

Online Labour Index: Measuring the Online Gig Economy for Policy and Research¹

Otto Kässi

otto.kassi@oii.ox.ac.uk

Vili Lehdonvirta

University of Oxford, 1 St Giles, Oxford, Oxon OX13JS, 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 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 this “online gig work”. We present the Online Labour Index (OLI), an experimental economic indicator that approximates the conventional labour market statistic of new open vacancies. It measures the utilization of online labour across countries and occupations by tracking the number of projects and tasks posted on major online gig platforms in near-real time. The purpose of this article is to introduce the OLI and describe the methodology behind it. We also demonstrate how it can be used to address previously unanswered questions about the online gig economy. To benefit policymakers, labour market researchers and the general public, our results are published in an interactive online visualisation which is updated daily.

Keywords: online freelancing, online labour markets, online gig economy, measurement, statistics, measurement of vacancies, web data collection

¹ This is the last working paper version of an article forthcoming in *Technological Forecasting and Social Change*

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 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 this “online gig work”. We present the Online Labour Index (OLI), an experimental economic indicator that approximates the conventional labour market statistic of new open vacancies. It measures the utilization of online labour across countries and occupations by tracking the number of projects and tasks posted on major online gig platforms in near-real time. The purpose of this article is to introduce the OLI and describe the methodology behind it. We also demonstrate how it can be used to address previously unanswered questions about the online gig economy. To benefit policymakers, labour market researchers and the general public, our results are published in an interactive online visualisation which is updated daily.

Labour markets are thought to be in the midst of a dramatic transformation, where standard employment is increasingly supplemented or substituted by temporary 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 this “online gig work”. We present the Online Labour Index (OLI), an experimental economic indicator that approximates the conventional labour market statistic of new open vacancies. It measures the utilization of online labour across countries and occupations by tracking the number of projects and tasks posted on major online gig platforms in near-real time. The purpose of this article is to introduce the OLI and describe the methodology behind it. We also demonstrate how it can be used to address previously unanswered questions about the online gig economy. To benefit policymakers, labour market researchers and the general public, our results are published in an interactive online visualisation which is updated daily.

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 utilisation of online labour platforms over time and across countries and occupations, providing a solid evidence base for future policy and research. The OLI is published online as an automatically updating open data set and interactive visualisation at <http://ilabour.oii.ox.ac.uk/online-labour-index>. The main contribution of this paper is to describe the index and the methodology underpinning it. We also illustrate how the index can be used to address three empirical questions concerning online gig work that existing data sources have difficulty addressing: Is online gig work growing? Where is it growing? In what occupations it is growing? We also briefly discuss the current limitations and planned extensions of the index.

It is worth emphasising that the main goal of this paper is exploratory: we assemble a new type of data set, and describe it. We publish our data set to stimulate further research, but refrain in this paper from elaborate hypothesis testing and theory development.

2 Background

The impacts of technological change on jobs have been a topic of much interest over recent decades. Some of the most influential studies have focused on the changing skill content of work (eg., Di Pietro, 2002; Frey and Osborne, 2017; Rumberger and Levin, 1985) and on the geography of job creation and destruction (eg., Howland, 1993). At least as important is to understand the impacts of technologies on the nature of the employment relationship itself. In particular, ICT adoption has been associated with increases in working remotely or ‘teleworking’ (Pratt, 1984), in temporally flexible working arrangements and in casual or project-based labour contracts (Aubert-Tarby et al., 2018). These three trends challenge standard full-time employment on multiple fronts.

The rise of online gig work can be seen as a conjunction of all three trends. Instead of undertaking full-time employment at the premises of a single employer, workers serve multiple clients at varying hours remotely from their homes or co-working spaces. If such a mode of working is becoming more common in some countries and occupations, it carries obvious and important policy and business implications. Yet 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 (Abraham et al., 2017). There are several reasons

for this. Existing economic statistics are in general prone to mismeasuring the value of digital activities and investments, because these are often not directly related to production, but to development, design, and marketing, whose value is harder to establish (Corrado and Hulten, 2015). 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 (Husmanns, 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 additional efforts are not captured in employment statistics. In addition, it is 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 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 local 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 Ernkqvist (2011), and Kuek et al. (2015) 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% 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 % of adults in the sample had earned income from the platform economy in each month, and that this figure had grown over tenfold 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. Following a similar approach, Katz and Krueger (2016) found that roughly 0.5 % of the U.S. labour force was employed through online labour platforms. The significant difference in the figures reported by Huws and Joyce (2016) and Katz and Krueger (2016) could be explained by the inclusion of non-traditional payment channels by the former, or by US-UK national differences, but they could also be explained by other methodological differences and differences in concepts and definitions. A notable methodological weakness in both Huws and Joyce's (2016) and Katz and Krueger's (2016) studies 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 restarted 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, the Contingent Worker Supplement will not address the significant limitation that these studies are national in scope. It will only produce statistics concerning work completed within the United States. This is problematic, because 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 might be poor at measuring the potentially rapid transformations associated with digitalization, making it hard for policy makers to respond in a timely manner.

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. This suggests 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 (Einav and Levin, 2014; Naccarato et al., 2017; Varian, 2014). 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. Previously, Agrawal et al. (2013) and Horton et al. (2017) gave detailed analyses on the geography, talent and money flows of online labour, but their attention is limited to a snapshot data from a single online labour market. 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 – comparable in scope and function to national labour market indices – is currently missing from digital economy research and policy. The approach in this paper, we believe, is a step in that direction.

3 Methodology

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. There are several reasons for limiting our attention to fully digital labour platforms. Most of the problems related to statistics coverage are pertinent to these types of platforms because of how they span national boundaries. Local statistics authorities have better chances of tracking local labour platforms such as Uber, because the transactions on these platforms always take place within a single country. Online gig platforms and local gig platforms are also conceptually quite different, making it desirable to track them separately.

The index is based on tracking all projects and tasks posted on a sample of platforms, using API access and web scraping. To develop the sample, we first collected a list of 40 prominent online labour platforms, and retrieved their monthly estimated unique visitor counts from Alexa. 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 behaviour.² Based on Alexa's figures, we then selected the top five platforms for which we were able to collect data as our initial sample, accounting for at least 70% of all traffic to online labour platforms³. The resulting sample is listed in Table 1. Constantly changing APIs and web interfaces make it infeasible to continuously track every platform in the market, but with this sample we are able to cover a majority of the market by traffic volume.

The traffic share calculation was conducted near the beginning of the data collection in July 2016 and again in February 2017 and January 2018 yielding roughly 70% each time. A time series plot of the traffic shares of different platforms in Figure 1 furthermore shows that the

² See <http://aws.amazon.com/alexa-top-sites/> (accessed 2016-07-19).

³ This calculation assumes that same users do not use multiple platforms. If there is overlap, the true coverage of our sample is higher.

market shares of the top platforms are quite stable, with the exception of few individual dates. This suggests that any growth observed in our index represents growth in the overall market rather than shifting coverage (ie. tracked platforms winning market share from untracked ones). Our sample is also very likely to represent the majority of the market in terms of gross transaction value (see Kuek et al. 2015, fig. 2.1), and the five sample platforms also represent a range of different market mechanisms and contracting styles, from online piecework to hourly freelancing. In all these different ways, our index should offer a good proxy for observing trends in the overall online labour market. But we also continue to monitor the coverage of the index and introduce new platforms to the sample if and when needed.

A notable limitation of the current sample is that it contains only English-language platforms. However, according to Alexa's figures, the traffic shares of non-English language platforms are quite negligible compared to the top English language platforms. Many platforms do exist in local languages, but it seems that the English-language online labour platforms have to some extent become the global standard, similarly to how social media websites originating in the United States have become globally dominant. In the future we would like to start tracking platforms in other languages as well, but for now we are content that the leading English-language platforms are a satisfactory proxy, as long as it is borne in mind that English-language countries are likely to be overrepresented.

Table 1: Traffic of the platforms currently included in the index (accessed 2016-06-29).

	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




Figure 1: Alexa traffic of platforms tracked.

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 changes in statuses 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 $0, \dots, t-1$, and is observed on period t .

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 generalising 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.

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. Another possible shortcoming of our approach is that we might be double-counting some vacancies that are posted on multiple platforms. In practice this is quite rare: fewer than 1.5% of vacancies posted in August of 2016 were posted on multiple platforms⁴.

We acknowledge that our measure of online labour utilisation is incomplete for two reasons. First, we are not collecting data from all online labour platforms. In addition, as we discuss above, there are some openings we miss even on the platforms we are tracking. Therefore, instead of reporting the absolute number of new vacancies, we report an index number normalised to 05/2016=100. If the share of unobserved to observed vacancies remains constant, the changes in the index measure the changes in new vacancies in online labour markets.

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
Mturk.com	Observed	Unobserved
Peopleperhour.com	Predicted	Unobserved
Upwork.com	Observed	Observed

3.3 Classification of work done on platforms

⁴ This notion is also supported by Horton (2017) who discusses a survey according to which 83% of employers on an online labour platform oDesk had only posted the vacancy on a single platform.

Following the practise of occupation classification in official statistics, we classify the collections of similar projects into occupations. The occupations are “set[s] of jobs whose main tasks and duties are characterised by high degree of similarity” (see ILO (2012) pp. 59-60). This venture has several practical difficulties. First, there is the obvious difficulty that there are many jobs that overlap several and jobs, which might make it impossible to unambiguously map a single worker to any unique occupation. Further, the processes, tasks and skill requirements within jobs are constantly changing, so any attempt to reduce this complexity to any fixed classification has 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 of its reliability. Nonetheless, it is evident that some operational classification is required to facilitate comparisons across countries, companies, time, or online labour market platforms.

We take a very practical approach to the classification of occupations and build on the occupation classification used in Upwork.com, which by all accounts captures the main contours of online work relatively well. For the other platforms, we manually mapped their occupation taxonomies to the Upwork.com’s classification. The classification is outlined in Table 3. Similar classification is also used in previous literature (see Kokkodis and Ipeirotis 2015; Kokkodis et al. 2015). In addition, using qualitative data from online freelancers’ interviews, Wood et al. (2016) finds broadly similar categories of online freelancing.

The typical microwork or human intelligence work vacancies which include tasks like data entry, image classification fall in the Clerical and data entry category. These tasks typically require only basic computer literacy and numeracy. The occupations in the professional services category, on the other hand, typically require formal education and knowledge about local institutions. 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. Writing, Software development and technology, and Creative and multimedia categories are mostly self-explanatory.

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 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 traditional labour market.

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 Software development and technology vacancy. This caveat is not specific to our occupation classification, but is present in all empirical studies studying occupational groups.

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
	Architecture
	Audio
	Logo design
	Photography
	Presentations
	Video production

Sales and marketing support	Voice acting
	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

3.4 Predicting the unobserved occupation classes

The platforms do not expose their occupation taxonomies for 15% of the vacancies. Applying the approach in Amato et al. (2015), we employed a machine learning approach combined with manually classified occupations. We discuss this next.

We took a random sample of 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, special symbols and numbers. The remaining words were stemmed. As a result, we ended up with a 1172×3259

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.

As is evident, our prediction model has more potential explanatory features (3259) than we have observations (1172). To reduce the dimension of our space of predictive features, and to increase the predictive power of our model we applied LASSO. The LASSO method is used to select only a subset of the explanatory features and regularizes their weights to reduce overfitting.

The LASSO method tends to perform well in problems where the set of potential predictive features is large but only a moderate amount of features have predictive power (Tibshirani, 1996). Our model falls in this category (for example names of programming languages and technologies are strong predictors for a project being a software development project).

We present the confusion matrix resulting from applying our model to validation data set along with more details in Appendix A.

3.5 Calculating the employer country distribution

We save the information on the home countries of the employers of posted vacancies in cases when this information is available. Column 3 of Table 2 lists the platforms for which we are able to collect the home country of the employer. To generalise from the subset of the platforms to all of them, we assume that the employer country within occupations are the same across all platforms.

In practise, we calculate the employer country distribution within each occupation on the two platforms we observe. Thereafter, we generalise the within-occupation country distributions to the three platforms for which we do not have this information. Our sampling scheme effectively assumes that the only feature of the platforms which drives the differences in the employers' home countries is the type of work being contracted. This might be a questionable assumption. Nonetheless, we do it for practical reasons, since we are not able to collect country level data from all platforms. Further, since we observe the employer home country for the biggest platform in our data set, Upwork.com, we know the country distribution to be correct for a majority of the new vacancies.

4 Emerging trends in the online gig economy

In this section we use the Online Labour Index to address some of the basic questions about the online gig economy, based on the initial 18 months of data. The focus is on descriptive findings and detecting potentially emerging trends and patterns that should be addressed with further research. Real-time time series data and visualizations are available online in the automatically updated Online Labour Index website.

4.1 Growth of online labour

Figures 2 and 3 plot Online Labour Index from May 2016 to January 2018. Figure 1 shows that the time series of new vacancies exhibit significant weekly variation: a noticeable dip takes place each weekend. This is consistent with the observations in Difallah et al. (2015), who note that there is strong weekly periodicity of arrival of new tasks on Mechanical Turk.

Figure 3 plots a 28-day moving average of the new vacancy time series. We can see from Figure 3 that the index has grown about 21 % (21 index points) from May of 2016 to mid-January 2018. The smoothed time series shows one stark growth period starting from early April and lasting until mid-June. This growth period can be traced back to concurrent marketing efforts by individual platforms.

The observed time series is too short to project any future growth of the online gig market, but the data at hand nonetheless suggests that the growth period starting from mid-January dominates any seasonal effects.




Figure 2: Online labour index time series, normalised to 2016/05 = 100




Figure 3: Online labour index time series, normalised to 2016/05 = 100 (28-day moving average)

4.2 Leading occupations of online work




Figure 4: Distribution of occupations in online labour

Previous studies provide estimates of the total market size, but only limited views of where exactly this market is emerging. Which skills and occupations is it affecting? Figure 4 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 (Carmel & Agarwal, 2002). 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 as discussed in Lacity et al. (2011), 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. Further, these types of tasks may not be as easily codified in rule-based as other types of tasks.

4.3 Geography of demand for online work




Figure 5: Buyer country distribution in online labour

Previous studies such as Kuek et al. (2015), Lehdonvirta et al. (2014), Agrawal et al. (2013) and Horton et al. (2017) 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 5. Since the employer country distribution is highly skewed with top 5 countries adding up to over 90% of the posted projects, we group the smaller countries into geographical groups for visualisation purposes.

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 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.

One potential concern is that since our data is collected exclusively from English language online labour platforms, we might be inadvertently biasing our data collection towards platforms with high U.S. labour demand. Nonetheless, as we argue in Section 3, the largest platforms are also the English language. Therefore, including the largest non-English platforms would only have a minimal influence on the results presented in Figure 5.

Finally, Figure 6 plots the occupation distribution within the countries in Figure 5. 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 occupational 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 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.




Figure 6: Occupation share by country

5 Discussion and conclusions

The online gig economy represents a conjunction of three major transformations: from local to remote, from full-time to temporally flexible, and from permanent to casual. In this article, we introduced the Online Labour Index, an experimental economic indicator that tracks the utilisation of online labour. We described how the OLI is constructed and justified the methodological choices behind it, arguing that it offers a good indicator for observing trends

in the overall online gig economy. We also illustrated how the resulting data can be used to address crucial policy issues that existing data sources are unable to address properly and detect emerging trends. As the OLI continues to accumulate data, we expect that each of these issues will be addressed more thoroughly in dedicated studies by ourselves and other researchers using the data. Nevertheless, our initial look into these questions yielded the following insights and indications for follow-up work.

All in all, the index suggests that the volume of new vacancies has grown roughly 20% from the start of data collection. The growth is largely accounted by a single marked hike in the April-May of 2017. During this period, the index rose by almost 35% (from 105 to 140 index points) and has since declined.

The time series of new vacancies also shows some cyclical patterns with clearly observable dips around Christmas and New Year. We find likely that this pattern represents a holiday effect since most of the largest online labour buying countries have public holidays around Christmas and New Year. The country and occupation shares, on the other hand, have remained rather stable over the past 18 months.

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 by far, representing roughly a third of all vacancies posted. They are followed by creative and clerical work. This finding seems to support the recent findings from, e.g., (Blinder & Krueger, 2013) who claim that educated white-collar workers hold the most off-shoreable jobs. This finding stands in contrast to the literature on outsourcing and geographic divisions of labour, in which it is emphasised that complex and high-skilled expert work is harder to situate remotely than clerical work, since the former tends to involve interdependencies with other tasks and require frequent and intense client interaction (Howland 1993). It is possible that improvements in bandwidth and remote collaboration technologies such as video conferencing are now changing this accepted truth. If so, the core-periphery model may be giving way to more diverse and skill-based international divisions of labour.

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 the data 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. A caveat to this finding is that the index currently only tracks English-language platforms; though these are by far the biggest by traffic volume, the results are likely to somewhat over represent English-language countries. Another striking feature of the geography of online work is that the occupational demand profiles for all of the employer countries are quite 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 in accordance with the needs of the locally dominant industries.

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. The raw data used to produce the visualizations is also available. 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. We will also continue to monitor the coverage of the index, and introduce new platforms as needed.

Acknowledgements: We are grateful to the two anonymous referees and the editor Scott Cunningham for useful suggestions. In addition, Fabian Braesemann, Stijn Broecke, Greta Corporaal, Mark Graham, Andrew Karpie, Gina Neff, Nigel Swier and Alex Wood provided helpful comments and support during the project. The interactive visualization was developed in collaboration with Martin Hadley from the University of Oxford Interactive Data Network. We are thankful for their continuous assistance. This research is funded by the ERC (grant number 639652). Declarations of interest: None.

References

- Abraham, Katharine, John Haltiwanger, Kristin Sandusky, and James Spletzer. 2017. "Measuring the Gig Economy: Current Knowledge and Open Issues." In *Measuring and Accounting for Innovation in the 21st Century*. University of Chicago Press. <http://www.nber.org/chapters/c13887>.
- Agrawal, A., Horton, J., Lacetera, N., & Lyons, E. (2013). Digitization and the Contract Labor Market: A Research Agenda. *NBER Working Paper Series*, 19525. <http://doi.org/10.3386/w19525>
- Amato, F., Boselli, R., Cesarini, M., Mercorio, F., Mezzanzanica, M., Moscato, V., Persia, F., and Picariello, A. (2015). Challenge: Processing web texts for classifying job offers. In Proceedings of the 2015 IEEE 9th International Conference on Semantic Computing, IEEE ICSC 2015, pages 460–463.
- Aubert-Tarby, Clémence, Octavio R. Escobar, and Thierry Rayna. 2018. "The Impact of Technological Change on Employment: The Case of Press Digitisation." *Technological Forecasting and Social Change* 128.
- Blinder, A. S., & Krueger, A. B. (2013). Alternative Measures of Offshorability: A Survey Approach. *Journal of Labor Economics*, 31(S1), S97–S128. <http://doi.org/10.1086/669061>
- BLS (2015). Labor Force Statistics from the Current Population Survey. <http://www.bls.gov/cps/faq.htm - Ques1>.
- Carmel, E., & Agarwal, R. (2002). The Maturation of Offshore Sourcing of Information Technology Work The Maturation of Offshore Sourcing of Information Technology Work. *MIS Quarterly Executive*, 1(2), 1–19.
- Corrado, C. A. and Hulten, C. R. (2015). How Do You Measure a "Technological Revolution"? *American Economic Review*, 100(2):99–104.

- 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.
- Di Pietro, G. (2002). Technological change, labor markets, and “low-skill, low-technology traps.” *Technological Forecasting and Social Change*, 69(9), 885–895. [http://doi.org/10.1016/S0040-1625\(01\)00182-2](http://doi.org/10.1016/S0040-1625(01)00182-2)
- Einav, L. and Levin, J. (2014). Economics in the age of big data. *Science*, 346(6210).
- 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>.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <http://doi.org/10.1016/j.techfore.2016.08.019>
- Horton, J., Kerr, W. R., and Stanton, C. (2017). Digital labor markets and global talent flows. Working Paper 23398, National Bureau of Economic Research.
- Horton, J. J. (2017). The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment. *Journal of Labor Economics*, 35(2).
- Howland, M. (1993). Technological change and the spatial restructuring of data entry and processing services. *Technological Forecasting and Social Change*, 43(2), 185–196. [http://doi.org/10.1016/0040-1625\(93\)90016-Z](http://doi.org/10.1016/0040-1625(93)90016-Z)
- Husmanns, R. (2007). Measurement of employment, unemployment and underemployment: Current international standards and issues in their application.
- 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(2):16–21.
- Katz, L. F. and Krueger, A. B. (2016). The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015.
- 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. Washington, D.C. : World Bank Group. <http://documents.worldbank.org/curated/en/2015/06/24702763/global-opportunity-online-outsourcing>.
- Lacity, M. C., Solomon, S., Yan, A., and Willcocks, L. P. (2011). Business process outsourcing studies: a critical review and research directions. *Journal of Information Technology*, 26(4):221–258.
- 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.
- 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.
- Naccarato, A., Falorsi, S., Loriga, S., & Pierini, A. (2017). Combining official and Google Trends data to forecast the Italian youth unemployment rate. *Technological Forecasting and Social Change*. <http://doi.org/10.1016/j.techfore.2017.11.022>
- Pratt, J. H. (1984). Home teleworking: A study of its pioneers. *Technological Forecasting and Social Change*, 25(1), 1–14. [http://doi.org/10.1016/0040-1625\(84\)90076-3](http://doi.org/10.1016/0040-1625(84)90076-3)

- Rumberger, R. W., & Levin, H. M. (1985). Forecasting the impact of new technologies on the future job market. *Technological Forecasting and Social Change*, 27(4), 399–417.
[http://doi.org/10.1016/0040-1625\(85\)90020-4](http://doi.org/10.1016/0040-1625(85)90020-4)
- 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(1):267–288.
- Varian, H. R. (2014). Big data: New tricks for econometrics. *Journal of Economic Perspectives*, 28(2):3–28.
- 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.

Appendix A. Supplementary material

Column 3 of Table 2 lists how we infer the type of the opening for each of the platforms we observe. We denote the probability that observation i is in the occupation category

M ($M=1, \dots, 6$), conditional on the observed stemmed word counts x , as

$$\Pr(\text{Occupation}_i = M \mid x_i) = \frac{e^{\alpha_M + x_i' \beta_M}}{\sum_{k=1}^6 e^{\alpha_k + x_i' \beta_k}} \quad (1)$$

To estimate the parameter vectors $\beta = \{\alpha_M, \beta_M\}$ for all $M=1, \dots, 6$, we maximize the following objective function

$$\max_{\beta} \left[\frac{1}{N} \sum_{n=1}^N \sum_{k=1}^6 Y_{ik} \log P_{M_i}(\beta \mid x_i) - \lambda \sum_{k=1}^6 \|\beta_k\| \right], \quad (2)$$

where $P_{M_i}(\beta \mid x_i) = \Pr(\text{Occupation}_i = M \mid x_i)$ 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.

$\|\beta_k\|$ is the vector norm of β_k . The first element of the maximand is the standard log-likelihood function, and the second element is a penalty term. The estimation boils down to choosing a value of λ , and the corresponding vector β which minimizes a cross-validated mean deviance of (2).

Table A1 describes the various classification accuracy metrics calculated from the holdout sample. 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 (Garcia, 2009). All in all, both the balanced accuracy and the aggregate precision measures demonstrate that the regularized 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 71%. To put the accuracy of 71% into context, it should be compared to the completely random classification. If the occupations were randomly classified, the occupation shares shown in Column 1 of Table 3 suggest that random classification would result in an accuracy of roughly 32%.

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 projects is subject to misclassification due to prediction uncertainty. The classification based on the platform taxonomies has no random error involved, so the 85% of the projects that are classified based on the platform

taxonomies have an agreement rate of 100%. On the other hand, the 15% of the vacancies are classified using the LASSO classifier, which is estimated to classify 71% of the projects correctly. We end up with an agreement rate of 96% ($85\% + (71\% \times 15\%) \approx 96\%$).

Figure A1 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 extent 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 A1 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 $\frac{13}{13+2+11+2+1} \approx 43\%$ 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 misclassification is that a vacancy is erroneously classified as a *software development and technology*.




Figure A1: 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).

Table A1: Classification precision metrics. [95% confidence interval].

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

References

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.