

NBER WORKING PAPER SERIES

DIGITIZATION AND THE CONTRACT LABOR MARKET:
A RESEARCH AGENDA

Ajay Agrawal
John Horton
Nicola Lacetera
Elizabeth Lyons

Working Paper 19525
<http://www.nber.org/papers/w19525>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2013

This research was funded by the Centre for Innovation and Entrepreneurship at the Rotman School of Management and the Social Sciences and Humanities Research Council of Canada. We thank Shane Greenstein, Catherine Tucker, Avi Goldfarb, and participants at the NBER pre-conference and conference meetings for valuable feedback. Errors remain our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w19525.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Ajay Agrawal, John Horton, Nicola Lacetera, and Elizabeth Lyons. 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.

Digitization and the Contract Labor Market: A Research Agenda
Ajay Agrawal, John Horton, Nicola Lacetera, and Elizabeth Lyons
NBER Working Paper No. 19525
October 2013
JEL No. J24,J61,O15,O33

ABSTRACT

Online contract labor globalizes traditionally local labor markets, with platforms that enable employers, most of whom are in high-income countries, to more easily outsource tasks to contractors, primarily located in low-income countries. This market is growing rapidly; we provide descriptive statistics from one of the leading platforms where the number of hours worked increased 55% from 2011 to 2012, with the 2012 total wage bill just over \$360 million. We outline three lines of inquiry in this market setting that are central to the broader digitization research agenda: 1) How will the digitization of this market influence the distribution of economic activity (geographic distribution of work, income distribution, distribution of work across firm boundaries)?; 2) What is the magnitude and nature of information frictions in these digital market settings as reflected by user responses to market design features (allocation of visibility, investments in human capital acquisition, machine-aided recommendations)?; 3) How will the digitization of this market affect social welfare (increased efficiency in matching, production)? We draw upon economic theory as well as evidence from empirical research on online contract labor markets and other related settings to motivate and contextualize this research agenda.

Ajay Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
ajay.agrawal@rotman.utoronto.ca

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

Nicola Lacetera
University of Toronto
105 St. George Street
Toronto, ON M5S 2E9
Canada
and NBER
nicola.lacetera@utoronto.ca

Elizabeth Lyons
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
liz.lyons09@rotman.utoronto.ca

1 Introduction

We motivate this chapter on the digitization of the market for contract labor with three observations. First, this market is growing rapidly in terms of the number and variety of participants and transactions. Second, in contrast to the highly localized exchange of services typical in the traditional offline market for contract work, the online market is dominated by long distance North-South (as defined below) trade. Third, the online platforms that facilitate trade in this market introduce seemingly small informational frictions that have significant effects on outcomes. We describe each of these market features in turn.

The growth of online markets for contract labor has been fast and steady. According to Horton (2010), workers in this market earned about \$700 million by 2009, and the Financial Times (2012) estimated this market to be worth \$1 billion annually by the end of 2012. Additional details from oDesk, the largest online marketplace for contract labor in terms of earnings, provide further insight into the growth of this market. The number of employers billing on the site per quarter increased by over 800% between 2009 and 2013 (Figure 1), and the number of working contractors per quarter increased by approximately 1,000% over the same period (Figure 3). In pecuniary terms, the quarterly wage bill on oDesk increased by approximately 900%, from \$10,000,000 to almost \$100,000,000 over the same four-year period (Figure 2).

North-South exchange dominates the pattern of trade in these markets (e.g., relative to North-North, South-South, and South-North). In other words, employers are predominantly from high-income countries¹ and contractors are mainly from lower-income countries. We classify countries as “high-income” using the 2012 World Bank list of high-income countries. We classify the remaining countries as “lower-income.” In Figure 1 we illustrate that not only are there more employers on oDesk from high- compared to lower-income countries, but the number from high-income countries is also growing at a faster rate. Similarly, the wage bill per quarter is significantly greater and growing faster for employers from high- versus lower-income countries (Figure 2). While the contrast is not quite as extreme on the contractor side of the market (a significant number of contractors are from high-income countries), there were approximately three times as many lower versus high-income contractors in 2009 and that difference increased to five times by 2013 (Figure 3). This does not simply reflect a growing volume of small jobs performed by contractors from lower-income countries. In Figure 4 we illustrate that the wage bill reflects a similar pattern in terms of contractors from high- versus lower-income countries.

A number of studies examine how seemingly small information frictions may significantly influence match-

¹We define high-income countries according to the World Bank classification available at <http://data.worldbank.org/income-level/HIC>.

ing outcomes in online markets for contract labor. Perhaps the most dramatic finding is the one reported by Pallais (2012). In this study, Pallais conducts a field experiment where she “treats” 952 randomly selected contractors by hiring them and then providing feedback on their performance. Then she compares the subsequent employment performance of these treated contractors with a set of 2,815 other contractors (“controls”) who applied for her posted jobs but whom she did not hire and therefore did not post information on. She reports that, for those with no prior work experience on oDesk, the subsequent income of treated contractors almost triples relative to the income of control contractors over the following two-month period. She then takes a number of steps to provide further evidence that the observed increase in employment performance is due to the information she posted to the platform about the contractor (i.e., rating and feedback), rather than due to other explanations such as human capital accumulation by the contractor due to the experience of doing the job. The reason this result is so dramatic is because the treatment is so small: the job is only a 10-hour data entry task, the rating is only a single score out of five, and the feedback is only a single sentence: “It was a pleasure working with [x].” In fact, for inexperienced workers, the marginal effect of a more detailed comment that specifies data entry speed, accuracy, following of directions, and timely task completion is not statistically distinct from zero. In other words, the trebling of income is caused by minimal information provided by the employer based on a remarkably small job. Although the observed effect is based on low-wage data-entry specialists who propose wages of \$3 per hour or less, the effect of such a seemingly small amount of information is striking. It points to an important market friction present in this online setting. The author draws a welfare implication from her finding: “Under plausible assumptions, the experiment’s market-level benefits exceeded its cost, suggesting that some experimental workers had been inefficiently unemployed.”

Similar information frictions are reported in other studies in this market setting. Stanton and Thomas (2012) estimate the effect of information from intermediaries on contractor employment. They find that inexperienced contractors affiliated with an intermediary have substantially higher job-finding probabilities (almost double) and wages (15%) at the beginning of their careers on oDesk. Agrawal et al. (2013) examine the relative role of information about experience on oDesk for contractors from high- versus low income countries. They find that information about platform-based work experience disproportionately benefits contractors from low-income countries (approximately 40% premium). In a related study, Mill (2011) finds that once an employer on Freelancer has a good experience with a contractor from a particular country, then the employer is more likely to hire someone else from that country. Also related, Ghani et al. (2012) report that members of the Indian diaspora hiring on oDesk are more likely to hire workers in India than are

other employers. Finally, Horton (2012) finds that recommendations increase the likelihood of a hire in job categories with fewer qualified candidates. In each of these cases, seemingly small amounts of information have significant effects on employment outcomes, suggesting that information frictions play an important role in the matching process online.

With these three market features in mind - rapid growth, North-South trade, and sensitivity to information-based frictions - we turn to analyzing the basic economics of online markets for contract labor in Section 2. In doing so, we consider the characteristics of both the demand and supply sides, stressing the incentives that lead employers as well as contractors to utilize this channel. The main trade-off that we consider is between the reduction in search, communication, monitoring, and transportation costs on the one hand and the potential for new sources of information-related frictions to arise on the other. We then describe the role that online contract labor platforms play in facilitating matches between demand and supply and in addressing some of these trade-offs. Again, we use evidence from oDesk to provide an in-depth illustration.

Drawing on these insights regarding the basic economic properties of online markets for contract labor, we outline a research agenda predicated on three lines of inquiry. These include: 1) distributional effects, 2) market design, and 3) welfare. We describe each in turn.

In Section 3, we ask: How will the digitization of this market influence the distribution of economic activity? We consider distribution along three dimensions. First, we contemplate the distribution of work across geographies. Will digitization shift the distribution of contract work towards lower-income countries? Second, we question the distribution of income within and across countries. Will digitization further accentuate income inequality by amplifying superstar-type distributions whereby only a small fraction of contractors capture a large fraction of rents (although some of these individuals may be in lower-income countries)? Finally, we raise the question of outsourcing. Will digitization lead to a shift in the distribution of work across firm boundaries, constricting the boundary of the firm due to a lowering cost of contracting out discrete jobs? The answers to these lines of inquiry regarding distribution-related effects of digitization will have important implications for understanding the effect of digitization on the overall organization of work and thus implications for social welfare.

Next, in Section 4, we raise this question: How might market design features influence matching in the digital setting? We describe above the impact of ratings and feedback, a market design feature common across most platforms. In the digital setting, platforms can add or change market features at reasonably low cost. However, the ease with which they can be added, deleted, or changed belies the influence they may have on matching outcomes. While contracting platforms employ many interesting market design features,

we focus our discussion on five: 1) performance feedback (e.g., ratings), 2) machine-aided recommendations (for employers and contractors), 3) the allocation of visibility, 4) pricing to reduce congestion, and 5) job category specification. Although platforms in the online contract labor market do not have the match-setting power that is typically analyzed in the market design literature (i.e., directly matching trading partners in settings such as kidney exchanges and medical student-hospital matching), they do influence which matches are ultimately formed and under what terms. Thus, as market design features evolve, so will the types of matches they facilitate.

Finally, in Section 5, we ask: How will the digitization of this market affect social welfare? In particular, we specify two channels through which digitization may generate efficiency gains: 1) better matching, and 2) better production. With regards to matching, the shift from local to global search along with the utilization of market design features enabled by the digitization of relevant information may lead to efficiency gains. With regards to production, a reduction in coordination costs that enables more flexibility in terms of the location and timing (asynchronous) of work as well as a finer division of labor due to the feasibility of contracting out smaller jobs, which enables more specialization, may lead to efficiency gains. At the same time, however, new frictions may lead to new forms of welfare losses.

We conclude by outlining three primary challenges to this research agenda. First, offline data for this sector is costly to obtain but is required to estimate the causal effect of digitization on changes in distributional properties (geography, income, firm boundaries) and welfare effects. Second, the economic salience of particular market design features may be fleeting since the market is evolving quickly and subject to rapid technological change. Finally, data ownership is concentrated among a few platforms that seem interested in engaging with the research community but have interests that are not fully aligned. Despite these challenges, this research agenda identifies opportunities to shed light on questions that are of first-order importance from both a scholarly and economic relevancy perspective.

2 The economics of online contract labor markets

Like other digitized markets, the most salient features of online labor markets are the potential for a large number of transactions and services to be provided by suppliers who may be geographically distant from buyers. What are the implications for the demand and supply of services in this context? Who supplies labor online? What entities search for online services, and what are the trade-offs they face? What institutions contribute to clearing these markets? To address these questions, we begin by discussing how oDesk works.

This will frame the ensuing discussion on labor supply, labor demand, and market-making platforms.

2.1 Work Process on oDesk

To post jobs on oDesk, employers have to register on the site by giving their contact details and information on their company, including name, owner, and location. Once registered, employers are free to post as many jobs as they like. Job postings include a description of the task, the location of the employer, and the type of contract being offered. oDesk supports two contract types—hourly wage and fixed price. Beyond the different payment structures, the contracts have different implications for monitoring and duration specifications. Specifically, when posting an hourly-wage job, employers have to specify the expected number of hours per week and the number of weeks required to complete the job. They stipulate a limit on the number of hours per week a contractor can work. When posting a fixed-price job, employers have to specify the budget and deadline. Employers can make job postings public (so that any contractor can apply to them) or private (so that only contractors they invite can apply to them).

To be hired on oDesk, workers similarly must register on the site by giving their contact details, name, and location as well as by setting up a profile page. Profile pages are meant to advertise contractors to potential employers and can include a description of skills, education, work experience outside of oDesk, oDesk-administered test scores, certifications, whether or not they belong to an agency, and oDesk-specific work histories and feedback scores. Once they have set up their profile pages, contractors can apply to jobs by submitting cover letters and bids to job postings. A bid indicates the amount a contractor is willing to be paid to work on a job.

Employers have the option to interview and negotiate over bids with applicants before hiring and to hire as many contractors as they like. Once hired, contractors complete tasks remotely. Contractors submit their work to employers online and are paid via oDesk. Employers have the option to give contractors bonuses and can also reimburse expenses through oDesk.

After each job, employers give contractors a rating out of five based on six criteria: skills, quality, availability, deadlines, communication, and cooperation. Each contractor also has an overall feedback score, which is a job-size-weighted average of the individual scores. Contractors can provide their employers feedback scores based on the same criteria; employers have a similarly constructed overall score. oDesk provides this service in exchange for 10% of every transaction made on the site.

In addition to oDesk, Elance, Freelancer, and Guru are among the largest online contract labor markets. Elance and Guru were both launched in 1999, followed by oDesk in 2005 and Freelancer in 2009. These sites

are similar in that they allow employers to find and hire short-term workers by registering on the platform and posting jobs to attract applicants. Similarly, they all allow registered contract workers from around the world to apply for jobs posted on the sites by bidding on them and to advertise themselves to employers with profile pages. These platforms earn revenue by charging a percentage of each transaction or member fees to workers and, in some cases, both. In addition to providing a (virtual) place for demand and supply to meet and for the market to clear, these platforms have evolved over time toward addressing some of the key challenges of labor markets in general and online labor markets in particular.

While the other major platforms in the industry share several features with oDesk, they differ on certain dimensions. The primary variations lie in the services they provide participants. For instance, some support contractor employment agencies while others do not, some offer guaranteed payment for hourly wage contracts while others do not, and at least one of the major platforms does not have a virtual office application. Perhaps the most significant difference concerns Freelancer, which supports both traditional hiring and crowdsourcing. Given that crowdsourcing has different implications for matching and production, findings from research based on traditional hiring may not generalize to crowdsourcing settings.

2.2 Labor supply

What are the incentives for individuals to supply labor online? One of the most important benefits to having access to online contract labor markets, especially for individuals participating from lower-income countries who are more constrained in terms of opportunities, is that these marketplaces dramatically increase the pool of available jobs. In addition to increasing the number of opportunities, they also increase the likelihood that contractors will find suitable matches for their skills and preferences.

Contractors also benefit from an increase in flexibility in this market setting. For the most part, these transactions are contract-based: workers are not employees and therefore have more control over their schedules and how they allocate time between the provision of these services and other activities (e.g., another job, family, leisure, etc. (The Economist, 2010)). In a survey of workers on oDesk, more than 80% state that the flexibility and freedom associated with working on the site is a major benefit of online work. Evidence also shows that the flexibility provided by telecommunication contributes to a significant increase in female labor force participation (Dettling, 2012). Thus, these online marketplaces may induce women previously out of the labor market to enter. Especially for contractors in the developing world, who make up the vast majority of workers, easier access to job opportunities from entities in higher-income countries might also imply higher earnings.

Some of the characteristics leading to benefits in participating in these markets may also be sources of costs and risks for contractors. In particular, the contractual nature of these labor relations might lead to more uncertainty about the duration and conditions of a work relationship. The dramatic increase in participation in these markets and the typical profile of participants as relatively highly educated suggest that on balance these markets represent viable and appealing opportunities for a large set of individuals.

2.3 Demand for contract labor

The online market for contract labor offers several benefits to employers relative to traditional offline markets. It lowers the cost of search, communication, and transportation, which benefits trade in various services, such as data entry, translation, and software development. This also enables access to a broader pool of prospective workers with potentially more suitable skills and possibly at more competitive wage rates. Although oDesk has a range of organization types and sizes that use the platform, the access to a large and diverse pool of contract workers provided by these platforms is particularly unique for small, entrepreneurial ventures. For instance, in a survey of employers using oDesk, more than half consider themselves start-ups.

However, the relative lack of face-to-face interactions might make it difficult for employers to extract high-bandwidth information (Autor, 2001). Furthermore, the increased heterogeneity of applicants make comparisons among them more challenging; for instance, comparing seemingly similar school degrees or job experiences of applicants from different countries may be problematic, particularly for novice recruiters. In addition to hidden-quality problems, an obvious issue for prospective employers is the difficulty in monitoring and verifying effort from a distance and through an Internet-mediated transaction.

2.4 Market-making platforms

Consistent with other two-sided markets, intermediaries in online markets take actions to ensure the participation of both suppliers and buyers (Armstrong, 2006; Rysman, 2009). As mentioned above, a key challenge in online contract labor transactions arises from the limited access to high-bandwidth information about both applicants and employers (Autor, 2001). Online contract labor platforms are increasingly providing features that attempt to solve these information problems.² First, platforms provide a verification and standardization device for some of this information; for example, although offline work experiences and educational attainments cannot be easily compared across individuals, especially if they come from very different institutional and cultural contexts, employers can more easily compare work experience accumulated by con-

²Dellarocas (2006) provides a review of reputation systems designed to solve information problems in online markets.

tractors on the platform (i.e., the number of jobs, duration, types, as well as performance as expressed by the rating given by the employers and workers). This information is available in online contract labor markets on contractor profiles, and platforms generally do not allow contractors to delete or block this information from their profiles, thus reducing selectivity issues and increasing the reliability of these signals. Platforms also offer the possibility for applicants to perform standardized tests that offer some easy-to-assess quality measures for prospective employers. Moreover, some platforms support contractor agencies or companies. Contractors in an agency can cooperate to apply for and complete jobs on the site. Some evidence illustrates that agencies help reduce information asymmetries (Stanton and Thomas, 2012).

In addition to providing quality information, online contract labor platforms also help solve challenges relative to the observability and verifiability of effort, on both the worker's and employer's sides, through various mechanisms. Direct monitoring is available on some platforms through virtual office applications³. Contractors who perform their work while logged into these virtual offices are monitored through regular screen shots and activity logs. To provide incentives for contractors to accept this degree of monitoring, some platforms guarantee contractor payment for hourly wage work only if it is performed while logged into the virtual office. Along with direct monitoring, workers' ratings represent a potentially powerful reputational mechanism for aligning their objectives with employer objectives.

Likewise and as in other online markets, moral hazard issues can arise on the part of employers (see, for instance, Resnick and Zeckhauser (2002) and Cabral and Hortacsu (2010) for a discussion of moral hazard in online markets). For example, employers could refuse to pay for work performed outside virtual offices or to reimburse expenses. However, contract workers can rate their experience with an employer on most platforms, thus reducing concerns about the risk of exploitative behavior and renegeing on previous agreements. Furthermore, both employers and contractors can file disputes if they feel they've been unjustly charged or underpaid. Platforms act as mediators in these disputes and ultimately decide how they should be resolved.

³Evidence shows that strict monitoring is important for the success of working from home. Bloom et al. (2013) study a Chinese travel agency that decided to try having some employees work from home. The study finds significant gains from working from home in terms of worker productivity and satisfaction. This may be partially a result of the firm's careful monitoring of telecommuting workers. Dutcher and Jabs Saral (2012) highlight the difficulties that may arise if telecommuting workers are not properly monitored by showing experimental evidence that non-telecommuting workers perceive that their telecommuting counterparts are shirkers.

3 Digitization and the distribution of work

Keeping in mind the incentives and frictions facing employers and contractors that we described above, we turn to contemplating how the digitization of this market may influence the distribution of work. We consider and describe in turn distributional effects along three dimensions: geography, contractor income, and firm boundaries.

3.1 Geographic distribution

The reduction in search, communication, and monitoring costs brought by the digitization of contract labor markets raises the possibility of improving employer-contractor matching and thus enhancing gains from trade. A consequence of this is a potential impact on the geographic distribution of work. Perhaps the most immediate and dramatic gains are those based on cross-region wage variation. Indeed, the dramatic growth in activity on oDesk seems to be primarily of this nature. Specifically, employers in high-income countries hire contractors from low-income countries, even for small jobs that were previously infeasible offline due to transaction costs. As reported in Figure 1, not only were there more than 10 times as many employers from high compared to lower-income countries by late 2012, but the growth rate of employers from high-income countries was much higher than that from lower-income countries. The gap was even greater when expressed in terms of the wage bill rather than the number of employees (Figure 2). Conversely, by 2013, approximately 4.5 times as many contractors were from lower- compared to high-income countries (Figure 3). The trends so far suggest that the spread will continue to increase over time since the number of contractors from lower-income countries is growing at a faster rate. Figure 4 confirms this trend also exists in terms of the total monthly wage bill, not just the number of contractors, despite the fact that, as one might expect, wages are higher for contractors in more developed countries.

Although access to lower-cost labor is one reason for recruiting distant contractors, employers report other reasons as well. In a survey of its users conducted by oDesk, 76% indicated that “remote is less expensive” was a primary reason they were interested in using the platform. However, 46% selected “can get work done faster remotely,” 31% selected “difficult to find local talent,” and 21% selected “no room/equipment.” Thus, in addition to the reduced cost of accessing lower-wage workers, enhanced matching seems to benefit from gains on multiple dimensions.

Countries vary in terms of their level of participation in online contract labor markets. For example, on oDesk approximately 10 times as many contractors are from the Ukraine as from Spain, even though the

two countries are similar in size (populations in 2013: Ukraine 45m; Spain 47m; however, Spain’s economy is approximately 10 times larger: 1.4 trillion USD compared to 0.165 trillion USD for the Ukraine). We illustrate this in Figure 5, where we plot the number of contractors on oDesk per country against population. Nations such as Mexico, Brazil, and China appear to be under-users (participation below what their population would predict), whereas the Philippines, Bangladesh, and India appear to be over-users.

The variation in usage of this digital marketplace may simply reflect offline employment opportunities. oDesk contractors from Bangladesh and the Philippines, for example, earn significantly more than local minimum wages, perhaps partly explaining their disproportionate use of the platform. However, contractors from China also earn significantly above the local minimum wage on average yet under-use the platform relative to other nations. Furthermore, contractors from several countries, like Australia, earn only slightly more than the local minimum wage, on average, and yet seem to be over-users. This variation reflects the relative benefits and costs, including opportunity costs, faced by the labor force in each country. Factors such as proficiency in English (the language used on the site), Internet access, and education levels all affect the returns to engaging with a digitized labor market platform such as this. As these online markets grow, they will provide researchers with useful data to better understand offline employment opportunities (particularly where reliable government data is sparse) and the relative returns to different forms of education in a global work environment. In addition, they will provide a setting for further analysis on the extent to which geographic, language, cultural, and other forms of distance influence flows of trade in labor.

The different composition of online contract workers across countries may also explain the unexpectedly high average wages received by contractors in certain countries, such as China, Poland, and Russia, as reported in Figure 7. Contractors from these three countries in particular are primarily concentrated in software development, information systems, and web development, which offer higher wages on average than most other types of work on oDesk: by 2013, the average wage in software development (\$16) was approximately double that of writing and translation (\$8) and more than triple that of administrative support (\$4) as well as customer support (\$5) and sales and marketing (\$5) (Figure 8). Furthermore, the quarterly spend in software development and web development is significantly greater than in any other category (Figure 9). We plot the concentration of total contractor wage bill by country over time in Figure 10. Russia and Ukraine stand out as especially concentrated in only a few sectors (software development in particular)⁴. In contrast, contractors from the US and the Philippines do work across many categories. This variation in the geographic distribution of work by category likely reflects language, education, and offline

⁴This is consistent with the geographic distribution of work on Kaggle, an online data science competition platform, where software programmers are disproportionately located in Eastern Europe.

work opportunities. That said, Figure 11 indicates that software is one of the least concentrated sectors in terms of the distribution of total wages across countries.

3.2 Income distribution

The digitization of contract labor markets may affect the distribution of income across workers. However, the direction of this effect is ambiguous. On the one hand, digitization could amplify income inequality by way of the superstar effect (Rosen, 1981), whereby the shift to lower search costs enables employers to identify and contract for the best workers (or workers supplying the best value) in a global rather than local context such that the distribution of the total wage bill skews further towards a minority of contractors. On the other hand, digitization could reduce inequality due to more information leading to less mainstream skills in the “long tail” being more efficiently matched (Anderson, 2006).

Researchers report evidence of both types of effects resulting from digitization. For example, Tucker and Zhang (2007) find that when consumers on a wedding vendor website are able to see the popularity of a given vendor, sales concentrate around the more popular vendors. This suggests that online feedback systems have the potential to increase skewness. Elberse and Oberholzer-Gee (2008) find similar support for video sales. In other cases, the reverse is true. Zentner et al. (2012) show that online video rentals are less concentrated around blockbusters than physical rentals, Peltier and Moreau (2012) show that online book sales in France are less concentrated around superstars than offline, and Brynjolfsson et al. (2011) find that Internet sales for women’s clothing are less concentrated than catalog sales. All of these papers identify search cost differences as a core explanation for the results.

Superstar and long tail effects are not necessarily mutually exclusive, and both in fact may be at work in the context of online markets for contract labor. This is because they are influenced by related but distinct characteristics of the services traded in this market. Vertical differentiation (quality) drives the superstar effect, whereas horizontal differentiation (variety) drives the long tail effect (Bar-Isaac et al., 2012). Therefore, subject to demand constraints, they may coexist. The superstar effect will result in increased income inequality as employers tend towards the highest quality (or best value) contractors based on a global rather than local search. Thus, income will shift from contractors supplying the best value locally to those supplying the best value globally. Increased demand will drive up the wages of the highest-quality workers, mainly in cases where the spread is greatest between local and global wages (i.e., low-income countries). The superstar effect may be exacerbated due to information asymmetries and features of the market.

At the same time, horizontally differentiated contractors (e.g., those who specialize in less common areas)

whose offline wages are lower due to limited local demand for their expertise may particularly benefit from digitization since the shift from local to global matching may disproportionately increase the demand for their skills relative to the supply. For example, a software developer in Malaysia who learns to program in a new cutting-edge language (e.g., django) may benefit from digitization since by connecting to the global market that contractor will likely face a greater increase in demand for that skill than an increase in competition for supplying that skill.

In summary, digitization may shift the income distribution in a manner that benefits contractors with skills that are vertically differentiated (i.e., higher quality), horizontally differentiated (i.e., scarce), or lower cost (due to fewer local offline opportunities) at the expense of those with skills that are neither differentiated nor low cost (i.e., mediocre quality, common skills, in high- or middle-income countries). The net effect of such a shift is ambiguous, both at the country level and the individual level. At the country level, although the immediate effect of digitization may be to decrease income inequality as the total wage bill shifts from high- to low-income countries due to expanded search for skills and lower wage rates in low-income countries, the resulting increase in productivity of firms in high-income countries may further increase offline wages there, offsetting the effect of offshoring. At the individual level, while digitization will favor the highly skilled relative to the less-skilled, particularly in high-income countries the services provided by a contractor have increasing marginal costs, unlike products with low marginal costs such as music, books, and software. Therefore, enhanced matching and constrained supply may at least partially offset increased competition and thus temper the extent to which digitization amplifies the skewness of income distribution at the individual level.

Information asymmetries may also affect income distribution. The available evidence shows that even small amounts of (employer- or platform-provided) information have a large effect on future employment prospects (Pallais, 2012; Agrawal et al., 2013). On the one hand, this may increase the skewness of income distribution because contractors who obtain a small lead early on, in terms of online work experience with a positive public employer review, may experience subsequent gains and benefit from increasing returns (at least in the short term). On the other hand, to the extent that online markets facilitate low-cost trials for employers to test working with novice contractors and then publicize their quality, the digitization of this market may decrease skew through the increased public revelation of contractor quality. The fact that a small amount of verified work experience online is associated with a disproportionate increase in winning subsequent jobs for contractors in low-income countries (Agrawal et al., 2013) seems consistent with this latter view.

3.3 Boundaries of the firm

How will the digitization of this marketplace influence the boundary of the firm? Economic theory suggests that because digitization lowers transaction costs (search, communication, and monitoring), the returns to contracting in the market increase relative to performing these services in-house. For example, Grossman and Rossi-Hansberg (2008) model the tension between the benefits (lower cost of labor) and costs (coordination and monitoring) of offshoring to examine precisely the effects of a decline in the cost of offshoring, focusing on the productivity effect of increased offshoring. Similarly, Antras and Helpman (2004) present a model of North-South trade where final-goods firms choose whether to vertically integrate into the production of intermediate goods or to outsource them. Their model offers an explanation for variation in firm boundary decisions (in equilibrium, some firms outsource while others do not, and those that do vary in their outsourcing location choice) based on the variation in firms' productivity levels. Although the authors do not focus on the effect of falling transaction costs associated with outsourcing per se, the influence of this on firm boundaries is a natural implication of their model.

Several studies report empirical evidence that digitization is associated with a contraction in the boundary of the firm. For example, Abramovsky and Griffith (2006) report that more ICT-intensive firms purchase a greater amount of services on the market (rather than vertically integrating) and are more likely to purchase offshore, Brynjolffson et al. (1994) report that investment in IT is correlated with a subsequent decrease in firm size, and Hitt (1999) shows that an increase in IT use is correlated with a decrease in vertical integration.

A recent survey conducted by oDesk of its users sheds further light on the relationship between digitization and firm boundary decisions. Two of the survey questions offer insight on how employers perceive the online platform relative to alternatives for performing contracted work. One of the survey questions asks: "If there had not been an appropriate oDesk contractor available for this project, then what would you most likely have done?" Of the 6,912 respondents, only 15% indicated they would have turned to a local hire, whereas 22% replied they would have worked extra hours, 9% replied they would have delayed or canceled the project, and 50% indicated that they would have used some other remote source. Although there is room for alternative interpretations of these responses (for example, "other remote sources" could refer to other online contract labor platforms such that the results under-represent the fraction who would hire locally in the absence of any online platforms), one possible explanation is that the digitization of this marketplace directly affected the boundary of the firm in only a minority (15%) of the cases. A second oDesk survey question asks: "Thinking about the last time that you hired a contractor through oDesk, what alternatives did you consider?" In this case, respondents were able to select more than one option. Again, only 15% selected "hiring an employee,"

whereas 58% selected “doing it myself.” Shifting from local to distant contractors appears to be a more significant economic effect from the digitization of this market than contraction in the boundary of the firm. Indeed, 40% of respondents indicated that a “local contractor” was an alternative they considered when they last hired a contractor through oDesk.

It is important to note that the majority of oDesk users are small businesses (90% of 7,098 survey respondents indicated that their business had 10 employees or less, with an overall average firm size of 2.6 employees). This raises the question of how the effect of digitizing this marketplace may vary across firm size. For example, do small firms benefit disproportionately from digitization? We cannot draw this conclusion simply from observing a high fraction of small-firm users. First, the 90% small-firm user population may just reflect the distribution of firm sizes in the economy (interestingly, respondents reveal that 68% are part-time businesses, 69% are home-based businesses, and the average firm age is 2.7 years). Second, the survey sample distribution may not reflect the population distribution. Perhaps small firms are more likely to respond to the survey. Still, one might conjecture that small firms are more likely to hire contract workers since large firms are better able to aggregate tasks into full-time jobs and thus avoid the contracting and discontinuity costs associated with task-based hiring.

4 Market design

Platforms in online contract labor markets do not have the match-setting power typical in other contexts that the market design literature has considered (e.g. Roth and Peranson, 1999; Roth, 2002; Milgrom, 2011) because, unlike kidney exchanges and medical student-hospital matching systems, they are not centralized. However, an inability to set matches explicitly does not imply an inability to influence which matches are ultimately formed and under what terms. The position of the platform vis-a-vis the marketplace is more like that of a government that sets policies to encourage efficient market outcomes without dictating trades. The platform decides how often and in what context participants are exposed to each other, what information is collected by parties, and how this information is displayed. Platforms also set policies about what trades are permissible, how entry is gained, what contracts and prices are allowed, and so on. The platform may also make recommendations and set defaults. A few market-design decisions in this softer match-making environment are worth considering to explore how these features affect matching.

First, platforms are in the position to provide standardized and verified information. For example, because oDesk does not permit contractors to delete ratings or comments provided by employers after a job

is completed, this information is possibly distinct from what contractors might include in their resumes and thus valuable to potential future employers. In the introduction above, we describe two studies that report findings indicating that online work history information has a significant influence on subsequent matching outcomes (Pallais, 2012; Agrawal et al., 2013). Furthermore, platforms can provide additional tools for contractors and employers to reveal standardized and verified information about themselves. For example, oDesk provides a series of standardized tests that contractors are able to take so that they may post their scores in order to communicate their proficiency in a specific domain.

Given that wading through too much information is costly for a potential employer, does a simple overall performance score convey an optimal amount of information? Would a more detailed scoring system enhance matching? Pallais (2012) reports that detailed feedback had no effect on subsequent outcomes, relative to simple feedback, for inexperienced contractors. However, for experienced contractors, the extra detail did make a difference. Furthermore, the Pallais result may underestimate the effect of a more detailed rating system since her feedback was conveyed via text rather than, for example, a simple ranking on five dimensions. Given the apparently high sensitivity to ratings and feedback, further research into market design features that address this particular type of information friction seems a fruitful direction for future research.

Second, because contractors have many decisions to make (such as what jobs to apply for, what wage to bid, what skills to learn), as do employers (who to hire, whether to use a fixed or variable fee contract, when to offer a bonus and how much), the digitized nature of these platforms, just like in other online markets, will likely lead to the development of algorithmic assistance with decision-making. These recommendation systems will augment human decision-making by, for example, reducing the search costs of market participants. One potential problem with recommending applicants is crowd-out. Recommending one worker presumably puts another worker at a disadvantage. However, Horton (2012) shows that the quantity and quality of matches can be improved via algorithmic recommendations to employers about candidates to recruit for their openings, without significant crowd-out effects.

Aside from the obvious recommendations about who to trade with and at what terms, the platform can also make other kinds of recommendations. It can, for example, advise parties of best practices in how to manage a working relationship, such as suggesting more communication, periodic raises, and performance evaluations. One interesting challenge of recommender systems is the trade-off between learning and recommending; recommender systems rely on natural decision-making to explore the space of alternatives to train models, but sufficiently good recommender systems that save their users substantial costs are likely to displace natural decision-making. So, maintaining some natural decision-making will eventually be costly,

at least to some users.

Another area where algorithmic recommendations might particularly influence matching is in helping individuals make good decisions about the accumulation of human capital, particularly around which skills to learn. Traditionally, such decisions are made a small number of times by relatively uninformed individuals who receive one-time feedback about their choices. In offline markets, decisions about human capital investments are difficult to observe. Online, these choices are more visible and measurable. On platforms like oDesk, an enormous amount of information illustrates which combination of skills command higher wages in any particular domain. This enables recommender systems to distill which skills are most valuable to learn given a contractor's existing capabilities; the system further learns by observing how contractors perform via experimentation.

Third, how should contract labor sites allocate visibility? Which applicants should be listed at the top versus the bottom on an employer's screen? The large size and value of the search engine optimization (SEO) market provides some indication of the importance of visibility. Should allocation preserve assortativity (e.g., contractors with higher feedback ratings or hours worked are given more visibility)? Should each worker be given at least some visibility? If visibility is auctioned off, what would be the efficiency and distributional properties of such an allocation? While this topic has received much attention from researchers in the private sector at companies like Google and eBay concerning other markets, it remains an open question in the context of the market for contract labor. Yet, this issue is important. Market design decisions concerning the allocation of visibility will surely influence matching outcomes, which in turn will influence both distribution and welfare effects. Moreover visibility relates to congestion, which we discuss next.

Fourth, platforms may need to control congestion due to the fact that posting (and applying for) a job is almost costless. The low cost of applications may lead to an everyone-applies-to-everything equilibrium in which each application also carries virtually no signal value. This was partly the motivation for introducing the AEA signalling mechanism (Coles et al., 2010), in which job-market participants are given two (and only two) signals to send to schools. The school's knowledge of the scarcity of signals makes those signals informative. Accordingly, platforms may consider job application quotas. However, as described above, this strategy might penalize new entrants with low probabilities of being hired (Pallais, 2012). It also ignores employer heterogeneity, with some employers preferring many applicants and others few. Another potentially interesting approach is to allow the employer to decide the cost of applying. These are additional areas for research that reflect the peculiarities of this market.

A fifth interesting market-design feature is the creation of submarkets and categories that are often defined

through some combination of geography and time to coordinate activities and thus create a sufficiently thick market (e.g., the creation of industrial districts for specific sectors). The platform must attempt to define at some level of detail the various services being supplied and then organize the market accordingly. In the language of machine learning, there is both a clustering task (finding the meaningful groups of jobs/contractors based on historical data) and a labeling task (being able to assign a new job to one of the identified clusters based on that job's attributes).

The five market features we describe above represent only a fraction of those that may be important for influencing the matching of employers and contractors as well as the way in which work is managed and produced online. The unique feature of this line of inquiry, relative to the one described above concerning distribution and the one below concerning social welfare, is that this research can be performed without offline data. That is because online features can be compared against each other with respect to the behaviour they elicit from users. So-called "A/B testing," which refers to controlled randomized experiments that allow for identifying causal relationships between variations in market design features and subsequent user behaviour, has already become a standard industry practise for determining the relative performance of competing market design features. That is likely the reason that the majority of research concerning online markets for contract labor relate to this line of inquiry whereas there is very little so far on the other topics.

5 Social Welfare

Two immediate consequences of digitization in this market may have important welfare implications. First, digitization may lead to better matching because the pool of prospective workers and employers increases dramatically due to the decline in costs associated with distance. Second, digitization may lead to efficiency gains from production due to lower coordination costs. We discuss both lines of inquiry below.

5.1 Matching made easier?

The ease of access to online contract labor markets, due to the development of platforms such as oDesk, Freelancer, Elance, and Guru, has the potential to considerably increase the pool of both job seekers and employers and to reduce search costs. Matching models, particularly as applied to labor markets, predict that this will lead to efficiency gains due to lower search costs and a lower likelihood of mismatches (Petrongolo and Pissarides, 2001; Wheeler, 2001).

However, opposite forces are also at play. While information technologies reduce the role of distance for

search and execution of work, they also lead to a more heterogeneous pool of both workers and employers. In addition, the absence of personal interactions typical of offline and more localized labor markets precludes access to soft or high-bandwidth information about both job seekers and prospective employers (Autor, 2001). This introduces uncertainty that in turn may lead to an overall reduction in the quality of workers (Akerlof, 1970) and/or to search frictions (Stigler, 1962). These search frictions could be exacerbated if quality is difficult to determine (Wilde, 1981), which is quite possible because of the diverse labor pool.

Although theories of search and matching specific to online labor markets have not yet been developed, a growing body of evidence, described above, points to the presence of these informational problems and the ways in which they are addressed in online contract labor platforms, (Horton, 2012; Pallais, 2012; Stanton and Thomas, 2012). A common pattern to a number of these studies is to look at the presence of preferences for certain geographic locations of workers as a way to alleviate uncertainty about workers' quality (Mill, 2011; Ghani et al., 2012; Agrawal et al., 2013). An implication here is that online contract labor platforms contribute to the alleviation of informational asymmetries by providing verifiable, standardized information (such as previous experience on the same platform) for all workers, regardless of their origin.

The broadening of the pool of workers and employers and, at least potentially, the increased likelihood of good matches, is also likely to have implications for wages and income distribution. The fact that in online contract labor markets the number of workers outweighs the number of employers in every job category suggests that while many workers may be left unemployed, employers have a relatively good chance of finding a worker who meets their criteria, with wages driven down (Petrongolo and Pissarides, 2006). However, because worker backgrounds may vary more than in traditional labor markets, a relatively small number of workers may meet the job requirements. As a result, wage offers could be higher than expected. This suggests that in job categories with many qualified workers, the wages will be lower than in those with few qualified workers relative to the number of job postings. As the market evolves, wage differences between job types should begin to disappear.

5.2 Efficiency gains from production?

Digitization may lead to efficiency gains in production due to lower coordination costs that enhance, for example, contractor flexibility, discretization of work into smaller jobs enabling more specialization, and remote team work. For instance, Dettling (2012) reports that flexibility provided by IT contributes to an increase in the female labor force participation. More broadly, digitization may enable efficiency gains in production through lowering the cost of outsourcing.

Of course, outsourcing and offshoring predates the development of online contract labor markets. Of particular relevance here are theories of service outsourcing and offshoring (e.g. Bhagwati et al., 2004; Francois and Hoekman, 2010). Combined, these theories predict that the gains to service outsourcing are potentially significant. However, they focus on relatively long arm’s-length contracts between relatively large firms rather than on the short contracts between small organizations and individuals, typical of online markets.

Outsourcing services to online contract labor markets is also likely to lead to geographically dispersed production, even within narrowly defined tasks. For example, work teams may be composed of individuals who are not necessarily co-located. Lazear (1999) argues that cultural diversity in work teams is costly and should only occur when skill complementarities exist between teammates to offset these costs. It may be harder to meet these conditions in very diverse online labor markets than it is in more traditional labor markets. Two recent studies based on online labor markets focus on task completion and the effects of team organization, communication structure, incentives, and motivation on performance. Lyons (2013) provides field experimental evidence on how nationally diverse communication impacts online team production and finds that nationally homogeneous teams benefit from working together but that diverse teams perform better when members work independently of one another. Related to the topic of online labor market partnerships, Horton (2011) uses survey data from the crowdsourcing site Mechanical Turk to show that workers believe employers on the site are more fair and honest than offline employers.

6 Conclusion

We identify three broad lines of inquiry as central to the digitization research agenda. All three focus on the effect of digitizing the market for contract labor. The first concerns welfare effects, the second distribution effects, and the third user behavior effects. All three are set in the context of the market for contract labor but have broader implications for digitization in other settings.

Access to data will pose a challenge to fully addressing these questions. In contrast to data from online platforms that collect information on hiring (as well as pre- and post-hiring) transactions at a granular level and at low cost, it is costly to obtain even a basic level of offline contracting data. Yet, to fully address the first and third lines of inquiry outlined above, offline data is required to estimate the causal effect of digitization on changes in distributional properties (geography, income, firm boundaries) and welfare. This is likely why most of the first wave of studies concerning the digitization of this market focuses on market design-related subjects (e.g., experience, agencies, ratings) since these questions only require observing within-platform

variation in user behavior and do not require linking these data to non-platform-participants.

While the second line of inquiry concerning market design, information frictions, and user behavior is largely spared from the requirement to link with offline data, the greatest challenge to this research in the short and medium term will likely be the rapid evolution of the industry. As illustrated above, the industry is growing rapidly. In addition, complementary technologies, such as those associated with mobile and social, are changing rapidly. As such, market design features that seem salient today may be less relevant relative to other features in the future. For example, monitoring technologies such as work rooms with screen shots were only recently introduced and are already standard practice across many platforms. Furthermore, they are likely to be replaced soon with better technology such as streaming screen video. While the ultimate goal of research of this type is obtaining a deeper understanding of human behavior rather than of a particular market design feature, the economic salience of the feature is often important for generalizability and yet may be fleeting due to the rapid pace of change in this setting. Still, insight into user response to informational frictions is an important contribution.

Whereas the distribution and welfare lines of inquiry are most likely to be led by scholars and policy makers, the market design-related research will almost surely include important contributions from industry since this issue is of first-order importance for product development and competition. This has already been the case with oDesk (Horton, 2010, 2012) as well as with other market design issues on platforms such as Google (Varian, 2007; Choi and Varian, 2012), eBay (Blake et al., 2013), and Yahoo (Ghosh and McAfee, 2011; Lewis and Reiley, 2011). Industry interest coupled with their access to high-quality data may significantly accelerate progress on this research frontier. At the same time, the competitive implications of market design insights may inhibit the dissemination of this type of research, and thus the overall impact of industry interest in this subject on the rate and direction of progress on this part of the agenda is ambiguous.

Given the role that platforms play as the central collectors of data in these markets, they will influence the direction of research on all three lines of inquiry through their decisions regarding providing researchers with access to their data. Early signs are promising for the research community since many of the most prominent platforms have established Chief Economists or similar types of research-friendly leadership positions and encourage employees to participate in the academic community by publishing their research and participating at conferences and other scholarly events.

Given the rapid growth rate of the online market for contract labor, this research agenda is economically important. The welfare-related line of inquiry will help us better understand the potential private and social benefits due to the digitization of this sector of the economy. The distribution-related line of inquiry

will shed light on how the benefits of digitization may be allocated across countries and individuals as well as its impact on the structure of the firm. Finally, the market design-related line of inquiry will provide further insight into the importance of particular information frictions and human behavior in the digital world as we explore user reactions to platform features, many of which are common across sectors outside of contract employment. Overall, these insights will be of great interest to scholars, policymakers, and industry participants alike.

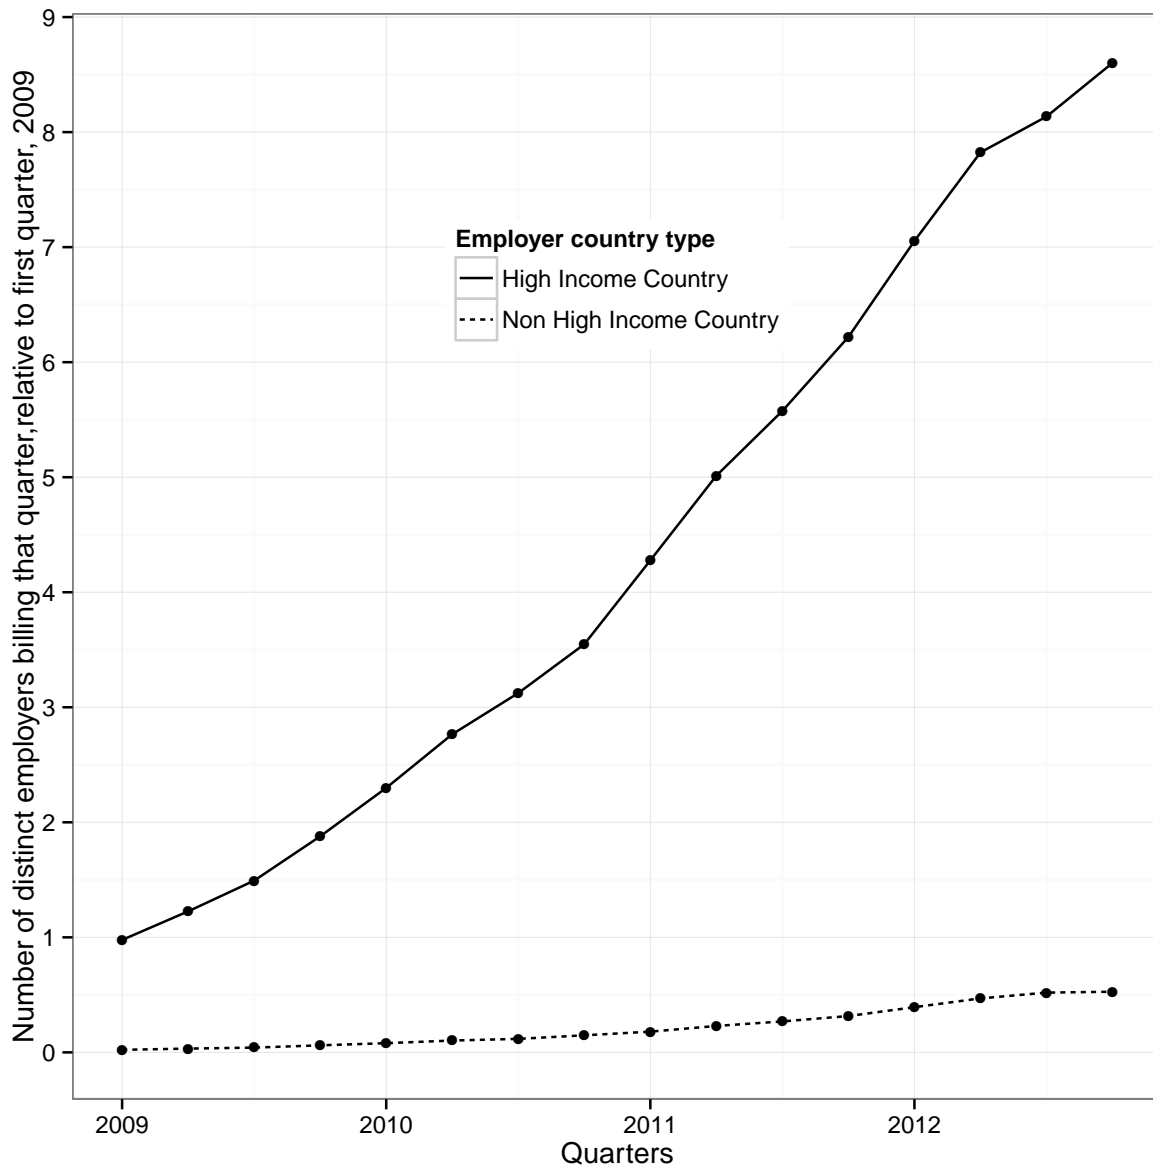
References

- Abramovsky, Laura and Rachel Griffith**, “Outsourcing and Offshoring of Business Services: How Important is ICT?,” *Journal of the European Economic Association*, 2006, 4 (2-3), 594–601.
- Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons**, “Does Information Help or Hinder Job Applicants from Less Developed Countries in Online Markets?,” 2013.
- Akerlof, George A.**, “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 1970, 84 (3), 488–500.
- Anderson, Chris**, “The Long Tail: How Endless Choice Is Creating Unlimited Demand,” 2006.
- Antras, Pol and Elhanan Helpman**, “Global Sourcing,” *Journal of Political Economy*, 2004, 112 (3), 552–580.
- Armstrong, Mark**, “Competition in two-sided markets,” *The RAND Journal of Economics*, 2006, 37 (3), 668–691.
- Autor, David H.**, “Wiring the Labor Market,” *Journal of Economic Perspectives*, 2001, 15 (1), 25–40.
- Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat**, “Information Gathering Externalities for a Multi-Attribute Good,” *The Journal of Industrial Economics*, 2012, 60 (1), 162–185.
- Bhagwati, Jagdish, Arvind Panagariya, and T. N. Srinivasan**, “The Muddles over Outsourcing,” *Journal of Economic Perspectives*, 2004, 18 (4), 93–114.
- Blake, Thomas, Chris Nosko, and Steven Tadelis**, “Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment,” 2013.
- Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying**, “Does Working from Home Work? Evidence from a Chinese Experiment,” 2013.
- Brynjolfsson, Eric, Thomas W. Malone, Vijay Gurbaxani, and Ajit Kambi**, “Does Information Technology Lead to Smaller Firms?,” *Management Science*, 1994, 40 (12), 1628–1644.
- , **Yu Hu, and Duncan Simester**, “Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales,” *Management Science*, 2011, 57 (8), 1373–1386.
- Cabral, Luis and Ali Hortacsu**, “The dynamics of seller reputation: Evidence from ebay,” *The Journal of Industrial Economics*, 2010, 58 (1), 54–78.
- Choi, Hyunyoung and Hal Varian**, “Predicting the Present with Google Trends,” *Economic Record*, 2012, 88 (s1), 2–9.
- Coles, Peter, John Cawley, Phillip B. Levine, Muriel Niederle, Alvin E. Roth, and John J. Siegfried**, “The Job Market for New Economists: A Market Design Perspective,” *The Journal of Economic Perspectives*, 2010, 24 (4), 187–206.
- Dellarocas, Chrysanthos**, “Reputation mechanisms,” *Handbook on Economics and Information Systems*, 2006, pp. 629–660.
- Dettling, Lisa J.**, “Opting Back In: Home Internet Use and Female Labor Supply,” 2012.
- Dutcher, E. Glenn and Krista Jabs Saral**, “Does Team Telecommuting Affect Productivity? An Experiment,” 2012.
- Elberse, Anita and Felix Oberholzer-Gee**, “Superstars and Underdogs: An Examination of the Long-Tail Phenomenon in Video Sales,” 2008.
- Financial Times**, “Virtual Working Takes off in EMs,” 2012.
- Francois, Joseph and Bernard Hoekman**, “Services Trade and Policy,” *Journal of Economic Literature*, 2010, 48 (3), 642–692.
- Ghani, Ejaz, William R. Kerr, and Christopher T. Stanton**, “Diasporas and Outsourcing: Evidence from oDesk and India,” 2012.
- Ghosh, Arpita and Preston McAfee**, “Incentivizing High-Quality User-Generated Content,” *Proceedings of the 20th International Conference on World Wide Web*, 2011, pp. 137–146.
- Grossman, Gene M. and Esteban Rossi-Hansberg**, “Trading Tasks: A Simple Theory of Offshoring,”

- American Economic Review*, 2008, 98 (5), 1978–1997.
- Hitt, Lorin M.**, “Information Technology and Firm Boundaries: Evidence from Panel Data,” *Information Systems Research*, 1999, 10 (2), 134–149.
- Horton, John J.**, “Online Labor Markets,” *Internet and Network Economics*, 2010, 6484, 515–522.
- , “The condition of the Turking class: Are online employers fair and honest?,” *Economic Letters*, 2011, 111 (1), 10–12.
- , “Computer-Mediated Matchmaking: Facilitating Employer Search and Screening,” 2012.
- Lazear, Edward P.**, “Globalisation and the Market for Team-Mates,” *The Economic Journal*, 1999, 109 (454), 15–40.
- Lewis, Randall A. and David Reiley**, “Does Retail Advertising Work? Measuring the Effects of Advertising on Sales via a Controlled Experiment on Yahoo!,” 2011.
- Lyons, Elizabeth**, “Team Production in International Labor Markets: Experimental Evidence from the Field,” 2013.
- Milgrom, Paul**, “Critical Issues in the Practice of Market Design,” *Economic Inquiry*, 2011, 49 (2), 311–320.
- Mill, Roy**, “Hiring and Learning in Online Global Labor Markets,” 2011.
- Pallais, Amanda**, “Inefficient Hiring in Entry-Level Labor Markets,” 2012.
- Peltier, Stephanie and Francois Moreau**, “Internet and the “Long Tail versus superstar effect” debate: Evidence from the French book market,” *Applied Economics Letters*, 2012, 19 (8), 711–715.
- Petrongolo, Barbara and Christopher A. Pissarides**, “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 2001, 39 (2), 390–431.
- and – , “Scale Effects in Markets with Search,” *The Economic Journal*, 2006, 116 (508), 21–44.
- Resnick, Paul and Richard Zeckhauser**, “Trust among strangers in Internet transactions: Empirical analysis of eBay’s reputation system,” *Advances in Applied Microeconomics*, 2002, 11, 127–157.
- Rosen, Sherwin**, “The economics of superstars,” *American Economic Review*, 1981, 71 (5), 845–858.
- Roth, Alvin E.**, “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics,” *Econometrica*, 2002, 70 (4), 1341–1378.
- and **Elliott Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 1999, 89 (4), 748–780.
- Rysman, Marc**, “The economics of two-sided markets,” *The Journal of Economic Perspectives*, 2009, 23 (3), 125–143.
- Stanton, Christopher T. and Catherine Thomas**, “Landing the First Job: The Value of Intermediaries in Online Hiring,” 2012.
- Stigler, George J.**, “Information in the Labor Market,” *Journal of Political Economy*, 1962, 70 (5), 94–105.
- The Economist**, “Work in the Digital Age: A Clouded Future,” 2010.
- Tucker, Catherine and Juanjuan Zhang**, “Long Tail or Steep Tail? A Field Investigation into How Online Popularity Information Affects the Distribution of Customer Choices,” 2007.
- Varian, Hal R.**, “Position Auctions,” *International Journal of Industrial Organization*, 2007, 25 (6), 1163–1178.
- Wheeler, Christopher H.**, “Search, Sorting, and Urban Agglomeration,” *Journal of Labor Economics*, 2001, 19 (4), 879–899.
- Wilde, Louis L.**, “Information Costs, Duration of Search, and Turnover: Theory and Applications,” *Journal of Political Economy*, 1981, 89 (6), 1122–1141.
- Zentner, Alejandro, Michael D. Smith, and Cuneyd Kaya**, “How Video Rental Patterns Change as Consumers Move Online,” 2012.

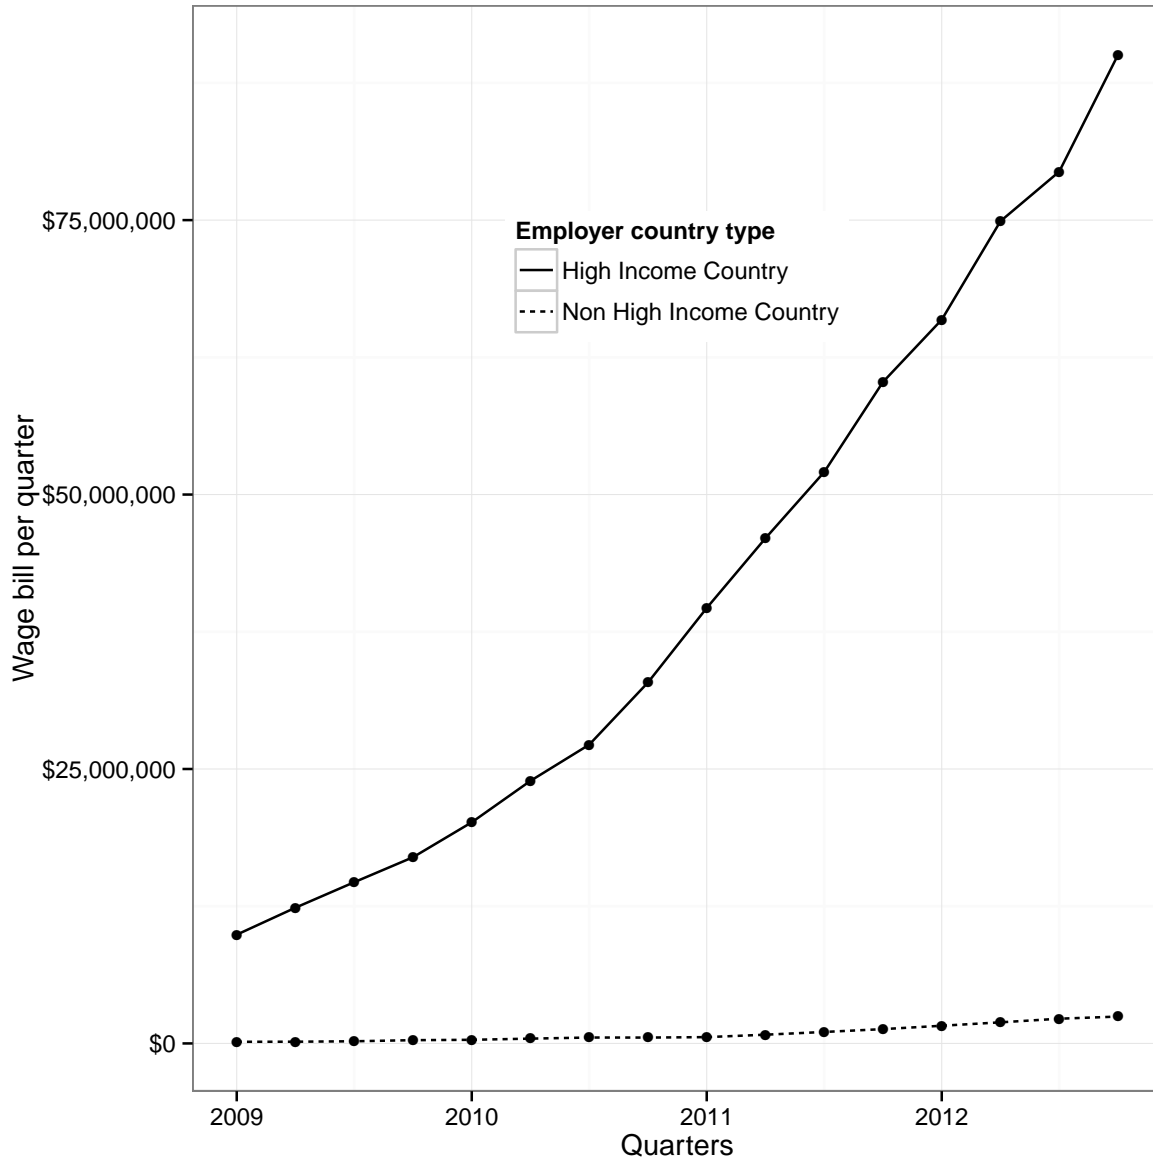
7 Figures

Figure 1: Number of billing employers per quarter on oDesk, relative to total number of employers in first quarter of 2009, by employer country income level



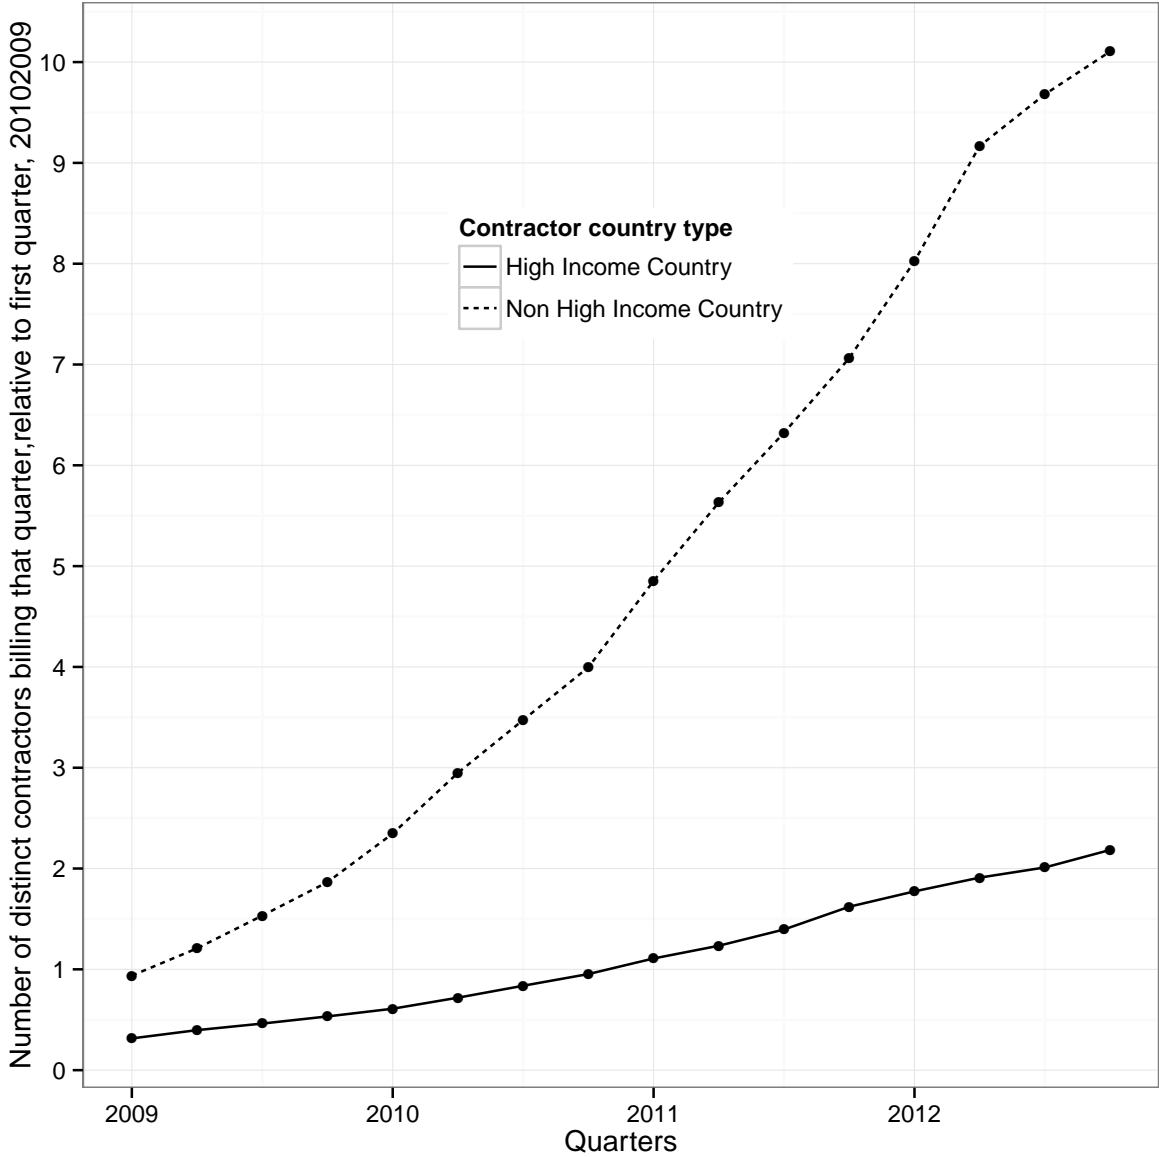
Notes: This figure uses data collected from oDesk to show the relative number of billing employers per quarter, by country income status. We use the 2012 World Bank list of high-income countries for our country classification. The base quarter is the first quarter (i.e., January, February, and March) for 2009. Although the count looks like it is exactly 1 for HIC in 2009, it is slightly below—there were a small but non-zero number of employers from lower-income countries during that quarter.

Figure 2: Quarterly wage bill on oDesk by employer country income level




Notes: This figure uses data collected from oDesk to show the quarterly wage bill by the employer's country's income status. We use the 2012 World Bank list of high-income countries for our country classification.

Figure 3: Number of working contractors per quarter on oDesk, relative to total number of contractors in first quarter of 2009, by contractor country income level




Notes: This figure uses data collected from oDesk to show the relative number of working contractors per quarter, by country income status. We use the 2012 World Bank list of high-income countries for our country classification. The base quarter is the first quarter (i.e., January, February, and March) for 2009.

Figure 4: Contractor quarterly earnings on oDesk by contractor country income level




Notes: This figure uses data collected from oDesk to show the quarterly wage bill, by the contractor's country's income status. We use the 2012 World Bank list of high-income countries for our country classification.

Figure 5: Number of contractors per country on oDesk versus country population, on a log-log scale



Notes: This figure uses data collected from oDesk to show the count of contractors who have ever worked on oDesk by country, versus the 2012 World Bank estimate of that country's population. Both axes are log-log scale. We only include countries with 500 or more ever-active contractors.


Figure 6: Contractor mean hourly wage on oDesk by country, relative to that country's estimated local minimum wage



Source: Wikipedia, May 22, 2013
en.wikipedia.org/wiki/List_of_minimum_wages_by_country


Notes: This figure uses data collected from oDesk to compare the mean hourly wage (log scale). To estimate hourly wages, we restrict our attention to hourly contracts in the first half of 2013. Harmonized minimum wage data is difficult to acquire. As a proxy, we use the Wikipedia estimates, as of May 2013.

Figure 7: Contractor mean hourly wages on oDesk, by country



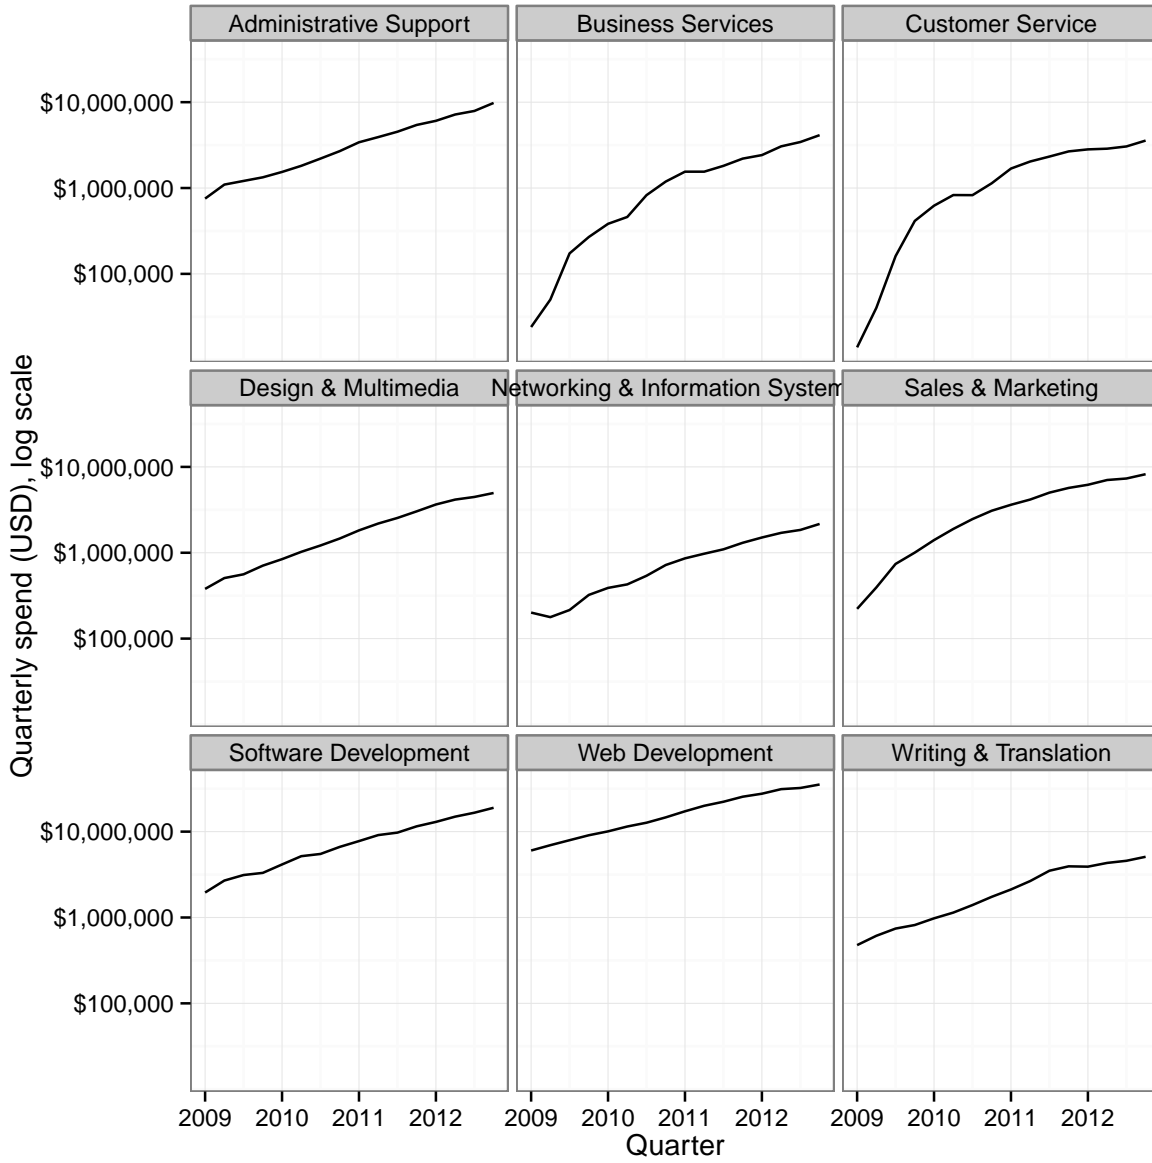
Notes: We estimate hourly wages using a sample of all hourly contracts in the first half of 2013. We exclude observations of less than 10 cents and more than \$100, as these observations are more likely to not be true hourly wages but rather individuals using the time tracking software provided by oDesk or approximating a fixed price contract of some kind with a high hourly wage. For each wage estimate, we include a 95% confidence interval. Note that for high population countries like India and Philippines, these confidence intervals are so narrow that they appear to be point estimates.

Figure 8: Average hourly wage on oDesk per quarter, by job category



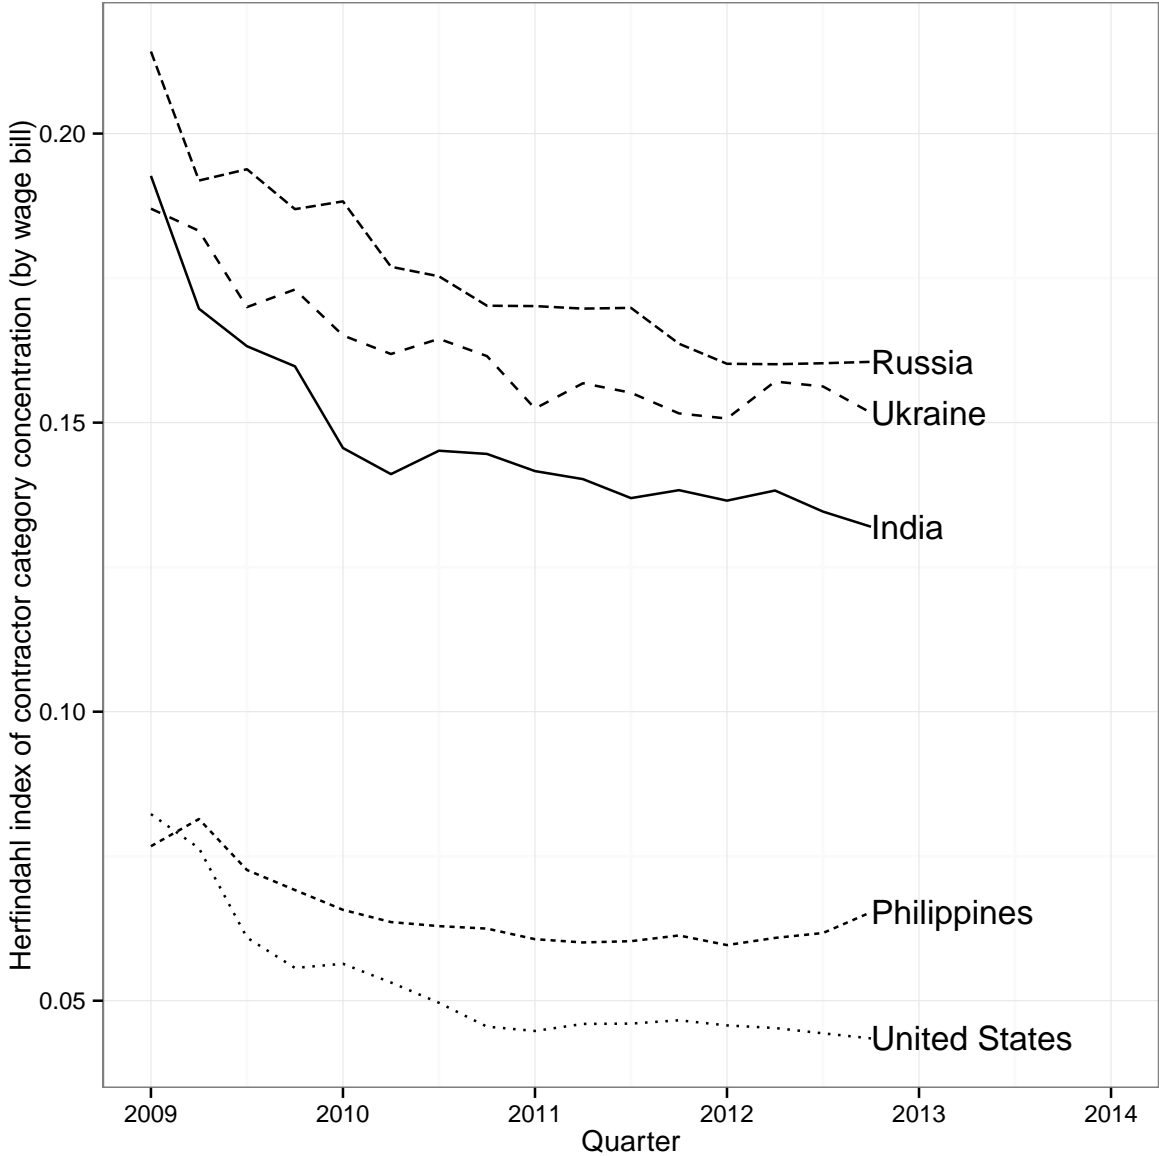
Notes: This figure uses oDesk data to show the mean hourly wages per quarter in each of the main oDesk categories of work.

Figure 9: Quarterly wage bill per job category on oDesk (log scale)




Notes: This figure uses oDesk data to show the total quarterly wage bill by job category.

Figure 10: Contractor job category concentration on oDesk by contractor country, over time



Notes: This figure uses oDesk data to compute a quarterly Herfindahl for a select number of oDesk contractor countries. We compute the index by treating oDesk job categories as “firms” and contractor countries as “industries.” To compute this measure, for each quarter, we estimate the share of dollars earned by contractors from a particular country, in each category. We then report the sum of the square of these shares. The higher the index, the more concentrated workers from that country. For example, if an index is near 1, it would mean that nearly all workers from that country work in a single category.

Figure 11: Job category concentration on oDesk by contractor, over time



Notes: This figure uses oDesk data to compute a quarterly Herfindahl for each job category, treating each country as a “firm” and each category as an “industry.” To compute this measure, for each quarter, we estimate the share of dollars within a category earned by contractors from each country. We then report the sum of the square these shares. The higher the index, the more that particular category is dominated by workers from a particular country. For example, if an index is near 1, it would mean that nearly all work in that category is completed by workers from a single country.