Christopher T. Stanton and *Catherine Thomas*

The Gig Economy Beyond Local Services and Transportation


Christopher T. Stanton

is Marvin Bower Associate Professor in the Entrepreneurial Management Unit at Harvard Business School.


Catherine Thomas

is an Associate Professor of Managerial Economics and Strategy at the London School of Economics.

The gig economy characterizes a wide variety of short-term freelance work, typically intermediated via online platforms that facilitate matching between buyers and providers. The widespread growth of ride-sharing platforms such as Uber and Lyft has led many to equate gig-economy work with tasks carried out face to face after matching on a platform. However, many gigs or tasks can be both contracted and performed remotely, particularly when the output can be delivered electronically. Platforms that enable this type of work are referred to as online labor markets.

The hiring process in these settings differs from those in traditional offline labor markets in that rapid matching happens for relatively short assignments. Most platforms have a core set of features, such as reputation systems, portfolios of past work, a partially standardized skills canon, payment handling, and provisions to prevent taking work off the platform. The precise contract form is often at the discretion of a buyer, ranging from those paying an hourly wage to those offering fixed fees negotiated for a specified output.

Blinder and Krueger (2013) estimate the extent to which occupations in the United States are amenable to online production, or, in their terminology "offshorable." They conclude that a reasonable estimate of the offshorable share of US employment given the technology available at the time was around 24%. This covers activities that required both skilled and unskilled labor. In information and professional, or scientific and professional, services, the offshorable share was even higher, at around 35%; office and administrative support occupations came in at 41%. Earlier work on services occupations by Jensen and Kletzer (2010) estimated that 93% of computer and mathematical and 64% of office and administrative support occupations could be offshored.

When it is technically feasible, the potential labor costs savings of remote work are substantial. The pay comparison website payscale.com reported that the annual salary in 2018 for a 25-year-old Software Developer with a Bachelor's degree and three years of experience was USD 112,000 in San Francisco, USD 24,000 in Warsaw, and USD 8,000 in Dhaka. A growing literature on market power in the labor market suggests local opportunities, rather than productivity differences alone, contribute to this wage gap (Ashenfelter et al. 2010; Caldwell and Danieli 2018). If wage differences of this order of magnitude can be realized at the level of tasks done online, then the variable cost savings would make it a very appealing option.

However, while several individual platforms have matured into liquid marketplaces (Kässi and Lehdonvirta 2018), aggregate adoption rates for online work remain low. This becomes apparent when considering the public earnings reports from the leading online platforms. At the time of writing, the combined revenues of publicly traded online labor platforms was only a fraction of the traditional staffing firm Manpower's annual USD 20 billion in revenue. Macroeconomic statistics yield similar conclusions. The US Census' Characteristics of Businesses Survey found that only 1.5% of all firms outsourced or transferred any business function and/or service to a company outside the US in 2015. The industry with the largest number of firms reporting this activity was Information (NAICS code 51, producing and distributing information and cultural products; providing the means to transmit or distribute these products as well as data or communications, and processing data), at 6.9%, and the US State with the largest reporting share was California at 2.6%.

Turning to the supply side, in the US over recent decades, there has been a decline in the share of self-employment among more highly educated workers relative to less-well educated workers. This suggests that those at the higher end of the labor market are not as likely to engage in gig work—which is, by definition, self-employment—at least as their primary form of employment. This is also likely consistent with analysis from tax records on the rise of the platform economy, which shows that the majority of new gig-like work arrangements tend to be coming from ride sharing or co-located services (Collins et al. 2019).

REASONS FOR THE LIMITED GROWTH OF ONLINE FREELANCE WORK

Considering the aforementioned trends together, the limited growth of online freelance work in spite of its technical feasibility and low cost is consistent with an existing barrier in the form of coordination costs or contracting frictions that exist at firm boundaries. This paper explores the composition of tasks that are conducted via online labor markets and offers some comments on why adoption is not yet as widespread as in more standardized gig economy sectors such as ride sharing.

Because of the differing trends in gig work by skills required and education level, it is interesting to first ask what types of skills are demanded online and by whom. Data show that the majority of potential buyers are in English-speaking countries and are looking for temporary freelance workers to complete discrete tasks. On most platforms, buyers must include a task description in the vacancy posting that goes into quite some detail about the skills required. Various sources show that most tasks are either of a technical nature (e.g., web or software development), or require specific skills such as design or translation. Nonetheless, vacancies in administrative support and data entry are also common.

Historical data from a leading platform reveals whether it is possible to predict which potential buyers will end up becoming frequent users of the platform from their observable characteristics. Data from 2008 to 2010 that contains information on over 60,000 buyers shows some evidence that buyers of technical services are relatively less likely to adopt the platform after trying it out. Buyers in larger enterprises are also less likely to adopt the platform than sole proprietors or smaller enterprises. Other than these factors, observable buyer characteristics or attributes of the vacancy posting have little explanatory power for platform adoption.

Given these two factors, it seems likely that buyer fit with online labor relates to the willingness or ability to carve out well-defined tasks that can be done by a specific individual and then integrated with other production activities. Complex technical jobs or those requiring integration into larger production processes may make the up-front investment in writing specifications and the onboarding and monitoring of arms-length contractors difficult to justify relative to a local, known alternative.

Much of the other work on these platforms has studied providers' careers or earnings (Horton 2010; Pallais 2014; Stanton and Thomas 2016), and the impact of the gig economy for the labor force as a whole – see Koustas (2020) in this volume, and Datta et al. (2018). The data on self-employment trends in the US show that more highly educated individuals, who tend to be in occupations that require interaction with other activities, are tending to remain within firm boundaries. Hence, it is plausible that the challenges associated with communication across tasks have proved to be a barrier to more rapid growth of task-based online work.

The Oxford Internet Institute's Online Labour Index[1] documents a 50% increase in job postings between May 2016 and March 2020 on the five largest online platforms. This paper concludes by providing results that suggest how the growth rates of online platforms may be impacted by the Covid-19 pandemic that hit the English-speaking world from March 2020 onward. Initial evidence shows that demand declined steeply until the first week of April, but rose to unprecedented heights up to the end of May. The increase was particularly steep in software development and technology tasks and from buyers located in the US. It is likely that many new buyers turned to these platforms for the first time, experimented with their use and learned how to hire and how to coordinate remote work. In this vein, recent work shows that gaining experience with this form of labor sourcing can help buyers to become long-term adopters (Stanton and Thomas 2020). This finding is analogous to recent literature suggesting that the pandemic revealed that, at least for some firms, workers were more productive working from home or in new arrangements (Bartik et al. 2020).

As countries ease out of lockdown, it will be interesting to observe whether the increased experimentation with online labor that has occurred during the last few months is sufficient to convince buyers of its overall appeal.

WHO BUYS ONLINE LABOR SERVICES?

The Online Labor Index tracks activity across the five largest English-language online labor platforms, which represent over 70% of the total market. The index measures supply and demand across countries and job types by tracking the number of projects posted across the different platforms in real time. According to these data, in the first week of July 2020, 38% of the value transacted online originated in the US, 9% in the UK, and 6% each in Canada and Australia. As the largest source of service providers globally, demand from India also made up 8% of this value. In the same week, 46% of the value was in software development and technology tasks, and 20% was in creative and multimedia tasks.

At the start of this index, in May 2016, 53% of vacancies were posted by buyers in the United States, and 35% were in software development and technology. Going back even earlier, in microdata from a single large platform dating from 2008 to 2010, 57%, or just over 67,000 buyers, were located in the US. 49% of all vacancy postings were in technology jobs, either web development or software development. At this time, the other main types of vacancies posted were tasks in Administrative Support (15%), Sales and Marketing (11%), Writing and Translation (10%), and Design and Multimedia (10%).

Around half of all vacancy postings were for tasks that lasted less than one month. This suggests potential buyers had defined discrete objectives to be delivered, and were not seeking to contract for ongoing deliverables. That is, the work content tended to be task- or gig-based. Technical tasks and tasks in design and multimedia were more likely to be short term,

[1] See http://ilabour.oii.ox.ac.uk/online-labour-index/.

consistent with the idea that the deliverable output is a discrete piece of work. Customer Service and Administrative Support vacancies were more likely to be longer term, consistent with these tasks being recurring for most businesses, rather than project-based.

WHEN DOES ONLINE PRODUCTION REALLY WORK?

Data from the large platform also reveals which potential buyers ended up being intensive users of this technology, and for what types of tasks. Amongst all buyers who posted at least one vacancy on the platform, outcomes were very heterogeneous. 11% hired five or more times over the next two years, while 73% didn't hire at all. For the rest of this paper, this 11% of buyers is referred to as the "adopters."

The data permits tracking all the buyers who tried out the platform during this time and, hence, at what stage the non-adopters opt out. In seeking to explain why some buyers adopt the technology and some do not, it is important to understand whether they differ from non-adopters in observable characteristics or in their early actions on the platform. Figure 1 compares adopters with all buyers. Adopters are slightly more likely to be located in the US, at 62% vs 57%. This may suggest that US buyers find it easier to contract at arm's length for the type of discrete tasks that are suited to these platforms. They may also have more flexibility in their labor-sourcing decisions or in regulations that enable this type of contracting.

Figure 1 also shows that the buyers who end up adopting the platform are less likely to have posted their first vacancy in a technical job category. Relative to expectations when deciding to try out the platform, it could be that these types of tasks require more communication and are harder to coordinate at arm's length. In contrast, there is no real difference between all buyers and those who adopt the platform as part of their first job posting, as measured by whether the first vacancy is likely to last over one month.

The final two pairs of bars in Figure 1 show that adopters differ from typical buyers in that they are more engaged with the platform right from the start. From the outset with their first posting, they are more likely to conduct applicant interviews, at 82% versus 63%, and they are much more likely to fill their first vacancy posting on the platform. In fact, 32% of adopters hire on the first posting as compared to 15% overall.

The figure therefore shows that early engagement is correlated with becoming an adopter. Adopters also conduct a larger number of interviews on the first posting and are more likely to hire when controlling for the week of first posting, the job category, and the country where the buyer is located. Even a buyer who does not hire on the first vacancy is more likely to become an eventual adopter if they conducted interviews for the first vacancy. Of course, these summary figures are consistent with either more engaged buyers doing more to figure out how the platform works and how to manage workers—actions that increase the utility from contracting with online labor—or due to unobserved heterogeneity across buyers. We now turn to the question of whether this selection either in or out of early engagement and eventual platform adoption is based on some unobservable buyer characteristic or is instead due to the early actions taken on the platform.

WHY DOES EARLY ENGAGEMENT VARY?

Ideally, buyer heterogeneity versus engagement with the platform could be disentangled using a well-designed experiment. In the absence of this kind of evidence, models of buyer engagement that account for the composition of the applicant pool, the prices offered, and details about the vacancy can be estimated. Accounting for the buyer and vacancy characteristics that are observed by potential applicants, seemingly random variation in the applicants seen early on may shape buyers' eventual adoption by affecting their initial experience. Furthermore, if some buyers adopt the platform despite receiving lower-quality applications and others opt out despite receiving high-quality applications, it is likely that heterogeneous buyer types are trying out the platform and eventual adoption is unrelated to their early experiences. Stanton and Thomas (2020) conduct this analysis, and the results put weight on both these explanations. That is, buyers' early actions play an important role in accelerating or hindering platform adoption even after attempting to account for differences in buyer types.

To illustrate how this procedure distinguishes between various explanations, imagine there are two types of buyers: one type places a high value on online hiring when trying it out, whereas the other type is skeptical. Because of their initial enthusiasm, type-one buyers are likely to hire even when the set of applicants is rather lousy. On the other hand, the

Figure 1
All Buyers vs Adopters

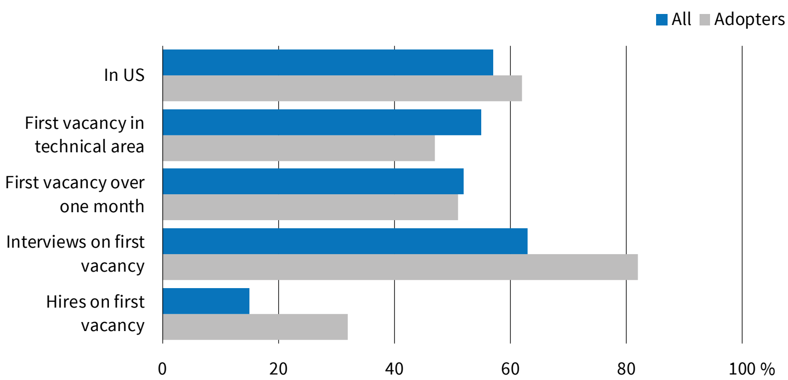

Source: Proprietary Data from a Large Online Labor Market.

© ifo Institute

type-two buyer will only hire when the applicant pool includes star providers. The distribution of buyer types can be inferred from buyer responses to receiving varying applicant pools. Still, within each type, some fraction of buyers will hire and some will not. The analysis shows that even when conditioning according to type, buyers that hire early on post more later vacancies and are more likely to become adopters. That is, early actions to engage with the market by hiring lead to long-term future use.

WHAT ELSE CHANGES WITH EXPERIENCES?

Having overcome any barriers to first using the platform, buyers appear to learn how the platform works as they hire labor services. The data contain two variables about the intensity of search to fill each vacancy: the length of the task description, in characters, as well as the number of interviews.

Regressions that include buyer fixed effects show that the length of the description of the work in the vacancy posting tends to shorten on successive posts. Figure 2 plots the average number of characters in the description, controlling for the week, the type of task, and the expected duration of the vacancy, along with buyer fixed effects. The grey bars are for those buyers who go on to become adopters, and the blue bars show non-adopters who are posting vacancies in the market after having made a number of hires, but will drop out before making five hires. Because buyer fixed effects are included, the average number for both groups is the same prior to doing any hires, and, because adopters are defined as those who make at least five hires, there are no non-adopters in the data at this experience level. The grey bars show significant declines after making one hire, with adopters posting descriptions that are 14% shorter after making five hires.

This reveals either that the buyers that become adopters learn to communicate their needs more effectively or learn that providing detail about the task does not improve matching applicants to the tasks.

Figure 3 plots the average number of interviews conducted at different levels of hiring experience, controlling for week, job category, and buyer fixed effects. Again, the grey bars are for those buyers who go on to be adopters and the blue bars represent all other buyers in the data after making the relevant number of hires. Both series show a reduction in the number of interviews on successive posts, and, while this decrease is even greater for the non-adopters, adopters decrease the number of interviews by 36% by the fifth hire.

This finding could mean that all buyers learn how to select applicants for interviews or how to differentiate between applicants more easily during interviews. It could also mean that buyers learn that interviews are not particularly helpful in differentiating among applicants. In either case, Figures 2 and 3 suggest the vacancy posting and search process is less time consuming once buyers have learned how the platform works.

Gains from experience hence appear to include reducing the cost of using the platform, but many buyers drop out before reaping these benefits. A key question for the future growth of these technologies, then, is whether it is possible to lower the costs of engaging with the platform when first trying it out. Recent work by John Horton addresses this question by studying interesting market design interventions that seek to lower these costs (Horton 2017 and 2019), and which are shown to be effective in doing so on the margin.

The fact that buyers opt out of the market after posting vacancies suggests they were initially uncertain about the extent of engagement costs. One way to overcome this uncertainty about the platform would be for buyers to run small tests to determine whether the market works for them. In fact, in most models of experimentation, one would expect buyers to post short tasks to evaluate whether the platform meets their needs, after which they would scale up hiring. The microdata, however, offer no evidence that buyers operate this way. For example, the propensity to post a long versus short vacancy varies little over the buyer life cycle. This suggests that one barrier to

Figure 2

Length of Job Description, Controlling for Expected Duration and Buyer Fixed Effects

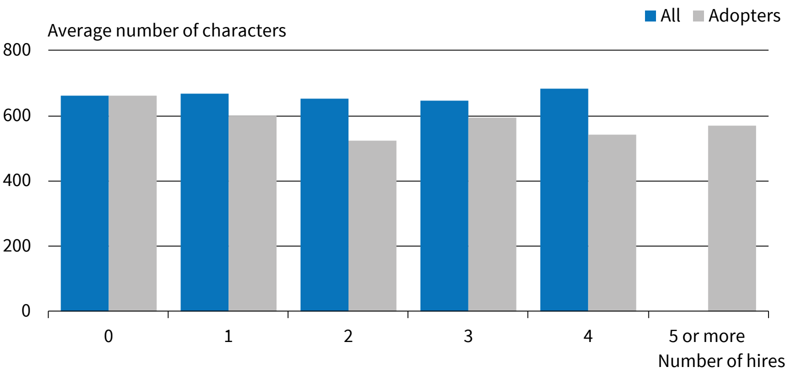

Source: Proprietary Data from a Large Online Labor Market. © ifo Institute

Figure 3

Number of Interviews by Number of Previous Hires, Controlling for Buyer Fixed Effects

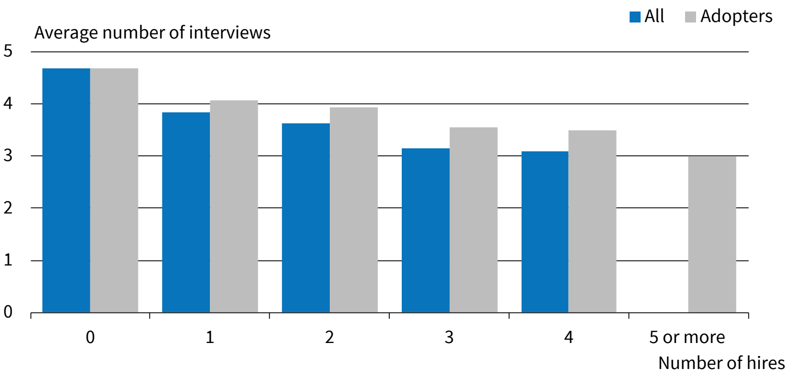

Source: Proprietary Data from a Large Online Labor Market. © ifo Institute adoption may be the difficulty in fragmenting tasks into their sub-parts, because, to conduct a short experiment on the platform, a buyer must first design a short and self-contained task. Instead, buyers appear to jump in headfirst with whatever task needs to be done. There is also evidence that buyers default to hiring providers who are more likely to be familiar to them (Ghani et al. 2014). When the first vacancy posting and first hire go well, and buyers find that the platform is a good fit, they post more vacancies on the platform and look for providers to do tasks that vary in work content but not in length.

Among the possibilities suggested by the data, one likely explanation for why some buyers opt out of platform use is that posting the vacancy and viewing applicants shows them that coordinating online work is relatively costly and they decide to go no further. Unlike the act of summoning an Uber, online labor markets require a buyer to be actively involved in project management, and this management is likely to be costly.

WHAT'S NEXT FOR ONLINE WORK?

Turning back to data on self-employment in the US, among the educated (those with a BA or higher degree), long differences in employment trends show that the highly educated—those whose skills are well-suited to online work, particularly to technical tasks—are much less likely to be self-employed than they were historically. This is displayed in Figure 4, which uses data from the CPS ASEC surveys for males over time. The left panel looks at changes in self-employment within the education cell and shows large drops for those with Bachelor's or higher degrees. The focus here is on males to avoid confusion with secular increases in female labor supply. Of course, educational attainment increased over this period, so the right-most figure considers changes to the population share of self-employment. Rising educational attainment partially offsets the decline in per-capita self-employment for those who hold at least a Master's degree. The direct implication of these trends is that more economic activity among skilled individuals is occurring within firms rather than via self-employment.

It is also notable that the population-level decline in self-employment ends around 2011 for those with a high-school degree or a lower level of education. The rate of self-employment even begins to turn upward in 2015, coincident with the rise of the driving economy (Collins et al. 2019). No such uptick exists for the highly educated.

Although there are a number of factors at play, differences in the task content of jobs may explain some of these differences. According to O-NET data, educated workers are less likely to be in routine jobs and are more likely to be in jobs that require math skills, social skills and interaction with others (Deming 2017). For these types of occupations, the need to combine technical skills and coordination on the job likely increases the overhead of managing contracts beyond firm boundaries.

Tying this back to online labor markets, why then are technical tasks such a large fraction of vacancy postings if these tasks require interaction and coordination? One possibility is that automation is putting pressure on more routine activities in administrative support, allowing employers to turn away from labor markets altogether and toward computing power for their needs. Another possibility is that tasks come with a fixed cost in terms of management time, and higher-value technical tasks may yield larger gains relative to fixed costs than doing routine tasks.

What do these findings mean for the future of remote work? The first half of 2020 has seen perhaps the largest shock to its prevalence, following national lockdowns in countries worldwide that meant that work should be done at home whenever feasible, even within firm boundaries. In a recent working paper, Stephany et al. (2020) document some very interesting patterns in the demand for arm's-length US online workers from buyers located in countries that went into lockdown. In the early days of a country's lockdown, demand for freelance online services fell, but, as the local lockdown continued, demand for online labor increased and soon overtook initial levels. For example, by the start of April, Korea had been in lockdown for several months and demand for online labor had risen to levels above those seen prior to the crisis. Germany saw a later upturn, and the US upturn was even later, as the crisis played out at different times in these countries.

It may well be that once a large share of the workforce is working from home, confronting all of the challenges of coordination and communication across diverse locations, undertaking such activities at arm's length, across firm boundaries, will also start to appear less daunting. If potential buyers are able to capitalize on what they have learned over the last few months, then demand for all types of remote online work may continue to grow, even when local economies return to more typical working conditions.

Figure 4

Male Self-employment by Education

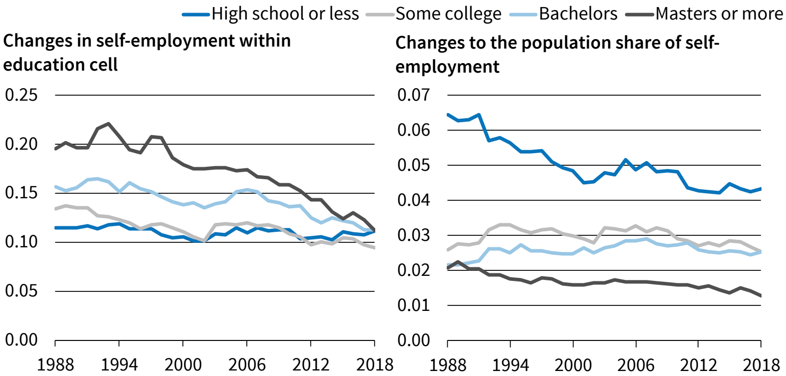

Source: CPS ASEC Universe for Ages 18–65 and Years 1988–2018.

© ifo Institute

REFERENCES

Ashenfelter, O. C., H. Farber and M. R. Ransom (2010), "Labor Market Monopsony", *Journal of Labor Economics* 28, 203–210.

Bartik, A. W., M. Bertrand, Z. B. Cullen, E. L. Glaeser, M. Luca and C. T. Stanton (2020), "How Are Small Businesses Adjusting to COVID-19? Early Evidence from a Survey", *NBER Working Paper* 26989.

Blinder, A. S. and A. B. Krueger (2013), "Alternative Measures of Offshorability: A Survey Approach", *Journal of Labor Economics* 31, S97–S128.

Caldwell, S. and O. Danieli (2018), *Outside Options in the Labor Market,* mimeo.

Collins, B., A. Garin, E. Jackson, D. Koustas and M. Payne (2019), *Is Gig Work Replacing Traditional Employment? Evidence from Two Decades of Tax Returns,* IRS SOI Joint Statistical Research Program, mimeo.

Datta, N., G. Giupponi and S. Machin (2019), "Zero Hours Contracts and Labour Market Policy", *Economic Policy* 34, 369–427.

Deming, D. J. (2017), "The Growing Importance of Social Skills in the Labor Market", *The Quarterly Journal of Economics* 132, 1593–1640.

Ghani, E., W. R. Kerr and C. Stanton (2014), "Diasporas and Outsourcing: Evidence from oDesk and India", *Management Science* 60, 1677–1697.

Horton, J. J. (2010), "Online Labor Markets", in A. Saberi (ed.), *Internet and Network Economics,* Springer, Berlin and Heidelberg, 515–522, https://doi.org/10.1007/978-3-642-17572-5_45.

Horton, J. J. (2017), "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment", *Journal of Labor Economics* 35, 345–385.

Horton, J. J. (2019), "Buyer Uncertainty about Seller Capacity: Causes, Consequences, and a Partial Solution", *Management Science* 65, 3518–3540.

Jensen, J. B. and L. G. Kletzer (2010), "Measuring Tradable Services and the Task Content of Offshorable Services Jobs", in K. G. Abraham, J. R. Spletzer and M. J. Harper (eds.), *Labor in the New Economy,* University of Chicago Press, Chicago, 309–335.

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

Koustas, D. (2020), "Insights from New Tax-Based Measures of Gig Work in the United States", *CESifo Forum* 21(3), forthcoming.

Pallais, A. (2014), "Inefficient Hiring in Entry-Level Labor Markets", *American Economic Review* 104, 3565–3599.

Stanton, C. T. and C. Thomas (2016), "Landing the First Job: The Value of Intermediaries in Online Hiring", *The Review of Economic Studies* 83, 810–854.

Stanton, C. T. and C. Thomas (2020), "The Gains from Trade in Online Labor Markets", *Working Paper.*

Stephany, F., M. Dunn, S. Sawyer and V. Lehdonvirta (2020), "Distancing Bonus or Downscaling Loss? The Changing Livelihood of US Online Workers in Times of COVID-19", *Journal of Economic and Social Geography*, https://doi.org/10.1111/tesg.12455.

The US Census (2015), The US Census' *Characteristics of Businesses Survey 2015,* https://www.census.gov/data/tables/2015/econ/ase/2015-ase-characteristics-of-businesses.html.