AI and Jobs: Evidence from Online Vacancies

Daron Acemoglu
MIT

David Autor
MIT

Jonathon Hazell
Princeton & LSE

Pascual Restrepo
BU

September 2020*

Abstract

We study the impact of AI on labor markets, using establishment level data on vacancies with detailed occupational information comprising the near-universe of online vacancies in the US from 2010 onwards. We classify establishments as "AI exposed" when their workers engage in tasks that are compatible with current capabilities of AI. We document rapid growth in AI related vacancies over 2010-2018 that is not limited to the Information Technology sector and is significantly greater in AI-exposed establishments. The labor market effects of these AI activities are less clear-cut, however. We find some evidence of negative hiring effects at the establishment level, but no discernible impact at the occupation or industry level. We interpret our results to indicate that it is still too soon to see most of the effects of AI in the US labor market, though the decline in non-AI hiring in exposed establishments suggests that AI adoption is likely substituting for humans in a subset of tasks.

Keywords: artificial intelligence, displacement, labor, jobs, tasks, technology, wages.

JEL Classification: J23, O33.

*We thank Bledi Taska for detailed comments and providing access to Burning Glass data, Jose Velarde for expert research assistance, and David Deming and Kadeem Noray for sharing code and data.

1 Introduction

The last decade has witnessed rapid advances in artificial intelligence (AI) based on new machine learning techniques and the availability of big data. The pace of change is expected to increase in the years to come (e.g., Neapolitan and Jiang, 2018, Russell, 2019), and AI applications have already started to impact businesses (e.g., Agarwal, Gans and Goldfarb, 2019).[1] Some commentators see this as a harbinger of a jobless future (e.g., Ford, 2015; West, 2018; Susskind, 2020), while others see the oncoming AI revolution as enriching human productivity and work experience (e.g., McKinsey Global Institute, 2017). The persistence of these contrasting visions is unsurprising given the limited evidence to date on the labor market consequences of AI. There are currently no representative data sets of AI, and hence we lack representative evidence on whether there has indeed been a major increase in AI adoption (as opposed to just talk of AI). Moreover, it is possible to find examples of AI technologies either replacing work or complementing workers, precisely because AI, as a broad technological platform, is capable of doing both, and it is thus partly a matter of societal and business choice how much job displacement AI will create (Acemoglu and Restrepo, 2019b).

This paper studies AI adoption in the US labor market and its implications. Our starting point is that AI adoption can be partially identified from the footprint it leaves at adopting establishments as they hire workers specializing in AI-related activities, such as supervised and unsupervised learning, natural language processing, machine translation, or image recognition. To put this idea into practice, we build an establishment-level data set of AI activity based on the near-universe of U.S. online job vacancy postings from Burning Glass Technologies for the years 2007 and 2010 through 2018. This data set, which has been used in several recent papers, contains detailed information on occupation and skills required for each posted vacancy,

[1] AI is a collection of algorithms that act intelligently by recognizing and responding to their environment. AI algorithms process and identify patterns in vast amounts of unstructured data (for example, speech data, text, or images), which allows them to perceive their environment and take actions to achieve some specific goal.

making it ideal for our purpose.[2]

We start with a task-based perspective, linking the adoption of AI and its possible implications to the task structure of an establishment. This perspective emphasizes the fact that current applications of AI are capable of performing specific tasks, and predicts that firms engaged in those tasks will be the ones to adopt AI systems.[3] To identify the tasks compatible with current capabilities of AI, we use three different but complementary measures: Felten, Raj and Seamans' (2018, 2019) AI Occupational Impact measure; Brynjolfsson, Mitchell and Rock's (2018, 2019) Suitability for Machine Learning (SML) index; and Webb's (2020) AI Exposure score. Each of these measures is computed based on different assumptions and identifies the occupations and tasks being impacted by AI technologies, as we detail below. Given the uncertainties about the exact time path and impact of AI, the fact that some of the measures (in particular Felten et al.) aim to capture complementarities as well as substitutability between AI and workers, and the potential synergies among these measures, we use all three throughout our analysis.[4]

Our first result is that there is a rapid takeoff in AI vacancy postings starting in 2010 and significantly accelerating around 2015-2016. Moreover, consistent with a task-based view of AI, this AI adoption is directed towards establishments with task structures that are compatible with current capabilities of AI. For instance, a one standard deviation increase in our baseline measure of AI exposure (which is approximately the difference in the average AI exposure between administrative & support versus finance industries) is associated with 15% more AI vacancy posting.

[2]A by no means exhaustive list of recent papers that use the Burning Glass data includes Hershbein and Kahn (2016), Azar et al. (2018), Modestino, Shoag, and Ballance (2019), Hazell and Taska (2019) and Deming and Noray (forthcoming).

[3]See Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018, 2019a). We outline the basic features of such a task-based framework and draw some of its implications in the Appendix. This is not the only approach one could take to AI. One could also think of AI as complementing some particular business models (rather than performing specific tasks within those models) or as allowing firms to generate and commercialize new products (see Agarwal, Gans and Goldfarb, 2019, and Bresnahan 2019).

[4]Figure 5 shows that the relationship between mean wage of an occupation and the three AI exposure measures are very different, which is the basis of our claim that each one of these indices captures a different aspect of AI exposure.

The fact that establishments with AI-suitable tasks hire workers into AI positions does not, of course, tell us whether AI is substituting for or complementing the workers engaged in these tasks. To gain insight into this question, we next investigate the consequences of the recent surge in AI for labor demand across establishments.

In principle, AI-exposed establishments may see an increase in hiring, if either: (1) AI complements workers in some tasks, increasing productivity and encouraging more hiring; (2) AI has a significant total factor productivity effect, increasing demand in non-exposed tasks and occupations (Acemoglu and Restrepo, 2019a); or (3) AI creates a competitive advantage for adopting firms, allowing them to expand at the expense of other firms in the market. Alternatively, AI adoption can reduce hiring if many tasks are replaced by AI and the additional hiring in non-automated tasks spurred by AI adoption does not make up for this reduction. Our results consistently show no positive effects of AI exposure on establishment hiring. Rather, we find evidence of lower hiring associated with greater AI exposure in some specifications, though the effect sizes are modest and not robust enough across all of our specifications to allow firm conclusions. This pattern of results, combined with our estimates showing that AI adoption is concentrated in establishments with more AI-exposed tasks, suggests that the recent AI surge is driven in part by task substitution whereby AI automates a subset of tasks formerly performed by labor.[5]

We also supplement these results with additional data to show:

- No significant employment impact on industries that have a task structure that exhibit greater exposure to AI.[6]
- No significant employment effect on occupations that are more exposed to AI.
- No significant effect of AI on the types of skills demanded at the establishment

[5] The negative or zero effects of AI on establishment hiring are consistent with the task displacement impact of AI being somewhat, but most likely imperfectly, offset by productivity gains (as clarified in our model in Appendix A).

[6] In particular, we calculate that even if the negative establishment-level effects are taken at face value, given the prevalence of AI in the average US industry, the employment effects would be too small to be detected at the moment. See Section 5.

level.

Our interpretation of these results is that, despite the notable surge in AI adoption, the impact of AI is still too small relative to the scale of the US labor market to have had first-order impacts on employment patterns — outside of AI hiring itself. Nevertheless, our finding that AI adoption is significantly driven by establishments that have a large fraction of tasks that are AI-suitable and the weak evidence for negative effects on hiring together imply that any positive productivity and complementarity effects from AI are at present small and likely less than its displacement consequences.

Our paper is related to a growing literature on the labor market effects of automation. The early literature in this area investigated the broad trends in terms of wages, employment polarization and wage inequality, and emphasized the inequality implications of automation (e.g., Autor, Katz, and Krueger, 1998; Autor, Levy and Murnane, 2003; Goos and Manning, 2007; Autor and Dorn, 2013; Goos, Manning, and Salomons, 2014; Michaels, Natraj and Van Reenen, 2014; Gregory, Salomons and Zierahn, 2016). More recent work has turned to estimating the impact of specific automation technologies, especially industrial robots, on employment and wages. Acemoglu and Restrepo (2020a) find significant negative effects on employment and wages, as well as blue-collar occupations, in local labor markets exposed to robots. Graetz and Michaels (2019) estimate an increasing productivity in industries adopting more robots, but no clear employment patterns. A more recent line of work has focused on firm-level use of robots and related technologies with mixed findings. For example, Acemoglu, Lelarge and Restrepo (2020) find evidence of positive firm-level effects combined with negative industry-level consequences from robots.[7]

There are fewer works on the effects of AI specifically, though this body of work is growing rapidly. Bessen et al. (2018) conduct a survey of AI startups, which

[7] Other papers using proxies of firm-level usage of robots include Dixon, Hong and Wu (2019), Bonfiglioli et al. (2019), Humlum, (2019). Also, other papers have relied on firm surveys of technology use and investments in automation for sub-sectors of manufacturing or available only for some European countries (see for example, Dinlersoz and Wolf, 2018; Koch, Manuylov and Smolka, 2019; Bessen et al., 2019).

among other things asked them to report the benefits that their products provide to customers. About 75% of AI startups report that their products help their clients make better predictions, manage data better, or provide better products; and 50% of startups report that their products help customers automate routine tasks and reduce labor costs. Grennan and Michaely (2019) study how the use of AI algorithms has affected security analysts. In line with the view that AI might substitute for some workers, they document that analysts are more likely to leave the profession when they cover stocks for which there is abundant data available. Distinct from these papers, we focus on the effects of AI on "AI-using sectors" while excluding the "AI-producing sectors" that are the focus of these works.

Most closely related to our paper are a few recent works also investigating the effects of AI on non-AI employment. Babina et al. (2020) study the relationship between AI adoption and employment and sales at both the firm and industry level. They document that, consistent with Alekseeva et al. (2020), AI investment is stronger among firms with higher cash reserves and higher mark-ups, and among firms with higher R&D intensity, and moreover, that these firms grow relatively more than non-adopters. A key difference between our approach and that of Babina et al. is that we focus on AI suitability based on firms' occupational structures rather than observed AI adoption, and this may explain why we arrive at distinctive results on employment and growth consequences. Also closely related is Deming and Noray (forthcoming), who use Burning Glass data (as we do here) to study the relationship between wages, technical skills, and skills obsolescence. Though their focus is not AI, their work provides strong evidence that Burning Glass data are suitable for detecting changes in job skill requirements at the occupational level, an angle of inquiry we take below. As noted above, our work exploits measures of AI-suitability developed by Felten, Raj and Seamans (2018, 2019), Brynjolfsson, Mitchell and Rock (2018, 2019), and Webb (2020). Our work is largely consistent with Felten, Raj and Seamans (2019), who find no significant relationship between AI-suitability and employment growth at the occupation level, and a positive relationship between AI-suitability and AI vacancy posting at the occupational level. Complementing their work, we evaluate outcomes at the establishment, occupation, and industry level. Our results confirm that AI suitability is not at present associated with more hiring at several levels of aggregation, although we detect some weak negative effects at the establishment level.

The rest of the paper is organized as follows. Section 2 describes the data and Section 3 presents our empirical strategy. Section 4 presents our main results on establishment AI exposure and AI hiring. Section 5 explores the effects of AI on hiring at the establishment, industry and occupation levels and also looks at the relationship between AI and new skills. Section 6 concludes. Appendix A presents a task-based model of AI activity, which is useful in interpreting some of our results. Appendix B includes additional empirical results.

2 Data

In this section, we describe the Burning Glass Technologies (Burning Glass or simply BG) data and document that it is broadly representative of employment and hiring trends across occupations and industries. We then describe our various AI exposure indices and show their distribution across occupations and evolution over time.

2.1 Burning Glass Data

Burning Glass collects data from roughly 40,000 company websites and online job boards, with no more than 5% of vacancies from any one source. Burning Glass then applies a deduplication algorithm, and converts the vacancies into a form amenable to data analysis. The coverage is the near-universe of online vacancies from 2010 onwards in the United States (with somewhat lower coverage in 2007). Our primary sample is data from the start of 2010 until October 2018, though we also make use of data from 2007. The vacancy data contains occupation, industry and region information; firm identifiers; and detailed information on skills required by vacancies, garnered from the text of the job posting.

A key question concerns the representativeness of Burning Glass (BG) data, given that the source of the vacancies is the Internet. Figure 1 shows that BG data tracks the evolution of overall vacancies in the US economy (from JOLTS) reasonably well. The exception is the downturn in BG postings data between 2015 and 2017.[8] Figure 2 shows that, over the 2010-2018 period, the occupational and industry composition in BG is closely aligned with both overall occupation employment shares from the OES and with industry vacancy shares from the BLS Job Openings and Labor Force Turnover Survey (JOLTS).[9] Nevertheless, it is worth remembering that BG data represent vacancy flows, while the OES reports employment stocks, and thus we should not expect the two data sources to align perfectly. Moreover, online vacancy postings tend to overrepresent technical and professional jobs relative to blue collar and personal service jobs (Carnevale et al., 2014).

We make use of Burning Glass' establishment and industry detail. When this information is available from the text of postings, vacancies are assigned a firm name and a location, typically at the city level. We define an establishment of a firm as collection of vacancies pertaining to a firm and commuting zone. Commuting zones are groups of counties with close commuting ties, that likely belong to the same labor market (Tolbert & Sizer, 1996). BG assigns an industry to each vacancy, when this information is available. We classify each firm as belonging to the industry in which it posts the most vacancies over our sample period.

Of particular importance for our paper are BG's skill and detailed occupation coding. Vacancies in BG data contain information on skill requirements, scraped from the text of the vacancy. The skills are organized according to several thousand standardized fields. Groups of related skills are collected together into "skill clusters". Regarding occupation information, over 95% of vacancies are assigned a six-digit

[8] We adjust the quantity of job openings in JOLTS to match the concept of vacancies in Burning Glass, using the approach developed by Carnevale et al. (2014). The difference in concept between JOLTS and Burning Glass vacancies likely accounts for the downturn in BG postings data between 2015 and 2017.

[9] Descriptive statistics and additional information on the BG data are provided in the Appendix.

(SOC) occupation code.[10]

Using this information, we constructed two measures of AI vacancies, narrow and broad. The narrow category includes a selection of skills relating to AI.[11] The broad measure of AI includes skills belonging to the broader skill clusters of Machine Learning and Artificial Intelligence. One concern with our broad AI measure is that it may include various IT functions that are separate from core AI activities. For this reason, we focus on the narrow AI measure in the text and show the robustness of our main results with the broad occupation measure in the Appendix. Figure 3 shows the evolution of postings of narrow and broad AI vacancies in the BG data, and highlights the rapid takeoff around 2016 already mentioned in the Introduction. The second panel of this figure depicts that this takeoff is particularly pronounced in the industry sectors of information; professional and business services; finance; and manufacturing. Still, a sharp uptick is visible in all industries.

Our primary focus is on "AI-using sectors". Therefore in our establishment level analysis, we drop establishments belonging to sectors likely to be producing AI-related products, namely the information sector (NAICS sector 51) and the professional and business services sector (NAICS sector 54). The former includes various information technology industries, likely to be selling AI products. The latter contains industries such as management consultancy, likely to be integrating AI into other industries' production processes.

[10] Six-digit occupation codes are highly granular, including occupations such as pest control worker.

[11] The skills are Machine Learning, Computer Vision, Machine Vision, Deep Learning, Virtual Agents, Image Recognition, Natural Language Processing, Speech Recognition, Pattern Recognition, Object Recognition, Neural Networks, AI ChatBot, Supervised Learning, Text Mining, Support Vector Machines, Unsupervised Learning, Image Processing, Mahout, Recommender Systems, Support Vector Machines (SVM), Random Forests, Latent Semantic Analysis, Sentiment Analysis / Opinion Mining, Latent Dirichlet Allocation, Predictive Models, Kernel Methods, Keras, Gradient boosting, OpenCV, Xgboost, Libsvm, Word2Vec, Chatbot, Machine Translation and Sentiment Classification.

2.2 AI Indices

We study three measures of AI exposure. Each measure is assigned at the 6 digit SOC occupation level. Each measure is designed to capture occupations concentrating in tasks that are compatible with the current capabilities of AI.

The first measure is from Felten et al. (2019). This measure is based on data from the AI Progress Measurement project, from the Electronic Frontier Foundation, starting in 2010. These data identify a set of tasks in which artificial intelligence has made progress over the recent years — for example, image recognition or machine translation. The authors then collect the tasks associated with each occupation, from O*NET data. Finally, occupations are exposed to AI if they specialize in the tasks at which AI has improved since 2010.

The second measure is from Webb (2020). The paper identifies an occupation as exposed to AI, if the occupation's tasks are similar to the tasks at which AI is capable. Similarity is measured from the text of patents in AI, and the text of the task descriptions in O*NET.

The third measure is Suitability for Machine Learning (SML), from Brynjolfsson et al. (2019). The authors construct a detailed, 21-item rubric of tasks that are suitable for machine learning/AI. They then identify AI exposed occupations as those containing tasks suitable for machine learning, once again mapped from O*NET data.

Figure 4 shows the distribution of our three indices by broad occupation categories and by one-digit industry.[12] Figure 5 represents the same information in a different manner, depicting the average AI exposure by occupational wage percentile. The figures confirm that these three measure capture different aspects of the effects of AI. The Felten et al. measure, for example, is particularly high for managers, professionals and office and administrative staff and is very low for service, production and construction workers, capturing the fact that these occupations involve various manual tasks that cannot be performed by algorithms. The Felten at al. measure

[12]The broad occupational categories are those utilized by Autor (2019) and aggregate six-digit occupations into 10, roughly one-digit categories.

is not particularly high in sales occupations and shows a strong positive relationship with occupational wage percentiles (Figure 5). In contrast, Brynjolfsson et al.'s SML measure is high for office and administrative occupations, and for sales occupations, and (perhaps surprisingly) above average for personal services, but is low for professional occupations and most blue-collar and service occupations. This pattern implies that SML has no systematic relationship with occupational wage percentiles. The contrast between the Felten et al. and SML measures may reflect the fact that Felten et al. attempt to capture tasks where AI is potentially complementary to human labor whereas SML attempts to capture tasks that can be accomplished by AI, as the acronym Suitability for Machine Learning implies. Webb's measure, based on textual similarity between US AI patents and job descriptions in the US O*NET database, is far more similar to the Felten et al. measure than the SML measure, with a key difference that Webb detects less AI-suitability than Felten et al. in highly paid professional and managerial occupations.[13]

3 Empirical Strategy

Our empirical strategy is to link measures of AI activity at the establishment level and various establishment-level job posting outcomes to the exposure of an establishment to AI, where AI exposure is measured using the three indices above.

Namely, we estimate the following type of regression model:

$$\Delta y_{i,t-t_0} = \beta AI_{i,t_0} + \mathbf{x}'_{i,t_0}\boldsymbol{\gamma} + \varepsilon_{i,t-t_0}, \tag{1}$$

where i denotes establishment, $\Delta y_{i,t-t_0}$ denotes the change in one of our establishment-level outcomes between 2010-12 and 2016-18, AI_{i,t_0} is one of our three measures of establishment AI exposure, as defined above and measured with data between 2010

[13] Another notable difference is that the Webb index finds very little AI-suitability in either office or sales occupations. Alongside his AI index, Webb (2020) creates a separate software exposure index, pertaining to traditional non-AI software, that detects substantial software-suitability in office, administrative, and sales occupations.

and 2012, and $\mathbf{x}_{i,t_0}$ is a vector of baseline controls, including industry dummies, firm size decile dummies, a dummy for the commuting zone (CZ) in which the establishment is located and, in some specifications, firm fixed effects.[14] Finally, $\varepsilon_{i,t-t_0}$ is an error term representing all omitted factors.

In all of the specifications, our main interest is with the coefficient β. We divide the measure of establishment AI exposure by its standard deviation, weighted by vacancies in 2010-2012. So, β is the change in the outcome variable associated with a standard deviation change in AI exposure.

The two main measures we focus on for $\Delta y_{i,t-t_0}$ are AI vacancies and overall non-AI hiring at the firm level. Our interest in investigating the effects of AI on non-AI activities also makes us focus on a sample excluding sectors 51 and 54 (information technology and professional and business services, which include many "AI producing" firms, that is, businesses that either provide AI services or AI-related consulting services to other firms).

We also aggregate this equation up to the occupational and industry level to explore hiring and other outcomes for more exposed occupations and industries.

4 AI Exposure and AI Vacancies

In this section, we show that AI exposure predicts establishment-level AI activity (proxied by our measure of narrow AI vacancies). Table 1 presents our main results. Panel A of this table shows the relationship between AI exposure based on Felten et al.'s Occupational AI Impact index and growth in AI vacancies (the other two panels are for our other measures of AI exposure). We estimate regression models based on (1), with the left-hand side variable defined as the change in the inverse hyperbolic sine of AI vacancies between 2010-12 and 2016-18.[15] Throughout, we

[14] We pool 2010-12 data and, separately, 2016-18 data to improve precision.

[15] The inverse hyperbolic sine transformation is given by:

$$\ln\left(x + \sqrt{(x^2+1)}\right).$$

focus on weighted specifications, using baseline vacancies as weights. As noted above, our sample excludes sectors 51 and 54 (information technology and professional and business services), which are "producers" of AI. We report heteroscedasticity robust standard errors that allow for arbitrary cross-sectional and over-time correlation across the establishments of each firm.

Column 1 is our most parsimonious specification and includes no covariates, thus depicting the unconditional bivariate relationship (inclusive of an intercept). The coefficient estimate in Panel A of $\beta = 15.96$ is precisely estimated (standard error = 1.73) and shows a sizable association between AI exposure and AI vacancies. Our estimate in this column implies that a one standard deviation increase in AI exposure (which also corresponds to the difference between the administrative & support and finance industries) is associated with a 15.96% increase in AI vacancies.[16]

The remaining columns explore the robustness of this predictive relationship. Column 2 controls for firm size decile and commuting zone fixed effects. The coefficient estimate of AI exposure declines slightly to 13.82 but is also more precisely estimated. Column 3 additionally adds three-digit (NAICS) industry fixed effects as controls. Reflecting the sizable variation in AI exposure across industries, shown in Figure 4 above, these controls are more important for our regression, and they reduce the magnitude of our estimate by about a third, to 9.19, but the standard error of the estimate also declines (to 1.21).

Column 4 goes one step further and includes a full set of firm fixed effects, so that now the comparison is between two establishments of the same firm that differ in their AI exposure as measured by their baseline vacancy postings across detailed occupations in 2010-12. Reassuringly, the estimate of β is quite similar to the bivariate relationship reported in column 1, 16.53, albeit slightly less precise, since all of the cross-firm variation is purged in these specifications.

For small values of x, this approximates a proportional change, but is well defined when $x = 0$, which is a frequent occurrence in our sample of establishments.

[16] The partial R^2 for our AI exposure measure in this regression is 2.5%. Though this appears small, given the vast cross-establishment variation in AI adoption, it is a non-trivial amount. For comparison, the partial R^2 of three-digit industry fixed effects in column 3 is 5.2%.

Figure 4 documented significant differences in AI exposure across occupations. This raises the concern that our results so far may be confounded by secular trends across broad occupational categories. Columns 5 and 6 add to the specifications in columns 3 and 4 the baseline shares of vacancies that are in sales and administration, two of the broad occupational categories that have been in decline for other reasons (e.g., Autor and Dorn, 2013). These controls have no major impact on the estimate in either column, which remain, respectively, at 9.75 (standard error = 1.20) and 16.87 (standard error = 1.86).[17]

Figure B1 in Appendix B shows the specification from column 5 in the form of a bin scatterplot (where each bin represents about 50,000 establishments). The relationship between AI exposure and AI vacancy postings is fairly close to linear across the distribution and is not driven by outlier observations.

Panels B and C of Table 1 repeat the same regressions using the Webb and SML measures of AI exposure. The results with the Webb measure, reported in Panel B, are very similar in the first four specifications, though they do not prove robust to controls for the baseline shares of sales and administration vacancies in columns 5 and 6. Quantitatively, the estimate in column 1, 6.59, implies that a one standard deviation increase in AI exposure is associated with a 6.59% increase in AI vacancies.

The results are more nuanced with the SML measure. There is a positive association between the SML measure and AI vacancy growth without any covariates, but when three-digit industry fixed effects are included, this relationship becomes negative. This is because, as noted above, sales and administration occupations have a high SML score and these broad occupational groups are negatively associated with the adoption of AI across establishments. Once we control for the baseline shares of these occupations in columns 5 and 6, the positive relationship in column 1 is restored. In light of the seemingly distinct components of AI exposure that these three measures are capturing (recall Figure 4 and Figure B2 in Appendix B), we find

[17] Table B5 in Appendix B shows that the results in Panel A are also robust if we include the baseline shares of ten broad occupational categories. For example, the coefficients in the specifications that parallel columns 5 and 6 are, respectively, 7.24 (standard error = 1.44) and 13.70 (standard error = 2.12). However, some of the results in other panels are sensitive to these controls.

it reassuring that they paint a broadly comparable picture of AI hiring occurring disproportionately at establishments that have a significant fraction of tasks that are AI-suitable.

In sum, our various task-based AI exposure measures robustly predict establishment-level growth in AI activity. The data point to a recent surge in AI-related hiring, and our regression evidence indicates that much of this is being directed towards establishments whose task structures enable the use of AI. This evidence is interesting in and of itself, since it suggests that an important component of AI activity is related to the types of tasks performed in an establishment. It does not, however, preclude the possibility that AI activity is fueled in part by other drivers, such as development of new products or complementing existing business models.

We next turn to analyzing the broader labor market implications of this upsurge in AI activity.

5 AI and Jobs

This section reports our results on the effects of our measures of AI exposure on broader labor market outcomes.

5.1 The Effects of AI on Non-AI Hiring

Table 2 turns to the hiring (vacancy posting) implications of AI exposure. The structure of the table is identical to that of Table 1, except that the left-hand side variable is now the change in the inverse hyperbolic sine of total non-AI vacancies (and there are two extra columns, which we describe below). This measure is chosen so as to focus on the effects of AI activity on establishment hiring exclusive of the mechanical (or direct) affect on AI hiring itself.

The picture that emerges from Table 2 is more nuanced than our results on AI vacancies. In Panel A, where we focus on Felten et al.'s measure, we see a fairly consistent negative association between AI exposure and subsequent non-AI hiring.

The estimate in column 1 is -13.80 (standard error = 4.22), indicating that a one standard deviation increase in AI exposure is associated with a 13.80% decline in overall non-AI vacancies. This is a sizable effect, but is less precisely estimated than our results in Table 1.[18]

This coefficient estimate remains similar when we control for firm size deciles, commuting zone controls and the three-digit industry fixed effects, but declines to a third of its size when we include firm fixed effects (-4.81, standard error = 1.44). This specification thus suggests a more modest, 4.81% decline in non-AI vacancies associated with a one standard deviation increase in AI exposure.

The relationship between AI exposure and non-AI vacancies remains comparable when we include the baseline shares of sales and administration occupations in columns 5 and 6: -12.42 (standard error = 4.01) and -4.04 (standard error 1.47), respectively. We also investigated whether these estimates are driven by establishments that posted jobs in 2010 and then stopped posting in 2018 (which may reflect true zero vacancy postings or the fact that they are no longer in the data set). Columns 7 and 8, therefore, limit the sample to establishments that posted in 2018. The estimates are now somewhat smaller, but still negative and marginally statistically significant: -8.38 (standard error = 3.46) in column 7, with three-digit industry fixed effects and -3.56 (standard error 1.86) in column 8, with firm fixed effects.[19]

Panel B turns to Webb's measure of AI exposure. The pattern is broadly similar to the one we see in Panel A, but somewhat less stable. The coefficient estimate without any covariates in column 1, -17.24 (standard error = 3.72), declines substantially to -2.22 (standard error = 0.93) in column 4 when we control for firm fixed effects, and is inconsistent in sign and magnitude in columns 5-8. When we use the SML

[18]Even if we take this estimate at face value, this should not be read as a 13.80% decline in employment, since vacancies are a flow variable, while employment is a stock variable. For example, if workers separate from firms at a rate of 10% per year, a one-year increase in postings of 13.83% raises employment by only 1.38%.

[19]Since there are positive effect on AI vacancies and some negative effects on non-AI vacancies, one question is whether the impact on total vacancies (including AI hiring) is also negative. We show in Table B6 in Appendix B that the answer is yes, which is not surprising since AI vacancies are a tiny share of total vacancies.

measure in Panel C, there is no consistent evidence for a negative impact from AI exposure (though there is a negative and statistically significant estimate in column 6), and in some specifications there is a positive association between AI exposure and non-AI vacancies.

In Table B7 in Appendix B we show results with average establishment size weights (rather than baseline establishment size weights as in Table 2). In this case, the negative estimates are somewhat smaller and less consistently negative. These results, combined with those reported in the next two subsections, which do not show statistically significant effects of AI exposure at the occupation and industry level, make us cautious in interpreting the negative establishment-level relationship between AI exposure and vacancy posting seen when using Felten et al.'s and Webb's measures.[20]

Summing the evidence, we find no support for the hypothesis that AI exposure is associated with greater hiring at the establishment level. Indeed, several of our specifications show negative effects, though they are not always fully robust to the entire range of AI exposure measures, specifications and controls we consider. We therefore view these results as merely suggestive of a negative impact from AI activity. Mostly, it is still too soon to see the full effects of AI adoption on non-AI hiring.

Even if tentative, these findings are consistent with the task-based approach to AI and are most likely driven by the fact that AI is substituting for some tasks previously performed by humans.[21] Recall, in particular, that if AI was complementing workers in tasks they already performed, we should see more hiring at more exposed establishments. Negative impacts of AI, on the other hand, would be consistent with task substitution leading to displacement of workers that is not compensated

[20] In Table B3 in Appendix B, we also look at specifications with the change in "at-risk" vacancies (those with high AI exposure) on the left-hand side. These specifications show consistent negative effects, but they suffer from potential mean reversion (since high AI exposure means that a high fraction of baseline vacancies were in the at-risk category, and an establishment that posts many vacancies in such areas in a given year may then post fewer of them in subsequent years). For this reason, we put little stock in these estimates.

[21] We caveat once again that our strategy does not identify AI adoption unrelated to baseline task structure. Such AI activity may be ongoing, and its effects may differ from those estimated here.

by productivity gains from AI (see Appendix A), while zero impacts of AI would be consistent either with the effects of AI still not being detectable in the data or task substitution impacts being partially balanced out by modest productivity gains. Our results thus weigh against significant complementarities or large productivity effects expanding employment in more exposed establishments and suggest at least the possibility of non-trivial displacement effects.

5.2 AI and Industry Employment

Associated with the surge in AI activity, there may also be a major industry-level reorganization. For example, Acemoglu, Lelarge and Restrepo (2020) find that robot-adopting French manufacturing firms expand, but this is at the expense of their competitors, and overall industry employment contracts significantly when firms adopt more robots. To investigate whether more exposed industries are contracting (or expanding), we aggregate our exposure measure to three-digit industries, and then use County Business Patterns data (CBP) between 2000 and 2016.[22]

The results are reported in the first three columns of Table 3, which again has three panels, one for each of our AI exposure measure. In these regressions, the observations are at the industry by commuting zone level, and throughout we include industry fixed effects, commuting zone fixed effects and baseline occupational shares in sales and administration. We again drop sectors 51 and 54 (as well as unclassified industries). The standard errors are are robust against heteroscedasticity and correlation within commuting zones.

Columns 1 and 2 look at employment before the major AI advances, 2003-2007 and 2007-2010, respectively. They show that industry AI exposure in 2010 does not predict differential employment behavior before 2010, meaning that three-digit industries with different levels of AI exposure are, roughly, on parallel trends before

[22]One important caveat is that because of changes to reporting in CBP after 2016 that create many incompatibilities, our sample here stops in 2016, and thus excludes the last several years of rapid AI expansion. In processing the CBP data, we use the harmonization and imputation procedures developed by Fabian Eckert, Teresa Fort, Peter Schott, and Natalie Yang, available at http://fpeckert.me/cbp/.

the pickup in AI activity in the late 2010s. Stated simply, industry AI exposure in 2010 does not predict differential employment behavior before 2010. But we do not see consistent positive or negative effects associated with AI exposure after 2010 either. For example, the estimate in column 3, which is for 2010-2016, is -0.05 (standard error = 0.08). The point estimate implies very small predicted effects from industry AI exposure: a one standard deviation increase in industry AI exposure should be associated with a tiny 0.049% decline in industry employment.[23]

Panels B and C of the table repeat this exercise using the Webb and SML measures of AI exposure in place of the Felten et al. measure. Again, we do not find a precise or stable relationship between these AI measures and industry-level employment growth, either before the current wave of AI (i.e., prior to 2010) or thereafter.

The evidence does not therefore indicatemajor industry-level effects from AI adoption so far. This may again be because it is too soon to see the impact of AI activity on industry reorganization or growth. It may also reflect the fact that much of the effects of AI will take place within industries.

5.3 Employment and Wages in AI-Exposed Occupations

Columns 4-9 of Table 3 explore variation across occupations, in particular, assessing whether more exposed occupations exhibit differential employment or wage trends after the onset of rapid AI hiring in the US. We use occupational employment and wage information from BLS' Occupational Employment Statistics (OES) data.

The observations in this table are at the six-digit occupation level, and we again exclude sectors 51 and 54, so that the dependent variable is the sum of employment in a six-digit occupation across all industries except sectors 51 and 54. In all columns, we control for three-digit occupation fixed effects and use baseline employment as weights. The standard errors are robust against heteroscedasticity. In columns 4 through 6, the dependent variable is occupational employment, while in columns 7

[23]Such small effects may not be completely surprising. Even if we take our establishment-level results at face value, they should not imply large industry-level impacts, since most establishments in an average industry has zero or very low AI exposure.

though 9 it is occupational (log) wages.

The results for employment and wage growth using each of the AI exposure measures appear to corroborate our industry-level results: we detect no differential employment or wage behavior in more affected occupations after 2010.[24] For example, the estimate in column 6, which is for 2010-2018, shows that a one standard deviation increase in AI exposure at the occupational level is associated with again a very small, 0.005%, increase in occupational employment. We see a negative and statistically significant effect on wages in column 9 for the same time period, but there is no similar result for employment or wages with the other measures.

In summary, we find no robust effects of AI activity on more AI-exposed occupations, which is consistent with the interpretation that it is premature to detect the effects of AI exposure on hiring. This pattern is also consistent with the results reported in Felten et al. (2020) relating occupation-level AI exposure (using the Felten et al. measure, logically) and occupational employment growth.

5.4 AI and New Skills

Finally, we explore whether AI is associated with changes in the nature of skills used in an occupation. Deming and Noray (forthcoming) have documented such changes associated with broader IT-related activity at the establishment level. An important question is whether similar or even more major changes are being brought about by AI.

To investigate this question, we follow Deming and Noray (forthcoming) and measure the gross change in skills demanded in each occupation as:

$$\text{gross skill change}_{o,t_2,t_1} = \sum_{s=1}^{S} \text{abs}\left[\left(\frac{\text{skill}^s_{o,t_2}}{\text{vacancies}_{o,t_2}}\right) - \left(\frac{\text{skill}^s_{o,t_1}}{\text{vacancies}_{o,t_1}}\right)\right].$$

[24] Differently from our industry results, with Felten et al.'s measure, there is a significantly faster increase in the employment of more exposed occupations before 2007, which may reflect fast expansion in some IT-related occupations which have high Felten et al.'s scores — or it may be the result of random variation given the number of point estimates report this table.

Here, $\text{skill}^s_{o,t}$ is the number of times skill s is posted in six-digit occupation o in year t, divided by the total number of vacancies posted in the occupation. This formula measures changes in the frequency with which skills appear in job postings for a particular occupation, and would capture changes resulting both from the increased importance of new skills (or skills that were not common in an occupation) or the obsolescence of previously common skills. We calculate this measure for all six-digit occupations, for all vacancies outside the information and business services sectors, for 2007-2010, 2010-2014 and 2014-2018.

Table 4 reports regressions linking the gross skill change measure to the occupation AI exposure measures. As in our previous occupation level specifications, all of our specifications in this table control for three-digit occupation fixed effects and use baseline employment weights. The standard errors are robust against heteroskedasticity. As before, the three panels correspond to the three measures of AI exposure, while different columns correspond to different time periods. The measures of occupation AI exposures are, as usual, standardized (divided by the weighted standard deviation) across occupations.

Overall, we do not find any significant relationship between AI and changes in the skills required in exposed occupations. Although for the 2007-2010 period, the correlation between occupation AI exposure and occupation net skill change is positive and significant, this is unlikely to be related to AI, as there was very little core AI activity before 2010, and most of the rapid increase does not take place until the second half of the 2010s. In all other specifications, the estimates are small and insignificant.[25]

These results suggest that, to date, AI adoption is not causing a considerable transformation of exposed occupations in terms of their skill requirements.

[25] In additional exercises not reported here, we decomposed the measure of gross changes into positive changes (new or uncommon skills gaining importance) and negative changes (existing skills becoming obsolete). We did not find any significant relationship between AI and these two components of the gross change in skills demanded by occupation.

6 Conclusion

There is much excitement and quite a bit of apprehension about AI and its labor market effects. In this paper, we documented that there is a recent surge in AI activity that is driven by establishments that engage in tasks that are suitable for AI-related technologies. There is some evidence of a slowdown in vacancy posting in the most AI-exposed establishments, but this evidence is not as robust as the association between AI exposure and AI vacancy postings.

We also do not find any effects of AI exposure at the occupation or industry level, nor do we detect any association between AI activity and new skills demanded in vacancies or systematic skill redundancies. Our interpretation of these findings is that it is mostly too soon to see the impact of AI on labor markets, though the existing evidence is consistent with a task-based approach, since AI activity is strongly directed towards firms with a high fraction of tasks that are exposed to AI and there is some evidence of workers being displaced from these tasks.

References

Acemoglu, Daron, and David Autor (2011) "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics*. Vol. 4. Elsevier, 1043–1171.

Acemoglu, Daron, and Pascual Restrepo (2018) "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.? *American Economic Review*, 108(6), 1488?-1542.

Acemoglu, Daron, and Pascual Restrepo (2019a) "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33(2): 3–30.

Acemoglu, Daron, and Pascual Restrepo (2019b) "The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand." *Cambridge Journal of Regions*, Economy and Society 13(1): 25–35.

Acemoglu, Daron, and Pascual Restrepo (2020a) "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128(6): 2188–2244.

Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo (2020) "Competing with Robots: Firm-Level Evidence from France." *AEA Papers and Proceedings*. Vol. 110.

Alekseeva, Liudmila, Jose Azar, Mireia Gine, Sampsa Samila, and Bledi Taska (2020) "The Demand for AI Skills in the Labor Market." CEPR Discussion Paper No. DP14320.

Autor, David, Lawrence Katz, and Alan Krueger (1998) "Computing Inequality: Have Computers Changed the Labor Market?" *The Quarterly Journal of Economics* 113(4): 1169–1213.

Autor, David, Frank Levy, and Richard Murnane (2003) "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118(4): 1279–1333.

Autor, David and David Dorn (2013) "The Growth of Low-skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*

103(5): 1553–97.

Autor, David (2019) "Work of the Past, Work of the Future." *AEA Papers and Proceedings* 109: 1?32.

Agarwal, Ajay, Joshua S. Gans and Avi Goldfarb (2018) *Prediction Machines: The Simple Economics of Artificial Intelligence.* Harvard Business Review Press, Cambridge

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska (2020) "Concentration in US Labor Markets: Evidence from Online Vacancy Data." *Labour Economics*, 101886.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson (2020) "Artificial Intelligence, Firm Growth, and Industry Concentration." Unpublished manuscript, Columbia University.

Bessen, James E., Stephen Michael Impink, Lydia Reichensperger, and Robert Seamans (2018) "The Business of AI Startups." Boston University, School of Law, Law and Economics Research Paper 18-28.

Bessen, James, Maarten Goos, Anna Salomons and Wiljan Van der Berge (2019) "Automatic Reaction-What Happens to Workers at Firms that Automate?" Unpublished manuscript. Boston University.

Bonfiglioli, Alessandra, Rosario Crino, Harald Fadinger, and Gino Gancia (2019) "Robot Imports and Firm Level Outcomes" Unpublished manuscript. Queen Mary University of London.

Bresnahan, Timothy (2019) "Artificial Intelligence Technologies and Aggregate Growth Prospects." Unpublished manuscript, Stanford.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018) "What can Machines Learn, and What Does it Mean for Occupations and the Economy?" *AEA Papers and Proceedings*. Vol. 108.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2019) "Machine Learning and Occupational Change." Unpublished manuscript, MIT.

Carnevale, Anthony, Tamara Jayasundera and Dmitri Repnikov (2014) "Understanding Job Ads Data: A Technical Report." Center on Education and the

Workforce McCourt School of Public Policy Georgetown University.

Deming, David J., and Kadeem Noray (forthcoming) "Earnings Dynamics, Changing Job Skills, and STEM Careers." *Quarterly Journal of Economics.*

Dixon, Jay, Bryan Hong and Lynn Wu (2019) "The Employment Consequences of Robots: Firm-level Evidence" Unpublished manuscript. Statistics Canada.

Felten, Edward, Manav Raj, and Robert Seamans (2018) "A Method to Link Advances in Artificial Intelligence to Occupational Abilities." *AEA Papers and Proceedings.* Vol. 108.

Felten, Edward, Manav Raj, and Robert Seamans (2019) "The Effect of Artificial Intelligence on Human Labor: An Ability-based Approach." *Academy of Management Proceedings* Vol. 1.

Ford, Martin (2015) *The Rise of the Robots.* Basic Books, New York.

Graetz, Georg, and Guy Michaels (2018) "Robots at Work." *The Review of Economics and Statistics* 100(5): 753–768.

Grennan, Jillian, and Roni Michaely (2019) "Artificial Intelligence and High-Skilled Work: Evidence from Analysts." Unpublished manuscript, Duke University.

Goos, Maarten, and Alan Manning (2007) "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *The Review of Economics and Statistics* 89(1): 118–133.

Goos, Maarten, Alan Manning, and Anna Salomons (2014) "Explaining Job Polarization: Routine-biased Technological Change and Offshoring." *American Economic Review* 104(8): 2509–26.

Hazell, Jonathon and Bledi Taska (2019) "Downward Rigidity in the Wage for New Hires" Unpublished manuscript, MIT.

Hershbein, Brad, and Lisa B. Kahn (2018) "Do Recessions Accelerate Routine-biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108(7): 1737–72.

Humlum, Anders (2019) "Robot Adoption and Labor Market Dynamics" Unpublished manuscript. Princeton University.

Koch, Michael, Ilya Manuylov and Marcel Smolka (2019) "Robots and Firms" Unpublished manuscript. Aarhus University.

McKinsey Global Institute (2017) "Artificial Intelligence: The Next Digital Frontier?" Discussion Paper.

Michaels, Guy, Ashwini Natraj, and John Van Reenen (2014) "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-five Years." *Review of Economics and Statistics* 96(1): 60-77.

Neapolitan, Richard E. and Xia Jiang (2018) *Artificial Intelligence: With an Introduction to Machine Learning.* Chapman and Hall/CRC, Second Edition

Russell, Stuart (2019) *Human Compatible: Artificial Intelligence and the Problem of Control.* Penguin.

Susskind, Daniel (2020) *A World Without Work: Technology, Automation and how We Should Respond.* Penguin, UK.

Tolbert, Charles and Molly Sizer (1996) "US Commuting Zones and Labor Market Areas: A 1990 update." ERS Staff Paper Number 9614. Washington, DC: Economic Research Service.

Webb, Michael (2020) "The Impact of Artificial Intelligence on the Labor Market." Unpublished manuscript, Stanford.

West, Darrell M (2018) *The Future of Work: Robots, AI, and Automation.* Brookings Institution Press.

Figure 1: Vacancies in Burning Glass and JOLTS

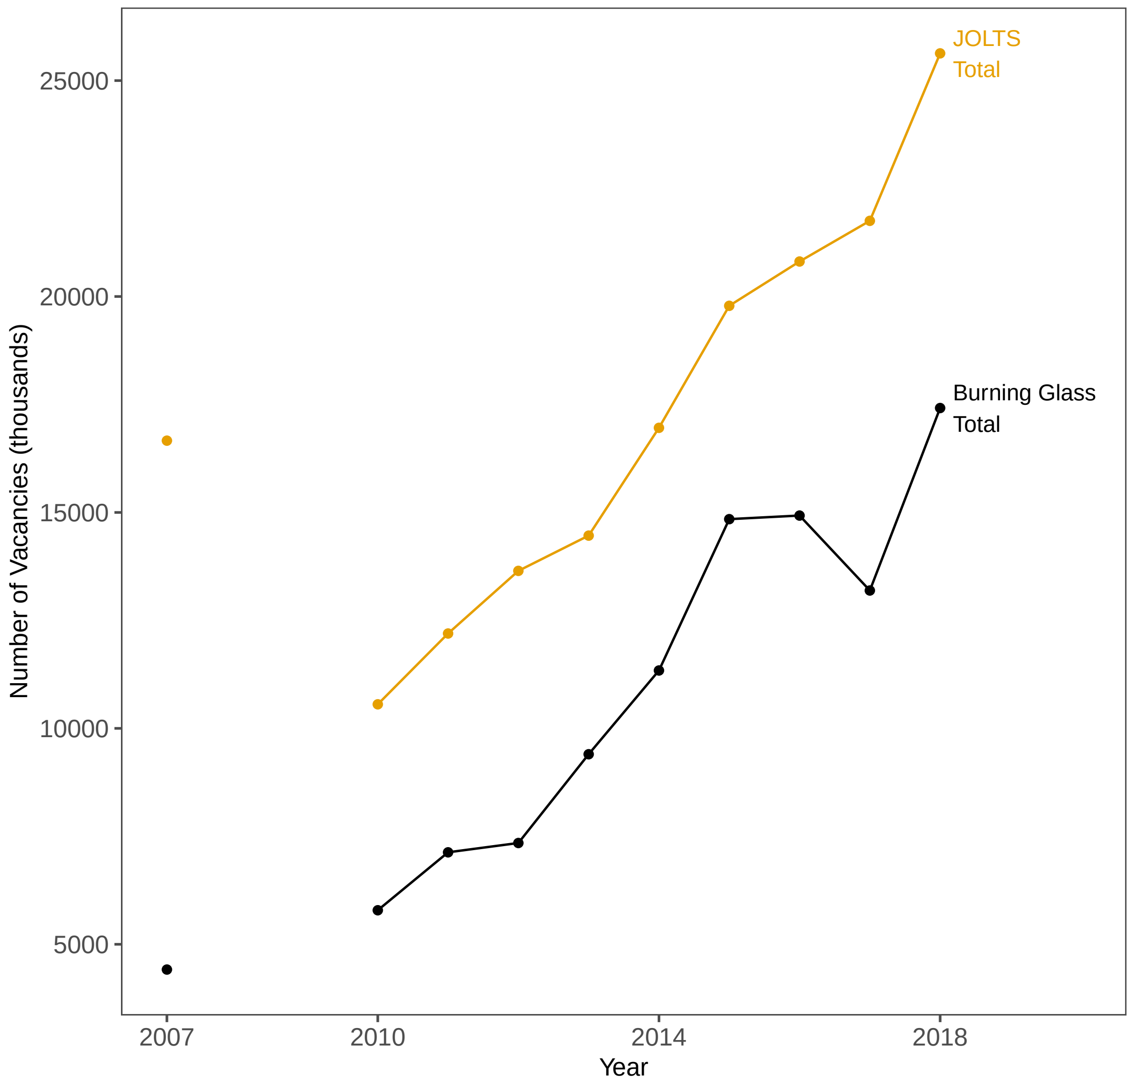

This figure plots the total number of vacancies in JOLTS, by year; and the total number of vacancies in Burning Glass, by year. We multiply the number of job openings in JOLTS by a constant factor, to arrive at a number of vacancies that matches the concept of a vacancy in Burning Glass. This method follows Carnevale et al. (2014).

Figure 2: Occupation and Industry Shares in Burning Glass

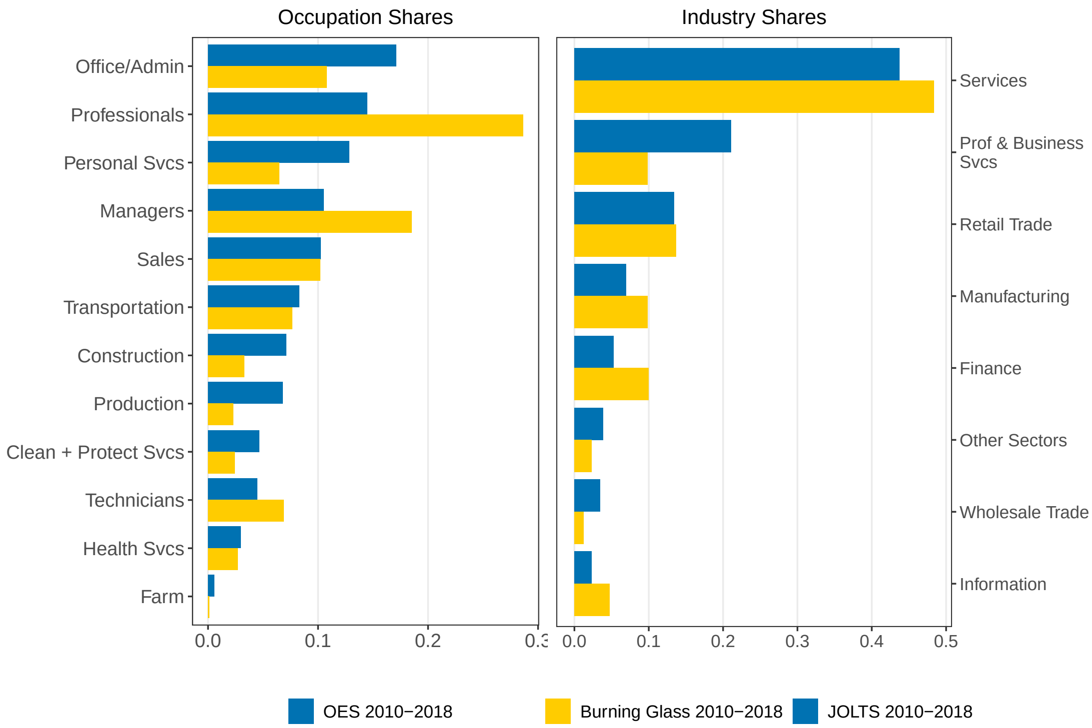

The left panel plots the share of vacancies by broad occupation in the 2010-2018 Burning Glass data, and the share of employment by broad occupation in the 2010-2018 Occupational Employment Statistics. The right panel plots the share of vacancies by broad industry in the 2010-2018 Burning Glass data, and also in JOLTS data.

FIGURE 3: Share of AI Vacancies in Burning Glass

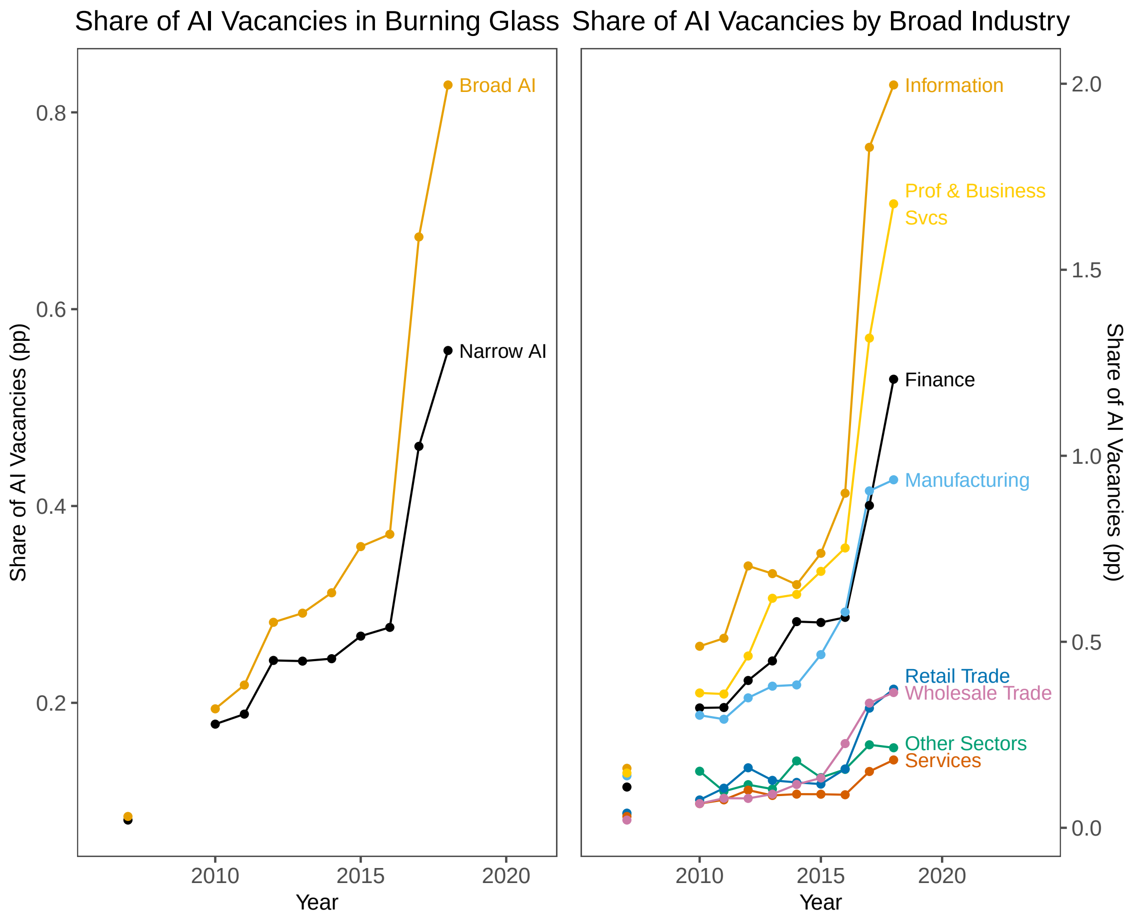

The left panel plots the share of vacancies in Burning Glass, that post a skill in the Broad or Narrow AI categories, as defined in the main text. The right panel plots the share of narrow AI vacancies in Burning Glass, by year, in each broad industry grouping.

FIGURE 4: AI Exposure by Broad Occupation and Sector

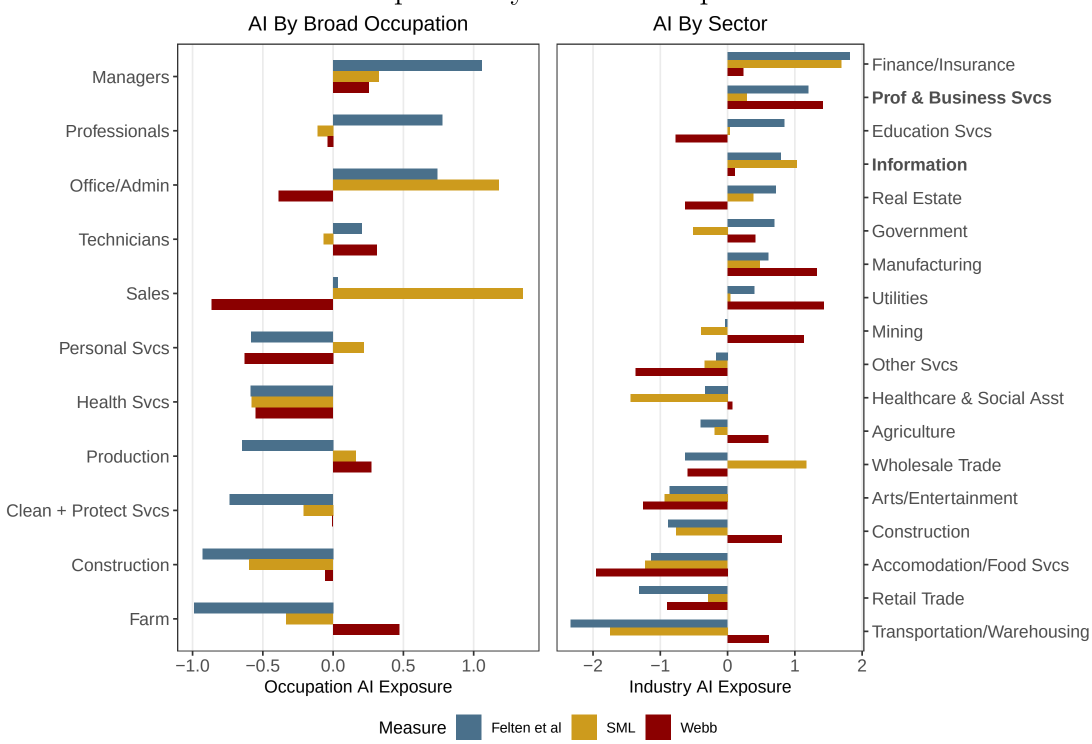

The left panel plots occupation AI exposure at the broad occupation level in Burning Glass, by taking the vacancy-weighted mean of occupation AI exposure at the 6 digit SOC occupation level, and then standardizing. The right panel plots industry AI exposure in Burning Glass, by taking the mean across the 6 digit SOC occupations posted in each 2 digit NAICS sector, weighted by the number of vacancies posted by each sector in each occupation, and then standardizing.

FIGURE 5: AI Exposure by Occupation Wages

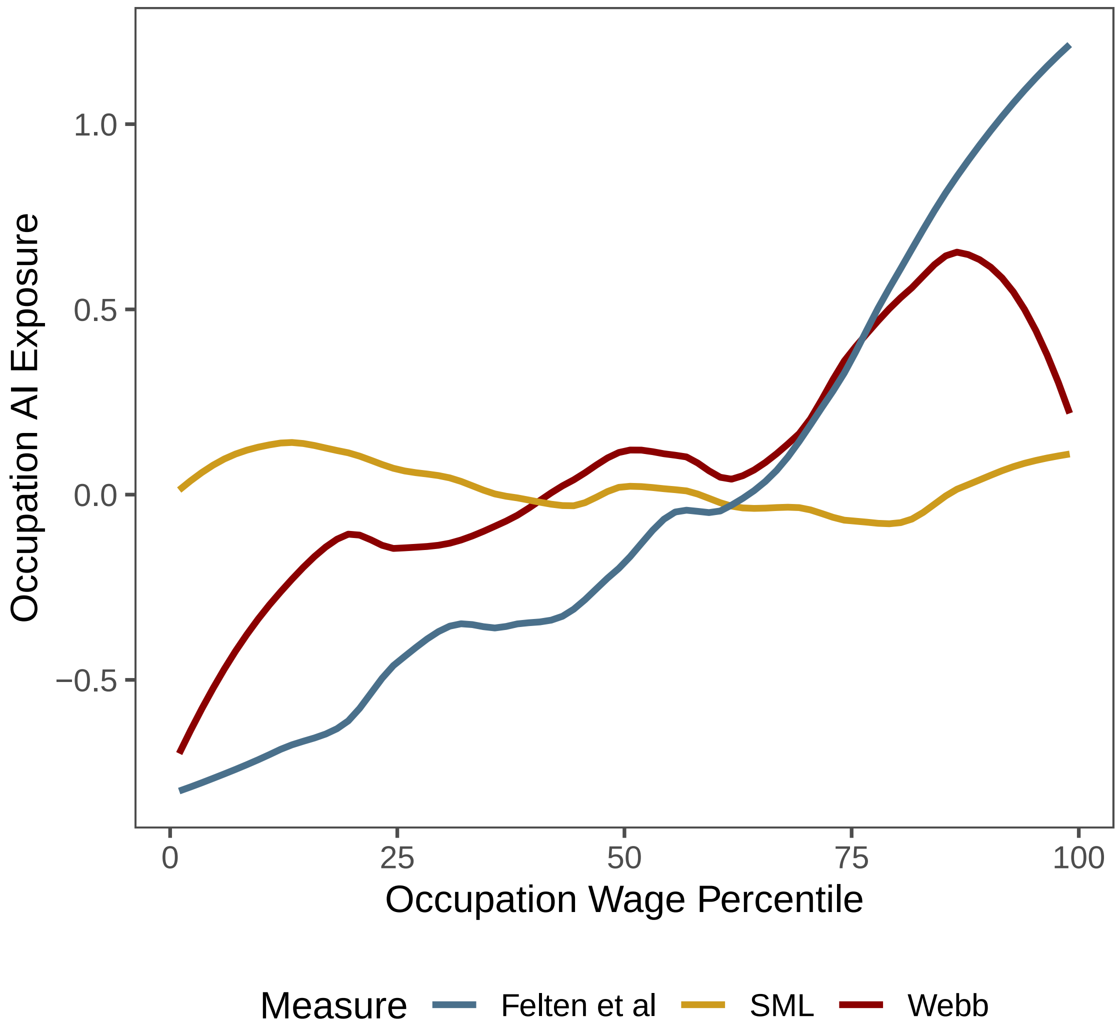

This figure plots occupation AI exposure, by 6 digit SOC occupation, against the occupation's percentile rank in the wage distribution. We study mean hourly wages from the Occupational Employment Statistics, averaged over 2010-2018. The measures of occupation AI exposure are standardized.

Figure 6: Establishment Share of AI Vacancies by Quartile of AI Exposure

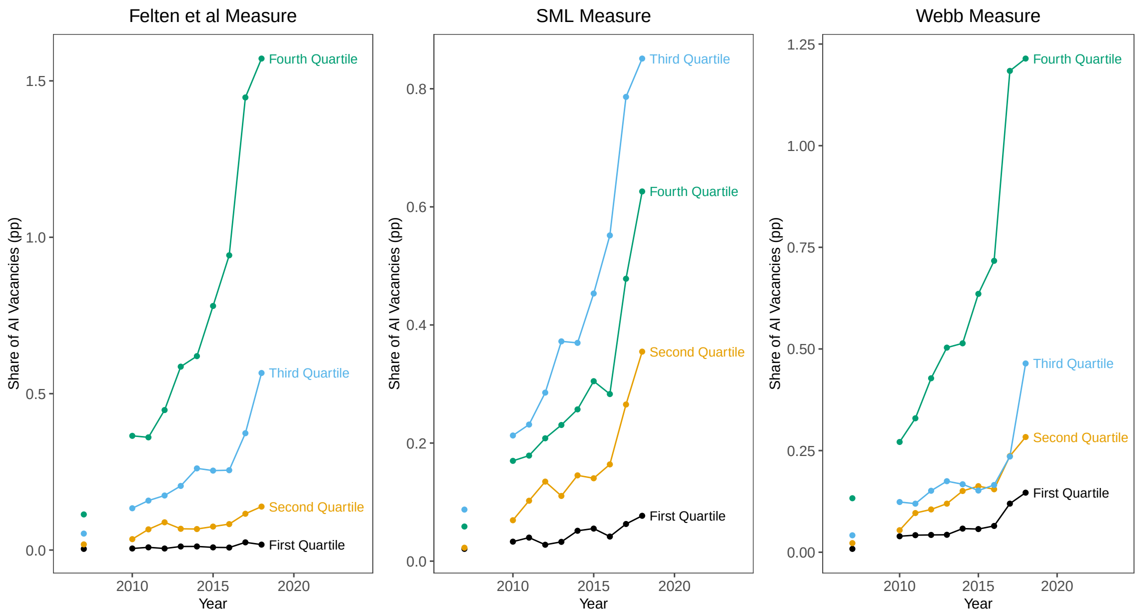

This figure plots establishments' share of AI vacancies in Burning Glass, by year, for each quartile of the distribution of 2010 establishment AI exposure. The distribution is weighted by establishment size in 2010-12. We study the distribution of residual establishment AI exposure, after controlling for the 2010-12 establishment share of vacancies in sales and administration. Establishment AI exposure is the mean of occupation AI exposure across the 6 digit occupations in which an establishment posts vacancies in 2010-12, weighted by the number of vacancies posted in each occupation. In the first panel, the measure of occupation AI exposure is from Felten et al. (2019). In the second panel, the measure is SML, from Brynjolfsson et al. (2019). In the third panel, the measure is from Webb (2020). We exclude vacancies in industry sectors 51 (Information) and 54 (Business Services).

FIGURE 7: Establishment Non-AI Vacancies by Quartile of AI Exposure

This figure plots non-AI vacancies for establishments in Burning Glass, by year, for each quartile of the distribution of 2010 establishment AI exposure. The distribution is weighted by establishment size in 2010-12. We study the distribution of residual establishment AI exposure, after controlling for the 2010-12 establishment share of vacancies in sales and administration. Establishment AI exposure is the mean of occupation AI exposure across the 6 digit occupations in which an establishment posts vacancies in 2010-12, weighted by the number of vacancies posted in each occupation. In the first panel, the measure of occupation AI exposure is from Felten et al. (2019). In the second panel, the measure is SML, from Brynjolfsson et al. (2019). In the third panel, the measure is from Webb (2020). We exclude vacancies in industry sectors 51 (Information) and 54 (Business Services).

Table 1: Effects of AI Exposure on Establishment AI Vacancy Growth, 2010-2018

	Growth of Establishment AI Vacancies, 2010-2018					
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Felten et al. Measure of AI Exposure					
Establishment AI	15.96***	13.82***	9.19***	16.53***	9.75***	16.87***
Exposure, 2010	(1.73)	(1.43)	(1.21)	(1.89)	(1.20)	(1.86)
Partial R^2	0.0250	0.0183	0.0042	0.0089	0.0046	0.0090
Observations	1,075,474	1,075,474	954,519	770,461	954,518	762,672
	Panel B: Webb Measure of AI Exposure					
Establishment AI	6.59***	5.08***	3.21***	5.91***	0.42	1.14
Exposure, 2010	(1.13)	(0.96)	(0.81)	(1.27)	(0.82)	(1.08)
Partial R^2	0.0043	0.0026	0.0008	0.0021	0.0000	0.0001
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	824,803
	Panel C: SML Measure of AI Exposure					
Establishment AI	3.76***	2.30**	-2.21**	-3.04**	1.95**	4.47***
Exposure, 2010	(1.19)	(1.04)	(0.96)	(1.38)	(0.89)	(1.34)
Partial R^2	0.0014	0.0005	0.0004	0.0005	0.0002	0.0007
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	824,803
Covariates:						
Share of Vacancies in Sales & Admin, 2010					✓	✓
Fixed Effects:						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓	✓	✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment AI vacancy growth. The outcome variable, constructed from Burning Glass data, is the growth rate of AI vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-6 include commuting zone fixed effects. Columns 3 and 5 include 3 digit NAICS industry fixed effects. Columns 4 and 6 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table 2: Effects of AI Exposure on Establishment Non-AI Vacancy Growth, 2010-2018

	Growth of Establishment Non-AI Vacancies, 2010-2018							
	Full Sample						Establishments Posting in 2018	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Felten et al. Measure of AI Exposure							
Establishment AI	-13.80***	-16.36***	-11.90***	-4.81***	-12.42***	-4.04***	-8.38**	-3.56*
Exposure, 2010	(4.22)	(4.11)	(4.08)	(1.44)	(4.01)	(1.47)	(3.46)	(1.86)
Observations	1,075,474	1,075,474	954,519	770,461	954,519	770,461	324,901	277,727
	Panel B: Webb Measure of AI Exposure							
Establishment AI	-17.24***	-18.21***	-6.73**	-2.22**	-8.30**	1.51	-4.70*	-1.44
Exposure, 2010	(3.72)	(3.63)	(3.01)	(0.93)	(3.70)	(0.98)	(2.66)	(1.36)
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	827,340	337,758	287,645
	Panel C: SML Measure of AI Exposure							
Establishment AI	7.02**	5.74*	2.05	0.95	2.21	-3.01**	0.01	-0.91
Exposure, 2010	(3.13)	(3.01)	(2.92)	(1.16)	(3.61)	(1.22)	(2.94)	(1.38)
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	827,340	337,758	287,645
Covariates:								
Share of Vacancies in Sales, Admin. in 2010					✓	✓		
Fixed Effects:								
Firm Size Decile		✓	✓		✓		✓	
Commuting Zone		✓	✓	✓	✓	✓	✓	✓
3 digit Industry			✓		✓		✓	
Firm				✓		✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment non-AI vacancy growth. The outcome variable, constructed from Burning Glass data, is the growth rate of non-AI vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The final two columns exclude establishments that do not post positive vacancies in 2018. The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2, 3, 5 and 7 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-8 include commuting zone fixed effects. Columns 3, 5 and 7 include 3 digit NAICS industry fixed effects. Columns 4, 6 and 8 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table 3: Effects of AI Exposure on Market Employment and Wage Growth

	Industry by CZ Employment Growth (CBP)			Occupation Employment Growth (OES)			Occupation Wage Growth (OES)		
	2003-2007	2007-2010	2010-2016	2004-2007	2007-2010	2010-2018	2004-2007	2007-2010	2010-2018
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				Panel A: Felten et al. Measure of AI Exposure					
Market AI Exposure,	0.03	0.10	-0.05	0.34	0.86***	0.51	-0.00	0.02	-0.17***
2010	(0.17)	(0.20)	(0.08)	(0.34)	(0.32)	(0.35)	(0.17)	(0.20)	(0.06)
Observations	10,937	10,926	10,929	736	700	680	680	648	629
				Panel B: Webb Measure of AI Exposure					
Market AI Exposure,	0.10	0.18	0.11	0.00	0.11	-0.17	0.11	-0.05	-0.02
2010	(0.15)	(0.17)	(0.09)	(0.17)	(0.21)	(0.29)	(0.08)	(0.10)	(0.04)
Observations	10,981	10,968	10,968	713	704	717	660	653	663
				Panel C: SML Measure of AI Exposure					
Market AI Exposure,	-0.14	0.37**	-0.01	0.00	-0.17	-0.37	-0.03	0.18	0.04
2010	(0.17)	(0.18)	(0.08)	(0.25)	(0.29)	(0.25)	(0.08)	(0.12)	(0.05)
Observations	10,981	10,968	10,968	713	704	717	660	653	663
Covariates:									
Share of Vacancies in Sales, Admin. in 2010	✓	✓	✓						
Fixed Effects:									
Commuting Zone	✓	✓	✓						
Sector	✓	✓	✓						
3 Digit Occupation				✓	✓	✓	✓	✓	✓

This table presents estimates of the effects of market AI exposure on market employment and wage growth. In columns 1-3, the outcome is the growth rate of sector (i.e. 2 digit NAICS industry) by commuting zone employment, measured in percentage points per year, from the County Business Patterns; for 2003-2007, 2007-2010 and 2010-2016, respectively. The sample excludes industry sectors 51 (Information) and 54 (Business Services). In columns 3-6, the outcome is the growth rate of 6 digit SOC occupation employment outside sectors 51 and 54, measured in percentage points per year, from the Occupation Employment Statistics; for 2004-2007, 2007-2010 and 2010-2018, respectively. In columns 7-9, the outcome is the growth of 6 digit SOC median hourly wages outside sectors 51 and 54, measured in percentage points per year, also from the Occupational Employment Statistics. In columns 1-3, the regressor is mean occupation AI exposure, across the 6 digit occupations posted in each sector by commuting zone cell, weighted by the number of vacancies posted in each occupation, for 2010-12. Vacancies are measured in Burning Glass. The regressions are weighted by baseline employment. In columns 4-9, the regressor is occupation AI exposure by 6 digit SOC occupation. In panel A, the measure of occupation AI exposure is from Felten et al. (2019); in panel B the measure is SML from Brynjolfsson et al. (2019); in panel C the measure is from Webb (2020). All regressions are weighted by baseline employment. All regressors are divided by the standard deviation, weighted by 2010 employment. The covariates included in each model are reported at the bottom of the table. Columns 1-3 contain sector and commuting zone fixed effects, and controls for the share of 2010-2012 vacancies in either sales or administration in each sector by commuting zone, measured from Burning Glass. Columns 4-9 control for 3 digit SOC occupation fixed effects. Standard errors are clustered by commuting zone in columns 1-3, and robust against heteroskedasticity in columns 4-9. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table 4: Effects of AI Exposure on Occupation Net Skill Change

	Occupation Net Skill Change		
	2007-2010	2010-2014	2014-2018
	(1)	(2)	(3)
	Panel A: Felten et al. Measure of AI Exposure		
Occupation AI	0.176**	-0.256	0.109
Exposure	(0.073)	(0.274)	(0.106)
Observations	695	700	702
	Panel B: Webb Measure of AI Exposure		
Occupation AI	-0.007	-0.038	0.002
Exposure	(0.038)	(0.096)	(0.081)
Observations	764	769	771
	Panel C: SML Measure of AI Exposure		
Occupation AI	0.018	0.006	0.016
Exposure	(0.047)	(0.083)	(0.080)
Observations	764	769	771
Fixed Effects:			
3 Digit Occupation	✓	✓	✓

This table presents estimates of the effects of occupation AI exposure on an occupation's net skill change, measured from Burning Glass. The units of our outcome variable is the change in share of vacancies posting a skill. We exclude vacancies in industry sectors 51 (Information) and 54 (Business Services). We calculate the measure for 2007-2010, in column 1; for 2010-2014 in column 2; and for 2014-2018 in column 3. The regressor is occupation AI exposure at the 6 digit SOC level. Each regression is weighted by number of posts in each occupation in the base year. n panel A, the measure of occupation AI exposure is from Felten et al. (2019); in panel B the measure is SML from Brynjolfsson et al. (2019); in panel C the measure is from Webb (2020). All regressors are divided by the standard deviation, weighted by the baseline number of vacancies. Each regression controls for 3 digit SOC occupation fixed effects. Standard errors are clustered by commuting zone in columns 1-3, and robust against heteroskedasticity in columns 4-9. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Appendix A: Theory

In this section, we provide a model to explain the task-based perspective on AI.

We focus on the production process of a single firm in this illustrative model. Output, y, is produced by combining the services, $y(x)$, of a measure of tasks $x \in [0, N]$ with unit elasticity (i.e., a Cobb-Douglas aggregator):

$$\ln y = \int_0^N \alpha(x) \ln y(x) dx, \tag{2}$$

where $\alpha(x) \geq 0$ designates the importance or quality of task x in the production process and $\int_0^N \alpha(x)dx = 1$.

Tasks are produced by human labor $\ell(x)$ or by AI-powered algorithms or machines

$$y(x) = \left[(\gamma_\ell(x)\ell(x))^{\frac{\sigma-1}{\sigma}} + (\gamma_a(x)a(x))^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}.$$

Workers have different occupations indexed by o and workers of occupation o specialized in disjoint subset of tasks $\mathcal{T}_o \subseteq [0, N]$ that partition the entire set of tasks used in production. Workers can be hired at a wage w_o and algorithms can be used at a cost of p at all tasks.

We will assume that workers in occupation o perform the same set of tasks across firms, which implies that

$$\alpha(x)/\alpha(x') \text{ is constant across firms for all } \{x, x'\} \in \mathcal{T}_o. \tag{3}$$

This assumption ensures that the occupational structure of a firm summarizes the set of tasks in which workers are engaged.

Within this framework, we think of new AI algorithms as enabling the production of additional tasks with algorithms or AI. That is, AI improvements map to an increase in $\gamma_a(x)$ (presumably starting at $\gamma_a(x) = 0$) for some subset of tasks.

Task Structure and Adoption

Define the occupational exposure to AI as

$$\text{occupational exposure to AI} = \frac{\int_{x \in \mathcal{T}_A \cap \mathcal{T}_o} \alpha(x) dx}{\int_{x \in \mathcal{T}_o} \alpha(x) dx}.$$

This measures the share of tasks that were assigned to workers in occupation o but that could now be performed at a lower cost by AI-powered software and machines. Our assumption in (3) implies that the occupational exposure to AI is constant across firms. We view the indices provided by Felten, Raj and Seamans, Webb, and Brynjolfsson, Mitchell and Rock as providing proxies for these exposure constructed under different assumptions about AI capabilities.

To illustrate how the task structure determines AI adoption, let us follow Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018, 2019a) and assume that $\sigma = \infty$ within tasks, so that algorithms and labor are perfectly substitutable within a task.

PROPOSITION 1 *Suppose that $\sigma = \infty$ and (3) holds. Consider an improvement in AI that allow AI-systems to be used for tasks in $\mathcal{T}_A$. Establishments that hire workers in occupations with a higher occupational exposure to AI will start producing more of their tasks with AI.*

Although the formal statement of the proposition applies to the case $\sigma = \infty$, a similar logic applies when AI complements workers in the tasks it is deployed and $\sigma < \infty$. In this case, we still have that establishments which use exposed occupations more intensively will tend to adopt more AI, but a precise proposition requires more notation and assumptions.

AI and within-task substitution

What does a shift in AI productivity does to hiring and employment of exposed firms? We will first study this question in the case in which $\sigma = \infty$ and AI systems substitute for workers in the tasks they are deployed.

Proposition 2 *Suppose that $\sigma = \infty$ and (3) holds. Consider an improvement in AI that allow AI-systems to be used for tasks in $\mathcal{T}_A$. This will increase the share of AI, and in particular of workers engaged in deploying AI in value added, and reduce the share of labor, and in particular of exposed occupations, in value added. However, the total employment effects for the firm are ambiguous.*

The main consequence of such a shift in technology will be to expand the set of tasks now performed by algorithms, $\mathcal{T}_A$, and shrink the set of tasks allocated to workers in exposed occupations. By the nature of the technology here—where algorithms and labor are perfect substitutes within tasks—the direct impact of AI will be to displace workers from tasks in $\mathcal{T}_A$.[26]

However, as emphasized in Acemoglu and Restrepo (2018), the reallocation of tasks from humans to machines or algorithms also reduces costs, creating a "productivity effect" in the form of expanding output, y, which causes an opposite expansion in hiring. This expansion in hiring is particularly pronounced in non-exposed occupations or in jobs directly related to the deployment of AI technologies.

6.1 AI and within-task complementarities

What about the human-complementary effect of AI? The possibility that AI will complement workers engaged in exposed tasks can be captured by assuming that $\sigma < 1$ so that algorithms and human labor are complementary within a task. This type of human-complementary AI will tend to increase labor demand, because algorithms are now complementary to human labor. In addition, one could also have AI expanding $\alpha(x)$ for exposed tasks, as firms start relying more intensively in these tasks.

[26]We can note that the same conclusion would apply even if the within-task elasticity of substitution between algorithms and labor, σ, were less than infinite but greater than unity. In this case, not all employment in the AI-exposed tasks would be displaced, but there would be a substitution away from labor to algorithms/machines within these exposed tasks, with the same quantitative consequences.

Evidence that AI is associated with greater establishment-level employment would be consistent with the human-complementary view, but could also be consistent with task-substitution associated with large productivity effects that nonetheless increase hiring by the exposed establishments. Conversely, evidence of negative or zero effects weigh against both the human-complementary view and the possibility of large productivity effects from AI.

7 Appendix B: Additional Tables and Figures

Figure B1: Binscatter of AI Growth and Establishment Felten et al. AI Exposure

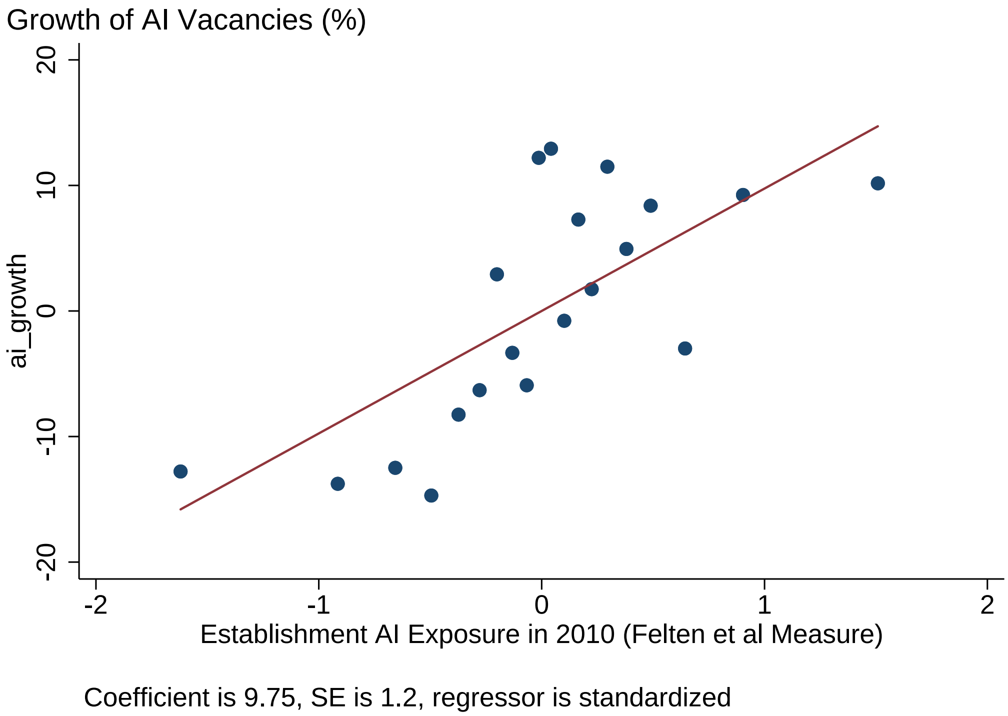

The figure presents the relationship between establishment AI Exposure in 2010 and the growth of AI establishment vacancies between 2010 and 2018. The covariates from column 5 of Table 1 are partialled out. The solid line corresponds to a regression with 2010 establishment vacancies as the weight. Each point is the mean of the y variable, conditional on the mean of the x variable, for each ventile of the x variable.

Figure B2: Binscatter of AI Growth and Establishment SML AI Exposure

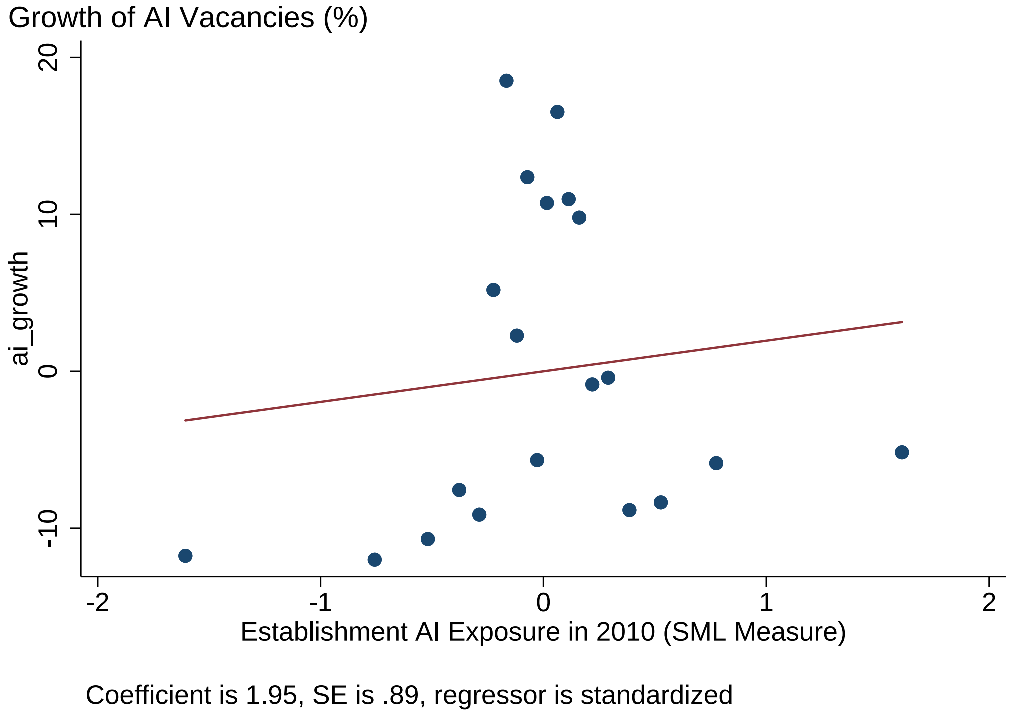

The figure presents the relationship between establishment AI Exposure in 2010 and the growth of AI establishment vacancies between 2010 and 2018. The covariates from column 5 of Table 1 are partialled out. The solid line corresponds to a regression with 2010 establishment vacancies as the weight. Each point is the mean of the y variable, conditional on the mean of the x variable, for each ventile of the x variable.

Figure B3: Binscatter of AI Growth and Establishment Webb AI Exposure

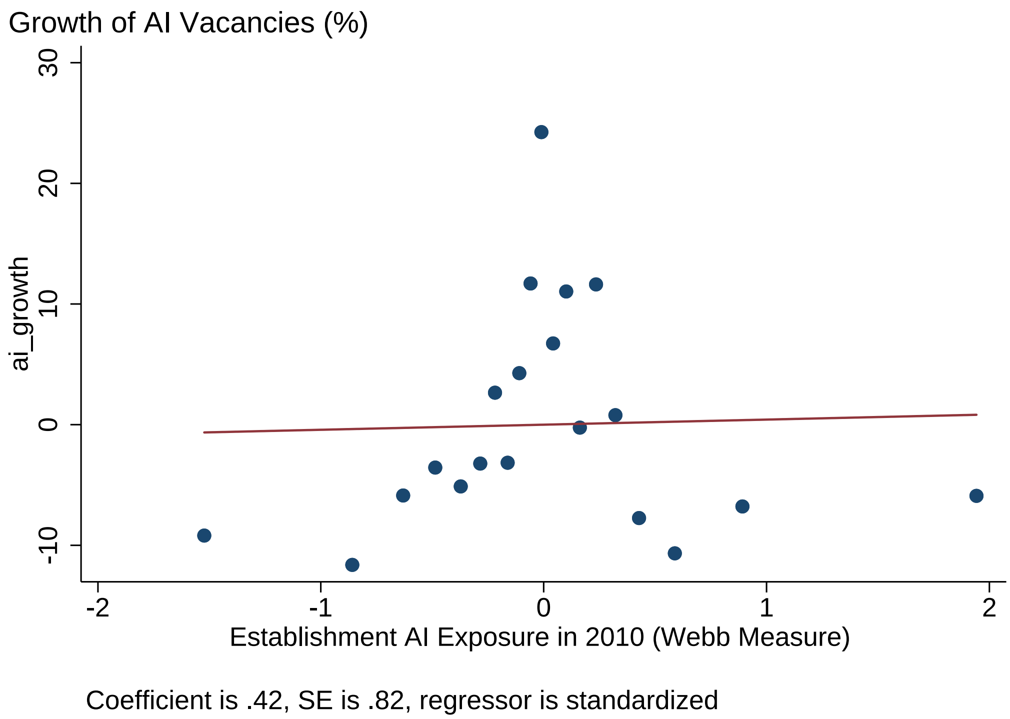

The figure presents the relationship between establishment AI Exposure in 2010 and the growth of AI establishment vacancies between 2010 and 2018. The covariates from column 5 of Table 1 are partialled out. The solid line corresponds to a regression with 2010 establishment vacancies as the weight. Each point is the mean of the y variable, conditional on the mean of the x variable, for each ventile of the x variable.

Table B1: Effects of 2007 AI Exposure on Establishment AI Vacancy Growth, 2010-2018

	Growth of Establishment AI Vacancies, 2010-2018					
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Felten et al. Measure of AI Exposure					
Establishment AI	23.32***	20.43***	12.20***	15.24***	12.07***	13.68***
Exposure, 2007	(2.33)	(1.98)	(1.78)	(2.04)	(1.74)	(1.88)
Partial R^2	0.0391	0.0301	0.0060	0.0099	0.0059	0.0079
Observations	102,783	102,783	101,553	99,078	101,524	94,866
	Panel B: Webb Measure of AI Exposure					
Establishment AI	8.87***	6.97***	4.49***	5.04***	1.92	2.48**
Exposure, 2007	(1.71)	(1.54)	(1.33)	(1.39)	(1.30)	(1.20)
Partial R^2	0.0057	0.0036	0.0012	0.0018	0.0002	0.0004
Observations	106,022	106,022	104,719	102,158	104,688	97,919
	Panel C: SML Measure of AI Exposure					
Establishment AI	7.46***	5.44***	-1.66	-3.39*	1.78	-0.68
Exposure, 2007	(1.99)	(1.78)	(1.57)	(1.82)	(1.44)	(1.64)
Partial R^2	0.0040	0.0022	0.0002	0.0007	0.0002	0.0000
Observations	106,022	106,022	104,719	102,158	104,688	97,919
Covariates:						
Share of Vacancies in Sales & Admin, 2010					✓	✓
Fixed Effects:						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓	✓	✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment AI vacancy growth. The outcome variable, constructed from Burning Glass data, is the growth rate of AI vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2007, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2007, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-6 include commuting zone fixed effects. Columns 3 and 5 include 3 digit NAICS industry fixed effects. Columns 4 and 6 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table B2: Effects of AI Exposure on Establishment Broad AI Vacancy Growth, 2010-2018

	Growth of Establishment Broad AI Vacancies, 2010-2018					
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Felten et al. Measure of AI Exposure					
Establishment AI	20.05***	17.17***	10.96***	18.72***	11.50***	18.80***
Exposure, 2010	(1.77)	(1.49)	(1.22)	(1.99)	(1.22)	(1.95)
Observations	1,075,474	1,075,474	954,519	819,014	954,519	819,014
	Panel B: Webb Measure of AI Exposure					
Establishment AI	8.47***	6.43***	3.72***	7.38***	0.99	2.53**
Exposure, 2010	(1.30)	(1.13)	(0.83)	(1.28)	(0.83)	(1.00)
Observations	1,159,789	1,159,789	1,021,673	879,954	1,021,673	879,954
	Panel C: SML Measure of AI Exposure					
Establishment AI	4.14***	2.18*	-2.22**	-2.90**	1.92**	4.90***
Exposure, 2010	(1.43)	(1.27)	(0.89)	(1.33)	(0.88)	(1.34)
Observations	1,159,789	1,159,789	1,021,673	879,954	1,021,673	879,954
Covariates:						
Share of Vacancies in Sales, Admin. in 2010					✓	✓
Fixed Effects:						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓	✓	✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment Broad AI vacancy growth. Broad AI is defined in the main text, as vacancies that post skills in the skill clusters "Artificial Intelligence" or "Machine Learning". The outcome variable, constructed from Burning Glass data, is the growth rate of AI vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-6 include commuting zone fixed effects. Columns 3 and 5 include 3 digit NAICS industry fixed effects. Columns 4 and 6 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table B3: Effects of AI Exposure on Establishment AI Share Change, 2010-2018

	Change in Establishment Share of AI Vacancies, 2010-2018					
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Felten et al. Measure of AI Exposure					
Establishment AI	0.29***	0.26***	0.20***	0.18***	0.22***	0.18***
Exposure, 2010	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Observations	341,525	341,525	324,901	299,602	324,901	299,602
	Panel B: Webb Measure of AI Exposure					
Establishment AI	0.25***	0.22***	0.14***	0.11***	0.12***	0.05**
Exposure, 2010	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Observations	355,529	355,529	337,758	311,012	337,758	311,012
	Panel C: SML Measure of AI Exposure					
Establishment AI	0.05***	0.03**	-0.06***	-0.08***	0.04**	0.05***
Exposure, 2010	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Observations	355,529	355,529	337,758	311,012	337,758	311,012
Covariates:						
Share of Vacancies in					✓	✓
Sales, Admin. in 2010						
Fixed Effects:						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓	✓	✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on the change in the share of AI vacancies. The outcome variable, constructed from Burning Glass data, is the change in the share of AI vacancies between 2010 and 2018, multiplied by 100. The shares are the ratio of AI vacancies to total vacancies in 2010-12 and 2016-2018. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-6 include commuting zone fixed effects. Columns 3 and 5 include 3 digit NAICS industry fixed effects. Columns 4 and 6 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table B4: Effects of AI Exposure on Change in Establishment Share of At Risk Vacancies, 2010-2018

	Change in Establishment Share of At Risk Vacancies, 2010-2018					
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Felten et al. Measure of AI Exposure					
Establishment AI	-1.85***	-2.00***	-4.38***	-7.47***	-4.32***	-7.19***
Exposure, 2010	(0.23)	(0.25)	(0.21)	(0.25)	(0.22)	(0.26)
Observations	341,525	341,525	324,901	299,602	324,901	299,602
	Panel B: Webb Measure of AI Exposure					
Establishment AI	-5.06***	-5.18***	-8.07***	-10.88***	-9.50***	-13.09***
Exposure, 2010	(0.26)	(0.26)	(0.33)	(0.26)	(0.32)	(0.28)
Observations	355,529	355,529	337,758	311,012	337,758	311,012
	Panel C: SML Measure of AI Exposure					
Establishment AI	-4.81***	-4.79***	-6.36***	-10.78***	-4.51***	-8.62***
Exposure, 2010	(0.43)	(0.45)	(0.34)	(0.34)	(0.44)	(0.32)
Observations	355,529	355,529	337,758	311,012	337,758	311,012
Covariates:						
Share of Vacancies in					✓	✓
Sales, Admin. in 2010						
Fixed Effects:						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓	✓	✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on the change in the share of at risk vacancies. The outcome variable, constructed from Burning Glass data, is the change in the share of at risk vacancies between 2010 and 2018, multiplied by 100. The shares are measured as the ratio of at risk vacancies to total vacancies in 2010-12 and 2016-2018. At risk vacancies are vacancies belonging to occupations in the top weighted quartile of occupation AI exposure, weighted by 2010-2018 vacancies. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-6 include commuting zone fixed effects. Columns 3 and 5 include 3 digit NAICS industry fixed effects. Columns 4 and 6 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table B5: Effects of AI Exposure on Establishment AI Vacancies, Occupation Share Controls

	Growth of Establishment AI Vacancies, 2010-2018			
	(1)	(2)	(3)	(4)
	Panel A: Felten et al. Measure of AI Exposure			
Establishment AI	10.12***	9.52***	7.24***	13.70***
Exposure, 2010	(1.58)	(1.50)	(1.44)	(2.12)
Observations	1,075,474	1,075,474	954,519	819,014
	Panel B: Webb Measure of AI Exposure			
Establishment AI	-1.55	-2.81***	-1.75**	-2.06*
Exposure, 2010	(1.06)	(0.98)	(0.84)	(1.08)
Observations	1,159,789	1,159,789	1,021,673	879,954
	Panel C: SML Measure of AI Exposure			
Establishment AI	1.52	1.56*	-0.99	0.42
Exposure, 2010	(0.95)	(0.89)	(0.91)	(1.26)
Observations	1,159,789	1,159,789	1,021,673	879,954
Covariates:				
2010 Vacancy Share in Broad Occupations	✓	✓	✓	✓
Fixed Effects:				
Firm Size Decile		✓	✓	
Commuting Zone		✓	✓	✓
3 digit Industry			✓	
Firm				✓

This table presents estimates of the effects of establishment AI exposure on establishment AI vacancy growth. The outcome variable, constructed from Burning Glass data, is the growth rate of AI vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). The regressor is the establishment's AI exposure. All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The covariates included in each model are reported at the bottom of the table. All columns control for the 2010-12 share of vacancies in each broad occupation, defined in the main text. Column 1 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-6 include commuting zone fixed effects. Columns 3 and 5 include 3 digit NAICS industry fixed effects. Columns 4 and 6 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table B6: Effects of AI Exposure on Establishment Total Vacancy Growth, 2010-2018

	Growth of Establishment Total Vacancies, 2010-2018							
	Full Sample						Establishments Posting in 2018	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Felten et al. Measure of AI Exposure							
Establishment AI	-13.61***	-16.19***	-11.79***	-4.71***	-12.29***	-3.94***	-8.16**	-3.39*
Exposure, 2010	(4.23)	(4.11)	(4.08)	(1.45)	(4.01)	(1.47)	(3.46)	(1.86)
Observations	1,075,474	1,075,474	954,519	770,461	954,519	770,461	324,901	277,727
	Panel B: Webb Measure of AI Exposure							
Establishment AI	-17.12***	-18.11***	-6.68**	-2.18**	-8.26**	1.50	-4.55*	-1.33
Exposure, 2010	(3.72)	(3.63)	(3.01)	(0.93)	(3.70)	(0.98)	(2.66)	(1.36)
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	827,340	337,758	287,645
	Panel C: SML Measure of AI Exposure							
Establishment AI	7.06**	5.77*	2.02	0.91	2.23	-2.98**	-0.04	-0.98
Exposure, 2010	(3.13)	(3.01)	(2.92)	(1.16)	(3.61)	(1.23)	(2.94)	(1.39)
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	827,340	337,758	287,645
Covariates:								
Share of Vacancies in Sales, Admin. in 2010					✓	✓		
Fixed Effects:								
Firm Size Decile		✓	✓		✓		✓	
Commuting Zone		✓	✓	✓	✓	✓	✓	✓
3 digit Industry			✓		✓		✓	
Firm				✓		✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment total vacancy growth. The outcome variable, constructed from Burning Glass data, is the growth rate of total vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The final two columns exclude establishments that do not post positive vacancies in 2018. The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2, 3, 5 and 7 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-8 include commuting zone fixed effects. Columns 3, 5 and 7 include 3 digit NAICS industry fixed effects. Columns 4, 6 and 8 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.

Table B7: Effects of AI Exposure on Establishment Non-AI Vacancy Growth, Average Establishment Size Weights

	Growth of Establishment Non-AI Vacancies, 2010-2018							
	Full Sample						Establishments Posting in 2018	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Felten et al. Measure of AI Exposure							
Establishment AI	-25.07***	-31.14***	-7.55**	-1.73	-6.01*	-0.54	-3.63	-1.35
Exposure, 2010	(4.62)	(4.51)	(3.39)	(1.82)	(3.46)	(1.92)	(3.29)	(2.06)
Observations	1,075,474	1,075,474	954,519	770,461	954,519	770,461	324,901	277,727
	Panel B: Webb Measure of AI Exposure							
Establishment AI	-8.96*	-10.64**	3.91	0.97	2.87	6.15***	5.66	1.66
Exposure, 2010	(4.72)	(4.58)	(3.84)	(2.09)	(4.53)	(2.22)	(4.02)	(2.50)
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	827,340	337,758	287,645
	Panel C: SML Measure of AI Exposure							
Establishment AI	-8.40**	-10.33***	-1.28	1.27	4.98	-1.66	-2.33	0.94
Exposure, 2010	(4.04)	(3.99)	(2.89)	(1.73)	(3.88)	(1.84)	(3.10)	(1.96)
Observations	1,159,789	1,159,789	1,021,673	827,340	1,021,673	827,340	337,758	287,645
Covariates:								
Share of Vacancies in Sales, Admin. in 2010					✓	✓		
Fixed Effects:								
Firm Size Decile		✓	✓		✓		✓	
Commuting Zone		✓	✓	✓	✓	✓	✓	✓
3 digit Industry			✓		✓		✓	
Firm				✓		✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment non-AI vacancy growth. The outcome variable, constructed from Burning Glass data, is the growth rate of non-AI vacancies between 2010 and 2018, multiplied by 100. We approximate this growth rate with the change in the inverse hyperbolic sine of the number of vacancies posted by the establishment in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the establishment posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Establishment AI exposure is divided by the standard deviation, weighted by average vacancies over 2010-12 and 2016-18. In Panel A, the measure of occupation AI exposure is from Felten et al. (2019). In Panel B, the measure of occupation AI exposure is SML, from Brynjolfsson et al. (2019). In Panel C, the measure of occupation AI exposure is from Webb (2020). All columns exclude establishments in NAICS sectors 51 (Information) and 54 (Business Services). The final two columns exclude establishments that do not post positive vacancies in 2018. The covariates included in each model are reported at the bottom of the table. Column 1 contains only establishment AI exposure. Columns 2, 3, 5 and 7 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-8 include commuting zone fixed effects. Columns 3, 5 and 7 include 3 digit NAICS industry fixed effects. Columns 4, 6 and 8 include firm fixed effects. Columns 5 and 6 control for the share of 2010-12 vacancies in each establishment, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total vacancies in 2010-2012. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.