

# Decoding the Online Labour Index (2016–2024): Methodology and Caveats

Version: May 25<sup>th</sup>, 2025

Marianne Furrer ✱, Otto Kässi ✧✉, Uma Rani 🌐

✱ Independent Researcher, Vevey, Switzerland

✧ Department of Computer Science, Aalto University, Espoo, Finland  
ETLA Economic Research, Helsinki, Finland

🌐 International Labour Office, Geneva, Switzerland

✉ otto.kassi@etla.fi

## Abstract

The Online Labour Index (OLI) was launched in 2016 to track the demand and supply of online freelance labour across multiple platforms. It collected real-time data on job postings and worker profiles, providing a valuable resource for researchers. However, data collection ceased in September 2024 due to the loss of access to a major platform, raising comparability concerns. This brief outlines the OLI's methodology, including demand- and supply-side data collection, and highlights key challenges related to data comparability, evolving platform structures, and the changing online labour market.

**JEL** J23, L86, C21

**Keywords:** online labour markets, gig economy, digital labour, data collection methods, labour market measurement

## Introduction

The Online Labour Observatory tracks the demand and supply of online freelance labour. On the demand side, it tracks the number of projects or tasks across platforms; and on the supply side, it collects data on workers, in almost real time across several platforms. The Online Labour Index (OLI), launched in 2016, was discontinued in September 2024. Data collection ceased due to the inability to collect data from one of the major platforms, raising concerns about data comparability. The OLI data is publicly available to researchers and institutions.<sup>1</sup> Although the complete OLI data collection process was not externally peer reviewed, internal review and testing starting in 2022 highlighted some strengths and weaknesses of the approach, leading to some modifications in the data collection and aggregation.

---

<sup>1</sup> The data and other resources can be found at: <http://onlinelabourobservatory.org/> and [https://figshare.com/articles/Online\\_Labour\\_Index\\_Measuring\\_the\\_Online\\_Gig\\_Economy\\_for\\_Policy\\_and\\_Research/3761562](https://figshare.com/articles/Online_Labour_Index_Measuring_the_Online_Gig_Economy_for_Policy_and_Research/3761562) [last accessed: 07.05.2025]

The OLI is built on two distinct data collection components. The demand-side data collection, launched in 2016, is documented in a peer-reviewed publication by Kässi and Lehdonvirta (2018), which details the methodology and scope of the data gathered. In addition, the supply-side data collection began in 2017. Unlike the demand-side data collection, the supply-side data underwent several structural changes and revisions, resulting in lower intertemporal comparability.

This brief complements existing methodological resources (Kässi and Lehdonvirta 2018; Stephany et al. 2021; Codebooks on Figshare<sup>2</sup>) by providing further detailed information on the data collection and aggregation, along with important caveats for researchers using the OLI data.

## **Demand-side data collection and aggregation**

Data collection on the demand-side began in 2016, focusing on five major English online labour platforms (Mturk, Guru, Freelancer, Peopleperhour, Upwork). These platforms were estimated to account for at least 60 per cent of all traffic to English-language online labour platforms (Kässi and Lehdonvirta 2018). As described by Kässi and Lehdonvirta (2018), the data collected included the projects or tasks posted (vacancies), the occupation classification (either observed or predicted), and – for some platforms – the employer country. To calculate the number of new vacancies, data for day  $t$  is compared to data on previous days (0, ...,  $t-1$ ). The OLI was constructed by aggregating the number of new vacancies across all five platforms and normalised so that the mean observed daily vacancies in May 2016 equals 100 index points.

In April 2020, six new platforms were added, and the Index was renamed “Online Labour Index 2020” (OLI 2020).<sup>3</sup> These additions included an English-language one (onlinejobs.ph), three Russian-language ones (freelance.ru and freelancehunt.com and weblancer.com), and two Spanish-language platforms (workana.es and twago.es<sup>4</sup>), expanding the index’s scope beyond the English-language online labour market.

Data were collected using data scrapers that crawled a given number of pages daily for each platform ( $N = 1, 2, \dots, n$ ), collecting information on all vacancies. The scrapers were set up in a way that the crawling stopped when either (i) an empty page was reached, or (ii) a page limit was met. For most platforms, the number of collected vacancies exceeded the number of new vacancies, indicating that the predefined

---

<sup>2</sup> See footnote 1.

<sup>3</sup> To ensure continuity between the OLI 5 and OLI 2020 time series, the method of chain linking data was used when adding the new platforms (see Stephany et al. 2021). Chain linking refers to the practice of “[j]oining together two indices that overlap in one period by rescaling one of them to make its value equal to that of the other in the same period, thus combining them into single time series” (OECD 2008: p. 71f).

<sup>4</sup> Twago.es stopped its operations in late October 2022, from November 2022 onward, the OLI 2020 demand side data include 10 platforms.

maximums allow for the collection of as many new vacancies as possible, which are usually posted on earlier pages and older vacancies are pushed back to the later pages.

The daily number of new vacancies varied significantly depending on the platform, from only a dozen of new vacancies to several hundreds or even thousands per day on others, reflecting the platforms' differences in market size.

Since its inception, the OLI 2020 has fluctuated between 95 and 172 index points (see Figure 1).<sup>5</sup> Some of the variation in the Index can be attributed to natural fluctuations in vacancy postings over time. For example, the number of new projects tends to drop every year around Christmas and New Year. The increase in the first half of 2020 is likely associated with firms adapting to the COVID-19 pandemic (Stephany et al. 2020).

However, internal review of the data collected and the data scraper history revealed that some variation in the Index and the skill distribution also stems from other factors, such as (i) changes in the data scrapers over time, (ii) website changes, or (iii) issues related to the data scrapers.

(i) **Changes in the data scrapers:** For example, in April 2017, for Upwork, the largest platform in the sample (see also Kässä and Lehdonvirta 2018), the data collection method was changed from crawling pages to using an API, which resulted in a significant increase in the number of daily new vacancies (see yellow bar in Figure 1).

**Figure 1:** Online Labour Index, 2016–2024



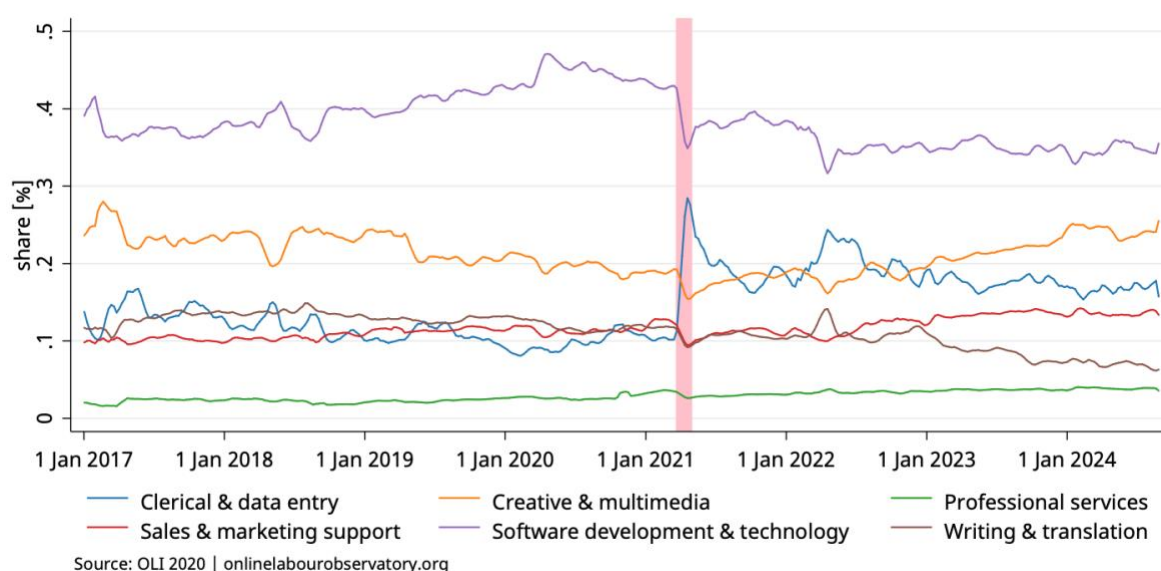
<sup>5</sup> These figures refer to the OLI 2020 data. The OLI 5 Index ranged between 95 and 226 points.

(ii) **Website changes:** Temporary spikes in the number of new projects, lasting several days or weeks, before returning to regular levels, have been observed on some platforms despite no apparent changes in the scrapers. This occurred on Peopleperhour in early April 2019 and from June to August 2019, and on Upwork from February to mid-April 2022 (see grey bars in Figure 1). Similarly, a late March 2021 spike in Mturk projects from an influx of relatively smaller requesters (see pink bar in Figure 1) resulted in a systematic increase in the OLI time series. This increase reflects an actual increase in job postings, not a data collection glitch or error.

(iii) **Issues related to the data scrapers:** Occasionally, the data scrapers encountered issues, such as pages being non-responsive, which can lead to fewer, or no data points being collected on certain days. This was relatively rare and addressed by imputing data from a previous point in time. For instance, when data collection failed for one of the major platforms for day  $t$ , data from  $t-7$  for the same platform was imputed.

Nevertheless, overall, the demand side data experienced relatively few disruptions. For instance, the most notable shift – the increase of administration and clerical vacancies in 2021 (see Figure 2) – seems to reflect a real change in the number of actual job postings on Mturk, rather than an artefact caused by the data collection.

**Figure 2:** OLI 2020 occupational breakdown, 2017–2024



## Supply-side data collection and aggregation

Data collection on the supply side began in 2017, covering fewer platforms than the demand side, because many platforms do not have a worker database that could be analysed. Initially, only Guru and Peopleperhour were included; Freelancer was added in 2019, and workana.es and freelance.ru were added in 2022.

While the concept of a new vacancy is relatively straightforward, defining a “worker” on online labour platforms presents several challenges. The total number of workers on any given platform is unknown, making it difficult to determine the representativeness of the collected data. Standard definitions of employment, such as at least one hour worked in a recent period - are difficult to apply, as platforms often have limited filtering options to identify active workers.<sup>6</sup> Consequently, the data may include a significant number of individuals who have never worked or have not been active for an extended period. Furthermore, many platforms allow workers to set their profiles to be invisible when they are not seeking new opportunities, regardless of their current work status. As a result, this could even exclude full-time platform workers from visibility if they are not available for additional jobs. These factors complicate even defining and developing methodologies to track workers in an appropriate way.

The supply-side data collection was initially conceived as a proof-of-concept to understand the usability of such data, and the technical implementation is provided in detail below. The supply-side data scrapers, like the demand-side counterparts, collected worker data daily for a specified number of pages for each platform ( $N = 1, 2, \dots, n$ ). The data included, where available, the workers’ self-declared country of residence and main occupation. The platform-level data was cleaned, workers were grouped by country and occupation, and then aggregated into a combined dataset. The aggregated data was weighted to reflect the share of workers from each platform, corresponding to the platform’s relative size based on vacancies (i.e. the OLI demand side).

As noted by Kässi (2020), this data cannot be used to infer changes in labour supply over time because the sample size (i.e. the number of pages from which information on workers is gathered) does not scale with the size of the workforce, which remains unknown, as discussed earlier. Therefore, the Online Labour Observatory website presents only occupational and regional breakdowns on the supply side, and not an Online Labour Supply Index.

An internal review and testing process, initiated in 2022 has revealed some potential issues with the data collected and their comparability over time. These issues can be grouped into three broad categories: scraper settings, data collection, and data aggregation.

**Scraper settings:** As noted earlier, for each platform, the scrapers are programmed to collect data from  $N$  number of pages. Important scraper settings include the target URL(s) from which data is collected, and the time-of-day data is collected.

For most platforms, some of these settings were modified multiple times over the course of the data collection period, impacting the number of workers collected from

---

<sup>6</sup> Kässi et al. (2021) estimated that over 90% of the registered worker profiles across 339 distinct online labour platforms had no work experience, indicating that they did not manage to get any jobs.

each platform, and the occupational and regional distribution of workers. For example, changes in the Guru and Peopleperhour scrapers in mid-August 2019, reduced the number of workers collected for these platforms by two-thirds and one-half, respectively, which influenced the occupational and regional breakdown (see pink vertical lines in Figures 3 and 4). This change coincided with the addition of Freelancer to the sample, which is discussed further below.

Some platforms list workers irrespective of their occupation or skill category, while others use occupation or skill-specific URLs. In the latter case, when a scraper is set to collect data for the same number (N pages) for the different occupations or skill categories, the occupational or skill shares will stay artificially fixed over time.<sup>7</sup> Three notable changes were made related to skill categorizations:

- i) When data collection began on Peopleperhour, workers offering services in the field of “business support” were classified based on their hourly rate into the OLI categories “clerical and data entry” or “professional services”. From December 2020, Freelancer no longer required the workers to set their hourly rate, resulting in many not available (NA) values in the hourly rate field. Consequently, the collection of the hourly rate field was discontinued, and all workers in “business support” were classified into “clerical and data entry”.
- ii) In late August 2022, the URLs for the scraper on Peopleperhour were changed to prevent double-counting workers with multiple listed services. Prior to the change, the scraper was collecting “services” listed on the platform by occupation category, and, as a worker could list multiple services, there was double counting in the dataset. After August 2022, the scraper was set to a URL where workers were listed, irrespective of the type of services they provided, ensuring that each worker was counted only once.
- iii) In early January 2023, the scraper on Guru was changed from gathering data from skill-specific URLs to a URL where workers were listed irrespective of their skills.

All three changes had an impact on occupational distributions. On Peopleperhour, the share of workers in “creative and multimedia” and “professional services” increased in August 2022, while “sales and marketing support” workers decreased. On Guru, the share of workers in “software development and technology” increased significantly,


---

<sup>7</sup> Platform occupational categorizations may differ from those presented in the OLI. For example, a platform could use different URLs for workers providing services in the fields of “design”, “music/audio”, and “video/photo/image”. In the OLI, these are all grouped under “creative and multimedia”. Conversely, the platform might have one single URL for workers providing “writing/translation” services, corresponding to the OLI category of “writing and translation”. If the same number of pages is scraped for each URL, the share of creative and multimedia workers would be three times that of those in writing and translation.

whereas that of workers in “professional services”, “sales and marketing support”, and “clerical and data entry” decreased.

Finally, regarding the time-of-day data is collected, internal testing in early 2022 showed that the regional composition of the workforce changed depending on the time of day at which the scraper is run. The tests determined that there was no need to make any adjustments regarding the time-of-day as the variation was negligible.

**Figure 3:** OLI supply side occupational breakdown, 2018–2024




**Data collection:** The scrapers were set up to run each day, but occasionally there were some days or time periods where no data were collected for one or several platforms. This was due to technical issues with the platforms, such as errors while loading pages; changes in the website structure due to which the scrapers could not collect data; or security measures put in place by platforms, such as captchas, that prevented the scraper from functioning without human intervention.

Missing data for a day or two on one or several platforms would be difficult to notice, as during data aggregation 28-day averages were used to smoothen out outliers. However, if data was missing for a somewhat longer period, this would be visible in the occupational and regional breakdown. For example, Peopleperhour data was missing from mid-October to mid-November 2017, as well as for December 2018 and January 2019 (see yellow bars in Figures 3 and 4).

In addition to completely missing data, there were also instances where partial data gaps occurred. For example, on Guru between late June 2021 and early April 2022, the scrapers failed to collect data from URLs in the fields of “software development and technology”, “sales and marketing support”, and “engineering and architecture” (part of “professional services”), affecting the occupational breakdown.

Since January 2023, the platform-level data is imputed if the number of workers falls outside a specified threshold, thereby systematically correcting for missing data.

**Figure 4:** OLI supply side country breakdown, selected top countries, 2017–2024



**Data aggregation:** Once the data for each platform has been collected, it is aggregated into a dataset containing data from all platforms. As mentioned earlier, the worker data collection began in 2017 with Guru and Peopleperhour; Freelancer was added in 2019 and workana.es and freelance.ru in August 2022. The addition of platforms to the aggregated dataset impacted the occupational and regional compositions. Freelancer’s inclusion in early August 2019 significantly impacted both distributions, while freelance.ru’s addition mainly affected the regional distribution (see green bar areas in Figures 3 and 4).

When the data is aggregated, weights are used to reflect platform size based on demand-side data collection (i.e. the OLI demand side). These weights were modified several times during the years of data collection, but not always consistently with observed changes in the market share of the platforms or the number of workers collected from each platform. This resulted in fluctuating platform representation within the aggregated data.

For example, Freelancer’s weight was halved in late November 2019 despite the number of workers being collected remained same. Conversely, changes in the Freelancer data scraper temporarily reduced the number of workers collected by about 30 per cent between July and December 2021, but the aggregation weight remained unchanged. In both cases, while visible from the figures, above, the impact of adjusting weights was relatively small. According to our calculations, regional and occupation shares changed by maximum less than 5 percentage points when weights are adjusted.



In August 2022, weights were re-adjusted to reflect the platform size (measured in terms of new projects, see table 1), and data collected from the platforms workana.es and freelance.ru were integrated into the aggregated data. However, due to human error the weight for workana.es was lower than it should have been; and there was an overrepresentation of workers from Guru (factor 5) between August and December 2022.

**Table 1:** Supply side target shares, as of January 2023

Platform	Target share
Freelancer	0.82
Freelance.ru	0.09
Guru	0.04
Peopleperhour	0.05
Workana.es	0.04

Source: OLI | [onlinelabourobservatory.org](https://onlinelabourobservatory.org)

Finally, when aggregating the data, the 28-day smoothing that is applied mitigates potential outliers and flattens out the data. However, it also had the effect to obscure sudden changes in the number of observations and the occupational or regional composition of the data, that might otherwise have raised concerns more easily among users of the data.

## Conclusion

This brief provides information on the OLI 2020's data collection and aggregation process, highlighting issues related to both demand- and supply-side data. The demand-side data are relatively stable, with consistent data collection since at least 2021. This facilitates reasonably good temporal, geographical, and occupational comparisons. However, there are some instances where the fluctuations in the index can be attributed to changes in the data scrapers, changes on the platform websites, or issues related to the data scrapers.

In contrast, the supply-side data remains at a proof-of-concept stage due to platform changes, data collection challenges, and varying data sources. Consequently, it suffers from poor intertemporal comparability. Observed shifts in occupational and geographical distributions likely reflect changes in data collection configurations, rather than genuine underlying trends. Users should exercise caution when interpreting the supply-side data due to these limitations.

Finally, while the platforms initially included in the dataset covered over two-thirds of the online labour platform market in 2016, the market has significantly evolved since then, with the emergence of several new players such as Fiverr, which the OLI has never tracked. Perhaps more significantly, the demand for generative AI model training workforces has surged since the onset of the OLO data collection began (see,

e.g., Tubaro et al. 2020). Although much of this work likely occurs through online platforms, these platforms often do not publicly disclose open vacancies or worker profiles, rendering this segment of platform labour entirely invisible to the OLI.

## Bibliography

Kässi, O. and Lehdonvirta, V., (2018). Online Labour Index: Measuring the Online Gig Economy for Policy and Research. *Technological Forecasting and Social Change*, 137: 241-248.

Kässi, O. (2020). Online labour Index worker data codebook, last updated 2020-11-23. Available at [https://figshare.com/articles/dataset/Online\\_Labour\\_Index\\_Measuring\\_the\\_Online\\_Gig\\_Economy\\_for\\_Policy\\_and\\_Research/3761562?file=25560452](https://figshare.com/articles/dataset/Online_Labour_Index_Measuring_the_Online_Gig_Economy_for_Policy_and_Research/3761562?file=25560452) [12.02.2025]

Kässi O., Lehdonvirta V., and Stephany F. (2021). How many online workers are there in the world? A data-driven assessment. *Open Research Europe* 2021, 1:53. <https://doi.org/10.12688/openreseurope.13639.3>

OECD (2008). OECD Glossary of Statistical terms. Paris: OECD. <https://doi.org/10.1787/9789264055087-en>

Stephany, F., Dunn, M., Sawyer, S. and Lehdonvirta, V. (2020), Distancing Bonus Or Downscaling Loss? The Changing Livelihood of Us Online Workers in Times of COVID-19. *Tijdschrift voor economische en Sociale Geografie*, 111: 561-573. <https://doi.org/10.1111/tesg.12455>

Stephany, F., Kässi, O., Rani, U., and Lehdonvirta, V. (2021). Online Labour Index 2020: New ways to measure the world's remote freelancing market. *Big Data & Society*. <https://doi.org/10.1177/20539517211043240>

Tubaro, P., Casilli, A. A., and Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1), 2053951720919776.