

NBER WORKING PAPER SERIES

DIGITAL LABOR MARKETS AND GLOBAL TALENT FLOWS

John Horton
William R. Kerr
Christopher Stanton

Working Paper 23398
<http://www.nber.org/papers/w23398>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2017

This paper is a revision of the paper prepared for the NBER Global Talent Summer Institute Conference meeting on July 10, 2016, and a draft chapter for the associated NBER volume edited by Gordon Hanson, William Kerr, and Sarah Turner. Comments are most welcome. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by John Horton, William R. Kerr, and Christopher Stanton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Digital Labor Markets and Global Talent Flows
John Horton, William R. Kerr, and Christopher Stanton
NBER Working Paper No. 23398
May 2017
JEL No. F15,F22,F23,J15,J31,J44,L14,L24,L26,L84,M55,O31,O32

ABSTRACT

Digital labor markets are rapidly expanding and connecting companies and contractors on a global basis. We review the environment in which these markets take root, the micro- and macro-level studies of their operations, their ongoing evolution and recent trends, and perspectives for undertaking research with micro-data from these labor platforms. We undertake new empirical analyses of Upwork data regarding 1) the alignment of micro- and macro-level approaches to disproportionate ethnic-connected exchanges on digital platforms, 2) gravity model analyses of global outsourcing contract flows and their determinants for digital labor markets, and 3) quantification of own- and cross-country elasticities for contract work by wage rate. Digital labor markets are an exciting frontier for global talent flows and growing rapidly in importance.

John Horton
Leonard N. Stern School of Business
Kaufman Management Center
44 West Fourth Street, 8-81
New York, NY 10012
john.horton@stern.nyu.edu

Christopher Stanton
210 Rock Center
Harvard University
Harvard Business School
Boston, MA 02163
and NBER
christopher.t.stanton@gmail.com

William R. Kerr
Harvard Business School
Rock Center 212
Soldiers Field
Boston, MA 02163
and NBER
wkerr@hbs.edu

1 Introduction

Globalization has been a mighty force over the last few decades. Compared to the movements of material goods and financial capital across countries, however, labor and talent have been much slower to globalize. This greater localization of labor and talent is perhaps not surprising given that it is easier to transmit financial capital in a disembodied form or build/ship a physical good for an exact purpose. People and their labor, however, have typically come as a collective and fully integrated package, so-to-speak, that makes location decisions more complex. If one seeks to access labor inputs available abroad, one option is to attract and host the individual, temporarily or permanently, near the location of the work to be performed. For a variety of reasons, this has proven politically unpopular and nearly all countries place restrictions on migrations. As a result, so only about 3% of the world’s population lives outside of their country of birth.¹

A second option is to identify how the required task can be exchanged at a distance, without necessitating a person’s physical migration. Offshoring—the performance of a specified task in another country—has become a substantial force in certain business functions where the tasks can be effectively located at a geographic distance. Thus, the focus shifts from “trade in goods” to “trade in tasks” needing to be performed (Grossman and Rossi-Hansberg, 2008). Prominent examples include low-end data entry and high-end back-office information technology (IT) in India for US and European companies. In a prominent study, Blinder and Krueger (2013) estimate that around one quarter of jobs could be offshored from the United States.²

Offshoring was initially best suited for large corporations due to the substantial fixed costs in establishing an overseas presence. Even if using an external outsourcing vendor, it only made sense for organizations to engage in trade in tasks if they had a sufficiently large ongoing volume of work to justify learning about overseas options, vetting contractors, negotiating terms and prices, and reorienting their own business processes to fit around the overseas work, etc. Similar to the Melitz (2003) model for international trade, firms entered into these overseas efforts when a large and sustained improvement that exceeded a threshold requirement was feasible. Helpman et al. (2004) develop a framework where the most productive firms launch overseas

¹Kerr et al. (2016, 2017) review the literature, data and policy environments for global talent flows. Clemens (2011) emphasizes the “trillion-dollar bills” that remain on the sidewalk due to this low rate of migration in light of productivity differences across countries.

²Offshoring closely relates to outsourcing—the performance of a specified task by an external party to the purchasing company—and the two terms are often used interchangeably in the press. Outsourcing is possible without offshoring (e.g., purchasing services from an external company in one’s own country), and offshoring is possible without outsourcing (e.g., setting up a company-owned data center or manufacturing plant abroad). For most of this paper’s discussion of digital labor markets, the two concepts overlap completely as the contracts are both externally sourced and abroad.

facilities, those with intermediate productivity engage in trade, and the least productive firms serve domestic markets only. Helpman (2014) provides a review.

Digital labor markets have the potential to radically alter this picture. These Internet-based platforms connect workers worldwide with companies seeking to have tasks completed. This chapter describes digital labor markets, evaluating their dramatic rise and global span, and reviews academic studies of how these markets function. We first discuss the persistent information frictions that have been a barrier to off-line global labor sourcing and how digital labor platforms address these barriers. Sections 2 and 3 provide both micro- and macro-level perspectives, respectively, and we present some new empirical analyses that link these two perspective together with respect to cross-border contract placement over countries. Our empirical discussion uses data from Upwork, the world’s largest digital labor platform, and its predecessor oDesk.³ We extend prior work by Ghani et al. (2014) on ethnic contract placement, and we provide new evidence regarding flows and substitution across countries.

Section 4 then considers the evolution of digital labor markets and provides case-based examples of other ways that digitization is extending the spatial reach of labor and talent inputs. For example, many corporations and governments are rushing to build “open innovation” platforms that expose their organizations to valuable external ideas. We discuss examples from Procter & Gamble, the National Aeronautics and Space Administration (NASA), and similar large organizations on how they are using open collaboration concepts for solving thorny innovation challenges. Digital platforms are also extending the use of global labor to many smaller start-ups, and overseas tech development has become the norm for many US and European entrepreneurs given the cost savings possible.

Only time will tell the ultimate impact of digital labor markets, online innovation contests, and similar collaborative activities for the globalization of labor markets and talent, but their strong potential is now evident. Moreover, they are becoming a powerful tool for researchers seeking to understand the functioning of labor markets. It is exceptional, for example, to observe a recorded history of the bids given for contracts, the traits of accepted bids versus the competition, the performance outcomes of projects, the prior and subsequent longitudinal history of workers and contracting firms, and so on. See Horton and Tambe (2016) for an overview of the research potential of computer-mediated labor markets. These platforms have also been the site

³Upwork is the result of a merger in 2014 of Elance and oDesk, which were founded in 1999 and 2003, respectively. In 2016, Upwork reports annually servicing over three million jobs that represent more than \$1 billion in work. Projects range from simple transcription work to high-end services, and Upwork records over 12 million registered freelancers and 5 million companies (<https://www.upwork.com/about/>, accessed June 21, 2016).

for multiple experimental studies of labor market behavior. Building on our research experience, the fifth section provides some perspectives for researchers about the advantages and pitfalls of using these types of data and platforms for economic studies, and we close with some open questions for the future about these platforms and the digitization of work.

2 The Environment of Digital Labor Markets

2.1 Upwork

Upwork is an online platform that connects workers who supply services with buyers who pay for and receive these services from afar. Examples include data entry and programming tasks. The platform is the result of a 2014 merger between Elance and oDesk, and the merged entity was rebranded as Upwork in 2015. In 2016, Upwork is the world’s largest platform for online outsourcing, and oDesk and Elance were the two largest platforms before the merger. To be consistent and reduce confusion, we favor using the name Upwork even when describing a period before the company was known by this name. When discussing and extending studies of earlier periods that use oDesk-specific data, we mention this alternative sample. The data used in this study were obtained directly from oDesk and Upwork for research purposes.

On the Upwork platform, any worker can contract with any firm directly, and all work takes place and is monitored via a proprietary online system.⁴ In exchange for a 10% transaction fee from the total wage bill⁵, Upwork provides a comprehensive management and billing system that records the time spent by the worker on the job, allows easy communication between workers and employers about scheduled tasks, facilitates simple document uploading and transfer, and takes random screenshots of workers’ computer terminals to allow electronic monitoring.⁶

These features facilitate easy, standardized contracting, and any company and any worker can form electronic relationships with very little effort. More advanced features provide tools for teams to collaborate on projects.

A worker who wants to provide labor services on Upwork fills out an online profile describing his/her skills, education level, and experience. A worker’s entire history of Upwork employment, including wages and hours, is publicly observable to potential employers. For contracts that have been completed, a feedback measure from the employer is publicly displayed. Figure 1 provides

⁴This section draws from Ghani et al. (2014).

⁵Upwork recently announced a new non-linear pricing structure in which fees would be gradually reduced as the match-specific wage bill increased.

⁶We use the terms “employer” and “employment” for consistency with the existing labor literature rather than as a comment on the precise legal nature of the relationships created on these sites.

an example of a worker profile.

Companies and individuals looking to hire on Upwork fill out a job description, including the skills required, the expected contract duration, and some preferred worker characteristics. In the first few years after the platform’s founding, most of the jobs posted were hourly positions for technology-related or programming tasks (e.g., web development), but postings for administrative assistance, data entry, graphic design, and smaller categories have become more prevalent as the platform has grown. Advanced tasks include search engine optimization, data analytics, and mobile app programming. Table 1 provides a distribution of contracts over job category. After a company posts a position opening, workers apply for the job and bid an hourly rate. Firms can interview workers via Upwork, and ultimately form a contract if both parties agree. In the past, this process was large decentralized, but in more recent years, Upwork has invested heavily in making algorithmic recommendations to both employers and workers about which worker to hire or which job to apply to, respectively. See Horton (2017) for evidence on the effectiveness of these algorithmic recommendations in increasing the quantity of matches formed in the market.

2.2 Micro-Evidence on Information Frictions

Most past studies of oDesk/Upwork are micro-based studies that tend to focus on matching or information frictions. Evidence on the existence of these frictions are evident in the data used here. The literature’s focus on these micro-based frictions is perhaps surprising at first glance, given that the core power of these platforms and their rising economic importance is the global information access and firm-worker matching process that the platforms enable, often for the first time. Yet, even though these platforms have removed many frictions from their labor markets (e.g., information access, document transfer, billing, etc.), some classic issues remain and perhaps become more evident, such as uncaptured externalities for the development of information about workers and firms for the market to observe or ethnic/racial biases people have in contract selection. Also, similar to other online environments like auction sites or e-commerce platforms, new issues can arise due to the platform’s features and aggregation of many buyers and sellers that are hard to anticipate. Here we review several studies and tie together what they mean for our understanding of matching frictions.

Many of the matching frictions that have previously been documented arise because employers hire discrete workers into particular slots (see Lazear et al. (2016)). Table 1 shows that there are many more applicants than slots available to contract. That there are many applicants relative to openings suggests that it may be hard for workers to determine what employers are looking

for or how an applicant will be assessed against other workers. On oDesk/Upwork, because of unobserved capacity to take on new projects, employers have the same problems when they pursue workers (see Horton (2016a)). These and other forms of information frictions result in sunk effort on both sides of the market before a successful match is formed.

Several factors contribute to these frictions, and many are also present in traditional labor markets. These include uncertainty and difficulty in assessing worker quality, leading to concerns about adverse selection. Other questions around direct contract enforcement are potentially relevant as well. For larger projects, team aggregation challenges appear to be compounded in the online setting.

Over time, the Upwork market has evolved to better provide features that mitigate these sources of friction. Reputation systems, prevalent in many peer-to-peer and electronic markets, were early features designed to mitigate adverse selection. However, as discussed by Horton and Golden (2016), these systems often provide only coarse information that results in “bunching” of scores either at the top or bottom of the rating scale. Many employers are reluctant to leave negative feedback, and so only “good” feedback is reported. It also appears that what is considered “good” has increased over time, leading to a kind of reputation inflation. As such, would-be employers have difficulty assessing ability *ex ante* (though this is far from a challenge unique to online settings). As a reaction to this problem, Upwork has moved to utilize the fact that experienced workers often transact with many employers, enabling the display of private feedback ratings that aren’t linked to an individual transaction. This has reduced the effect of bunching on market frictions by providing additional gradation between workers.

While reputation systems provide information about past performance, new workers face the problem of how to break into the market. Hiring a novice worker produces two outputs: the direct work product and information about that worker’s quality. However, because these are spot markets with somewhat limited full-time repeated contracting, the information about worker output is not particularly valuable to an employer. As a result, there is under hiring of unknown workers because employers do not internalize the value of generating knowledge about workers that is revealed once they start work (Tervio, 2009). In the data, very few novice workers are hired relative to the experienced cohort. Pallais (2014) demonstrates through an experiment that a major contributing factor to the low share of novice hires is the Tervio mechanism where employers do not to internalize the value of information. To do so, she randomly hires novice workers and leaves them honest feedback. This initial feedback has profound effects on treated workers’ online careers. Future employers are much more likely to hire workers in the treatment

group who receive a rating than control workers who did not receive the rating.

Stanton and Thomas (2016b) then show that the market has evolved to include intermediation as a response to the worker startup problem. Intermediaries, called agencies, have entered online labor markets and have altered hiring patterns for novice workers. These agencies tend to be small groups consisting of several online workers, and employers can observe agency affiliation and an agency-level feedback score on each affiliated worker's profile. Most agency workers are co-located, suggesting some role for offline ties in the formation of these groups. A key factor for overcoming the information problem is an incentive to invest, and intermediaries are provided with this incentive because they own the reputation of their affiliated workers. Stanton and Thomas show that novice workers who enter the market with intermediary affiliation are much more likely to find work than workers who enter without affiliation. They identify the information effect of intermediation by comparing outcomes over workers' careers; the initial intermediary advantage fades-out as workers gain experience. The entry of intermediary agencies has improved the prospects of novice affiliated workers and has reduced frictions for affiliated novice workers who seek to enter the market.

The earliest frictions explored in the literature were due to adverse selection concerns because of employers' difficulty distinguishing worker quality. More recent literature explores the consequences either of uncertainty about the environment that employers face or switching frictions when changing from a familiar offline environment. Stanton and Thomas (2016a) explore uncertainty about the market as the result of employers being unfamiliar with the value of the market. Because employers' interviews are observed in the data, a measure of search effort is available. Stanton and Thomas document that employer interviewing falls dramatically with experience, suggesting an important role for learning about the distribution of matches through the process of hiring. If some factors cause new employers to forego initial hiring, strong experience effects suggest that these factors limit market size by the failure to move new users along the experience curve. Stanton and Thomas suggest that the nature of how workers bid for jobs is a significant factor that has limited the takeup of new users. Because workers can observe employer inexperience, it is possible for them to tailor wage bids to what employers are likely to know about the market. In most markets inexperienced users receive lower prices to draw them in, but in online labor markets inexperienced employers receive wage bids that are approximately 10% higher than their experienced alter-egos. The spot nature of contracting means that workers do not participate in the employer gains from learning the market. Workers' higher bids limit take up of the market and hinder the expansion of online work. The failure of decentralized actors to

internalize the consequences of how their own behavior affects information for trading partners has the potential to limit the growth of online exchange. Differences in pricing policy may be necessary to counteract some of these incentives.

Other work suggests that offline familiarity influences online hiring behavior. For example, Ghani et al. (2014) document the prevalence of ethnic-linked exchanges online by studying the hiring patterns of the Indian diaspora on oDesk/Upwork. Importantly for identification, applicants do not know the ethnic identity of the employer; this minimizes concerns about sorting as a confound. Even with access to workers from all over the world, they find that the ethnic Indian diaspora is much more likely to hire in India than employers of other ethnicities. Whether due to preferences or information problems, this may limit the amount of trade conducted through opening labor markets online. On the other hand, the reliance on familiarity may, in theory, grease the wheels of transactions and help employers to overcome uncertainty about workers. In the Upwork context, the size of the Indian diaspora hiring online suggests this role for encouraging the sourcing of online work is likely to be a small factor in encouraging market growth.

For those employers who do take an initial jump, several strategies may be used to deal with an uncertain environment. For example, many employers appear to use hiring tournaments in which small pieces of a project are done by multiple workers; the best workers are retained. This process can be repeated until a satisfactory set is found. This strategy is likely to make sense for production processes like software engineering where there are multiple ways to solve a problem. For tasks where accuracy is important, sourcing redundant projects and using error checking across workers to find mistakes may be more appropriate. Both of these strategies help to resolve uncertainty. Employers also appear to use pattern matching after successful outcomes. For example, Ghani et al. (2014) report that employers who initially choose to source work in India are more than 11.5% more likely to choose India on their next contract upon success compared to employers with unsuccessful first contracts.

That employers use workers' countries as an important source of information has been documented in several sources. Mill (2013) studies statistical discrimination and employer learning through experience with hiring in particular countries. Guo (2016), using data from an early online labor market called rentacoder, shows that employers update their beliefs about all workers from a country after hiring from that country. Agrawal et al. (2014) examine the structure of information and how this affects workers differently depending on their country. An interesting and important finding of this paper is that although at least some employers behave in a way

consistent with statistical discrimination, information about actual worker productivity seems to be a remedy: with more information, employers engage in less crude statistical discrimination. Using an experiment, Lyons (2016) also examines cross-country versus intra-country differences in team production when hiring online, extending many of these results to more complicated production.

2.3 Ethnic Diasporas and Contract Placement

While micro-frictions have been the literature’s main focus, we turn in the next section towards a more macro-oriented analysis of contract placement, providing some first evidence regarding flows and substitution across countries. In preparation for the macro perspective, we first provide an example of how the micro and macro lens connect with each other. We do this by extending the work of Ghani et al. (2014), who quantify how members of the Indian diaspora are more likely to place an outsourcing contract into India, compared to non-Indians, and have some important differences as to how these contracts are structured. While this analysis shows micro-connectivity, it differs from the standard analysis in the macro literature. Rauch and Trindade (2002), for example, relate trade flows to the distribution of the ethnic Chinese population across countries, rather than the greater likelihood that two observed traders are Chinese. We thus extend our earlier work to now mirror the approach of Rauch and Trindade (2002). To keep the analysis in line with Ghani et al. (2014), we use oDesk data covering 2005-2010.^{7,8}

In this analysis, as well as the one to come in Section 3, we use the gravity framework from the international trade literature to guide our work. Similar to planetary pull, these trade models suggest that countries should engage more in trade to the degree that they are larger and also closer together. There are several theoretical ways that one can derive a gravity model, and the appendix to this paper outlines the Eaton and Kortum (2002) model that is most aligned with our work. The Eaton and Kortum (2002) model considers countries having a range of technological productivities for various activities. Each country purchases the activity from the country that

⁷The oDesk data do not record a person’s ethnicity or country of birth, so Ghani et al. (2014) use the names of company contacts to probabilistically assign ethnicities. This matching approach exploits the fact that individuals with surnames like Chatterjee or Patel are significantly more likely to be ethnically Indian than individuals with surnames like Wang, Martinez, or Johnson. The matching procedure exploits two databases originally developed for marketing purposes, common naming conventions, and hand-collected frequent names from multiple sources like population censuses and baby registries. The process assigns individuals a likelihood of being Indian or one of eight other ethnic groups. Kerr (2007, 2008) and Kerr and Lincoln (2010) provide extended details on the matching process, list frequent ethnic names, and provide descriptive statistics and quality assurance exercises. Ghani et al. (2014) provide an extended discussion and analysis of this match in the oDesk-specific context.

⁸More broadly, recent research emphasizes the importance of immigrants in frontier economies for the diffusion of technologies and ideas to their home countries (e.g., Saxenian 2002, 2006; Kerr 2008; Agrawal et al. 2011). Kerr (2016) reviews this literature and its connection to trade more completely and provides appropriate references.

can be the lowest cost provider of the activity, including the purchasing country itself. This cost considers the price levels and wage rates in countries, the productivity of countries for tasks, and distances between nations. Thinking of these activities as tasks on a digital labor platform is a natural extension, and our empirical analysis relates the volume of contracting between countries. The appendix provides a more rigorous introduction.

The dependent variable in Columns 1-7 of Table 2 is the share of contracts originating from a country on oDesk that are outsourced to India. We focus on shares of contracts, rather than contract volumes, as the adoption of oDesk across countries as a platform for e-commerce is still underway and somewhat idiosyncratic to date. Shares allow us to consider the choice of India for outsourcing independent of this overall penetration of oDesk. The core regressor is taken from the World Bank’s Bilateral Migration and Remittances 2010 database. This database builds upon the initial work of Ratha and Shaw (2007) to provide estimates of migrant stocks by country. We form the Indian diaspora share of each country’s population by dividing these stocks by the population levels of the country. We complement this diaspora measure with distances to India calculated using the great circle method, population and GDP per capita levels taken from the United Nations, and telephone lines per capita in 2007 taken from World Development Indicators. We also calculate a control variable of the overall fit of the country’s outsourcing needs with the typical worker in India.⁹

Column 1 presents our base estimation. We have 92 observations, and we weight by the log number of worldwide contracts formed on oDesk. The first row shows the connection of digital outsourcing to the diaspora population share, which is quite strong. A 1% increase in the Indian diaspora share of a country is associated with a 1% increase in the share of oDesk contracts outsourced to India. The country-level placement of digital contracts in India systematically followed the pre-existing levels of Indian diaspora communities. Looking at the other covariates, spatial distance does not matter in the digital labor context like it does in many estimates of economic exchanges. In fact, the share of contracts sent to India increases with spatial distance.¹⁰

⁹We calculate this control by first measuring the share of contracts outsourced from the country in nine job categories indicated. We likewise measure the distribution of oDesk work performed in India across the nine job categories, independent of where the company contact is located. We then calculate the sum of the squared deviations of these two distributions to measure how closely the work typically filled in India matches the needs of a given country. We subtract this sum of deviations from one, so that positive values represent a better fit, and we transform the measure to have unit standard deviation to aid interpretation.

¹⁰Unreported estimations also find that time zones do not play a strong role in contract placement. The coefficient values suggest a negative effect of being further apart in terms of time zone, but these results are very small in magnitude and not statistically significant. Two important details to note are 1) many digital contracts (e.g., data entry) do not require extensive synchronous interaction, and 2) for those that do, many Indian workers are willing to work the originating country’s business day if that is needed for securing the job. Appendix Figure 1 provides a more detailed application timezone analysis taken from Horton (2016a). This figure shows the shifting of schedules more broadly.

The overall fit of a country’s outsourcing needs with the skillsets of Indian workers predicts that greater shares of work are sent to India. On the other hand, country population levels and telephone penetration do not play an important role. We likewise find similar weakness in Internet penetration measures, but they are not as uniformly available. Finally, countries with higher GDP per capita send less of their work to India conditional on the other covariates.

Many countries have been slower to develop on digital labor platforms, and nations with very few contracts can generate noisy share estimates. Our main estimations thus weight by contract volume to focus attention on better measured data and more meaningful observations; we utilize log weights to not overly emphasize the United States experience in particular. Columns 3 and 4 show similar results when we weight by log country population or when we exclude the weights. In both cases, the coefficients decline somewhat and the standard errors grow given the greater emphasis placed on noisy outcomes, but the role of diasporas remains economically and statistically significant. Column 5 shows similar results when adding a control for the total worldwide count of contracts on oDesk by a country. This variable picks up the negative effect earlier associated with GDP per capita. Column 6 tests whether this connection is simply following on existing business relationships that countries have with India. We measure the extent to which India is a trading partner of the focal country by the total volume of trade in 2007 between India and the country divided by the country’s GDP. Introducing this as a control does not affect our results.

Column 7 shows that the elasticity declines when excluding an outlier firm in the United Arab Emirates that outsourced an enormous number of contracts to India, but overall the pattern remains similar and statistically significant.¹¹ Column 8 finds similar results when examining the dollar share of contracts being sent to India rather than the count share. This estimation naturally puts more weight on contracts that have higher wages and longer durations. The coefficient declines compared to Column 1 but remains economically important and statistically significant. Finally, Column 9 provides an important connection to our earlier estimation approaches. The dependent variable is the share of company contacts using oDesk in the focal country that are of ethnic Indian origin (independent of whether or not the work is contracted with India). Larger Indian diaspora shares in a country’s general population are highly correlated with a larger

¹¹The results are not overly-dependent upon a single country, and we find very similar results when excluding the United States, Pakistan, and similar. Excluding the UAE has the largest effect, resulting in a point estimate of 0.878 (0.660), which is not very surprising given that the Indian diaspora’s share of 35% in the UAE is by far the largest, twice that of the next-highest states of Qatar (18%) and Oman (17%). As a second approach, we find a point estimate of 1.629 (0.654) when winsorizing outlier diaspora shares to Oman’s value to cap the UAE’s extreme value. The role of the diaspora community is also very similar when including a control for English language proficiency, which we are able to assemble for about half of the countries in our sample.

share of oDesk company contacts for the country being of ethnic Indian origin. The coefficient measures that a 1% increase in the relative size of the India diaspora to host country population (e.g., from 1% to 2%) is correlated with a 2.6% increase in the share of oDesk company contracts in that host country who have Indian ethnic names (e.g., from 10% to 13%).¹²

To summarize, in a spirit similar to Rauch and Trindade’s (2002) analysis of Chinese diaspora and flows of trade in manufactured goods, we find clear evidence linking the Indian diaspora to the placement of digital outsourcing contracts into India. This complements the micro-level perspective taken by Ghani et al. (2014). This is encouraging more broadly as it provides greater assurance that micro- and macro-level approaches are providing complementary perspectives on the functioning of digital labor markets.

3 Macro-Level Perspective

3.1 Contract Flows on Digital Labor Platforms

Figure 2 displays the asymmetric distribution of contract flows on Upwork. The most striking features of contract flow on Upwork are 1) the North-South nature of placements and 2) the very limited degree that countries provide services to themselves, with the United States being a major exception.

Table 3A ranks the top 20 hiring countries by aggregate wage bill from cross-border contracts from the launch of the platform through 2015. The United States is by far the largest hiring economy, with a cumulative wage bill for cross-border contracts that is almost seven times higher than second-ranked Australia. In addition to placing more jobs abroad, US employers have contracts that average 35% more in wage bills compared to the other countries given in Table 3A. By contrast, the United States is only the seventh ranked country from a worker perspective, and only four of the top 20 worker countries are present on this employer country list (i.e., United States, United Kingdom, Canada, and Germany). This emphasizes the exceptionally strong North-South nature of contract placements on digital labor markets.

Table 3B provides a mirror image of Table 3A from the worker perspective. India is the largest country by worker wage bill, with \$340 million in cumulative wages received through 2014. This is about 19% larger than the cumulative wage bill for the Philippines, the second-

¹²Considering partitions of the data, the diaspora coefficient is 0.893 (0.263) for 2008 and prior, 1.085 (0.240) for 2009 and later, 0.798 (0.238) for high-end contracts, 0.592 (0.113) for low-end contracts, 0.448 (0.232) for initial contracts, and 1.134 (0.334) for subsequent contracts. Ghani et al. (2014) analyze further how overseas ethnic Indians show higher rates than other ethnic groups of outsourcing initial contracts to workers in India and the path dependence that follows for subsequent contracts.

ranked country. After the Philippines, the gap is more dramatic; the Ukraine is ranked third, with a wage bill of \$118 million, or about 35% of the Indian total. Figures 3A and 3B depict the top bilateral routes by contract volume and dollar value, respectively.

Table 4 ranks the top suppliers of contract labor to the United States, again using cumulative wage bills over the oDesk/Upwork history. The United States edges out India and the Philippines as the largest provider of contract labor to itself. Behind this aggregate statistic, India and the Philippines record greater contract volume in Column 3, but the average wage bill for US-sourced work has been higher (\$943 vs. \$666 for India and \$538 for Philippines, respectively). Two other South Asian countries (Pakistan, Bangladesh), Russia and Ukraine round out the next largest providers of digital labor and talent for US employers.

The pattern of flows are quite unique to digital labor markets. Excluding the United States, there is a -0.08 correlation among the remaining 19 countries in terms of aggregate wage bill supplied (Column 4) and total US imports of manufactured goods (Column 7). China is the eighth-ranked provider of services, at only 10% of the level of India or the Philippines. Quite noticeably, Japan and Mexico are not even listed on Table 4, suggesting the negative correlation would further strengthen in their presence. The correlation is a similar -0.09 when comparing Column 4 against the total US imports of services in Column 8. While not shown in this table, it is again quickly evident upon reflection that the global sourcing of Upwork contracts is also quite different from global sources for immigrants to the United States.

Figure 4 provides a summary statistic of the distribution of US source countries for workers on Upwork compared to America's distribution of source countries for traded goods and services. We calculate the sum of the squared deviations between the share of a country's Upwork wage bill paid by US employers and the equivalent share in traditional accounts of traded goods and services. Goods imports data come from the census, and services imports data come from the World Bank TSD database and are last reported by country in 2011. To avoid compositional changes in the series over time, the goods and services series are restricted to be balanced. Deviations of Upwork shares are calculated against the balanced series. A level of zero would indicate perfect alignment of source countries, while a level of two is the theoretical maximum.

In the earliest phases of the platform, circa 2006, there was substantial divergence of source countries for digital labor work compared to typical patterns for both trade in manufactured goods and trade in services. Since this time, the squared deviations of source countries for oDesk/Upwork have further diverged from the source countries for manufactured goods, while convergence towards source countries for trade in services is evident until 2011. Consistent with

earlier tables, the largest deviations for the goods series are for China, India, the Philippines, and Russia. China is a large trading partner offline but has little online share. For the other countries, their online share exceeds their offline share.

3.2 Gravity Models of Contract Flows

Stepping beyond the example of the United States, Table 5 next examines digital outsourcing patterns across all country pairs using the familiar gravity model. Beyond the information that we derive directly from the Upwork database, most covariates used in this section come from the bilateral gravity and TRADHIST CEPII datasets. We consider a cross-sectional estimation of bilateral country pairs using the pseudo maximum likelihood estimator of Santos Silva and Tenreyro (2006). This conservative approach also allows us to retain bilateral routes on which zero contract placement occurs on Upwork.

The dependent variable is the wage bill from cross-country contracts paid by employer country to the worker country. We include employer country and worker country fixed effects in estimations that account for overall levels evident in Tables 3a and 3b. Employer country fixed effects are concentrated out, and worker country fixed effects are included as unreported indicator variables. Estimations are weighted by total contracts paid by employer country to reflect the global distribution of trade and robust standard errors are used to account for heteroskedasticity.

To allow for non-linear effects, we model most explanatory variables as indicator variables for various points in the distribution of a covariate. The omitted category for each indicator is the smallest/least category (e.g., shortest bilateral distance or GDP/capita difference between employer and worker country smaller than \$5000). Coefficients for each explanatory factor show the conditional differential compared to the omitted group. All models also include unreported indicator variables for quartiles of the product of GDP between countries.

Column 1 provides a base estimation that does not use recent offline trade flows as an explanatory variable. We first find that distance still matters in shaping the broad distribution of outsourcing contracts. We did not observe this pattern in the special case of India, examined in Table 2, but it is more systematically present when considering global contract placements. On the other hand, contiguous countries often show stronger links and economic integration, but we do not find evidence of a border effect in these data. A common country language and sharing a time zone also appear to boost contract placement.¹³ Finally, we observe that the

¹³The common country language result, however, is not robust across the multiple language variants developed by Melitz and Toubal (2014) and should be treated with caution. The choice about these language variants does not affect the other coefficients reported in Table 5.

largest differences in GDP per capita between the employer and worker countries increase the wage bill of contracts.

These basic findings continue to hold in Column 2 when also including the level of recent bilateral trade flows. Recent offline trade patterns have modest power for predicting services trade online. We are unable to parse whether the act of trading physical goods has a causal effect in this regard, by for example boosting business connections and reputations for this work, or whether these past trade relationships reflect more primal determinants that we have not modeled or did not measure well. Potential examples include geographic and economic interactions that are more fine-grained than our gravity covariates could pick up or idiosyncratic relationships across countries (good and bad) that are not included in the framework but impact business interactions.

Columns 3 and 4 compare the periods before and after 2011. The role of distance is becoming less pronounced, while GDP differences are becoming more pronounced. As a whole it looks like a typical trade model performs better after 2011, suggesting that platform maturity is somewhat leading digital labor patterns to look more like those observed for other international exchanges.

Finally, Column 5 considers persistence in past online trade, which would be expected as a result of the information friction and path-dependence models reviewed earlier in the literature. Does an initial high share of wage bill pre 2011 continue to explain flows in the later period? The answer is a clear yes, even after controlling for offline conditions that may affect the initial distribution. The elasticity is 0.369, so a 10% increase in pre-2011 trade implies a 3.7% increase in post-2011 digital trade. This connection is consistent with the micro results in Ghani et al. (2014) which show that employers replicate their approach to contracting if it works the first time. The estimates may also be consistent with employers who exploit a locale after resolving uncertainty about its fit or after developing some location-specific knowledge.

3.3 Substitution Elasticities

The results to this point lead us to ask to what extent changes in relative prices overcome some frictions. The first attempt at addressing this question explores substitution patterns across countries. Table 4 suggests that American employers are home-biased and are likely to hire US workers despite their high prices. Here we attempt to quantify how variation in relative prices affects substitution by American employers away from US workers and towards workers from the rest of the world. To do so, we restrict the sample to US employers and estimate how contract

shares vary with mean wage bids. The regression is

$$\ln(s_{jkt}) - \ln(s_{0kt}) = \alpha_0 \bar{W}_{jkt} + \alpha_1 \bar{W}_{jkt} \times US_j + \text{country}_j + \text{time}_t + \text{jobCategory}_k + \varepsilon_{itk},$$

where s_{jkt} is the share of contracts relative to total job openings posted by US employers in job category k filled by workers from country j in time period t , s_{0kt} is the share of openings without a contract, and \bar{W}_{jkt} is the mean hourly wage bid in that cell. The interaction $\bar{W}_{jkt} \times US_j$ allows the coefficient on price to differ for workers from the United States. To account for endogeneity of wage bids, we instrument for bids by non-US workers using the z-score of the log of the local currency to dollar exchange rate. This instrument comes from Stanton and Thomas (2016a) and exploits the fact that all contracts are in US dollars but non-US workers' outside wages are paid in the local currency. The z-score normalization is necessary to account for different scales relative to the dollar across countries. A second instrument is necessary for US workers. Here we use an instrument that is based on common cost shocks across markets, taking the average wage bid for UK workers interacted with a dummy that the bid in question is from the United States.

The estimating equation is the linear IV analogue of a logit model, but the parameters α_0 and α_1 allow for some additional flexibility in assessing substitution patterns across countries relative to a model where the coefficient on price is constrained to be constant across all alternatives. The own-price elasticity for non-US workers, denoted “row” for “rest of world,” is $\alpha_0(1 - s_{row})\bar{W}_{row}$, where s_{row} is the share of contracts to job openings coming from the rest of the world. The own price elasticity for US workers is $(\alpha_0 + \alpha_1)(1 - s_{US})\bar{W}_{US}$. The cross-price elasticity for the rest of the world with respect to US bids is $-(\alpha_0 + \alpha_1)s_{US}\bar{W}_{US}$, and the cross-price elasticity for the United States with respect to bids from the rest of the world is $-\alpha_0 s_{row}\bar{W}_{row}$.

Table 6 provides estimates of substitution patterns across countries using these expressions, along with first-stage regressions for the linear IV estimates. In all specifications, demand for workers from the rest of the world is more elastic than for workers from the United States. The base own-bid elasticity for the rest of the world is -4.61. This says that a 1% increase in average bids leads to a 4.61% decrease in contract share for the rest of the world. Surprisingly, the elasticity is larger in magnitude for technical categories, -8.29, than for non-technical categories, -3.06. The elasticity has also fallen over time. The base own-bid elasticity for US workers, in contrast is -2.14. It is also larger for technical categories and displays a similar decline over time.

The cross-elasticities are of even more interest. We believe this is the first place to document that these elasticities are tiny; suggesting limited substitution across places based on price-related considerations. The cross-elasticity for the rest of the world with respect to US bids is 0.039. This says that a 1% increase in US bids leads to a 0.039% increase in contract share for the rest of the world. This rises to 0.044 in technical categories and has fallen over time. The magnitude of these cross-elasticities is even smaller when looking at the elasticity of US share relative to rest of world bids. Figure 5 provides a visual comparison.

These results suggest limited substitution between the United States and other countries. This lack of substitution suggests that frictions may be quite persistent. Even in a global labor market with limited switching costs, there is very little substitution between the United States and other countries. Instead, given the magnitude of own-bid elasticities, this suggests employers leave the platform in response to bid increases rather than substitute away from their target search location.

4 Additional Digital Collaborations

Our paper mostly concentrates on an empirical depiction of the Upwork platform, but we now turn to some case examples to describe the range of other ways that digital capabilities are extending access to talent over long distances. First, before leaving digital labor markets, it is important to recognize the multiple types of two-sided labor platforms being developed. Founded in 2013, HourlyNerd (not called Catalant) has built an innovative marketplace for management consulting work. It focuses on business consulting and has over 20,000 independent consultants registered for project-based work. Originally targeting ways to connect freelancers with small companies that would not otherwise use consultants, HourlyNerd has grown into fielding enterprise-level solutions that are used by many large companies too. Like management consulting, many areas that appear today to be protected from digital competition may soon become targets of entrepreneurs seeking to build platforms in these areas. Examples from the legal industry are UpCounsel and InCloudCounsel.

Second, online contests and crowd-based mechanisms provide ways for companies to solicit ideas from many unexpected sources. For instance, pharmaceuticals giant Merck designed an eight-week contest in 2012 to aid its drug development process. It released data on chemical compounds that it had previously tested, and then challenged participants to identify which held the most promise for future testing. The winner would receive \$40,000. The contest attracted 238

teams that submitted more than 2,500 proposals. The winning solution came from computer scientists (not professionals in the life sciences) who were using machine-learning approaches previously unknown to Merck. This opened up opportunities for Merck that would not have otherwise been feasible.

Contests held by NASA also illustrate the worldwide span of these human capital inputs (Lakhani et al. 2013). NASA launched a set of pilot projects in 2008 to evaluate the use of global contests and similar crowd-sourced approaches for solving thorny technical challenges that were proving difficult for its internal team. NASA used three for-profit platforms (InnoCentive, TopCoder, Yet2.com) during its early phases for challenges like forecasting of solar events, improved food barrier layers, and compact aerobic resistive device design. Seven challenges posted on InnoCentive illustrate the global engagement, with 2900 problem solvers from 80 countries participating. In many cases, the developed crowd-based solutions were twice as good or better than what the organization had achieved internally. These contests continue to be an important way that NASA sources global talent for its work.

Boudreau and Lakhani (2014) describe further the many ways that contests are used to access far flung ideas and insights. Similar to Upwork and HourlyNerd, many digital platforms like InnoCentive and TopCoder are positioning themselves to be the platforms for companies to reach talented people with ideas, no matter where they live. This breadth of the crowdsourcing platforms, moreover, is critically important for the value they can deliver to clients like Merck and NASA. This is because the quality of the outcome depends not on the average quality of the responses assembled, but instead on the extreme tail of the ideas generated. While internal experts may on average deliver better quality ideas, the extreme values when pulling ideas from a very large external contractor pool are likely to be higher. If it is only the best idea or solution that matters, access to a huge global developer pool can be very advantageous.

Third, as described in introduction, some companies are seeking to establish porous organizational boundaries directly for their businesses. When P&G developed its Connect + Develop platform, it had 7,500 employees worldwide working on innovation-related activities. But, P&G estimated that there were 200 people outside of P&G working on the same topics for each of its scientists, or about 1.5 million people, and it launched its Connect + Develop to be this global outreach. One of its earliest successes was an important innovation for its Pringles line that came from a technology developed in a small bakery in Italy (Huston and Sakkab, 2006). In a similar spirit, companies and developers engaged in open-source software depend upon and contribute to a global common good, where national borders are second-order.

5 Perspectives for Researchers

Data on digital labor markets provide some special advantages for researchers interested in empirical labor topics. First, they often can provide a unique or rare angle on an important topic through their records of bidders and the outside options of both parties, their record of performance outcomes, the ability to construct longitudinal careers for workers, the conduct of skills assessments for workers, and so on. This often allows researchers to attack very complex problems in new ways, providing a unique edge to papers. For example, studies of discrimination have often been perplexed by how to best determine levels of discrimination when observing realized wage differentials in the market, whereas these platforms could allow one to make inference from the observed bids given to an employer and the characteristics of the chosen worker. On the other hand, weighing against this advantage is the fact that these powerful approaches often bring their own complex problems to solve. Continuing the discrimination example, how do you correctly capture the employer’s perceptions of the various performance histories of the bidders?

One limitation of Upwork data for some labor topics is that it is not straightforward to identify corporate firms due to the lack of a unique company identifier. The person hiring within General Motors company, for example, could list many variations on the company’s name or even the name of the subsidiary that they work for. For researchers familiar with patent data, this structure is operationally quite similar to ambiguities with patent assignee codes/names. This structure limits the ability of researchers to describe outsourcing behavior very well across the firm size distribution on Upwork, but for most applications this has limited consequence. Longer-term, it would be very interesting to match digital labor markets data to confidential administrative sources of employer-employee information, like the Longitudinal Employer-Household Database (LEHD) that is developed and maintained by the Census Bureau. Another possibility is the Venture Xpert records on start-ups backed by venture capital. Obviously, overseas freelancers would not be captured, but such merger would allow interesting depictions of local hiring versus outsourced contracts.

Third, these platforms allow experiments to be run in labor markets that are not otherwise feasible (e.g., Pallais, 2014; Cullen and Pakzad-Hurson, 2016). Some of these experiments are conducted at the platform level, changing fundamentally how some aspect of the market “works”—a type of intervention that would be very difficult to conduct in other contexts. For example, Horton (2016b) reports the results of a true minimum wage experiment, while Horton

and Johari (2016) report the results of an experiment in which employers were required to publicly signal their relative preferences over price and quality to would-be applicants. They were able to experimentally manipulate whether the employer’s preferences were communicated to would-be applicants, allowing them to estimate how much additional sorting of workers to the “right” kind of employer occurs when employer preferences are made explicit.

6 Open Questions

The analyses and examples of digital labor markets provided in this review bear witness to an exciting phenomenon in its earliest stage of development. With a focus on high-skilled talent, Freeman (2013) argues increasing globalized knowledge creation and transfer could become the “one ring to rule them all” with respect to international trade in services, financial and capital mobility, and people flows. Perhaps so, and the evidence collected in this review suggests digital labor and talent access could be a central part of such a future. On the other hand, this fate is far from guaranteed, and the ultimate importance of these global forces will only be revealed over the next decade and beyond. We close this review with some open questions to this end.

First, several interesting questions exist about the platforms themselves. Perhaps most important, platforms are still experimenting with the technical designs and algorithms that govern how their labor markets operate, what information is provided to firms and workers, and so on. Many small tweaks are implemented, but some redesigns are quite impactful, such as when oDesk began requiring firms and workers to use a similar skill vocabulary, with implications for the matching efficiency of the platform. The digital platforms have clear incentives to make adjustments that improve their efficiency and competitiveness, and researchers likewise may uncover top-notch natural experiments if they can be closely integrated into these adjustments and their design/implementation. On a related note, complementary tools like Dropbox, Slack, GoogleDocs, etc. are improving the functioning and accelerating the development of digital labor exchanges. We need to learn more about the symbiotic relationship between other collaborative tools and digital labor markets and how the complementary products coevolve. Ownership of data and privacy have not been major concerns thus far, but may take on bigger roles in the future.

Next, many questions exist about how these rapidly-expanding digital labor platforms will impact the broader labor markets and economy around them. At present, the modest size of these labor platforms has not delivered local consequences in advanced economies like those

associated with Uber and Airbnb. As such, there has been less attention to regulatory structures and tax policies for these markets, especially compared to other parts of the shared economy. It is an open and important question about how the policy environment surrounding these companies will adjust as they scale. Similarly, the future interactions—competitive battles, mergers and acquisitions, etc.—with offline outsourcing or temporary help companies or online platforms in adjacent domains will be intriguing to watch. Recent start-ups that focus on online-to-offline work tasks (e.g., Hello Alfred) suggest the current perceived gaps might close faster than expected.

While small in advanced economies from a contracting perspective, the economic impacts in terms of freelancers and their local economies are already more accentuated in some special settings in developing and emerging economies. For example, some remote Russian towns have an abundance of technical talent due to the Cold War and utilize these digital labor platforms to obtain good-paying work globally when none is available in the local economy. Due to local spillovers and the development of agencies, as discussed in Section 2, remote places can even become somewhat known for a certain type of outsourced task, similar to the specialized manufacturing towns in China. More research should go into studying the development of these contractor pools and their local operations. Moreover, comparative studies across specialized places in the face of exchange rate movements and similar shocks will be interesting. On these and similar fronts, studies can be both leading edge in terms of describing an emerging global phenomenon and also on the leading edge in terms of academic insights about important broader economic questions.

References

- Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons, "Digitization and the Contract Labor Market: A Research Agenda," in Avi Goldfarb, Shane Greenstein, and Catherine Tucker (eds) *Economic Analysis of the Digital Economy* (Chicago, IL: University of Chicago Press, 2015), 219-250.
- Agrawal, Ajay, Devesh Kapur, John McHale, and Alexander Oettl, "Brain Drain or Brain Bank? The Impact of Skilled Emigration on Poor-Country Innovation," *Journal of Urban Economics*, 69:1 (2011), 43-55.
- Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons, "Does Information Help or Hinder Job Applicants from Less Developed Countries in Online Markets?," NBER Working Paper 18720 (2014).
- Autor, David, "Wiring the Labor Market," *Journal of Economic Perspectives* 15:1 (2001), 25-40.
- Blinder, Alan, and Alan Krueger, "Alternative Measures of Offshorability: A Survey Approach," *Journal of Labor Economics* 31:S1 (2013), S97-S128.
- Boudreau, Kevin, and Karim Lakhani, "Using the Crowd as an Innovation Partner," *Harvard Business Review* (2014).
- Clemens, Michael, "Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?," *Journal of Economic Perspectives* 25:3 (2011), 83-106.
- Costinot, Arnaud, David Donaldson, and Ivana Komunjer, "What Goods Do Countries Trade? New Ricardian Predictions," *Review of Economic Studies*, 79:2 (2012), 581-608.
- Cullen, Zoe, and Bobak Pakzad-Hurson, "Equal Pay for Equal Work? Evidence from the Renegotiations of Short-Term Work Contracts Online," Working Paper (2016).
- Eaton, Jonathan, and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 70:5 (2002), 1741-1779.
- Freeman, Richard, "One Ring to Rule Them All? Globalization of Knowledge and Knowledge Creation," NBER Working Paper 19301 (2013).
- Ghani, Ejaz, William Kerr, and Christopher Stanton, "Diasporas and Outsourcing: Evidence from oDesk and India," *Management Science*, 60:7 (2014), 1677-1697.
- Grossman, Gene, and Esteban Rossi-Hansberg, "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review* 98:5 (2008), 1978-1997.
- Helpman, Elhanan, "Foreign Trade and Investment: Firm-Level Perspectives," *Economica* 81 (2014), 1-14.
- Helpman, Elhanan, Marc Melitz, and Stephen Yeaple, "Export Versus FDI with Heterogeneous Firms," *American Economic Review* 94 (2004), 300-316.
- Horton, John, "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment," *Journal of Labor Economics*, 35(2), (2017).
- Horton, John, "Online Labor Markets," *International Workshop on Internet and Network Economics*, (2010). 515-522.
- Horton, John, "Buyer Uncertainty about Seller Capacity: Causes, Consequences, and Partial Solution," Working Paper, (2016a).
- Horton, John, "Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment," Working Paper, (2016b).
- Horton, John and Johari, Ramesh, "At What Quality and What Price? Eliciting Buyer Preferences as a Market Design Problem," Working Paper, (2016).

- Horton, John and Tamba, Prasanna, "Labor Economists Get Their Microscope: Big Data and Labor Market Analysis," *Big Data*, 3(3), (2015), 130-137.
- Huston, Larry, and Nabil Sakkab, "Connect and Develop," *Harvard Business Review* (2006).
- Kerr, Sari Pekkala, William Kerr, Çağlar Özden, and Christopher Parsons, "Global Talent Flows," *Journal of Economic Perspectives*, 30:4 (2016), 83-106.
- Kerr, Sari Pekkala, William Kerr, Çağlar Özden, and Christopher Parsons, "High-Skilled Migration and Agglomeration," *Annual Review of Economics*, (Forthcoming).
- Kerr, William, "Ethnic Scientific Communities and International Technology Diffusion," *Review of Economics and Statistics*, 90:3 (2008), 518-537.
- Kerr, William, "Heterogeneous Technology Diffusion and Ricardian Trade Patterns," Working Paper (2016).
- Kerr, William, "The Ethnic Composition of US Inventors," HBS Working Paper 08-006 (2007).
- Kerr, William, and William Lincoln, "The Supply Side of Innovation: H-1B Visa Reforms and US Ethnic Invention," *Journal of Labor Economics*, 28:3 (2010), 473-508.
- Kugler, Maurice, and Hillel Rapoport, "Migration, FDI and the Margins of Trade," Harvard Kennedy School Working Paper (2012).
- Lakhani, Karim, Hila Lifshitz-Assaf, and Michael Tushman, "Open Innovation and Organizational Boundaries: Task Decomposition, Knowledge Distribution and the Locus of Innovation," in Anna Grandori (ed) *Handbook of Economic Organization: Integrating Economic and Organization Theory* (Northampton, MA: Edward Elgar Publishing, 2013), 355-382.
- Lazear, Edward, Kathryn Shaw, and Christopher Stanton, "Who Gets Hired? The Importance of Finding an Open Slot," NBER Working Paper No. 22202 (2016).
- Lyons, Elizabeth, "Team Production in International Labor Markets: Experimental Evidence from the Field," Working Paper (2016).
- Melitz, Jacques, and Farid Toubal. "Native Language, Spoken Language, Translation and Trade," *Journal of International Economics*, 93(2), (2014), 351-363.
- Melitz, Marc, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71 (2003), 1695-1725.
- Mill, Roy, "Hiring and Learning in Online Global Labor Markets," Working Paper (2013).
- Pallais, Amanda, "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review* 104:11 (2014), 3565-3599.
- Ratha, Dilip, and William Shaw, "South-South Migration and Remittances," World Bank Working Paper 102 (2007).
- Rauch, James, and Vitor Trindade, "Ethnic Chinese Networks in International Trade," *Review of Economics and Statistics* 84:1 (2002), 116-130.
- Santos Silva, J.M.C., and Silvana Tenreyro, "Log of Gravity," *Review of Economics and Statistics* 88:4 (2006), 641-658.
- Saxenian, AnnaLee, with Yasuyuki Motoyama and Xiaohong Quan, *Local and Global Networks of Immigrant Professionals in Silicon Valley* (San Francisco, CA: Public Policy Institute of California, 2002).
- Saxenian, AnnaLee, *The New Argonauts* (Cambridge, MA: Harvard University Press, 2006).
- Stanton, Christopher, and Catherine Thomas, "Information Frictions and Observable Experience: The New Employer Price Premium in an Online Market," Working Paper (2016a).

Stanton, Christopher, and Catherine Thomas, "Landing the First Job: The Value of Intermediaries in Online Hiring," *Review of Economic Studies* 83:2 (2016b), 810-854.

Tervio, Marko, "Superstars and Mediocrities: Market Failure in the Discovery of Talent," *Review of Economic Studies* 76:2 (2009), 829-850.

Waugh, Michael, "International Trade and Income Differences," *American Economic Review*, 100 (2010), 2093-2124.

Xu, Gao, "How Does Collective Reputation Affect Hiring? Selection and Sorting in an Online Labour Market," Working Paper (2016).

Appendix: Conceptual Framework for Gravity Model

This section reviews the Eaton and Kortum (2002) model as a theoretical background for a gravity specification for trade.¹⁴ The world consists of N countries producing and consuming a continuum of goods or services $j \in [0, 1]$. In our setting, we think of j as tasks or services that are completed on a digital platform, but we will keep the simple label of “good” throughout this appendix for consistency. Consumers maximize utility in each period by purchasing these goods in quantities $Q(j)$ according to a constant elasticity of substitution (CES) objective function,

$$U = \left(\int_0^1 Q(j)^{(\sigma-1)/\sigma} dj \right)^{\sigma/(\sigma-1)}, \quad (1)$$

subject to prices determined below. $\sigma > 0$ is the elasticity of substitution across goods for the consumers. Consumers earn wage w and consume their full wages in each period. Accordingly, time subscripts are omitted.

Countries are free to produce or trade all goods. Inputs can move among industries within a country but not across countries. Industries are characterized by identical Cobb Douglas production functions employing labor with elasticity α and the continuum of produced goods, also aggregated with (1), with elasticity $1-\alpha$. Factor mobility and identical production functions yield constant input production costs across goods within each country, $c_i(j) = c_i \forall j$.

Technology differences exist across countries, so that country i 's efficiency in producing good j is $z_i(j)$. With constant returns to scale in production, the unit cost of producing good j in country i is $c_i/z_i(j)$. While countries are free to trade, geographic or cultural distance results in “iceberg” transportation costs so that delivering one unit from country i to country n costs $d_{ni} > 1$ units in i . Thus, the delivery to country n of good j made in country i costs

$$p_{ni}(j) = \left(\frac{c_i}{z_i(j)} \right) d_{ni}. \quad (2)$$

An increase in country i 's efficiency for good j lowers the price it must charge. Perfect competition allows consumers to buy from producers in the country offering the lowest price inclusive of shipment costs. Thus, the price that consumers in country n pay for good j is

$$p_n(j) = \min[p_{ni}(j); i = 1, \dots, N]. \quad (3)$$

The technology determining the efficiency $z_i(j)$ is modelled as the realization of a random variable Z_i drawn from a country-specific probability distribution $F_i(z) = \Pr[Z_i < z]$. Draws are independent for each industry j within a country. A core innovation of Eaton and Kortum's model is to use the Fréchet functional distribution to model technologies,

$$F_i(z) = e^{-T_i \cdot z^{-\theta}}, \quad (4)$$

where $T_i > 0$ and $\theta > 1$. The country-specific parameter T_i determines the location of the distribution, while the common parameter θ determines the variation within each country's distribution. By the law of large numbers, a larger T_i raises the average efficiency of industries for country i , and therefore its absolute advantage for trade. A larger θ , on the other hand, implies a tighter distribution for industries within every country and thereby limits the scope for comparative advantage across nations.

The Fréchet distribution (4) allows prices from equations (2) and (3) to be determined. The probability that country i is the lowest-cost producer of an arbitrary good for country n is

¹⁴Costinot et al. (2012) extend Eaton and Kortum (2002) to articulate appropriate industry-level estimations of Ricardian advantages as a source of trade among countries.

$\pi_{ni} = T_i(c_i d_{ni})^{-\theta} / \sum_{k=1}^N T_k(c_k d_{nk})^{-\theta}$.¹⁵ With a continuum of goods, π_{ni} is also the fraction of goods country n purchases from country i . Country n 's average expenditure per good does not vary by source country, so that the fraction of country n 's expenditure on goods from country i is also

$$\frac{X_{ni}}{X_n} = \frac{T_i(c_i d_{ni})^{-\theta}}{\sum_{k=1}^N T_k(c_k d_{nk})^{-\theta}}, \quad (5)$$

where X_n is total expenditure in country n . Holding input prices constant, technology growth in country i increases its exports to country n through entry into industries in which it was previously uncompetitive. Looking across import destinations for an industry in which it already exports, country i also becomes the lowest-cost producer for more distant countries it could not previously serve due to the markup of transportation costs. Condition (5) also shows how trading costs d lead to deviations in the law of one price.

Defining Q_i to be the total sales of exporter i , Eaton and Kortum (2002) show how bilateral exports can be expressed as

$$X_{ni} = \frac{(d_{ni}/p_n)^{-\theta} X_n}{\sum_{k=1}^N (d_{ki}/p_k)^{-\theta} X_k} Q_i, \quad (6)$$

where p_i is the price level of a country i . This equation shows how the trade connects with the aggregate size of the importer (X_n), the exporter (Q_i), and the price-adjusted distances between them (d_{ni}/p_n). The allocation of trade has an intuitive feel. The share of total exports of country i (Q_i) that go to country n is determined by how country n 's size, bilateral distance, and prices compare to the other countries in the world, with the latter being summarized in the denominator through the summation of countries.

Rearranging this for the purposes of estimation, we have

$$\log(X_{ni}) = \log(Q_i) - \theta \log(d_{ni}/p_n) + \log(X_n) - \log\left(\sum_{k=1}^N (d_{ki}/p_k)^{-\theta} X_k\right)$$

or

$$\log(X_{ni}/X_n) = \log(Q_i) - \theta \log(d_{ni}/p_n) - \log\left(\sum_{k=1}^N (d_{ki}/p_k)^{-\theta} X_k\right).$$

The last term is a worldwide constant term that would be captured by intercepts or fixed effects in estimation.

Reflecting on this model, there are parts of it that are not well suited to thinking about a digital labor market. For example, the model assumes balanced trade across goods and that all goods are represented, but we are examining only a small slice of economic activity and there is no trade balance. On the other hand, the choice to contract on these platforms may be closer to the perfect competition and distance assumptions than other settings. This provides some context and grounding for applying the gravity equation in our empirical work.

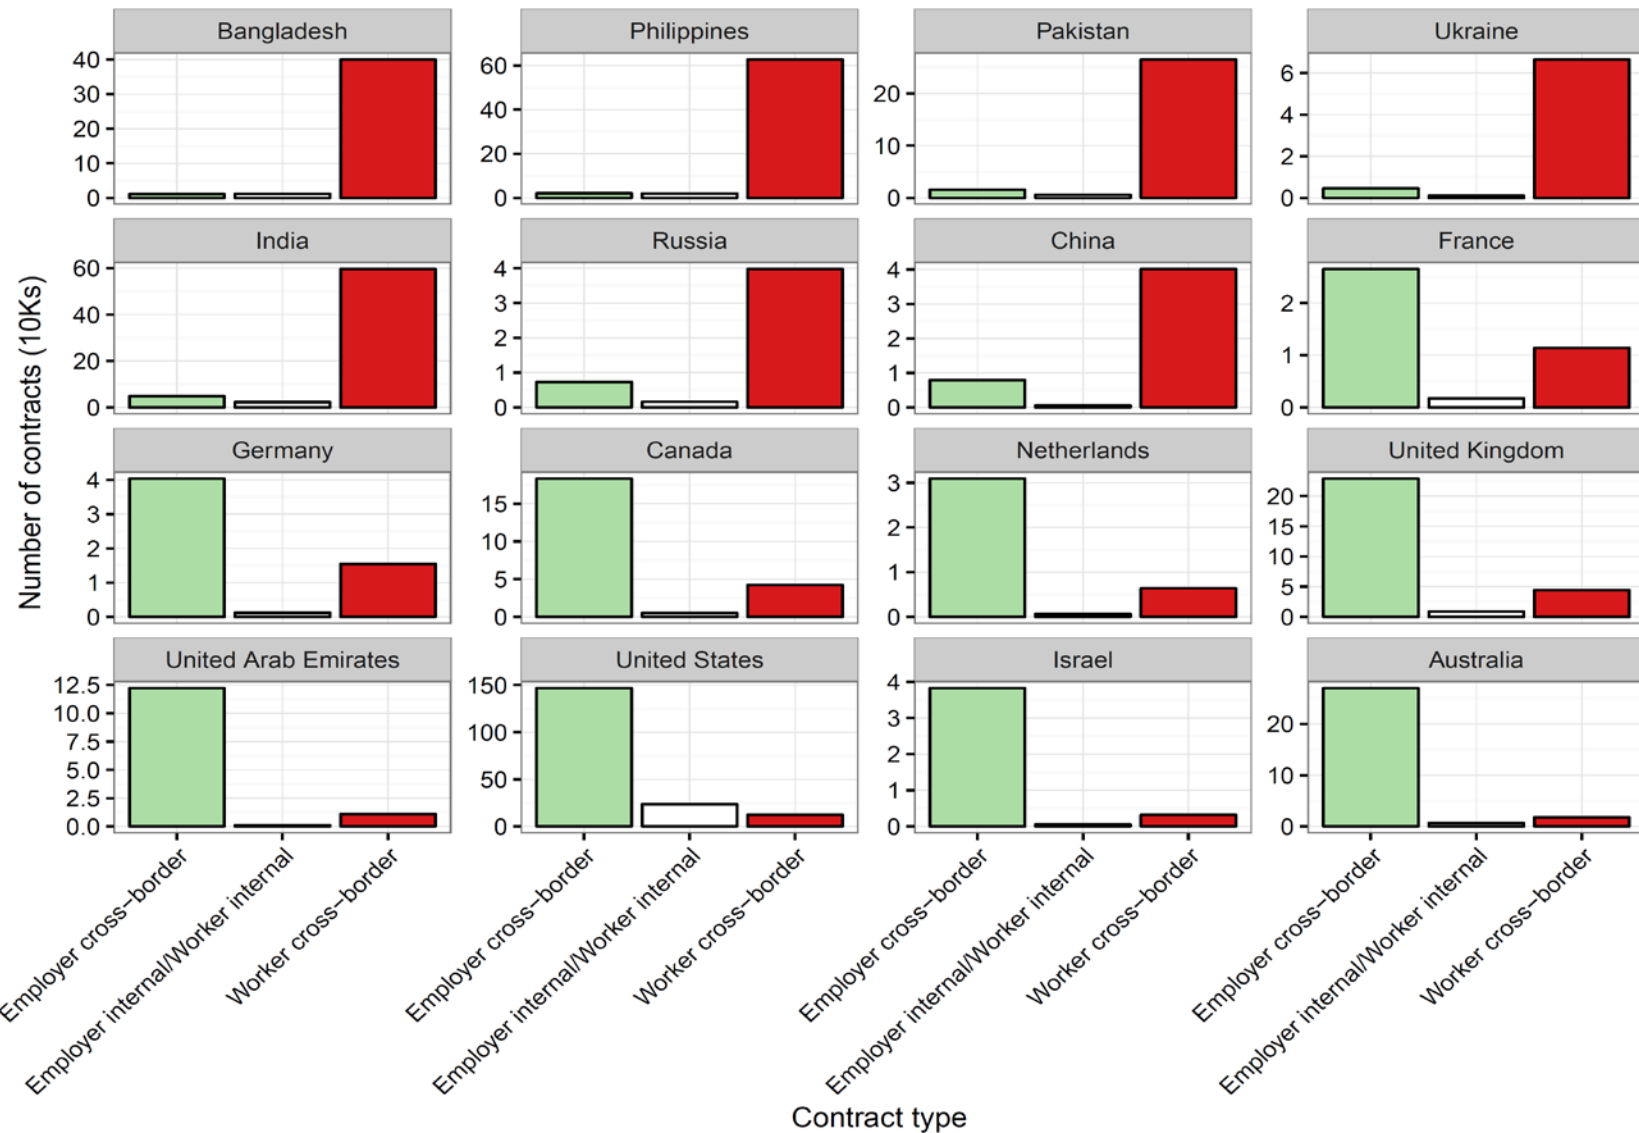
¹⁵The distribution of prices country i presents to country n is $G_{ni}(p) = \Pr[P_{ni} \leq p] = 1 - F_i(c_i d_{ni}/p) = 1 - \exp(-T_i(c_i d_{ni})^{-\theta} p^\theta)$. Country n buys from the lowest cost producer of each good, so that its realized price distribution is $G_n(p) = \Pr[P_n \leq p] = 1 - \prod_{i=1}^N [1 - G_{ni}(p)] = 1 - \exp(-p^\theta \sum_{i=1}^N T_i(c_i d_{ni})^{-\theta})$. The probability is $\pi_{ni} = \Pr[P_{ni}(j) \leq \min\{P_{ns}(j); s \neq i\}] = \int_0^\infty \prod_{s \neq i} [1 - G_{ns}(p)] dG_{ni}(p)$. See Eaton and Kortum (2002) for the full derivation of the price index.

Figure 1: Example of a worker profile from a digital labor platform

The screenshot shows a worker profile for Naveen Bali on the oDesk platform. The header includes the oDesk logo with the tagline "Love the way you work.", navigation links for "Hire Contractors", "Find Work", and "How It Works", and buttons for "Log in" and "Sign up". A "Back to search results" link is also present. The profile card for Naveen Bali features a profile picture, his name, a list of skills (PHP, MySQL, JS, Ajax, web scraping, JQuery, codeignitor), his location (Panchkula, India), his hourly rate (\$22.22/hr), a 5-star rating, and statistics showing 30 jobs completed and 1,234 oDesk hours. Action buttons for "Post a job" and "Contact" are visible, along with a "What's the difference" link. Below the profile card, the "Skills" section lists php, mysql, ajax, jquery, and html, with a "more..." link. The "Tests Taken" section shows "PHP4 Test" at the 92nd percentile and "PHP5 Test" at the 90th percentile. The "Overview" section contains a detailed bio: "Master in Project development i have deep knowledge of every stage of software development lifecycle . My last job was at Yahoo banglore my designation was senior software engineer in web development department. my area of skills is PHP5.0,mysql,javascript,jQuery,Ajax,JSON,HTML5.0 i have good understanding of MVC Architect i have worked in codeigniter and Zend framework . i have completed lot of projects in web scraping as well". Below this, it states "Naveen Bali has added 7 portfolio pieces. Create an account to review them." The "Work History & Feedback" section lists three jobs: "PHP Development & Support" (May 2012 - Jun 2012, \$3,030.56, 124 hrs @ \$24.44/hr, 5-star feedback), "Data Scraping" (Mar 2012 - May 2012, \$315.44 fixed-price, 5-star feedback), and "Web Scraping / Data Harvesting" (Sep 2011 - Mar 2012, \$582.05 fixed-price, 5-star feedback). On the right, a "Search for others" section includes a search box and a list of categories: PHP Developers, Zend Developers, CSS Designers, Smarty Developers, and MySQL Administrators.


Notes: Example of a worker profile from the oDesk platform.

Figure 2: The asymmetric nature of Upwork global contract flows



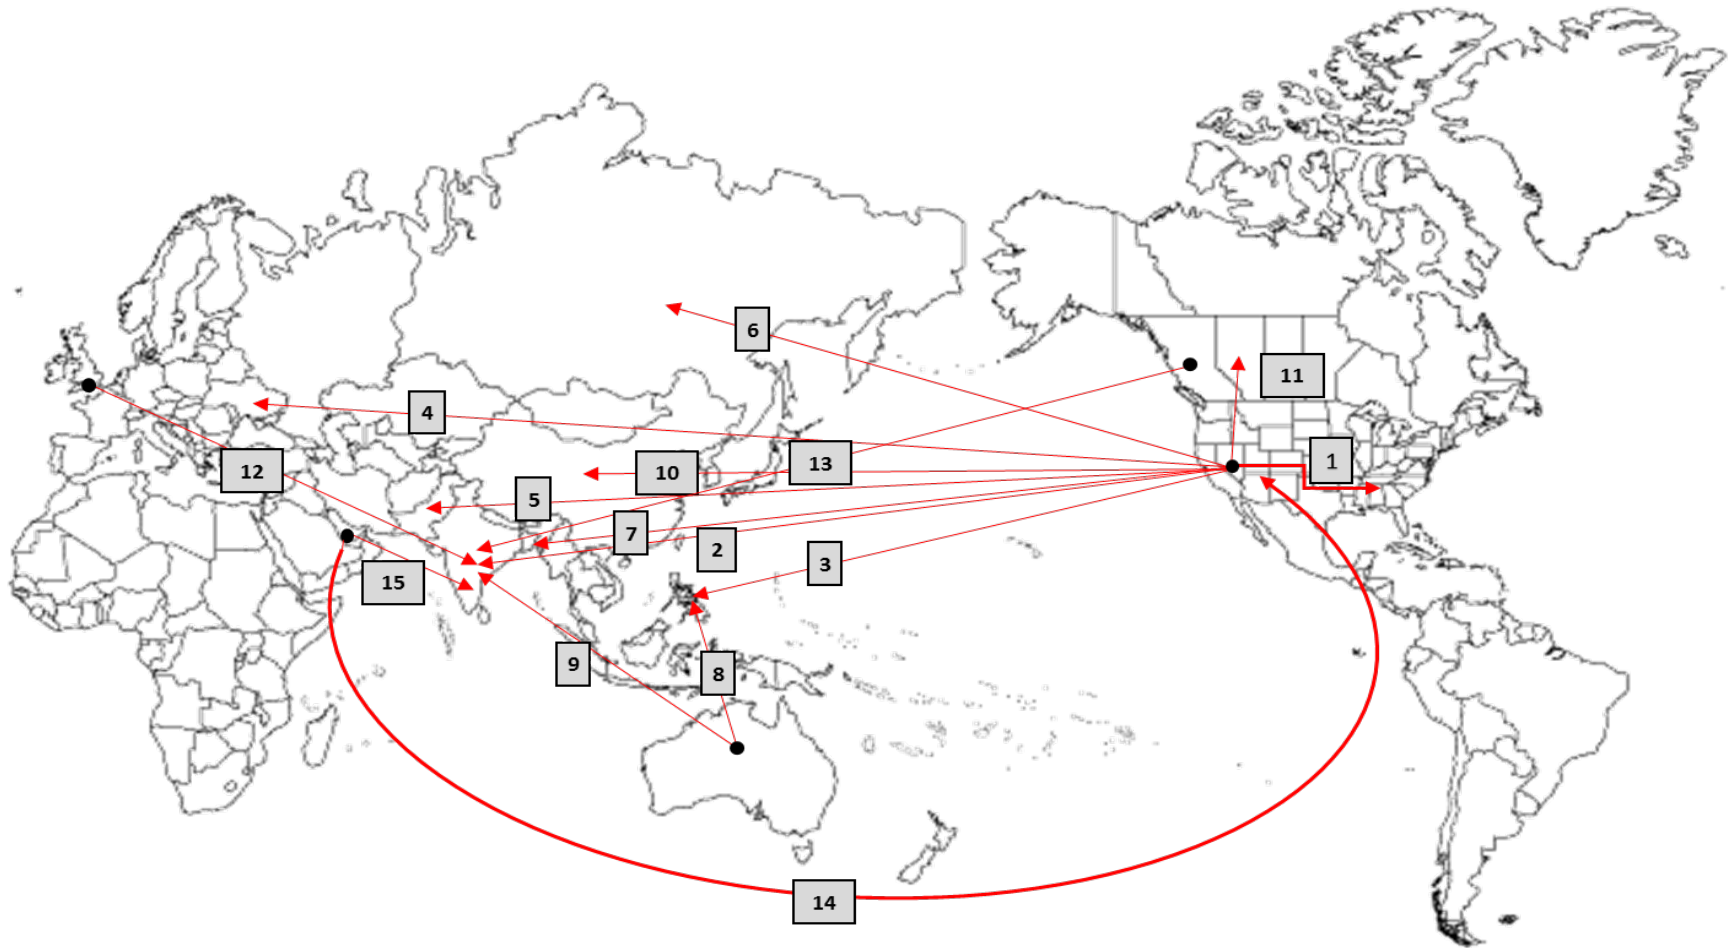
Notes: The figure quantifies the number of contracts for selected countries by their cross-border nature. The vertical axis is specific to each country and measured in tens of thousands. The first bar in each triplet captures outbound contracts made by employers in the country to workers in another country. The second bar reflects contracts made by employers to workers in their own country. The third bar measures the contracts completed by workers in the country where the employer is in another country.

Figure 3a: Top employer-worker contract flows on Upwork platform by contract volume




Notes: The figure shows the top 15 employer-worker contract flows on oDesk/Upwork by contract volume through 2014. The arrow points to the location of the worker. The number indicates the rank of the flow, and Appendix Table 1 provides values. For the #3 case, the third largest flow is the United States to itself.

Figure 3b: Top employer-worker contract flows on Upwork platform by wage bill



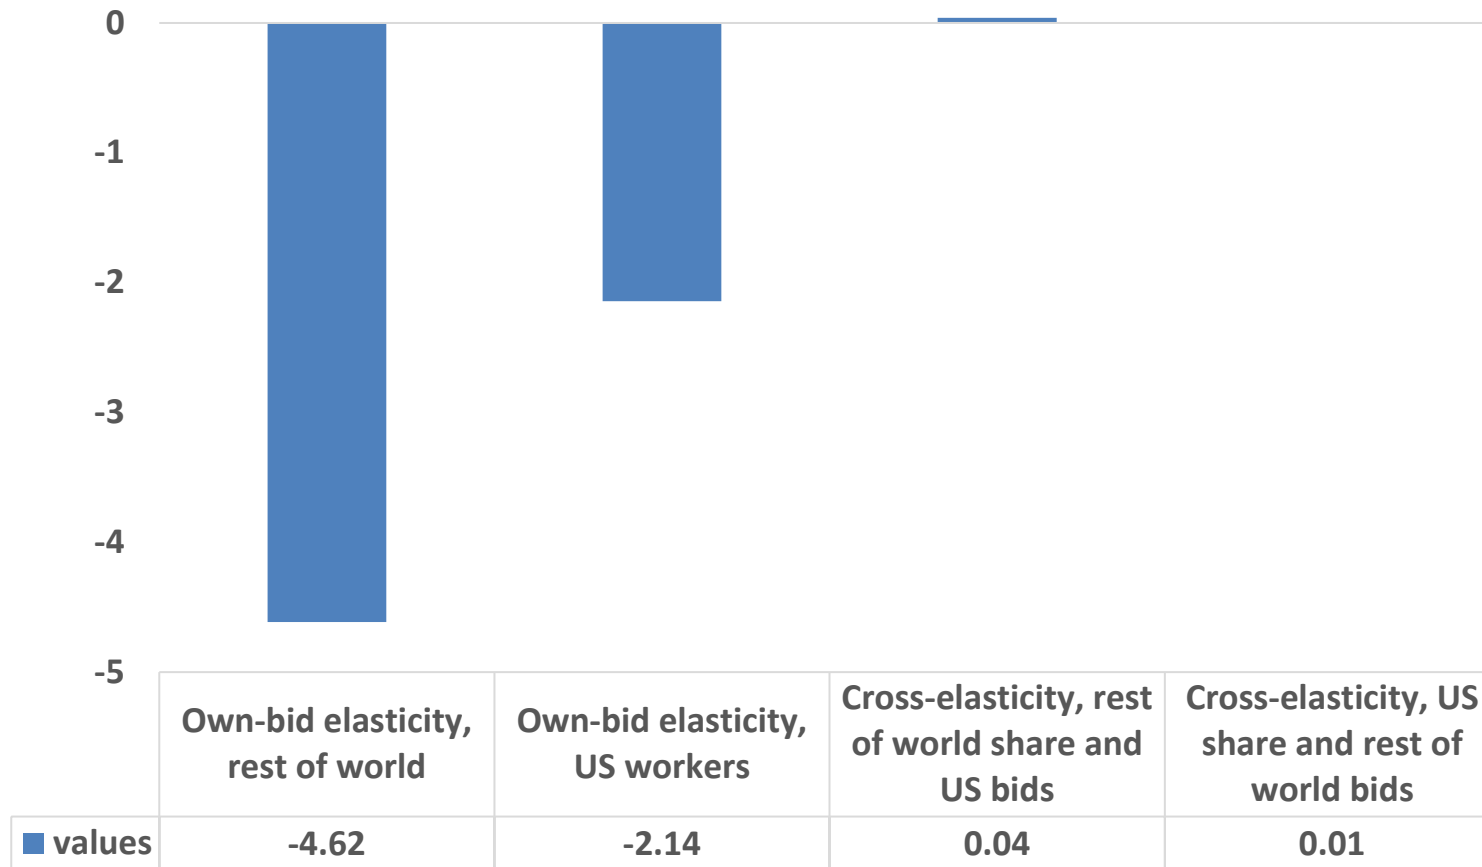
Notes: The figure shows the top 15 employer-worker contract flows on oDesk/Upwork by wage bill through 2014. The arrow points to the location of the worker. The number indicates the rank of the flow, and Appendix Table 1 provides values. For the #1 case, the largest flow is the United States to itself.

Figure 4: Comparison of Upwork's global sourcing distribution for US employers to that for goods and services imports




Notes: Squared deviations of the share of Upwork wage bill paid by US employers to a country against the US share of imports of goods and services from that country. Services imports data come from the World Bank TSD database and are last reported by country in 2011. Goods imports data come from the census. To avoid compositional changes in the series over time, the goods and services series are restricted to be balanced. Deviations of Upwork shares are calculated against the balanced series.

Figure 5: Elasticities of work to own-bid and cross-bids



Notes: See Panel A of Table 6.

Appendix Figure 1: Hours of work on Upwork platform



Notes: Taken from Horton (2016a). Panel A describes application timing from the perspective of Pacific Standard Time on the West Coast of the United States. Panel B maps from the perspective of the contractor. Many countries have workers who adjust their work schedules to that of the United States.

Table 1: Summary statistics by job category

Job Category	Number of job openings	Number of unique applicants	Number of contracts	Number of cross-border contracts	Cross-border contract share	Wage bill, (\$ mil)	Wage bill from cross-border contracts, (\$ mil)	Cross-border wage bill share
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Administrative Support	3,741,563	22,443,920	484,244	435,450	90%	184.7	155.5	84%
Business Services	853,436	2,741,608	95,699	79,552	83%	78.3	59.4	76%
Customer Service	436,211	2,232,582	50,541	43,665	86%	89.3	70.9	79%
Design & Multimedia	5,339,055	14,267,825	498,949	457,700	92%	106.2	92.9	87%
Networking & Info. Systems	646,733	1,431,241	47,445	43,108	91%	39.3	31.8	81%
Sales & Marketing	2,797,382	15,416,597	376,295	342,496	91%	141.8	121.0	85%
Software Development	2,736,048	8,886,761	230,407	217,343	94%	307.5	285.0	93%
Web Development	7,918,424	30,365,649	772,828	724,107	94%	607.6	561.5	92%
Writing & Translation	4,163,107	10,566,662	598,715	480,517	80%	167.4	107.2	64%
Total	28,631,959	108,352,845	3,155,123	2,823,938	90%	1,722	1,485	86%

Notes: Data come from oDesk/Upwork from the launch of the platform through 2014. Wage bill is in millions of US dollars at the time of recording.

Table 2: Estimates of contract volumes formed on oDesk with workers in India

	Dependent variable is share of oDesk contracts formed with workers in India							DV is India's share of dollar value of contracts for country	DV is share of company contacts with Indian ethnic name
	Base estimation	Including distance covariates only	Weighting by log population	Unweighted estimation	Adding worldwide oDesk contracts	Adding trade with India control	Excluding UAE outlier firm		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Indian diaspora share of country population	1.090*** (0.197)	0.728*** (0.156)	0.969*** (0.364)	0.850* (0.429)	1.135*** (0.218)	1.004*** (0.236)	0.531** (0.204)	0.694* (0.360)	2.577*** (0.188)
Indicator for geographical distance to India of 5,000-10,000 kilometers	0.071** (0.030)	0.041 (0.026)	0.090** (0.043)	0.087** (0.044)	0.069** (0.031)	0.077** (0.033)	0.059* (0.030)	-0.033 (0.056)	0.004 (0.058)
Indicator for geographical distance to India of >10,000 kilometers	0.095*** (0.029)	0.088*** (0.030)	0.100** (0.039)	0.092** (0.041)	0.119*** (0.027)	0.102*** (0.033)	0.087*** (0.028)	0.058 (0.063)	-0.074 (0.047)
Log population	-0.009 (0.007)		-0.016* (0.009)	-0.017* (0.010)	0.010 (0.008)	-0.009 (0.007)	-0.010 (0.006)	0.011 (0.013)	0.000 (0.007)
Log GDP per capita	-0.042** (0.022)		-0.051* (0.027)	-0.045 (0.029)	-0.008 (0.025)	-0.041* (0.023)	-0.034 (0.022)	0.010 (0.039)	-0.078** (0.035)
Log telephone lines per capita	0.004 (0.034)		0.002 (0.039)	-0.004 (0.038)	0.005 (0.030)	0.004 (0.034)	-0.005 (0.034)	-0.056 (0.055)	0.052 (0.041)
Overall fit of project profile with India's worker profile	0.078** (0.039)		0.070 (0.047)	0.054 (0.054)	0.085** (0.039)	0.076* (0.040)	0.074* (0.039)	0.015 (0.057)	-0.032 (0.083)
Log count of oDesk contracts worldwide					-0.027*** (0.009)				
Trade with India as share of GDP						0.660 (1.046)			
Observations	92	92	92	92	92	92	92	92	92
Mean of DV	0.341	0.341	0.341	0.341	0.341	0.341	0.338	0.372	0.096
Relative effect (1% diaspora share)	0.032	0.021	0.028	0.025	0.033	0.029	0.016	0.019	0.268

Notes: Country-level regressions in Columns 1-7 estimate traits associated with a larger share of work being contracted to India. Column 8 considers shares based upon dollar values. Column 9 considers the share of company contacts placing contracts that have an ethnically Indian name. Regressions weight by log number of worldwide contracts formed on oDesk, unless otherwise noted, and report robust standard errors.

Table 3A: Hiring and working patterns for top hiring countries on oDesk/Upwork

Country	Employer wage-bill rank from cross-border contracts	Worker wage-bill rank from cross-border contracts	Number of cross-border hiring contracts	Wage-bill from cross-border hiring, (\$ mil)	Number of cross-border worker supply contracts	Wage-bill from cross-border worker supply, (\$ mil)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
United States	1	7	1,468,476	964.6	123,157	56.0
Australia	2	24	269,941	138.0	17,499	8.5
United Kingdom	3	11	229,056	92.6	44,201	23.0
Canada	4	9	183,206	86.6	42,332	30.4
United Arab Emirates	5	45	122,343	69.7	10,939	3.5
Germany	6	19	40,392	19.1	15,456	10.7
France	7	33	26,494	18.5	11,356	5.7
Netherlands	8	40	30,933	14.8	6,379	4.6
Israel	9	57	38,285	13.3	3,202	2.3
Ireland	10	58	17,984	13.0	4,525	2.2
Denmark	11	72	13,119	11.4	2,179	1.4
Switzerland	12	64	18,428	11.1	1,532	1.7
Sweden	13	46	13,980	10.0	6,045	3.4
Spain	14	23	19,295	9.1	12,803	8.9
Singapore	15	56	31,820	9.0	4,713	2.3
New Zealand	16	55	16,772	8.4	4,422	2.4
Hong Kong	17	60	15,320	8.2	3,013	2.0
Norway	18	66	9,344	7.4	2,075	1.7
Belgium	19	61	11,263	6.1	2,547	2.0
Italy	20	27	13,373	4.3	12,039	6.9

Notes: See Table 1. The top 20 countries by hiring employer wage bill are displayed.

Table 3B: Hiring and working patterns for top working countries on oDesk/Upwork

Country	Employer wage-bill rank from cross-border contracts	Worker wage-bill rank from cross-border contracts	Number of cross-border hiring contracts	Wage-bill from cross-border hiring, (\$ mil)	Number of cross-border worker supply contracts	Wage-bill from cross-border worker supply, (\$ mil)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
India	22	1	48,236	3.4	595,980	340.3
Philippines	41	2	20,573	1.2	627,497	286.9
Ukraine	37	3	4,526	1.4	66,436	118.3
Russia	25	4	7,292	3.1	39,754	89.2
Pakistan	45	5	15,480	0.9	265,127	87.3
Bangladesh	75	6	11,078	0.3	399,845	62.5
United States	1	7	1,468,476	964.6	123,157	56.0
China	30	8	7,962	2.2	40,153	38.1
Canada	4	9	183,206	86.6	42,332	30.4
Poland	38	10	3,967	1.4	13,529	25.5
United Kingdom	3	11	229,056	92.6	44,201	23.0
Belarus	119	12	356	0.1	9,799	18.6
Romania	46	13	5,523	0.9	32,769	17.8
Vietnam	89	14	1,832	0.2	16,929	13.3
Indonesia	55	15	2,941	0.6	26,272	11.5
Argentina	64	16	2,043	0.5	10,228	10.9
Serbia	78	17	2,253	0.3	20,196	10.8
Armenia	100	18	734	0.1	8,918	10.7
Germany	6	19	40,392	19.1	15,456	10.7
Egypt	56	20	5,288	0.6	26,445	10.1

Notes: See Table 1. The top 20 countries by worker wage bill are displayed.

Table 4: Top countries supplying work to American employers

Country	Worker wage-bill rank	Number of work supply contracts, total	Wage-bill from work supply, total, (\$ mil)	Wage-bill from work supply, 2005-2011, (\$ mil)	Wage-bill from work supply, 2012-2014, (\$ mil)	Total US imports of goods, (\$ mil)	Total US imports of services, (\$ mil)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
United States	1	235,225	221.7	67.7	154.0	n.a.	n.a.
India	2	317,731	211.6	59.2	152.4	158,462	53,945
Philippines	3	358,671	193.0	48.4	144.6	51,737	8,362
Ukraine	4	30,612	67.4	20.0	47.4	8,230	n.r.
Pakistan	5	140,552	58.4	15.4	42.9	21,345	10,118
Russia	6	19,305	50.1	16.7	33.4	144,435	17,658
Bangladesh	7	218,882	39.7	6.5	33.1	23,322	n.r.
China	8	20,055	23.4	3.8	19.7	2,007,688	46,240
Canada	9	25,264	21.2	6.3	15.0	1,778,196	177,874
Poland	10	6,208	16.5	4.8	11.7	16,446	9,476
United Kingdom	11	22,265	14.2	3.5	10.6	317,506	286,063
Belarus	12	4,444	9.8	2.9	6.9	3,753	275
Romania	13	14,447	9.8	2.1	7.7	6,474	n.r.
Argentina	14	5,516	8.1	2.5	5.6	26,484	6,612
Vietnam	15	7,836	7.7	1.5	6.2	76,744	198
Indonesia	16	12,735	7.1	1.8	5.3	92,053	1,598
Brazil	17	4,773	6.6	1.7	5.0	158,228	21,621
Egypt	18	11,534	6.4	1.3	5.1	13,498	n.r.
Armenia	19	3,949	6.2	1.8	4.5	368	n.r.
Australia	20	9,444	6.1	2.1	4.0	54,245	28,384

Notes: The top 20 countries by worker wage bill paid by US employers are displayed. Data come from oDesk/Upwork from the launch of the platform through 2014. Columns 7 and 8 use external data from the Census and the World Bank TSD database and take totals over data from 2006-2011. 2011 is the last year of services imports data with a country breakdown. Although trade in goods data is available through later periods, data ends in 2011 to maintain comparability between the goods and services series. Missing services data are not reported in the TSD database (n.r.).

Table 5: Estimates of wage-bill volume formed on Upwork

	DV is wage bill from cross-country contracts paid by employer country				
	Estimation without historical trade	Base estimation	Base estimation, pre 2011	Base estimation, post 2011	Base estimation, post 2011 with lag
	(1)	(2)	(3)	(4)	(5)
(0,1) Geographic distance, quartile 2	-0.281*** (0.00150)	-0.245*** (0.00146)	-0.324*** (0.00673)	-0.274*** (0.00242)	-0.162*** (0.00299)
(0,1) Geographic distance, quartile 3	-0.425*** (0.00184)	-0.386*** (0.00188)	-0.551*** (0.00912)	-0.389*** (0.00323)	-0.172*** (0.00424)
(0,1) Geographic distance, quartile 4 (longest)	-0.539*** (0.00214)	-0.449*** (0.00243)	-0.752*** (0.0108)	-0.438*** (0.00412)	-0.142*** (0.00528)
(0,1) Employer - Worker GDP/Cap > \$5k	0.343*** (0.00368)	0.367*** (0.00366)	1.254*** (0.0172)	0.135*** (0.00464)	-0.133*** (0.00442)
(0,1) Employer - Worker GDP/Cap > \$10k	1.117*** (0.00498)	1.087*** (0.00466)	0.419*** (0.0149)	1.087*** (0.0112)	0.825*** (0.00943)
(0,1) Past trade volume, quartile 2		0.0100** (0.00466)	-0.105*** (0.0370)	0.0773*** (0.00515)	0.194*** (0.00476)
(0,1) Past trade volume, quartile 3		-0.0836*** (0.00501)	-0.328*** (0.0427)	0.0145*** (0.00555)	0.123*** (0.00560)
(0,1) Past trade volume, quartile 4 (largest)		0.0867*** (0.00491)	-0.0520 (0.0421)	0.131*** (0.00550)	0.170*** (0.00563)
(0,1) Zero historical trade		-0.357*** (0.0197)	-0.495 (0.309)	-0.324*** (0.0206)	-0.0622*** (0.0175)
(0,1) Common country border	-0.449*** (0.00555)	-0.414*** (0.00559)	-0.477*** (0.0221)	-0.430*** (0.00864)	-0.336*** (0.00767)
(0,1) Common country language	0.198*** (0.00319)	0.204*** (0.00308)	0.295*** (0.00785)	0.226*** (0.00478)	0.116*** (0.00294)
Time zone difference	-0.0114*** (0.000153)	-0.0177*** (0.000137)	0.0249*** (0.000841)	-0.0268*** (0.000195)	-0.0381*** (0.000172)
(0,1) Both the employer and worker countries are in the WTO	1.645*** (0.266)	1.835*** (0.127)	1.916*** (0.529)	1.828*** (0.173)	1.144*** (0.108)
Lag of log wage bill for country pair					0.369*** (0.00237)
Observations	19,430	18,143	13,330	17,485	17,485
Mean of dependent variable (\$ mil)	0.0758	0.0811	0.0288	0.0622	0.0622

Notes: Estimates are weighted by total employer country wage bill. Robust standard errors are reported. Estimation is via Poisson pseudo maximum likelihood and includes employer country and worker country fixed effects. Employer country fixed effects are concentrated out, and worker country fixed effects are included as dummies. All models also include dummies for quartiles of the product of gap between countries (not reported). Historical trade volume is flows taken from the cepii TRADHIST dataset and averaged over 2001-2004. Other gravity covariates come from the cepii Gravity dataset. Column 5 includes an indicator for zero trade prior to 2011 in the oDesk data and the lag of log wage bill is set to zero in these cases.

Table 6: Estimates of contract elasticities for US employers

	See panel headings for dependent variable				
	Base	Technical categories	Non-technical categories	Base estimation 2005-2011	Base estimation 2011-2014
	(1)	(2)	(3)	(4)	(5)
Panel A: Main Regression. Dependent variable is log share of contracts less the no-hire share					
Wage bid average	-0.316*** (0.0298)	-0.416*** (0.0596)	-0.261*** (0.0313)	-0.629*** (0.186)	-0.268*** (0.0247)
Wage bid average x 1(US worker)	0.211*** (0.0189)	0.215*** (0.0340)	0.204*** (0.0216)	0.257*** (0.0682)	0.188*** (0.0171)
Own-bid elasticity, rest of world	-4.615	-8.291	-3.056	-7.990	-4.140
Own-bid elasticity, US workers	-2.144	-5.557	-0.953	-6.581	-1.808
Cross-elasticity, rest of world share and US bids	0.0387	0.0442	0.0222	0.0981	0.0370
Cross-elasticity, US share and rest of world bids	0.00685	0.00964	0.00507	0.0135	0.00577
Panel B: First stage regression for wage bid average					
Z-score of log local currency to dollar exchange rate	-0.597*** (0.0525)	-0.447*** (0.0900)	-0.604*** (0.0610)	-0.651*** (0.0870)	-0.515*** (0.0548)
Log wage bid average in the United Kingdom x Worker in United States	0.361*** (0.0379)	0.278*** (0.0551)	0.423*** (0.0409)	0.222*** (0.0489)	0.566*** (0.0328)
Panel C: First stage regression for wage bid x 1(US worker)					
Z-score of log local currency to dollar exchange rate	-0.0164*** (0.00403)	-0.0289*** (0.00782)	-0.00630 (0.00383)	-0.0214*** (0.00601)	-0.00439*** (0.000804)
Log wage bid average in the United Kingdom x Worker in United States	0.556*** (0.0496)	0.518*** (0.0870)	0.595*** (0.0438)	0.336*** (0.0595)	0.877*** (0.0331)
Observations	33,604	11,862	21,742	10,311	23,293

Notes: Table provides estimations of contract elasticities for US employers. The unit of observation is the country-job category-month of contracts with US employers. Regressions include worker country-by-job category and year-by-month fixed effects. Z-scores of the local currency to US dollar exchange rate are used as instruments for the mean of the bid. The log of the average UK wage bid interacted with a dummy for workers from the United States is an instrument for the wage bid average for US workers. Robust standard errors reported.

Appendix Table 1: Top employer-worker routes on Upwork

Contract volume				Wage bill (\$ mil)			
Rank	Worker country	Employer country	Contracts	Rank	Worker country	Employer country	Wage bill
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	Philippines	United States	358,671	1	United States	United States	221.7
2	India	United States	317,731	2	India	United States	211.6
3	United States	United States	235,225	3	Philippines	United States	193.0
4	Bangladesh	United States	218,882	4	Ukraine	United States	67.4
5	Pakistan	United States	140,552	5	Pakistan	United States	58.4
6	Philippines	Australia	71,139	6	Russia	United States	50.1
7	India	Australia	59,339	7	Bangladesh	United States	39.7
8	India	United Kingdom	47,279	8	Philippines	Australia	39.3
9	Philippines	United Kingdom	40,178	9	India	Australia	32.1
10	Philippines	Canada	36,838	10	China	United States	23.4
11	India	Canada	36,500	11	Canada	United States	21.2
12	Bangladesh	United Kingdom	32,957	12	India	United Kingdom	21.1
13	Bangladesh	Australia	30,667	13	India	Canada	18.5
14	Ukraine	United States	30,612	14	United States	United Arab Emirates	17.2
15	India	United Arab Emirates	28,515	15	India	United Arab Emirates	16.9
16	Canada	United States	25,264	16	Poland	United States	16.5
17	Bangladesh	Canada	24,688	17	Philippines	Canada	14.7
18	India	India	23,294	18	United Kingdom	United States	14.2
19	Philippines	United Arab Emirates	23,200	19	Philippines	United Kingdom	11.2
20	Pakistan	United Kingdom	22,919	20	Belarus	United States	9.8
21	United Kingdom	United States	22,265	21	Romania	United States	9.8
22	Pakistan	Australia	20,804	22	United States	Australia	8.2
23	Philippines	Philippines	20,250	23	Argentina	United States	8.1
24	China	United States	20,055	24	Vietnam	United States	7.7
25	United States	Australia	19,750	25	United States	Canada	7.6
26	Russia	United States	19,305	26	Russia	United Kingdom	7.3
27	United States	Canada	19,270	27	Ukraine	United Kingdom	7.2
28	United States	United Kingdom	18,360	28	Indonesia	United States	7.1
29	Pakistan	Canada	17,200	29	Brazil	United States	6.6
30	Kenya	United States	16,461	30	Ukraine	Canada	6.6

Notes: See Figures 3A and 3B.