Online Labour Index: Measuring the Online Gig Economy for Policy and Research

Otto Kässi and Vili Lehdonvirta

Oxford Internet Institute, Oxford Internet Institute

3 November 2016

Online Labour Index: Measuring the Online Gig Economy for Policy and Research

Otto Kässi E-mail: otto.kassi@oii.ox.ac.uk

Oxford Internet Institute, 1 St Giles, OX13JS Oxford, United Kingdom.

Vili Lehdonvirta

Oxford Internet Institute, 1 St Giles, OX13JS Oxford, United Kingdom.

Abstract. Labour markets are thought to be in the midst of a dramatic transformation, where standard employment is increasingly supplemented or substituted by temporary gig work mediated by online platforms. Yet the scale and scope of these changes is hard to assess, because conventional labour market statistics and economic indicators are ill-suited to measuring online gig work. We present the Online Labour Index (OLI), a new economic indicator that provides the online gig economy equivalent of conventional labour market statistics. It measures the utilization of online labour across countries and occupations by tracking the number of projects and tasks posted on platforms in near-real time. We describe how the OLI is constructed and demonstrate how it can be used to address previously unanswered questions about the online gig economy; in particular, we show that the online gig economy grew at an annualized rate of 14 percent. To benefit policymakers, labour market researchers, and the general public, the index is available as an open data set and interactive online visualization, which are automatically updated daily.

Keywords: online labour, online gig work, measurement of vacancies, web data collection, occupation prediction

1. Introduction

Labour markets are thought to be in the midst of a dramatic transformation, where standard employment is increasingly supplemented or substituted by temporary gig work mediated by online platforms. Instead of hiring a standard employee or contracting with a conventional outsourcing firm, companies are using online labour platforms to find, hire, supervise, and pay workers on a project, piece-rate, or hourly basis. Enterprises from small to large are using these platforms to access skills and flexible labour, assisted by specialized consultants and online outsourcing firms. Dozens of platforms have emerged to cater to different types of clients, workers, and projects, ranging from deskilled microtasks to complex technical projects and professional services. Tens of millions of workers are thought to have sought employment through such platforms (Kuek et al., 2015).

The potential policy implications of this emerging 'online gig economy', 'platform economy', or 'on-demand economy' are deep and wide-ranging, but not yet fully understood. It may create significant new earning opportunities in countries and occupations suffering from unemployment, but also erode labour protections and contribute to economic insecurity. It may alleviate local labour shortages, but also generate demand for new skills and training. It may contribute to the temporal flexibility of work, but also to the unpredictability of working life, and further undermine social policies based on binary notions of employment and unemployment, breadwinners and dependants. Yet the real scale and scope of these implications is hard to assess, because conventional labour market statistics and economic indicators are ill-suited to measuring work that is transacted via online platforms. The entire digital transformation of labour markets remains largely unobservable to policy makers and labour market researchers.

The purpose of this paper is to introduce the *Online Labour Index (OLI)*, a new economic indicator that provides an online labour market equivalent of conventional labour market statistics. The Online Labour Index is an index that measures the utilization of online labour platforms over time and across countries and occupations. It provides a solid evidence base for future policy and research. The OLI is published online as an automatically updating open data set and interactive visualization at `http://ilabour.oii.ox.ac.uk/online-labour-index/`. In this paper we describe how the OLI is constructed and illustrate how it can be used to address crucial policy issues that existing data sources are unable to address. We also briefly discuss the current limitations and planned extensions of the index.

The remainder of this paper is structured as follows. Section two discusses related literature. Section three discusses the methdology: sample selection, data collection, occupation classification and sampling of employer countries. Section four presents our results, and section five concludes.

2. Background

Both policy makers and researchers (Sundararajan, 2016; Parker et al., 2016; Evans and Schmalensee, 2016) are paying an increasing amount of attention to the online gig economy. A recent EU Commission flagship strategy paper notes that "online platforms are playing an ever more central role in social and economic life" (European Commission, 2015). American policy makers have likewise held several workshops and hearings on the topic. Courts on both sides of the Atlantic have been asked to rule on gig workers' employment rights.

At the same time, it is widely recognized both in the research literature as well as among policy makers that existing economic statistics are not well suited to measuring the online gig economy, in terms of both capturing its full extent as well as distinguishing its impact from other activities. There are several reasons for this. In general, existing economic statistics are prone to mismeasuring the value of digital activities and investments. This is because these activities are often not directly related to production, but to development, design, and marketing, whose values are harder to establish (Corrado and Hulten, 2015; Brynjolfsson and McAfee, 2014; Coyle, 2015, 2016). Existing labour market statistics in particular are missing online work because of definitional and measurement issues. A standard ILO definition of employment used by statistical agencies counts as employed anyone gainfully employed for at least one hour either in a week or a day (Hussmanns, 2007). This measure fails to capture any incremental effects of online work – if someone already has a job and does a second job online, their efforts are not captured in employment statistics. It is also not clear to what extent online workers choose to report their earnings to tax agencies, especially if the earnings are small. This might be an especially relevant concern for the many online gig workers thought to be living in developing countries, where the informal economy often dominates and tax underreporting is common (Kuek et al., 2015). Even when online earnings are duly reported, the existing statistical categories do not allow such earnings to be distinguished from contingent income earned from the traditional labour market.

Previous studies have used a variety of methods to attempt to address the paucity of statistics on the online gig economy. Lehdonvirta and Ernkvist (2011), Kuek et al. (2015), and Groen and Maselli (2016) used a combination of expert interviews and data disclosed by online labour platforms to estimate total market sizes and future growth rates. Kuek and colleagues estimated that the global annual gross market size, including workers' earnings and fees charged by platforms, was approximately $2 Bn in 2013, reaching $4.8 Bn in 2016. They also estimated that there were a total of approximately 48 million registered workers on the platforms, of whom 10 percent were active. Estimates based on expert interviews and platform disclosures are useful, but their sources and methods are often opaque, and they are difficult to repeat regularly in a way that would produce comparable statistics over time. For business reasons, online labour platforms tend to disclose statistics selectively at best; detailed and repeated disclosures could be used to derive market shares, earnings, and growth rates, which early-stage companies often prefer to keep confidential and publicly listed companies may be legally held back from publishing.

Studies can also potentially use data from other intermediaries. Farrel and Gregg (2016) used proprietary data on JPMorgan Chase's American customers' bank account transactions to estimate participation in the platform economy, defined as including both labour platforms and capital platforms such as Airbnb. They found that roughly 1 percent of adults in the sample had earned income from the platform economy in each month, and that this figure had grown over 10-fold from 2013 to 2015. These are useful statistics and the methodology is reliable and repeatable, though only by those with access to the bank's data. The methodology misses transactions outside the traditional banking system, paid with media such as PayPal or Amazon vouchers; these are likely to be non-trivial in volume.

Many traditional labour market statistics are produced by surveying workers and establishments on a regular basis. A recent survey of UK adults by Huws and Joyce (2016) found that as many as 11 percent had successfully earned income through gig work platforms, while three percent said they were doing so at least weekly. These are significantly higher figures than those reported by Farrel and Gregg (2016), which could be explained by the inclusion of non-traditional payment channels or by US-UK national differences, but also by other methodological differences and differences in concepts and definitions. A notable methodological weakness in Huws and Joyce's study is that the respondents were recruited via a commercial online panel whose members participate in surveys against compensation; such respondents seem likely to be more engaged in online work than the general population.

Official labour market statisticians have also started efforts to address the online gig economy. The U.S. Department of Labor has announced that it plans to restart the Contingent Worker Supplement of the Current Population Survey in 2017. It was previously collected in 2005. It will address many of the limitations of studies such as Farrel and Gregg (2016) and Huws and Joyce (2016). However, a significant limitation of these studies that the Contingent Worker Supplement will not address is that the resulting statistics are national in scope. That is, it only concerns work completed within the United States. This is probelmatic since the online gig economy is highly transnational, with 89 percent of transactions crossing national borders on one large platform (Lehdonvirta et al., 2014). Many of the policy issues hinge on understanding the global dynamics of the economy, which is difficult with statistics drawn from a patchwork of national initiatives and methodologies. A further limitation of survey-based approaches and especially telephone and postal surveys is that they are relatively costly, and as a result likely to be repeated only infrequently (BLS, 2015). The resulting statistics are thus poor at measuring the potentially rapid changes in the online gig economy, which are relevant to many policy questions.

In summary, previous studies have used a variety of methods to examine the total size of the online gig economy, its growth, and the incidence of its use in national populations. Their findings suggest that the absolute size of the market remains small by national economy standards, but that it is growing rapidly and involves measurable fractions of national populations. The findings suggest that the online gig economy may already be having non-trivial impacts on labour markets and societies, but are not detailed enough to reveal where the impacts are being felt the most. Important questions are left unanswered or answered only with unreliable one-off statistics. Which countries and occupations are being affected? In which countries and occupations is the use of online labour platforms – and thus its impacts – growing? How stable or volatile is online employment in different occupations? New statistics are needed if these really quite elementary questions about the online gig economy are to be addressed in a satisfactory manner.

If the digital economy presents new challenges for statistics production, it also presents new opportunities. Many digital platforms provide application programming interfaces (APIs) intended to be used by third-party software developers to integrate the platform with other applications. Such APIs can frequently be used to access and automatically collect data on the platform's contents. If an API is unavailable or unsuitable for data collection, it is frequently possible to collect relevant data by 'scraping' or automatically accessing and downloading the platform's web user interface. There are examples of such data collection approaches being used to create labour market indices. For instance, the Conference Board 'Help Wanted OnLine Index' tracks vacancies posted on Internet job boards in the United States on a monthly basis ((The Conference board, 2016)). It measures the number of new vacancies and vacancies reposted from the previous month for over 16,000 Internet job boards and corporate boards, broken down by state/city and occupation. It provides excellent statistics, but covers conventional employment only, not gig online work. Similarly, the MTurk Tracker project (Difallah et al., 2015; Ipeirotis, 2010) tracks new and completed tasks on Amazon Mechanical Turk, an online labour platform. It produces interesting statistics in almost real time, but is limited to a single platform that is not a very good representative of the online gig economy more generally. A general online labour index – comparable in scope and function to national labour market indices – is currently missing from digital economy research and policy.

3. Methodology

To address the gap in existing statistics, we have developed the Online Labour Index, an index that measures the utilization of online labour platforms over time and across countries and occupations. Online labour platforms are here understood as websites and apps through which buyers and sellers of labour and services transact fully digitally. That is, we require that the worker and employer are matched digitally, the payment is conducted digitally via the platform, and that the result of the work is delivered digitally. This definition includes platforms for online freelancing, microwork, and similar activities, but excludes platforms for local gigs, such as Uber and Deliveroo. Though both the online gig economy and the local gig economy can be understood as components of an overall gig economy, they are likely to involve rather different dynamics, and measuring them are also rather different endeavours methodologically. The OLI is designed to measure what is understood as the online gig economy.

3.1. Sample selection

The OLI is based on tracking all projects and tasks posted to selected online labour platforms, using API access and web scraping. In the current version of the index, we define the sample as the five largest English-language online labour platforms, as indicated by the unique visitor estimates provided by Alexa.com. Alexa is the only publicly available source of traffic meas-

Table 1. Traffic of the platforms currently included in the index.

	Alexa rank	Monthly unique visitors (estimate)
Freelancer.com	1,308	75,755,378
Guru.com	7,742	12,617,987
Mturk.com	5,144	19,052,971
Peopleperhour.com	6,563	14,904,412
Upwork.com	488	204,657,137

Table 2. Summary of types of data collected from each of the platforms.

	Occupation classification	Employer countries observed
Freelancer.com	Observed	Unobserved
Guru.com	Predicted	Observed
Upwork.com	Observed	Unobserved
Mturk.com	Observed	Unobserved
Peopleperhour.com	Predicted	Unobserved
Upwork.com	Observed	Observed

urements for all major websites around the world, based on a voluntary plugin that observes browsing behaviour.† This sample is listed in Table 1. Using our informal census of online labour platforms combined with Alexa's figures, we estimate that these five platforms account for at least 60 percent of all traffic to English-language online labour platforms. These five platforms also represent a range of different market mechanisms and contracting styles, from online piecework to hourly freelancing. An index based on this sample is therefore likely to provide a reasonable proxy for the dynamics of the whole market.

3.2. Data collection

The data from which the OLI is calculated is collected by periodically crawling the list of vacancies available on each of the sample platforms. As in conventional labour markets, a vacancy refers to a job, project, or task offered by a firm that wishes to hire a worker. For each crawl, we save the list of open vacancies. Comparing the lists allows us to calculate the number of new vacancies between two crawls. A new vacancy for day t is defined to be a vacancy which has not been observed for any period $0, ..., t-1$, and is observed on period t.

The main shortcoming of this approach is that we do not observe vacancies which were either posted and completed between two crawls, or which were completed without a vacancy being posted. The latter might happen if a vacancy is filled without it being posted on a platform. These hidden vacancies also exist and remain unmeasured in traditional vacancy statistics, so the problem is not unique to our work. Instead of reporting the absolute number of new vacancies, we report an index number normalised so that the mean observed daily vacancies in May 2016 equals 100 index points. If the share of unobserved to observed vacancies remains constant, the changes in the index accurately measure the changes in new vacancies.

Besides vacancy status, we also seek to observe the occupation classification and employer country for each vacancy. The platforms differ in what pieces of information they make available through API access and scraping, with the consequence that these dimensions of the index are based partly on prediction and on generalizing from a subset of the sample. The data collected from each platform is summarized in Table 2 and discussed in more detail in the following sections.

3.3. Classifying online gig work by occupation

Following the practice of occupation classification statistics, we classify similar vacancies into occupations, allowing the index to be used to track different types of work. Occupations are

†See `http://aws.amazon.com/alexa-top-sites/` (accessed 2016-07-19).

"set[s] of jobs whose main tasks and duties are characterised by high degree of similarity" (see ILO (2012) pp. 59-60).

There are several practical challenges in performing occupational classification. First, occupational boundaries are ambiguous, so that it can be impossible to unambiguously map a vacancy to a single occupation. Second, the processes, tasks, and skill requirements within jobs are constantly changing, so that any attempt to reduce this complexity to a fixed classification will encounter problems. Finally, there are considerable differences in the contents of similar occupations across countries, industries, and establishments‡. For all of these reasons, occupational classification is a difficult process, subject to criticisms concerning its reliability (Elias, 1997). Nonetheless, it is clear that some operational classification of occupations is required to facilitate comparisons of different types of work in the OLI across time and countries. 3.

After experimenting with a variety of approaches, we settled on a very practical approach to the occupation classification problem: we adopt the top-level occupation classes being used in Upwork.com in May 2016, modifying it so that certain smaller classes are combined. This results in a high-level classification system consisting of six classes. By all accounts these classes capture the main contours of online gig work relatively well. Similar classifications are used in previous literature (see Kokkodis and Ipeirotis 2015; Kokkodis et al. 2015), and qualitative research based on online gig workers' interviews also finds broadly similar categories of work (Wood et al., 2016). For the other platforms in the sample, we manually map their occupation taxonomies to the six-class system.

The resulting classification is summarized in Table 3. The Clerical and data entry class consists of so-called microwork or human intelligence tasks ('HITs'), which include tasks like data entry and image classification. These tasks typically require only basic computer literacy and numeracy. In contrast, the vacancies in the Professional services class typically require formal education and knowledge of local formal institutions such as accounting regulations. The Sales and marketing support class consists largely of support tasks related to online advertising. They are distinguished from the two aforementioned classes because they form a large and distinct portion of online freelancing. The Writing, Software development and technology, and Creative and multimedia classes are mostly self-explanatory. Like any classification system, this system will not allow all possible vacancies to be classified unambiguously. For example, a website design project that comprises both graphic design and programming could be classified as either Creative or Software development. However, in our case this problem may be smaller than in standard classification systems. In typical labour force surveys, the classification of vacancies into occupations is performed retrospectively after a vacancy is posted on a job board, but in our case it is in the interest of the employer posting the vacancy to classify it in a correct fashion to reach the best-matching pool of applicants.

Due to limitations in the data we work with we do not attempt to separate professional occupations (those requiring a university level education) from associate professional (those requiring a non-university degree) and clerical occupations (those requiring at most a secondary level degree). This is perhaps a weakness of our approach in comparison to to standard occupation classifications such as the International Standard Classification of Occupations (ILO, 2012). Nonetheless, as we argue above, these standard occupational classifications are not necessarily free of their own problems. Further, it seems that the formal educational qualifications play a relatively small role in online gig labour markets compared to a conventional labour market; casually browsing through vacancies shows that none make any explicit mention of required educational qualifications.

3.4. Predicting the unobserved occupation classes

Not all of the sample platforms expose their occupation taxonomies, with the consequence that the occupation class cannot be directly observed for approximately 15 percent of the vacancies.

‡For instance, using detailed survey data, Autor and Handel (2013) find that there are considerable differenes in tasks completed by workers even if they are classified into the same detailed occupation classes.

Table 3. Classification of occupation types on platforms

Occupation class	Examples of projects
Professional services	Accounting
	Consulting
	Financial planning
	Legal services
	Human resources
	Project management
Clerical and data entry	Customer service
	Data entry
	Transcription
	Tech support
	Web research
	Virtual assistant
Creative and multimedia	Animation
	Architechture
	Audio
	Logo design
	Photography
	Presentations
	Video production
	Voide acting
Sales and marketing support	Ad posting
	Lead generation
	Search engine optimization
	Telemarketing
Software development and technology	Data science
	Game development
	Mobile development
	QA and testing
	Server maintenance
	Software development
	Web development
	Web scraping
Writing and translation	Academic writing
	Article writing
	Copywriting
	Creative writing
	Technical writing
	Translation

We follow Amato and colleagues' (2015) machine learning approach to predict the occupations of these vacancies. We took a random sample of 1172 vacancies from the set of vacancies with an unobserved occupation class and manually classified them into the six occupations. We used these manually classified vacancies as the training data for the classifier. To produce the features on which the classifier acts, we processed the vacancies' titles and descriptions by removing stopwords, special characters, and numbers. The remaining words were stemmed. As a result, we ended up with a 1172×2951 matrix where rows represent vacancies, columns represent stemmed words, and cells record word counts, which are the predictive features in our model.

As is evident, the model has more potential explanatory features (2951) than we have observations (1172). Further, a relatively small subset of the features tend to be the best predictors for each occupation class; for example, all projects mentioning a programming language by name are in the Software development and technology category. To reduce the dimension of our space of predictive features, and to increase the predictive power of the model, we applied the LASSO method (Hastie et al. (2008), pp. 68-70). The LASSO method tends to perform well in problems where the set of potential predictive features is large but only a moderate number of features have predictive power (Tibshirani, 1996). We used a multinomial logistic LASSO implemented in R's glmnet package (Friedman & Hastie, 2010a).

We evaluated the predictions by randomly splitting our training data into two, fitting a model with half of the data, predicting occupation classes for the other half, and comparing the predictions with the manually assigned classes. The confusion matrix along with additional details is presented in Appendix A. Table 3 presents various accuracy metrics calculated from the holdout sample. Since the occupation categories in the training data are unbalanced, our preferred accuracy metric is the balanced accuracy, which accounts for the unbalanced occupation proportions (García et al., 2009). Both the balanced accuracy and the aggregate precision measures demonstrate that the regularised multinomial regression performs well in our data. It reaches a balanced accuracy of over 75 percent in five of the six occupation classes, and a total accuracy of 71 percent. This is a significant improvement over a random classification, which would result in a total accuracy of roughly 32 percent.

Another way to study the impact of the classification on the outcome is to look at the reliability of the classification – that is, how large a share of the vacancies is subject to misclassification

Table 4. Classification precision metrics. [95% confidence interval].

	Prevalence	Precision	Recall	Balanced accuracy
Clerical and data entry	7%	69%	97%	83%
Professional services	5%	39%	96%	69%
Creative and multimedia	15%	77%	90%	84%
Sales and marketing support	7%	70%	96%	83%
Software development and technology	53%	71%	91%	81%
Writing and translation	13%	86%	93%	90%
Aggregate precision of classifier	71% [67%, 75%]			

due to prediction uncertainty. The classification based on the platform taxonomies has no random error, so the approximately 85 percent of the vacancies that are classified based on the platform taxonomies have an agreement rate of 100 percent. The remaining 15 percent are classified using the LASSO classifier, which is estimated to classify 71 percent of the vacancies correctly. We end up with an overall agreement rate of 96 percent $((.85+(.71\times.15))\approx.96)$).

3.5. Estimating the employer country distribution

Besides occupations, we also wish to use the OLI to track countries where online gig work is being used. As indicated in Table 2, not all of the platforms allow us to observe vacancies' posters' countries. We therefore estimate the overall employer country distribution using data from those platforms that do allow the country information to be observed. This is done using the assumption that cross-platform differences in employer country distributions are driven only by cross-platform differences in occupation distributions. In reality there are no doubt many other factors as well that cause employer country distributions to vary across platforms, but in practice we can only control this one. Fortunately the largest platform allows the employer country to be observed directly, and this among other things leads us to believe that our estimate will reasonably track the employer country distribution of the market.

The employer country distribution is estimated as follows. We take a random sample of all new vacancies posted to the two platforms that allow the country information to be observed. A random sample rather than a census is used to reduce the number of requests made to the platforms' servers. These samples are then weighted to match the occupation distribution of all the platforms. This process involves four steps:

(a) We collect a random sample of vacancies from the two platforms. We calculate the share of occupation m on platform p as o_{mp};
(b) we calculate the observed occupation shares across our full sample as O_m;
(c) thereafter, we we multiply each sampled observation by $\frac{O_m}{o_{mp}}$; and
(d) finally we weight the samples from the two platforms according to the relative size of the platforms.

As a result, the occupation distribution of the subsample for which we can observe the employer's country matches the occupation distribution of the full sample. Further, by weighting the two random samples by the relative sizes of the platforms we take into account the differences in the sizes of our data sources. The weights are dynamically recalculated each time the data is updated.

4. Results

In this section, we demonstrate some of the insights that can be derived from the Online Labour Index. Since the main purpose of this paper is to introduce the index as a tool that can subsequently be applied in a variety of further research and policy making, the results are fairly

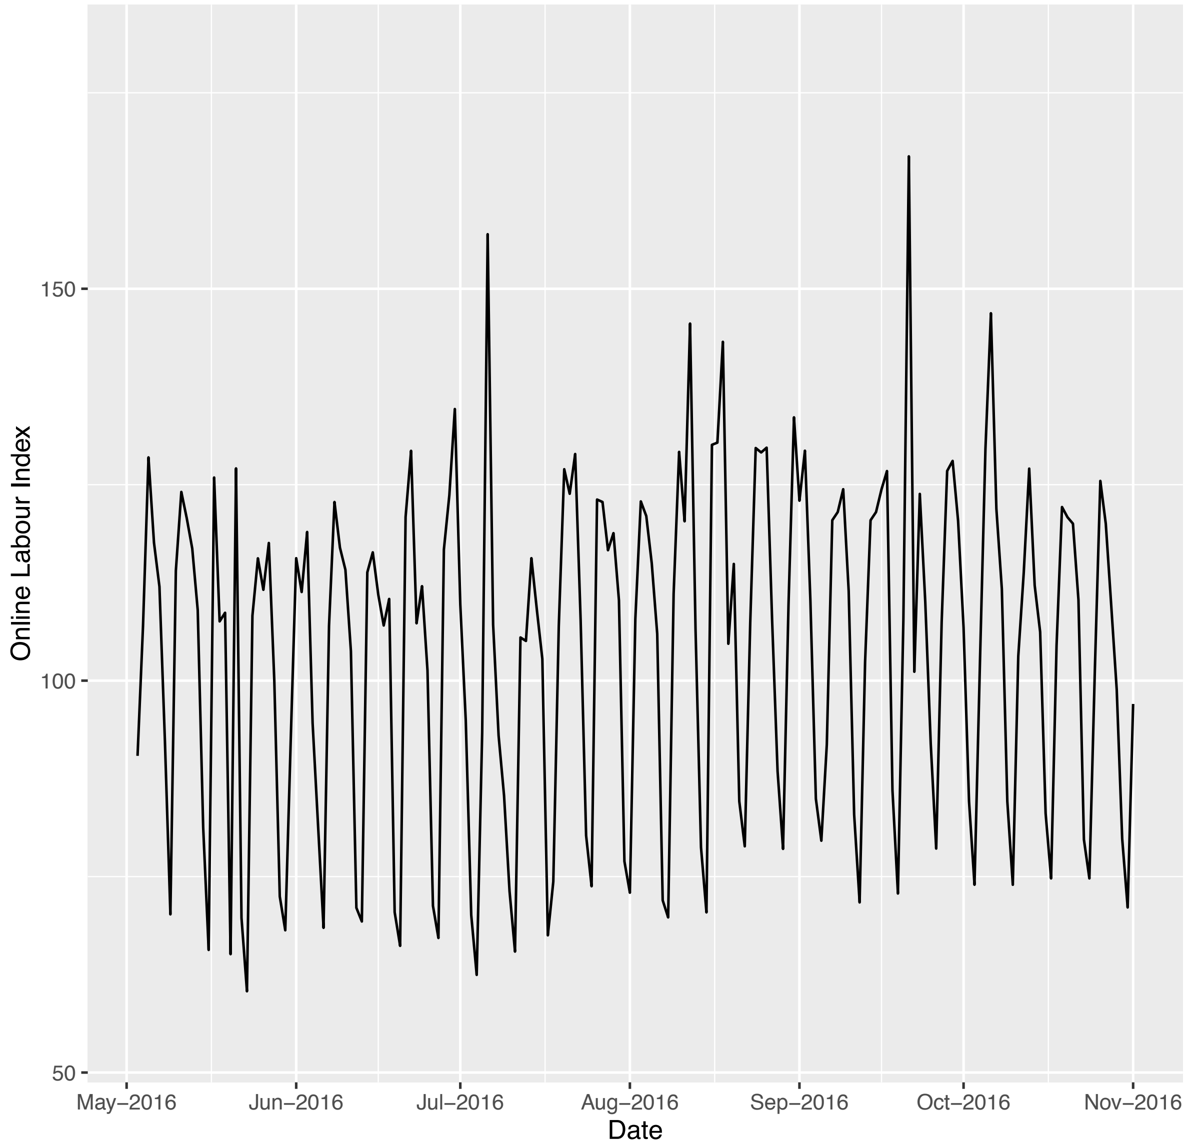

Figure 1. Online Labour Index time series, normalised to 2016/05 = 100

descriptive in nature and no attempt is made in this paper to set them in the full context of existing debates. Nonetheless, when appropriate, we discuss potential theoretical explanations for our findings.

4.1. Growth of online labour

There is a lot of discussion about the emerging online gig economy in the media, but is it actually growing? This is an elementary question for any policymaker or researcher approaching the phenomenon, yet existing data sources only provide incomplete snapshots based on incommensurate methodologies.

Figures 1 and 2 plot the OLI from May to early October 2016, the period for which data is available at the time of writing. Figure 1 shows that the time series exhibits significant weekly variation: a noticeable dip takes place each weekend. This is consistent with the observations inDifallah et al. (2015), who note that there is strong weekly periodicity of arrival of new tasks on Mechanical Turk. To allow us to see the overall trend in the market without the weekend effect, Figure 2 plots a 28-day moving average of the index. We can see from Panel (b) that the

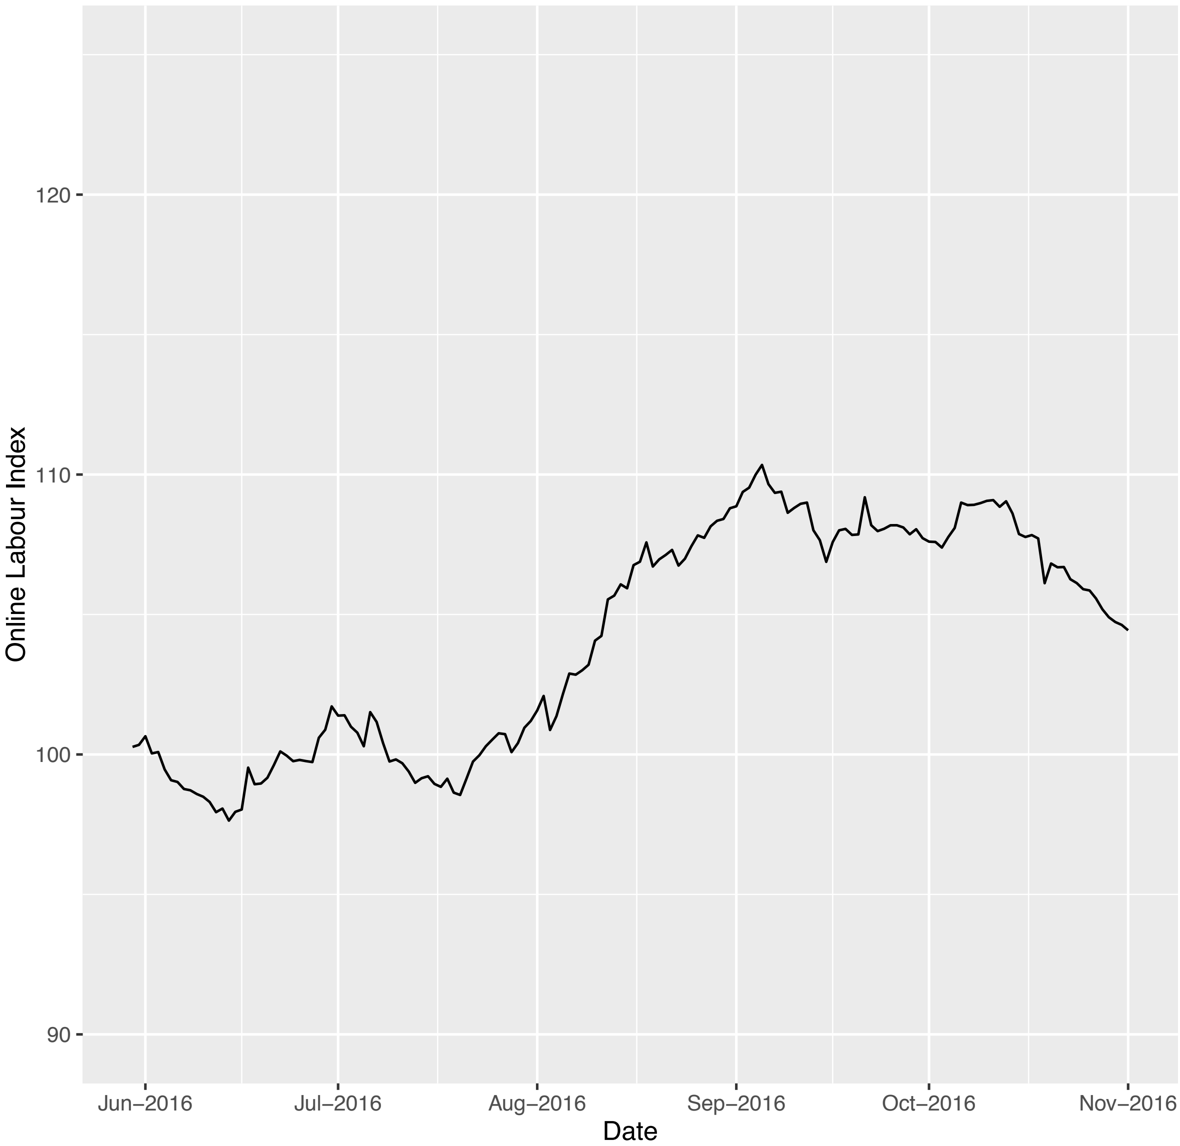

Figure 2. Online Labour Index, normalised to 2016/05 = 100. (28-day moving average)

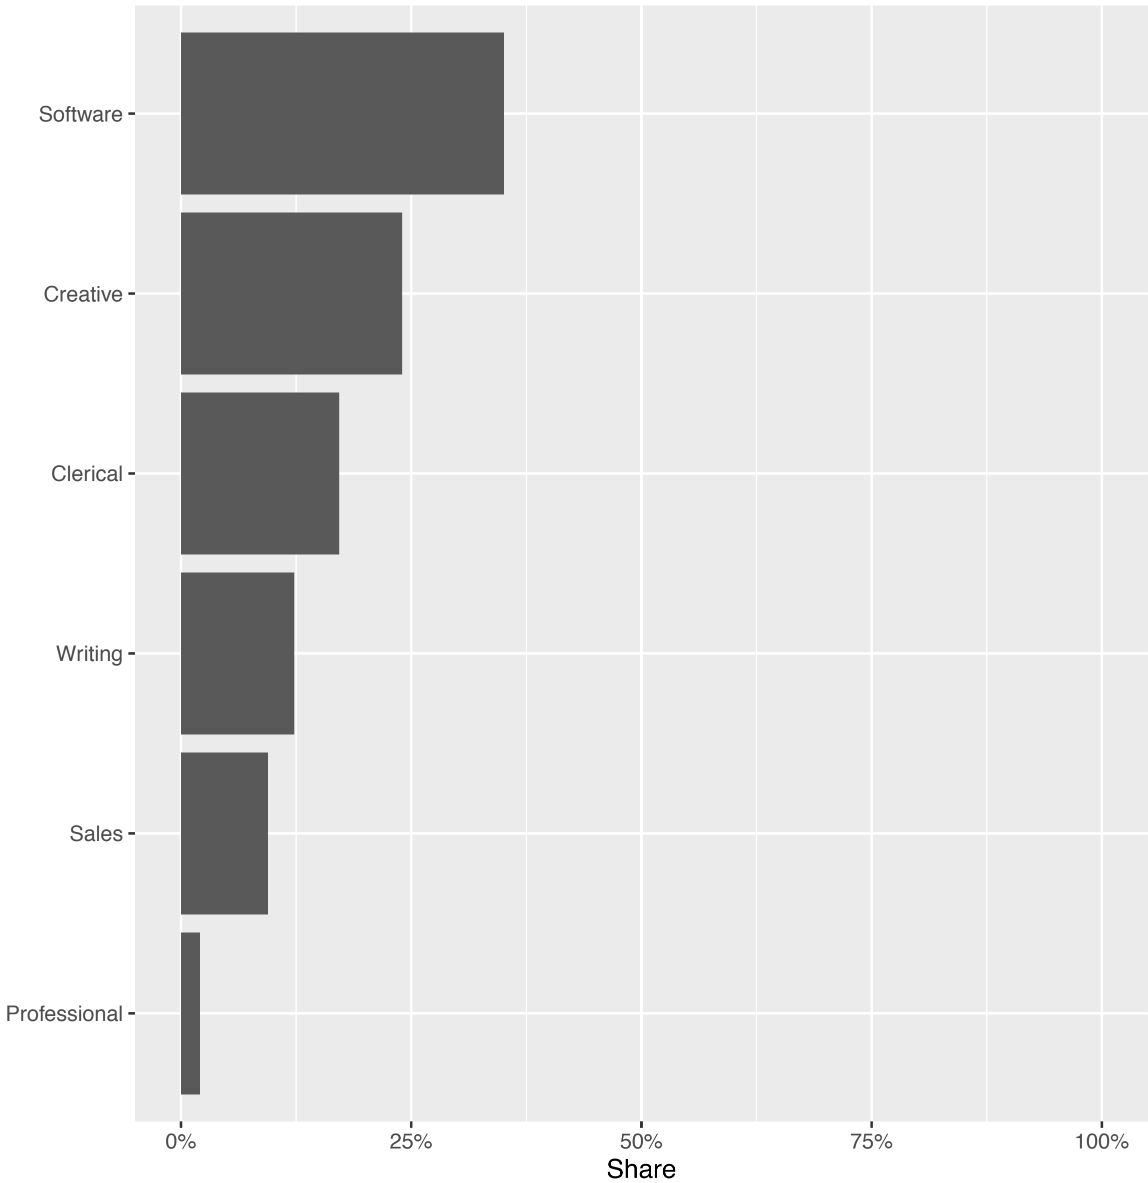

Figure 3. Distribution of occupations in online labour

index grew by about 9 index points (9%) from May to early October 2016.

The 9 percent overall growth over the observation period corresponds with an approximately 18 percent annual growth rate. In labour market terms this represents extremely rapid growth. Most of this growth took place during the month of August. Over time, as the OLI continues to accumulate data, it will reveal whether this growth spurt represents a lasting trend in the online gig economy, or simply annual variation or some other short-term dynamic. The reader is invited to consult the real-time online visualization of the index at the URL indicated in the introduction.

4.2. Leading occupations in online gig work

Previous studies provide estimates of the total market size, but only limited views of where exactly this market is emerging. Which occupations is it affecting? Figure 2 reveals that the highest demand is for software development and technology skills, with roughly one third of the vacancies belonging to that category. Software development and technology are followed by creative and multimedia work, followed by clerical and data entry work.

The relative prominence of software development and technology vacancies in the online labour market can perhaps be explained by the relatively long history of the outsourcing and offshoring of IT services, and the standardised processes associated with it. The use of online labour for repetitive clerical tasks such as data entry similarly follows on the footsteps of conventional business process outsourcing (BPO) practices, except that the work is being sent directly to individual online workers rather than to BPO firms with conventional offices and employees. Conversely, the relatively small amount of professional services being contracted on platforms (2 percent of the total market) could be explained by the fact that these types of services often require a high level of trust and tacit communication that may not be as easily achieved via online communications. They may also require familiarity with the client's local institutional environment, which distant online service providers may not possess. Still, the fact that professional services such as legal services are now regularly bought via online platforms at all is quite remarkable, given that the established professions have not always been at the forefront of technology adoption (Susskind and Susskind, 2015).

4.3. Geography of demand for online gig work

Previous studies such as Kuek et al. (2015) and Lehdonvirta et al. (2014) provide glimpses of how workers on specific platforms are situated around the world, and but there is even less information on how employers are situated. The distribution of employers by country and occupation as revealed by the OLI is presented in Figure 3. Since the employer country distribution is highly skewed in such a way that the top five countries add up to over 90 percent of the vacancies, we group the smaller countries into geographical groups for visualization purposes.

The figure shows that roughly 52 percent of all the vacancies are posted by employers from the United States. Other top employer countries are United Kingdom (6.3%), India (5.9%), Australia (5.7%), and Canada (5%). Non-UK European employers together account for approximately 10 percent of the market, so Europe's total market share is about 16 percent. The largest European country after the UK is Germany. It might seem surprising that a developing country such as India would be so prominent on the hiring side. One potential explanation for this is that workers who win projects sometimes hire other online workers to do the work in their stead, acting as project managers or simply salespersons (Lehdonvirta et al., 2015). But India also has a large IT sector of its own, which is likely to be generating domestic demand for online workers.

Finally, Figure 4 plots occupation distributions within the countries. A striking feature of the geography of online gig work utilization is that the occupational demand profiles of the leading employer countries are rather similar. Employers from all the leading buyer countries post most vacancies in the software development and technology category, followed in most cases by creative and multimedia, and so on. This is surprising, because the sectoral and industry structures of these countries are very different, as are the occupational profiles of their conventional domestic labour markets. The fact that their online labour demand profiles nevertheless resemble each other suggests that the demand largely comes from the same industry within each country: information technology, broadly defined. If and when other industries and sectors start making use of online labour in greater quantities, the OLI should begin to show employer countries' occupational demand profiles diverging. Once again, we invite the reader to consult the real-time web version of the index.

5. Discussion and conclusions

In this paper, we introduced the Online Labour Index, a new economic indicator that provides an online gig economy equivalent of conventional labour market statistics. We described how the OLI is constructed and illustrated how it can be used to address crucial questions about the online gig economy that existing data sources are unable to address. The most fundamental finding from the OLI is that the online gig economy is growing rapidly, at an annualized rate of approximately 14 percent. This is a striking figure when it is contrasted with growth rates in conventional

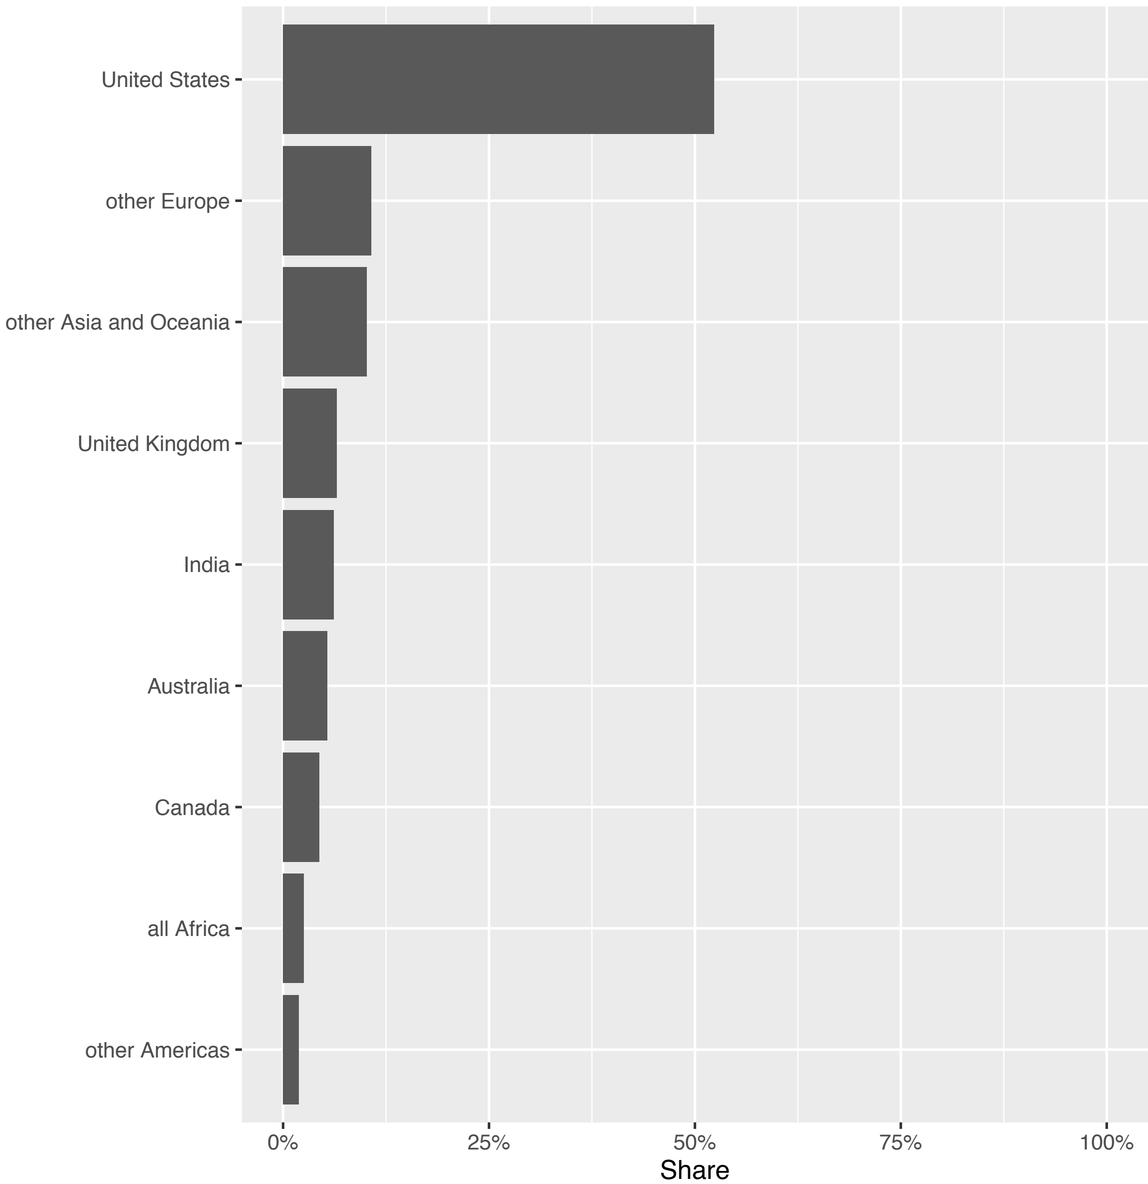

Figure 4. Online Labour Index share by employer country, May-October 2016

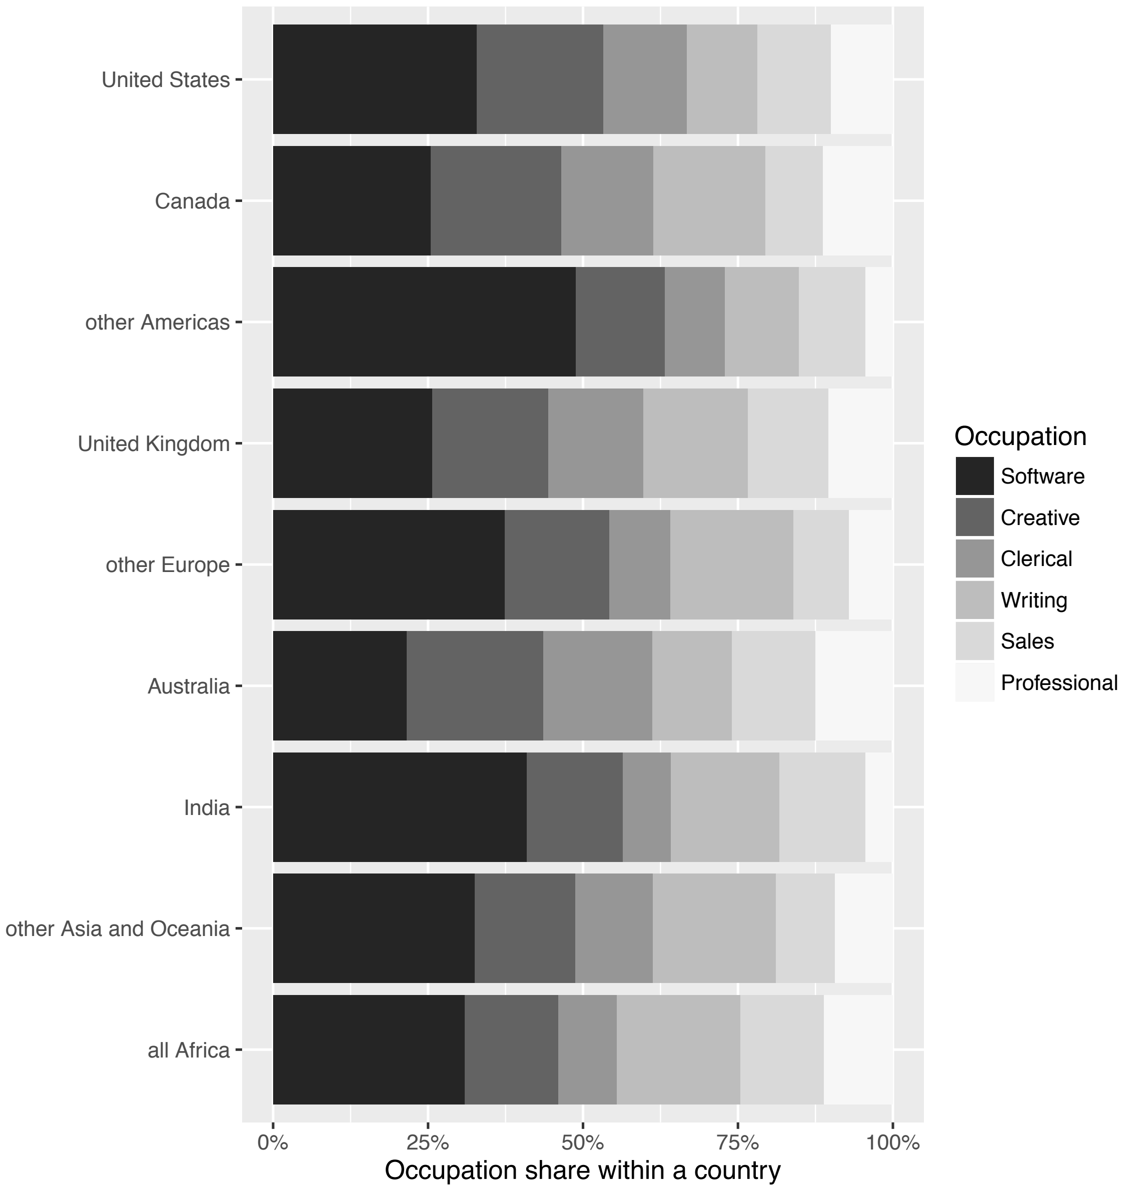

Figure 5. Online Labour Index share by country and occupation, May-October 2016

labour markets, which remained close to stagnant in the UK and US according to latest national statistics. As the OLI accumulates more data, it will reveal whether this represents a lasting trend. As for who is affected by this rise of the online gig economy, the OLI showed that software development and technology are currently the most sought-after skills, followed by creative and clerical work. Any future dips in conventional labour market statistics for these occupations should be checked against the OLI to see if employers are in effect substituting online gig work for standard employment.

The OLI also showed that employers in the United States are currently by far the biggest buyers of online gig work, representing over half of the market. However, UK and Europe-based employers' growth outpaced the growth of US employers. It will be interesting to see to what extent the US can maintain its considerable lead, especially as US-based online labour platforms are starting to expand their marketing efforts to other countries. Another striking feature of the geography of online work is that the occupational demand profiles for all of the employer countries are remarkably similar. This suggests that it is mainly the information technology industry in each country that is currently making use of online labour. If and when employers in other industries enter the online labour market, the OLI should show the national demand profiles diverging.

Beyond the static picture presented in this paper, our results are published online and updated in near-real time at http://ilabour.oii.ox.ac.uk/online-labour-index/. The interactive visualization tool allows anyone to produce graphics similar to the ones presented in this paper. The raw data used to produce the visualizations is also available as an open data set. We believe that the OLI will be a useful tool for policy makers, researchers, and investors striving to make sense of how the platform economy is developing and where its effects are being felt. An important advantage of the index over existing work is that it is continuously updated, yielding over time a methodologically consistent time series similar in power to conventional labour market statistics.

At the time of writing, one important dimension missing from the OLI is the geography of supply: where are the workers located who are filling online vacancies across different occupations? This would be important information for understanding how online labour platforms are contributing to new international divisions of labour. It would also yield further insight on the reasons behind online labour platforms' growth in different industries and occupations, whether it be cost-cutting or tapping specialized skills.

Another important limitation is that the OLI is currently limited to tracking English-language online labour platforms. The English-language market is currently the largest, probably by far Kuek et al. (2015), and English-language platforms are used across the world. However, in future updates we plan to augment OLI with the capacity to track platforms in other languages.

Acknowledgements

We are grateful to our colleagues Gretta Corporaal, Sumin Lee, Gina Neff, Sanna Ojanperä, Godofredo Ramizo, Alex Wood and the participants at the IPP 2016 conference and at the Symposium on Big Data and Human Development in Oxford for various useful questions and comments. Any remaining errors are our own.

References

Autor, D. H. and Handel, M. J. (2013) Putting Tasks to the Test: Human Capital, Job Tasks, and Wages. *Journal of Labor Economics*, **31**, 59–96.

BLS (2015) Labor Force Statistics from the Current Population Survey. http://www.bls.gov/cps/faq.htm#Ques1.

Brynjolfsson, E. and McAfee, A. (2014) *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, 1st edn.

Corrado, C. A. and Hulten, C. R. (2015) How Do You Measure a "Technological Revolution"? *American Economic Review*, **100**, 99–104.

Coyle, D. (2015) Commentary: Modernising Economic Statistics: Why It Matters. *National Institute Economic Review*, **234**, F4–F7.

— (2016) The Sharing Economy in the UK. *Tech. rep.*, Sharing Economy UK.

Difallah, D. E., Catasta, M., Demartini, G., Ipeirotis, P. G. and Cudre-Mauroux, P. (2015) The Dynamics of Micro-Task Crowdsourcing: The case of Amazon MTurk. In *International World Wide Web Conference (WWW)*, 238–247.

Elias, P. (1997) Occupational Classification (ISCO-88): Concepts, Methods, Reliability, Validity and Cross-National Comparability. OECD Publishing.

European Commission (2015) Online Platforms. https://ec.europa.eu/digital-single-market/en/online-platforms-digital-single-market.

Evans, D. and Schmalensee, R. (2016) *Matchmakers: The New Economics of Multisided Platforms*. Brighton, MA: Harvard Business Review Press.

Farrel, D. and Gregg, F. (2016) Paychecks, Paydays, and the Online Platform Economy: Big Data on Income Volatility. https://www.jpmorganchase.com/corporate/institute/report-paychecks-paydays-and-the-online-platform-economy.htm.

Friedman, J. H. and Hastie, T. (2010) Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, **33**, 1–22.

García, V., Mollineda, R. A. and Sánchez, J. S. (2009) Index of balanced accuracy: A performance measure for skewed class distributions. In *Lecture Notes in Computer Science*, vol. 5524 LNCS, 441–448.

Groen, W. P. D. and Maselli, I. (2016) The Impact of the Collaborative Economy on the Labour Market. URL: `https://www.ceps.eu/system/files/SR138CollaborativeEconomy_0.pdf`.

Hastie, T., Tibshirani, R. and Friedman, J. (2008) *The Elements of Statistical Learning*. New York, NY: Springer, 2nd edn.

Hussmanns, R. (2007) Measurement of employment, unemployment and underemployment: Current international standards and issues in their application. URL: `http://www.ilo.org/global/statistics-and-databases/WCMS_088394/lang--en/index.htm`.

Huws, U. and Joyce, S. (2016) Crowd Working Survey: Size of the UK's "Gig Economy" revealed for the first time.

ILO (2012) International Standard Classification of Occupations ISCO-08.

Ipeirotis, P. G. (2010) Analyzing the Amazon Mechanical Turk marketplace. *XRDS: Crossroads*, **17**, 16–21.

Kokkodis, M. and Ipeirotis, P. G. (2015) Reputation Transferability in Online Labor Markets. *Management Science*, **1909**, 1–20.

Kokkodis, M., Papadimitriou, P. and Ipeirotis, P. G. (2015) Hiring Behavior Models for Online Labor Markets. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining - WSDM '15*, 223–232.

Kuek, S. C., Paradi-Guilford, C., Fayomi, T., Imaizumi, S. and Ipeirotis, P. (2015) The Global Opportunity in Online Outsourcing. Washington, D.C. : World Bank Group. http://documents.worldbank.org/curated/en/2015/06/24702763/global-opportunity-online-outsourcing.

Lehdonvirta, V., Barnard, H., Graham, M. and Hjorth, I. (2014) Online labour markets - leveling the playing field for international service markets? URL: `http://ipp.oii.ox.ac.uk/sites/ipp/files/documents/IPP2014_Lehdonvirta_0.pdf`.

Lehdonvirta, V. and Ernkvist, M. (2011) Knowledge Map of the Virtual Economy: Converting the Virtual Economy into Development Potential.

Lehdonvirta, V., Hjorth, I., Graham, M. and Barnard, H. (2015) Online Labour Markets and the Persistence of Personal Networks: Evidence From Workers in Southeast Asia. *Tech. rep.*

Parker, G., Van Alstyne, M. and Choudary, S. P. (2016) *Platform revolution : how networked markets are transforming the economy and how to make them work for you.*

Sundararajan, A. (2016) *The Sharing Economy.* MIT Press.

Susskind, R. and Susskind, D. (2015) *The Future of the Professions - How Technology Will Transform the Work of Human Experts.* Oxford, United Kingdom: Oxford University Press, 1st edn.

The Conference board (2016) The Conference Board Help Wanted OnLine® (HWOL). https://www.conference-board.org/data/helpwantedonline.cfm.

Tibshirani, R. (1996) Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological) Journal of the Royal Statistical Society. Series B J. R. Statist. Soc. B*, **58**, 267–288.

Wood, A., Graham, M. and Lehdonvirta, V. (2016) The New Frontier of Outsourcing: Online Labour Platforms and the Consequences for Poverty in the Global South. In *Work, Employment and Society Conference 6-8.9.2016, Leeds, U.K.*

A. Appendix

Following Friedman and Hastie (2010), we denote the probability that observation i is in the occupation category M ($M = 1, ..., 6$), conditional on the observed stemmed word counts x, as

$$Pr\left(\text{Occupation}_i = M \mid x_i\right) = P_{Mi}(\beta \mid x_i) = \frac{e^{\alpha_M + x_i' \beta_M}}{\sum_{k=1}^{6} e^{\alpha_k + x_k' \beta_k}} \tag{1}$$

To estimate the parameter vectors $\beta = \{\alpha_{0M}, \beta_M\}$ for all $M = 1, ...6$, we maximise the following objective function

$$\underset{\beta}{max} \left[\frac{1}{N} \sum_{i=1}^{N} \sum_{k=1}^{6} Y_{ik} log P_{Mi}(\beta \mid x_i) - \lambda \sum_{k=1}^{6} \parallel \beta_k \parallel \right], \tag{2}$$

where $P_{Mi}(\beta_k \mid x_i) = Pr\left(\text{Occupation}_i = M \mid x_i\right)$, and Y_{ik} is an $(N \times 6)$ indicator response matrix, where each row has value 1 for the column where $\text{Occupation}_i = M$, and zero otherwise. $\parallel \beta_k \parallel$ is a of vector norm of β_k. The first element of the maximand is the standard log-likelihood function. The estimation boils down to choosing a value of λ, and the corresponding vector β which minimises the cross-validated mean deviance.

Figure 5 presents the confusion matrix of our classifier. The shares of correctly predicted classes – i.e. the *precision* of the classifier – are visible from the diagonal of the confusion matrix. By far, our accuracy is the highest in the *software development and technology* occupation.This is to some extend driven by the fact that our training data is unbalanced; over 50% of the projects are in the *software development and technology* category, whereas only roughly 5% of the projects are in the *professional services* category.

Figure A.1 gives an indication of how much confusion there is within the occupation categories. This can be read from the columns of the confusion matrix. For instance, in our training set, we see that $\frac{13}{1+13+2+11+2+1} \approx 43\%$ of the *professional services* vacancies were classified as *software development and technology* vacancies. Since the training set is a random sample of the vacancies, our best estimate is that 22% of the true *professional services* vacancies are misclassified as *Software development and technology* vacancies. In general, we see that the most common type of misclassificaion is that a vacancy is erroneously classified as a *software development and technology.*

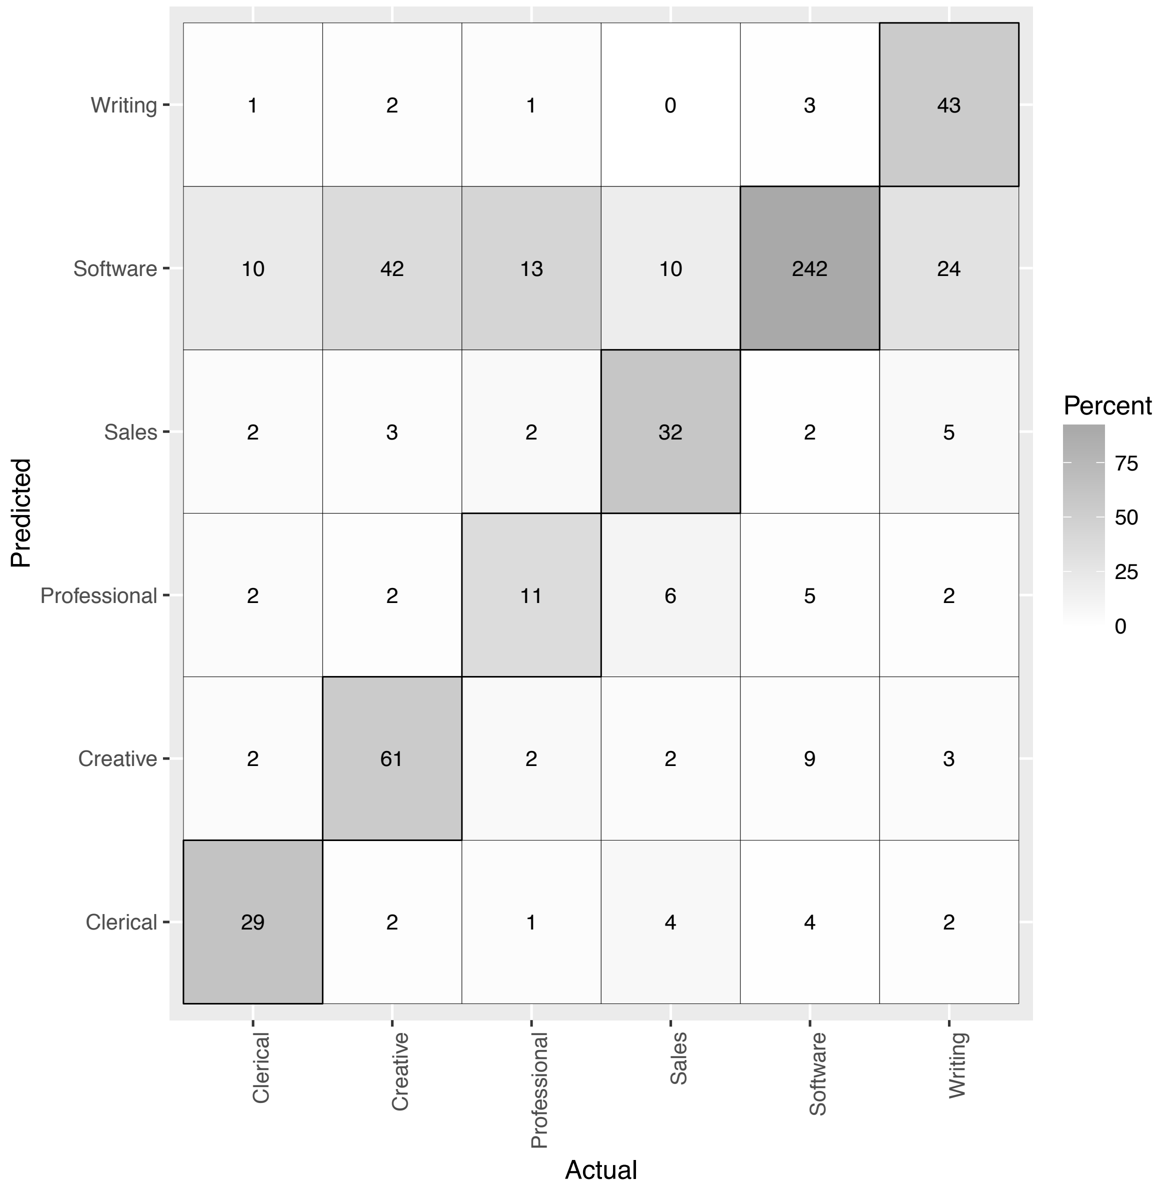

Figure A.1. Confusion matrix of the regularised multinomial logistic classifier. The cell colouring corresponds to percentage shares relative to column sums (i.e. sensitivity of the classifier).