undertaken by the author for future research and beyond the scope of this paper.

be endogenous and sresult in highly unequal outcomes among individuals from different social categories even if they possess identical underlying abilities. For example, under this scenario, there would be little reason for other members of those maligned social groups to acquire the appropriate education or skills to seek out those positions.

Third, my paper also clearly contributes to the more nascent stream of work that has highlighted the increase in temporary contract work, or the “gig” economy, as a consequential labor market category of the workforce (Cappelli 1999, Barley and Kunda 2004, Katz and Krueger 2016) and more specifically, the impact of virtual platforms on both workers and employers. Temporary contract work induces employees to be constantly searcher for new jobs to apply to. An understanding into how employees may process these applications and the perspectives they may bring to evaluating them provide us with insight into the challenges a large temporary contract workforce may face. For example, the paradigm that complete transparency will allow employers to make the ‘best’ decisions may be called into question by my findings because employers do not seem to ‘use’ the information to their best outcomes. Instead, this could suggest that guided application matching on the part of the virtual platform to better identify appropriate workers though the use of algorithms may be a viable point to investigate.

Fourth, as the labor market becomes increasingly global, an understanding of the dynamics of hiring as a function of the diversity in the labor market comes to the fore. Though not as well-established as the work on discriminatory hiring due to race and gender, the nascent work on how employers view and respond to job applicants from different countries will continue to increase in importance (Rissing and Castilla 2014). Relatedly, cognitive psychologists are only beginning to develop an understanding as to how stereotypes of individuals from particular regions develop (Lee and Fiske 2006). The potential spillover effects of the experiences employers have with employees from a particular nation and how it may affect beliefs regarding a region are valuable to consider.

References

Allport, G. W. 1954. *The nature of prejudice*. Cambridge, MA: Perseus Books

Altonji, Joseph G., and Charles R. Pierret. 2001. Employer learning and statistical discrimination. *Quarterly Journal of Economics* 116:313-50.

Arrow, Kenneth J. 1972. Models of job discrimination. In *Racial discrimination in economic life*, ed. Anthony H. Pascal, 83-102. Lexington, MA: D.C. Heath.

Barley, Stephen R. and Gideon Kunda. 2004. *Gurus, Hired Guns and Warm Bodies: Itinerant Experts in a Knowledge Economy*. Princeton, NJ: Princeton University Press.

Baron, JN and DM Kreps. 1999. Strategic Human Resources: Frameworks for General Managers (John Wiley & Sons, New York).

Becker, G. S. 1971. *The economics of discrimination*. Chicago, IL: University of Chicago Press.

Bidwell, M.J. 2016. Unpacking Human Capital: Exploring the Role of Experience and Education in Shaping Labor Market Opportunity. Working Paper, University of Pennsylvania.

Bidwell, M.J. 2013. What happened to long term employment? The role of worker power and environmental turbulence in explaining declines in worker tenure. *Organization Science*, 24, 4, 1061.

Bidwell, M.J. and Briscoe, F.S. 2009. Who contracts? Determinants of the decision to work as an independent contractor among information technology workers. *Academy of Management Journal*, 52(6), 1148–1168.

Bielby, William, and James N. Baron. 1986. Men and women at work: Sex segregation, and statistical discrimination. *American Journal of Sociology* 91:759-99.

Bills, David B. 1988. Educational Credentials and Promotions: Does Schooling do More than Get You in the Door? *Sociology of Education*, Vol. 61, No. 1, pp. 52-60.

Bills, David B. 2003. Credentials, Signals, and Screens: Explaining the Relationship between Schooling and Job Assignment. *Review of Educational Research*, Vol. 73, No. 4, pp. 441-469.

Borjas, George J. 1987. "Self-Selection and the Earnings of Immigrants." *American Economic Review,* 77(4):531–53.

Burns, Peter and James G. Gimpel. 2000. "Economic Insecurity, Prejudicial Stereotypes, and Public Opinion on Immigration Policy." *Political Science Quarterly,* 115(2):201–225.

Cappelli, P. 1999. *The new deal at work: Managing the market-driven workforce*. Boston: Harvard Business School Press.

Chiswick, Barry R. 1978. "The Effect of Americanization on the Earnings of Foreign-Born Men." *Journal of Political Economy* 86(5):897–921.

Coate, Stephen, and Glenn C. Loury, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, 83 (1993), 1220–1240.

Correll, S. J. 2001. Gender and the career choice process: The role of biased self-assessments. *American Journal of Sociology, 106,* 1691-1730.

Correll, S. J. 2004. Constraints into preferences: Gender, status and emerging career aspirations. *American Sociological Review, 69,* 93-133.

Correll, S. J., and Benard, S. 2006. Biased estimators? Comparing status and statistical theories of gender discrimination. In S. R. Thye & E. J. Lawler (Eds.), *Advances in group processes*. Vol. 23, pp. 89–116. New York, NY: Elsevier Science.

Denrell, J. and G. Le Mens (2007). "Interdependent Sampling and Social Influence." *Psychological Review*, 114 (2): 398-422.

England, Paula, Melissa Herbert, Barbara Stanek Kilbourne, Lori Reid, and Lori McCreary Megdal. 1994. “The Gendered Valuation of Occupation and Skills: Earnings in the 1980 Census Occupations.” *Social Forces*, 73:65–98.

Erev, I. and A. Roth 2014 Maximization, learning, and economic behavior. *PNAS*, vol. 111, suppl. 3, p. 10825.

Estes, WK 1950. Toward a statistical theory of learning. *Psychol Rev* 57(2):94–107.

Farmer, Amy, and Dek Terrell. 1996. Discrimination, Bayesian updating of employer beliefs and human capital accumulation. *Economic Inquiry* 34:204-19.

Fernandez, Roberto M. and Jason Greenberg. 2013. “Race, Gender and Statistical Discrimination.” *Research in the Sociology of Work* 24:81–102.

Fernandez, Roberto M. and Marie Louise Mors. 2008. Competing for jobs: Labor queues and gender sorting in the hiring process. *Social Science Research* 37: 1061–1080.

Fernandez-Mateo, Isabel. 2009. Cumulative Gender Disadvantage in Contract Employment. *American Journal of Sociology*, Vol. 114, No. 4, pp. 871-923.

Freelancer Union 2015. *Freelancing in America Report*. October, 10, 2015.

Gilboa, I. and Schmeidler, D. 1995. Case-based decision theory. *Q. J. Econ*. 30, 605–639

Gonzalez, C. et al. 2003. Instance-based learning in real-time dynamic decision making. *Cogn. Sci.* 27, 591–635.

Hertwig R and Erev I 2009. The description-experience gap in risky choice. *Trends Cogn Sci* 13(12):517–523.

Ho, Pensri. 2003. “Performing the ‘Oriental’: Professionals and the Asian Model Minority Myth.” *Journal of Asian American Studies* 6(2):149–75.

Hofstede, Geert 2001. *Culture’s Consequences: Comparing Values, Behaviors, Institutions, and Organizations Across Nations*. Second Edition, Thousand Oaks CA: Sage Publications.

http://hostadvice.com/blog/elance-freelancers-research-2015/ Accessed March 22, 2016

http://www.williamshaker.com/how-to-effectively-use-elance-and-odesk-to-outsource-quality-people/ Accessed March 22, 2016

http://blog.timedoctor.com/2011/02/22/the-top-6-outsourcing-sites-and-how-to-use-them Accessed March 22, 2016

Jacobs, Jerry A. 1989. "Long Term Trends in Occupational Segregation by Sex." *American Journal of Sociology*, 95(1):160-73.

Kalev, A., F. Dobbin, and E. Kelly. 2006. “Best practices or best guesses? Assessing the efficacy of corporate affirmative action and diversity policies.” *American Sociological Review*, 71: 589-617.

Kahneman, Daniel and Amos Tversky. 1979. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), pp. 263-291.

Kaiser, C. R., B. Major, I. Jurcevic, T. L. Dover, L. M. Brady, and J. R. Shapiro. 2013. “Presumed fair: Ironic effects of organizational diversity structures.” *Journal of Personality and Social Psychology*, 104: 504-519.

Katz, L.F. and A.B. Krueger 2016. The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015. NBER Working paper.

Lee, Tiane L. and Susan T. Fiske. 2006. “Not an Outgroup, Not Yet an Ingroup: Immigrants in the Stereotype Content Model.” *International Journal of Intercultural Relations* 30(6):751–68.

Levine, Robert. 1997. *A Geography of Time: the temporal misadventures of a social psychologist, or how every culture keeps time just a little bit differently*. New York : BasicBooks.

Liu, John M. 1992. “The Contours of Asian Professional, Technical and Kindred Work Immigration, 1965–1988.” *Sociological Perspectives* 35(4):673–704.

Lundberg, Shelly J., and Richard Startz. 1983. "Private Discrimination and Social Intervention in Competitive Labor Markets." *American Economic Review*, 340-47.

Mattoo, Aaditya, Ileana Cristina Neagu, and Caglar. Ozden. 2007. “Brain Waste? Educated Immigrants in the US Labor Market.” *Journal of Development Economics* 87(2):255–69.

National Research Council. (2004). Measuring racial discrimination. Panel on methods for assessing discrimination. In R. M. Blank, M. Dabady & C. F. Citro (Eds.), Committee on national statistics, division of behavioral and social sciences and education. Washington, DC: The National Academies Press.

Pager, Devah, and Diana Karafin. 2009. “Bayesian Bigot? Statistical Discrimination, Stereotypes, and Employer Decision Making”. *Annals of the American Academy of Political and Social Sciences* 621 (January):70-93.

Pager, Devah and David S. Pedulla. 2015. Race, Self-Selection, and the Job Search Process. *American Journal of Sociology*, Vol. 120, No. 4, pp. 1005-1054.

Pager, D., and Shepherd, H. 2008. The sociology of discrimination: Racial discrimination in employment, housing, credit, and consumer markets. *Annual Review of Sociology*, 34, 181–209.

Pager, Devah, Bruce Western, and Bart Bonikowski. 2009. Discrimination in a Low-Wage Labor Market: A Field Experiment. American Sociological Review, Vol. 74, October: 777–799.

Pedulla, D.S. 2016. Penalized or Protected? Gender and the Consequences of Nonstandard and Mismatched Employment Histories. *American Sociological Review*, Vol. 81(2): 262–289.

Petersen, Trond and Laurie A. Morgan. 1995. Separate and Unequal: Occupation-Establishment Sex Segregation and the Gender Wage Gap. *American Journal of Sociology*, Vol. 101, No. 2, pp. 329-365.

Petersen, T., and Saporta, I. 2004. The opportunity structure for discrimination. *American Journal of Sociology*, 109, 852–901.

Phelps, Edmund. 1972. The statistical theory of racism and sexism. *American Economic Review* 62:659-61.

Reskin, Barbara F., and Patricia A. Roos. 1990. *Job Queues, Gender Queues: Explaining Women’s Inroads into Male Occupations*. Philadelphia: Temple University Press.

Rissing, B.A. and E.J. Castilla 2014. House of Green Cards: Statistical or Preference-Based Inequality in the Employment of Foreign Nationals. *American Sociological Review*, Vol. 79(6) 1226–1255.

Ross, Lee; Anderson, Craig A. 1982. "Shortcomings in the attribution process: On the origins and maintenance of erroneous social assessments", in Kahneman, Daniel; Slovic, Paul; Tversky, Amos, *Judgment under uncertainty: Heuristics and biases*, Cambridge University Press, pp. 129–152.

Rubineau, B., and Kang, Y. 2012. Bias in white: A longitudinal natural experiment measuring changes in discrimination. *Management Science*, 58, 660–677.

Sterling, Adina D. 2014. Friendships and Search Behavior in Labor Markets. *Management Science* 60(9):2341-2354.

Tilcsik, András. 2011. Pride and Prejudice: Employment Discrimination against Openly Gay Men in the United States. *American Journal of Sociology*, 117: 2, 586-626.

Tubergen, Frank van, Ineke Maas, and Henk Flap. 2004. “The Economic Incorporation of Immigrants in 18 Western Societies: Origin, Destination, and Community Effects.” *American Sociological Review*, 69(5):704–727.

FIGURE 1

JOB POSTING

Design plans for Residential Exterior Front Patio

Help

Engineering & Manufacturing > Architecture

Job Actions | Edit Job & Settings | Share

Job Status: Selecting Candidate
Budget: Less than $500
Posted: Thu, Mar 5, 2015
Fixed Price Job
Location: Anywhere
Elance Escrow Protection
Start: Immediately
W9 Not Required
My Team (1)
30 invited

We would like to re-design the front entrance area of our house. We have pictures and some rough measurements and some ideas. We need someone to help us with design renderings and drawings of what you may propose and ultimately a set of drawings that we can build from.

Currently, the front of the house is hidden from view from the sidewalk by a brick wall. Behind the wall is a raised area of dirt and plants. We would like to get rid of ALL the dirt, flatten out the front and make it a patio. We also would like to remove the brick wall and replace it with a fence so the house can be seen better from the street.

We would like the patio to be level with the Double Doors. We would need to have a retaining wall at least up to the level of the sidewalk - which is approximately 3 feet above the bottom of the double doors. The tree in front of the front door needs to stay there - which may mean it needs to be boxed in. Everything else can go.

I attach several pictures of the house as well as some basic measurements. Here's a link to some pictures of patios that we like:

https://www.pinterest.com/itsazooatthezoo/patio/

In your proposal, please answer the following questions:
1. What programs would you use to produce your drawings?
2. What types of drawings or rendering will you deliver?
3. What is the process by which you work on these projects remotely?

Job ID: 69720721

44 total proposals from 7 regions.

Proposals (44)

7 North America
11 India/Southern Asia
3 Africa
2 Eastern Asia
11 Eastern Europe
5 Western Europe
5 Central & South America

FIGURE 2
JOB APPLICANT LIST

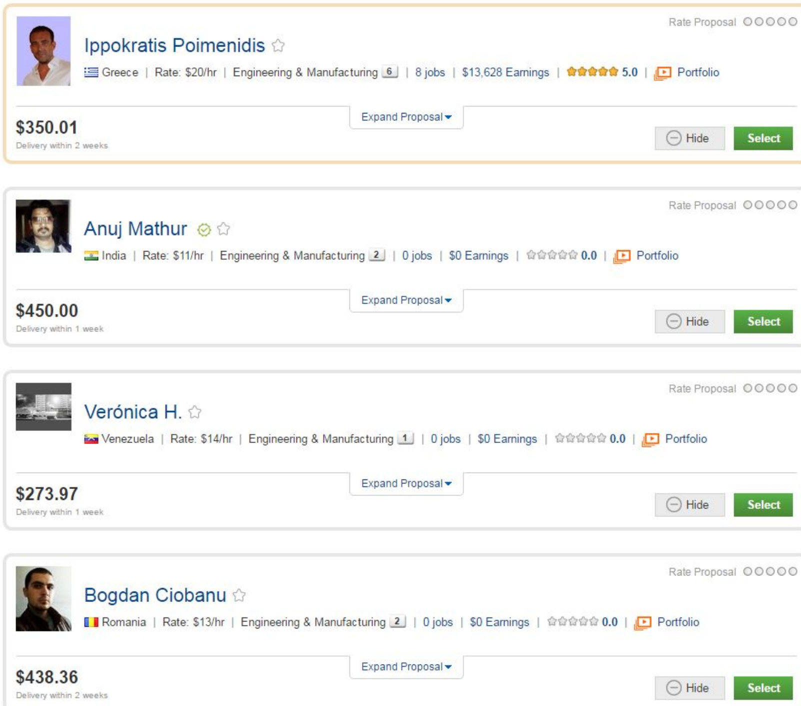

FIGURE 3
FREELANCER PROFILE PAGE

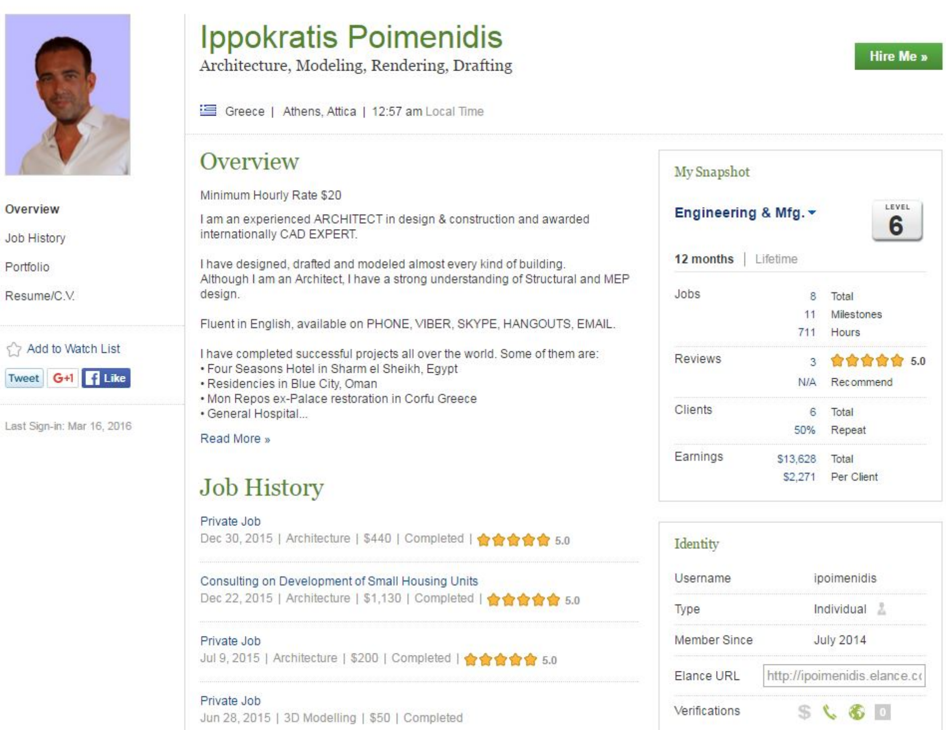

FIGURE 4

FREELANCER JOB FEEDBACK PAGE

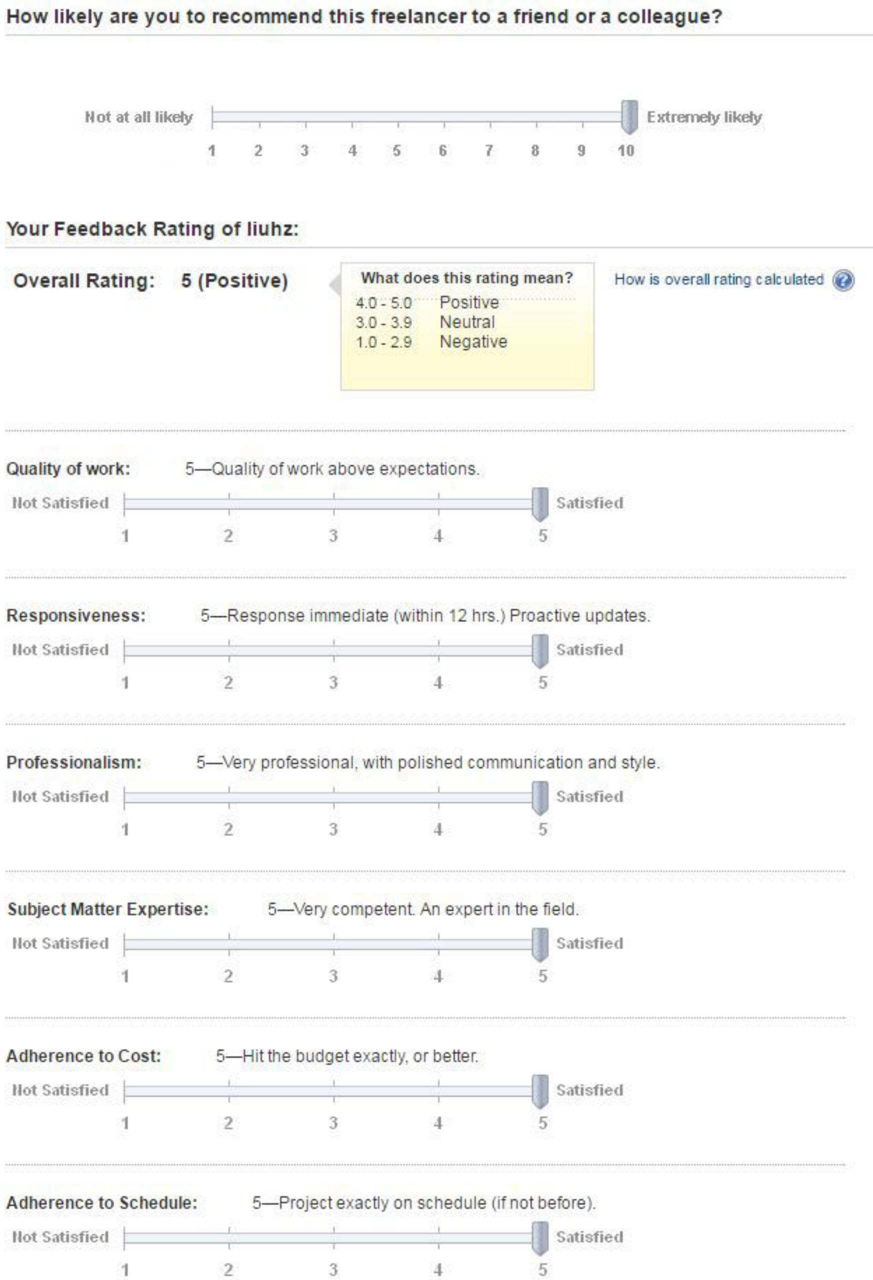

TABLE 1

SUMMARY STATISTICS AND CORRELATIONS

N=26,421,726

Variable	**Mean**	**Std. Dev.**	**Min**	**Max**
Hired =1	0.105	0.404	0	1
Individual bidder (=1)	0.375	0.484	0	1
Number of previous jobs (logged)	2.571	3.318	-4.605	9.661
Amount of bid (logged, in USD)	4.511	2.904	-4.605	29.386
Star rating	4.294	0.406	1	5
Level Score	5.333	3.129	1	18
Number bidders from same country	8.828	11.144	1	391
Average level score of bidders from this country	2.442	1.803	0	16.145
Number previously hired (logged)	1.553	2.387	-4.605	9.681
Number previously hired from bidder country (logged)	1.204	3.130	-4.605	8.610
Employer/Freelancer from same country (=1)	0.035	0.184	0	1
Number bad experiences from bidder country (logged)	-4.485	0.762	-4.605	3.332
Number bad experiences from bidder country in job category (logged)	-4.294	0.147	-4.605	3.095

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
(1)	Hired =1	1											
(2)	Individual bidder (=1)	0.029	1										
(3)	Number of previous jobs (logged)	0.060	-0.121	1									
(4)	Level Score	0.045	-0.127	0.721	1								
(5)	Star rating	-0.084	-0.037	0.167	0.165	1							
(6)	Amount of bid (logged, in USD)	0.082	0.004	-0.030	-0.051	-0.030	1						
(7)	Number bidders from same country	-0.144	0.014	0.000	-0.015	-0.028	0.026	1					
(8)	Average level score of bidders from this country	-0.024	0.011	0.193	0.214	0.063	-0.063	-0.043	1				
(9)	Number previously hired (logged)	0.186	0.034	-0.002	-0.001	-0.038	0.011	0.007	0.011	1			
(10)	Number previously hired from bidder country (logged)	0.457	0.023	0.053	0.034	-0.055	0.046	0.189	-0.04	0.498	1		
(11)	Employer/Freelancer from same country (=1)	-0.001	-0.016	-0.046	-0.032	-0.006	0.004	-0.056	-0.044	-0.017	-0.001	1	
(12)	Number bad experiences from country (logged)	0.031	0.006	0.002	-0.002	-0.017	0.009	0.060	-0.009	0.150	0.197	-0.027	1
(13)	Number bad experiences from country in job category (logged)	0.032	0.006	-0.001	-0.004	-0.017	0.007	0.017	0.007	0.050	0.071	-0.006	0.352

TABLE 2

LOGISTIC REGRESSION ESTIMATES OF THE LIKELIHOOD OF BEING HIRED
(JOB FIXED-EFFECTS)

	(1)	(2)	(3)
Individual bidder (=1)	-0.3613***	-0.3607***	-0.4678***
	(0.0055)	(0.0055)	(0.0066)
Number of previous jobs (logged)	0.1123***	0.1120***	0.0877***
	(0.0014)	(0.0014)	(0.0018)
Level score	0.0178***	0.0177***	0.0144***
	(0.0011)	(0.0011)	(0.0014)
Star rating	-.2160***	-.2162***	-.2166***
	(0.0063)	(0.0060)	(0.0070)
Amount of bid (logged, in USD)	0.1313***	0.1314***	0.0299***
	(0.0010)	(0.0010)	(0.0015)
Number bidders from same country	-0.0719***	-0.0713***	-0.1376***
	(0.0006)	(0.0006)	(0.0010)
Average level score of all bidders from this country	0.1227***	0.1213***	0.0205*
	(0.0041)	(0.0041)	(0.0093)
Number previously hired (logged)	3.0277***	3.0337***	3.0388***
	(0.0159)	(0.0163)	(0.0155)
Number previously hired from bidder country (logged)	0.8653***	0.8683***	1.1557***
	(0.0018)	(0.0018)	(0.0034)
Employer/Freelancer from same country (=1)	0.0452***	0.0280*	0.5727***
	(0.0118)	(0.0118)	(0.0029)
Count of bad experiences from bidder country (logged)		-0.1112***	-0.0968***
		(0.0041)	(0.0067)
Count of bad experiences from bidder country in this job category (logged)			-0.8040***
			(0.0432)
Number Bids	4287426	4287426	4287426
Number Jobs	527226	527226	527226
Pseudo-R^2	0.3937	0.3942	0.4483
Log-Likelihood	-476517.5	-476152.3	-476002.3
χ^2	618956.8	619687.1	622774.1

Notes: Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 3

COMPARISON OF PERCENT GOOD FREELANCERS
FROM NEGATIVE EXPERIENCE COUNTRY TO SUBSEQUENT COUNTRY OF HIRE
(PAIRED T-TESTS)

	(1) All Negative Experiences	(2) First Negative Experience	(3) Negative Experiences After First One
Negative Experience Country Mean	0.97663	0.97627	0.97762
	(0.00002)	(0.00002)	(0.0002)
Subsequent Country Mean	0.97663	0.97625	0.97765
	(0.00002)	(0.00002)	(0.0002)
Difference (Country mean – Subsequent Country Mean)	2.0e-06	1.5e-05	-3.4e-05
	(0.0002)	(0.0003)	(0.0003)
Observations	16462	12040	4422
D.F.s	16421	12039	4421

Notes: Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

FIGURE 5

HISTOGRAM OF FREELANCERS BY PERCENT 'GOOD' TYPE

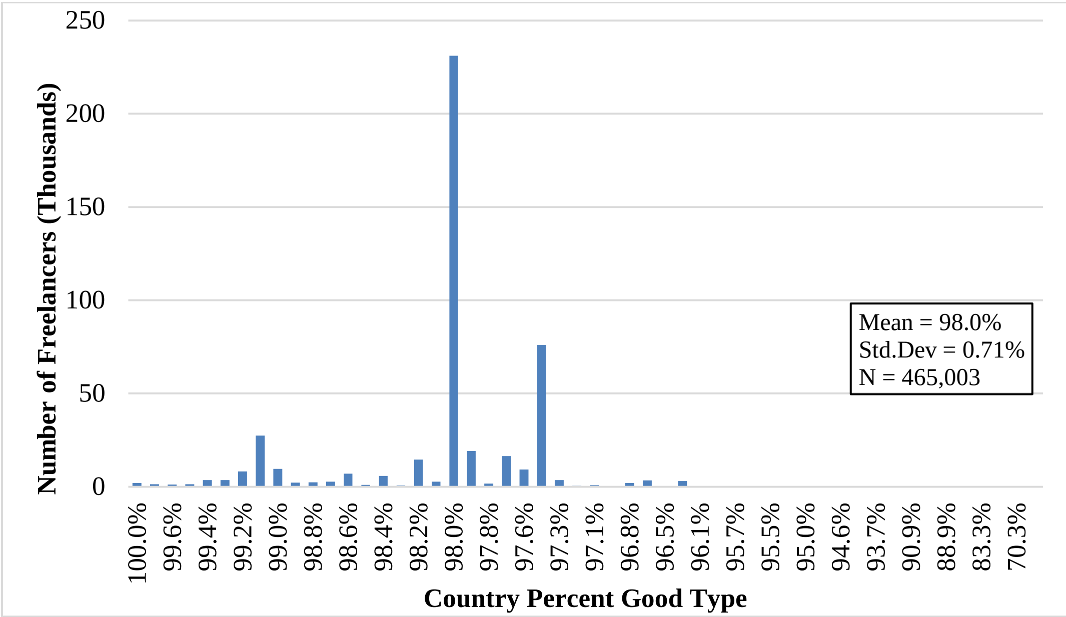

FIGURE 6

HISTOGRAM OF REALIZED % GOOD TYPE MEAN

RESULT OF 10,000 COUNTRY EXCLUSION SIMULATIONS

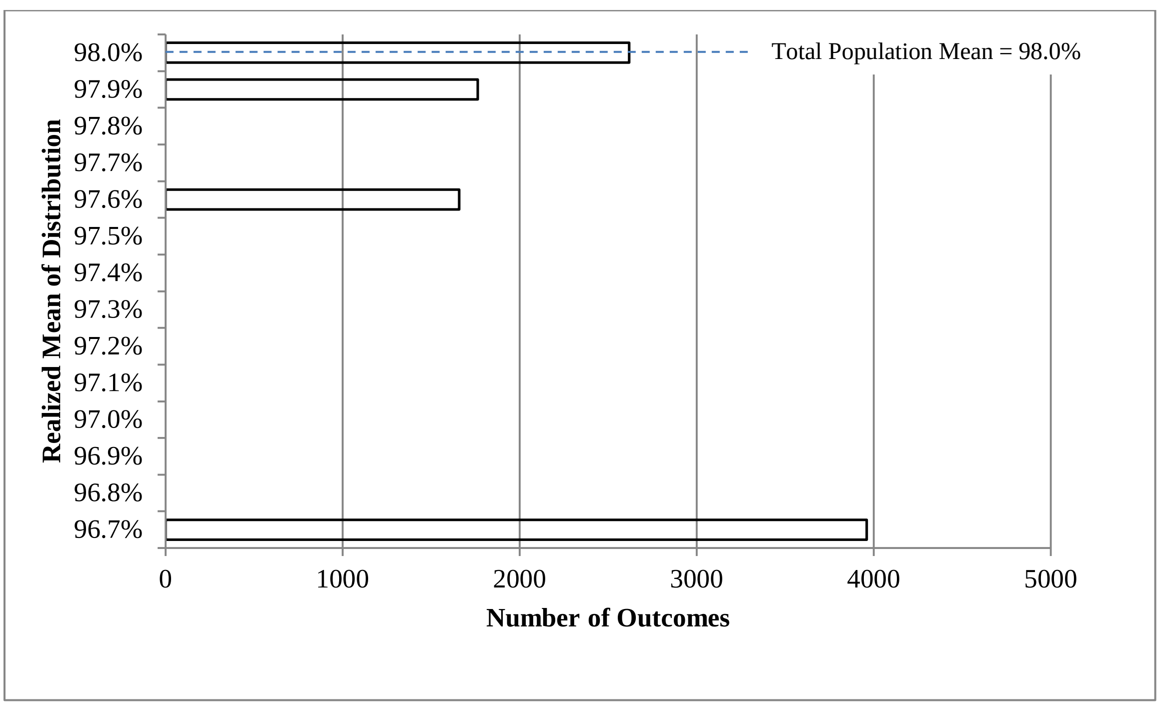

FIGURE 7
DISTRIBUTION OF OUTCOMES FROM 10,000 SWITCHED TO COUNTRIES
PERCENTAGE GOOD TYPE FREELANCERS

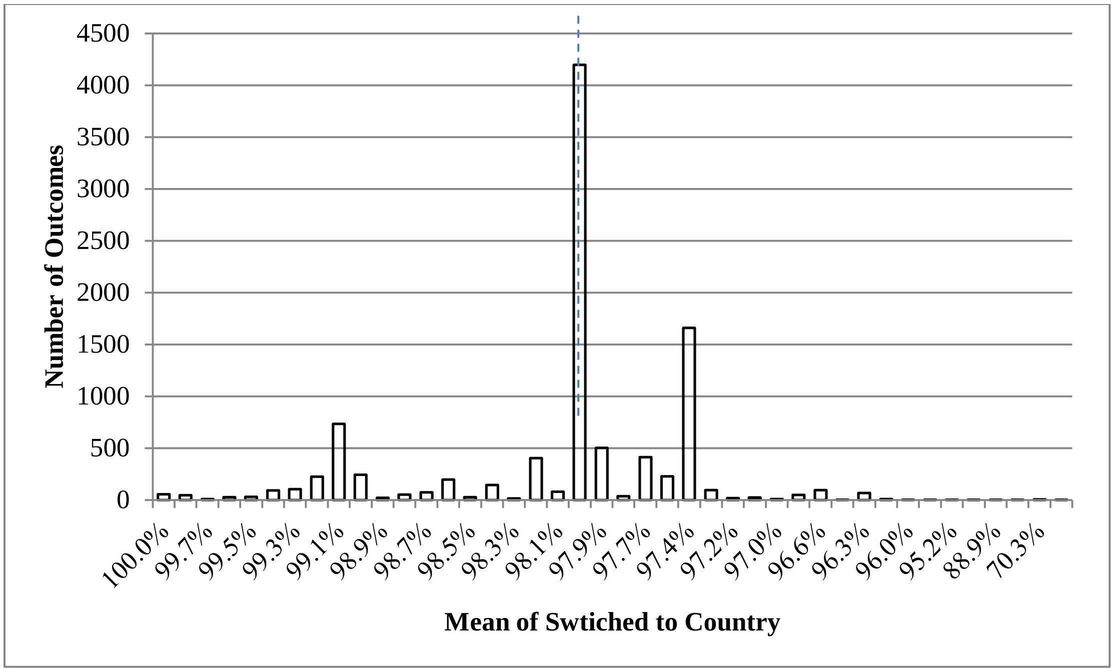

TABLE 4
POISSON ESTIMATES OF NUMBER OF BIDDERS FROM A COUNTRY
OLS ESTIMATES OF THE AVERAGE LEVEL SCORE OF APPLICANTS FROM A COUNTRY
(EMPLOYER FIXED-EFFECTS)

	(1) Number Bidders from this Country	(2) Average Level Score of Bidders from this Country
Number previously hired (logged)	-0.1183***	0.1715***
	(0.0016)	(0.0012)
Number previously hired from bidder country (logged)	0.4870***	0.0589***
	(0.0010)	(0.0007)
Employer/Freelancer from same country (=1)	2.0543***	0.0187*
	(0.0331)	(0.0265)
Count of bad experiences from bidder country (logged)	0.1317***	0.0042
	(0.0047)	(0.0036)
Constant	5.3872***	3.1076***
	(0.0210)	(0.0161)
Number Bids from Unique Countries	3192219	3192219
Number Employers	265286	265286

Notes: Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 5

LOGISTIC REGRESSION ESTIMATES OF THE LIKELIHOOD OF BEING HIRED
(JOB FIXED-EFFECTS, TOP 20 COUNTRY DUMMIES)

	(1)		(2)	
Individual bidder (=1)	-0.3607***	(0.0055)	-0.4914***	(0.0054)
Number of previous jobs (logged)	0.1120***	(0.0014)	0.0935***	(0.0015)
Level score	0.0177***	(0.0011)	0.0310***	(0.0013)
Star rating	-0.2160***	(0.0063)	-0.2560***	(0.0066)
Amount of bid (logged, in USD)	0.1314***	(0.0010)	0.1427***	(0.0012)
Number bidders from same country	-0.0713***	(0.0006)	-0.0280***	(0.0006)
Average level score of all bidders from this country	0.1213***	(0.0041)	0.0719***	(0.0044)
Number previously hired (logged)	3.0277***	(0.0159)	3.4277***	(0.0229)
Number previously hired from bidder country (logged)	0.8683***	(0.0018)	1.0386***	(0.0023)
Employer/Freelancer from same country (=1)	0.0280*	(0.0118)	-1.0347***	(0.0178)
Count of bad experiences from bidder country (logged)	-0.1112***	(0.0041)	-0.0810***	(0.0045)
Asia				
Freelancer from India			-0.2456***	(0.0079)
Freelancer from Pakistan			0.2413***	(0.0119)
Freelancer from Philippines			2.0800***	(0.0252)
Freelancer from Bangladesh			1.4892***	(0.0257)
Freelancer from China			1.4447***	(0.0297)
Freelancer from Indonesia			2.6025***	(0.0386)
Freelancer from Sri Lanka			1.8493***	(0.0442)
Western Europe				
Freelancer from Great Britain			1.5268***	(0.0156)
Freelancer from France			2.1039***	(0.0416)
Freelancer from Italy			1.9315***	(0.0519)
Freelancer from Spain			2.4461***	(0.0490)
Eastern Europe				
Freelancer from Romania			1.7704***	(0.0206)
Freelancer from Ukraine			1.7429***	(0.0261)
Freelancer from Russian Federation			1.4832***	(0.0308)
Freelancer from Serbia			1.8966***	(0.0483)
Freelancer from Bulgaria			2.1562***	(0.0402)
Americas				
Freelancer from Canada			1.4115***	(0.0169)
Freelancer from Argentina			0.8615***	(0.0267)
Oceania				
Freelancer from Australia			1.1380***	(0.0220)
Freelancer from Other Countries			1.7948***	(0.0170)
Number Bids	4287426		2981678	
Number Jobs	527226		294282	
Pseudo-R^2	0.3932		0.4602	
Log-Likelihood	-476957.82		-341139.38	
χ^2	669713.17		408076.29	

Notes: Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

TABLE 6

LOGISTIC REGRESSION ESTIMATES OF THE LIKELIHOOD OF BEING HIRED (JOB FIXED-EFFECTS)

	(1)	(2)
Individual bidder (=1)	-0.4421***	-0.2564***
	(0.0051)	(0.0063)
Number of previous jobs (logged)	0.0780***	0.1299***
	(0.0014)	(0.0015)
Level score	0.0355***	0.1234***
	(0.0014)	(0.0011)
Star rating	-0.2180***	-0.2660***
	(0.0063)	(0.0066)
Amount of bid (logged, in USD)	0.1444***	0.0057***
	(0.0011)	(0.0012)
Number bidders from same country	-0.0942***	-0.0658***
	(0.0006)	(0.0006)
Average level score of all bidders from this country	0.1694***	0.1091***
	(0.0047)	(0.0043)
Number previously hired (logged)	3.0066***	3.0007***
	(0.0159)	(0.0162)
Number previously hired from bidder country (logged)	0.9427***	0.8401***
	(0.0019)	(0.0019)
Employer/Freelancer from same country (=1)	-0.4123***	-0.0097
	(0.0358)	(0.0124)
Count of bad experiences from bidder country (logged)	-0.1750***	-0.0903***
	(0.0055)	(0.0064)
Number of days since last bad experience	0.3405***	
	(0.0078)	
Bad Experience X Number of Days ago	0.0078***	
	(0.0016)	
Cumulative $ value of bad experiences (logged)		-0.0091**
		(0.0028)
Number Bids	4287426	4287426
Number Jobs	527226	527226
Log-Likelihood	-476012.2	-476015.4
χ^2	621974.1	622694.1

Notes: Standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$