



AI and Jobs: Evidence from Online Vacancies

Daron Acemoglu
David Autor
Jonathon Hazell
Pascual Restrepo

December 2020

NBER WORKING PAPER SERIES

AI AND JOBS: EVIDENCE FROM ONLINE VACANCIES

Daron Acemoglu
David Autor
Jonathon Hazell
Pascual Restrepo

Working Paper 28257
<http://www.nber.org/papers/w28257>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
December 2020, Revised January 2021

We thank Bledi Taska for detailed comments and providing access to Burning Glass data, Andreas Mueller, Rob Seamans, and Betsey Stevenson for discussion and comments, Jose Velarde and Zhe Fredric Kong for expert research assistance, and David Deming and Kadeem Noray for sharing their code and data. Acemoglu and Autor acknowledge support from Accenture LLP, IBM Global Universities, Schmidt Futures, and the Smith Richardson Foundation. Acemoglu also acknowledges support from Google, the National Science Foundation, the Sloan Foundation, and the Toulouse Network on Information Technology, and Autor also thanks the Carnegie Fellows Program. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

AI and Jobs: Evidence from Online Vacancies
Daron Acemoglu, David Autor, Jonathon Hazell, and Pascual Restrepo
NBER Working Paper No. 28257
December 2020, Revised January 2021
JEL No. J23,O33

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 AI capabilities. We document rapid growth in AI related vacancies over 2010-2018 that is not limited to the Professional and Business Services and Information Technology sectors and is significantly greater in AI-exposed establishments. AI-exposed establishments are differentially eliminating vacancy postings that list a range of previously-posted skills while simultaneously posting skill requirements that were not previously listed. Establishment-level estimates suggest that AI-exposed establishments are reducing hiring in non-AI positions even as they expand AI hiring. However, we find no discernible impact of AI exposure on employment or wages at the occupation or industry level, implying that AI is currently substituting for humans in a subset of tasks but it is not yet having detectable aggregate labor market consequences.

Daron Acemoglu
Department of Economics, E52-446
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
daron@mit.edu

Jonathon Hazell
Department of Economics
Princeton University
214 Julis-Rabinowitz Center
Princeton, NJ 08544
jadhazell@gmail.com

David Autor
Department of Economics, E52-438
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
dautor@mit.edu

Pascual Restrepo
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
pascual@bu.edu

1 INTRODUCTION

The last decade has witnessed rapid advances in artificial intelligence (AI) based on new machine learning techniques and the availability of massive data sets. 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. Data collection efforts have only recently commenced to determine the prevalence of commercial AI use, and hence, aside from Burning Glass data used here, we lack systematic evidence on whether there has indeed been a major increase in AI adoption—as opposed to just extensive media coverage of AI.² 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 and its implications. Our starting point is that AI adoption can be partially identified from the footprints 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 US online job vacancy postings and their detailed skill requirements from Burning Glass Technologies (Burning Glass or BG, hereafter) for the years 2007 and 2010 through 2018.³

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 that current

¹AI 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), and this allows them to perceive their environment and take actions to achieve some specific goal.

²Papers by Alekseeva et al. (2020) and Babina et al. (2020), discussed below, also apply Burning Glass data to study AI use.

³The BG data have been used in several recent papers. A non-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).

applications of AI are capable of performing specific tasks and predicts that firms engaged in those tasks will be the ones that adopt AI technologies.⁴ To identify the tasks compatible with current AI capabilities, 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 indices is computed based on different assumptions and identifies different sets of tasks and occupations as being most impacted by AI technologies, and we construct an establishment’s AI exposure from its baseline (2010) occupational structure according to each one of these indices. We use all three of these proxies for AI exposure throughout our analysis.⁵ Since our goal is to study the impact of AI on *AI-using* firms rather than *AI-producing* firms, our empirical analysis excludes firms in the Professional and Business Services and Information Technology sectors (NAICS 51 and 54), both of which are primary suppliers of AI services

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 driven by establishments with task structures that are compatible with current AI capabilities. For instance, a one standard deviation increase in our baseline measure of AI exposure based on Felten et al.—which is approximately the difference in the average AI exposure between finance versus the mining and oil extraction sector—is associated with 15% more AI vacancy posting. The strong association between AI exposure and subsequent AI hiring is robust to numerous controls and specification checks when using the Felten et al. and the Webb measures, but not robust with SML. Because the SML measure is not strongly predictive of AI adoption, we place greater emphasis on the other two measures, Felten et al. and Webb.

We then investigate whether AI exposure is associated with changes in the skill content of posted vacancies. With the Felten et al. and Webb measures (and with the SML measure to a lesser degree), we find that AI exposure is associated with both a significant decline in some of the skills previously demanded in vacancies and the emergence of new skills.

⁴See Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018, 2019a). 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). We explain below why the task-based approach is particularly well-suited to our empirical approach and how it receives support from our empirical findings.

⁵Figure 4 below 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.

This evidence bolsters the case that AI is altering the task structure of jobs, replacing some human-performed tasks while simultaneously generating new tasks accompanied by new skill demands.

The finding that establishments with AI-suitable tasks hire workers into AI positions and change their demand for certain types of skills 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 (non-AI) hiring, if *either* AI directly complements workers in some tasks, increasing productivity and encouraging more hiring, *or* AI substitutes for workers in some tasks, but increases (total factor) productivity sufficiently to raise demand in non-exposed tasks and occupations via a productivity effect (Acemoglu and Restrepo, 2019a). Alternatively, AI adoption may reduce hiring if many tasks are replaced by AI while the additional hiring in non-automated tasks spurred by AI adoption does not make up for this displacement.

Our results consistently show no positive relationship between AI exposure and establishment hiring. Rather, we find evidence of lower hiring associated with greater AI exposure in almost all of the specifications using the Felten et al. measure and in most specifications with the Webb measure. The timing of these relationships is also plausible, concentrating in the time window between 2014 and 2018, during which AI vacancy postings surged. 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. We find no evidence for either the view that there are major human-AI complementarities in these establishments or the expectation that AI will increase hiring because of its large productivity effects—though of course we cannot rule out that other applications of AI not captured here could have such effects.

In contrast to the establishment-level patterns, we do not detect any relationship between AI exposure and overall employment or wages at the industry or occupation level. There are no significant employment impacts on industries that have a task structure that exhibit greater exposure to AI, and also no employment or wages effects for occupations that are more exposed to AI.

We conclude 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, combined with the evidence for a negative association with establishment-level non-AI hiring, implies that any positive productivity and complementarity effects from AI are at present small compared to AI's displacement consequences.

Our paper builds on Alan Krueger's seminal work on the effects of new digital technologies on workers and wages (Krueger, 1993; Autor, Katz and Krueger, 1998). Subsequent literature has investigated the implications of automation technologies, focusing on wages, employment polarization and wage inequality (e.g., 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). Recent work has studied the impact of specific automation technologies, especially industrial robots, on employment and wages (Graetz and Michaels, 2019; Acemoglu and Restrepo, 2020a; Acemoglu, Lelarge and Restrepo, 2020).⁶

There are fewer studies of the effects of AI specifically, though this body of work is growing rapidly. Bessen et al. (2018) conduct a survey of AI startups, which 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 higher-quality; 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 are abundant data available. Differently from these papers' focus on AI-producing sectors and specific applications of AI such as finance, we study AI's effects on typical AI-using establishments and their non-AI workers.

Most closely related to our paper are a few recent works also investigating the relationship between AI and firm-level outcomes. 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, higher mark-ups and higher R&D intensity, and moreover, that these firms grow relatively more than non-adopters. A contrast between our approach and Babina et al.'s is that we focus on AI suitability based on establishments' occupational structures

⁶Other 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).

rather than observed AI adoption, and this may explain why we arrive at different results on for hiring. Also closely related are Dillender and Forsythe (2019) and Deming and Noray (forthcoming). Dillender and Forsythe (2019) use Burning Glass data to examine the effect of new software adoption on skill requirements at the firm level. Deming and Noray (forthcoming) use Burning Glass data to study the relationship between wages, technical skills, and skill obsolescence at the occupation level. Though neither paper focuses specifically on AI, both papers provide strong evidence that Burning Glass data are suitable for detecting changes in job skill requirements, an angle of inquiry we pursue 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 the occupational level, but we estimate significant changes in job skill demands and reduced non-AI hiring at establishments with AI-suitable occupational structures.

The rest of the paper is organized as follows. Section 2 presents a model motivating our empirical strategy and interpretation. Section 3 describes the data, and Section 4 presents our empirical strategy. Section 5 presents our main results on establishment AI exposure and AI hiring, while Section 6 looks at changes in the types of skills AI-exposed establishments are looking for. Section 7 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 8 concludes. Additional robustness checks and empirical results are presented in the Appendix.

2 THEORY

In this section, we provide a model that motivates our empirical approach and interpretation.

2.1 Tasks, Algorithms and Production

Establishment e 's output, y_e , is produced by combining the services, $y_e(x)$, of tasks $x \in \mathcal{T}_e \subset \mathcal{T}$ with unit elasticity (i.e., a Cobb-Douglas aggregator):

$$\ln y_e = \int_{\mathcal{T}_e} \alpha(x) \ln y_e(x) dx, \tag{1}$$

where \mathcal{T} is the set of feasible tasks, a subset \mathcal{T}_e of which is used in the production process of establishment e , and $\alpha(x) \geq 0$ designates the importance or quality of task x in the production process and are common across establishments. Throughout, we assume that each establishment faces a downward-sloping demand curve for its product, and also impose $\int_{\mathcal{T}_e} \alpha(x) dx = 1$ for all feasible \mathcal{T}_e , which ensures that all establishments have constant returns to scale.

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

$$y_e(x) = A_e \left[(\gamma_\ell(x)\ell_e(x))^{\frac{\sigma-1}{\sigma}} + (\gamma_a(x)a_e(x))^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad (2)$$

where A_e is an establishment-specific productivity term, σ is the elasticity of substitution between labor and algorithms or machines, and $\gamma_\ell(x)$ and $\gamma_a(x)$ are assumed to be common across establishments.

The common technology assumption can be relaxed easily but simplifies the exposition by ensuring that differences in factor demands across establishments are driven entirely by task structures, making the link between the model and our empirical approach more transparent.

Workers are homogeneous, can be hired at wage w , and can be assigned to any task. Only a subset of tasks, denoted by \mathcal{T}^A , can be performed by algorithms/AI, which can be purchased at user cost R .

We assume that the set of tasks in the economy, \mathcal{T} , is partitioned into tasks performed by a set of distinct occupations and denote the set of tasks performed in occupation $o \in O$ by \mathcal{T}^o . Each establishment e 's task structure is represented by the set of occupations that the establishment employs, denoted by $O_e \subset O$, and so $\mathcal{T}_e = \cup_{o \in O_e} \mathcal{T}^o$. In words, each establishment employs workers in a subset of occupations, and its task structure is given by (all of) the tasks performed in these occupations. For example, some establishments will employ accountants and their production will use the set of tasks accountants perform, while others require the tasks performed by security analysts or retail clerks and thus hire workers into these occupations.

These assumptions imply that advances in AI technology that improve the ability of algorithms to perform certain tasks—corresponding to an increase $\gamma_a(x)$ for some x —will have heterogeneous impacts on establishments depending on their task/occupation structure, represented by O^e for establishment e . For example, an increase in $\gamma_a(x)$ for text recognition will impact occupations in which workers perform significant text recognition tasks and will change the factor demands of “exposed establishments” meaning those that employ workers in these occupations. To make these ideas precise, analogously with our

empirical work, we define establishment e 's exposure to AI as

$$\text{exposure to AI}_e = \frac{\int_{x \in \mathcal{T}_e \cap \mathcal{T}^A} \ell_e(x) dx}{\int_{x \in \mathcal{T}_e} \ell_e(x) dx}. \quad (3)$$

This measure represents the share of tasks performed in an establishment that can be performed by AI-powered algorithms.⁷ High exposure establishments will be more impacted by advances in AI technologies. In our empirical work, we operationalize AI exposure using occupational indices provided by Felten, Raj and Seamans (2018, 2019), Webb (2020) and Brynjolfsson, Mitchell and Rock (2018, 2019), as noted above.⁸

We next explore under what conditions greater establishment exposure to AI will lead to AI activity and to more or less hiring of (non-AI) workers.

2.2 Task Structure and AI Adoption

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 $\sigma = \infty$, so that algorithms and labor are perfectly substitutable within a task.

PROPOSITION 1 *Suppose that $\sigma = \infty$. Consider an improvement in AI technologies that increases $\gamma_a(x)$ in \mathcal{T}^A . Establishments with greater exposure to AI will increase their share of AI/algorithms in value added.*

The proof of this proposition follows from Acemoglu and Restrepo (2018) and is omitted.

Although the formal statement of the proposition applies to the case $\sigma = \infty$, a similar logic applies when AI does not fully replace workers in the tasks it is deployed, meaning that $\sigma > 1$. In this case, we still have that more exposed establishments will tend to adopt more

⁷When $\sigma = \infty$ as in Propositions 1 and 2 below, $\ell_e(x)$ is proportional to $\alpha(x)$ in all establishments (provided that task x is performed by labor). In that case, equation (3) can be written in terms of $\alpha(x)$'s.

⁸Using our notation, we can represent these indices as

$$OExp^o = \frac{\int_{x \in \mathcal{T}^o \cap \mathcal{T}^A} \ell(x) dx}{\int_{x \in \mathcal{T}^o} \ell(x) dx},$$

where $\ell(x)$ is average employment in task x and we denote average employment in occupation o by $\ell^o = \int_{x \in \mathcal{T}^o} \ell(x) dx$. When $\ell(x) = \ell_e(x)$, which follows from our common technology assumption, the exposure to AI measure is equal to the employment weighted average of the occupation AI exposure measure:

$$\frac{\sum_{o \in O_e} OExp^o \ell^o}{\sum_{o \in O_e} \ell^o} = \frac{\sum_{o \in O_e} \frac{\int_{x \in \mathcal{T}^o \cap \mathcal{T}^A} \ell(x) dx}{\int_{x \in \mathcal{T}^o} \ell(x) dx} \ell^o}{\sum_{o \in O_e} \ell^o} = \frac{\int_{x \in \mathcal{T}_e \cap \mathcal{T}^A} \ell(x) dx}{\int_{x \in \mathcal{T}_e} \ell(x) dx} = \text{exposure to AI}_e.$$

AI, and this will raise the share of AI/algorithms in value added. When $\sigma < 1$, there will again be greater AI use following improvements in AI technology, but this may not increase the share of AI/algorithms in value added, because of strong complementarities.

In our empirical work, we will identify greater use of AI with the posting of more AI vacancies.⁹ We also note that if different tasks require different skills that are included in job postings, the adoption of AI technologies for tasks previously performed by labor will be associated with a reduction in some of the skills that were previously demanded in vacancies. Conversely, if working with AI requires new skills, we may also see additional skills demanded in vacancies after the adoption of AI technologies. We explore these implications empirically as well.

2.3 AI, Task Displacement and Hiring

How do improvements in AI technology impact hiring? We first study this question in the case in which $\sigma = \infty$ and AI systems substitute for and displace workers in the tasks they are deployed. The next proposition also follows from the analysis in Acemoglu and Restrepo (2018) and its proof is omitted.

PROPOSITION 2 Suppose that $\sigma = \infty$. Consider an improvement in AI technologies that increases $\gamma_a(x)$ in \mathcal{T}^A . This will increase the share of AI/algorithms in value added and will reduce the share of labor in value added. The total employment effects for the establishment are ambiguous.

The direct consequence of an improvement in AI technology is to expand the set of tasks now performed by algorithms, \mathcal{T}^A , and to shrink the set of tasks allocated to workers in exposed occupations. Because $\sigma = \infty$, better AI technology will displace workers from tasks in \mathcal{T}^A , reducing the share of employment in exposed occupations in overall employment.¹⁰

However, as emphasized in Acemoglu and Restrepo (2018), the reallocation of tasks from humans to machines or algorithms also reduces production costs, creating a “productivity effect”. This productivity effect increases establishment output, y_e , and may generate an expansion in hiring driven by tasks that are not displaced by AI. Establishment employment will increase if this expansion is greater than the displacement of workers from tasks taken over by AI. It will decrease otherwise.

⁹We do not incorporate AI workers explicitly into the model, both for simplicity and to avoid confusion between their employment and the employment of (non-AI) workers performing tasks suitable for AI.

¹⁰The same conclusion would apply even when σ is not infinite but still greater than 1. 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 qualitative consequences.

Our empirical work will explore the relationship between AI adoption and employment. Specifically, we ask how AI exposure affects (non-AI) hiring, and in particular, whether establishments with baseline task structures that are well-suited to AI technologies according to (3) post more vacancies for non-AI positions.

2.4 Human-Complementary AI

What about the human-complementary effects of AI? The possibility that AI will complement workers engaged in exposed tasks can be captured by assuming that $\alpha(x)$ increases for exposed tasks or, alternatively that $\sigma < 1$, so that algorithms and human labor are complementary within a task (or both). This type of human-complementary AI will increase labor demand because algorithms raise human productivity in exactly the tasks in which AI is being adopted.

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 at the exposed establishments. Conversely, evidence of negative, or even zero, effects would weigh against both the human-complementary view and the possibility of large productivity effects from AI—since AI-human complementarity or employment-generating productivity effects are needed for AI adoption to raise establishment-level employment. Our evidence below suggests that establishments with AI-suitable task structures reduce hiring of non-AI workers, providing support for the view that the current generation of AI technologies is predominantly task-replacing and generates productivity effects that are not large enough to counterbalance any resulting worker displacement. We note again that this conclusion applies to the variation in AI adoption identified by baseline task structure. It remains possible that other AI technologies than the ones we are proxying here could have different effects.

3 DATA

We next describe the Burning Glass Technologies data, document that it is broadly representative of employment and hiring trends across occupations and industries, present our AI exposure indices, and document their distribution across occupations and their evolution over time.

3.1 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 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 more limited coverage in 2007. Our primary sample comprises data from the start of 2010 until October 2018, though we also make use of data from 2007. The vacancy data enumerate 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 online job postings. Figure 1 shows that BG data closely track the evolution of overall vacancies in the US economy as recorded by the nationally representative Bureau of Labor Statistics (BLS) Job Openings and Labor Turnover Survey (JOLTS). The exception is the downturn in BG postings data between 2015 and 2017.¹¹ Appendix Figure A1 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 JOLTS.¹²

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, as well as an industry code. We classify each firm as belonging to the industry in which it posts the most vacancies over our sample period. We define an establishment of a firm as the collection of vacancies pertaining to a firm and commuting zone. Commuting zones are groups of counties that, due to their strong commuting ties, approximate a local labor market (Tolbert and Sizer, 1996).

Of particular importance for our paper are BG’s detailed skill and 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”. Over 95% of vacancies are assigned a six-digit (SOC) occupation code.¹³

¹¹We adjust the numbers 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.

¹²Descriptive statistics and additional information on the BG data are provided in the Appendix. We note that BG data represent vacancy flows, while the OES reports employment stocks, and thus we do 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).

¹³Six-digit occupation codes are highly granular, including occupations such as pest control worker, college

We construct two measures of AI vacancies, narrow and broad, using this information. The narrow category includes a selection of skills relating to AI.¹⁴ The broad measure of AI includes skills belonging to the broader skill clusters of Machine Learning and Artificial Intelligence. A 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 2 shows the evolution of postings of narrow and broad AI vacancies in the BG data, highlighting the rapid takeoff of AI vacancies after 2015, as noted in the Introduction. Our empirical work exploits this sharp acceleration to test whether the estimated relationship between AI adoption and non-AI hiring matches the timing of adoption. While a sharp uptick is visible in all industries, the the second panel of Figure 2 shows that the takeoff is particularly pronounced in the information, professional and business services, finance, and manufacturing sectors.

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, while the latter contains industries such as management consultancy, likely to be integrating AI into other industries’ production processes.

3.2 AI Indices

We study three measures of AI exposure. Each is assigned at the 6 digit SOC occupation level, and each 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). It is based on data from the AI Progress Measurement project, from the Electronic Frontier Foundation. The Electronic Frontier data identify a set of nine application areas in which artificial intelligence has made progress since 2010—for example, image recognition or language modeling. Felten et al. use Amazon MTurk to collect crowdsourced assessments of the relevance of each of these application

professor in physics, and home health aide.

¹⁴The 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, Un-supervised 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.

areas to the 52 O*NET ability scales, e.g., depth perception, number facility, and written comprehension. The authors finally construct the AI Occupational Impact (AIOI) for each O*NET occupation as the weighted sum of the 52 AI application-ability scores over each occupation, where weights are equal to the O*NET-reported prevalence and importance of each ability in the occupation.

The second measure is from Webb (2020). Webb’s analysis identifies occupations that are exposed to AI based on whether the occupation’s tasks, measured using O*NET as in Felten et al., overlap with capabilities described in AI patents. Occupations that have a larger fraction of such overlapping tasks are classified as more exposed.

The third measure is Suitability for Machine Learning (SML) from Brynjolfsson et al. (2019). To build this measure, Brynjolfsson et al. (2019) develop a 23-item rubric for systematically scoring the suitability of any given job task for machine learning/AI. They apply this rubric to the full set of O*NET occupations using CrowdFlower, a human intelligence task crowdsourcing platform, to score the O*NET textual description of each occupation, thus producing an SML score.

Figure 3 shows the distribution of our three indices by broad occupation categories and by one-digit industry.¹⁵ Figure 4 relates this same information to wages by plotting average AI exposure by occupational wage percentile for each index. The figures underscore that these three measure capture different aspects 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 Webb measure is not particularly high in sales occupations and shows a strong positive relationship with occupational wage percentiles. 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. Consequently, SML has no systematic relationship with occupational wage percentiles.¹⁶

¹⁵The broad occupational categories are those utilized by Autor (2019) and aggregate six-digit occupations into 10, roughly one-digit categories.

¹⁶Another 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. We use this index in our robustness checks as a control.

4 EMPIRICAL STRATEGY

Our empirical strategy links measures of AI activity and job posting outcomes to AI exposure, where both outcome and exposure variables are measured at the establishment level and AI exposure is measured using the three indices above.

We estimate the following regression model:

$$\Delta y_{e,t-t_0} = \beta AI_{e,t_0} + \mathbf{x}'_{e,t_0} \boldsymbol{\gamma} + \varepsilon_{e,t-t_0}, \quad (4)$$

where e denotes establishment, $\Delta y_{e,t-t_0}$ denotes the change in one of our establishment-level outcomes between 2010-12 and 2016-18, $AI_{e,t_0} = \sum_o \text{share postings}_{oe,t_0} \times \text{AI score}_o$ is one of our three measures of establishment AI exposure calculated using establishment data for 2010 through 2012, and \mathbf{x}_{e,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, a set of firm fixed effects.¹⁷ Finally, $\varepsilon_{e,t-t_0}$ is an error term representing all omitted factors.

Our primary interest is in the coefficient β , which captures the relationship between AI exposure and the outcome variable. We standardize the establishment AI exposure measure by dividing it by its standard deviation, weighted by vacancies in 2010-2012. Hence, β is the change in the outcome variable associated with a standard deviation change in AI exposure. We also aggregate this equation to the occupational and industry level to explore hiring and other outcomes for more exposed occupations and industries.

The three main outcome measures we focus on for $\Delta y_{e,t-t_0}$ are AI vacancies, changes in job skill requirements of posted vacancies, and overall non-AI hiring, all measured at the establishment level.

5 AI EXPOSURE AND AI VACANCIES

We first document that AI exposure predicts establishment-level AI activity, as proxied by our measure of narrow AI vacancies. Table 1 presents the main estimates. Panel A of this table shows the relationship between AI exposure based on Felten et al.’s index and growth in AI vacancies, while the subsequent two panels present results for our other measures of AI exposure. We estimate regression models based on (4), with the left-hand side variable defined as the change in the inverse hyperbolic sine of AI vacancies between 2010-12 and

¹⁷We pool 2010-12 data and, separately, 2016-18 data to improve precision.

2016-18.¹⁸ We focus on weighted specifications, using baseline establishment vacancies as weights. We report heteroscedasticity-robust standard errors that allow for arbitrary cross-sectional correlation across the establishments of each firm.

Column 1 is our most parsimonious specification and includes no covariates, thus depicting the unconditional bivariate relationship. 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. This estimate implies that a one standard deviation increase in AI exposure—which corresponds to the difference between the finance and the mining and oil extraction sector—is associated with approximately a 16% increase in AI vacancies.¹⁹

The remaining columns explore the robustness of this 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 now more precisely estimated. Column 3 additionally adds three-digit (NAICS) industry fixed effects. Reflecting the sizable variation in AI exposure across industries shown in Figure 3 above, these controls are more important for our regressions, 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. 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.

Figure 3 documented significant differences in AI exposure across occupations. This raises the concern that our results may be confounded by secular trends across broad occupational categories. Columns 5 and 6 additionally control for 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 do not substantially change the estimate in either column, which remain, respectively, at 9.75 (standard error = 1.20) and 16.87 (standard error = 1.86).²⁰

¹⁸The inverse hyperbolic sine transformation is given by:

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

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.

¹⁹The partial R^2 for our AI exposure measure in this regression is 2.5%. Given the vast cross-establishment variation in AI adoption, this is a sizable share of variation. For comparison, the partial R^2 of three-digit industry fixed effects in column 3 is 5.2%.

²⁰Table ?? in the Appendix shows that the results in Panel A are also robust if we include the baseline

The first panel of Figure 5 shows the specification from column 3 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 does not appear to be driven by outliers. The first panel of Figure 6 provides a complementary visualization, depicting the evolution of AI vacancies for the four quartiles of the establishment AI exposure measure. It shows that the top two quartiles post significantly more AI vacancies and drive the surge in AI vacancies after 2015.

Panels B and C of Table 1 repeat the Panel A regressions using the Webb and SML indices of AI exposure, respectively. Estimates using the Webb measure, reported in Panel B, are similar to those in Panel A in the first four specifications, though they are not fully robust to controls for the baseline shares of sales and administration vacancies in columns 5 and 6. The second panel of Figure 5 shows that the bin scatter plot with the Webb measure looks similar to the one in the first panel with the Felten et al. measure, and Figure 6 confirms that the surge in AI vacancies is again driven by the top two quartiles. Given that the Felten et al. and Webb indices capture different components of AI exposure (recall Figure 3), the broadly similar picture that emerges from these two measures is reassuring.

Results with the SML index in Panel C are significantly weaker, however. There is a positive association between the SML-based measure of exposure to AI and AI vacancy growth without any covariates, but when three-digit industry fixed effects are included, this relationship becomes negative. The proximate explanation for this pattern is that, as noted above, the sales and administration occupations have a high SML score and are negatively associated with the adoption of AI. When we control for the baseline shares of these occupations in columns 5 and 6, the positive relationship in column 1 is restored. The third panels of Figures 5 and 6 show a less clear pattern as well. The bin scatterplot confirms the lack of a robust relationship between exposure to AI based on the SML measure and AI vacancies (from the specification in column 3), and the evolution of AI vacancy growth by exposure quartiles in Figure 6 no longer shows a monotone pattern. This pattern motivates our greater emphasis on the results using the Felten et al. and Webb measures in the remainder of the paper.

One concern is that the AI measures may be proxying for exposure to non-AI digital technologies. If so, this would cloud the interpretation of our estimates as primarily capturing the impacts of AI exposure on establishment outcomes. We check for this possibility in

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 for the Webb and SML AI exposure measures are sensitive to these controls.

Table 2 by additionally controlling for Webb’s measure of exposure to “software”, which is calculated analogously to his AI exposure measure, but focusing on occupations and tasks suitable for traditional software and digital technologies. The inclusion of the software exposure measure has little impact on the coefficients of interest, particularly in the case of the Felten et al. measure. For example, in the most exhaustive specification (column 6), the point estimate is now 17.47 with a standard error of 1.90, as compared to 16.87 and a standard error of 1.86 in Table 1. The software exposure measure itself does not have a consistent association with AI vacancy growth: it is positive and statistically significant in some specifications, small and insignificant in others, and negative and significant in still others. This set of estimates bolsters our confidence that the AI exposure variable identifies meaningful variation in the suitability of establishment task structure for AI, and that this variation is distinct from exposure to traditional software and digital technologies.

We provide several robustness checks on these basic patterns in the Appendix. In Table A1, we report estimates for AI vacancy growth using AI exposure calculated using establishments’ occupational structures in 2007 rather than 2010. Although this greatly reduces the sample size as many establishments operating in both 2010 and 2018 are not found in the 2007 data, we obtain qualitatively similar results to those in Table 1.

In Table A2, we replace the narrow AI vacancy measures used in Table 1 with the broad AI vacancy indices discussed above (see Figure 2), and in Table A3, we use the change in the *share* of AI vacancies among all vacancy postings as the dependent variable. The results in both tables corroborate our main findings in Table 1 and, if anything, are stronger and more stable, especially with the Webb measure. The quantitative magnitudes implied by the estimates in these tables are comparable and plausible. For example, with the Felten et al. measure and the specification in column 6, a one standard deviation higher AI exposure is associated with a 0.18 percentage point increase (standard error 0.02) in the share of AI vacancies between 2010 and 2018.

We also explored firm-level variants of the establishment-level models above. Because of the many zeros in the data, the establishment-level estimates do not aggregate cleanly to firm-level estimates. As shown in columns 1 and 2 of Table A4, these estimates are generally imprecise and inconsistently signed. However, when we estimate models for the firm-level mean of establishment AI vacancy growth (columns 3 and 4) or models with share of AI vacancies (columns 5 and 6), the estimates are very similar to our main results in Table 1.

In summary, the data point to a recent surge in AI-related hiring, and our regression evidence finds that establishments whose task structures enable the use of AI. have substantially increased their postings of AI-related vacancies. This evidence suggests that an important

component of AI activity is related to the types of tasks performed in an establishment—though it does not preclude the possibility that AI activity has other drivers, such as the development of new products or new business models. Having established the link between AI suitability and AI hiring at the level of establishments, we next turn to the broader labor market implications of growing AI adoption.

6 AI AND NEW SKILLS

Since AI is intended to supplement, replicate, and in some cases, exceed human-level intelligence in a variety of tasks, we anticipate that establishments with task structures that are suitable for AI will tend to change the types of skills they demand. We explore here whether this occurs.

To investigate whether AI exposure predicts change in the nature of skills that an establishment demands in *non-AI* jobs, we build on work by Deming and Noray (forthcoming) who document such changes associated with broader IT-related activity. We adapt Deming and Noray’s (forthcoming) measure of change in skill demands to establishments, and we also separate their gross skill change measure into negative and positive changes, capturing the disappearance of existing skills and the emergence of new skills, respectively:

$$\text{negative skill change}_{e,t_2,t_1} = - \min \left\{ \sum_{s=1}^S \left[\left(\frac{\text{skill}_{e,t_2}^s}{\text{vacancies}_{e,t_2}} \right) - \left(\frac{\text{skill}_{e,t_1}^s}{\text{vacancies}_{e,t_1}} \right) \right], 0 \right\}, \text{ and}$$

$$\text{positive skill change}_{e,t_2,t_1} = \max \left\{ \sum_{s=1}^S \left[\left(\frac{\text{skill}_{e,t_2}^s}{\text{vacancies}_{e,t_2}} \right) - \left(\frac{\text{skill}_{e,t_1}^s}{\text{vacancies}_{e,t_1}} \right) \right], 0 \right\},$$

Here, $\text{skill}_{e,t}^s$ is the number of times skill s is posted by establishment e in year t , which we normalize by dividing it by the total number of vacancies posted by that establishment. The negative skill change measure therefore represents a decline in the frequency with which some of the skills that were formerly posted appear in vacancies, while the positive skill change measure captures increases in the frequency with which other skills are posted in vacancies—which may include the addition of skills that were not previously posted. We calculate these measures for non-AI vacancies and, as before, for all establishments except those in the Professional and Business Services and Information Technology sectors (51 and 54).

Tables 3 and 4 show that establishment AI exposure is robustly associated with both more negative and more positive skill changes. For example, in the first panel of Table

3, which focuses on negative skill changes, the column 1 estimate of 0.83 indicates that a one standard deviation in the Felten et al. AI exposure measure is associated with a 0.83 absolute *decline* in the per-vacancy frequency with which skills previously demanded are posted (standard error 0.09). This is a large change compared to the mean negative skill change in our sample, 4.70, and thus suggests that the deployment of AI technologies goes hand-in-hand with significant skill redundancies. Equally interesting is the pattern in Table 4, which shows that AI exposure is associated with demand for new skills. Column 1 of this table shows that a one standard deviation increase in the Felten et al. AI exposure measure is associated with 0.95 absolute increase (standard error = 0.08) in the per-vacancy frequency of skills that were either demanded less frequently previously or were not previously demanded. This too is a sizable impact as compared to the mean positive skill change in our sample of 6.30.

These patterns are quite robust, as is shown in the remaining columns and panels of Tables 3 and 4. Each of the three AI exposure measures—Felten et al., Webb, and SML—predict both negative and positive establishment-level skill changes between 2010 and 2018. Paralleling our findings at many points in the paper, the Felten et al. measure proves to have the most stable and the largest quantitative relationship to the outcome variable, followed by Webb and SML. In particular, all measures prove robust to the inclusion of firm size deciles, commuting zone dummies, and controls for initial establishment vacancy structures in sales and administrative occupations. All three are robust to the inclusion of firm fixed effects as well. The association between AI exposure and positive skill changes is no longer present when we include firm fixed effects, however, presumably because the addition of new skills is not localized to the establishments that have higher AI exposure but rather occurs throughout the firm.

The finding that AI exposure is associated with significant changes in the set of skills posted in vacancies bolster our confidence that AI adoption has real effects on the job content of non-AI jobs—enabling firms to replace some of the tasks previously performed by workers, making certain skills redundant, while simultaneously generating demand for new skills. These results are also consonant with our theoretical model in Section 2, which suggests that AI adoption will induce churn of job tasks performed by workers as some tasks previously performed by humans are taken over by algorithms, causing labor reallocation. While the churn prediction is unambiguous, the theory leaves open the possibility that hiring of workers in AI-exposed establishments could either rise or fall: if algorithms are a complement to labor, hiring will increase; conversely, if algorithms are a substitute for labor, hiring will fall unless the productivity gains from AI adoption are sufficiently strong

to offset the direct reduction in labor demand stemming from algorithm-labor substitution. The following section explores the net impact of AI exposure on (non-AI) hiring.

7 AI AND JOBS

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

7.1 AI Exposure and Establishment Hiring

Table 5 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). The non-AI vacancy measure is chosen so as to focus on the effects of AI activity on establishment hiring exclusive of already-reported impact on AI hiring itself. As before, we drop the primary sectors 51 and 54.

In Panel A, where we focus on Felten et al.’s measure, we see a robust 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 roughly 14% decline in overall non-AI vacancies. (We interpret the economic magnitudes of these point estimates below.) This coefficient estimate remains stable when we control for firm size decile, commuting zone and three-digit industry fixed effects in columns 2 and 3.

In column 4, we replace firm-level covariates with firm fixed effects while retaining the commuting zone dummies from the prior column. This is a stringent specification since we are now comparing between establishments of the same firm that differ in their AI exposure as measured by their baseline (2010-2012) vacancy postings across detailed occupations. In this specification, the point estimate for AI exposure is -4.81, which is about half of its prior magnitude. Simultaneously, the estimates become more precise as the standard error falls from 4.08 to 1.44.

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 2010-2012 and then stopped posting in 2016-2018 (which may reflect either true zero

vacancy postings or establishment exits, perhaps for sampling reasons). Columns 7 and 8, therefore, limit the sample to establishments that posted in 2016-2018. The estimates are now somewhat smaller, but still negative and statistically significant at 5%: -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.²¹

How large are the effects reported in Panel A? The interpretation of the regression coefficients is not straightforward because our outcome variable is vacancy flows, which differs from the stock of employment. To estimate the implied impact on employment, we cumulate vacancies between 2010 and 2018 to create a measure of 2018 employment for each establishment. Then we regress our measure of cumulative hiring between 2010 and 2018 on establishment AI exposure in 2010 exactly as in Table 5.²² Table A6 in the Appendix reports the results of this exercise. In Panel A, with Felten et al.’s measure of AI exposure, the regression coefficient in column 1 is -7.24, which implies that a one standard deviation in AI exposure is associated with a 7% decline in implied non-AI employment between 2010 and 2018. Since a one standard deviation increase in AI exposure is quite large, this is a sizable but not implausible relationship. We also note that because this coefficient estimates the *relative* change in non-AI hiring at more versus less AI-exposed establishments, it does not imply an aggregate reduction in total hiring.

Panel B turns to Webb’s measure of AI exposure. The pattern is broadly similar to the one we see in Panel A but less stable. The coefficient estimate without any covariates in column 1 is -17.24 (standard error = 3.72). It is comparable in column 2 when we control for commuting zone and three-digit industry fixed effects. However, the estimate 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 through 8. When we use the SML measure in Panel C, there is no consistent evidence for a negative association between AI exposure and non-AI vacancy postings. While there is a negative and statistically significant estimate in column 6, there is a positive association in other columns.

If AI is indeed distinct from traditional (“procedural”) software in that it can potentially

²¹Since there are positive effects on AI vacancies, it is not self-evident whether the impact on total vacancies (including AI hiring) is also negative. We show in Table A5 in the Appendix that the answer is yes. This is as expected, since AI vacancies are a tiny share of total vacancies.

²²More specifically, we assume that establishment employment, $\ell_{e,t}$, follows the law of motion $\ell_{e,t+1} = f v_{e,t} + (1 - s) \ell_{e,t}$, where $v_{e,t}$ denotes the establishment’s vacancies, f is the vacancy fill rate and s is the employment separation rate. We calculate employment in 2010 by assuming that the establishment is in steady state at the state, and we compute employment in 2018 by iteratively solving the law of motion forward. We set $s = 0.4$ to match the 2018 annual separation rate from JOLTS. We thank Andreas Mueller for suggesting these exercises.

perform job tasks requiring pattern recognition, learning, and prediction, we should expect AI exposure and traditional software exposure to have different relationships with non-AI hiring. We put this idea to the test in Table 6, controlling as before for Webb’s measure of exposure to software. These estimates find that inclusion of the software exposure variable in our main specification has modest effects on the estimates using the Felten et al. and Webb AI exposure measures, which remain very similar to those in Table 5. The software exposure measure itself has no consistent association with non-AI hiring in these models, with magnitude, sign, and significance varying substantially across specifications and panels.²³

We showed in Figure 2 that AI adoption sharply accelerated after 2015, after having grown comparatively slowly in the prior five years. This discontinuous growth provides an opportunity to test whether any potential association between AI exposure and non-AI hiring fits this timing. We perform this exercise in Table 7, where we break the outcome period of 2010-2018 into two subperiods, 2010-2014 and 2014-2018. We then estimate a subset of the specifications in Table 5 separately for these sub-intervals.

Focusing on Felten et al.’s AI exposure measure in Panel A, we see no economically or statistically significant relationship between AI exposure in 2010 and non-AI hiring during the years 2010-14. This is reassuring since AI adoption had not yet taken off in this period. For 2014-18, we instead see a very different pattern: column 5 shows a strong negative association between AI exposure and subsequent non-AI hiring. The estimate is -11.94 (standard error = 3.80), indicating that a one standard deviation increase in AI exposure is associated with an approximately 12% decline in overall non-AI vacancies. This estimate remains quite stable and actually becomes more precise when we control for firm size deciles, commuting zone controls, and three-digit industry fixed effects in column 6. The relationship also remains comparable when we include in column 7 the baseline shares of sales and administration occupations. The point estimate is now -10.60 with a standard error of 2.82. Column 8 shows that this result is robust to firm fixed effects, though, as before, their inclusions reduces the magnitude of the relationship. Panels B and C report the same specifications for the Webb and SML AI measures, respectively. As in Table 5, the relationships between non-AI hiring and AI exposure are less consistent and robust for these AI exposure indexes.

²³Another prediction of our conceptual framework is that we should also see a decline in postings of vacancies for occupations that have themselves high exposure to AI. When we include the growth of “at-risk” vacancies on the left-hand side, we see consistent negative effects with all three measures. Nevertheless, we are cautious in interpreting these specifications, and do not report them, because they suffer from potential mean reversion. In particular, because the exposure measure is equal to the share of establishment postings that are ‘at-risk’ in 2010-2012, any mean reversion in this measure will induce a spurious negative relationship between an establishment’s at-risk vacancy share in 2010-2012 and its subsequent change.

Table A7 in the Appendix presents results at the firm level, which are, on the whole, quite similar to the establishment-level results. Table A8, on the other hand, depicts similar, even if slightly smaller, estimates with average establishment size weights (rather than baseline establishment size weights as in Table 5).

In sum, we find no evidence that AI exposure is associated with greater hiring at the establishment level. Rather, several of our specifications show statistically significant and economically meaningful negative effects, and these effects concentrate in the years between 2015 and 2018, precisely when we see the surge in AI vacancies.

7.2 AI and Industry Employment

Associated with the surge in AI activity, there may also be changes in industry-level organization that potentially offset or augment the establishment-level consequences. To investigate whether more exposed industries are contracting (or expanding), we aggregate our exposure measure to the industry sector by Commuting Zone level, and then use County Business Patterns data (CBP) between 2000 and 2016. To do this aggregation, we take the 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. Unfortunately, our CBP sample ends in 2016—thus excluding the last several years of rapid AI expansion—because changes to reporting in CBP after 2016 create many incompatibilities.²⁴

The results are reported in the first three columns of Table 8, 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. The standard errors are robust to heteroscedasticity and correlation within commuting zones.

Columns 1 and 2 examine trends in industry employment during 2003-2007 and 2007-2010, before the major AI advances that followed. These columns show that industry AI exposure in 2010 does not predict differential employment behavior before 2010. This result implies that three-digit industries with different levels of AI exposure were on roughly parallel trends before the pickup in AI activity in the late 2010s. Yet, 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 effects from industry AI exposure: a one standard deviation increase in industry AI exposure should be associated with a small 0.049% decline in industry employment.

²⁴In 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/>.

Panels B and C of the table show similar results using the Webb and SML measures of AI exposure in place of the Felten et al. measure.

This set of null results may indicate that it is too soon to see the impact of AI activity on industry reorganization or growth. It could also indicate that much of the effect of AI on employment, if eventually present, will occur within industries.

7.3 Employment and Wages in AI-Exposed Occupations

Columns 4 through 9 of Table 8 explore variation across occupations, in particular, assessing whether more occupations with greater AI exposure exhibit differential employment or wage trends after the onset of rapid AI hiring. For this analysis, we use occupational employment and wage information from the US BLS Occupational Employment Statistics (OES) data series. This establishment-based data series is likely to provide more accurate estimates of changes in employment and wages within occupations than is available from household surveys.

The observations in this table are at the six-digit occupation level, and the dependent variable is the sum of employment in a six-digit occupation across all industries (excluding 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 change in occupational employment, while in columns 7 through 9 it is change in occupational (log) wage.

The results for employment and wage growth using each of the AI exposure measures parallel our industry-level results: we detect no differential employment or wage behavior in more AI-exposed occupations after 2010.²⁵

In summary, we find no robust effects of AI activity on more AI-exposed industries or occupations. These occupation results are also consistent with the patterns Felten et al. (2019) report.

7.4 Interpretation

What do the totality of the results reported about the labor market effects of AI? Our interpretation of the evidence is that AI is having real effects: there is a significant surge in vacancies for AI workers in establishments with task structures that are more suitable

²⁵Differently from our industry results, we detect a significantly faster increase in the employment of more exposed occupations between 2007 and 2010 when using the Felten et al. AI measure. This 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 large number of point estimates report in this table.

for AI; skill churn increases differentially at AI-exposed establishments, with both greater retirement of previously-posted skills and greater introduction of not-previously-posted skills; and finally, AI-exposed establishments show some evidence of reduced non-AI hiring. These patterns are consistent with the hypothesis that AI-powered algorithms are being used, at least in part, for substituting for human skills—though of course, our evidence concerns the component of AI adoption driven by baseline task structure, and is therefore silent on the consequences of other types of AI technologies.

Our evidence is also fairly clear that there is no systematic aggregate relationship between industry and occupation-level AI exposure and either changes in industry or occupation-level employment or changes in occupational wage levels. Our overall interpretation is that, while AI technologies appear to be changing task and skill composition at exposed establishments and firms, any aggregate effects of AI, if present, are too small to detect.

8 CONCLUSION

There is much excitement and quite a bit of apprehension about AI and its labor market effects. In this paper, we explored the nature of AI activity in the US labor market and its consequences for skill change, hiring, and industry and occupation level changes in employment and earnings. We have four main findings:

1. We see a surge in AI activity, particularly after 2014, proxied by vacancies seeking workers with AI skills, and this surge is driven by establishments with high exposure to AI—meaning that their task structure in 2010 is suitable for the introduction of AI technologies. This pattern is highly robust with two of our three AI-exposure measures (those based on the indices constructed by Felten et al., 2018, 2019, and Webb, 2020) and still present but less robust with our third measure, the SML index (based on the work by Brynjolfsson, Mitchell and Rock, 2018, 2019).
2. We estimate consistent and robust changes in the skills demanded by high-exposure establishments. In particular, establishments with task structures suitable for AI cease to post vacancies that list a range of previously-posted skills and start posting additional skill requirements. This evidence suggests that some of the tasks that workers used to perform in these establishments are no longer being performed, while simultaneously they introduce new tasks.
3. With two of our three measures, we also estimate that establishments with high exposure to AI reduce their non-AI hiring. These results are statistically significant,

economically sizable, and robust with the Felten et al. measure and robust in most specifications with the Webb measure. We do not detect such negative employment effects with the SML index. Because the relationship between AI exposure and AI hiring is also much less robust and consistent with the SML measure, we view this evidence as suggestive but inconclusive.

4. We do not detect any negative relationship between AI exposure and employment or wage levels at more aggregated levels. Specifically, we find no AI-employment relationship in more exposed industries or occupations and no AI-wage relationship in more exposed occupations.

Taken together, we interpret these results as indicating that there is a genuine increase in the deployment of AI technologies throughout the US economy and this increase is much more pronounced in establishments that have task structures suitable for deploying AI-powered algorithms. Reflecting the potential for substitution between algorithms and humans, we detect nontrivial reductions in non-AI hiring in more AI-exposed establishments, though these estimates are not uniformly robust with all three of our AI exposure measures. We do not detect any aggregate relationships between AI exposure and either employment or wages, plausibly because AI technologies are still in their infancy and have spread only to a limited part of the US economy.

Our results raise several important questions and also have evident shortcomings. First, it is important to further explore and understand the juxtaposition of negative relationships between AI exposure and non-AI hiring at the level of establishments that are not present at the level of industries and occupations. Second, our focus on AI adoption that is driven by the task structure of establishments may leave out other types of AI impacts that are less related to task structures, such as the use of AI to launch new products and services. Our estimates are not informative about AI applications that are missed by our AI exposure measures. Finally, because the next generation of AI-powered technologies will likely have different capabilities and applications from the current generation, our results do not foretell whether future AI technologies will prove more complementary or more substitutable with human capabilities.

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 (2020) “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 Eco-*

nomics, 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*.

Dillender, Marcus, and Eliza Forsythe (2019) “Computerization of white collar jobs.” Unpublished manuscript. University of Illinois at Urbana-Champaign.

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.

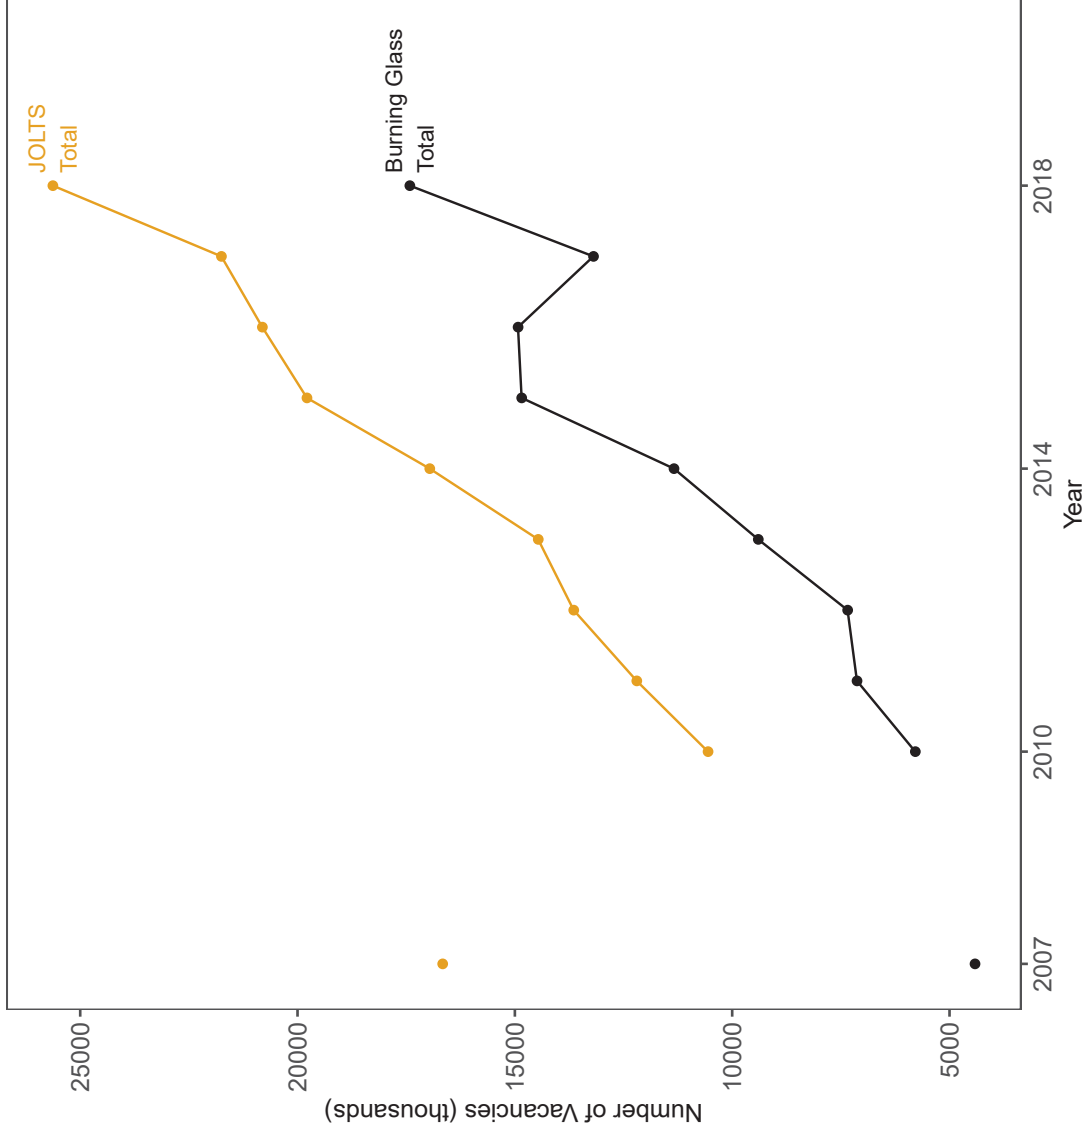
Felten, Edward, Manav Raj, and Robert Seamans (2019) “The Effect of Artificial Intelligence on Human Labor: An Ability-based Approach.” *Academy of Management Proceedings* Vol. 1.

- 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.
- Krueger, Alan B. (1993)** “How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989” *The Quarterly Journal of Economics* 108(1): 33–60.
- 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.

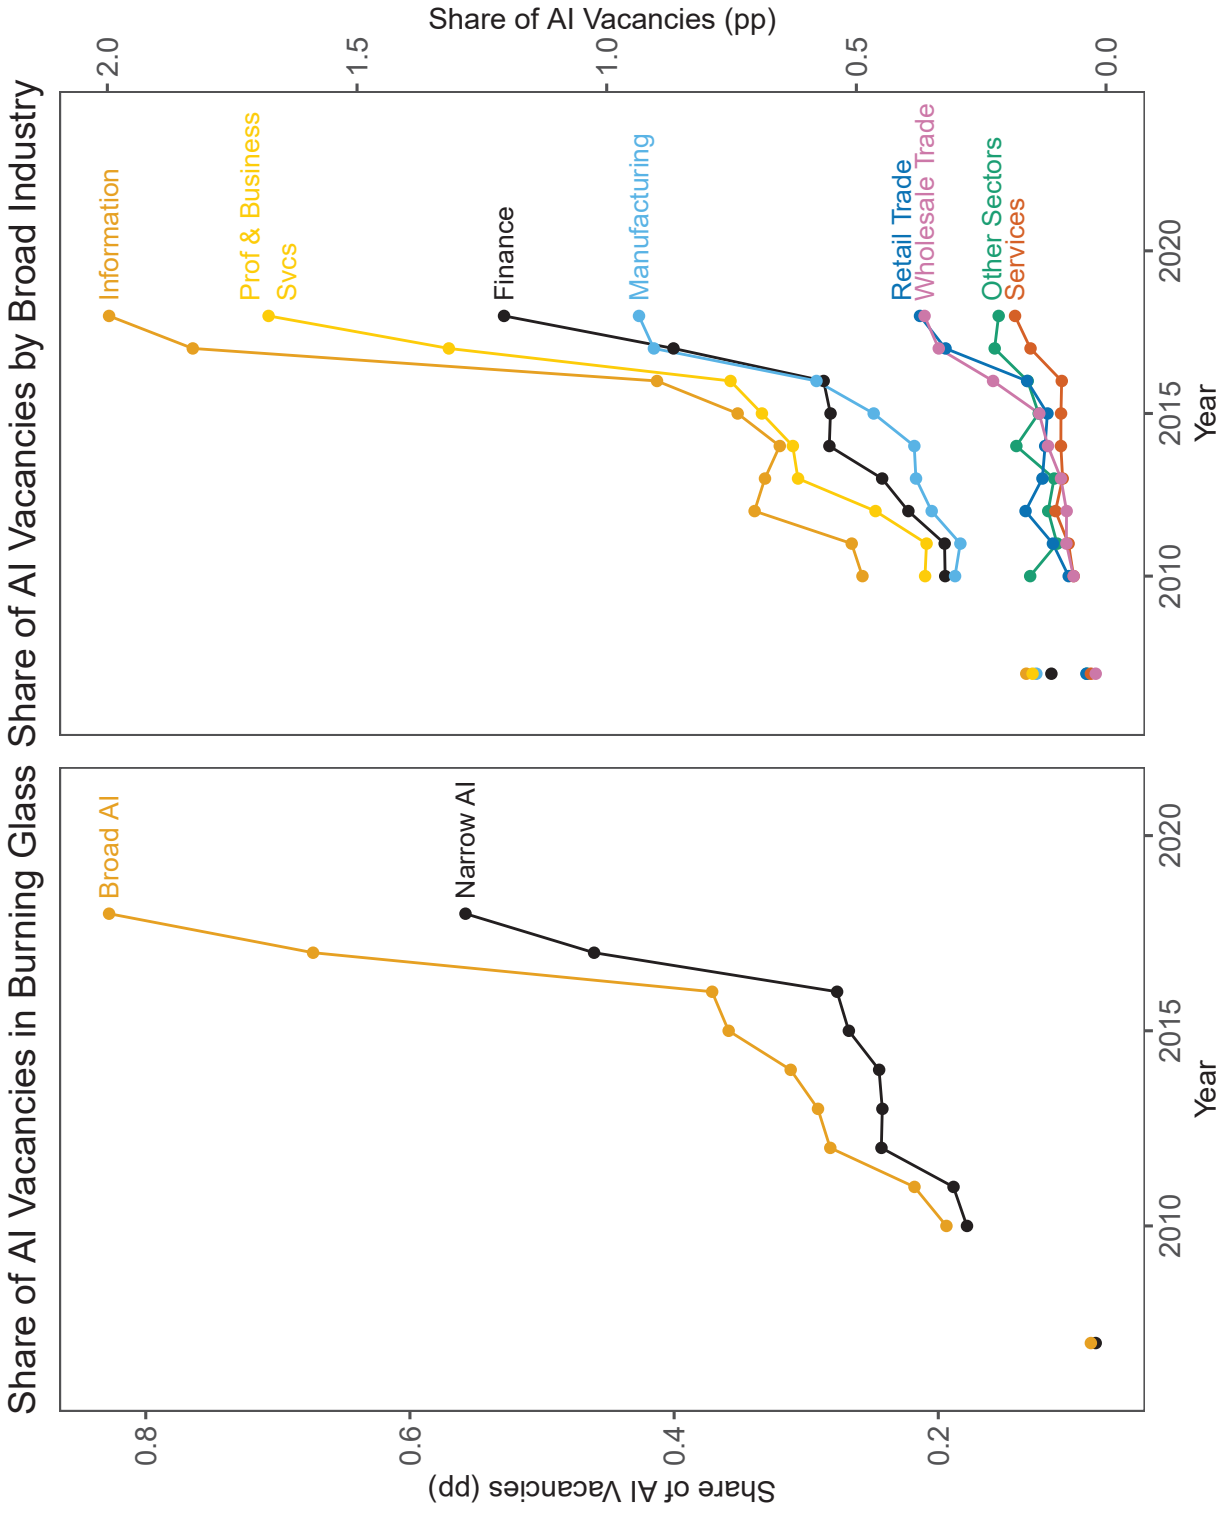
West, Darrell M (2018) *The Future of Work: Robots, AI, and Automation*. Brookings
Institution Press.

FIGURE 1: Vacancies in Burning Glass and JOLTS



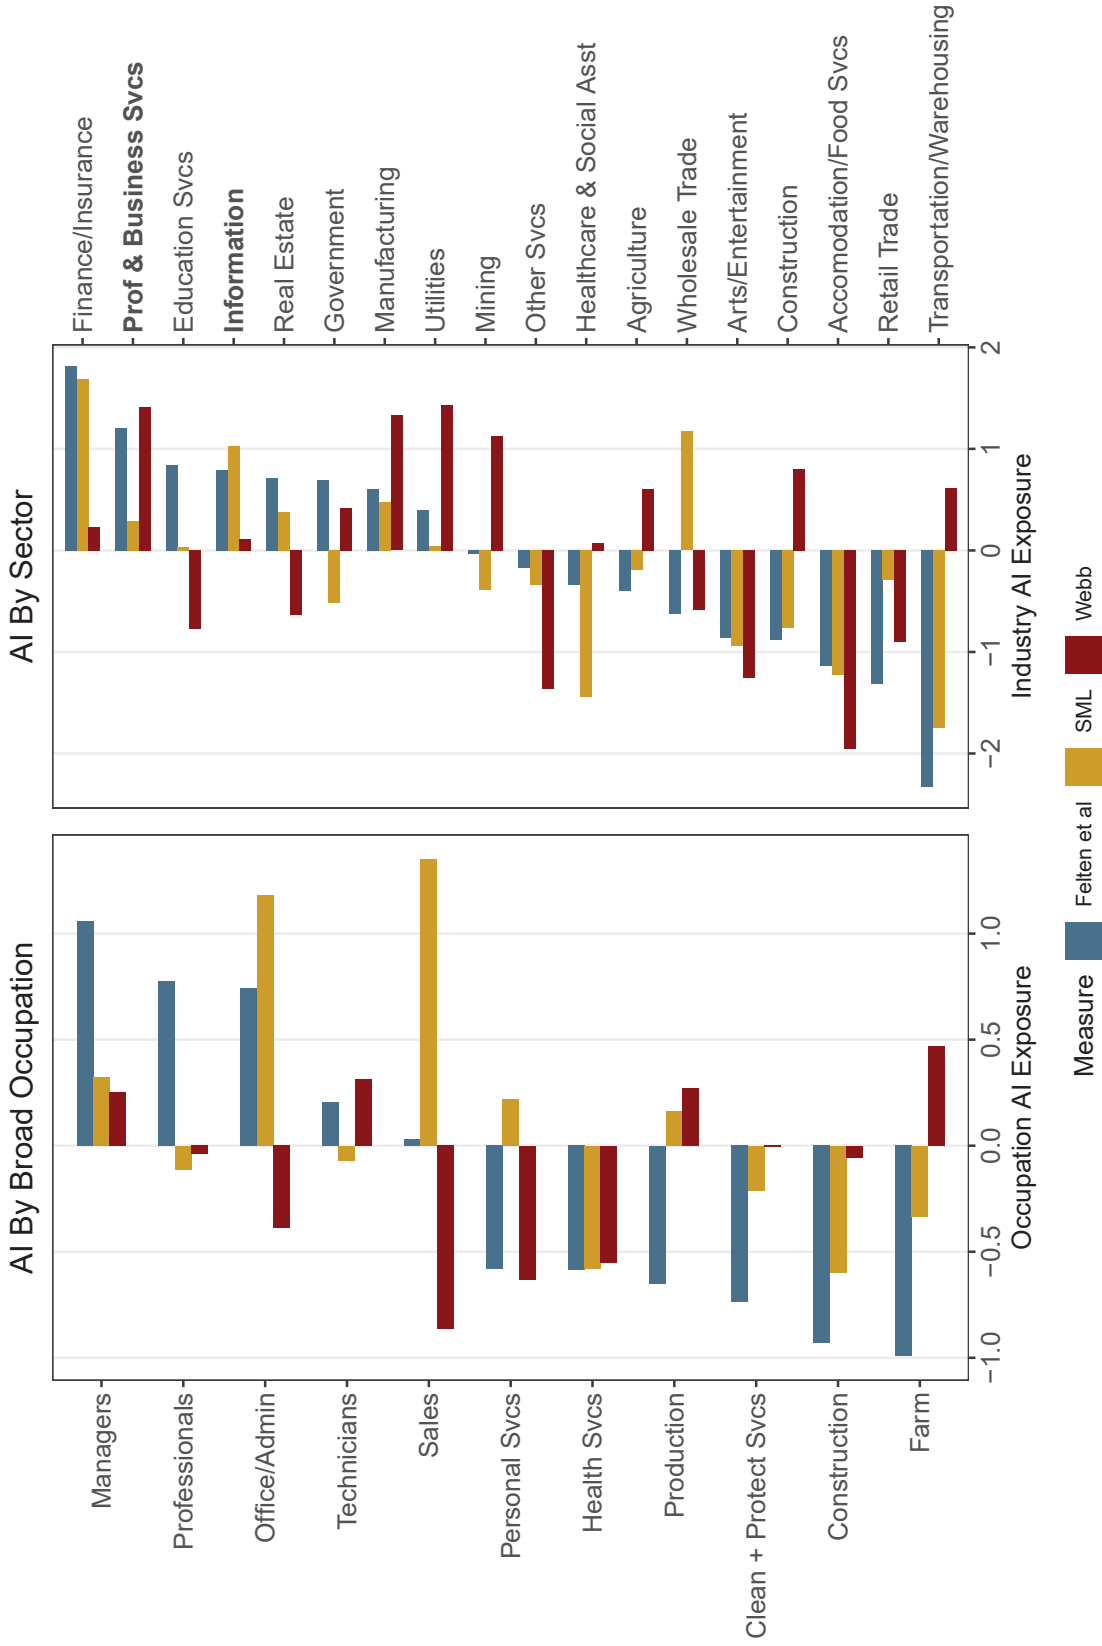
This figure plots the total number of vacancies in JOLTS and the total number of vacancies in Burning Glass, by year. We multiply the number of job openings in JOLTS by a constant factor, to arrive at a number of vacancies that matches the concept of a vacancy in Burning Glass. This method follows Carnevale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

FIGURE 2: Share of AI Vacancies in Burning Glass



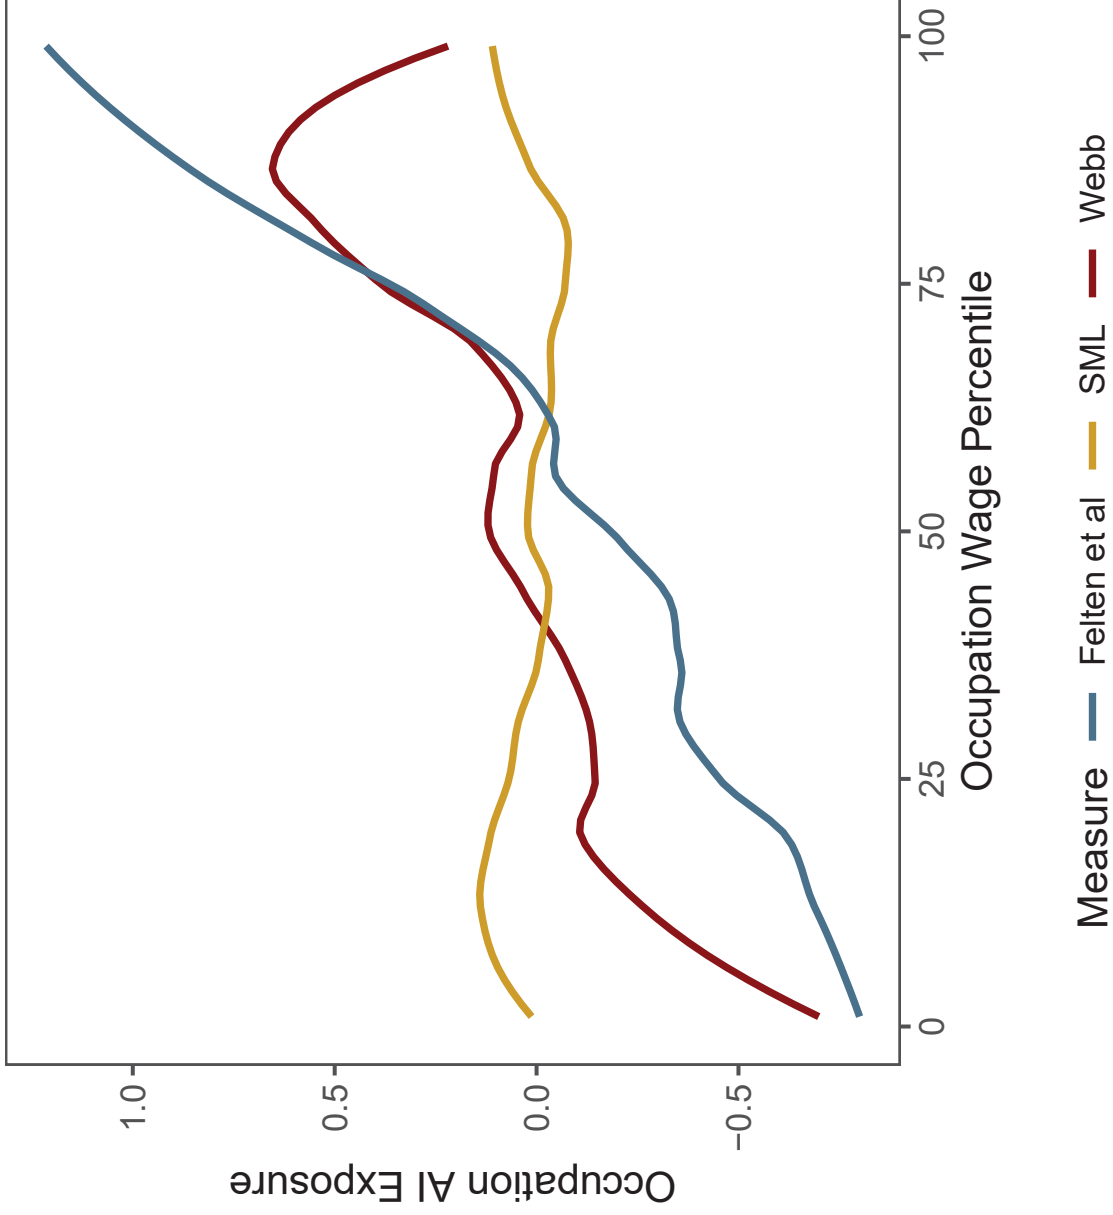
The left panel plots the share of vacancies in Burning Glass, that post a skill in the Broad or Narrow AI categories, as defined in the main text. The right panel plots the share of narrow AI vacancies in Burning Glass, by year, in each broad industry grouping.

FIGURE 3: AI Exposure by Broad Occupation and Sector



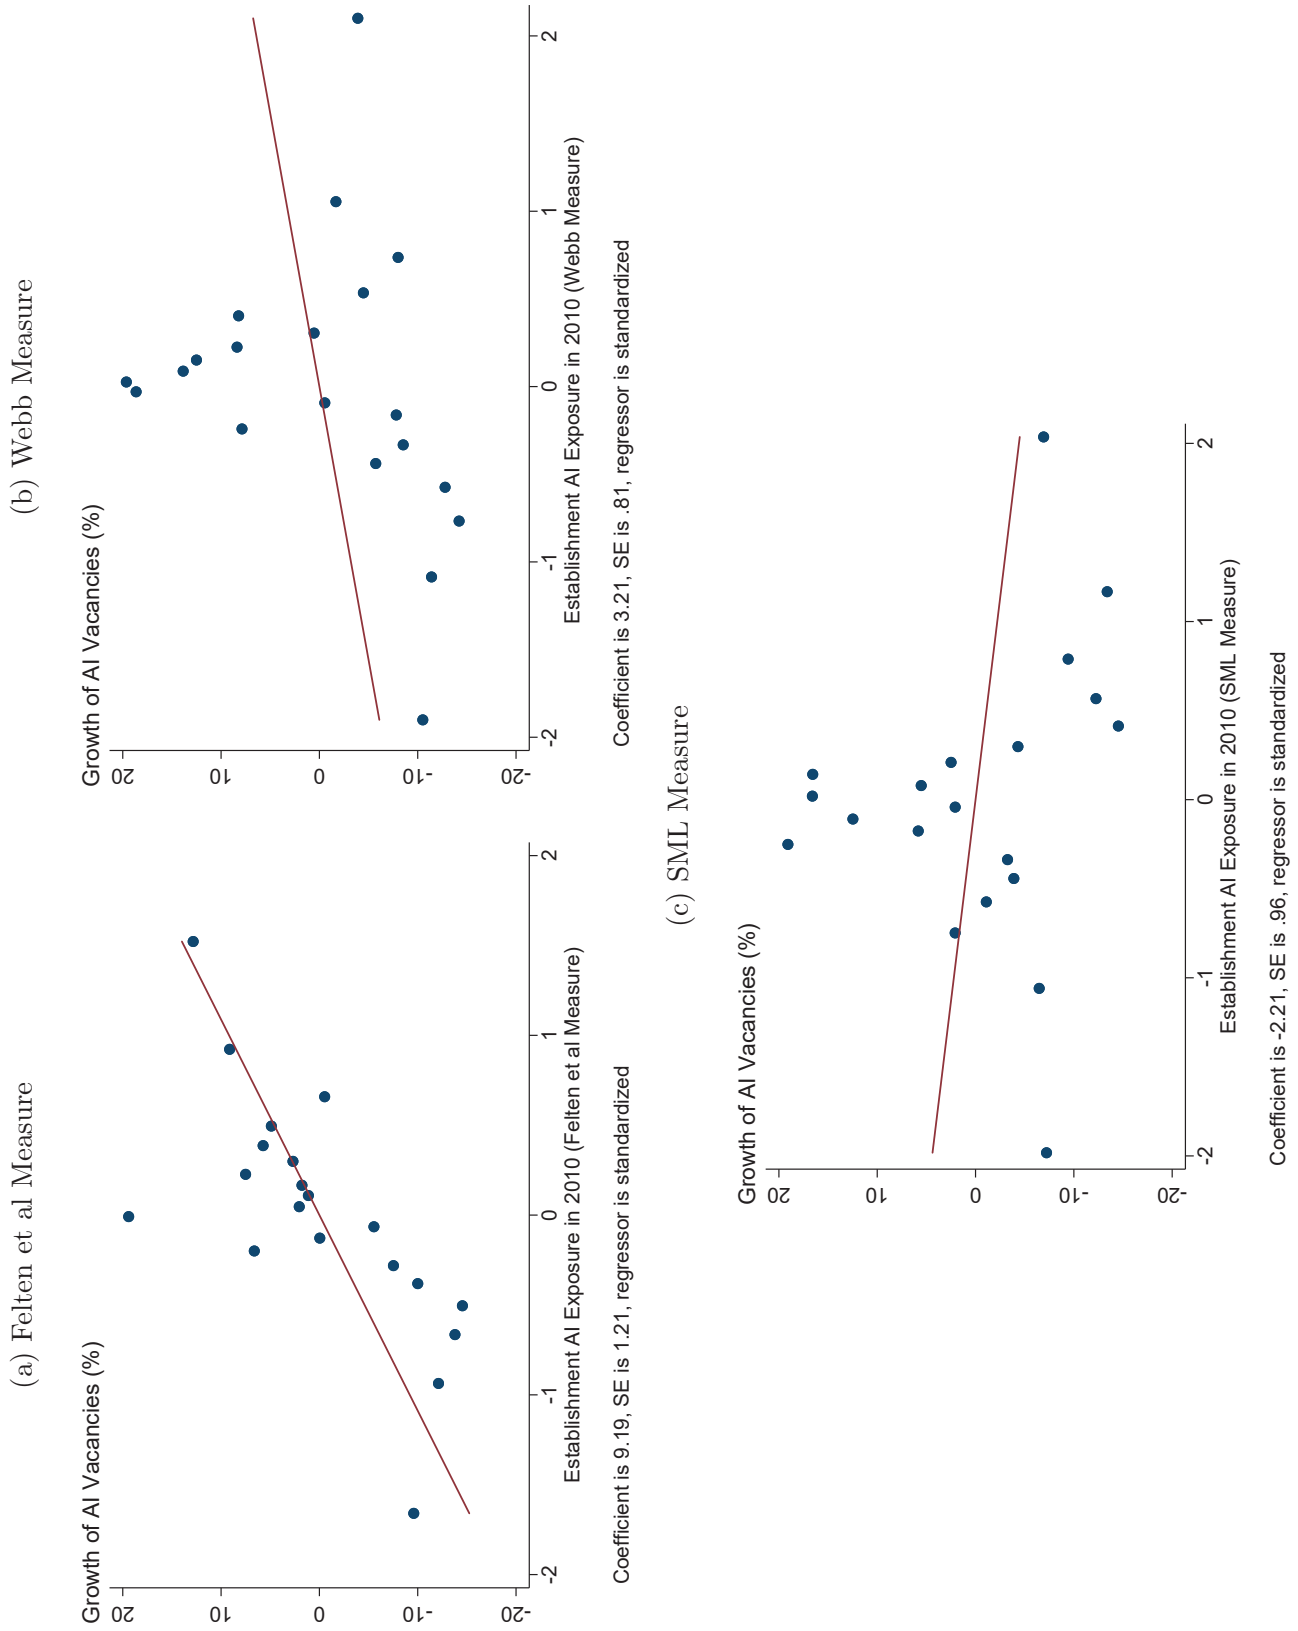
The left panel plots the average of the standardized measures of AI exposure across broad occupations. The right panel plots the average of the standardized measures of AI exposure across 2-digit NAICS sectors, by taking the mean across the 6 digit SOC occupations posted in each 2 digit NAICS sector, weighted by the number of vacancies posted by each sector in each occupation.

FIGURE 4: AI Exposure by Occupation Wages



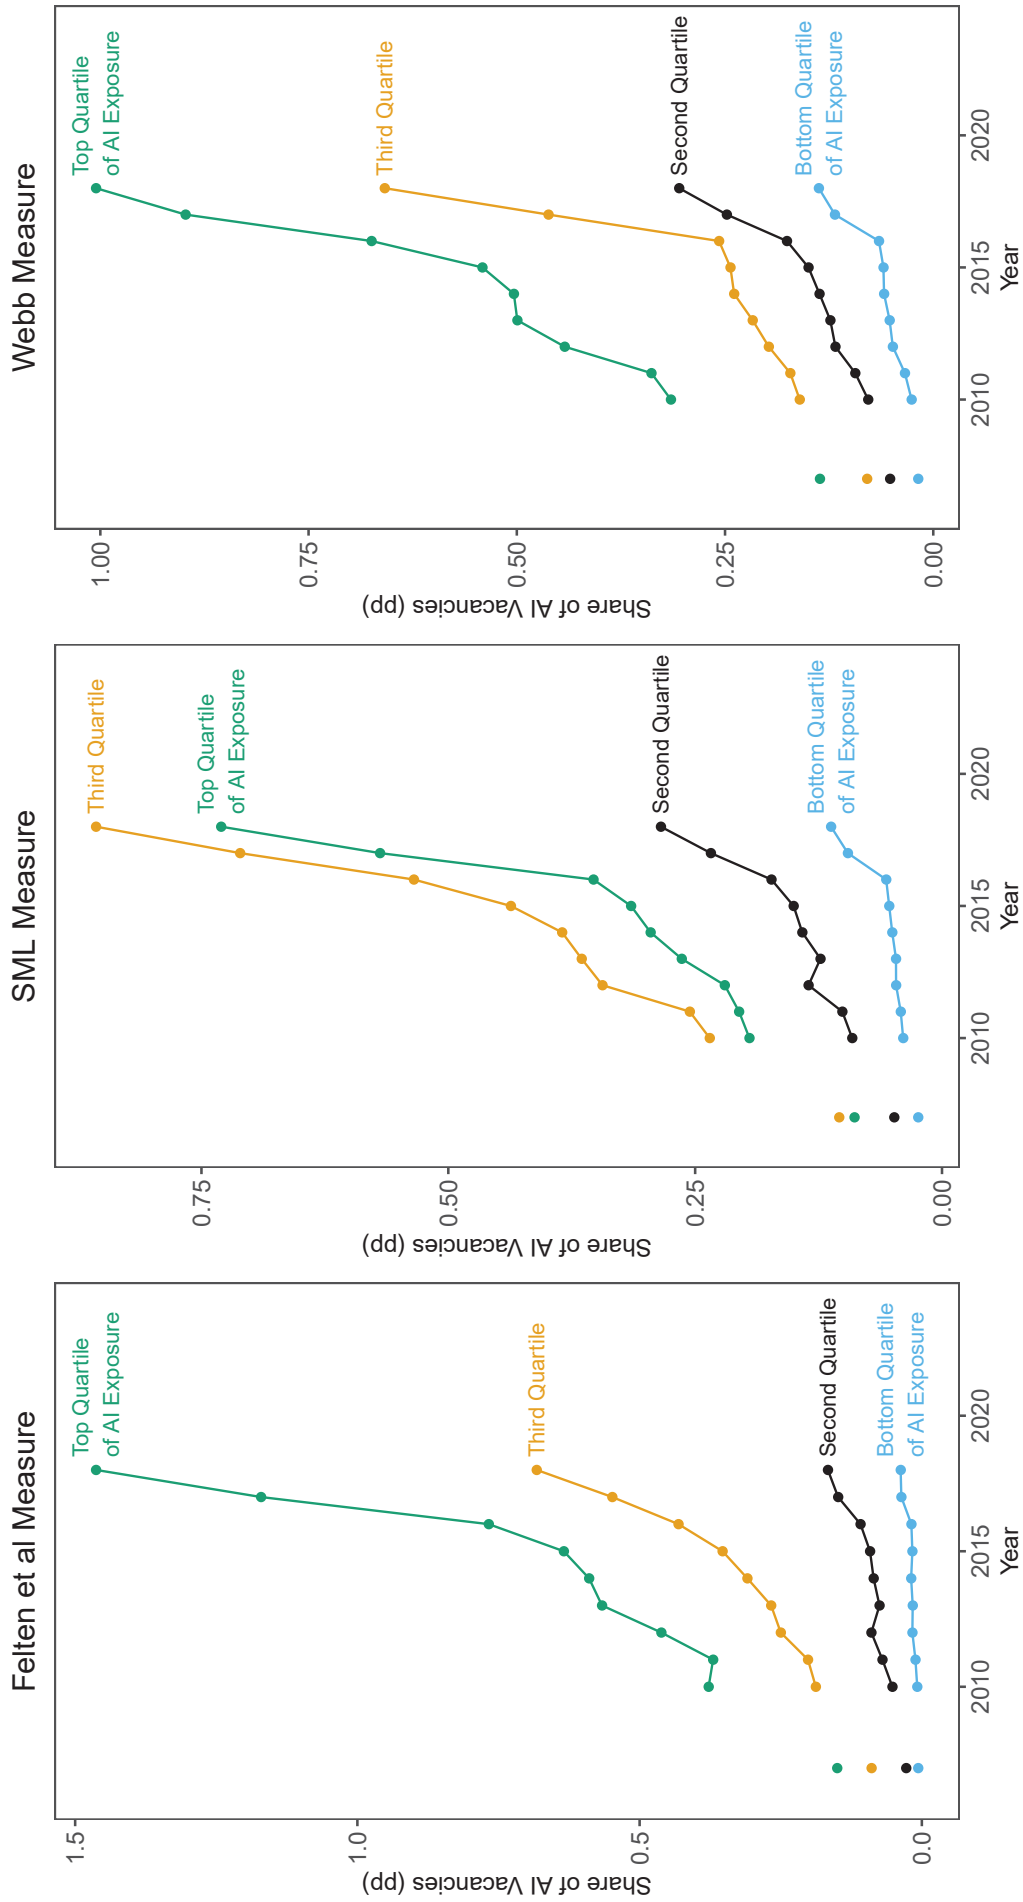
The figure plots a smoothed polynomial regression of the (standardized) measures of AI exposure in each 6-digit SOC occupation against its rank in the wage distribution. We rank occupations according to their mean hourly wage for 2010-2018, obtained from the Occupational Employment Statistics.

FIGURE 5: Binscatter of AI Growth and Establishment AI Exposure



The figure presents binned scatter plots that summarize the relationship between establishment AI Exposure in 2010 and the growth of AI establishment vacancies between 2010 and 2018. The covariates from column 3 of Table 1 are partialled out. The solid line corresponds to a regression with 2010 establishment vacancies as the weight. The corresponding point estimates and standard errors are reported at the bottom of each panel. Panel A uses the measure of AI exposure from Felten et al. (2019). Panel B uses the Webb (2020) measure. Panel C uses the SML measure, from Brynjolfsson et al. (2019).

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



This figure plots establishments' share of AI vacancies in Burning Glass, for each quartile of the distribution of 2010 establishment AI exposure, after partialling out their 2010-2012 share of vacancies in sales and administration. 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

	(1)	(2)	(3)	(4)	(5)	(6)
Growth of Establishment AI Vacancies, 2010-2018						
<i>Panel A: Felten et al. Measure of AI Exposure</i>						
Establishment AI Exposure, 2010	15.96*** (1.73)	13.82*** (1.43)	9.19*** (1.21)	16.53*** (1.89)	9.75*** (1.20)	16.87*** (1.86)
Observations	1,075,474	1,075,474	954,519	1,075,474	954,518	762,672
<i>Panel B: Webb Measure of AI Exposure</i>						
Establishment AI Exposure, 2010	6.59*** (1.13)	5.08*** (0.96)	3.21*** (0.81)	5.91*** (1.27)	0.42 (0.82)	1.14 (1.08)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	824,803
<i>Panel C: SML Measure of AI Exposure</i>						
Establishment AI Exposure, 2010	3.76*** (1.19)	2.30** (1.04)	-2.21** (0.96)	-3.04** (1.38)	1.95** (0.89)	4.47*** (1.34)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	824,803
<i>Covariates:</i>						
Share of Vacancies in Sales & Admin, 2010		✓			✓	✓
<i>Fixed Effects:</i>						
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 sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). 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 standardized 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. 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 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 size (defined as the total vacancies posted by an establishment's firm in 2010-2012). 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 establishment 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 AI Vacancy Growth, Controlling for Software Exposure

	(1)	(2)	(3)	(4)	(5)	(6)
Growth of Establishment AI Vacancies, 2010-2018						
	<i>Panel A: Felten et al. Measure of AI Exposure</i>					
Establishment AI Exposure, 2010	16.28*** (1.74)	14.10*** (1.44)	9.63*** (1.23)	17.43*** (1.95)	9.96*** (1.24)	17.47*** (1.90)
Estab. Software Exposure	2.36*** (0.76)	2.24*** (0.71)	2.62*** (0.76)	4.83*** (1.23)	0.66 (0.82)	1.83 (1.19)
Observations	1,059,620	1,059,620	941,046	1,059,620	941,046	1,059,620
	<i>Panel B: Webb Measure of AI Exposure</i>					
Establishment AI Exposure, 2010	10.64*** (1.83)	7.88*** (1.50)	3.85*** (1.10)	6.81*** (1.52)	1.50 (1.07)	2.57* (1.41)
Estab. Software Exposure	-6.28*** (1.51)	-4.27*** (1.26)	-0.96 (0.96)	-1.44 (1.34)	-1.81* (1.00)	-2.54* (1.49)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789
	<i>Panel C: SML Measure of AI Exposure</i>					
Establishment AI Exposure, 2010	4.14*** (1.18)	2.64** (1.04)	-1.96** (0.96)	-2.42* (1.28)	1.90** (0.88)	4.37*** (1.28)
Estab. Software Exposure	1.56* (0.80)	1.44* (0.77)	1.09 (0.69)	2.40** (1.04)	-0.90 (0.78)	-0.76 (1.11)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789
<i>Covariates:</i>						
Share of Vacancies in Sales & Admin, 2010					✓	✓
<i>Fixed Effects:</i>						
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, controlling for establishment software exposure. Our measure of software exposure is from Webb (2020). Establishment software exposure is the standardized mean of occupation software 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. The sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). 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 standardized 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. 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 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 size (defined as the total vacancies posted by an establishment's firm in 2010-2012). 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 establishment 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 Establishment Negative Skill Change, 2010-2018

	(1)	(2)	(3)	(4)	(5)	(6)
Establishment Negative Skill Change, 2010-2018						
<i>Panel A: Felten et al. Measure of AI Exposure</i>						
Establishment AI Exposure, 2010	0.83*** (0.09)	0.83*** (0.09)	0.97*** (0.07)	0.50*** (0.05)	1.00*** (0.07)	0.54*** (0.05)
Observations	339,282	339,282	322,901	339,282	322,901	339,282
<i>Panel B: Webb Measure of AI Exposure</i>						
Establishment AI Exposure, 2010	0.62*** (0.11)	0.60*** (0.11)	0.45*** (0.06)	0.20*** (0.04)	0.68*** (0.11)	0.34*** (0.04)
Observations	353,107	353,107	335,589	353,107	335,589	353,107
<i>Panel C: SML Measure of AI Exposure</i>						
Establishment AI Exposure, 2010	0.53*** (0.08)	0.52*** (0.07)	0.32*** (0.07)	0.26*** (0.04)	0.46*** (0.09)	0.36*** (0.04)
Observations	353,107	353,107	335,589	353,107	335,589	353,107
<i>Covariates:</i>						
Share of Vacancies in Sales & Admin, 2010					✓	✓
<i>Fixed Effects:</i>						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓		✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment negative skill change. The sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). The outcome variable, constructed from Burning Glass data, is establishment negative skill change for 2010-2018, as defined in the main text. The sample is establishments posting in both 2010 and 2019. The regressor, establishment AI exposure in 2010, is the standardized 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. 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 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 size (defined as the total vacancies posted by an establishment's firm in 2010-2012). 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 establishment 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 4: Effects of AI Exposure on Establishment Positive Skill Change, 2010-2018

	(1)	(2)	(3)	(4)	(5)	(6)
Establishment Positive Skill Change, 2010-2018						
<i>Panel A: Felten et al. Measure of AI Exposure</i>						
Establishment AI Exposure, 2010 Observations	0.95*** (0.08) 339,282	0.94*** (0.09) 339,282	0.58*** (0.09) 322,901	0.02 (0.04) 339,282	0.62*** (0.09) 322,901	0.05 (0.04) 339,282
<i>Panel B: Webb Measure of AI Exposure</i>						
Establishment AI Exposure, 2010 Observations	0.69*** (0.09) 353,107	0.66*** (0.09) 353,107	0.26*** (0.08) 335,589	-0.01 (0.03) 353,107	0.43*** (0.08) 335,589	0.13*** (0.04) 353,107
<i>Panel C: SML Measure of AI Exposure</i>						
Establishment AI Exposure, 2010 Observations	0.62*** (0.09) 353,107	0.59*** (0.09) 353,107	0.19** (0.09) 335,589	0.10** (0.04) 353,107	0.26*** (0.09) 335,589	0.03 (0.04) 353,107
<i>Covariates:</i>						
Share of Vacancies in Sales & Admin, 2010					✓	✓
<i>Fixed Effects:</i>						
Firm Size Decile		✓	✓		✓	
Commuting Zone		✓	✓	✓	✓	✓
3 digit Industry			✓		✓	
Firm				✓		✓

This table presents estimates of the effects of establishment AI exposure on establishment positive skill change. The sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). The outcome variable, constructed from Burning Glass data, is establishment positive skill change for 2010-2018, as defined in the main text. The sample is establishments posting in both 2010 and 2019. The regressor, establishment AI exposure in 2010, is the standardized 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. 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 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 size (defined as the total vacancies posted by an establishment's firm in 2010-2012). Column 4 contains only establishment AI exposure. Columns 2-3 and 5 include fixed effects for the decile of 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 establishment 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 5: Effects of AI Exposure on Establishment Non-AI Vacancy Growth, 2010-2018

Growth of Establishment Non-AI Vacancies, 2010-2018								
Full Sample								
	(1)	(2)	(3)	(4)	(5)	(6)	(7) (8)	
Establishment AI Exposure, 2010	-13.80*** (4.22)	-16.36*** (4.11)	-11.90*** (4.08)	-4.81*** (1.44)	-12.42*** (4.01)	-4.04*** (1.47)	-8.38** (3.46)	-3.56* (1.86)
Observations	1,075,474	1,075,474	954,519	1,075,474	954,519	1,075,474	324,901	341,525
<i>Panel B: Webb Measure of AI Exposure</i>								
Establishment AI Exposure, 2010	-17.24*** (3.72)	-18.21*** (3.63)	-6.73** (3.01)	-2.22** (0.93)	-8.30** (3.70)	1.51 (0.98)	-4.70* (2.66)	-1.44 (1.36)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789	337,758	355,529
<i>Panel C: SML Measure of AI Exposure</i>								
Establishment AI Exposure, 2010	7.02** (3.13)	5.74* (3.01)	2.05 (2.92)	0.95 (1.16)	2.21 (3.61)	-3.01** (1.22)	0.01 (2.94)	-0.91 (1.38)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789	337,758	355,529
<i>Covariates:</i>								
Share of Vacancies in Sales, Admin. in 2010					✓			
<i>Fixed Effects:</i>								
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 sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). 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 establishments in 2010-12 and 2016-18. The regressor, establishment AI exposure in 2010, is the standardized 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. 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 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 size (defined as the total vacancies posted by an establishment's firm in 2010-2012). 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 6: Effects of AI Exposure on Establishment Non-AI Vacancy Growth, Controlling for Software Exposure

Growth of Establishment Non-AI Vacancies, 2010-2018								
Full Sample								
	(1)	(2)	(3)	(4)	(5)	(6)	(8)	
Establishment AI Exposure, 2010	-14.62*** (4.28)	-17.21*** (4.17)	-12.02*** (4.00)	-5.43*** (1.49)	-12.47*** (3.94)	-3.93** (1.54)	-8.68*** (3.36)	-3.73* (1.90)
Estab. Software Exposure	-7.50** (3.68)	-7.35** (3.60)	0.66 (3.09)	-1.77 (1.08)	1.07 (3.28)	2.03* (1.17)	-1.67 (3.41)	-0.41 (1.40)
Observations	1,059,620	1,059,620	941,046	1,059,620	941,046	1,059,620	322,187	338,645
<i>Panel B: Webb Measure of AI Exposure</i>								
Establishment AI Exposure, 2010	-23.04*** (4.61)	-25.36*** (4.47)	-14.04*** (4.68)	-3.06*** (1.05)	-14.95*** (5.59)	-0.29 (1.09)	-7.30 (4.51)	-2.56 (1.56)
Estab. Software Exposure	9.01** (4.32)	10.88*** (4.22)	10.98** (4.79)	1.36 (1.14)	11.16** (5.11)	3.19*** (1.18)	3.80 (5.26)	1.74 (1.54)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789	337,758	355,529
<i>Panel C: SML Measure of AI Exposure</i>								
Establishment AI Exposure, 2010	5.97* (3.15)	4.70 (3.03)	2.62 (2.93)	0.84 (1.12)	2.40 (3.63)	-2.67** (1.20)	-0.17 (2.89)	-0.94 (1.39)
Estab. Software Exposure	-4.41 (3.66)	-4.37 (3.60)	2.46 (3.10)	-0.43 (0.95)	3.04 (3.30)	2.80*** (1.03)	-0.85 (3.36)	-0.13 (1.34)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789	337,758	355,529
<i>Covariates:</i>								
Share of Vacancies in Sales & Admin, 2010					✓	✓		
<i>Fixed Effects:</i>								
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, controlling for establishment software exposure. Our measure of software exposure is from Webb (2020). Establishment software exposure is the standardized mean of occupation software 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. The sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). 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 standardized 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. 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 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 size (defined as the total vacancies posted by an establishment's firm in 2010-2012). 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 7: Effects of AI Exposure on Establishment Non-AI Vacancy Growth, 2010-2014 and 2014-2018

	Growth of Establishment Non-AI Vacancies							
	2010-2014 Growth			2014-2018 Growth				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Establishment AI Exposure, 2010 Observations	-1.86 (4.77) 1,075,474	-0.59 (3.52) 954,519	-1.82 (3.46) 954,519	0.39 (1.11) 1,075,474	-11.94*** (3.80) 1,075,474	-11.32*** (2.93) 954,519	-10.60*** (2.82) 954,519	-5.21*** (1.02) 1,075,474
Establishment AI Exposure, 2010 Observations	-7.51** (3.38) 1,159,789	-2.57 (2.17) 1,021,673	-6.04** (2.66) 1,021,673	-0.35 (0.64) 1,159,789	-9.73*** (2.42) 1,159,789	-4.16** (1.83) 1,021,673	-2.26 (2.16) 1,021,673	-1.86*** (0.71) 1,159,789
Establishment AI Exposure, 2010 Observations	3.73 (2.66) 1,159,789	1.17 (2.46) 1,021,673	2.79 (3.08) 1,021,673	1.90*** (0.73) 1,159,789	3.30 (2.09) 1,159,789	0.88 (2.30) 1,021,673	-0.58 (2.75) 1,021,673	-0.95 (0.91) 1,159,789
<i>Covariates:</i>								
Share of Vacancies in Sales, Admin. in 2010			✓				✓	
<i>Fixed Effects:</i>								
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, separately for 2010-2014 and 2014-2018. The sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). In the first four columns, the outcome variable, constructed from Burning Glass data, is the growth rate of non-AI vacancies between 2010 and 2014, 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 2013-15. In the last four columns, the outcome variable is the change in the inverse hyperbolic sine of the number of vacancies posted by the establishments in 2013-15 and 2016-18. The regressor, establishment AI exposure in 2010, is the standardized 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. 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 covariates included in each model are reported at the bottom of the table. Columns 1 and 5 contains only establishment AI exposure. Columns 2, 3, 6 and 7 include fixed effects for the decile of firm size (defined as the total vacancies posted by an establishment's firm in 2010-2012); and also include 3 digit NAICS industry fixed effects. All columns other than 1 and 5 include commuting zone fixed effects. Columns 4 and 8 include firm fixed effects. Columns 3 and 7 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.

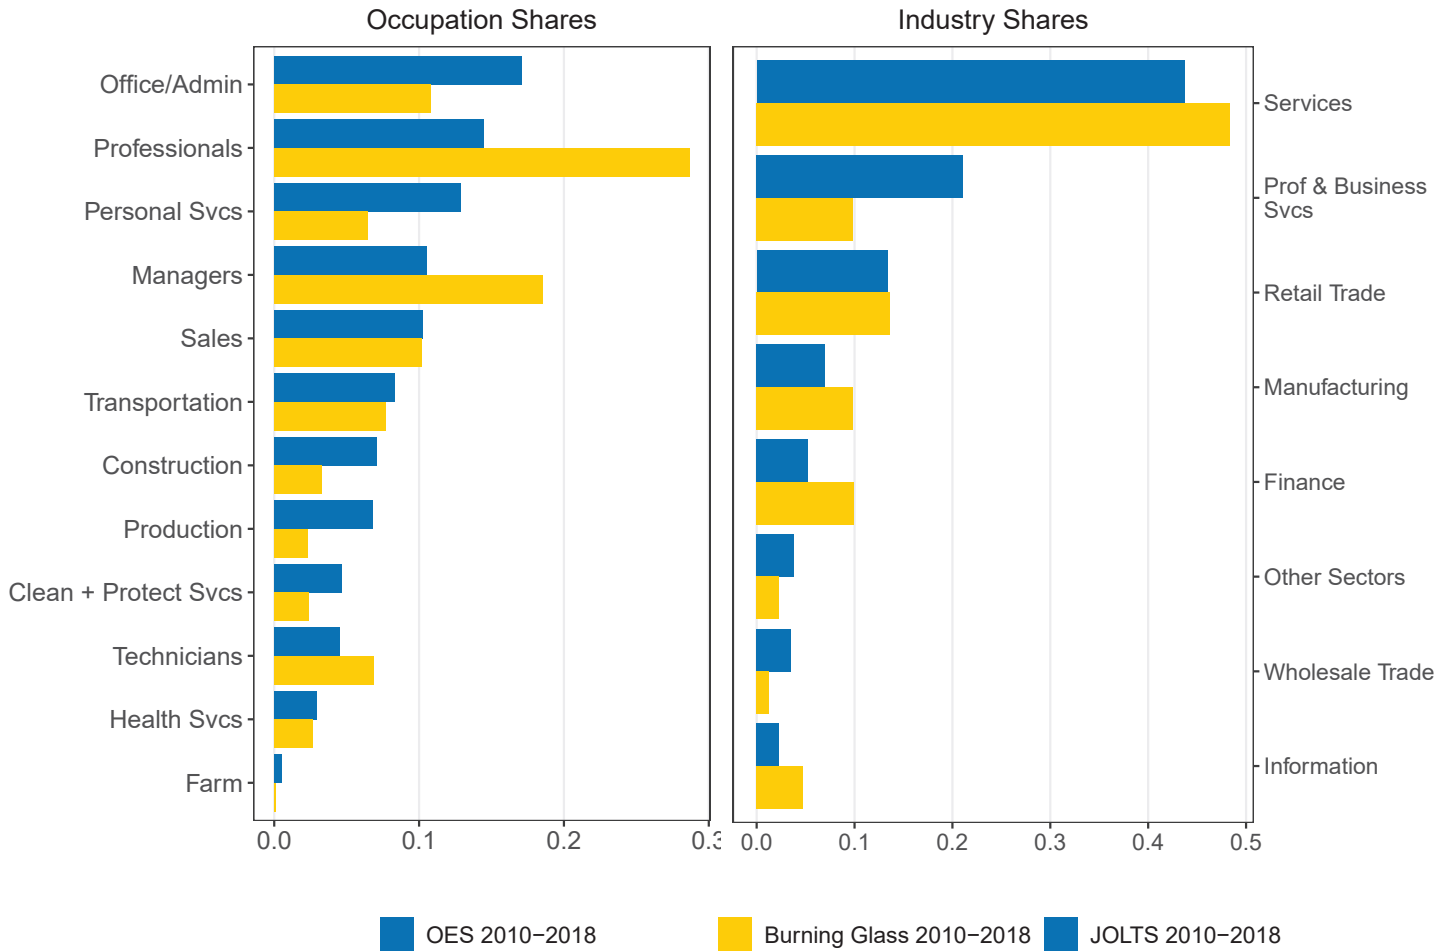
TABLE 8: Effects of AI Exposure on Market Employment and Wage Growth

	Industry by CZ Employment Growth (CBP)			Occupation Employment Growth (OES)			Occupation Wage Growth (OES)		
	2003-2007 (1)	2007-2010 (2)	2010-2016 (3)	2004-2007 (4)	2007-2010 (5)	2010-2018 (6)	2004-2007 (7)	2007-2010 (8)	2010-2018 (9)
Market AI Exposure, 2010 Observations	0.03 (0.17) 10,937	0.10 (0.20) 10,926	-0.05 (0.08) 10,929	<i>Panel A: Felten et al. Measure of AI Exposure</i>					
				0.34 (0.34) 736	0.86*** (0.32) 700	0.51 (0.35) 680	-0.00 (0.17) 680	0.02 (0.20) 648	-0.17*** (0.06) 629
Market AI Exposure, 2010 Observations	0.10 (0.15) 10,981	0.18 (0.17) 10,968	0.11 (0.09) 10,968	<i>Panel B: Webb Measure of AI Exposure</i>					
				0.00 (0.17) 713	0.11 (0.21) 704	-0.17 (0.29) 717	0.11 (0.08) 660	-0.05 (0.10) 653	-0.02 (0.04) 663
Market AI Exposure, 2010 Observations	-0.14 (0.17) 10,981	0.37** (0.18) 10,968	-0.01 (0.08) 10,968	<i>Panel C: SML Measure of AI Exposure</i>					
				0.00 (0.25) 713	-0.17 (0.29) 704	-0.37 (0.25) 717	-0.03 (0.08) 660	0.18 (0.12) 653	0.04 (0.05) 663
<i>Covariates:</i>									
Share of Vacancies in Sales, Admin. in 2010	✓	✓	✓	✓	✓	✓	✓	✓	✓
<i>Fixed Effects:</i>									
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 4-6, 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 the standardized mean occupation AI exposure, across the 6 digit occupations posted in each sector by commuting zone cell, based on the distribution of vacancies by detailed occupation in each zone and industry in 2010-2012. The regressions are weighted by baseline employment in each sector by commuting zone. In columns 4-9, the regressor is standardized 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. 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.

APPENDIX FIGURES AND TABLES

FIGURE A1: 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.

TABLE A1: Effects of 2007 AI Exposure on Establishment AI Vacancy Growth, 2010-2018

	(1)	(2)	(3)	(4)	(5)	(6)
Growth of Establishment AI Vacancies, 2010-2018						
<i>Panel A: Felten et al. Measure of AI Exposure</i>						
Establishment AI Exposure, 2007	23.32*** (2.33)	20.43*** (1.98)	12.20*** (1.78)	15.24*** (2.04)	12.07*** (1.74)	13.68*** (1.88)
Observations	102,783	102,783	101,553	102,783	101,524	94,866
<i>Panel B: Webb Measure of AI Exposure</i>						
Establishment AI Exposure, 2007	8.87*** (1.71)	6.97*** (1.54)	4.49*** (1.33)	5.04*** (1.39)	1.92 (1.30)	2.48** (1.20)
Observations	106,022	106,022	104,719	106,022	104,688	97,919
<i>Panel C: SML Measure of AI Exposure</i>						
Establishment AI Exposure, 2007	7.46*** (1.99)	5.44*** (1.78)	-1.66 (1.57)	-3.39* (1.82)	1.78 (1.44)	-0.68 (1.64)
Observations	106,022	106,022	104,719	106,022	104,688	97,919
<i>Covariates:</i>						
Share of Vacancies in Sales & Admin, 2010					✓	✓
<i>Fixed Effects:</i>						
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 sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). 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, established 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). 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 establishment 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 A2: Effects of AI Exposure on Establishment Broad AI Vacancy Growth, 2010-2018

	(1)	(2)	(3)	(4)	(5)	(6)
Growth of Establishment Broad AI Vacancies, 2010-2018						
	<i>Panel A: Felten et al. Measure of AI Exposure</i>					
Establishment AI Exposure, 2010 Observations	20.05*** (1.77)	17.17*** (1.49)	10.96*** (1.22)	18.72*** (1.99)	11.50*** (1.22)	18.80*** (1.95)
	1,075,474	1,075,474	954,519	1,075,474	954,519	1,075,474
	<i>Panel B: Webb Measure of AI Exposure</i>					
Establishment AI Exposure, 2010 Observations	8.47*** (1.30)	6.43*** (1.13)	3.72*** (0.83)	7.38*** (1.28)	0.99 (0.83)	2.53** (1.00)
	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789
	<i>Panel C: SML Measure of AI Exposure</i>					
Establishment AI Exposure, 2010 Observations	4.14*** (1.43)	2.18* (1.27)	-2.22** (0.89)	-2.90** (1.33)	1.92** (0.88)	4.90*** (1.34)
	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789
<i>Covariates:</i>						
Share of Vacancies in Sales, Admin. in 2010		✓			✓	✓
<i>Fixed Effects:</i>						
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. The sample is establishments posting vacancies in 2010-12 or 2016-18, outside NAICS sectors 51 (Information) and 54 (Business Services). 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). 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 establishment 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 A3: Effects of AI Exposure on Establishment AI Share Change, 2010-2018

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Change in Establishment Share of AI Vacancies, 2010-2018</i>						
<i>Panel A: Felten et al. Measure of AI Exposure</i>						
Establishment AI Exposure, 2010 Observations	0.29*** (0.03) 341,525	0.26*** (0.02) 341,525	0.20*** (0.02) 324,901	0.18*** (0.02) 341,525	0.22*** (0.02) 324,901	0.18*** (0.02) 341,525
<i>Panel B: Webb Measure of AI Exposure</i>						
Establishment AI Exposure, 2010 Observations	0.25*** (0.03) 355,529	0.22*** (0.03) 355,529	0.14*** (0.02) 337,758	0.11*** (0.02) 355,529	0.12*** (0.02) 337,758	0.05*** (0.02) 355,529
<i>Panel C: SML Measure of AI Exposure</i>						
Establishment AI Exposure, 2010 Observations	0.05*** (0.02) 355,529	0.03** (0.02) 355,529	-0.06*** (0.02) 337,758	-0.08*** (0.02) 355,529	0.04** (0.02) 337,758	0.05*** (0.02) 355,529
<i>Covariates:</i>						
Share of Vacancies in Sales, Admin. in 2010					✓	✓
<i>Fixed Effects:</i>						
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. Column 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 establishment 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 A4: Effects of AI Exposure on Firm AI Vacancy Growth, 2010-2018

	Firm AI Vacancy Growth, 2010-2018	Firm Mean, Establishment AI Growth, 2010-2018	Firm AI Vacancy Share Change, 2010-2018
	(1)	(2)	(3)
	(1)	(2)	(3)
	(4)	(5)	(6)
	<i>Panel A: Felten et al. Measure of AI Exposure</i>		
Firm AI Exposure, 2010	-3.77 (4.81)	-0.49 (4.13)	13.60*** (1.85)
Observations	401,399	320,628	401,399
	<i>Panel B: Webb Measure of AI Exposure</i>		
Firm AI Exposure, 2010	-6.83 (4.23)	-7.02 (4.86)	5.29*** (1.25)
Observations	431,195	339,512	431,195
	<i>Panel C: SML Measure of AI Exposure</i>		
Firm AI Exposure, 2010	6.01 (4.17)	3.44 (3.63)	4.96*** (1.59)
Observations	431,195	339,512	431,195
<i>Covariates:</i>			
Share of Vacancies in Sales & Admin, 2010	✓	✓	✓
<i>Fixed Effects:</i>			
Firm Size Decile	✓	✓	✓
3 digit Industry	✓	✓	✓

This table presents estimates of the effects of firm AI exposure on firm AI vacancy growth. In the first two columns, the outcome variable, constructed from Burning Glass data, is the change in the inverse hyperbolic sine of AI vacancies posted by the firm, between 2010-12 and 2016-18, multiplied by 100. In the middle two columns, the outcome variable is the firm mean of the change in the inverse hyperbolic sine of AI vacancies posted by each establishment in the firm, between 2010-12 and 2016-18, multiplied by 100. The firm mean is weighted by vacancies in each establishment in 2010-2012. In the last two columns, the outcome variable is the change in the ratio of AI vacancies to total vacancies between 2010-12 and 2016-2018, multiplied by 100. The regressor, firm AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the firm posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Firm AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies in the firm. 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. Columns 1, 3 and 5 contains only firm AI exposure. Columns 2, 4 and 6 include fixed effects for the decile of firm total vacancies in 2010-12, 3 digit NAICS industry fixed effects, and controls for the share of 2010-12 vacancies in each firm, belonging to the broad occupations of either sales or administration. In each regression, observations are weighted by total firm 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 A5: Effects of AI Exposure on Establishment Total Vacancy Growth, 2014-2018

	(1)	(2)	(3)	(4)
Growth of Establishment Total Vacancies, 2014-2018				
<i>Panel A: Felten et al. Measure of AI Exposure</i>				
Establishment AI Exposure, 2010	-11.80*** (3.80)	-11.24*** (2.93)	-10.51*** (2.81)	-5.13*** (1.02)
Observations	1,075,474	954,519	954,519	1,075,474
<i>Panel B: Webb Measure of AI Exposure</i>				
Establishment AI Exposure, 2010	-9.62*** (2.42)	-4.11** (1.83)	-2.21 (2.16)	-1.81** (0.71)
Observations	1,159,789	1,021,673	1,021,673	1,159,789
<i>Panel C: SML Measure of AI Exposure</i>				
Establishment AI Exposure, 2010	3.32 (2.09)	0.85 (2.30)	-0.57 (2.75)	-0.99 (0.92)
Observations	1,159,789	1,021,673	1,021,673	1,159,789
<i>Covariates:</i>				
Share of Vacancies in Sales, Admin. in 2010			✓	
<i>Fixed Effects:</i>				
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 2014 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 2013-15 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 and 3 include fixed effects for the decile of firm total vacancies in 2010-12 to which the establishment's firm belongs. Columns 2-4 include commuting zone fixed effects. Columns 2-3 include 3 digit NAICS industry fixed effects. Column 4 includes firm fixed effects. 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 A6: Effects of AI Exposure on Implied Establishment Non-AI Employment Growth, 2010-2018

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Growth of Establishment Non-AI Employment, 2010-2018								
	<i>Panel A: Felten et al. Measure of AI Exposure</i>							
Establishment AI Exposure, 2010	-7.24 (4.66)	-8.91** (4.53)	-5.62 (3.61)	-4.86*** (1.39)	-6.36* (3.59)	-4.80*** (1.39)	-4.57 (5.32)	-9.11*** (2.00)
Observations	1,075,474	1,075,474	954,519	1,075,474	954,519	1,075,474	324,901	341,525
	<i>Panel B: Webb Measure of AI Exposure</i>							
Establishment AI Exposure, 2010	-12.54*** (3.52)	-13.59*** (3.49)	-7.53*** (2.48)	-0.19 (0.81)	-12.17*** (3.25)	0.68 (0.93)	-7.47** (3.66)	-1.29 (1.37)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789	337,758	355,529
	<i>Panel C: SML Measure of AI Exposure</i>							
Establishment AI Exposure, 2010	2.13 (3.13)	1.45 (3.04)	1.56 (2.58)	-1.27 (0.94)	4.06 (2.83)	-2.80*** (1.05)	-0.72 (4.36)	-2.38 (1.47)
Observations	1,159,789	1,159,789	1,021,673	1,159,789	1,021,673	1,159,789	337,758	355,529
<i>Sample:</i>							✓	✓
Posts in 2018:								
<i>Covariates:</i>					✓			
Share of Vacancies in Sales & Admin, 2010								
<i>Fixed Effects:</i>								
Firm Size Decile	✓		✓		✓		✓	
Commuting Zone	✓		✓	✓	✓		✓	
3 digit Industry			✓		✓		✓	
Firm				✓		✓		✓

This table presents estimates of the effects of establishment AI exposure on implied establishment non-AI employment growth. The outcome variable, constructed from Burning Glass data, is the growth rate of establishment non-AI employment between 2010 and 2018, multiplied by 100. We measure establishment employment growth using the procedure for cumulating vacancies into a stock of employment, as described in the main text. 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 A7: Effects of AI Exposure on Firm Non-AI Vacancy Growth, 2010-2018

	(1)	(2)	(3)	(4)	(5)
Growth of Firm Non-AI Vacancies, 2010-2018					
<i>Panel A: Felten et al. Measure of AI Exposure</i>					
Firm AI Exposure, 2010	-14.06*** (5.40)	-12.09** (5.25)	-5.94 (8.09)	-5.85 (9.54)	-6.92 (7.79)
Observations	401,399	401,399	320,628	99,109	320,628
<i>Panel B: Webb Measure of AI Exposure</i>					
Firm AI Exposure, 2010	-23.16*** (4.98)	-22.74*** (4.88)	-4.95 (5.65)	0.56 (6.76)	-14.39** (6.82)
Observations	431,195	431,195	339,512	102,917	339,512
<i>Panel C: SML Measure of AI Exposure</i>					
Firm AI Exposure, 2010	7.38 (4.65)	7.59* (4.49)	3.68 (5.36)	-3.05 (6.00)	14.71** (6.72)
Observations	431,195	431,195	339,512	102,917	339,512
<i>Sample:</i>					
Posts in 2018:				✓	
<i>Covariates:</i>					
Share of Vacancies in Sales & Admin, 2010					✓
<i>Fixed Effects:</i>					
Firm Size Decile		✓	✓	✓	✓
3 digit Industry			✓	✓	✓

This table presents estimates of the effects of firm AI exposure on firm non-AI vacancy growth. The outcome variable, constructed from Burning Glass data, is the change in the inverse hyperbolic sine of non-AI vacancies posted by the firm, between 2010-12 and 2016-18, multiplied by 100. The regressor, firm AI exposure in 2010, is the mean of occupation AI exposure, over the 6 digit SOC occupations for which the firm posts vacancies in 2010-2012, weighted by the number of vacancies posted per occupation. Firm AI exposure is divided by the standard deviation, weighted by 2010-12 vacancies in the firm. 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 firm AI exposure. Columns 2-5 include fixed effects for the decile of firm total vacancies in 2010-12. Columns 3-5 include 3 digit NAICS industry fixed effects. Column 5 controls for the share of 2010-12 vacancies in each firm, belonging to the broad occupations of either sales or administration. Column 5 restricts the sample to firms with positive vacancies in 2018. In each regression, observations are weighted by total firm 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 A8: Effects of AI Exposure on Establishment Non-AI Vacancy Growth, Average Establishment Weights

	(1)	(2)	(3)	(4)
Growth of Establishment Non-AI Vacancies, 2014-2018				
<i>Panel A: Felten et al. Measure of AI Exposure</i>				
Establishment AI Exposure, 2010	-10.69*** (2.45)	-4.89** (2.22)	-3.42 (2.21)	-4.08*** (1.17)
Observations	1,075,474	954,519	954,519	1,075,474
<i>Panel B: Webb Measure of AI Exposure</i>				
Establishment AI Exposure, 2010	-2.89 (1.94)	1.71 (1.85)	3.98* (2.09)	-1.14 (0.94)
Observations	1,159,789	1,021,673	1,021,673	1,159,789
<i>Panel C: SML Measure of AI Exposure</i>				
Establishment AI Exposure, 2010	-2.60 (1.90)	1.63 (1.77)	3.41 (2.09)	0.75 (0.85)
Observations	1,159,789	1,021,673	1,021,673	1,159,789
<i>Covariates:</i>				
Share of Vacancies in Sales, Admin. in 2010			✓	
<i>Fixed Effects:</i>				
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. In each regression, observations are weighted by average vacancies in 2010-2012 and 2016-2018. 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. Standard errors are clustered by firm. One, two and three asterisks denote significance at the 10, 5 and 1 percent level.