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

Otto Kässi, Vili Lehdonvirta*

21st September 2016

otto.kassi@oii.ox.ac.uk, vili.lehdonvirta@oii.ox.ac.uk

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 on 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 questions about the online gig economy that are crucial for policy and research. To benefit policymakers, digital labour market researchers and the general public, our results are published in an interactive online visualisation which is updated daily.

*Both authors are at University of Oxford, Oxford Internet Institute, United Kingdom.

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 appeared 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. 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 temporal flexibility, 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. By tracking the utilization of online labour platforms across countries and occupations in near real time, 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.

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.

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. Existing economic statistics are in general prone to mismeasure the value of digital activities and investments, because they are often not directly related to production, but to development, design, and marketing, whose value is 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. In addition, it is not clear to what extend 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 large share of online workers living in developing countries, where the informal economy dominates and tax evasion is common. 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. 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 addressing the potentially rapid temporal patterns of online work, relevant to many social 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. 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) for 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. 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 online work. 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 online labour platforms more generally. A general online labour index – something comparable in scope and functions to national labour market indices – is currently missing from digital economy research and policy.

3 Sample selection and data collection

3.1 Sample selection

The Online Labour Index is an index that measures the utilization of online labour platforms over time and across countries and occupations. Online labour platforms are here understood as platforms through which buyers and sellers of labour or 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, excluding platforms for local services such as Uber and Airbnb. The index is based on tracking all projects and tasks posted into a selected sample of platforms, using API access and web scraping.

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 measurements for all major websites around the world, based on a voluntary plugin that observes browsing behavior.[1]

To approximate the coverage of the OLI, we collected a list of 40 prominent English-language online labour platforms, retrieved their monthly estimated unique visitor counts from Alexa, and selected the top five. This sample is listed in Table 1. Using Alexa's figures, we estimate that these five account for at least 60% of all traffic to English-language online labour platforms. They also represent a range of different market mechanisms and contracting styles, from online piecework to hourly freelancing.

[1] For details, see http://aws.amazon.com/alexa-top-sites/ (accessed 2016-07-19).

	Alexa rank	**Monthly unique visitors (est)**
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 1: Traffic of the platforms currently included in the index.

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 status of each vacancy: open, in progress, or completed. Comparing changes in statuses allows us to calculate the number of new and filled vacancies between two crawls. 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 exist, and remain unmeasured, in traditional vacancy statistics as well. Notwithstanding these caveats, we believe that our measure fulfils its purpose of tracking the volume of work transacted on the platforms.

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 for 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 available on each platform is summarised in Table 2 and discussed in more detail in the following sections.

	Occupation classification	Employer countries observed
Freelancer.com	Based on platform taxonomy	-
Guru.com	Predicted	Observed
Upwork.com	Based on platform taxonomy	-
Mturk.com	Predicted	-
Peopleperhour.com	Predicted	-
Upwork.com	Platform taxonomy	Observed

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

3.3 Classification of work done on platforms

In order to classify the work done on the various platforms, the disparate classifications utilised across platforms need to be normalised. For the platforms that provide a taxonomy, we manually map the occupation taxonomies to 6 broadly similar occupation classes outlined in Table 3. The 6 classes are adopted from the existing classification used in Upwork.com, and by all accounts capture the main contours of online work relatively well. Similar classification is also used in previous literature (see, e.g., Kokkodis and Ipeirotis 2015; Kokkodis et al. 2015).

We start by briefly discussing some details of the classification adopted. First, the typical 'microwork' vacancies which include tasks like data entry, image classification fall in the *Clerical and data entry* category. The main difference between *Clerical and data entry* and *Professional services* categories is that the latter mostly require formal education and knowledge about local institutions, whereas the former have less stringent skill requirements. The *Sales and marketing support* are largely support tasks related to online advertising. They are separated from the two other aforementioned categories because they form a large and distinct portion of online freelancing.

The mapping follows the philosophy of standard international occupation classifications such as the International Standard Classification of Occupations (ISCO) which define an occupation as a "*set of jobs whose main tasks and duties are characterised by high degree of similarity*" (see ILO (2012) pp. 59-60). In contrast to ISCO we are not able to measure the skill level

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

Table 3: Classification of occupation types on platforms

required to accomplish a particular task within an occupation.

The mapping has an obvious caveat. Namely, there are some occupations whose class is not clear. For example, take a web site design vacancy which includes both graphical design and programming of the web site. In this case, the vacancy could either be either classified as a design and creative vacancy, or as a programming vacancy. This caveat is not specific to our occupation classification, but is present in all empirical studies studying occupational groups (see e.g. Sullivan (2010) pp. 568-569 for discussion). Nonetheless, we argue that in our case this problem is smaller because in the case of typical labour force surveys the classification of occupations is done retrospectively after a vacancy is posted on a job board, but in our case it is in the interest of the employer posting a job ad to classify it in a correct fashion to get the best matching pool of applicants to their vacancy.

3.4 Predicting the unobserved occupation classes

The platforms do not expose their occupation taxonomies for some 15% of the vacancies. In order to classify these vacancies, we employ a machine learning process, which is discussed next.

We took a random sample of about 1172 vacancies from the set of vacancies with an unobserved occupation class. The 1172 sampled job vacancies were manually classified to 6 occupations. We used the 1172 manually classified projects as the training data set for our classifier. We first processed the projects' titles and descriptions by removing stopwords, and stemming the other observed words. As a result, we ended up with a 1172×2951 matrix where each row represents a project, and every column represents the word count of different stemmed words, which are the predictive features in our model.

We used a regularised multinomial logistic classifier implemented in R's glmnet package (Friedman and Hastie, 2010).[2] We evaluated our prediction by randomly splitting our training sample in two, and fitted a model with

[2] We also experimented with a Support Vector Machine classifier, but it performed worse in a validation data set.

	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%]			

Table 4: Classification precision metrics. 95% confidence interval for total precision in brackets.

half of the data, predicted the occupation classes for the other half, and compared our predictions to the observed occupation classes. We present the confusion matrix of our classification model along with more details of the classification in Appedix A.

Table 3 describes the various classification accuracy metrics calculated from the learning data. Since the occupation categories in the learning data are unbalanced, our preferred accuracy metric is the *balanced accuracy*, which accounts for the unbalanced occupation proportions in the learning data (García et al., 2009). All in all, both the balanced accuracyand the aggregate precision measures demonstrate that the regularised multinomial regression performs well in our data. It reaches a balanced accuracy of over 75% in 5 of the 6 occupation categories, and a total accuracy of 70%.

Column 3 of Table 2 lists how we infer the type of the opening for each of the platforms we observe.

3.5 Employer country distribution

We save the information on the home countries of the employers of posted vacancies in cases when this information is available. Since the employer country distribution is highly skewed with 5 top countries adding up to over 90% of all posted projects. For visualisation purposes, we group the smaller countries into geographical groups. Further, to reduce the number of requests made to the platforms, we only fetch the country info for a random sample of all projects.

Column 4 of Table 2 lists the platforms for which we are able to collect the home country of the employer.

4 Applying the Online Labour Index

4.1 Leading occupations in online work

In this section, we illustrate how the OLI can be used to address some of the crucial policy questions associated with the rise of the online gig economy. 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 1 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. Any future dips in conventional employment statistics in these occupations should be checked against the OLI to see if they are being offset by corresponding increases in online work, suggesting that employers are moving their vacancies online.

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

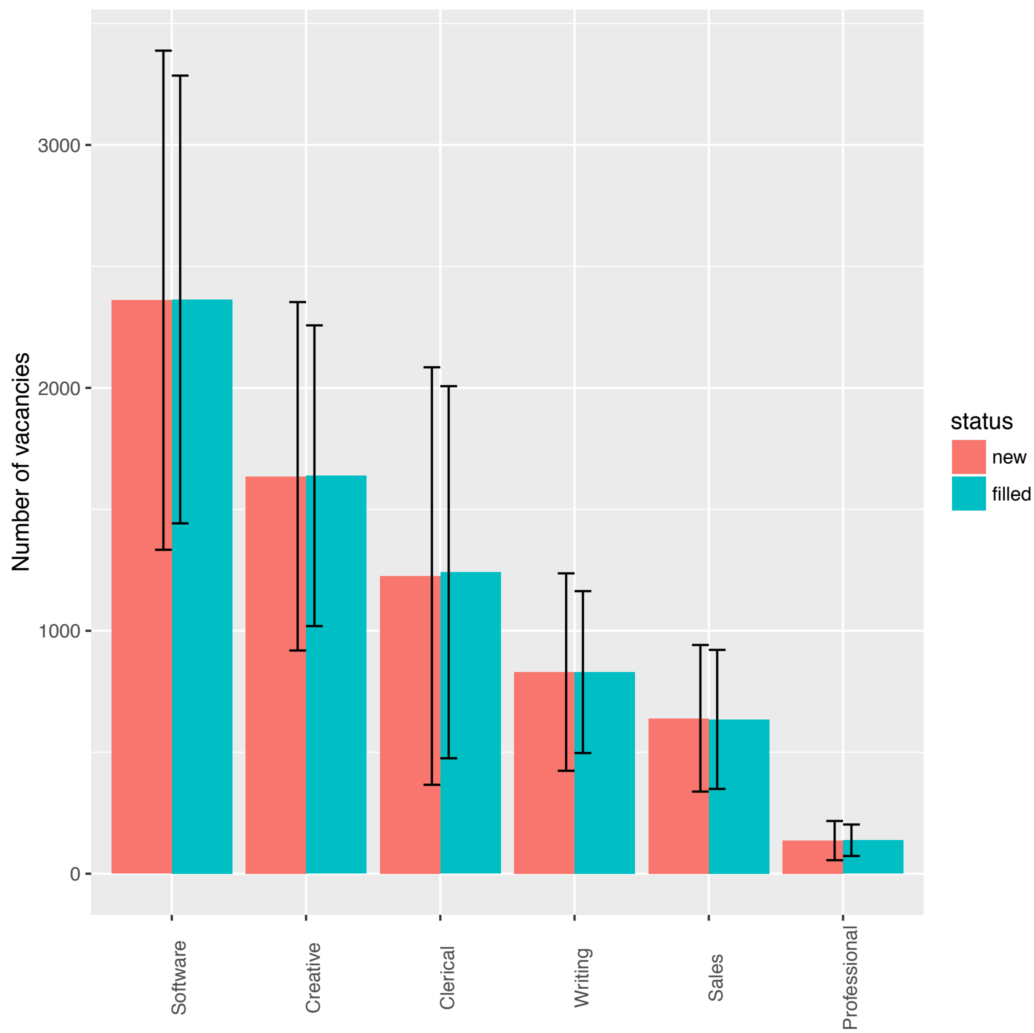

Figure 1: New and filled vacancies by occupation class.

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). Since the OLI is constantly updated, it will soon reveal whether the use of online labour platforms for procuring professional services is growing, and by how much.

4.2 Geography of demand for online work

Another crucial question is the geography of online work. Which countries are affected? 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 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 2. Across occupations, roughly half of the vacancies are posted by employers from the United States. Other prominent employer countries include the United Kingdom, India, Australia, and Canada. 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 (2015). But India also has a large IT sector of its own, which is likely to be generating domestic demand for online workers.

A striking feature of the geography of online labour 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 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 occupa-

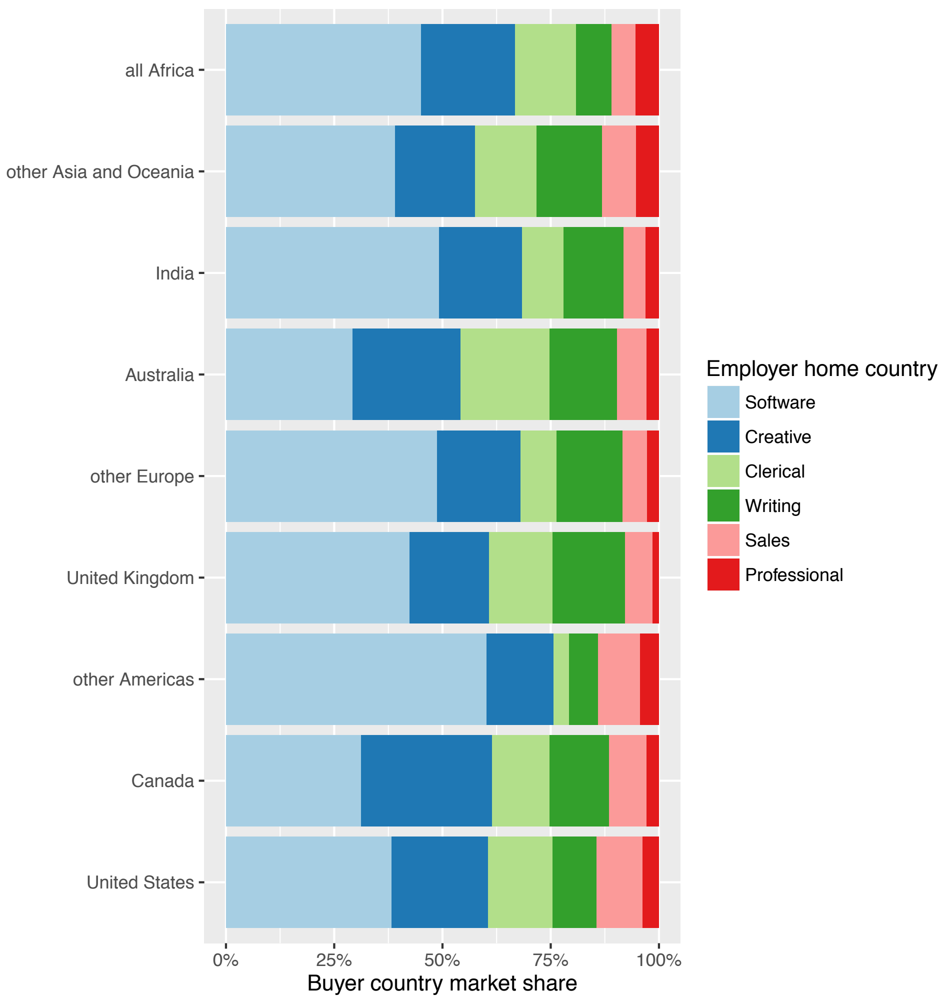

Figure 2: Employer market shares by country and occupation class.

tional profiles of their conventional domestic labour markets. The fact that they nevertheless resemble each other rather much in online labour demand profiles suggests that the demand probably 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.

4.3 Growth and growth potential in online work

We saw that most of the volume in today's online gig economy is in software, creative, and clerical work, and that most of the demand is coming from the United States. In time we will be able to produce OLI time series plots that span months and years, and use them to compare the growth trajectories of different occupations and countries, and to make future predictions. But at the time of writing the data only spans four months, too short a period for detecting any larger trends. However, the data allows us to assess the growth potential of different occupations in the online gig economy via a different method: examining vacancy filling times.

Figure 2 shows that the supply and demand bars for all the occupations were roughly in balance. This implies that the projects posted on platforms eventually get completed with a high probability. But there are actually considerable occupational differences in the average survival times, that is, the time that it takes for an employer to find a suitable contractor. These differences are depicted in Figure 3, which plots the distribution of days it takes to fill a vacancy in different occupations. Most occupations have modal filling times of less than a week. That is, a typical vacancy is filled in less than 7 days after its posting. However, there are considerable differences across occupations. In particular, a considerable proportion of Clerical and data entry and Software development and technology vacancies are filled within a week, while there is much more mass in the upper tail of the time-to-fill distribution of Professional services and Sales and marketing support. This

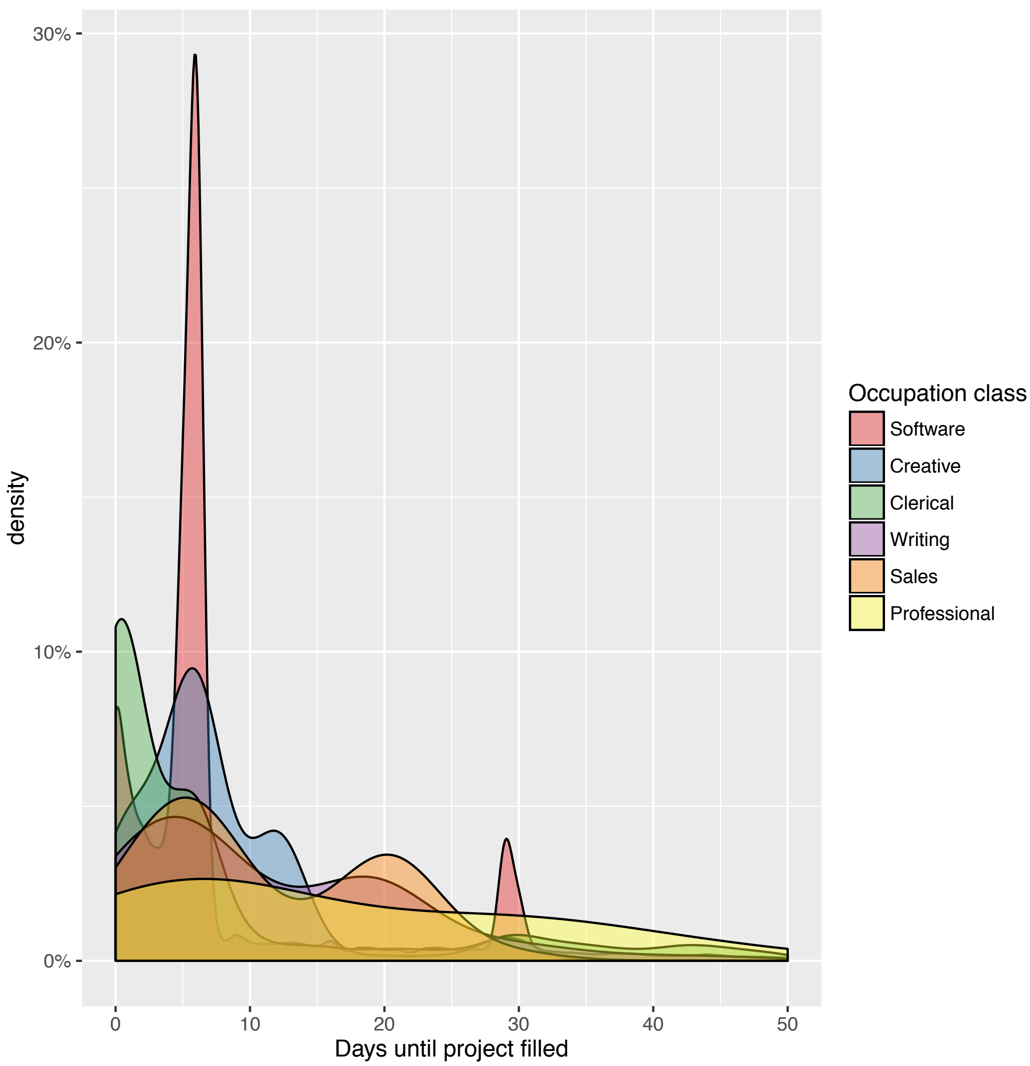

Figure 3: Empirical densities of project filling times.

	10%	30%	50%	70%	90%	Propoprtion open after 60 days
Professional services	1	6	16	30	53	4.9%
Clerical and data entry	0	0	3	6	34	0.9%
Design and creative	0	5	6	8	13	0.2%
Sales and marketing services	1	6	6	20	21	0.2%
Software development and technology	1	5	6	6	29	1.4%
Writing and translation	0	6	6	19	29	1.9%

Table 5: Quantiles of empirical cumulative distribution of days until project is filled.

suggests that there is more room for new workers specialising in professional services and sales to enter the market, whereas the market for clerical and software workers is comparatively more saturated.

To dig deeper into the filling times in different occupations, the empirical cumulative distributions of filling times are given in Table 5. For instance, the table shows that in Professional services, 10 percent of vacancies are filled in less than a day; 30 percent are filled in less than 6 days; 50 percent are filled in less than 16 days, and so on. There are still 4.9 percent of vacancies open after 60 days, at which point it is likely that the vacancy will not be filled at all. Table 5 supports the observations we made from Figure 3: Clerical and data entry vacancies are filled much quicker than other project types. Conversely, Professional services vacancies take most time to fill. The expected closing time of a vacancy seems to bear some relationship with the skills involved. If a project requires specialised skills such as writing, programming, or accounting, it takes on average longer to fill it than if a project requires only relatively basic computer literacy and numeracy skills, as is often the case with clerical and sales support tasks. Overall, this suggests that from the demand perspective, the growth potential for skilled online labour is currently greater than for less skilled labour.

4.4 Temporal patterns of online work

Finally, we briefly examine the temporal patterns of online work. Temporal patterns are interesting for policy and research, because they are directly

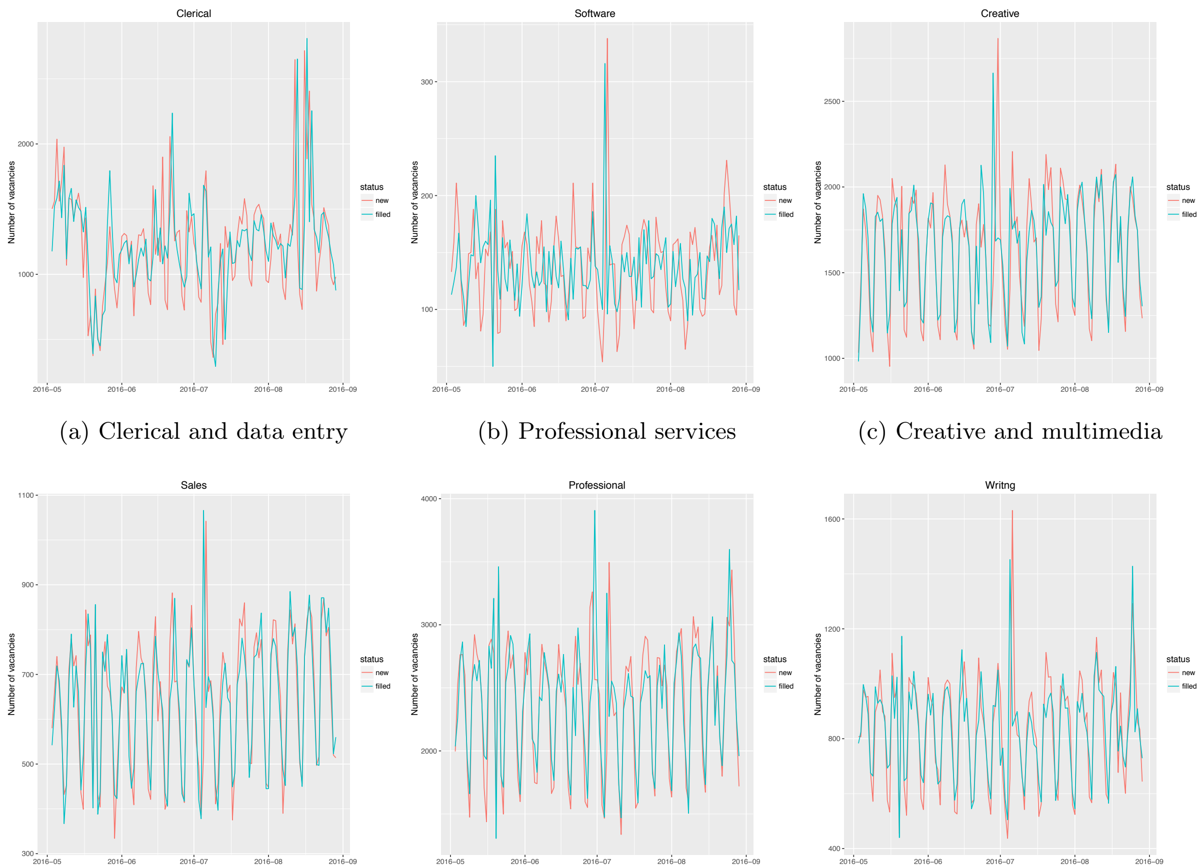

(a) Clerical and data entry

(b) Professional services

(c) Creative and multimedia

(d) Sales and marketing support

(e) Software development and technology

(f) Writing and translation

Figure 4: Dynamics of new and filled vacancies by occupation class

related to questions about workers' ability to combine work with other commitments, such as studies or caring duties. Going back to Figure 1, an observant reader might have noticed that the standard deviations depicted as T-bars were greater for new vacancies than for filled vacancies. This suggests that day-to-day variation in new vacancies is greater than the variation in completed of vacancies. One explanation for this could be that employers mostly post new vacancies on weekdays, whereas contractors are active also during weekends. This, in turn, is consistent with the observation presented in earlier literature that many contractors work on platforms for additional income outside of their regular working hours. Difallah et al. (2015) likewise note that there is strong weekly periodicity of arrival of new tasks on Mechanical Turk, but much less periodicity in task completion.

Figure 4 shows how the OLI can be used to dig deeper into this issue, by plotting new and filled vacancies as a time series. It shows that both new and filled vacancies exhibit significant weekly variation: a noticeable dip takes place each weekend. The dip in the filled vacancies means that online workers are getting some respite from work during weekends, even if some work still does get done. But the dip in the new vacancies means that this may not always be by choice: even if a worker wanted to work more because of other commitments during weekdays, less work is available in the online gig economy during weekends.

5 Discussion and conclusions

In this paper, we introduced the Online Labour Index, a new economic indicator that provides an online labour market equivalent of conventional labour market statistics. We described how the OLI is constructed and illustrated how it can be used to address crucial policy issues that existing data sources are unable to address. On the question of who is affected by the rise of online labour markets, 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 the conventional labour market statistics for these occupations should be checked against the OLI to see if employers are moving their vacancies online. The online market for professional services occupations remains smaller in comparison, but shows growth potential, judging by the fact that half of open vacancies take over two weeks to fill. As the OLI is constantly updating with new data, it will soon be apparent whether the market is growing and by how much.

The OLI also showed that employers in the United States are by far the biggest users of online labour at the moment, followed by the United Kingdom, India and Australia. In the future it will show to what extent the U.S. can maintain its considerable lead, as U.S.-based online labour platforms 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, OLI should show the national demand profiles diverging.

We also used the OLI to address the question of how the online gig economy is influencing everyday life. Across all occupations, there is a strong relationship between the quantity of vacancies posted in a day and the quantity of vacancies filled, lending support to the idea that online work is 'on-demand' work that workers must adapt their schedules to. However, there are occupational differences in the vacancy filling times, with clerical and data entry work currently appearing to be the most 'on-demand' in nature, and professional services the least. On a weekly level, the market is roughly twice as busy or more on weekdays as it is on weekends. Despite being technically open every day, the online labour market does not seem to have entirely erased the concept of weekends. In fact, employers are stricter about not posting vacancies during weekends than workers are about filling them over weekends. This lends support to the idea that online gig work is used at least to some extent as a secondary job, outside main working hours or studies. Over time the OLI will allow us to monitor how weekend vacancies and other temporal aspects of online work continue to develop.

Beyond the static picture of online labour markets 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, and far more. The raw data used to produce the visualizations is also available. We belive 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 to reaching out to specialized skills. Existing studies such as Kuek et al. (2015) and Lehdonvirta et al. (2014) provide glimpses of this information, but we also plan to add this dimension to the OLI in a future update.

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.

References

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 edition.

Corrado, C. A. and Hulten, C. R. (2015). How Do You Measure a "Technological Revolution"? *American Economic Review*, 100(2):99–104.

Coyle, D. (2015). Commentary: Modernising Economic Statistics: Why It Matters. *National Institute Economic Review*, 234(1):F4–F7.

Coyle, D. (2016). The Sharing Economy in the UK. Technical report, 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)*, pages 238–247.

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*. Harvard Business Review Press, Brighton, MA.

Farrel, D. and Gregg, F. (2016). Paychecks , Paydays , and the Online Platform Economy: Big Data on Income Volatility.

Friedman, J. and Hastie, T. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, 33(1):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*, volume 5524 LNCS, pages 441–448.

Groen, W. P. D. and Maselli, I. (2016). The Impact of the Collaborative Economy on the Labour Market.

Hussmanns, R. (2007). Measurement of employment, unemployment and underemployment: Current international standards and issues in their application. Technical report, ILO Bureau of Statistics.

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*, I:1–420.

Ipeirotis, P. G. (2010). Analyzing the Amazon Mechanical Turk marketplace. *XRDS: Crossroads*, 17(2):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*, pages 223–232.

Kuek, S. C., Paradi-Guilford, C., Fayomi, T., Imaizumi, S., and Ipeirotis, P. (2015). The Global Opportunity in Online Outsourcing.

Lehdonvirta, V. (2015). Do platforms connect clients directly to providers? The new network patterns of digital work.

Lehdonvirta, V., Barnard, H., Graham, M., and Hjorth, I. (2014). Online labour markets - leveling the playing field for international service markets?

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

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.*

Sullivan, P. (2010). Empirical evidence on occupation and industry specific human capital. *Labour Economics*, 17(3):567–580.

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 University Press, Oxford, United Kingdom, 1st edition.

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

Appendix A

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.

Further, Figure 5 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.*

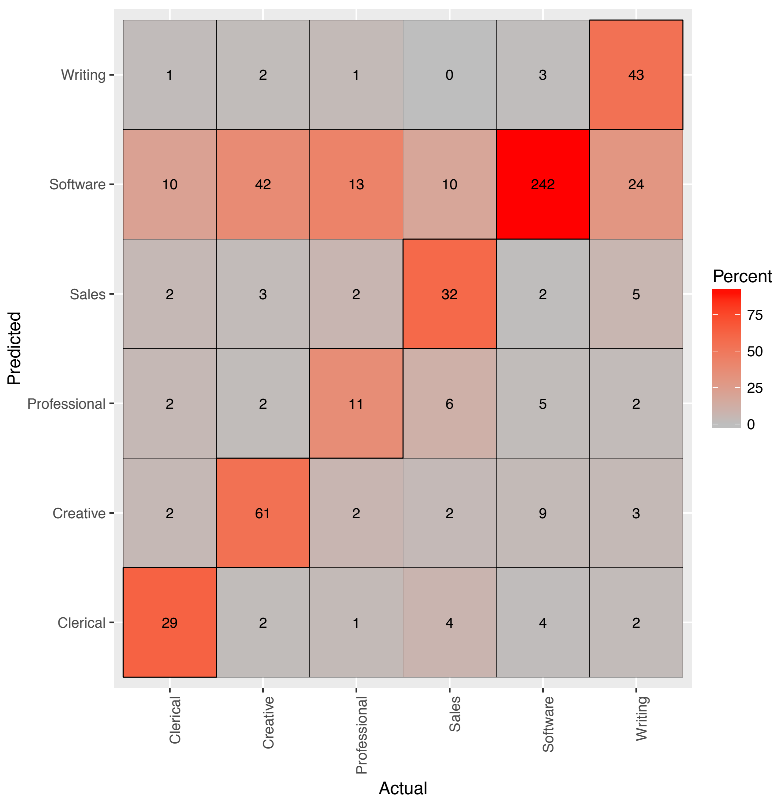

Figure 5: Confusion matrix of the regularised multinomial logistic classifier. The cell colouring corresponds to percentage shares relative to column sums (i.e. sensitivity of the classifier).