

C A G E

The New Wave? The Role of Human Capital and STEM Skills in Technology Adoption in the UK

CAGE working paper no. 726

September 2024

Mirko Draca,
Max Nathan,
Viet Nguyen-Tien,
Juliana Oliveira-Cunha,
Anna Rosso,
Anna Valero



Economic
and Social
Research Council

The New Wave? The Role of Human Capital and STEM Skills in Technology Adoption in the UK *

Mirko Draca^{1,4}

Max Nathan^{1,4,5*}

Viet Nguyen-Tien^{1,2}

Juliana Oliveira-Cunha^{1,3}

Anna Rosso^{1,6}

Anna Valero^{1,2}

¹CEP, London School of Economics

²POID, London School of Economics

³IGC, London School of Economics

⁴Warwick University and CAGE Research Centre

⁵University College London

⁶University of Insubria

*Corresponding author: max.nathan@ucl.ac.uk

Abstract

Which types of human capital influence the adoption of advanced technologies? We study the skill-biased adoption of information and communication technologies (ICT) across two waves in the UK. Specifically, we compare the ‘new wave’ of cloud and machine learning / AI technologies during the 2010s - pre-LLM - with the previous wave of personal computer adoption in the 1990s and early 2000s. At the area-level we see the emergence of a distinct STEM-biased adoption effect for the second wave of cloud and machine learning / AI technologies (ML/AI), alongside a general skill-biased effect. A one-standard deviation increase in the baseline share of STEM workers in areas is associated with around 0.3 of a standard deviation higher adoption of cloud and ML/AI. We find similar effects at the firm level where we are able to test for the influence of a wide range of skills. In turn, this STEM-biased adoption pattern has encouraged the concentration of these technologies, leading to more acute differences between high-tech and low-tech areas and firms. In contrast with classical technology diffusion, recent cloud and ML/AI adoption in the UK seems more likely to widen inequalities than reduce them.

Keywords: Technology Diffusion, ICT, Human Capital, STEM

JEL codes: D22, J24, O33, R11

*Many thanks to seminar participants at AIEL, University of Aston, Banque De France/OECD LEED, CASA, CReAM, ESCoE, OECD and POID-LSE for helpful comments. We thank the ESRC for funding under grant ES-T002506/1, the CAGE Research Centre (grant ES/L011719/1), the Centre for Economic Performance (ES/T014431/1), and the Programme on Innovation and Diffusion (ES/V009478/1). The usual disclaimer applies.

1 Introduction

Many advanced economies have experienced a productivity slowdown since the 2008 financial crisis. In the UK, this has been particularly acute, widening the longstanding gap with its main peers (Oliveira Cunha et al., 2021; Van Reenen and Yang, 2023). Moreover, there are large and persistent gaps in productivity between firms and regions (De Loecker, Obermeier, and Van Reenen, 2022; Brandily et al., 2022). Investment in fixed capital, skills and innovation are key to addressing these gaps. Beyond invention, raising innovation - and thus productivity - requires firms to successfully adopt new technologies.

Recent years have seen rapid advances in the development of potential general-purpose technologies (GPT) with applications across the economy, in particular artificial intelligence (AI) and the current crop of generative AI variants, notably large language models.¹ These technologies are still in the early stages of diffusion, and while many expect this latest wave of technological change to bring with it significant positive impacts on productivity, uncertainty remains on what this will look like for firms (Baily, Brynjolfsson, and Korinek, 2023) as well as for workers. In part this is because adoption patterns are uneven and are impacted by the availability of complementary factors, including worker skills and managerial capabilities (David, 1990; Brynjolfsson, Rock, and Syverson, 2018, 2021).

In this context it is important to build a deeper understanding of the drivers of technology adoption, how these have changed over different waves of digitisation, and the implications for national and regional growth policies as digital technologies continue to advance.

This paper focuses on what are generally considered to be ‘skill-biased’ technologies. This means they are complementary with skilled workers at least in the adoption phase (Violante, 2008). We consider two different eras in the evolution of information and communications technology (ICT) development: the cloud computing and machine learning/AI

¹See, for example, Goldfarb, Taska, and Teodoridis (2022) who provide empirical evidence that a cluster of technologies comprised of cloud, machine learning and related data science is most likely to have general purpose characteristics compared to other emerging technologies; Brynjolfsson, Rock, and Syverson (2021) who argue that AI is a potential GPT; and Eloundou et al. (2023) who project GPT-like spreads of current and future waves of generative AI across tasks and jobs, using human and GPT-4 based assessment.

wave during the 2010s (hence ‘second wave’), and the personal computer (PC) wave of the 1990s-2000s (‘first wave’). At present, our main contribution is that we contrast the roles of general human capital (as proxied by the degree share) to that of science, technology, engineering, and mathematics (STEM) skills (as proxied by workforce occupational structure) in explaining differential adoption across these waves.

To do this, we construct new datasets containing measures of technology adoption and skills availability at the firm and region level. Given the challenges of obtaining granular measures of technology adoption in firms via standard survey or administrative data sources, our measures of technology adoption in the second cloud and ML/AI wave are based on the well-known Lightcast online job vacancies database.² Using the UK version of the database covering the period from 2012-2019, we derive technology adoption measures from the text of vacancies following the methodology of [Bloom et al. \(2021\)](#), and focus on cloud and machine learning/AI technologies (hence ‘ML/AI’) in the period immediately before generative AI applications. We also use data for the first wave of PC adoption derived from the database of Harte-Hanks, an ICT marketing company with a long historical presence in ICT equipment supply industry.

Figure 1 shows how measures of technology adoption that we will outline in this paper are strongly correlated with UK area-level productivity (measured by GVA/hour), even after controlling for industry and firm size composition in places. We explore the drivers behind these varying adoption patterns. We look at complementarities between skilled labour and different types of digital technology by estimating ‘demand equations’ ([Brynjolfsson and Milgrom, 2013](#)), relating technology adoption to the availability of skills across places ([Beaudry, Doms, and Lewis, 2010](#)) and within organisations ([Garicano and Heaton, 2010](#)).

We begin with an area-level analysis conducted at the level of 206 Travel-to-Work Areas (TTWAs). We consider the associations between lagged measures of regional human capital (degree share), the supply of STEM skills, and our technology measures, controlling for key covariates such as industrial structure, urbanisation and the underlying digital infrastructure

²Previously known as Burning Glass Technologies.

in places. The structure of the Lightcast data also allows us to credibly define firm-specific proxy measures of skill structure over time, based on the initial composition of vacancies for a large group of firms. For the second wave, this allows us to extend the adoption analysis to the firm level.

We find that there are significant differences in the nature of technology adoption across the two waves. Both are skill-biased in the sense that they are adopted more intensively where skilled workers are relatively more abundant, even after controlling for other regional covariates (including a London dummy). However, comparative advantage in technical skills (STEM) becomes much more important in technology adoption for the second cloud and ML/AI wave. In our area-level TTWA models, both general human capital and STEM skills have separate effects on PC adoption, but STEM skills dominate as the main explanation for the second wave technologies. We estimate that a one-standard deviation increase in the local share of STEM employees is associated with around 0.3 of a standard deviation increase in adoption of cloud computing and ML/AI technologies, respectively.

The firm-level analysis allows for more granular perspectives. The level of initial firm-specific STEM skills has a strong relationship with second wave technology adoption that is over and above the effects of high skill generally. Specifically, the coefficient for high skill is reduced by almost half when a STEM-specific skills measure is included in our technology adoption regressions. The dominant role of STEM skills is also supported by an auxiliary set of results that directly compares their effects to those of non-STEM professional skills (proxied by legal and administrative skills) as correlates of adoption.

The STEM-biased technology adoption effect that we uncover in these regional and firm-level models is in turn reflected in the concentration of technology. At the regional level, we summarise this in terms of the difference between the top quintile of STEM-intensive areas and the rest of the UK. The gap in average rates of adoption between STEM-intensive areas and the rest of the UK increased sharply over the period that we consider, particularly for ML/AI where STEM-intensive areas have adoption rates that are nearly three times higher than the rest of the UK.

The firm-level results indicate that this regional effect is likely to be rooted in a highly skewed adoption pattern for cloud and ML/AI technologies. The distribution of STEM skills is very concentrated amongst firms making the the contribution of the STEM-biased adoption effects especially high for the analogous top quintile of STEM-intensive firms. Technology adoption probabilities are 18-24% higher for these firms and, on the intensive margin, approximately 25-35% of their final levels of technology can be explained by the STEM effect.

This accelerated rate of adoption by the top STEM-intensive firms means technology gaps widen over the period that we consider. This even applies when comparing the fifth (top) and fourth quintiles of firms by initial STEM intensity. Finally, we also explore how this concentrated pattern of STEM-biased adoption translates in terms of the volume of vacancies. At the start of our sample in 2012 the top 20 firms account for 40% and 60% of all cloud and ML/AI vacancies respectively. By 2019, while cloud diffusion has not shifted, ML/AI technologies have diffused more widely: the top 20 firms now account for 40% of all ML/AI-related vacancies.

Overall, both area and firm-level evidence point to skill-biased, and in particular, STEM-biased adoption for cloud and ML/AI technologies. The time coverage of our data means that we are comparing a plausibly ‘complete’ adoption wave for PCs circa the early 2000s to an ongoing adoption wave for cloud and ML/AI technologies. However, the crucial point here is that gaps in adoption rates *actually widened* for cloud and ML/AI technologies as the 2010s progressed. The classical pattern of technology diffusion is for adoption to increase to a saturation point across a given population - but this did not take place. If cloud and, in particular, AI technologies continue to be unevenly distributed, this presents ongoing challenges in addressing regional inequalities.

Related Literature. Our work contributes to three main literatures. First, we contribute to the evidence on skill-biased technology and technology-skill complementarity. [Nelson and Phelps \(1966\)](#) argue that educated workers facilitate the diffusion of new technologies, and [Milgrom and Roberts \(1990\)](#) emphasise how ICT has shifted organisational de-

sign, benefiting more highly skilled workers, and this has been formalised by [Garicano and Rossi-Hansberg \(2006\)](#). In models of endogenous technology adoption ([Basu and Weil, 1998](#); [Zeira, 1998](#); [Caselli, 1999](#)), when a major technology becomes available, it tends to be adopted more quickly in environments where complementary factors are plentiful and cheap. [Brynjolfsson and Milgrom \(2013\)](#) provide an overview of the theory and empirics of organisational complementarities. In our geographic analysis, we follow [Beaudry, Doms, and Lewis \(2010\)](#) who find that cities in the United States with a more abundant supply of skills adopted computers more intensively. We perform a similar analysis in the UK for the first wave, and extend this to include STEM skills and to consider the early stages of 2010s technologies as well. In a similar vein, [Feng and Valero \(2020\)](#) use international data on manufacturing firms to show that firms in higher skill environments are more likely to adopt productivity enhancing modern management practices ([Bloom and Van Reenen, 2007](#)), arguing that this provides evidence in support of a complementarity between human capital and management practices in firms. [Stephany and Teutloff \(2024\)](#) provide an even more disaggregated analysis, using detailed online job platform data to study skill complementarities in terms of network-linked bundles, including an important cluster related to ML/AI technologies.

Second, we contribute to the broader economic geography literature on innovation and path dependencies. [Bryan and Williams \(2021\)](#) emphasise the role of information asymmetries in technology diffusion. Cities, especially large urban cores, mitigate such frictions. As well as providing deep pools of skilled human capital, big cities help new ideas diffuse through knowledge spillovers - as embodied in worker moves, firm-firm linkages and face-to-face interaction. Observed levels of innovation are highest in urban areas ([Carlino and Kerr, 2015](#)), partly driven by the clustering of complex activities ([Balland et al., 2020](#); [Davis and Dingel, 2019](#)), and knowledge spillovers are also highly localised ([Audretsch and Feldman, 1996](#); [Atkin, Chen, and Popov, 2022](#)). Within cities, urban cores generate higher levels of novel innovations ([Berkes and Gaetani, 2020](#)). Our analysis incorporates both area complements and features of these large agglomerations.

While today's innovation geographies are highly concentrated, the past century saw substantial shifts, associated with the diffusion of past technology waves ([Andrews and Whalley, 2022](#); [Esposito, 2023](#)). [Brezis and Krugman \(1997\)](#) suggest that 'leapfrogging' takes place when new technologies are non-overlapping with existing leaders' knowledge bases. Similarly, [Mestieri, Berkes, and Gaetani \(2021\)](#) show that urban industrial diversity makes US cities more economically resilient to successive technology waves. For our two technology waves, we similarly explore changing patterns of adoption at city-region level, and relate this to area-level differences in human capital and industry mix.

Third, we contribute to a recent literature that creates new measures of advanced technology adoption at the firm-level, with a special focus on cloud, machine learning and early AI technologies. A set of these focus on using the data in job ads to create measures of diffusion of technologies into firms. A pair of papers by [Babina et al. \(2022\)](#) use vacancy and CV data for the US to understand firm-level technology adoption, and this approach is also taken by [Acemoglu et al. \(2020\)](#) [Goldfarb, Taska, and Teodoridis \(2022\)](#) and [Bloom et al. \(2021\)](#). [Dahlke et al. \(2024\)](#) look at European firms using text-based measures of AI adoption with a focus on epidemic effects driven by direct between-firm linkages and mimetic imitation. Other papers are based on new survey questions. [Zolas et al. \(2021\)](#) outline the structure and the first results for a new survey module attached to the US Census Bureau's Annual Business survey circa 2019. They found that AI use was widespread across sectors but concentrated amongst large firms. [Calvino and Fontanelli \(2023\)](#) also utilise survey data but look at Europe and also a strong presence of AI amongst large firms alongside important complementarities related to skills and digital infrastructure.

The paper proceeds in the typical way. Section 2 describes the data that we use and how it is put together. Section 3 outlines our empirical modelling strategy. Section 4 reports our results. In Section 5, we conclude with the aim of positioning our contribution in the literature.

2 Data and measurement

2.1 Data

2.1.1 ‘Second wave’ technologies: cloud, machine learning and AI

Lightcast: Our main dataset, which we use for measuring the adoption of ‘second wave’ digital technologies, comes from Lightcast.³ Lightcast is a leading vendor of online job vacancy information for both commercial and academic usage. They webscrape information across online sources and de-duplicate entries in order to capture the universe of vacancies in a given country as comprehensively as possible. Academic studies using these data are growing in number. Relevant studies include the already-noted work on AI in firms by Babina et al. (2022) and research by Adams-Prassl, Balgova, and Qian (2023) on flexible work practices, using the UK version of the data.

The UK data that we use begins in 2012 and comprises approximately 59.9 million vacancies in total. Online vacancy data is a rich resource but has some limitations, and we briefly discuss three issues here. First, the name of the firm or organisation posting a vacancy can be directly identified for 33.2% of all vacancies, the remainder being vacancies advertised via a third-party recruiter. We use all of the vacancies to construct area-level and occupation-level datasets, but restrict to the 33.2% subset when doing firm or establishment-level analysis.⁴ Second, the volume of Lightcast vacancy data collected per time period can often vary dramatically and this is related to the varying efficiency of the tools that the company uses to scrape the web. We deal with this issue by normalising different data series according to context. For example, we look at the share of observations in a given area or industry that have adopted a technology rather than the volume. Third, online job vacancies also tend to be biased in their coverage against low-skilled, manual occupations such as those

³Previously known as Burning Glass Technologies. In June 2021 Burning Glass merged with EMSI and the merged firm was rebranded as Lightcast.

⁴Lightcast data contains a mixture of city/town and county information. We describe the consolidation of this information into a TTWA-level code in Appendix A.2.

in hospitality or high turnover segments of retail trade. That said, the relationship between the Lightcast data and population survey data is constant over time: [Adams-Prassl, Balgova, and Qian \(2023\)](#) report that occupation-county regressions of vacancy and wage measures on employment and wage measures from the Annual Survey of Hours and Earnings (ASHE) are characterised by stable R-squared measures.

Measures of technology adoption: We adapt the measurement approach of [Bloom, Hassan, Kalyani, Lerner, and Tahoun \(2021\)](#) [hence BHKLT] as part of our method for building vacancy-based indicators of technology adoption. BHKLT’s approach identifies common technology keywords across different domains. Specifically, they use a combination of patent data and company earnings call text to validate their choice of economically important bigrams (two-word phrases) that are then grouped as ‘technologies’. The earnings calls text come from the quarterly presentations by company management to investors and journalists, with an existing literature indicating that these calls are indicative of ‘some of the most important issues facing these organizations’. In short, BHKLT (2021) disciplines the choice of keywords using the overlap of scientific and business information. We describe BHKLT in more detail in [Appendix 5](#).

Our implementation uses the ‘end result’ of 29 BHKLT technologies and feeds the associated bigrams into the pre-processed UK Lightcast text of job vacancy advertisements. Some pre-processing that we apply is standard⁵ but we also conduct some additional cleaning to remove irrelevant content.⁶ As a final validation step, we inspect vacancy posts with unusual 4-digit SIC or 4-digit SOC codes (e.g. cloud computing expertise amongst florists). After this inspection we build a list of common causes for the false positives and set the bigram indicator to zero for these spurious vacancies.

⁵First we lowercase the content of each job post. Then we remove all words that contain any special characters other than alphanumeric characters, underscores, whitespace characters (spaces, tabs, line breaks) and hyphens (-). Each removed word is replaced by a space character.

⁶To detect cloud computing technology, for instance, we remove phrases such as “upload your cv upload from your computer or import from cloud storage”, “cloud based portal”, and “add your cv upload from your computer or import from cloud storage”. These phrases describe the application process rather than tasks or requirements for candidates and were found to lead to a large number of false positives (e.g. exposure to cloud computing for manual jobs such as chefs and cleaners).

We then define a binary variable that flags the presence of a given technology according to whether any one of the candidate bigrams is present in a job advertisement. Two out of the 29 BHKLT technologies (fingerprint sensors and stentgrafts) are associated with a negligible number of job postings, so we drop them from our analysis, leaving us with 27 technologies. We focus on two of these: cloud and machine learning/ artificial intelligence ('ML/AI'). As Appendix Figure A1 shows, these are clearly the most 'mature' or widely diffused technologies out of the 27 detected in the UK data.

Regional data: To construct regional data, we utilize the 'County/UA' information available in Lightcast and harmonize it with the official geospatial dataset from ONS before aggregating it to the TTWA level. This aggregation involves accounting for cases where one county/UA spans across multiple TTWAs. For more information, refer to the details provided in Appendix A.2.

Firm level data: To construct the firm-level panel, we group the vacancy data by variables that identify a type of unit and then apply different restrictions. We start with a sample of 1,051,820 observations, each identified by a unique triplet of the pre-processed, non-missing employer name, SIC code and year (between 2012-2020). The dyad of employer name and SIC code is then used to identify firms (we refer to this in our in Table A.4 as 'Sample 1'). Note that this necessarily identifies a firm as a multi-plant or multi-location unit since we do not use any geographical information. We then make a distinction between 'small' and 'large' firms based on the threshold of posting an average of 100 vacancies or more per year over the observed lifetime of the firm.

Our main analysis focuses on these large firms who approach nearly 2,000 in number (1,855 firms). Having a large pool of vacancies allows us to define less noisy firm-specific measures of baseline skills. The balanced panel of large firms also accounts for a large fraction of technology adoption activity. Specifically, while it only covers 1.59% of identifiable firm observations, the balanced panel accounts for 55.65% of cloud computing vacancies and 77.06% of ML/AI vacancies. We report on the characteristics of different samples of large versus small and balanced versus unbalanced firms in Appendix Table A1.

2.1.2 ‘First wave’ technologies: personal computers

Following [Beaudry, Doms, and Lewis \(2010\)](#) and [Autor, Dorn, and Hanson \(2013\)](#), we use PCs per person as our measure of technology adoption for the first wave of ICT development in the 1990s and 2000s. Our measures are derived from the ICT equipment database maintained by Harte-Hanks (HH), a multinational company that surveys equipment stocks at business sites in order to provide market intelligence information to ICT hardware and software vendors. Harte-Hanks surveys large establishments (100 employees or more) on an annual basis. Establishments have a site identifier allowing for consistent tracking over time as well as a postcode.

We construct area-level TTWA measures of PCs per person for the UK using a pooled establishment panel across 3 years from 2000-2002. Details of the aggregation process are given in Appendix [A.2](#). By measuring PCs per person in these years, we aim to capture the state of PC technology at the end of the 1990s wave of adoption. In line with [Beaudry, Doms, and Lewis \(2010\)](#), in our preferred regressions, we use a measure of PCs per person that pre-adjusts for firm size, industry interactions, year in the underlying firm micro-data, but also compare this adjusted measure to alternative measures that relax various steps (including straightforward weighted and unweighted area means of PCs per person).

2.1.3 Skill measures and area-level characteristics

Our regional skills measures are based on UK Census data. For general human capital, we use the 1991, 2001 and 2011 Census to construct a general skills measure based on the share of university graduates in a region’s population. In turn, we measure regional STEM skills with reference to the size of the STEM-related workforce, using occupational classifications. Specifically, our STEM share measure relates to the share of Science and Engineering Professionals and Associate Professionals in the workforce. We define this as those workers in the 2-digit SOC code categories of: 21 (Science Research, Engineering and Technology Professionals) and 31 (Science, Engineering and Technology Associate Professionals).

We aggregate to TTWA level from small area data using postcode weights, as described in Appendix A.2.

Firm-specific skills measures are based on the skill requirements of firms as found in their job vacancy listings. This is of course a proxy measure that assumes a strong correlation between vacancies (skills demanded) and the skills actually found in firms. Our main measure of interest is the share of STEM occupations in all vacancies posted by a firm as following the STEM job classification by Bakhshi , Davies and Freeman (Bakhshi et al., 2015), which encompasses 44 4-digit SOC codes. In our regressions, we also break up the overall STEM classified group into occupational sub-groups (eg. ICT versus R&D professionals) to understand within-STEM differences. Furthermore, we construct the share of postings by skill level using the first digit SOC codes defined by ONS including high skill (1-3), medium skill (4-7), and manual-intensive skill (8-9).

In addition to these skill variables, we also extract a variety of other information from the 2011 and 1991 Censuses. These include the resident adult population, share of people working in manufacturing, and share of unemployed workforce. To construct a proxy for digital infrastructure, we use information available Ofcom’s 2011 Communications Infrastructure Report which provides data on various aspects of broadband services delivered over fixed telecom networks.⁷ These small-area data are aggregated to TTWA using the same crosswalk strategy employed for the Census data when measuring skills.

3 Area-level descriptives, models and results

3.1 Descriptive statistics

Table 1 provides the descriptive statistics for our area-level sample of interest, which is comprised of 206 TTWAs covering England and Wales. Cloud and ML/AI adoption reached

⁷Available at: <https://webarchive.nationalarchives.gov.uk/ukgwa/20200803095351/>, <https://www.ofcom.org.uk/research-and-data/multi-sector-research/infrastructure-research>.

respective rates of 1% and 0.2% per vacancy in 2019, which are approximately triple and 13-fold increases since our initial measurement year of 2012.

However, despite this large increase in levels Figure 2 shows that new wave technology adoption is highly clustered across space, with relatively little spatial diffusion evident over our sample period. Each panel in Figure 2 shows TTWA-level location quotients (LQs) of adoption for 2012 and 2019, for cloud (left) and ML/AI (right). LQs over one indicate local adoption shares greater than the national share for that year, a measure of co-location. Most places have LQs under one at both the start and end of the 2010s. Compared with cloud, ML/AI adoption is rather more spatially concentrated, with less diffusion across locations over the decade.

Table 1 also shows that there has been an increase in our human capital and skills measures since the 1990s. Human capital, as measured by the labour force share of graduates, increased dramatically from only 6.2% in 1991 to 25.0% by 2011. As we will discuss, this increase was not proportionate across types of qualification with the STEM share of the workforce increasing by less than one-third, from 3.8% in 1991 to 5.0% in 2011.

The basic relationships between our area-level technology adoption measures and the two measures of human capital and skills are then plotted in Figure 3. The top panel shows the relationship between degree-level skills and technology adoption at the TTWA-level. As per [Beaudry, Doms, and Lewis \(2010\)](#), we use the (log) ‘skilled-to-unskilled’ ratio to measure skill (x-axis) and technology counts per employee for technology (y-axis). The bottom panel repeats this with a focus on STEM skills, again measured in terms of a skill ratio.

For both wave 1 and wave 2 technologies, the Figure 3 plots show strong unconditional correlations, which are higher for general degree-level skills than STEM skills. However, there is no overt difference in the strength of the correlations between technology adoption and human capital across the two waves. That is, there is no indication of differential skill bias in adoption patterns. Of course, this does not account for omitted variables or the fact

that the separate effects of general and STEM skills cannot be distinguished in a bivariate analysis - but these correlations motivate the regression analysis that follows.

3.2 Regression models

To assess the conditional correlation at the area level (TTWAs) we estimate the following regression model:

$$\Delta Y_{jt} = \gamma_0 + \gamma_1 \ln\left(\frac{S}{U}\right)_{j,t_0} + \Pi_1 z_{j,t_0} + \varepsilon_{jt} \quad (1)$$

where ΔY_{jt} is the change in ‘technology intensity’ (per employee or vacancy) in TTWA j at time t (relative to some t_0 base period), $\left(\frac{S}{U}\right)$ is the ratio of skilled to unskilled workers in TTWA j at t_0 , z_{j,t_0} is a vector of controls, and ε_{jt} is an error term. This regression model is implemented separately for first wave PC adoption and for ‘new wave’ cloud and ML/AI technologies. In each case we use measures of skills that lag current technology levels by 8-11 years, as dictated by data availability. For example, since PC adoption levels are measured in 2000-2002 we use measures of skills from the 1991 Census, while for Cloud and ML/AI (measured in 2019) we use skills indicators from the 2011 Census.

The inclusion of control variables z_{j,t_0} in this model is aimed at dealing with the omitted variables problem that naturally underlies the unconditional correlation. However, there is also another fundamental issue related to the overlap between degree-level and STEM skills. To distinguish the effects of each type of skill we extend this model as follows:

$$\Delta Y_{jt} = \gamma_0 + \gamma_1 \ln\left(\frac{S}{U}\right)_{j,t_0} + \gamma_2 \ln\left(\frac{S^{STEM}}{U^{STEM}}\right)_{j,t_0} + \Pi_1 z_{j,t_0} + \varepsilon_{jt} \quad (2)$$

where $\left(\frac{S^{STEM}}{U^{STEM}}\right)_{j,t_0}$ is the initial ratio of STEM-classified workers to non-STEM workers in TTWA j at time t_0 . Again this is measured 8-11 years before the current period t for both of our technology waves. Given the human-capital intensity of STEM occupations, we

expect some correlation between the degree share and STEM intensities. Hence a finding of $\gamma_2 > 0$ would indicate that area STEM intensity has a distinct relationship with technology adoption that cannot be conflated with degree-level qualification intensity.

Two challenges for the interpretation of this regression model are area-level confounders and reverse causality (Beaudry, Doms, and Lewis, 2010). Confounders are what BDL call ‘innovative city factors’ - such as local clusters/milieux effects, unobserved worker quality and stocks of supporting equipment (including earlier IT hardware/infrastructure). These will be correlated with the stock of skilled workers and may also help raise technology adoption. Reverse causality arises due to area-specific long-run trends in the demand for skilled workers - for example, broader patterns of urban revival - that could either attract those workers and/or directly raise technology adoption.

As a response to this challenge, the z_{j,t_0} controls are designed to account for confounders. In all regressions we introduce agglomeration proxies (TTWA population density in the base year), milieu controls (London and Oxbridge dummies) and industry/labour market controls (share of manufacturing and ILO unemployment rate, both in the base year). In Wave two regressions, we also include controls for enabling technologies and infrastructure (TTWA shares of super-fast broadband connections in 2011) and measures of earlier PC adoption. We account for differences in local industrial structure by pre-adjusting our adoption measures by firm size-SIC3 cells in each TTWA. We further weight all regressions by TTWA working-age population.

Finally, note that in the case of the ΔY_{jt} dependent variable we assume that the base value for technology adoption is zero at t_0 and effectively use the level of adoption in year t as the measure of the ‘change’. In practice, this amounts to an assumption that PCs per person are zero in 1991 and that the second wave technologies (cloud computing and ML/AI) are in a similar state circa 2011. We benchmark our PC specification against that of Beaudry, Doms, and Lewis (2010), inclusive of adjustments for industry composition and cleaning rules (eg. winsorizing). The results for this are reported in appendix Table A2 and, encouragingly, our

estimates are easily within the range of BDL's(2010) analogous models.⁸

3.3 Area-level results

Table 2 reports the main estimation results for the models represented by equations (1) and (2). We standardize the variables to allow for comparison across the two waves. Panel (A) focuses on the general degree-level specification outlined in equation (1). The first three columns represent conditional models that add controls for such variables as population density, industry composition, unemployment rates, and specific area dummies. These show an additional sensitivity of cloud adoption relative to both PC adoption and ML/AI adoption. Specifically, a one-standard deviation increase in degree-level skills is associated with +0.58 standard deviations higher adoption rate for cloud relative to (approximately) +0.3 for the other two technologies. Furthermore, as column (4) indicates, this additional sensitivity of cloud technologies is robust to the inclusion of controls for related, potentially complementary technologies, namely broadband penetration and the level of PCs per worker. Of these, the coefficient on broadband is positive and significant and reduces our coefficients of interest slightly.

Panel (B) then implements the equation (2) model that compares the role of general degree versus STEM skills. These estimates indicate a finding of $\gamma_2 > 0$ for all three technology measures, that is, a distinct association of STEM skills with adoption that cannot be conflated with general degree-level skills. Furthermore, across the various specifications in Panel (B) it is evident that STEM skills dominate general degree-level skills across the three technologies, especially for the second wave technologies. Specifically, the Panel (B) estimates indicate that a one standard deviation increase in STEM skills is associated with a robust 0.35 standard deviation increases in the adoption of ML/AI technologies with only weaker associations apparent for cloud and STEM skills (0.27), and for general degree-level

⁸Beaudry, Doms, and Lewis (2010) show that the log ratio of college equivalents to non-college equivalents in 1980 - comparable to our log ratio of skilled to unskilled workers - is positively correlated with PCs adoption. Specifically they find that 10% increase in skills is associated with 1.5 extra PCs per 100 population, while we find that a 10% increase in skilled in associated with 1.4 extra PCs per 100 population.

skills across both technologies.

The consequences of this STEM-biased adoption effect for the regional distribution of technology are illustrated in Figure 4. Here we plot mean adoption rates according to five quintiles of initial STEM intensity. This shows that a steady increase in the top quintile's rate of technology adoption compared to the other broad regions. In the case of ML/AI, the top quintile's rate is nearly three times higher than the lower quintiles. Interestingly, there is a clear levels gap between the top quintile and the rest. This points back to the evidence in Figure 2 which showed a clustered and highly persistent pattern of technology adoption in the UK.

4 Firm-level descriptives, models and results

4.1 Descriptive statistics

The area-level models presented above point to a STEM-biased adoption pattern for second wave technologies defined in terms of local skill supplies. We now look at this pattern more closely using the firm-level information derived from the vacancy data. Overall, there was rapid adoption of cloud and ML/AI technologies during the 2010s. Figure 5 shows this for our panel of 1,855 firms. Cloud adoption grew from approximately 17% of firms in 2012 to 48% in 2019 while ML/AI adoption increased from 5% to 28% of firms over the same period. On the intensive margin (panel (b)), cloud vacancies represented approximately 2% of vacancies by 2019 while ML/AI vacancies constituted around 0.7%, with especially strong growth from 2016.

The role of STEM skills within this adoption pattern is then illustrated in Figure 6. We divide firms according to quintiles of initial STEM intensity in 2012 and plot average intensive margin values by year. This shows the top quintile 'breaking away' sharply from the lower quintiles. The most STEM-intensive firms begin the period with a higher share

of job ads specifying the ‘new wave’ technologies and then seek them at a faster rate than less (initially) intensive firms. Again, while this points to a strong relationship between STEM skills and adoption these are unconditional relationships that are subject to omitted, correlated factors.

4.2 Regression models

4.2.1 The role of STEM skills in firm-level adoption

We now test for the role of firm-specific skills using the following regression model:

$$Y_{ikt} = \theta_0 + \theta_1 \text{ShareSTEM}_{i,k,t_0} + \theta_2 \text{ShareHigh}_{i,k,t_0} + \theta_3 \text{ShareMid}_{i,k,t_0} + \Pi_1 z_{i,k,t_0} + \tau_t + \mu_k + \epsilon_{jt} \quad (3)$$

where Y_{ikt} is technology adoption measured either as a 0-1 indicator (extensive margin) or a rate per overall number of vacancies (intensive margin). Our key variable of interest $\text{ShareSTEM}_{i,k,t_0}$ is the fraction of vacancies classified as STEM in the baseline year of 2012 while $\text{ShareHigh}_{i,k,t_0}$ and $\text{ShareMid}_{i,k,t_0}$ are similar fractions for the share of vacancies classified as high or low skill according to the ONS’ standard 1-digit taxonomy based on occupation codes⁹. The z_{i,k,t_0} vector contains non-skill based measures of baseline firm and industry characteristics, μ_k are industry k fixed effects, τ_t are time effects and ϵ_{jt} is an error term.

We report the results for this specification in Table 3. The first two columns in each panel present specifications that focus solely on either STEM or non-STEM skills with the third column reporting the full model. The results indicate a robust association between STEM skills and adoption that exists over and above that of high or medium skills. Furthermore, these results indicate that STEM skills are a major driver of the overall effect of high skills. As an example, in the extensive margin cloud adoption model the coefficient on high skills falls by nearly half (from 0.38 to 0.21, both significant at 1 percent) after controlling for

⁹We net out STEM vacancies from the High and Mid categories to avoid double counting.

STEM skills.

In terms of magnitudes, these estimates are best understood by comparing firms by their quintile of initial STEM intensity. A firm in the top quintile has a 0.336 initial STEM share compared to zero for firms in the bottom quintile. In turn, this means that a top quintile firm has a 23.7% higher probability of adopting cloud computing and an 18.8% higher probability of adopting ML/AI relative to firms in the bottom quintile. Note that this effect will still be large even compared to firms in the top quintile since their mean initial STEM share is only 0.108.

On the intensive margin, this top quintile effect translates into increases in the share of technology-rated vacancies by 0.02 and 0.005 for cloud and ML/AI respectively. To put this into context, these effects can explain 35% and 25% of the observed levels of cloud computing and ML/AI technologies in top quintile firms circa the end of our sample in 2019. STEM skills have therefore played a leading role in the divergence in ML/AI adoption patterns seen in Figure 6. We present a discretised, quintile specification of equation 3 in Figure A4 which shows the dominant role played by the top quintile of STEM-intensive firms for our main firm-level results.

4.2.2 Different types of STEM skills

As discussed, the shifts in coefficient estimates in Table 3 suggest that STEM skills drive a major portion of the overall effect of high (i.e. professional) skills. We break this down further in Figure 7 where we decompose STEM skills according to its occupational sub-groups and define various managerial, legal and sales-related skills as a non-STEM benchmarks.

These results are consistent in showing a limited role for non-STEM professional skills, including the advanced ‘soft skills’ associated with legal, administrative, sales and executive occupations. In terms of within-STEM comparisons there is evidence of heterogeneity across the two technologies. ICT professional skills (e.g. programming) are more important for cloud adoption than other types of STEM skills while scientific skills (as represented

by formal scientific occupations) play the leading role in ML/AI adoption. In short, these results confirm that ‘hard’ STEM skills have a stronger association with ML/AI adoption than even the most sophisticated subsets of ‘soft’ skills.

4.2.3 How concentrated is adoption?

The fact that STEM skills are both correlated with adoption and that they are subset of overall skills creates the potential for a concentrated pattern of adoption. Our analysis so far has focused on firm-level means and shown that a pattern of STEM-intensive firms ‘breaking away’ with higher levels of adoption is robust to controlling for industry structure, non-STEM professional skills, and firm size. We now study this in terms of vacancy-weighted numbers to understand how technology adoption is distributed across firms.

In Figure 8 we first rank firms according to the share of technology-related vacancies that they account for in 2012 and 2019. We then calculate the rank-sorted cumulative sum. As an example, panel A shows that the top 20 firms accounted for around 60% of all ML/AI vacancies in 2012 and 40% in 2019. We also plot the rank-sorted cumulative sum of all STEM-related vacancies as a benchmark.

Concentration levels fall for both cloud and ML/AI technologies between 2012 and 2019, indicating that substantial diffusion did take place in the 2010s. That said, concentration is still very high in 2019 with 20 firms accounting for 40% of all technology-related vacancies in 2019. And, by the time that the top 100 firms are considered, this reaches a remarkable 80%. Furthermore, this cannot be explained as a general STEM effect since the concentration of STEM vacancies is around half that if the ‘new wave’ technologies in 2019. Nor can it be explained as a function of industry or firm size heterogeneity, that is, the clustering of adoption into specialised industries. In Figure A5 we purge the raw vacancy shares of industry and firm size effects before plotting the 2019 concentration curve and this only reduces concentration by one-quarter. Overall, our firm results for cloud and ML/AI technologies show both a distinctive STEM-bias of adoption as well as a very acute cluster-

ing of adoption in a relatively small number of firms that holds even on a within-industry basis.

5 Conclusions

We test for skill-biased and STEM-biased technological adoption associated with ML/AI and cloud technologies, along with a quantification of its importance. In the UK, a pattern of STEM-biased adoption is apparent at both the area and firm-level for ML/AI and cloud technologies at the early stages of adoption. At area level, we also compare adoption behavior in the 2010s with the PC wave of the 1990s and early 2000s. This indicates that general human capital (as proxied by degree shares) is much less important for the modern wave of cloud and ML/AI adoption. So far, the STEM effect has dominated as a correlate of adoption.

Importantly, the firm-level evidence for the 2010s allows us to isolate the source of the STEM effect as emanating from the top group of STEM-intensive firms. Insofar that ‘new wave’ cloud and ML/AI technologies are associated with higher productivity, this points to these technologies exerting pressure to increase differences in productivity between firms and regions rather than helping to reduce them. Also, when considering how the volume of vacancies is distributed, the high degree of concentration of cloud and ML/AI vacancies amongst a small number of large firms reinforces this finding regarding the potential direction of future productivity differences.

We must be clear about the limitations in our analysis. First, at present our analysis cannot rule out all sources of endogeneity, and we therefore interpret the relationships we present as robust associations. Second, our measure of technology adoption reflects one channel - firm hiring activity - and cannot differentiate between firms developing or using technologies, nor delineate how these are applied to the business. Finally, when building the firm-level variables, we construct base year skills measures that represent a flow and not a stock. This can create ambiguity because a firm with high baseline skills capabilities might

hire more or less skilled workers in any given year depending on their strategy.

As discussed, the time coverage of our data means that we are comparing a complete wave of PC adoption in the 2000s to an incomplete or ‘in progress’ wave of ML/AI and cloud adoption - prior to the proliferation of generative AI since 2022. This is an important consideration but it is notable that gaps in cloud and AI adoption have significantly widened with time rather than been being reduced (as would occur with a classical diffusion pattern). A possible explanation is that, per [Eloundou et al. \(2023\)](#), a wave of user application technologies could emerge from the STEM-intensive firms we consider and become more widely adopted by the less STEM-intensive firms. However, even in this scenario, as long as there is a correlation between different levels of cloud and AI expertise and different levels of productivity, the new wave technologies are still more likely to increase inequality than reduce it.

Bibliography

- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2020. “AI and jobs: Evidence from online vacancies.” Tech. rep., National Bureau of Economic Research.
- Adams-Prassl, Abi, Maria Balgova, and Matthias Qian. 2023. “Firm concentration job design: the case of schedule flexible work arrangements.” Working paper, Institute for Fiscal Studies.
- Andrews, M. J. and A. Whalley. 2022. “150 years of the geography of innovation.” *Regional Science and Urban Economics* 94:103627.
- Atkin, D., M. K. Chen, and A. Popov. 2022. “The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley.” Tech. Rep. 30147, National Bureau of Economic Research. URL <http://www.nber.org/papers/w30147.pdf>.
- Audretsch, D. and M. Feldman. 1996. “RD Spillovers and the Geography of Innovation and Production.” *American Economic Review* 86:630–640.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103 (6):2121–68. URL <https://www.aeaweb.org/articles?id=10.1257/aer.103.6.2121>.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2022. “Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition.” In *NBER Volume on Technology, Productivity, and Economic Growth*. NBER.
- Baily, Martin Neil, Erik Brynjolfsson, and Anton Korinek. 2023. “Machines of mind: The case for an AI-powered productivity boom.” .
- Bakhshi, Hasan, John Davies, Alan Freeman, and Peter Higgs. 2015. “The geography of the UK’s creative and high-tech economies.” .
- Balland, P.-A., C. Jara-Figueroa, S. G. Petralia, M. P. A. Steijn, D. L. Rigby, and C. A. Hidalgo. 2020. “Complex economic activities concentrate in large cities.” *Nature Human Behaviour* 4 (3):248–254.
- Basu, Susanto and David N. Weil. 1998. “Appropriate Technology and Growth.” *The Quarterly Journal of Economics* 113 (4):1025–1054.
- Beaudry, Paul, Mark Doms, and Ethan Lewis. 2010. “Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas.” *Journal of Political Economy* 118 (5):988–1036.
- Berkes, E. and R. Gaetani. 2020. “The Geography of Unconventional Innovation.” *Economic Journal* 131 (May):1466–1514.
- Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun. 2021. “The Diffusion of Disruptive Technologies.” Working Paper 28999, National Bureau of Economic Research. URL <http://www.nber.org/papers/w28999>.
- Bloom, Nicholas and John Van Reenen. 2007. “Measuring and explaining management practices across firms and countries.” *The quarterly journal of Economics* 122 (4):1351–1408.

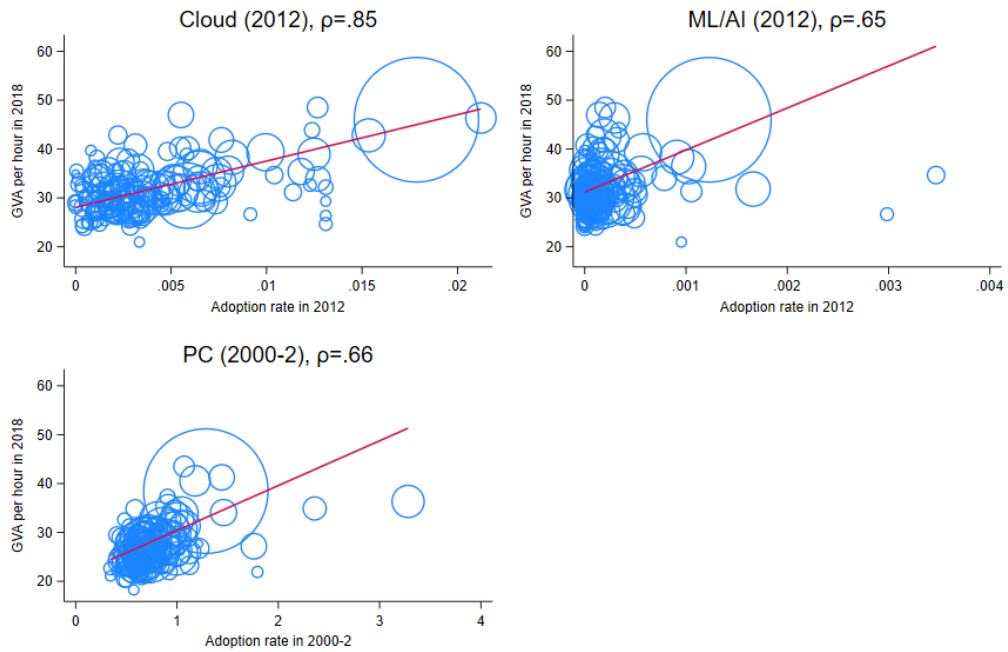
- Brandily, Paul, Mimosa Distefano, Helene Donnat, Immanue Feld, Henry G Overman, and Krishan Shah. 2022. “Bridging the productivity gap between different areas in the UK.” .
- Brezis, E. S. and P. R. Krugman. 1997. “Technology and the Life Cycle of Cities.” *Journal of Economic Growth* 2 (4):369–383.
- Bryan, Kevin A and Heidi L Williams. 2021. “Innovation: market failures and public policies.” In *Handbook of Industrial Organization*, vol. 5. Elsevier, 281–388.
- Brynjolfsson, E. and P. Milgrom. 2013. “Complementarity in Organisations.” In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, chap. 1. Princeton, 11–55.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018. “Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics.” In *The economics of artificial intelligence: An agenda*. University of Chicago Press, 23–57.
- . 2021. “The productivity J-curve: How intangibles complement general purpose technologies.” *American Economic Journal: Macroeconomics* 13 (1):333–72.
- Calvino, Flavio and Luca Fontanelli. 2023. “A portrait of AI adopters across countries: Firm characteristics, assets’ complementarities and productivity.” Working Paper 2023/02, OECD. URL https://www.oecd-ilibrary.org/science-and-technology/a-portrait-of-ai-adopters-across-countries_0fb79bb9-en.
- Carlino, G. and W. R. Kerr. 2015. *Agglomeration and Innovation*, vol. 5. Elsevier, 349–404.
- Caselli, Francesco. 1999. “Technological revolutions.” *American Economic Review* 89 (1):78–102.
- Dahlke, J, M Beck, J Kinne, D Lenz, R Dehghan, M Wörter, and B Ebersberger. 2024. “Epidemic effects in the diffusion of emerging digital technologies: evidence from artificial intelligence adoption.” *Research Policy* 53 (2):104917.
- David, Paul A. 1990. “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox.” *The American Economic Review* 80 (2):355–361. URL <http://www.jstor.org/stable/2006600>.
- Davis, D. R. and J. I. Dingel. 2019. “A Spatial Knowledge Economy.” *American Economic Review* 109 (1):153–170. URL <http://www.nber.org/papers/w18188.pdf>.
- De Loecker, Jan, Tim Obermeier, and John Van Reenen. 2022. “Firms and inequality.” .
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. “GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models.” *arXiv preprint arXiv:2303.10130v5* .
- Esposito, C. R. 2023. “The geography of breakthrough invention in the United States over the 20th century.” *Research Policy* 52 (7):104810.
- Feng, Andy and Anna Valero. 2020. “Skill-biased management: evidence from manufacturing firms.” *The Economic Journal* 130 (628):1057–1080.
- Garicano, L. and Paul Heaton. 2010. “Policemen, managers, lawyers: New results on complementarities between organization and information and communication technology.” *International Journal of Industrial Organization* 28 (4):355–358.

- Garicano, Luis and Esteban Rossi-Hansberg. 2006. "Organization and inequality in a knowledge economy." *The Quarterly journal of economics* 121 (4):1383–1435.
- Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis. 2022. "Could Machine Learning be a General Purpose Technology? A Comparison of Emerging Technologies Using Data from Online Job Postings." Working Paper 29767, National Bureau of Economic Research. URL <http://www.nber.org/papers/w29767>.
- Mestieri, M., E. Berkes, and R. Gaetani. 2021. "Cities and Technological Waves." Tech. Rep. DP16794, Centre for Economic Policy Research.
- Milgrom, P. and J. Roberts. 1990. "The Economics of Modern Manufacturing - Technology, Strategy, and Organization." *American Economic Review* 80 (3):511–528.
- Nelson, Richard R and Edmund S Phelps. 1966. "Investment in humans, technological diffusion, and economic growth." *The American economic review* 56 (1/2):69–75.
- Oliveira Cunha, Juliana, Jesse Kozler, Pablo Shah, Gregory Thwaites, and Anna Valero. 2021. "Business time: how ready are UK firms for the decisive decade?" .
- Stephany, F and O Teutloff. 2024. "What is the price of a skill? The value of complementarity." *Research Policy* 53 (1):104898.
- Van Reenen, John and Xuyi Yang. 2023. "Cracking the Productivity Code: An international comparison of UK productivity." .
- Violante, Giovanni L. 2008. "Skill-biased technical change." *The new Palgrave dictionary of economics* 2:1–6.
- Zeira, Joseph. 1998. "Workers, machines, and economic growth." *Quarterly Journal of Economics* 113 (4):1091–1117.
- Zolas, Nikolas Jason, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David Beede, Catherine Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. 2021. "Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey." Working Paper 28290, National Bureau of Economic Research.

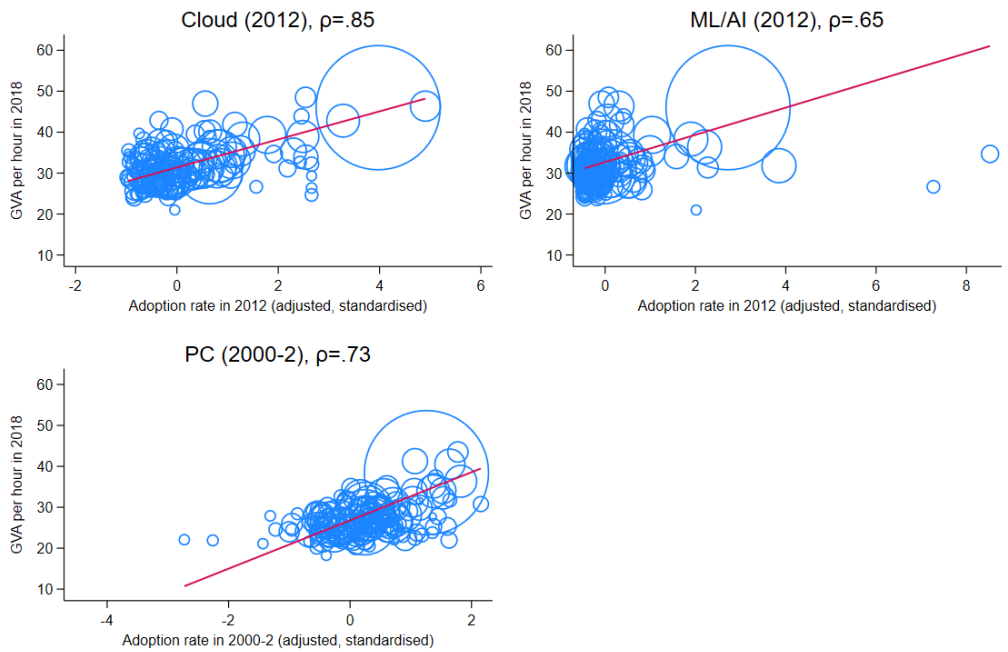
Figures

Figure 1: Productivity disparities and technology adoption

Panel A: Raw Adoption



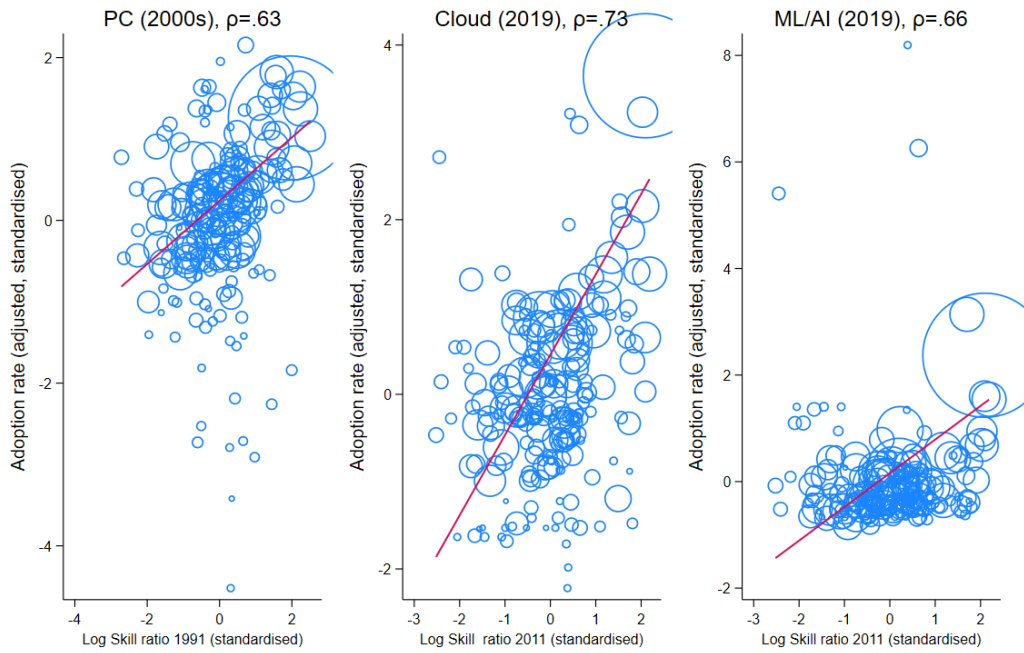
Panel B: Adoption adjusted by firm size and industry



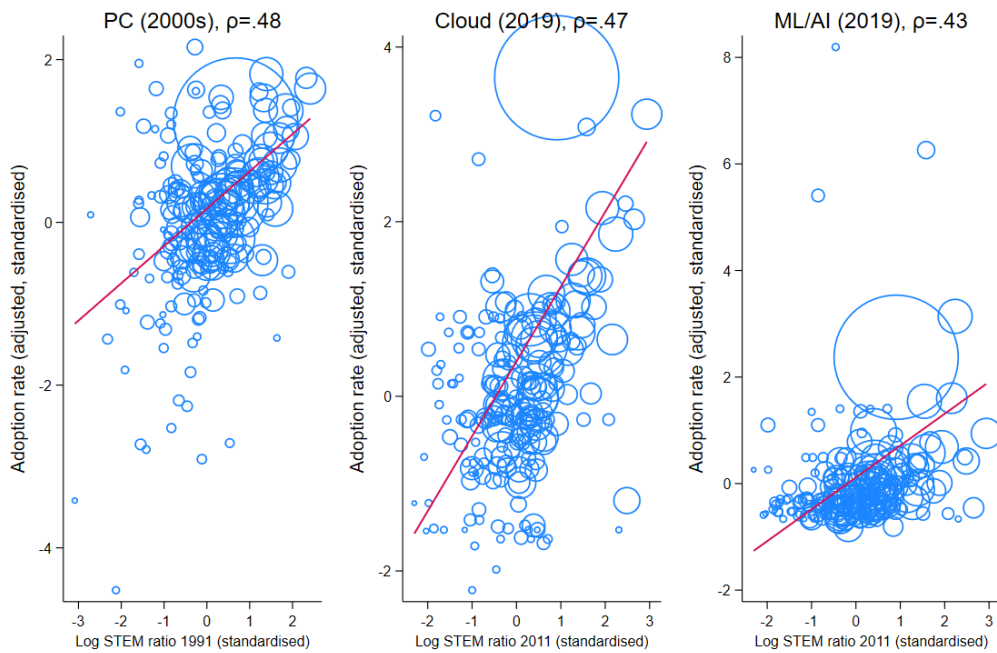
Source: ONS, Lightcast, Harte-Hankes. Notes: Each figure shows the raw correlation of GVA per hour against our lagged technology adoption measures, adjusted for 3-digit industry and firm size. Each circle in a figure is a UK Travel To Work Area, weighted by population aged 16+ in 2011. PC adoption is computers/worker. Cloud and ML/AI adoption is relevant job ads per 1000 ads. The graph titles report correlation coefficient ρ .

Figure 3: Skills in Two Waves

Panel A: General skill

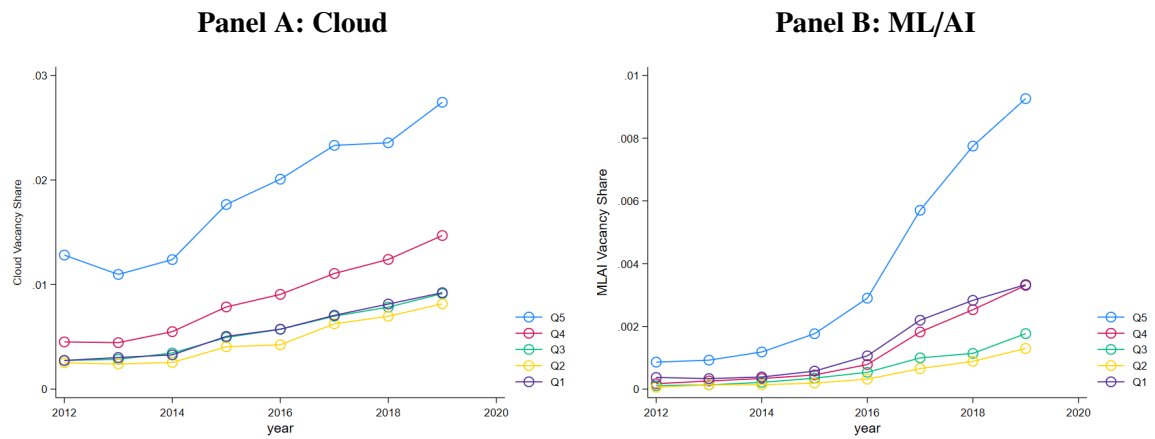


Panel B: STEM skill



Source: ONS, Lightcast, Harte-Hankes. Notes: Dependent variables are number of PC per employee (2000-02) and share of vacancies exposed to cloud computing and ML/AI. All dependent variables are regression-adjusted (for establishment size and industry) and standardised. The graph titles report correlation coefficient ρ .

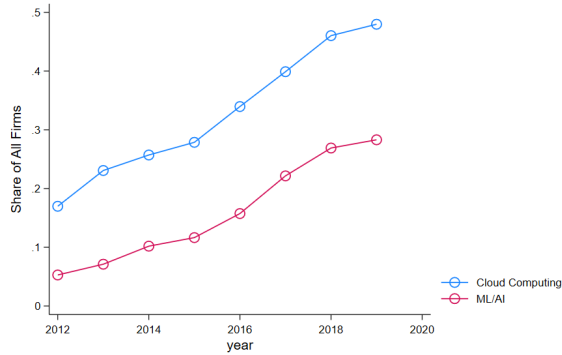
Figure 4: Cloud and ML/AI Adoption by STEM Quintiles



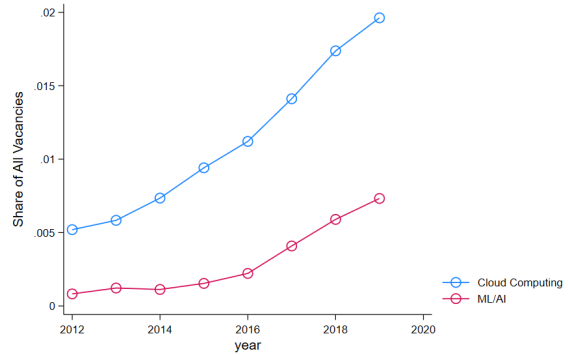
Source: Lightcast. Notes: This figure plots average levels of technology adoption across five quintiles of area STEM workforce shares as defined in 2011 (UK Census data). N = 206 TTWAs. These are unweighted means across the approximately 40 TTWAs in each quintile. Q5 is the quintile with the highest STEM intensity.

Figure 5: Cloud and ML/AI adoption trend (balanced panel)

Panel A: Extensive margin



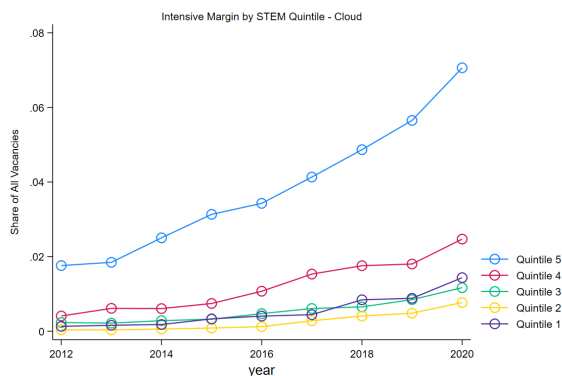
Panel B: Intensive margin



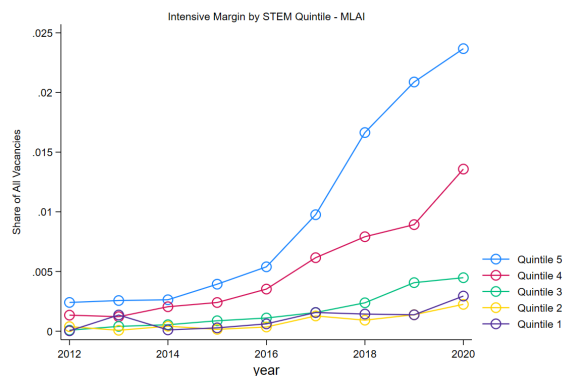
Source: Lightcast. Notes: This graph plots the hiring trends of a balanced panel of 1,855 firms with an average of 100+ posts per year. The Extensive Margin in Panel A is measured as the fraction of firms who report any Cloud or ML/AI vacancy while the Intensive Margin measures the fraction of overall vacancies related to the technology .

Figure 6: Intensive margin by initial STEM intensity

Panel A: Cloud

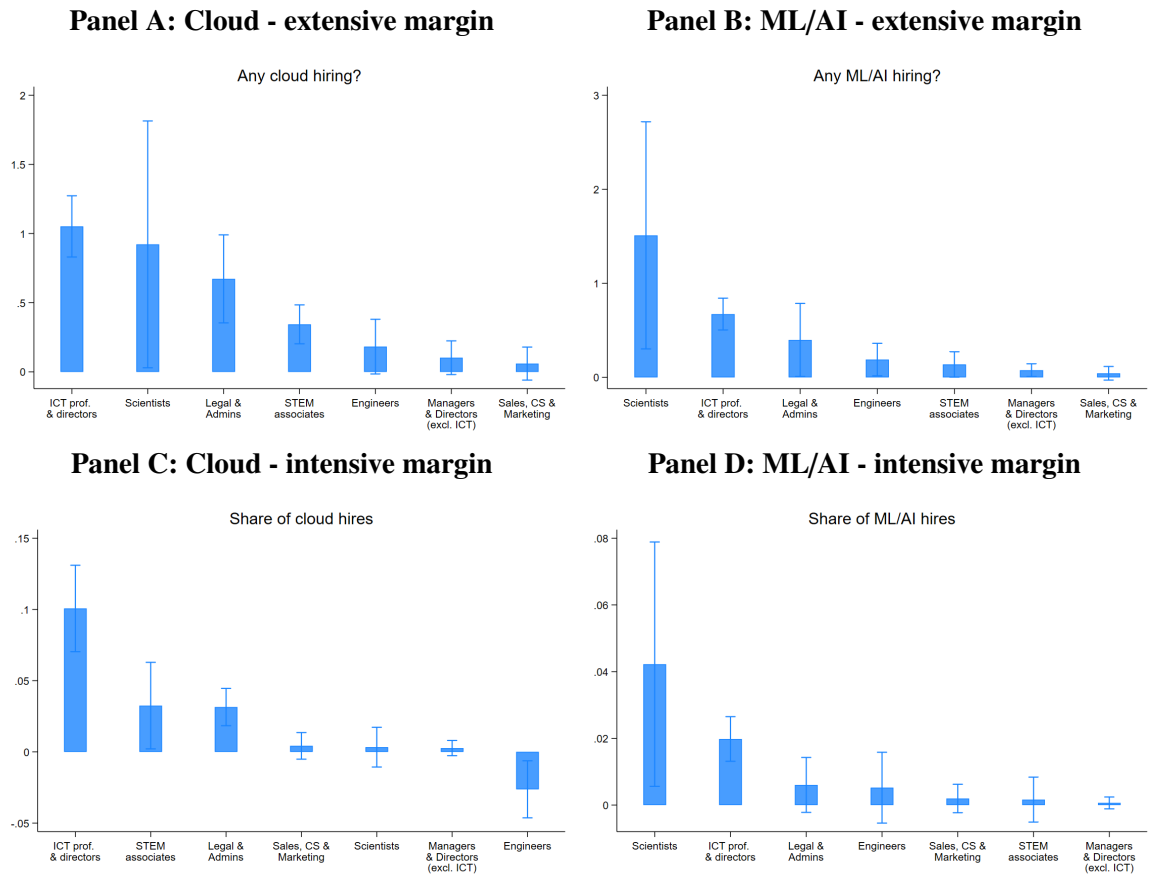


Panel B: ML/AI



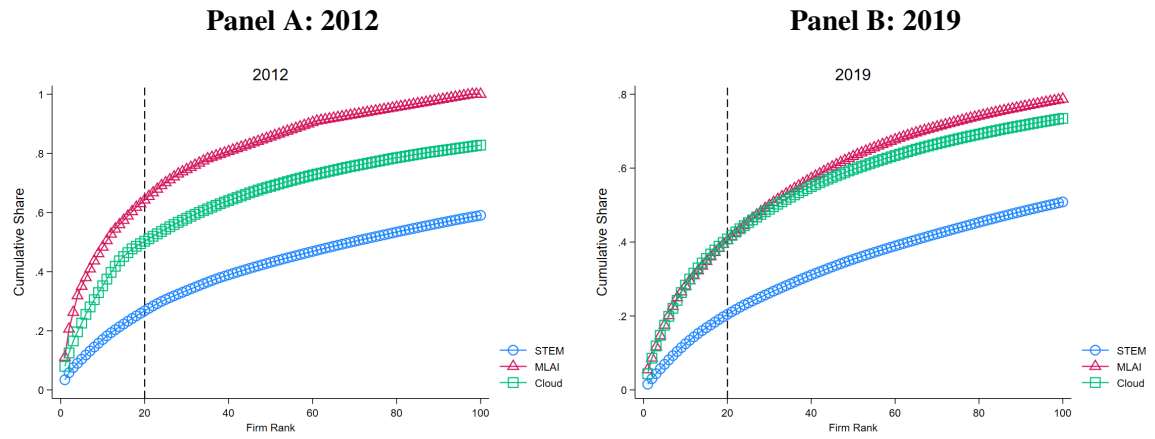
Source: ONS, Lightcast. Notes: This figure plots the average technology adoption by the firm-level STEM intensity in 2012, where STEM intensity is measured as using the STEM job classification by (Bakhshi et al., 2015). Firms are allocated to different quintiles of the STEM intensity distribution across all firms in 2012. Sample and definitions of extensive/intensive margins as per Figure 5.

Figure 7: Technology adoption by skill type



Source: ONS, Lightcast. Notes: This figure plots the coefficients (with 95% confidence intervals) of technology adoption models analogous to equation (3) but with initial skill groups disaggregated by the occupation of vacancies. The solid blue bars are meant to index the size of coefficients so that they can be visually compared. The relevant 2-digit SOC groups are ICT professionals (213x,1136), Legal and Admins (241x,242x), STEM Associates (31xx), Engineers (212x), Managers/Directors (1xxx, excluding 1136) and Sales, Customer Service and Marketing (354x,71xx,72xx). All specifications include year, firm SIC2 and initial size controls.

Figure 8: Concentration of Technology-Related Vacancies



Source: ONS, Lightcast. Notes: This figure plots the cumulative sum of technology-related vacancies according to the the rank of firms based on their vacancy volume. That is, the firm offering the largest number of vacancies for a technology type is ranked first, the next largest is second and so on. The top 100 firms are shown.

Tables

Table 1: Descriptive Statistics

	Mean	sd	Min	Max
(1)				
(A) Technology				
PC per employee	0.721	0.377	0.034	3.282
Cloud Adoption 2012	0.003	0.004	0.000	0.021
Cloud Adoption 2019	0.010	0.007	0.000	0.045
ML/AI Adoption 2012	0.000	0.003	0.000	0.003
ML/AI Adoption 2019	0.002	0.003	0.000	0.025
Broadband	0.366	0.287	0.000	0.857
(B) Human Capital				
Share graduates 1991	0.062	0.021	0.023	0.134
Share graduates 2011	0.250	0.049	0.143	0.369
Share STEM 1991	0.039	0.014	0.012	0.087
Share STEM 2011	0.050	0.016	0.023	0.115
(C) Demographics				
Pop. density 91	2.476	2.989	0.026	20.663
Pop. density 11	2.895	3.471	0.031	26.774
Share Manuf.1991	0.172	0.057	0.053	0.358
Share Manuf.2011	0.100	0.035	0.037	0.233
UR 1991	0.084	0.024	0.035	0.173
UR 2011	0.057	0.016	0.026	0.118
Observations	206			

Source: ONS, Ofcom, Lightcast, Harte-Hankes. Notes: Cloud and ML/AI Adoption rates per vacancy. Share of Manuf. 91/11 is the share of employed working in manufacturing in 1991/2011. UR91/11 is the unemployment rate in 1991/2011. Broadband is the share of super-fast broadband coverage in 2011 and is from Ofcom. PC per employee data is from Hart-Hanks. Cloud and ML/AI data is from BGT. Other data is from Census. The sample include Scotland, England and Wales TTWAs with non-missing STEM data.

Table 2: Technology, lagged general and STEM skills (standardised), comparing waves

	(1)	(2)	(3)	(4)	(5)
	PC	Cloud	ML/AI	Cloud)	ML/AI
Panel A: General Skills					
Baseline degree share	0.301*** (0.0507)	0.584*** (0.116)	0.328*** (0.126)	0.438*** (0.125)	0.239* (0.124)
London	0.157 (0.227)	1.675*** (0.404)	1.848*** (0.286)	2.412*** (0.434)	2.338*** (0.346)
Share Superfast Broadband				1.074*** (0.351)	0.709*** (0.227)
PC per Employees (adj win std)				0.132 (0.0888)	0.0541 (0.0589)
Controls	Y	Y	Y	Y	Y
Observations	206	206	206	206	206
Panel B: General and STEM Skills					
Baseline degree share	0.150** (0.0650)	0.259* (0.156)	0.0280 (0.106)	0.276* (0.149)	0.0317 (0.102)
Baseline STEM skills share	0.274*** (0.0838)	0.407*** (0.125)	0.375*** (0.118)	0.270* (0.148)	0.346** (0.144)
London	0.501** (0.252)	2.260*** (0.383)	2.388*** (0.374)	2.568*** (0.416)	2.538*** (0.384)
Share Superfast Broadband				0.730* (0.440)	0.268 (0.262)
PC per Employees (adj win std)				0.0617 (0.0801)	-0.0355 (0.0680)
Controls	Y	Y	Y	Y	Y
Observations	206	206	206	206	206

Source: ONS, Ofcom, Lightcast, Harte-Hankes. Notes: Outcome variables are PC per employee data is from Hart-Hanks (column 1) and Cloud and ML/AI Adoption rates per 1,000 vacancies from LightCast (columns 2 to 5). PC per employees is winsorised at the top and bottom 1% and then adjusted by removing the effects of employee number (8 bins) interacted with industry (3-digit SIC codes) and year. Cloud and ML/AI adoption is adjusted by removing the effect of the interaction of vacancy count (8 bins) with industry (3-digit SIC codes), separated by year. Baseline = 1991 for PC and 2011 for Cloud and ML/AI. Broadband is the share of super-fast broadband coverage in 2011. Other controls are measured at corresponding baseline and include population density, the share of employed working in manufacturing, the unemployment rate, a dummy variable for Scotland and a dummy variable for Wales. Each regression is weighted using the baseline population in the TTWA. Robust standard errors in parentheses. * p < 0.1, ** p < 0.05, *** p < 0.01.

Table 3: Firm Technology adoption (balanced panel)

	Cloud Computing			ML/AI		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Extensive margin						
% STEM vac. (2012)	0.859*** (0.0846)		0.706*** (0.0839)	0.651*** (0.102)		0.560*** (0.0974)
% High skill vac. (2012)		0.383*** (0.0447)	0.213*** (0.0430)		0.225*** (0.0378)	0.0900** (0.0362)
% Middle skill vac. (2012)		0.0321 (0.0446)	0.0353 (0.0415)		-0.0307 (0.0410)	-0.0282 (0.0407)
Controls & FE	Y	Y	Y	Y	Y	Y
Observations	16695	16695	16695	16695	16695	16695
ymean			0.344			0.174
Panel B: Intensive margin						
% STEM vac. (2012)	0.0715*** (0.0109)		0.0656*** (0.0111)	0.0214*** (0.00393)		0.0195*** (0.00422)
% High skill vac. (2012)		0.0253*** (0.00465)	0.00952*** (0.00304)		0.00600*** (0.00122)	0.00133 (0.00141)
% Middle skill vac. (2012)		0.00270 (0.00395)	0.00299 (0.00382)		-0.00139 (0.00128)	-0.00131 (0.00128)
Controls & FE	Y	Y	Y	Y	Y	Y
Observations	16695	16695	16695	16695	16695	16695
ymean			0.0129			0.00374

Source: ONS, Lightcast, Harte-Hankes. Notes: All regressions include year fixed effects and control for 8 bins of firm size (proxied by number of job posts in 2012). Initial skill level variables (high, middle and labour skill) are constructed using the 1-digit SOC code with the reference group being low-skill vacancies. Standard errors clustered at 2-digit SIC code level. * p <0.1, ** p <0.05, *** p <0.01.

Appendix

Appendix A

A.1 Construction of cloud and ML/AI measures

A.1.1 Measurement of disruptive technologies

We adopt the BHKLT (2021) approach to technology measurement. BHKLT combines earnings call and patent text in order to identify new, economically significant technologies that could be considered ‘disruptive’.

The main features of the BHKLT design are as follows. The authors start their analysis with a dataset of 17 million unique bigrams extracted from the text (after pre-processing) of the nearly three million utility patents awarded by USPTO to US assignees or inventors between 1976 and 2016. To target ‘novel’ inventions, they retain 1.5 million ‘technical’ bigrams that do not appear in the common vocabulary of everyday English up to 1970 using a historical corpus.¹⁰ To target ‘influential’ inventions, they keep the top 35,063 technical bigrams based on patent citation, considering the heterogeneity of citation across technology classes and over time.¹¹

To target ‘economically important’ inventions, they narrow down to 2,181 technical bigrams that appear in at least 100 earning conference calls between 2002 and 2019. The final set of 305 bigrams covers ‘disruptive’ inventions that increased their popularity in earning conference calls by at least ten times during the sample period.¹²

In the next step, they group the 305 bigrams into 29 technologies using a supervised approach. Firstly, they manually read all bigrams and filter out those that refer to problems

¹⁰The Corpus of Historical American English.

¹¹To do so, they use a threshold of 1,000 to shortlist key technical bigrams based on their citations normalized by the technology-by-year mean citation.

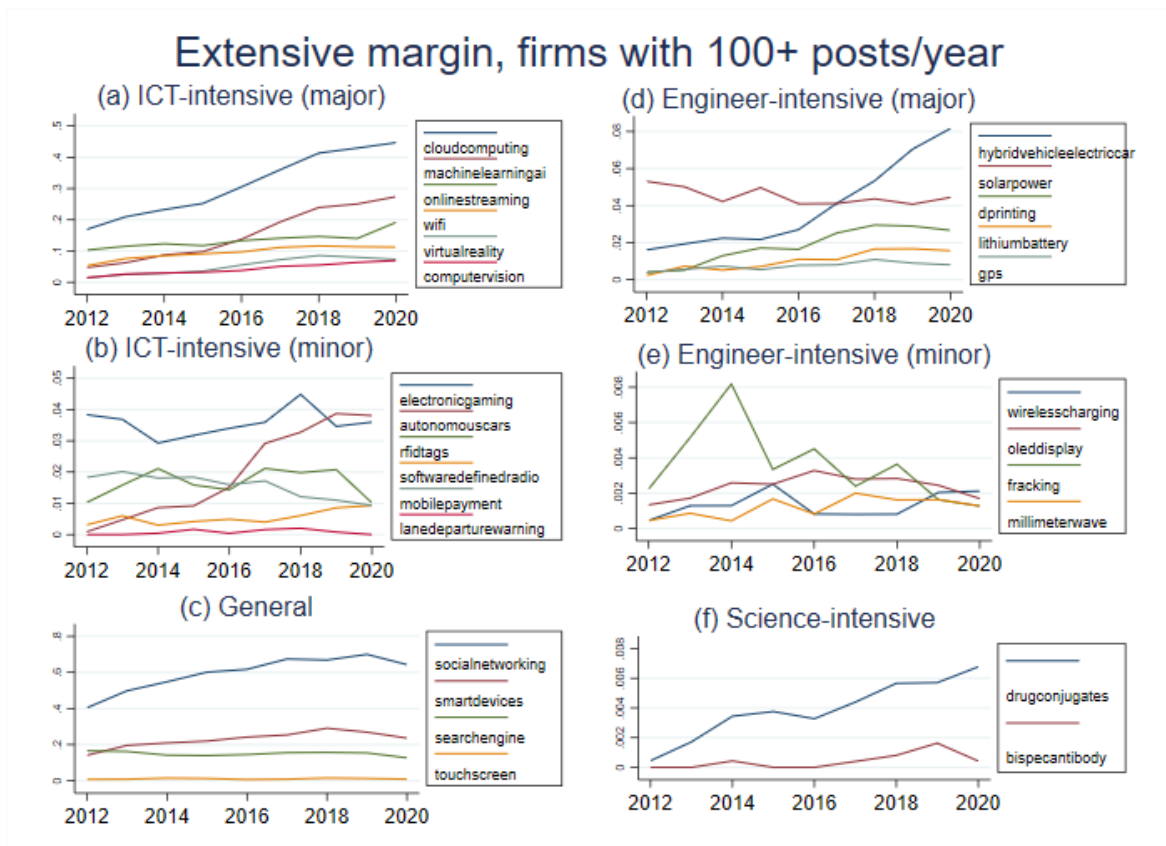
¹²These include bigrams that never mentioned in the initial year (2002).

(such as power outages) rather than solutions, old technologies that enjoyed renewed interest (such as smart electricity grids), vague concepts, or multiple innovations. This produces 105 bigrams, which they then put into 29 groups based on definitions sourced from Wikipedia.

To address the concern that business vocabulary may differ from technical vocabulary, they extend the list of bigrams that define technology groups using their own judgment, based on the most similar bigrams recommended by an embedding vector algorithm. Finally, to ensure these bigrams are compatible with the language of job postings, they use human auditing on a random sample of job postings and shortlisted those that help unambiguously detect the technology in 80% of cases or more. After this process, they are left with 221 audited bigrams that refer to 29 disruptive technologies.

As Section 2 of the paper explains, our analysis focuses on two of these disruptive second wave technologies, cloud and ML/AI. Figure [A1](#) shows these are clearly the most ‘mature’ or widely diffused technologies out of the 27 that are detected in the UK data.

Figure A1: Extensive margin of new technologies



Source: Lightcast. Notes: This figure shows the share of firms who have adopted a technology according to our vacancy-based measure. A discrete 0-1 indicator is defined to flag whether a firm posts vacancies related to a given technology. The y-axis in each panel therefore represents the share of firms who have ‘adopted’ a technology according to this vacancy-based definition. The firms included here are those with a non missing SIC codes and employer names and an average of at least 100 vacancies across the years that they appear in the data. We have grouped the technologies by applying a k-means clustering algorithm to technology-occupation (4-digit) data. We have additionally split the ICT-intensive and Engineer-intensive group into two (major and minor) for clearer exposition. Similarly, the panels use different scales on the y-axis in order to show the time series trends more clearly.

A.1.2 Validation of cloud and ML/AI technology measures

The validity of our measurements rest on the assumption that hiring activities mentioning specific technologies are a reasonable proxy of the adoption of those technologies in firms. To assess the plausibility of this assumption, we leverage rich ONS firm survey data. The ONS E-Commerce and ICT Activity Survey covers around 11,000 UK businesses across the manufacturing, production, construction and distribution sectors as well as parts of the service sector.¹³ As noted by ONS, results for businesses with 10 or more employees are more robust, so we use these estimates to make a comparison with aggregated measures from our BGT dataset, where we restrict the sample to firms with at least 10 posts per year.

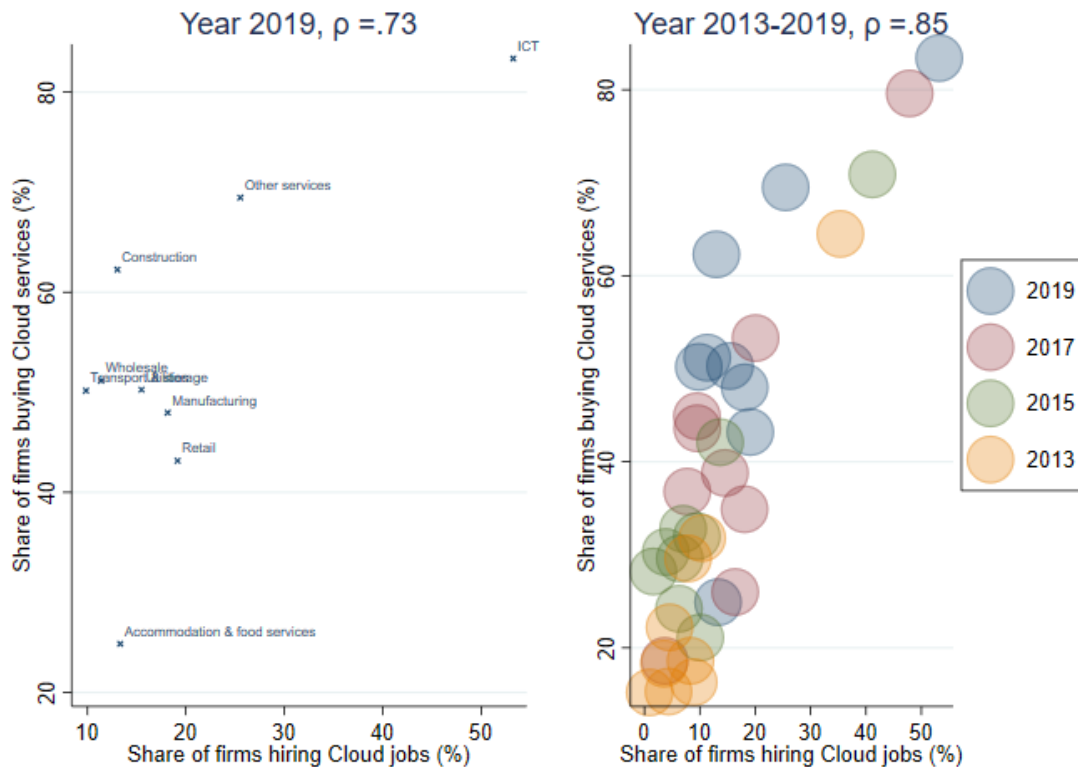
We focus our validation exercises on cloud and ML/AI, the technologies that are the main focus of the paper. Figure A2 compares the share of firms hiring cloud jobs computed from BGT data (x-axis) with the share of firms buying cloud services in ONS surveys (y-axis). The left panel provides a snapshot of 9 broad sectors in 2019 and suggests that two measures are in broad agreement, with a correlation coefficient of 0.73. At the sector level, sectors with a higher share of cloud jobs are indeed more to have a higher share of firms buying cloud services. Similarly, as illustrated in the right panel, when we look at a panel of broad sectors across four survey waves (2013/15/17/19) the correlation coefficient of these measure is 0.85. Overall, the orange scatters (year 2013) move northeast toward blue scatters (year 2019). This means over time, firms of different broad sectors consume more cloud services and simultaneously are hiring for more cloud-related jobs.

It is not possible to perform a directly-equivalent validation exercise for ML/AI overall as ONS data do not include specific estimates for machine learning or AI adoption. The most relevant metric available is perhaps the share of firms using ‘big data analysis’, reported separately for 4 types of data in 2019 including Smart devices or sensor data; Geolocation data, Social media data, and Other big data. They are illustrated in the four panels of Figure A3, which suggests that there is a positive correlation between these big data activities (y-

¹³See <https://www.ons.gov.uk/businessindustryandtrade/itandinternetindustry/bulletins/ecommerceandictactivity/latest> for more details. Again, we assume firms in BGT can be distinguished by employer name and SIC code.

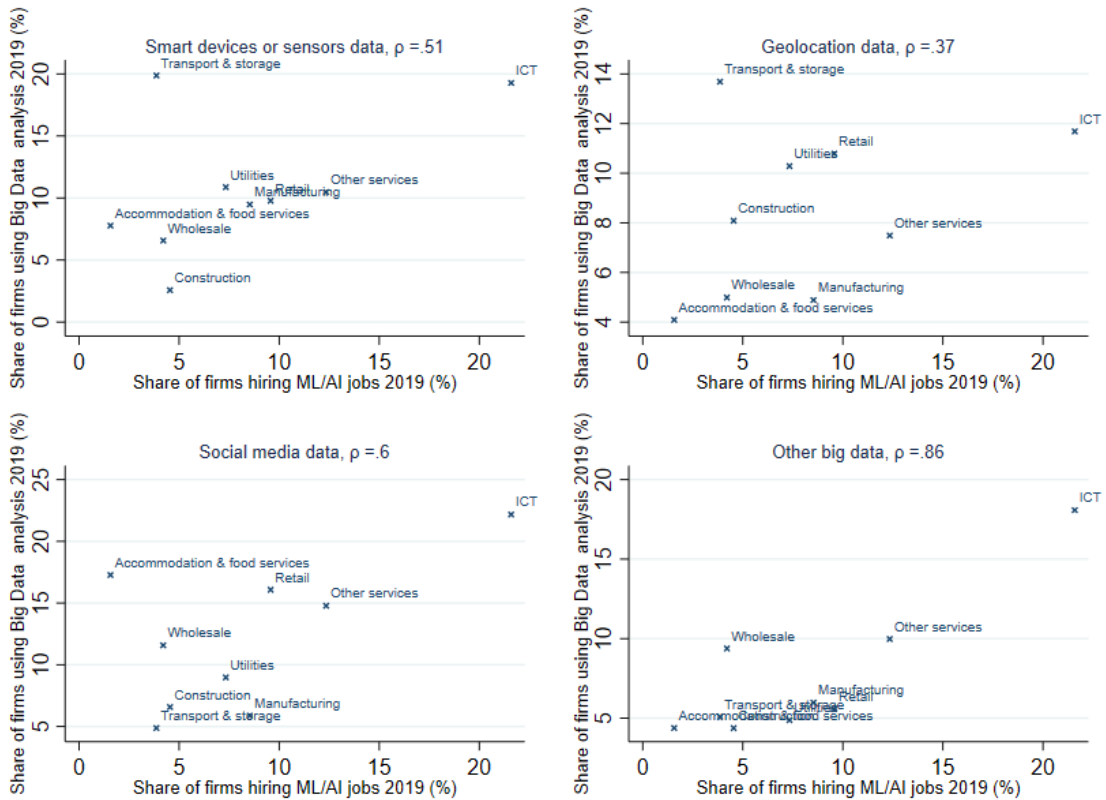
axis) and our measure of MLAI hiring (in x-axis) at the broad sector level. The highest correlation coefficient is 0.86 for Other big data, followed by Social media data (0.6), Smart devices or sensor data (0.51) and Geolocation data (0.37). These sector level correlations provide evidence of a link between hiring activities and the adoption of Cloud Computing and ML/AI in firms.

Figure A2: Cloud adoption: Lightcast vs ONS



Source: ONS, Lightcast. Note: This figure compares the share of firms buying Cloud Services reported in ONS E-commerce and ICT activity annual surveys and the share of firms hiring Cloud Computing jobs computed from Lightcast. The comparison include firms have at least 10 employers or post at least 10 online vacancies. We assume firms in Lightcast can be identified by employer name and SIC code. Correlation coefficients (ρ) are reported for 2019 (Left panel) and 2013/15/17/19 (Right Panel).

Figure A3: ML/AI Adoption: Lightcast vs ONS



Source: ONS, Lightcast. Note: This figure compares the share of Firms Using Big Data Analysis 2019 reported in ONS E-commerce and ICT activity annual surveys and the share of Firms Hiring ML/AI Jobs 2019 computed from Lightcast. The comparison include firms have at least 10 employees or post at least 10 online vacancies. We assume firms in Lightcast can be identified by employer name and SIC code. Correlation coefficients (ρ) are reported for each type of big data.

A.2 Construction of regional data

The datasets that we use have a range of original-source geographical codes that we cross-walk to the TTWA level.

Lightcast: For the location variable in the Lightcast dataset, we use the ‘County/UA’ field in the vacancy-level records. It is worth noting that the information in this field sometimes differs from the names given in official geographical units, leading to the need for some adjustments.

To delineate the Lightcast County/UA field, we match the location name in Lightcast with the geospatial dataset ‘Counties and Unitary Authorities (December 2018) Generalised Clipped Boundaries UK’ from the ONS. We choose the December 2018 version to minimise the need for location name harmonisation and to settle the few discrepancies between Lightcast with the ONS source.¹⁴ It is important to note that we need to pool all related Lightcast county/UA records for areas within London (e.g. City of London, Camden, Hackney) into a single ‘Greater London’ cell, since the location name ‘Greater London’ appears in many Lightcast records.¹⁵ Finally, we overlay the County/UA shapefile with the TTWA boundaries shapefile from ONS to assign LAD/County/UA to TTWA. When a County/UA is spread over several TTWAs, we partition it to each TTWA using the area ratio.¹⁶

Harte-Hanks. Establishments are tagged with physical location postcode in the source data. We directly map establishments’ postcodes to TTWAs using the ONS National Statistics Postcode Lookup (NSPL) crosswalk version November 2021.

Census/NOMIS.¹⁷ We aggregate our Census data to TTWA level, using labour force data from NOMIS. For 2011, we use Outputs Areas (OAs), the smallest geographical unit avail-

¹⁴For example, we replace "Rhondda Cynon Taff" with "Rhondda Cynon Taf," replace "The Vale of Glamorgan" with "Vale of Glamorgan," and combine several Lightcast county/UAs (in Northern Ireland) into one county/UA in the ONS dataset where needed.

¹⁵Similarly, we aggregate "Bournemouth" with "Poole" to reflect the recent establishment of the "Bournemouth, Christchurch and Poole" UA.

¹⁶We pool "London" with "Slough and Heathrow Airport" TTWAs to accommodate the large share of "Greater London" vacancies in Lightcast.

¹⁷NOMIS (National Online Manpower Information System) is a web-based database of labour market and Census statistics from the ONS maintained by Durham University.

able. OAs can be directly nested in TTWAs. In 1991, OAs are not made available at NOMIS, and we therefore leverage variables at the Local Authority District (LAD) level. LAD and TTWA boundaries are not nested, however, and LADs often sit across more than one TTWA. We therefore use counts of postcodes in each LAD and TTWA to make weights, using the ONS National Statistics Postcode Lookup (November 2020).¹⁸ We use these weights to aggregate data from LAD to TTWA level. For example, suppose a TTWA consists of parts of three LADs. The TTWA has 100 postcodes, 60 of which are in LAD A, 30 in LAD B and 10 in LAD C. The relevant LAD weights are 0.6, 0.3 and 0.1, respectively.

¹⁸The NSPL is available at <https://geoportal.statistics.gov.uk/datasets/4df8a1a188e74542aebce164525d7ca9/about>.

A.3 Information on firm and establishment-level panels

A.4 Additional results

Table A1: Firm panel data (2012-2020)

Variables	Sample 1 (Identifiable firms)	Sample 2 (Small firms - unbalanced panel)	Sample 3 (Large firms - unbalanced panel)	Sample 4 (Large firms - balanced panel)
	(1)	(2)	(3)	(4)
Share of firms hiring cloud computing jobs	.0281	.0221	.3155	.3442
Share of cloud computing jobs	.0098	.0098	.0125	.0129
Share of firms hiring ML/AI jobs	.0083	.0052	.1562	.1744
Share of ML/AI jobs	.0019	.0018	.0034	.0037
Share of STEM jobs	.0576	.0568	.0992	.0989
Share of high-skill jobs	.4186	.4157	.5528	.5565
Share of middle-skill jobs	.4353	.4374	.3305	.3302
Share of labour-skill jobs	.1297	.1305	.0941	.0893
Share of missing-SOC jobs	.0165	.0164	.0226	.024
Share of London jobs	.2242	.2236	.2536	.2547
Share of Golden Triangle jobs	.248	.2473	.2813	.2827
Number of unique firms	539073	536173	2900	1855
Observations	1051820	1030293	21527	16695
Share of total vacancies	1	.9795	.0205	.0159
Share of total cloud computing vacancies	1	.3706	.6294	.5563
Share of total ML/AI vacancies	1	.2294	.7706	.7207

Sample 1 in Column (1) describes all the ‘firms’ identified by the dyad of employer name and SIC code. The next Sample 2 then conditions on the ‘small’ firms with an average of less than 100 vacancies per year. Sample 3 is the ‘large’ firm counterpart of firms with an average of 100 or more vacancies. Finally, Sample 4 is the balanced panel of large firms where each firm appears in all years between 2012-2019. The share variables are the unweighted means across firms. STEM share is calculated using the NESTA classification ([Bakhshi et al., 2015](#)). Skill level variables are defined by the first digit SOC codes using the ONS classification including high skill (1-3), medium skill (4-7), labour skill (8-9)

Source: Lightcast. Notes: Sample 1 in Column (1) describes all the ‘firms’ identified by the dyad of employer name and SIC code. The next Sample 2 then conditions on the ‘small’ firms with an average of less than 100 vacancies per year. Sample 3 is the ‘large’ firm counterpart of firms with an average of 100 or more vacancies. Finally, Sample 4 is the balanced panel of large firms where each firm appears in all years between 2012-2019. The share variables are the unweighted means across firms. STEM share is calculated using the NESTA classification ([Bakhshi et al., 2015](#)). Skill level variables are defined by the first digit SOC codes using the ONS classification including high skill (1-3), medium skill (4-7), labour skill (8-9)

Table A2: Technology, lagged general and STEM skills (Unadjusted), comparing waves

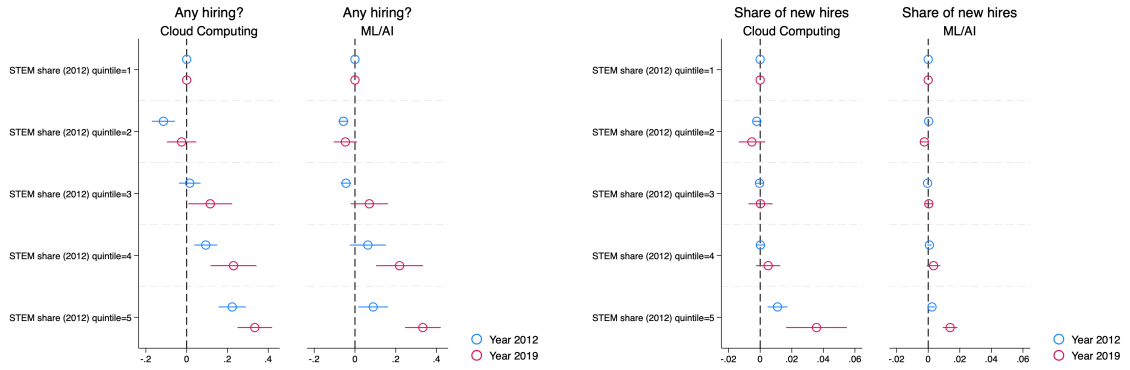
	(1)	(2)	(3)	(4)	(5)
	PC	Cloud	ML/AI	Cloud)	ML/AI
Panel A: General Skills					
Baseline General Skills	0.121*** (0.0301)	0.00538*** (0.000857)	231.6*** (54.76)	0.00428*** (0.000774)	215.0*** (55.01)
London	0.0378 (0.165)	0.00444 (0.00351)	16698.4*** (388.1)	0.0105*** (0.00375)	16843.9*** (411.6)
Share Superfast Broadband				0.00881*** (0.00226)	205.3** (99.44)
PC per employee (adj win std)				0.000691 (0.000679)	-20.09 (27.82)
Controls	Y	Y	Y	Y	Y
Observations	206	206	206	206	206
Panel B: STEM Skills					
Baseline General Skills	0.0549* (0.0300)	0.00243** (0.000983)	169.4*** (56.94)	0.00253*** (0.000941)	169.8*** (57.06)
Baseline STEM Skills	0.120*** (0.0387)	0.00369*** (0.00112)	77.91 (58.00)	0.00292** (0.00125)	75.20 (69.83)
London	0.188 (0.154)	0.00976*** (0.00371)	16810.6*** (449.3)	0.0122*** (0.00386)	16887.2*** (442.8)
Share Superfast Broadband				0.00508* (0.00281)	109.4 (107.8)
PC per employee (adj win std)				-0.0000656 (0.000510)	-39.57 (32.38)
Controls	Y	Y	Y	Y	Y
Observations	206	206	206	206	206

Source: ONS, Ofcom, Lightcast, Harte-Hankes. Notes: This tables replicate Table 2 with outcome variables unadjusted including PC per employee data from Hart-Hanks (column 1) and Cloud and ML/AI Adoption rates per 1,000 vacancies from LightCast (columns 2 to 5). Baseline = 1991 for PC and 2011 for Cloud and ML/AI. Broadband is the share of super-fast broadband coverage in 2011. Other controls are measured at corresponding baseline and include population density, the share of employed working in manufacturing, the unemployment rate, a dummy variable for Scotland and a dummy variable for Wales. Each regression is weighted using the baseline population in the TTWA. Robust standard errors in parentheses. * p <0.1, ** p <0.05, *** p <0.01.

Figure A4: Quintile Specification

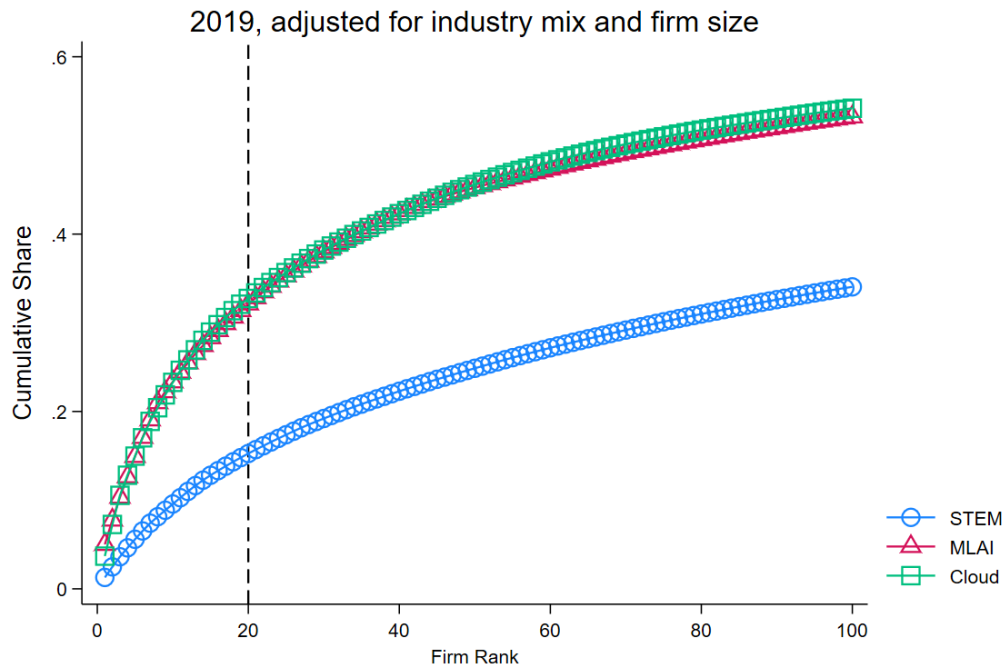
Panel A: Extensive margin

Panel B: Intensive margin



Source: Lightcast. Notes: This figure illustrates the coefficients of firms' initial STEM-intensity (2012) categorized into five quintiles when regressing cloud/ML/AI adoption in 2012/2019 on them, along with other controls (2-digit SIC code, firm size, and the share of high-skill and middle-skill vacancies in 2012). The adoption is measured either as the extensive margin (panel A) or the intensive margin (panel B). Standard errors are clustered at the industry level (2-digit SIC code).

Figure A5: Adjusted technology concentration 2019 [REP]



Source: Lightcast. Notes: Building on Figure 8, this figure plots the cumulative sum of technology-related vacancies according to the the volume-rank of firms. However, we regression-adjust the firm vacancy shares for SIC4 industry and firm size effects before calculating the cumulative sum. Firm size is specified as dummies for 8 quantiles.