



**Centre for
Economic
Performance**

Discussion Paper

ISSN 2042-2695

No. 2038
October 2024

Is distance from innovation a barrier to the adoption of artificial intelligence

Jennifer Hunt
Iain Cockburn
James Bessen



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■



**Economic
and Social
Research Council**

Abstract

Using our own data on artificial intelligence publications merged with Burning Glass vacancy data for 2007-2019, we investigate whether online vacancies for jobs requiring AI skills grow more slowly in US locations farther from pre-2007 AI innovation hotspots. We find that a commuting zone which is an additional 200km (125 miles) from the closest AI hotspot has 17% lower growth in AI jobs' share of vacancies. This is driven by distance from AI papers rather than AI patents. Distance reduces growth in AI research jobs as well as in jobs adapting AI to new industries, as evidenced by strong effects for computer and mathematical researchers, developers of software applications, and the finance and insurance industry. 20% of the effect is explained by the presence of state borders between some commuting zones and their closest hotspot. This could reflect state borders impeding migration and thus flows of tacit knowledge. Distance does not capture difficulty of in-person or remote collaboration nor knowledge and personnel flows within multi-establishment firms hiring in computer occupations.

Keywords: Artificial intelligence, Technology adoption and diffusion

JEL: O33, R12

This paper was produced as part of the Centre's Trade Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

We thank Bledi Taska and Lightcast (formerly known as Burning Glass Technologies) for access to data and Erich Denk formerly of the Technology and Policy Research Initiative at Boston University for extensive data work on the Burning Glass Technologies files. We thank Dany Bahar, Shai Bernstein, Lindsey McGowen, and Gregor Schubert for providing other data. We are grateful for comments from Flavio Calvino, Aureo de Paula, Ina Ganguli, Sabrina Genz, Britta Glennon, Tarek Alexander Hassan, Sabrina Howell, Ben Jones, Bill Kerr, John Landon-Lane, Dennis Novy, Gregor Schubert, Ruonan Xu and participants in the Australian OVERS seminar, the University of Sydney Microeconometrics and Public Policy Working Group seminar, the Technology and Policy Research Initiative seminar, the University of Southern Switzerland economics seminar, the 2021 NBER Artificial Intelligence Conference, the 2023 Artificial Intelligence and the Economy conference and the 2024 Royal Economic Society conference for helpful comments. All three authors are affiliated with the Technology and Policy Research Initiative at Boston University, and Hunt is also affiliated with the University of Sydney, IZA (Bonn), CEPR (London) and DIW (Berlin) and thanks the Centre for Economic Performance at the London School of Economics for hospitality while working on the paper.

Jennifer Hunt, Rutgers University. Iain Cockburn, Questrom School of Business, Boston University. James Bessen, Technology and Policy Research Initiative, Boston University School of Law.

Published by

Centre for Economic Performance

London School of Economic and Political Science

Houghton Street

London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© J. Hunt, I. Cockburn, and J. Bessen, submitted 2024.

The extent to which geographic distance is a barrier to technological knowledge transfer is of interest to governments of countries distant from centers of knowledge creation or technology production; to entrepreneurs deciding where to locate a new firm that will need to remain abreast of technological developments; and to national or local policy-makers seeking to influence the decisions of such entrepreneurs. These agents may value knowledge transfer as an input to further knowledge creation, or as a prerequisite for the adoption of new technology practices. In this paper, we provide insight into a little-studied aspect of knowledge transfer, by examining the geography of U.S. firms' adaptation and adoption of Artificial Intelligence (AI) in response to AI innovation.

The importance of distance for the diffusion of inventive and research activity has received considerable attention. Theoretically, distance could reduce inventors' and researchers' ability to source knowledge or their ability to collaborate, by reducing the probability of serendipitous meetings or raising the cost of planned meetings. The reduced probability of serendipitous meetings could reduce the probability of collaborations being initiated, while the higher cost of planned meetings could make sustaining a collaboration more expensive.¹ Because knowledge has been shown to be transferred when an inventor moves to a new firm, distance could also be a barrier to knowledge transfer because it is a barrier to migration.²

Such considerations may seem unimportant in the face of technological progress including the telephone, modern means of transportation, email, texting, the worldwide web and video conferencing. These are likely to have reduced the role of distance in both knowledge sourcing and especially sustaining collaboration. Indeed, several papers have shown that cross-location collaboration or citing of academic papers and patents has been increased by shorter and cheaper travel times.³ However, initiations of collaborations

¹ Esposito (2023); Catalini (2018). The World Intellectual Property Organization (2019) discusses the creation of contacts and networks in an international context.

² Empirical evidence for the importance of inventors' changing firm has been found for within-country firm to firm moves by Agrawal, Cockburn and McHale (2006); Almeida and Kogut (1999); Rahko (2017); and Sonmez (2017). For international moves see Kerr (2008); Briggs (2016); and Bahar, Choudhury and Rapoport (2020).

³ See Bahar et al. (2023), Berger and Prawitz (2023), and Hu et al. (2022) for air travel and Pauly and Stipanovic (2022) for rail travel.

appear sensitive even to small changes in distance – Catalini (2018) finds that existing collaborations persisted after the 1997–2014 shuffling of research laboratory locations on a Paris university campus, while collaboration between newly proximate laboratories increased greatly – and a more general empirical literature indicates that distance remains a barrier to the diffusion of inventive activity to potential inventors.⁴

A related literature examines how the adoption of technology, often across countries, is affected by the proximity of other adopters. One hypothesis is that it is advantageous for a potential adopter of a technology to be proximate to an earlier adopter because this makes adoption less risky: the later adopter could discuss adoption with the early adopter, observe the early adopter’s methods and outcomes, and poach the early adopter’s experienced workers. Another hypothesis is that firms could learn about distant technology through trade or their region’s receiving direct investment, and distance is a barrier to trade and direct investment. The empirical adoption literature confirms that distance is a barrier to the diffusion of adoption⁵, but finds the barrier to be lower for multiestablishment or multinational firms, which presumably have internal communication channels and coordination.⁶

Our paper examines whether distance constitutes a barrier between technology production (innovation) and technology adoption or adaptation, focusing on the technology of artificial intelligence (AI). We choose to examine AI in part because the rapid growth in AI research papers and patents began only recently, allowing an examination of its geographic diffusion from early in the process. It is also of particular interest because it is potentially important for future economic growth.⁷ Because AI is still immature, with few off-the-shelf applications yet available, we seek evidence for the effect of distance on

⁴ For analysis of patents, see Henderson, Jaffe and Trajtenberg (1991, 2005); Keller (2004); Peri (2005) Blit and Packalen (2018); Ganguli, Lin and Reynolds (2019); and Bernard, Moxnes and Saito (2020). Thompson and Fox-Kean (2005) have a contrary view. For analysis of country R&D as a proxy for innovation, see Keller (2002) and papers in Keller’s (2004) survey.

⁵ Little and Triest (1996); Comin, Dmitriev and Rossi-Hansberg (2012). See also papers on trade and innovation cited in Akcigit and Melitz (2021)

⁶ See Branstetter, Blenon and Jensen (2018). Giroud, Liu and Müller (2024) show that manufacturing plants in “tech clusters” are influenced by patenting by other firms in other tech clusters if their parent company has a plant in both clusters.

⁷ Aghion, Jones and Jones (2017); Goldfarb, Taska and Teodoridis (2019).

the adaptation of AI to a new environment, such as a new industry, in addition to the effect on adoption.⁸

We are the first in the small literature studying geographic links between innovation and technology adoption to study distance explicitly, and we are the first in the large literature on geographic diffusion of knowledge and technology to compare the roles of scientific papers and patents. The technology adoption literature provides evidence that an industry adopting an innovation tends either to be established in the location of the innovation, or, in the case of a mature industry, to move to the location of the innovation.⁹ The paper most closely related to ours is by Bloom et al. (2021), who are also the first to analyze the geographic diffusion of AI. They consider a group of 29 “disruptive” technologies including AI, showing they emerge through patents in concentrated “pioneer locations”, before spreading geographically as measured by convergence across locations in the share of job advertisements involving the technology group.¹⁰ Baruffaldi and Poege (2023) demonstrate that conferences are a channel through which scientific knowledge is passed from academics to firms.

To measure innovation, we create a dataset of AI publications, using Microsoft Academic Graph (MAG) to count journal articles, conference proceedings and patents identified in MAG as relevant to “deep learning”. We measure AI adaptation or adoption using job vacancy information from U.S. online job advertisements scraped by Burning Glass Technologies from 2007–2019. We divide the United States into 741 commuting zones and use 722 of them as a panel after having aggregated the variables to this level and excluded Alaska and Hawaii.

Our main approach to the question involves designating as AI innovation hotspots those commuting zones whose cumulative AI publications before our study period were

⁸ McEleran et al. (2024) describe which firms are adopting AI.

⁹ Duranton 2007; Kerr 2010; Zucker, Darby and Brewer 1998.

¹⁰ Other related papers are by Andersson, Quigley and Wilhelmsson (2009), and Dittmar and Meisenzahl (2022), who look at the impact of universities on local innovation. Acemoğlu, Autor, Hazell and Restrepo (2021) examine the growth of AI job advertisements in the Burning Glass Technologies data and Babina et al. (forthcoming) combine resume information with Burning Glass data, but these papers do not consider geography.

over a certain threshold. Our outcome of interest is subsequent growth in AI job advertisements as a share of all job advertisements, with the key covariate being the distance to the closest AI innovation hotspot. We assume that companies are able to fill the vacancies they post, and interpret a negative effect of distance as a barrier to hiring AI workers. Distance could be a barrier to the hiring of AI workers by companies already operating in distant commuting zones, or to the establishment in distant commuting zones of companies anticipating requiring AI workers. That U.S. firms mentioning AI on their website tend to be young (51% less than five years old) suggests the latter mechanism is likely to be important.¹¹ We complement this analysis by comparing coefficients on distance to the closest hotspot and on the radius of the circle around the commuting zone which encloses more than a certain threshold of cumulative AI publications before our study period.

We demonstrate that over 2007-2019, U.S. commuting zones more distant from established AI innovation hotspots had slower growth in AI-related hiring, whether for AI research or adapting AI for new purposes. The magnitude is substantial, with an additional 200km (125 miles) from the closest AI hotspot with at least 1000 papers and patents reducing a commuting zone’s seven-year growth in AI jobs (as a share of vacancies) by 17% of median growth.

The greatest effect of distance on AI job growth in percentage point terms is for computer and mathematical research occupations, the group with the fastest AI job growth, indicating that distance slows growth in AI research and innovation. The greatest effect as a percent of AI job growth is for developers of software applications, the occupation with the second-fastest AI job growth: an additional 200km (125 miles) from the closest hotspot reduces a commuting zone’s seven-year growth in such jobs by 27% of median growth. This suggests an important role for distance in slowing the adaptation of AI for new purposes. Consistent with this, distance slows AI job growth in finance and insurance, the industry with the fastest growing AI job growth, and in other industries likewise not associated with fundamental innovation in AI.

On the other hand, the evidence for whether distance to the closest hotspot slows

¹¹ Denis et al. (2023). See also Acemoğlu et al. (2022) for related statistics.

adoption of AI is mixed. We present weak evidence that distance may slow the extensive but not the intensive margin for the education and health sectors, business and finance occupations and image processing.

We find that the effect of distance from a hotspot of AI papers (designated based on journal articles and conference proceedings) is two to four times more negative than the effect of distance from a hotspot of AI patents. The positive correlation between the two distances means distance to a patent hotspot appears to have a much more negative effect when paper hotspots are ignored. This suggests that studies focusing on spillover effects or other geographic aspects of AI patents alone may be mistaking the effect of scientific papers for an effect of patents.

Our results are inconsistent with the hypothesis that distance is reducing AI job growth by making in-person or remote collaboration difficult, since travel time and time zone differences do not affect AI job growth, conditional on distance. These results also exclude the possibility that distance increases the difficulty of traveling to AI conferences, often held in AI hotspots, to obtain knowledge and encounter possible collaborators.¹² We also rule out the possibility that the distance effect reflects rapid internal knowledge transfer between geographically clustered establishments in the same multi-establishment firm: though commuting zones whose establishments are hiring in computer occupations both at home and in the hotspot have faster AI job growth, controlling for this does not change the effect of distance.

Rather, we find that 20% of the effect of distance is explained by commuting zones that are more distant from a hotspot being more likely to be in a different state from their hotspot. Consistent with our finding, Singh and Marx (2013) find political borders, including those within the United States, to be larger barriers to citations of patents than distance. Borders have been found to reduce exchanges of various types between jurisdictions. This is logical for international borders, where explicit barriers like tariffs

¹² For example, the Association for the Advancement of Artificial Intelligence’s annual Conference on Artificial Intelligence has rarely been held outside an AI hotspot. See <https://aaai.org/conference/aaai/> for conferences since 1980.

and quotas reduce trade and migration, and where institutional, language and cultural differences are large,¹³ but more surprising for domestic borders. That state borders are a barrier to trade within the United States is well known (e.g. Coughlin and Novy 2013), but possibly more relevant for our paper is the recent finding that Americans are three times more likely to move to a county in the same state than to an equally distant county in another state (Wilson 2023). Wilson convincingly rules out all explanations except a sense of identity Americans have with their state (particularly their birth state).

Thus, a state border may hinder a commuting zone’s ability to entice AI workers from hotspots, slowing the AI innovation and adaptation in its existing companies or deterring companies intending to use AI from setting up in the commuting zone. Although our finding for state borders is conditional on migration from the hotspot to the commuting zone, apparently disproving this theory, the apparently weak role for migration may reflect the coarse nature of the migration data.

1 Data

We have created our own database of AI papers and patents, and merge it with Burning Glass Technologies job advertisement information, as described in this section; more details are in the Data Appendix. We describe samples and variables in this section, with further details provided in the Data Appendix.

1.1 AI publications database and designation of innovation hotspots

Using the January 2020 release of Microsoft Academic Graph (Sinha et al. 2015), we have compiled a database of journal articles, conference proceedings and patents related to machine learning and neural networks, the areas that have led to a surge in commercial applications. These publications were selected using the coding with one or more fields of study from Shen et al.’s (2018) “hierarchical concept structure”, which is based on keyword and text analysis of publications and the graph structure of the database’s authorship and

¹³ See Clemens (2011) for migration and Havranek and Irsova (2016) for trade.

citation linkages. The authors of these publications work at firms and research institutes as well as universities. The location of each author was carefully geo-coded using information on their organizational affiliation at the time of publication. Of the 442,563 publication-author pairs which we identified as having a U.S. location, less than 0.5% could not be further geo-coded to the city-state level and were excluded from further consideration. Among the pairs in U.S. locations, 2.7% represent patents rather than journal articles or conference proceedings.

Using the city and state of each author, we obtain the county FIPS code, and then aggregate papers and patents into 741 commuting zones for each year.¹⁴ Each author is thus the source of potential spillovers, whether in the same or a different location from his or her co-authors. While we refer to the commuting zones' publications, these are really author-publication pairs.

We use these data to designate certain commuting zones as innovation hotspots, based on the cumulative number of AI papers+patents (a sum we refer to as publications) through 2006, the year before our study period. We assume that it is the total rather than per capita number of publications that matter for spillovers to other locations, and experiment with different absolute thresholds. To distinguish between the importance of papers (journal articles or conference proceedings) and patents, we also designate paper hotspots and patent hotspots, based on a commuting zone's number of papers or patents.

1.2 Lightcast (Burning Glass Technologies) job advertisements

Lightcast, formerly Burning Glass Technologies, is an employment analytics and labor market information firm which since 2007 has daily scraped the web's online U.S. job postings and produces files with duplicates eliminated and standardized information for each advertisement. Unfortunately, there are no data for 2008 and 2009, so our sample period is February-December 2007, all years and months from 2010-2018, and January-July

¹⁴ We match cities to counties using the "Pro" file provided at <https://simplemaps.com/data/us-cities>, accessed 18 February 2022. Of 128,692 publications, 34 have missing city; 770 have a city not in the simplemaps database, of which 750 are manually assigned a county, in some cases using wikipedia.

2019. Data collection in 2007 differs somewhat from that in later years, but we include 2007 because it is desirable to have data from the period when AI job advertisements were very uncommon.¹⁵ Of the variables available for each of the 200 million job advertisements, we use the location, the NAICS 2-digit industry code, the standard occupation classification code, classifications of keywords for required skills, and the employer name, which we harmonize across job advertisements.

We designate a job advertisement as being an AI job advertisement if the required skills include the Burning Glass keywords Artificial Intelligence, Machine Learning, Image Processing or any of the more specific keywords listed in Appendix Table 1; this is the set of terms used by Alekseeva et al. (2021). Ideally we would distinguish innovation from adoption based on the AI skills required, but the best we can do is examine three mutually exclusive categories: unspecified AI skills only (the 37% of job advertisements mentioning nothing beyond AI or ML); image processing, whether requested along with other AI skills or not; and the remaining AI skills or skill combinations. Image processing is clearly an adoption of AI, frequently requested for health occupations.¹⁶

We aggregate the total job advertisements, AI job advertisements and IT job advertisements to the commuting zone-year level using the county of the employer, and calculate the share of the commuting zone's total advertisements which are AI or IT advertisements in each year. Finally, we merge the data with the publication data. Our dependent variable is based on the share of advertisements that require AI skills, so that small commuting zones may experience as large an effect of distance as large commuting zones. We also calculate shares based on samples of job advertisements in 2-digit occupations and industries, and for selected 6-digit occupations.

¹⁵ We elected to concentrate on the United States only, because data for other countries (UK, Canada and Singapore) are available only from 2012 onwards.

¹⁶ See Burning Glass Technologies (2019) for a description of how required skills are codified. The examination of the raw Burning Glass text files by Bloom et al. (2021) allows them to divide the job postings according to whether the job will use, develop or produce the technology of interest.

1.3 Distance calculations

The files provided by Burning Glass provide the latitude and longitude of the employer, and we calculate the location of the commuting zone by averaging the latitude and longitude of all job advertisements over all years. Then we calculate the distances between commuting zones using Stata command `geodist` (based on Vicenty’s reference ellipsoid formula). For each commuting zone, we average the distances to all other commuting zones to compute the node centrality, and we calculate the distance to the nearest commuting zone. We also calculate travel times between commuting zones and their closest hotspot based on Google Maps and Google Flights.

To construct the independent variable we emphasize, we combine the distances with the hotspot information to compute the distance to the closest AI innovation hotspot for each commuting zone. Unless there is only one hotspot (a case we do not consider), even hotspots have a closest hotspot. We also compute the distance to the closest populous commuting zone for each commuting zone, with the definition of a populous commuting zone depending on the definition of hotspot being used: if a given AI threshold yields h commuting zones defined as hotspots, we define a populous commuting zone as one of the h most populous commuting zones. We also exploit an independent variable that does not rely on the concept of a hotspot: for each commuting zone, we calculate the radius of the circle around it which encompasses a given number of other commuting zones’ pre-2007 AI publications.

1.4 Other data and variables

We choose as our primary dependent variable differences (length k) in AI jobs’ share of job advertisements in commuting zone c whose closest innovation hotspot is commuting zone h : $\Delta^k AI_{cht}^s = \frac{AI\ job\ ads}{All\ job\ ads}_{cht} - \frac{AI\ job\ ads}{All\ job\ ads}_{cht-k}$, where the s superscript indicates a share. We use shares to avoid having the variation in AI reflect variation in commuting zone population. Our secondary dependent variable is the probability that $\Delta^k AI_{cht}^s$ is positive: this is insensitive to outliers and provides information on the extensive margin of growth

(particularly geographic).

We construct several variables for use in tests of the mechanism through which distance to the closest hotspot matters, beginning with the winter time difference between each commuting zone and its hotspot. To measure historical migration patterns between commuting zones, we use data from the IRS on movements of tax files.¹⁷ We obtain certain commuting zone characteristics from Opportunity Insights.¹⁸

We wish to assess whether the effect of distance to a hotspot reflects an exchange of information between establishments of multi-establishment firms which are close together. To do so, we create a variable measuring the number of 2007 computer and mathematical job advertisements in a commuting zone placed by firms which also post such vacancies in the closest AI hotspot in 2007, divided by the total number of 2007 job advertisements. We choose these occupations because they account for 63% of AI advertisements (see Appendix Table 2).

2 Methods

Selecting an identification approach that avoids the pitfalls outlined in Gibbons and Overman (2012), we estimate the effect of distance to the closest AI publication hotspot on AI job share growth, before distinguishing between paper and patent hotspots. We then assess the size of hotspots for which distance is important and distinguish between the effect of distance on AI innovation and AI adoption or adaptation. Finally, we investigate the mechanisms through which distance could be having an effect.

¹⁷ These data, aggregated to between-commuting zone flows, were generously provided by Gregor Schubert.

¹⁸ <https://opportunityinsights.org/data/?geographic.level=101&topic=0&paper.id=0#resource-listing>, accessed September 19, 2024.

2.1 Basic equation

The basic equation for estimating the effect of our key explanatory variable, distance to the nearest AI innovation hotspot (D_c^{Hot}) is:

$$\begin{aligned}
\Delta^k AI_{cht}^s = & \alpha + \sigma D_c^{Hot} \\
& + \beta_1 AI\ Papers_{>0_{c,t^*}} + \beta_2 AI\ Papers_{ct^*} + \beta_3 (AI\ Papers_{ct^*})^2 \\
& + \beta_4 AI\ Patents_{>0_{c,t^*}} + \beta_5 AI\ Patents_{ct^*} \\
& + \theta_1 AI\ Pubs_{ht^*}^{Hot} + \theta_2 (AI\ Pubs_{ht^*}^{Hot})^2 \\
& + \phi_1 D_c^{Big} + \phi_2 Pop_{ht^*}^{Big} + \phi_3 (Pop_{ht^*}^{Big})^2 \\
& + \gamma_1 \log(All\ job\ ads_{ct^*}) + \gamma_2 \log(Pop_{ct^*}) + \gamma_3 IT_{ct^*}^s \\
& + \delta_1 \bar{D}_c + \delta_2 D_c^{min} \\
& + \rho_1 \Delta^k AI\ Papers_{ct} + \rho_2 \Delta^k AI\ Patents_{ct-2} + \rho_3 \Delta^k \log(All\ job\ ads_{ct}) + \rho_4 \Delta^k IT_{ct}^s \\
& + \eta_t + \Delta^k \epsilon_{cht},
\end{aligned} \tag{1}$$

where t^* indicates a variable measured in 2007 or before (through 2006 in the case of AI publications) and that is therefore time-invariant. The coefficient of interest is σ . If σ is negative, distance constitutes a barrier to the posting of AI job vacancies. If it is zero, however, this could reflect either that distance is no barrier, or that distance is such a barrier that only innovation in the commuting zone affects a commuting zone's AI job vacancies. The conceptual randomization is the distribution of pre-2007 AI publications among commuting zones.

We mitigate the problem of outliers in the change in AI share by using longer differences and, where possible, median regression rather than OLS.¹⁹ Median regression is also helpful because of the large number of zeros in the dependent variable. However, median regression downweights both outliers we would like to downweight and those based on

¹⁹ At a given point in time, a large share of commuting zones have no AI job advertisements and many have only one or two, and the number of job advertisements is often also small in such commuting zones. We avoid using commuting zone fixed effects (including Poisson fixed effects), which might use short-run variation for identification, and which would also be problematic due to the absence of 2008 and 2009 data.

large changes in AI advertisements we might prefer not to downweight.²⁰ We also estimate equation 1 for the probability that $\Delta^k AI_{cht}^s$ is positive: this is insensitive to outliers and provides information on the extensive margin of growth (particularly geographic). We estimate these equations with linear probability: in a few cases, unreported results from probits show coefficients that differ slightly.

The first set of additional controls in equation 1 captures the commuting zone’s own AI innovation prior to 2007: a quadratic in the commuting zone’s own cumulative AI papers through 2006, $AI\ Papers_{ct^*}$ (quadratic rather than log due to the presence of zeros); a dummy for any such paper, $AI\ Papers > 0_{ct^*}$; a linear term in the commuting zone’s own cumulative AI patents through 2006, $AI\ Patents_{ct^*}$; and a dummy for any such patent, $AI\ Patents > 0_{c,t^*}$. The second pair of additional controls is a quadratic in the hotspot’s publications (when hotspot status is based on publications; otherwise a quadratic in hotspot papers or patents).²¹

The third set of controls ensures that the proximity of an AI hotspot is not proxying for the proximity of a populous commuting zone: the distance to the nearest populous commuting zone (D_c^{Big}), and a quadratic in the populous commuting zone’s population (D_c^{Big}). The fourth set of controls is for initial conditions: the 2007 number of job advertisements of all types, $\log(All\ job\ ads_{c,t^*})$; the population in the 2000 census, $\log(Pop_{c,t^*})$, despite the fact that the dependent variable is scaled, to control for variation in the size of online job boards relative to population; and IT’s share of job advertisements in 2007 (IT_{c,t^*}^s), to avoid the AI publication covariates picking up variation in non-AI IT. The fifth set of controls contains other distances that could be confounders of distance to the closest hotspot: node centrality \bar{D}_c (the average distance to all other commuting zones), for which network theory would predict a positive effect, and the distance to the closest commuting zone D_c^{min} .

Finally, we control for contemporaneous changes in some key variables: the number of

²⁰ A different solution would be to perform least squares weighting by commuting zone total job advertisements. But Solon, Haider and Woodridge (2015) recommend against weighting in such situations.

²¹ The hotspot-specific covariates vary only by h . Were we interested in their standard errors we would have to adjust for this, but since we are not, we do not.

the commuting zone’s own AI papers and patents, $\Delta^k AI Papers_{ct}$ and $\Delta^k AI Patents_{ct-2}$; the change in log job advertisements, $\Delta^k \log(All\ job\ ads_{ct})$; and the change in the IT job advertisements’ share in all advertisements, $\Delta^k IT_{ct}^s$.²² This could constitute overcontrolling: some or all of these could be the result of growth in AI job advertisements, rather than the cause, and their inclusion could bias $\hat{\sigma}$ upward toward zero. On the other hand, they may be causal and/or control for unobserved factors influencing growth in the adoption or adaptation of AI. Our preferred specification is therefore the one including all the covariates in the equation above.

We focus on seven-year differences (2007–2014, 2010–2017, 2011–2018, and 2012–2019), each of which spans the year (2016) when AI accelerates. In robustness checks, we report three-year differences, the only difference length to use all the years’ data (except 2007) as the first year of a difference. We cluster standard errors by commuting zone.²³ To justify our modeling of the effect of distance to the closest hotspot as linear, we replace the linear term in distance with dummies based on dividing distance to the closest hotspot into eight ranges with equal numbers of commuting zones in each.

2.2 Distinguishing distance to paper versus patent hotspots

In order to distinguish the effects of the distance to the closest AI paper hotspot and the distance to the closest AI patent hotspot, we estimate a variant of equation 1 in which covariates are duplicated to refer to each type of hotspot (for conciseness, we shorten the

²² We lag the change in patents to account for their reflecting patents applied for rather than granted.

²³ To cluster the standard errors we use the Stata `qreg2` command written by Parente, Santos Silva (2016).

equation by grouping the covariates which are unchanged from equation 1):

$$\begin{aligned}
\Delta^k AI_{cht}^s = & \alpha + \sigma_1 D_c^{HotPap} + \sigma_2 D_c^{HotPat} \\
& + \theta_1 AI\ Papers_{ht^*}^{HotPap} + \theta_2 \left(AI\ Papers_{ht^*}^{HotPap} \right)^2 + \theta_3 AI\ Patents_{ht^*}^{HotPat} \\
& + \phi_1 D_c^{BigPap} + \phi_2 Pop_{ht^*}^{BigPap} + \phi_3 \left(Pop_{ht^*}^{BigPap} \right)^2 + \phi_3 D_c^{BigPat} + \phi_4 Pop_{ht^*}^{BigPat} \\
& + AI\ Papers_{c,t^*} \beta_A + AI\ Patents_{c,t^*} \beta_B + Z_{ct^*} \gamma + D_c^{Other} \delta \\
& + \Delta^k X_{ct} \rho + \eta_t + \Delta^k \epsilon_{cht},
\end{aligned} \tag{2}$$

where D_c^{HotPap} is the distance to the closest hotspot with more than a certain number of papers and D_c^{HotPat} is the distance to the closest hotspot with more than a certain number of patents. The covariates concerning the closest populous commuting zone are also duplicated, since the number of populous commuting zones is based on the number of paper hotspots and patent hotspots.

2.3 Distinguishing larger and smaller hotspots

In addition, we investigate different thresholds for publication, paper or patent hotspots. To distinguish the influence of distance from a hotspot with at least 100 papers from the influence of distance from a hotspot with at least 1000 papers, for example, we include controls for the distance to the closest hotspot of at least 100 papers $D_c^{HotPap100}$ and the distance to the closest hotspot of at least 1000 papers $D_c^{HotPap1000}$, and duplicated other hotspot-specific covariates as well:

$$\begin{aligned}
\Delta^k AI_{cht}^s = & \alpha + \sigma_1 D_c^{HotPap100} + \sigma_2 D_c^{HotPap1000} \\
& + \theta_1 AI\ Papers_{ht^*}^{HotPap100} + \theta_2 \left(AI\ Papers_{ht^*}^{HotPap100} \right)^2 \\
& + \theta_3 AI\ Papers_{ht^*}^{HotPap1000} + \theta_4 \left(AI\ Papers_{ht^*}^{HotPap1000} \right)^2 \\
& + \phi_1 D_c^{BigPap100} + \phi_2 Pop_{ht^*}^{BigPap100} + \phi_3 \left(Pop_{ht^*}^{BigPap100} \right)^2 \\
& + \phi_3 D_c^{BigPap1000} + \phi_4 Pop_{ht^*}^{BigPap1000} + \phi_3 \left(Pop_{ht^*}^{BigPap1000} \right)^2 \\
& + AI\ Papers_{c,t^*} \beta_A + AI\ Patents_{c,t^*} \beta_B + Z_{ct^*} \gamma + D_c^{Other} \delta \\
& + \Delta^k X_{ct} \rho + \eta_t + \Delta^k \epsilon_{cht},
\end{aligned} \tag{3}$$

Note, however, that the distances to hotspots with a higher threshold are longer on average than distances to hotspots with a lower threshold, so the effect of distance is identified by variation at different distances for the two groups of hotspots. Although this problem cannot be overcome fully, we can estimate the effects in a different way for robustness. We can define hotspots based on a lower threshold only, then interact distance to the closest hotspot with the number of publications in the hotspot (in practice we choose to interact with a dummy for a large hotspot). Although in this specification the size of the closest hotspots will be smaller on average due to the lower threshold, there will be some closest hotspots which are large and fairly close, and if the coefficient on the interaction is negative, it implies that AI job growth is depressed more by being distant from larger hotspots than from smaller hotspots when the identification is from similar distances for the two groups of hotspots. When controlling for this interaction, we also control for the interaction of distance to the closest populous commuting zone and a dummy for a very populous closest populous commuting zone.

We gain insight further insight into the the impacts of small versus large hotspots by adding to the controls the radius of the circle around the commuting zone which encloses N or more pre-2007 AI publications (R_c^N) to specifications above. When doing so, we also add $\log(Pop_c^{RN})$, the log of the 2000 population within the circle with that radius, exclusive of the commuting zone's own population, and the number of AI publications within the circle, since this varies due to the lumpy geographic nature of AI publications at the commuting zone level.

2.4 Distinguishing the effect of distance on AI innovation versus AI adoption or adaptation

Further analysis is designed to distinguish whether the barrier is to the adaptation or adoption of AI, or merely to additional innovation in AI. For this purpose, we investigate the role of distance by industry and occupation, using as the outcomes the number of AI job advertisements in a particular occupation or industry, divided by job advertisements

in that occupation or industry. Finally, we examine the effect of distance to the closest hotspot on different AI types separately. In all these cases, we retain the covariates described above.

2.5 Investigating mechanisms

After this extensive analysis of the role of distance to AI hotspots, we add controls that might yield insight into why distance matters, focusing on hotspots defined as having at least 1000 publications. Distance may simply represent travel time, so we begin by controlling for this to see if its inclusion increases the coefficient on distance towards zero. To investigate the role of tacit knowledge brought by immigrants, we control for annual 1990–1999 per-capita immigration to the commuting zone from the closest hotspot. To assess whether distance reflects an exchange of information between establishments of multi-establishment firms which are close together, we control for the share of the commuting zone’s job advertisements which for computer and mathematical workers and placed by a firm which is also hiring such workers in the commuting zone’s hotspot.²⁴

3 Descriptive statistics

The national time-series of AI job advertisements is plotted in Figure 1. The increase over time from 9000 in 2007 to 190,000 in 2018 (and 135,000 in the first six months of 2019) far outstrips the 50% increase in the total number of job advertisements online. Figure 2 shows that the AI jobs share in all advertisements rises from 0.07 percent to 0.75 percent, rising linearly with a break in the slope in 2016, when growth increases (black squares, left scale). The IT jobs share is much higher (red circles, right scale) and evolves quite differently, rising from 2007–2012, then changing non-monotonically but ending lower in 2019 than 2012. The shares of the three types of AI job advertisements in all job advertisements are shown in Figure 3: unspecified AI and “other AI” have grown equally

²⁴ The calculation is necessarily based on the job advertisements for which the firm name is in the data.

quickly over the whole period and begin and end at the same shares, but “other AI” grew faster in the 2007–2012 period. Image processing, on the other hand, has not grown over the period. The Figure 4 maps indicating commuting zones’ AI job advertisement shares show how the fraction of commuting zones with no AI job advertisement (white) has shrunk with time, and how the non-zero shares has risen with time (as represented with darker shading) to a maximum of 4.0% in San Jose in 2019 (and one other small commuting zone).

The mean increase in the three-year AI job advertisement increase is 0.06 percentage point (first row of Table 1 panel A), while the median increase is lower at 0.03 percentage point. The minimum value of -2.46 percentage points and the maximum value of 4.70 percentage points confirm the existence of the outliers mentioned above. The mean seven-year increase is 0.14 percentage point and the median increase is 0.09 percentage point (second row). The share of observations for which growth is positive is 63% for three-year differences and 75% for seven-year differences (panel B), indicating broad growth. The lower panels of Table 1 shows the means of key covariates. The mean number of pre-2007 patents (5.7) is much lower than the mean number of pre-2007 papers (154); the median for both is zero, although 48% of commuting zones had had an AI publication by 2007, compared to only 17% having had an AI patent (panel D).

The national time-series for AI papers and patents from 1950 onwards (a few publications are pre-1950) are shown in Figure 5. Papers (times the number of authors), plotted in black squares using the left scale, increased from 6 in 1950, to 11,620 in 2007, to 49,484 in 2018 and to 65,411 in the first half of 2019. Patents rose from 1 in 1950 to 469 in 2007 to 1525 in 2017; the numbers fall in 2018 and 2019, reflecting the dating using patent applications rather than patent awards.²⁵

One definition of an innovation hotspot we use is having at least 1000 cumulative publications by 2006, and Figure 6 depicts the cumulative number of papers and patents for each of the 31 commuting zones satisfying this requirement: commuting zones are ordered by cumulative publications. The three top AI publication hotspots are Los An-

²⁵ This is why in regressions we lag patents by two years.

geles, Boston and Arlington, V.A. (the area around Washington, D.C.), each with more than 6000 publications, followed by the trio of New York, Pittsburgh and San Jose, with more than 4000 publications each. The highest publishing commuting zone outside the Northeast (including Pittsburgh) and California is Seattle in ninth place. Some of the hotspots are recognizable as technology and centers, others as university towns, and others as centers of military activity (Los Angeles is all three). New York, San Jose and Seattle stand out as having a large number of patents, while Pittsburgh stands out among the top ten as having a small number of patents. The five AI “pioneer locations” designated by Bloom et al. (2021) are all among our top nine AI hotspots, though notably do not include Los Angeles or Pittsburgh.²⁶ Figure 7 shows each commuting zone’s distance to its closest hotspot for the hotspots in Figure 6.

The map in Figure 8 shows the distribution of cumulative pre-2007 publications, while the four maps in Figure 9 show that there is very slow diffusion of publications through 2014, but faster diffusion afterwards.

4 Regression analysis

We begin our analysis by establishing that distance from an AI hotspot has a robust negative effect on the growth of AI job advertisements, driven by hotspots’ AI papers rather than by AI hotspots’ patents. We then examine the effects by industry, occupation and AI skill type, to establish whether distance from an AI hotspot affects (further) innovation or rather adaptation and adoption. Finally, we investigate possible mechanisms.

²⁶ Bloom et al. (2021) use 917 Core-Based Statistical Areas as their geographic units. The pioneer locations are (in order): Seattle; San Jose; San Francisco; New York–Newark and Boston. Based on AI patents alone, our top hotspots would be (in order): New York; Seattle; San Jose; Arlington (Washington); Newark; Houston and Boston. The Arlington V.A./Washington D.C. area is the main discrepancy between the two patent-based lists.

4.1 Effect of distance to the closest hotspot on all AI job advertisements

The first estimates of the effect of distance from a hotspot on growth in AI job advertisements are presented in Table 2 panel A, where a hotspot is defined as a commuting zone with at least 1000 publications (papers plus patents). All regressions in the table are median regressions. Column 1 includes only year dummies in addition to the distance to the closest hotspot. The statistically significant coefficient of -0.149 indicates that an increase in distance of 1000km (625 miles) reduces AI job advertisement growth by 0.149 percentage point (taking into account that the coefficients in the table are multiplied by 1000). The coefficient rises to -0.099 when we include the commuting zone's own pre-2007 AI papers and patents information and its hotspot(s)' pre-2007 papers and patents information in column 2.

This magnitude is fairly robust to adding more covariates in subsequent columns: covariates pertaining to other initial conditions in column 3; the distance to the closest populous commuting zone and a quadratic in its population in column 4; the average distance to other commuting zones, the distance to the closest commuting zone, and the population in the closest commuting zone in column 5; and finally the seven-year differences in key covariates in column 6. Although the inclusion of the latter covariates might appear to constitute over-controlling, it does not affect the coefficient much. The standard deviation of the distance to the closest hotspot of this type is 207km, and the column 6 coefficient of -0.079 can therefore be divided by five to obtain an effect on AI job share growth of -0.016 percentage point due to an increase in distance of 200km (125 miles), approximately a standard deviation. This represents a sizeable 17% of the median seven-year growth of 0.094 percentage point.

In panel B we consider the distance to the closest paper-based hotspot of at least 1000 papers (rather than publications), with results very similar to those for the publication hotspot in panel A. This is not surprising, since most publications are papers. In panel C, we instead consider the distance to the closest patent-based hotspot of at least 20 patents.

We choose this threshold for patents because the average distance to the closest hotspot of at least 20 patents is similar to the average distance to the closest hotspot of at least 1000 papers. The coefficients are statistically significantly negative in all columns, with the magnitude similar to the panel B paper coefficient in the early columns, but becoming considerably less negative in the later columns (-0.047 in column 6). However, the distances to the paper and patent hotspots used in panels B and C have a correlation of 0.68, so to distinguish the effects of paper and patent hotspots separately we must control for both, as in panel D.

In the panel D specification, we include the distance to the closest AI hotspot of each type and hotspot-specific covariates in corresponding pairs.²⁷ The column 1 coefficients on the two distances are less negative than when each distance was included individually in the upper panels, and are similar to each other in magnitude (-0.082 and -0.074). The column 2 coefficient on the distance to the closest paper hotspot is -0.075, compared to -0.096 in panel B, but remains statistically significant. On the other hand, the -0.022 coefficient on the distance to the closest patent hotspot is much less negative than in column 1 (-0.074) or in panel C (-0.085) and statistically insignificant.²⁸ The addition of further covariates leads to a statistically significant coefficient of -0.081 for the distance to the paper hotspot (column 5), essentially the same as in panels A and B, and a statistically insignificant coefficient of -0.026 for the distance to the patent hotspot.

We next test the robustness of the Table 2 panel D results to different estimation methods and other changes, maintaining the extensive covariates of column 5, as in all subsequent tables and figures. Column 1 of Table 3 reproduces the Table 2 column 5 results, while column 2 provides the equivalent results with the paper hotspot defined to have at least 500 papers (instead of 1000) and the patent hotspot defined to have at least 10 patents (instead of 20). The 10-patent threshold is chosen because its average associated distance is similar to the distance associated with the 500-paper threshold. The

²⁷ These include the distance to the closest populous commuting zone and its population, since the designation of a commuting zone as large depends on the number of hotspots.

²⁸ The coefficient on distance to the closest patent hotspot is -0.044 and statistically significant controlling only for 2006 publications and year dummies. This coefficient is not reported in the table.

two distance coefficients are slightly more negative than in column 1, but the coefficient of -0.042 on the distance to the closest patent hotspot remains statistically insignificant. The similarity of the columns 1 and 2 results is not surprising, since the two paper hotspot distances have a correlation of 0.74 and the two patent distances are similarly correlated. In column 3 we revert to the original hotspot definitions, but drop the seven-year difference of 2014-2007, as the 2007 sample differs from that of other years in some ways, such as the high share of IT job advertisements. The coefficient on distance to a paper hotspot is essentially unchanged compared to column 1, but the coefficient on distance to a patent hotspot is more negative, at -0.047, and statistically significant.

In column 4, we revert to using all years but estimate the regressions with OLS instead of median regression. The coefficient on distance to a paper hotspot is considerably more negative than earlier estimates, at -0.121, while the coefficient on distance to a patent hotspot is small, positive and statistically insignificant. The larger standard errors in OLS regression were expected due to the noise in the change in mean AI job advertisement share. The more negative coefficient on distance to a paper hotspot compared to the coefficient from median regression could represent the influence of larger commuting zones with AI job growth well above the median, or an asymmetric effect of outliers (and vice versa for the coefficient on the distance to a patent hotspot). Finally, column 5 shows that median regression using three-year differences also yields a statistically negative coefficient for the distance to the closest paper hotspot, less negative at -0.034 as would be expected given the shorter time span, and a coefficient of zero for the distance to the closest patent hotspot.

In further sets of robustness checks, we find that distance to smaller paper hotspots (e.g. commuting zones with at least 100 papers, see equation 3) may matter over shorter distances but its effect is hard to estimate precisely, and that the effect of distance is more negative the larger the AI hotspot (see method on p.15). Despite the latter result, many commuting zones are unaffected by being far from very large (2000 or more papers) hotspots because these large hotspots are so far away that it is smaller, closer hotspots that are relevant. We also find that it is more damaging to a commuting zone's AI job

growth to be distant from a concentration of AI papers (hotspot) than to require a large radius to enclose the same number of (more dispersed) papers. This could reflect a higher quality of papers in hotspots, or greater ease of communication and interaction when innovators are grouped in one place. The results of these robustness checks are presented in Appendix Tables 3 and 4 and analyzed in the Robustness Appendix.

In the many regressions we have run, the coefficient on distance from a patent hotspot is statistically significant in only one case. We conclude that the effect of being far from a patent hotspot is possibly zero, but more likely simply too small to identify precisely, and after the next figure we focus on hotspots defined based on the sum of papers and patents. Thus, studies analyzing spillovers from AI patents only may be largely picking up the effects of AI papers.

Thus far, we have modeled the effect of distance to the closest hotspot as linear. This is motivated by regressions where we instead use dummies based on dividing distance to the closest hotspot into eight ranges with equal numbers of commuting zones in each. These regressions are for pairs of paper and patent hotspots we have already used with linear distance. In Figure 10 we plot the resulting coefficients on the distance dummies (with the shortest distance set to zero), for the hotspot pairs of 500 papers and 10 patents, 1000 papers and 20 patents, and 1000 papers and 50 patents (we omit the very large standard errors from the plot), for seven-year differences, median regression and extensive covariates. For five of the six sets of coefficients, the effect of distance looks approximately linear (the exception is one of the sets for the distance to a patent hotspot). This does not necessarily mean the effect would be linear if the distances were longer, and indeed we have argued that the effect is approximately zero at the distances involved for a 2000-paper hotspot.

4.2 Distinguishing between the effect of distance on AI innovation and adoption

Having established that commuting zones farther from an AI paper hotspot have slower growth in AI job advertisements, we dig deeper to find what is driving this and to judge whether the slower growth represents slower growth in further innovation in AI or whether it represents slower AI adoption. We do so by estimating effects by industry, occupation and AI skill. We focus on distance to the closest 1000–publication AI hotspot, implicitly emphasizing spillovers from AI papers because we have found them to be more influential than AI patents, allowing the effect of distance to be captured in a single coefficient. We select the threshold of 1000 rather than 500 because in unreported regressions including both, it is the distance to the hotspot with 1000 or more publications that is statistically significantly negative, while the coefficient for the 500–publication hotspot is positive.

4.2.1 The effect of distance to an AI hotspot on different industries

For the analysis by industry and occupation, our commuting zone-level statistics are aggregated from industry and occupation-specific job advertisements. Table 4 shows that seven-year growth in AI job share in advertisements is fastest in the finance and insurance industry (0.27 percentage point on average, column 1), which clearly represents adoption or adaptation.²⁹ The third fastest AI growth is in the category containing real estate, professional and scientific services and administration (0.25 percentage point), which also represents adaptation or adoption as long as the occupation is not significantly influenced by advertisements for workers who will be employed by a staffing agency while working in (for example) the information sector. The information sector, where innovation would take place, has the fourth fastest AI job growth (0.21 percentage point). Growth in AI job advertisements with missing industry has the second-fastest AI job growth (0.25 percentage point). Almost all job advertisements with missing industry (16% in the advertisement level microdata) are also missing the firm name, which means that the job

²⁹ If a commuting zone has no advertisements in an industry at $t=0$ we set AI's share of advertisements in that industry to zero.

has been posted by an employment agency.³⁰

Because industries and industry-specific AI are not spread evenly across the United States, the median growth is zero for all industries except the category of real estate, professional services etc. and the missing industry category (column 2). This suggests the job advertisements with missing industry are in many different industries, as does the seven-year growth in 66% of missing industry observations (column 3) Despite its fast mean growth, AI's share in job advertisements in the information sector grew in only 23% of observations over seven years (column 3), suggesting geographically concentrated growth, whereas growth in the broader category of real estate, professional services etc. was positive in 53% of observations. Corresponding means for three-year differences are in Appendix Table 5.

We estimate the effect of distance to the closest 1000-publication hotspot by industry in Table 5. For most industries, median regression does not converge, so we begin with OLS, for seven-year differences, in column 1. The coefficients for the industries with the fastest growth are also those with the most negative coefficients on the distance to the closest hotspot, more negative than the all-industry coefficient of -0.12 in the first row. The largest is for missing industry (-0.39), followed by real estate, professional services etc. (-0.27), and then the statistically insignificant coefficients of -0.23 for information and -0.15 for finance and insurance. As a percent of the mean AI job growth, the effects of a 200km increase in distance to the closest hotspot are 31%, 22%, 22% and 11% respectively, higher in the first three cases than for all industries in the first row (18 %).

For those industries for which median regression is possible (column 2), the estimated coefficients are statistically significant and much smaller than for OLS (as is the case for all industries together). Notably, the coefficient for real estate, professional services etc. rises from -0.27 (OLS) to -0.07 (median regression). It is again difficult to judge whether the OLS coefficients are influenced by outliers due to small changes in small commuting

³⁰ Burning Glass Technologies, personal communication. Burning Glass aims to have the industry code reflect the industry of the ultimate employer, since otherwise the NAICS 2-digit code would always be 56 (the category including employment services) for job advertisements with missing employer name, which is not the case.

zones, or whether the outliers are economically important changes it is desirable to capture in an estimate. However, the fact that median regression yields a less negative coefficient than OLS in every regression in the paper where both estimation methods are possible suggests the latter. As a share of median AI job growth, the effect of a 200km increase is only slightly smaller (20%) for real estate, professional services etc. than the OLS effect.

In the last two columns, we estimate the probability of any growth in the AI job share, a measure of the geographic extent of growth is that ignores the intensity of growth in a given commuting zone. Column 3 continues to use seven-year differences, while column 4 uses three-year differences. This provides a different sort of information that like median regression is insensitive to outliers, possibly revealing an effect of distance in an industry that could not be detected using OLS or median regression. In the event, this is only true for education and health: despite a small positive coefficient in column 1, the column 4 coefficient is a statistically significant negative -0.1, similar to the coefficients for finance and insurance, and real estate, professional services etc. It is noteworthy, by contrast, that the point estimates for the information industry in these columns are small and statistically insignificant.

For all industries, the -0.17 coefficient in column 3 implies a 1000km increase in distance reduces the probability of positive growth by 17 percentage points; this means a more realistic 200km increase reduces the probability by $17/5=3.4$ percentage points, only 4.5% of the 76% of observations experiencing an increase in AI job share. Overall, the distance to the closest AI hotspot does not contribute to the extensive margin to nearly the extent that it contributes to the intensive margin. The analysis by industry supports the hypothesis that distance from AI innovation is a barrier to the adoption or adaptation of AI, though not that it is a barrier to (further) AI innovation, since the coefficients for the information industry are all statistically insignificant.

4.2.2 The effect of distance to an AI hotspot on different occupations

We gain more insight into the matter by studying four 2-digit occupations with fast AI growth, including computer and mathematical occupations, and by studying the 6-digit

computer and mathematical occupations for which we have more than one million job advertisements.³¹

Summary statistics by occupation for seven-year differences are given in Table 6 (three-year difference statistics are in Appendix Table 6). The growth in AI job share is far higher in computer and mathematical occupations than in other 2-digit occupations, with 1.2 percentage points growth on average compared to 0.35 percentage point growth in architecture and engineering, 0.25 for management, 0.15 for business and finance, and only 0.02 for other 2-digit occupations pooled (panel A column 1). The median is non-zero only for computer and mathematical occupations (0.88 percentage point) and for other occupations pooled (0.001 percentage point), as shown in column 2. Column 3 shows that the growth in the AI job share in computer and mathematical occupations is the most extensive, occurring in 72% of observations, while growth in business and finance occupations is the least extensive, occurring in only 35% of observations.

In the first row of panel B, we have aggregated small 6-digit occupations within the 2-digit computer and mathematical occupation whose workers seem very likely to be innovating. Based on the micro job advertisement data, 70% of the group are statisticians while 72% of the group's advertisements requiring AI are for research computer and mathematical workers. The growth in the AI job share for this group, which also includes mathematicians and operations researchers, is an enormous 7.3 percentage points per year (column 1), while the median is 2.3 percentage points (column 2). The next fastest growth is for developers of software applications (1.3 percentage points in column 1 and 0.4 percentage point in column 2) whom one would expect to be integral to adapting AI for new industries or purposes. Fast AI job growth is also occurring among computer programmers and computer and mathematical workers not elsewhere classified as well as web developers but not among network and computer systems administrators nor computer user support specialists. AI use among web developers reflects an application of AI, and therefore adaptation or adoption.³²

³¹ Large numbers of missing values for industries at a more detailed level than 2-digit preclude a similar analysis by detailed industry.

³² The research occupation group represents 5.1% of computing and mathematical workers, the devel-

Regressions by occupation are presented in Table 7, whose layout mirrors that of the Table 5 analysis by industry. Panel A shows statistically significant OLS coefficients on distance to the closest hotspot only for computer and mathematical workers (column 1). The coefficient for computer and mathematical workers is very negative at -1.02, implying that an increase in distance of 200km reduces the AI job share growth by $1.02/5=0.2$ percentage point. Of course, because the growth rate is high, this represents a more modest sounding 17% of the 1.2 percentage points average growth rate. The coefficient for architecture and engineering workers, more likely to be adopting than adapting or innovating in AI, is -0.30, statistically significant only at the 10% level, while the coefficients for the remaining occupations in the panel are much closer to zero and statistically insignificant.

Computer and mathematical occupations are the only 2-digit category for which median regression may be performed, and the coefficient is much smaller than for OLS (-0.55 column 2). An increase in distance of 200km reduces AI job share growth by 0.11 percentage point, equivalent to 13% of the median growth of 0.88 percentage point. Measured this way, the effect is smaller than the 17% effect for AI overall in Table 2 panel A column 5. In the regressions for the extensive margin in Table 7 panel A columns 3 and 4, we do detect a statistically significant negative effect of distance for business and finance occupations (-0.16 in column 3), whose OLS coefficient was -0.083 and statistically insignificant, and notably for the residual ‘other’ category, whose coefficients are as negative (-0.19 for seven-year differences and -0.11 for three-year differences) as those for architecture and engineering, despite a positive (statistically insignificant) OLS coefficient in column 1.

In panel B, we run regressions based on job-level samples of the larger of the 6-digit computer and mathematical occupations. The OLS coefficient on distance for the occupations likely to be engaged in innovation is very negative, at -4.0 (column 1). This implies a 200km increase in distance would reduce the AI job share growth by 0.8 percentage point or 11% of the mean growth. The median coefficient is half as large at -2.0 but statistically

opers of software applications 28%, the computer programmers 6.6%, the computer and mathematical workers not elsewhere classified 25%, the web developers 7.9%, the network and computer systems administrators 4.8% and the computer user support specialists 7.9%.

insignificant (column 2), while the statistically significant coefficients on the probability of any growth (-0.21 and -0.16 in columns 3 and 4) are similar to the coefficients for all occupations (-0.17 and -0.21 in Table 5 first row). The very negative coefficients in columns 1 and 2 thus reflect not that distance disproportionately reduces the likelihood of any computer and mathematical AI job growth compared to other occupations, but that it disproportionately slows the amount of growth in growing commuting zones, especially in commuting zones with well above median AI job growth.

The precisely estimated negative effects of distance to the closest AI hotspot on innovation in AI is in contrast to the imprecisely negative OLS effect for the information industry in Table 5 column 1 and the industry's extensive margin effects close to zero in Table 5 columns 3 and 4. Since the 2-digit information industry includes 6-digit industries that do not involve AI innovation, and since our innovative computer occupations likely includes AI innovators (professors) in the 2-digit education sector, overall we conclude that distance is a barrier to the growth of AI innovation.

The results for developers of software applications (Table 7 panel B), likely engaged in adapting AI, are qualitatively similar to those for research-related computer and mathematical occupations. The OLS coefficient of -1.8 implies a 200km increase in distance would reduce AI job share growth by 0.36 percentage point, equivalent to 28% of the 1.3 percentage points mean growth, the highest fraction in the analysis so far. The corresponding fraction for the median coefficient of -0.6 is 27%. The large impact of distance appears to be due in part to a strong reduction in the share of observations with positive growth (-0.37 and -0.32, columns 3 and 4). Network administrators make little use of AI, but would be expected to be adopting the AI they do use. The coefficients for this occupation are positive, and even statistically significant in two columns, so clearly this is a case where distance to an AI hotspot does not reduce growth of AI adoption.

For different reasons, the information gleaned from the results for other 6-digit occupations is more limited. The results for computer and mathematical workers not elsewhere classified are qualitatively similar to those for developers of software applications, though the coefficients slightly less negative in each column, but it is not possible to judge whether

workers in this group are likely to be predominantly involved in innovation, adaptation or adoption. The same problem of interpretation applies to the general computer programmers category, whose coefficients are less negative and less precisely estimated than those of the not elsewhere classified group. On the other hand, we would expect the AI use among web developers and the very limited AI use among computer user support specialists to represent adoption, but estimates are imprecise.

Overall, the results by occupation provide evidence for distance being a barrier to innovation in and adaptation of AI, but not for adoption of AI, since the results for web developers and computer user support specialists are imprecisely estimated and the results for network administrators reject the hypothesis.

4.2.3 The effect of distance to the closest AI hotspot on growth in various skills

We now exploit in more detail the skills required in the job advertisements. The most important results are for image processing, which is clearly an adoption of AI. Distance has only a small and statistically insignificant coefficient in seven-year difference OLS regressions for AI job growth and the probability of any AI growth, although the coefficient for the latter outcome is statistically significant for three-year difference regressions. Overall, the effect of distance on image processing appears small at best, perhaps because the technology is mature. Means and regression coefficients are shown in panel A of Appendix Tables 7–9.

We compare these results with three non-AI skills: computer-aided design; solar energy skills excluding installation, sales and management; and quantum computing skills. Quantum computing is intensive in AI skills and computer occupations, but has few advertisements (means and regression coefficients are shown in panel B of Appendix Tables 7–9). The coefficients on distance are always statistically insignificant, although the quantum computing coefficient is negative and significant at the 10% level.

4.3 Mechanisms

The previous section has established that distance to the closest AI hotspot is a brake on AI innovation as well as adaptation in finance and insurance, and likely in a broader set of industries including real estate and professional and scientific services. Distance operates at least in part by reducing innovation by computer and mathematical researchers and adaptation by developers of software applications to industries like finance and insurance. We turn now to investigating the mechanism through which distance could be a brake on AI job growth. In the first column of Table 8, we reproduce the overall effect of distance using median regression and seven-year differences (a coefficient of -0.079 from Table 2 column 5), before adding more covariates in the next following columns.

The most obvious hypothesis is that distance is proxying for travel time, and that greater travel time impedes in-person collaboration and networking and knowledge sourcing at conferences. However, when we add travel time as a covariate in column 2, its coefficient is positive (though statistically insignificant) and the coefficient on distance to the closest hotspot becomes somewhat more negative. Conversely, in the absence of the distance covariate, travel time has a negative and statistically significant coefficient (column 3). In the next several columns we omit the travel time covariate.³³

A second hypothesis is that a time difference between a commuting zone and its hotspot impedes remote collaboration. This seems somewhat unlikely, because no commuting zone has more than one hour time difference with its hotspot of 1000 publications, and indeed the coefficient on time difference in column 4 is statistically insignificant. Its point estimate implies that a one hour time difference reduces AI job growth by only 0.008 percentage point or 9% of median growth. A third hypothesis is that an important channel for information or personnel flows is between establishments of multi-establishment firms, and that these do not tend to locate far apart, causing distance to be a barrier. In column 4, we also control for the share (%) of the commuting zone’s advertisements which are computer worker advertisements posted by firms which are also seeking to hire com-

³³ We have also experimented with the price of the shortest flights, but found no evidence it matters.

puter workers in the hotspot. The coefficient on this variable is statistically significantly positive, indicating that a one percentage point increase in the share increases seven-year AI job share growth by 0.01 percentage point.

A fourth hypothesis is that distance matters because the greater the distance between two commuting zones, the less likely workers with tacit knowledge are likely to migrate or commute between the two. We control for a quadratic in the average annual per capita immigration to the commuting zone from its closest hotspot commuting zone over the 1990-1999 period, when it was unlikely to be influenced by AI, also in column 4. This yields a statistically insignificant positive effect on AI job share growth at the mean (0.033, see lower panel). The additional covariates of column 4 together do not explain much of the column 1 distance coefficient, though migration might matter were we able to measure flows of computer workers.

It is possible that distance is capturing an effect of state borders, and we find evidence to support this in columns 4 and 5. We add a dummy for the commuting zone and its closest hotspot being in the same state (allocating the commuting zone to the state in which most of it is located if it contains counties from more than one state) to the base specification of column 1. Doing so raises the coefficient on distance to the closest hotspot from -0.079 to -0.043, or 46%. The coefficient on the dummy is large and statistically significant, indicating that seven-year AI job share growth is higher by 0.034 percentage point or 36% of median growth when the commuting zone's closest hotspot is in the same state. However, this dummy is in effect an interaction term, and in column 6 we add a dummy for whether there is any hotspot of 1000 publications or more in the state. Though its coefficient is statistically insignificant and small (0.012), its inclusion reduces the coefficient on the dummy for same-state closest hotspot from 0.034 to 0.023.

We would like to gauge the contribution of each covariate to the change in the coefficient on distance to the closest hotspot, and as a first step, we control in column 7 for all the mechanism-related coefficients except travel time, which causes little change in the coefficient on distance to the closest hotspot (-0.045). No decomposition of the contributions is possible with median regression, however, so as a second step, we repeat

the column 1 and column 7 regressions using OLS (see Appendix Table 10) to establish that the results are qualitatively similar to those in Table 8. Based on the imprecise OLS results, we can then use a Gelbach (2016) decomposition to calculate that the dummy for same-state closest hotspot is responsible for 0.025 (equivalent to 18% of AI job growth) of the 0.055 rise (45%) in the coefficient on distance to the closest hotspot. The dummy for any hotspot in the state accounts for for 0.016 of the rise (see Appendix Table 10 column 6). If we assume that in median regression too, the same-state closest hotspot component is responsible for 60% of the joint effects of a commuting zone’s having a same-state closest hotspot and having any hotspot in the state, it would be responsible for approximately 60% of the 0.033 rise in the distance to closest hotspot coefficient from -0.079 (Table 8 column 1) to -0.046 (Table 8 column 6), or 0.020, equivalent to 20% of median seven-yearly AI job share growth.

Finally, in column 8, we add the control for travel time to the controls of column 7. Comparing these two columns confirms that controlling for travel time makes the coefficient on the distance to the closest hotspot more negative, this time considerably so. It is because this is difficult to interpret that we have based our calculation of contributions of covariates on specifications without this control.

These results appear to rule out the theory that distance to the closest hotspot is a brake on AI job growth because it reduces collaboration. The covariate with the most explanatory power is whether there are the commuting zone’s closest hotspot is in the same state. Although this result is conditional on immigration to the commuting zone from its closest hotspot, we do not rule out the possibility that the state border impedes flows of tacit knowledge by discouraging migration.

5 Conclusions

We have demonstrated that over 2007-2019, U.S. commuting zones more distant from established AI innovation hotspots had slower growth in AI-related hiring, whether for AI research or adapting AI for new purposes. The magnitude is substantial, with an

additional 200km (125 miles) from the closest AI hotspot with at least 1000 papers and patents reducing a commuting zone's seven-year growth in AI jobs (as a share of vacancies) by 19% of median growth. We find that the effect is driven by distance from AI scientific paper hotspots rather than by distance from AI patent hotspots. The positive correlation between the two distances means distance to a patent hotspot appears to have a much more negative effect when paper hotspots are ignored. This suggests that studies focusing on spillover effects or other geographic aspects of AI patents alone may be mistaking the effect of scientific papers for an effect of patents. Our analysis by occupation and industry suggests that distance is important for slowing both AI innovation and the adaptation of AI for new purposes. We find mixed evidence that distance is slowing AI adoption, with some weak evidence that it is slowing the extensive but not intensive margin for the health and education industries, business and finance occupations, and image processing.

We show that distance is not reducing AI job growth by making in-person or remote collaboration or networking difficult. Nor is it reflecting the locations of establishments in multi-establishments which use efficient internal knowledge transfer channels. Rather, we find that 20% of the effect of distance is explained by commuting zones that are distant from a hotspot being more likely to be in a different state from their hotspot. Since Americans are disinclined to move across state lines, this may reflect difficulty in hiring AI experts from the hotspot, which either slows AI jobs growth in existing firms or deters new firms wishing to use AI from setting up. More refined migration data than we have available are necessary to confirm this conjecture.

References

- Acemoglu, Daron, Gary W. Anderson, David N. Beede, Cathy Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, Pascual Restrepo and Nikolas Zolas. 2022. “Automation and the Workforce: A Firm–Level View From The 2019 Annual Business Survey”. NBER Working Paper 30,659.
- Daron Acemoglu, David Autor, Jonathon Hazell and Pascual Restrepo. 2022. “Artificial Intelligence and Jobs: Evidence from Online Vacancies”. *Journal of Labor Economics*, 40(S1): S293–S340.
- Aghion, Philippe, Benjamin F. Jones and Charles I. Jones. 2019. “Artificial Intelligence and Economic Growth”. In Ajay Agrawal, Joshua Gans and Avi Goldfarb eds. *The Economics of Artificial Intelligence*, Chicago: University of Chicago Press.
- Acemoglu, Daron, David Autor, Jonathon Hazell and Pascual Restrepo. 2021. “AI and Jobs: Evidence from Online Vacancies”. *Journal of Labor Economics*, 40(S1): S293–S340.
- Agrawal, Ajay, Iain Cockburn and John McHale. 2006. “Gone but not forgotten: knowledge flows, labor mobility, and enduring social relationships”. *Journal of Economic Geography*, 6: 571–591.
- Akcigit, Ufuk and Marc Melitz. 2022. “International Trade and Innovation”. In Gita Gopinath, Elhanan Helpman and Kenneth Rogoff eds. *Handbook of International Economics*, 5: 377–4004.
- Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila and Bledi Taska. 2021. “The Demand for AI Skills in the Labor Market”. *Labour Economics*, 71: 102002.
- Almeida, Paul and Bruce Kogut. 1999. “Localization of Knowledge and the Mobility of Engineers in Regional Networks”. *Management Science*, 45(7): 905–1024.
- Andersson, Roland, John M. Quigley and Mats Wilhelmsson. 2009. “Urbanization, productivity, and innovation: Evidence from investment in higher education”. *Journal of Urban Economics*, 66: 2–15.
- Babina, Tatian, Anastassia Fedyk, Alex He and James Hodson. Forthcoming. “Investments in Artificial Intelligence Technologies and Changes in Workforce Composition”. In Susanto Basu, Lucy Eldridge, John Haltiwanger and Erich Strassner eds. *Technology, Productivity and Economic Growth*, Chicago: University of Chicago Press.
- Bahar, Dany, Prithwiraj Choudhury, Do Yoon Kim and Wesley W. Koo. 2023. “Innovation on Wings: Nonstop Flights and Firm Innovation in the Global Context”. *Management Science*, 69(10): 6202–6223.

- Bahar, Dany, Prithwiraj Choudhury and Hillel Rapoport. 2020. “Migrant Inventors and the Technological Advantage of Nations”. *Research Policy*, 49(9): 103947.
- Baruffaldi, Stefano and Felix Poege. 2023. “Like Stars: How Firms Learn at Scientific Conferences”. Polytechnic of Milan working paper.
- Berger, Thor and Erik Prawitz. 2023 “Collaboration and Connectivity: Historical Evidence from Patent Records”. CEPR Discussion Paper 18,031.
- Bernard, Andrew B, Andreas Moxnes and Yukiko U. Saito. 2020. “The Geography of Knowledge Production: Connecting Islands and Ideas”. Dartmouth College working paper.
- Bernstein, Shai, Xavier Giroud and Richard R. Townsend. 2015. “The Impact of Venture Capital Monitoring”. *Journal of Finance*, 71(4): 1591–1622.
- Blit, Joel and Mikko Packalen. 2018. “A Machine Learning Analysis of the Geographic Localization of Knowledge Flows”. 2018. University of Waterloo working paper.
- Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner and Ahmed Tahoun. 2021. “The Diffusion of Disruptive Technologies”. NBER Working Paper 28999.
- Branstetter, Lee, Britta Glennon and J. Bradford Jensen. 2018. “Knowledge Transfer Abroad: The Role of U.S. Inventors within Global R&D Networks”. NBER Working Paper 24453.
- Burning Glass Technologies. 2019. “Mapping the Genome of Jobs: The Burning Glass skills taxonomy”. https://www.burning-glass.com/wp-content/uploads/2019/09/Burning_Glass_Skills_Taxonomy.pdf, accessed 2 June 2021.
- Carnevale, Anthony P., Tamara Jayasunder, and Dmitri Repnikov. 2014. “Understanding online job ads data”. Georgetown University, Center on Education and the Workforce, Technical Report.
- Catalini, Christian. 2018. “Microgeography and the Direction of Inventive Activity”. *Management Science*, 64(9): 4348–4364.
- Catalini, Christian, Christian Fons-Rosen and Patrick Gaulé. “How Do Travel Costs Shape Collaboration?” *Management Science*, 66(8): 3340–3360.
- Clemens, Michael. 2011. “Economics and Emigration: Trillion-Dollar Bills on the Sidewalk?”. *Journal of Economics Perspectives*, 25(3): 83–106.
- Comin, Diego, Mikhail Dmitriev and Esteban Rossi-Hansberg. 2012. “The Spatial Diffusion of Technology”. NBER Working Paper 18534.

- Coughlin, Cletus C. and Dennis Novy. 2013. “Is the International Border Effect Larger than the Domestic Border Effect? Evidence from US Trade”. *CESifo Economic Studies*, 59(20): 249–276.
- Deming, David and Lisa B. Kahn. 2018. “Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals”. *Journal of Labor Economics*, 36(S1): S337–S369.
- Dernis, H el ene, Flavio Calvino, Laurent Moussi egt, Daisuke Nawa, Lea Samek and Maria-grazia Squicciarini. 2023. “Identifying artificial intelligence actors using online data”. OECD Science, Technology and Industry Working Papers 2023/01.
- Dittmar, Jeremiah and Ralf Meisenzahl. 2022. “The Research University, Invention and Industry: Evidence from German History”. CEPR Discussion Paper 17383.
- Duranton, Gilles. 2007. “Urban Evolutions: The Fast, the Slow, and the Still”. *American Economic Review*, 97(1): 197–221.
- Esposito, Christopher J. 2023. “The geography of breakthrough invention in the United States over the 20th century”. *Research Policy*, 104810.
- Ganguli, Ina, Jeffrey Lin and Nicholas Reynolds. 2020. “The Paper Trail of Knowledge Spillovers: Evidence from Patent Interferences”. *American Economic Journal: Applied Economics*, 12(2): 278–302.
- Gelbach, Jonah. “When Do Covariates Matter? And Which Ones, and How Much?” 2016. *Journal of Labor Economics*, 34(2): 509–543.
- Gibbons, Stephen and Henry G. Overman. 2012. “Mostly Pointless Spatial Econometrics?” *Journal of Regional Science*, 52(2): 172–191.
- Giroud, Xavier, Ernest Lui and Holger M uller. 2024. “Innovation Spillovers Across U.S. Tech Clusters”. CEPR Discussion Paper 19206.
- Goldfarb, Avi, Bledi Taska and Florenta Teodoridis. 2023. “Could Machine Learning Be a General Purpose Technology? A comparison of emerging technologies using data from online job postings”. *Research Policy*, 52(1): 104653.
- Havranek, Tomas and Zuzana Irsova. 2016. “Do Borders Really Slash Trade? A Meta-Analysis”. *IMF Economic Review*, 65(2): 365–396.
- Henderson, Rebecca, Adam Jaffe and Manuel Trajtenberg. 1993. “Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations”. *Quarterly Journal of Economics*, 108(3): 577–598.
- Henderson, Rebecca, Adam Jaffe and Manuel Trajtenberg. 2005. “Patent Citations and the Geography of Knowledge Spillovers: A Reassessment: Comment”. *American Economic Review*, 95(1):461–464.

- Hershbein, Brad and Lisa B. Kahn. 2018. “Do Recessions Accelerate Routine–Biased Technological Change? Evidence from Vacancy Postings”. *American Economic Review*, 108(7): 1737–1772.
- Hu, Guangyuan, Rong Ni and Li Tang. 2022. “Do international nonstop flights foster influential research? Evidence from Sino-US scientific collaboration”. *Journal of Informetrics*, 16: 101348.
- Keller, Wolfgang. 2002. “Geographic Localization of International Technology Diffusion”. *American Economic Review*, 92(1): 120–142
- Keller, Wolfgang. 2004. “International Technology Diffusion”. *Journal of Economic Literature*, 42: 752–782.
- Kerr, William R. 2008. “Ethnic Scientific Communities and International Technology Diffusion”. *Review of Economics and Statistics*, 90(3): 518–537.
- Kerr, William R. 2010. “Breakthrough inventions and migrating clusters of innovation”. *Journal of Urban Economics*, 67:46–60.
- Little, Jane Sneddon and Robert K. Triest. 1996. “Technology Diffusion in U.S. Manufacturing: The Geographic Dimension”. *Proceedings of the Boston Federal Reserve Bank Conference on Technology and Growth*.
- McEleran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia S. Foster and Nikolas Zolas. 2023. “AI Adoption in America: Who, What, and Where”. *Journal of Economics and Management Strategy*, 33(2): 375–415.
- Parente, Paulo and João Santos Silva. 2016. “Quantile Regression with Clustered Data”. *Journal of Econometric Methods*, 5(1):1–15.
- Pauly, Stefan and Fernando Stipanovic. 2022. “The Creation and Diffusion of Knowledge: Evidence from the Jet Age”. <https://sciencespo.hal.science/hal-04067326>
- Peri, Giovanni. 2005. “Determinants of Knowledge Flows and Their Effect on Innovation”. *Review of Economics and Statistics*, 87(2): 308–322.
- Rahko, Jaana. 2017. “Knowledge spillovers through inventor mobility: the effect on firm–level patenting”. *Journal of Technology Transfer*, 42: 585–614.
- Singh, Jasit and Matt Marx. 2013. “Geographic Constraints on Knowledge Spillovers”. *Management Science*, 59(9): 2056–2078.
- Shen, Zhihong, Hao Ma and Kuansan Wang. 2018. “A Web–scale system for scientific knowledge exploration”. *2018 Meeting of the Association for Computational Linguistics*, pp 87-92 DOI: 10.18653/V1/P18-4015

- Sinha, Arnab, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June (Paul) Hsu, and Kuansan Wang. 2015. “An Overview of Microsoft Academic Service (MAS) and Applications”. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15 Companion)*, ACM, New York, NY, USA, 243-246. DOI=<http://dx.doi.org/10.1145/2740908.2742839>
- Solon, Gary, Steven J. Haider and Jeffrey M. Wooldridge. 2015. “What Are We Weighting For?” *Journal of Human Resources*, 50(2): 301–316.
- Sonmez, Zafer. 2017. “Inventor mobility and the geography of knowledge flows: evidence from the US”. *Science and Public Policy*, 44(5): 670–682.
- Thompson, Peter and Melanie Fox–Kean. 2005. “Patent Citations and the Geography of Knowledge Spillovers: A Reassessment”. *American Economic Review*, 95(1): 450–460.
- Wilson, Riley. 2023. “The Isolated States of America: Home State Bias and the Impact of State Borders on Mobility”. CESifo Working Paper 10724.
- World Intellectual Property Organization. 2019. *World Intellectual Property Report 2019. The Geography of Innovation: Local Hotspots, Global Networks*. https://www.wipo.int/edocs/pubdocs/en/wipo_pub_944_2019.pdf , accessed 29 September, 2023.
- Zucker, Lynne G., Michael R. Darby and Marilyn B. Brewer. 1998. “Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises”. *American Economic Review*, 88(1): 290–306.

A Data Appendix

This appendix provides some details not given in the main text.

A.1 AI publications database and designation of innovation hotspots

Our geo-coding was based on the text string containing the name of that author’s organizational affiliation, for example “Boston University, Boston, MA USA”. Of the 3.46 million publication-author pairs worldwide, 1.12 million could not be geo-coded: in the great majority of these cases, this was because we were unable to identify even the country of the author’s organizational affiliation because this text field was missing, corrupted, or was an ambiguous acronym. But our focus is on publications attributable to U.S. locations, and we are confident that our exhaustive search accurately captures the great majority of these in this set of AI publications. We obtain 1.14 million such publications worldwide, with an average of just over 3 authors per paper. 99% of the publications in this sample had 10 authors or fewer, though the distribution of authors per publication has a very long tail. We used all available information, including the apparent language or script of the text string (e.g. Cyrillic, Katakana), the top level domain of any email address or URL provided, the international calling code of any phone number, the linkage between the internal affiliation id and the GRID identifier developed by Microsoft, hand lookups using web searches, and (as a default) the geo-coding returned by the Google Maps API.

A.2 Lightcast (Burning Glass Technologies) job advertisements

Lightcast’s database has been widely used by labor economists (e.g. Deming and Kahn 2018). Hershbein and Kahn (2018) show that aggregate vacancy trends are consistent with those in administrative data, and while postings for college graduates and for industries with skilled workers are overrepresented (Carnevale, Jayasundera and Repnikov 2014), this is not a problem for our study.

In the raw data, 24% of job advertisements are missing industry, but we are able to reduce this share to 16% by replacing missing values with the modal industry code available for the same firm in the same year, when available. The remaining advertisements with missing industry therefore almost uniformly have a missing value for the employer name, which means the advertisement is posted by an employment agency (Burning Glass Technologies, personal communication). For a small proportion of postings, the county is missing, but as state is never missing, missing counties are assigned randomly within the state; for 2007 randomly to one of three counties with most advertisements.

We designate job advertisements as being IT job advertisements if the advertisement requires a skill other than Microsoft Office that is coded as “Information Technology” in the Skill Cluster Family (most aggregate) field (and the advertisement is not also an AI advertisement, though there is almost no overlap).

A.3 Distance calculations

To calculate travel times, we obtained driving times from Google Maps using the Stata `georoute` command, calculating driving times both to the closest hotspot and to the closest major (primary hub) airport. We manually queried Google Flights for the fastest one-way flights between major airports on Tuesday 29 August 2024. The travel time to the closest hotspot is then the driving time, if it is three hours or less, or the lesser of the driving time and the total time with a flight leg (allowing for two hours' airport time). We submitted the Google Maps requests as a batch, and the results included some missings and some driving distances much longer than would be implied by the Vicenty distance. We made individual queries in these cases, which yielded valid times (except for Hawaii and some Alaskan commuting zones), and shorter distances and times in the cases of the excessively long driving distances.

B Robustness Appendix

Following the analysis in Table 3, we pursue tests of robustness to further different hotspot definitions, using median regression, seven-year differences and extensive covariates throughout.

B.1 Distance to multiple hotspots

We consider thresholds for paper hotspots of 100, 1000 and 2000 papers, chosen because their average associated distances have correlations of 0.53 or lower, and thresholds for patent hotspots of 1, 20 and 200 patents, chosen because their associated distances have correlations of 0.49 or lower and because the average associated distances are pairwise approximately equal. In column 1 of Appendix Table 3, we use the pair of hotspots with thresholds 100 papers and a single patent to illustrate that for papers, this threshold yields the largest coefficient on distance found so far: a statistically significant -0.146. However, column 2 shows that it is not the case that the threshold of 100 papers captures the effect of distance better than the threshold of 1000 papers. When we add covariates associated with the pair of hotspots with thresholds of 1000 papers and 20 patents, we find that the point estimates of the coefficients on distance to the closest 100-paper hotspot and the closest 1000-paper hotspot are similar (a now statistically insignificant -0.068 for the former and a statistically significant -0.081 for the latter). The results suggest that distance to smaller paper hotspots may matter over shorter distances but its effect is hard to estimate precisely, while distance to larger paper hotspots matters over longer distances. The coefficients on distances to closest patent hotspots are statistically insignificant in columns 1 and 2.

In column 3, we add covariates associated with another pair of hotspots: commuting zones with at least 2000 papers and commuting zones with at least 200 patents. Because median regression does not converge when we include the three pairs of hotspots, we drop the covariates associated with hotspots of 100 papers and one patent. Compared to

column 2 or Table 2 panel D, this regression yields a slightly less negative coefficient on the distance to a paper hotspot of 1000 papers (-0.068), and a more negative (-0.040) but still statistically insignificant coefficient on the distance to a patent hotspot of 20 patents. The coefficient on the distance to a patent hotspot of 200 patents is approximately zero. More surprising is the small positive, statistically insignificant coefficient on the distance to a paper hotspot of 2000 papers. We investigate this result further in the next appendix table.³⁴

Before doing so, in column 4 we show results from one more pair of hotspots in Appendix Table 3 column 4: hotspots with thresholds of 1000 papers and 50 patents. In this specification, the two coefficients on distance are more similar than in other regressions: -0.054 for distance to the paper hotspot and -0.036 for distance to the patent hotspot. But as in almost all our regressions, the coefficient for the patent hotspot is statistically insignificant.

B.2 Distance from very large hotspots

In Appendix Table 4, we show that the positive coefficient on the distance to the closest 2000-paper hotspot of Appendix Table 3 column 3 is not due to collinearity, since when only this hotspot definition is considered, the coefficient on distance to the closest hotspot is a statistically significant 0.029 (column 1). When thinking about this paper hotspot threshold, it is useful to know that the average distance to the closest paper hotspot jumps 200km when the threshold is raised so as to exclude Austin, TX, which has 1827 AI papers and is the largest hotspot far from the east or west coast. We add controls for the radius around each commuting zone which encloses 2500 AI papers (as well as a control for the enclosed population), a radius chosen because 2500 is approximately the average number of papers in a hotspot with at least 2000 papers. For commuting zones in Texas and neighboring states, this radius is much shorter than the distance to the nearest hotspot of 2000 papers. Column 2 shows that the radius has a statistically significant coefficient of -0.092, and the coefficient on the distance to the paper hotspot is unchanged. Thus, it is not that once paper numbers get large they have no spillover effect; it is likely instead that paper hotspots of 2000 papers or more are on average too far away to exert any influence.³⁵

If we also control for distance to the closest paper hotspot of 1000 or more papers (and associated covariates, column 3), its estimated coefficient is a statistically significant -0.083, very similar to estimates in earlier tables, while the coefficient on the radius flips sign and becomes statistically insignificant (the two distance covariates have a correlation of 0.77). This suggest that it is more damaging to a commuting zone's AI job growth to

³⁴ An unreported OLS regression including the three pairs of hotspots yields coefficients of the same sign as their counterparts in columns 1 and 2 but all six are statistically insignificant. The coefficient on distance to the closest 1000-paper hotspot of -0.109 is statistically significant at the 10% level, and the point estimate on the distance to the closest 20-patent hotspot is more negative than in earlier estimates, at -0.088.


³⁵ The equivalent jump for patents is when the hotspot threshold is raised above Houston's 165 cumulative patents, leaving only Arlington VA, New York, Newark, San Jose and Seattle as hotspots.

be distant from a concentration of AI papers than to require a large radius to enclose the same number of (more dispersed) papers. This could reflect a higher quality of papers in hotspots, or greater ease of communication and interaction when innovators are grouped in one place. The point estimate on distance to the closest 2000-paper hotspot remains unchanged, however. The column 3 coefficients are unchanged by controlling for the distance to a patent hotspot of 200 patents and the radius enclosing 200 patents (and associated covariates) in column 4; the coefficients on the distance to the patent hotspot and on the patent radius are statistically insignificant.

We follow up on our hypothesis concerning hotspots of at least 2000 papers by identifying their effect differently. We control for the distance to the closest hotspot of at least 1000 papers, but also control for an interaction of this with a dummy for the hotspot in question having at least 2000 papers.³⁶ Unlike when the hotspot threshold is set at 2000 papers, commuting zones whose closest hotspot has at least 2000 papers do not have a closer hotspot of 1000–1999 papers, and the average distance to a closest hotspot with 2000 or more papers is consequently shorter. We exploit variation in distance to relatively close 2000–paper hotspots and variation in distance to 1000–1999 paper hotspots instead of relying on variation in distance to relatively distant 2000-paper hotspots. Column 5 shows that the coefficient on the distance to the closest hotspot of 1000 papers is a statistically significant -0.074, while the coefficient on the interaction term has a comparatively negative point estimate of -0.038, statistically significant at the 10% level. Here, the effect of distance to a hotspot of 2000 papers is thus unambiguously negative, and possibly much more negative than the distance to the closest hotspot of 1000–1999 papers. The coefficient of close to zero on the distance to the closest hotspot of at least 2000 papers in the previous columns is therefore very likely to reflect that most commuting zones are too far from such hotspots for variation in distance to matter.

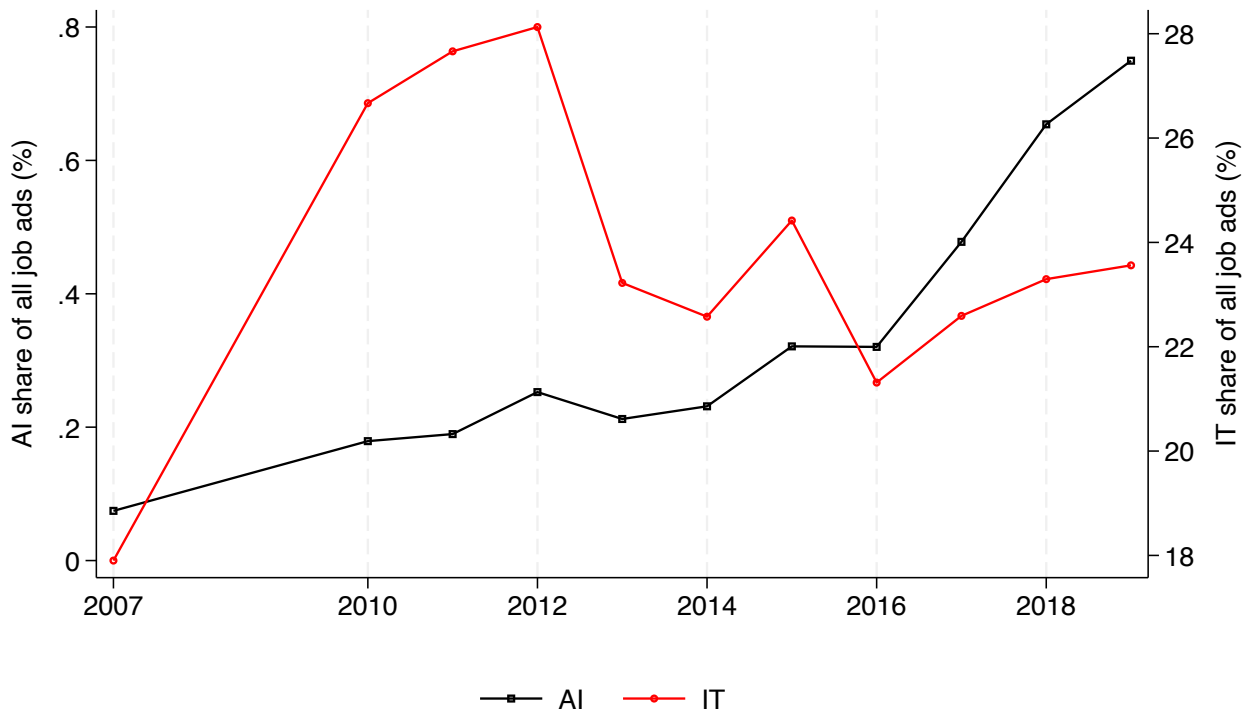
³⁶ We also add a control for the interaction between distance to the closest populous commuting zone and a dummy for its being a particularly populous commuting zone; the number of particularly populous commuting zones is the same as the number of hotspots with at least 2000 papers. We retain the main effect as a quadratic in hotspot papers.

Figure 1: Number of online AI job advertisements 2007–2019



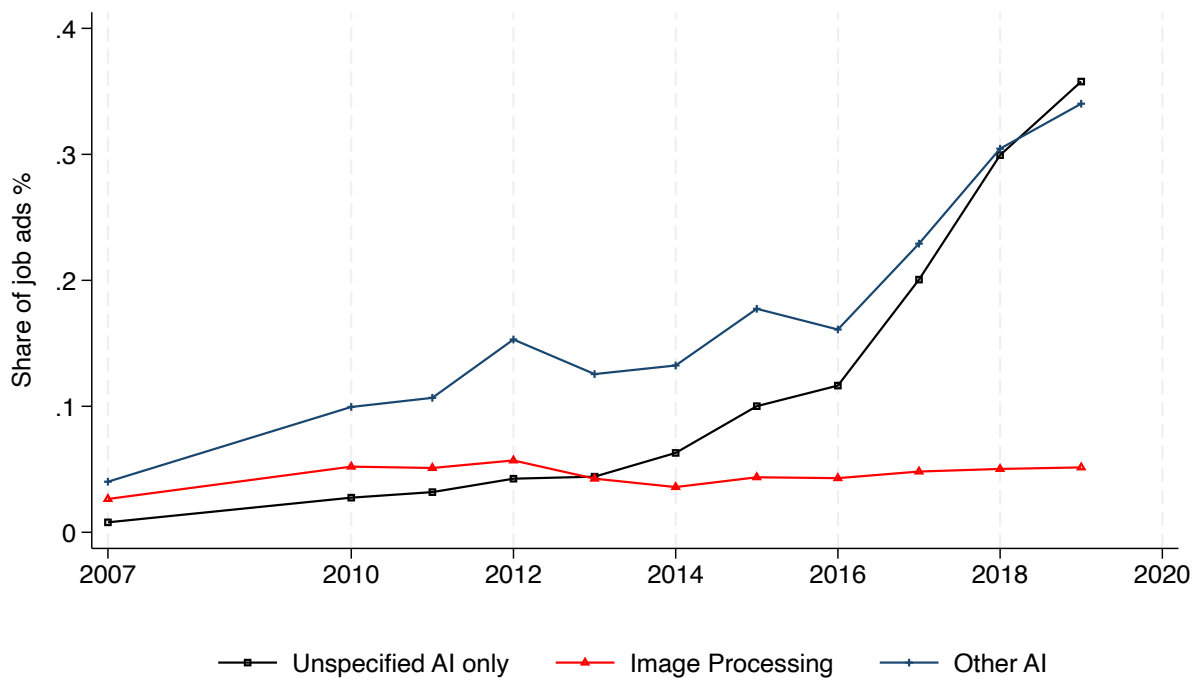
Notes: Data for 2019 are for January–July. Data for 2008 and 2009 are not available.
Source: Burning Glass Technologies.

Figure 2: AI share of job ads (%)



Source: Burning Glass Technologies.


Figure 3: Growth in share of job advertisements accounted for by different types of AI (%)




Note: Unspecified AI job advertisements require “Artificial Intelligence” and/or “Machine Learning” skills with no further detail given. Image processing AI job advertisements require image processing as one of the required skills. The “other” category is defined so the three categories are mutually exclusive.

Figure 4: AI job advertisements as percent of jobs advertisements in given year


(a) 2007



(b) 2010



(c) 2014



(d) 2018


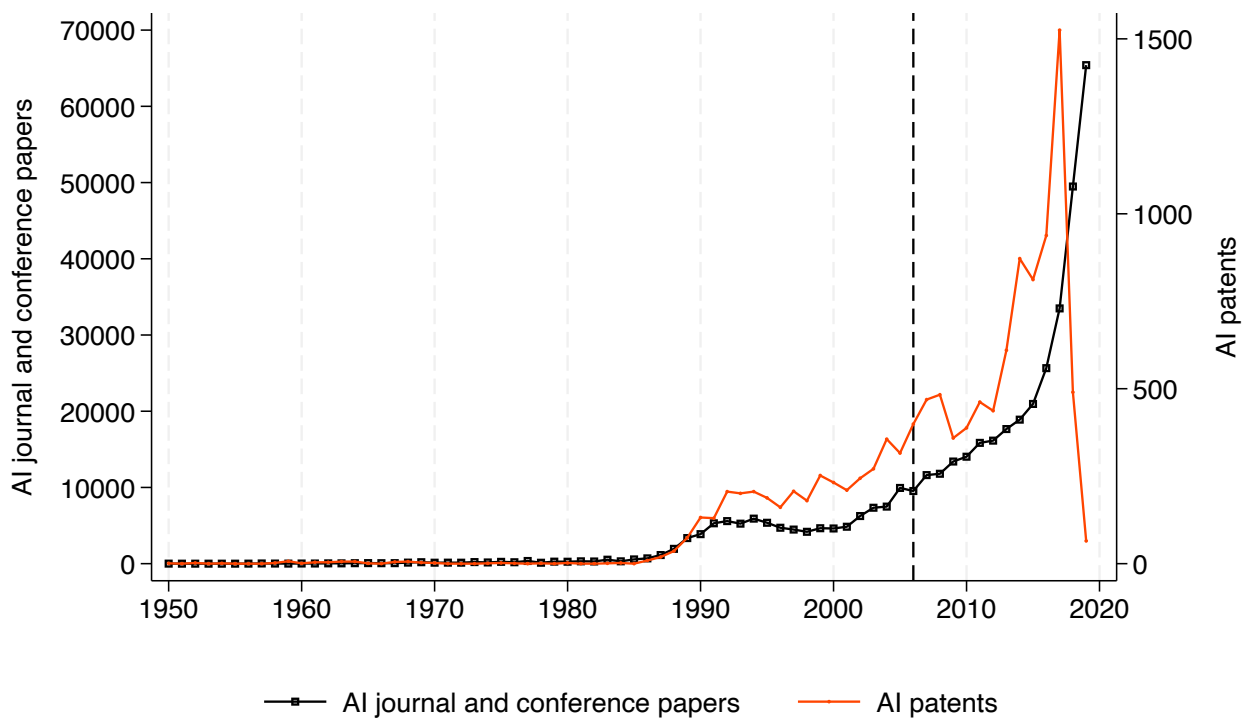
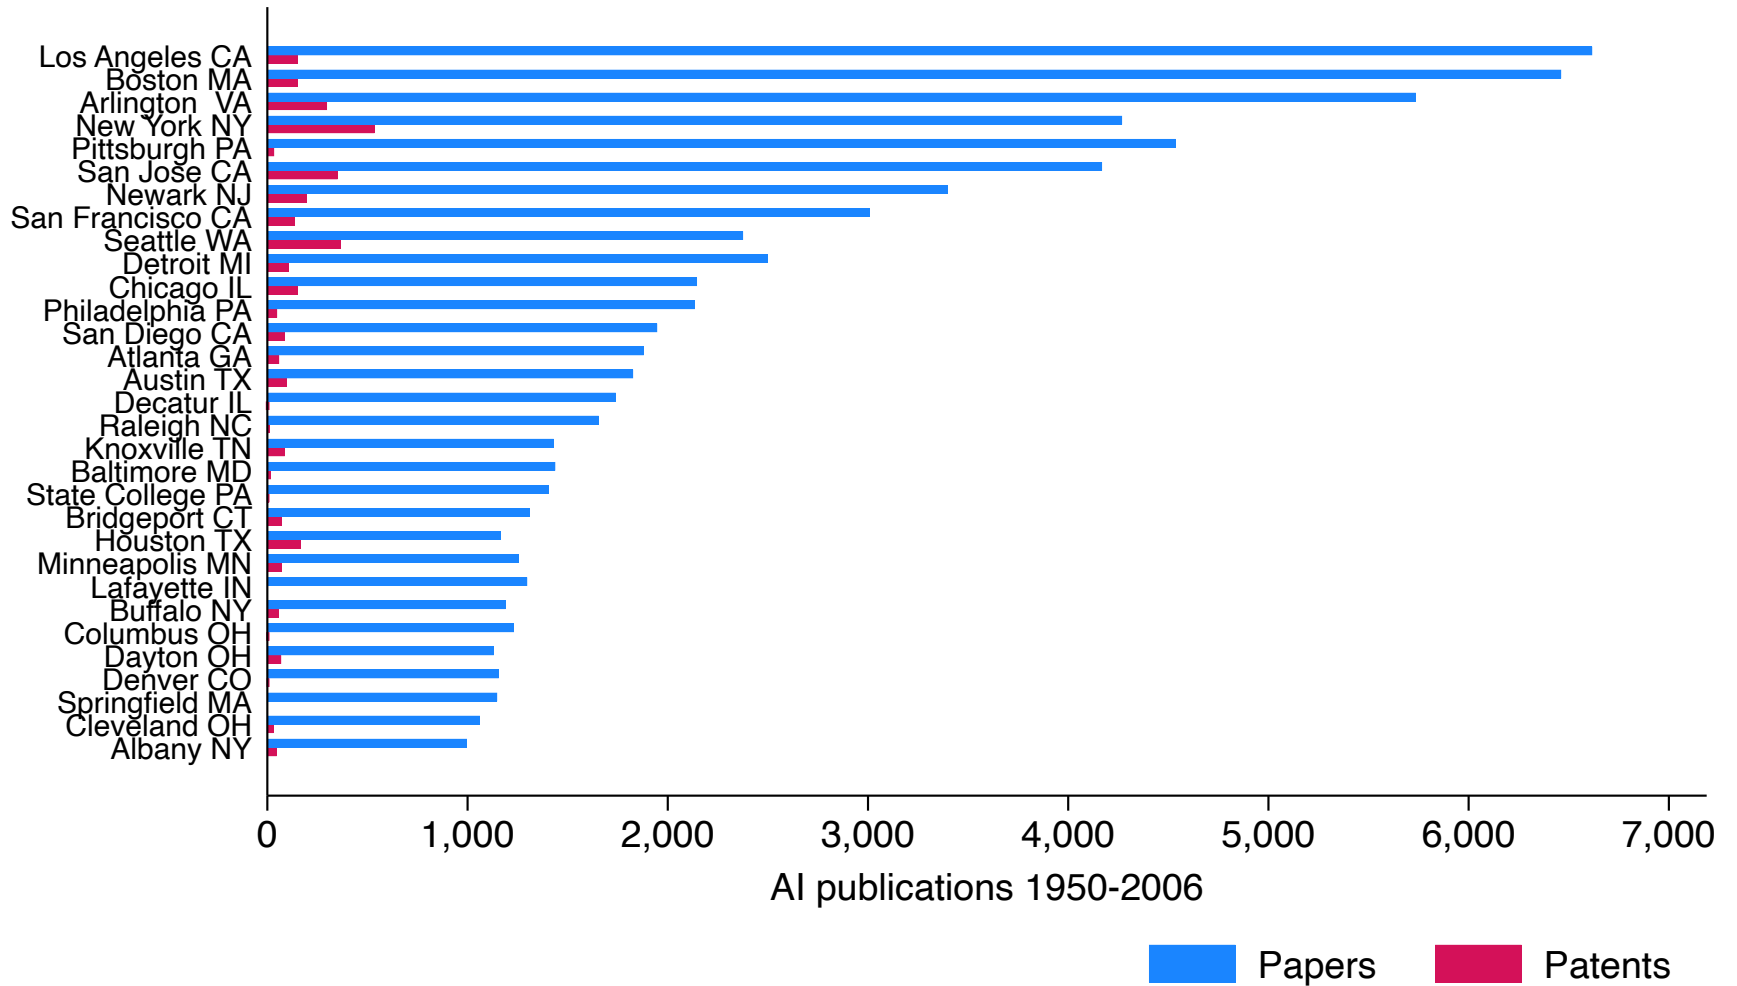


Figure 5: AI publications 1950-2019



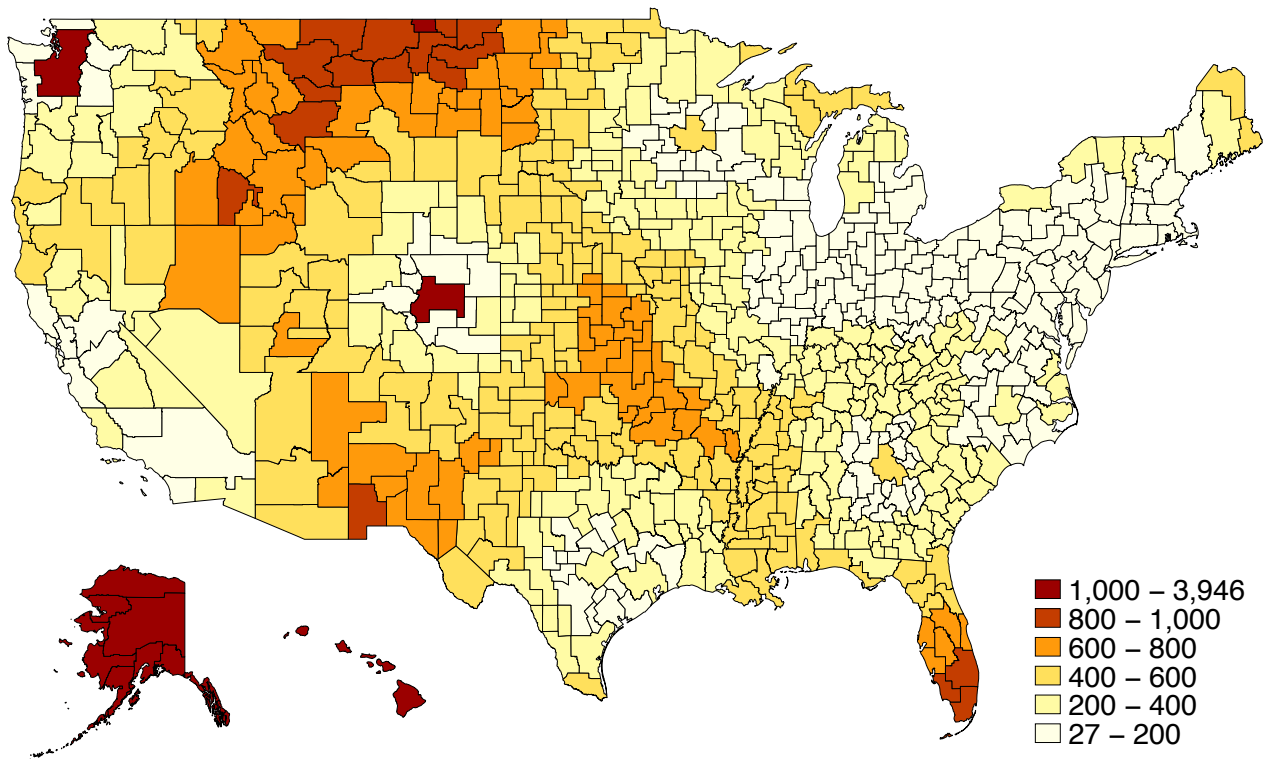
Source: Authors' dataset.

Figure 6: Innovation hotspots' AI publications through 2006



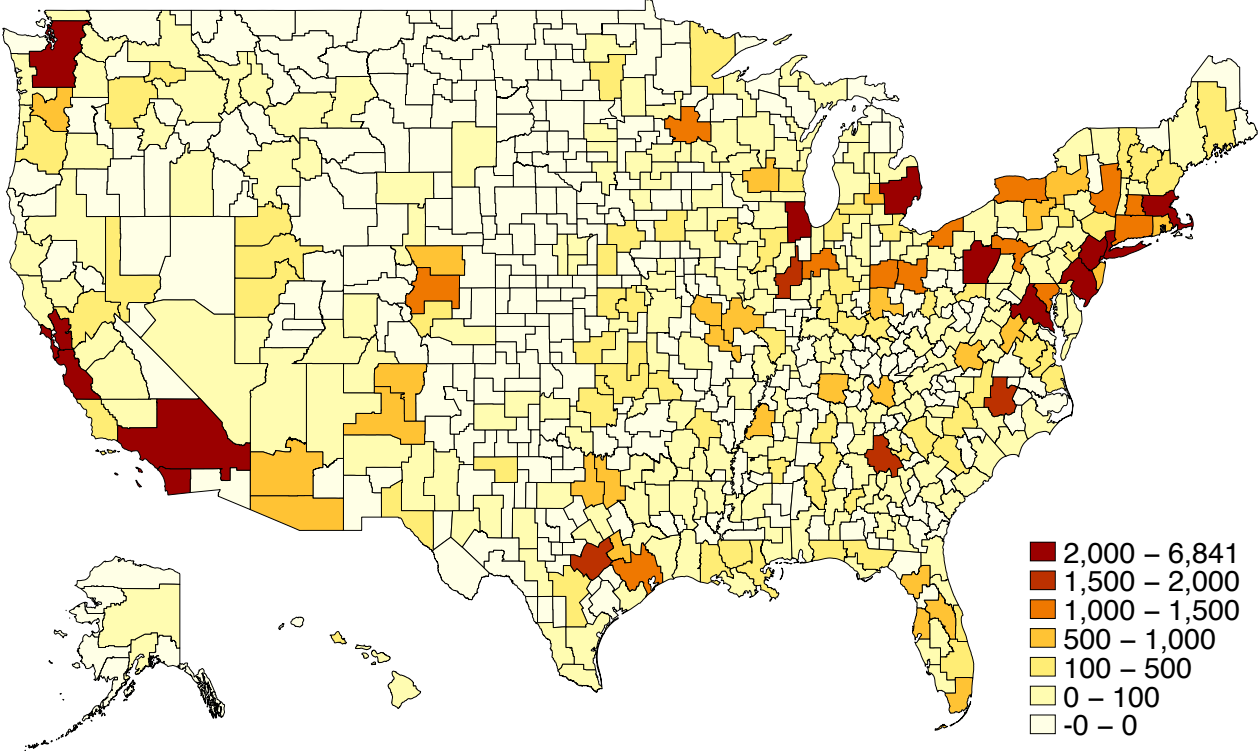
Note: The definition of a hotspot is a commuting zone with at least 1000 AI publications (papers+patents) through 2006.

Figure 7: Commuting zone distance to closest hotspot of 1000 or more publications (km)



Note: The definition of a hotspot is a commuting zone with at least 1000 AI publications (papers+patents) through 2006.


Figure 8: Commuting zones' AI publications through 2006




Note: Publications are the sum of patents and papers.

Figure 9: Commuting zones' AI publications in given year


(a) 2007: 240 commuting zones with any publication




(b) 2010: 245 commuting zones with any publication



(c) 2014: 252 commuting zones with any publication

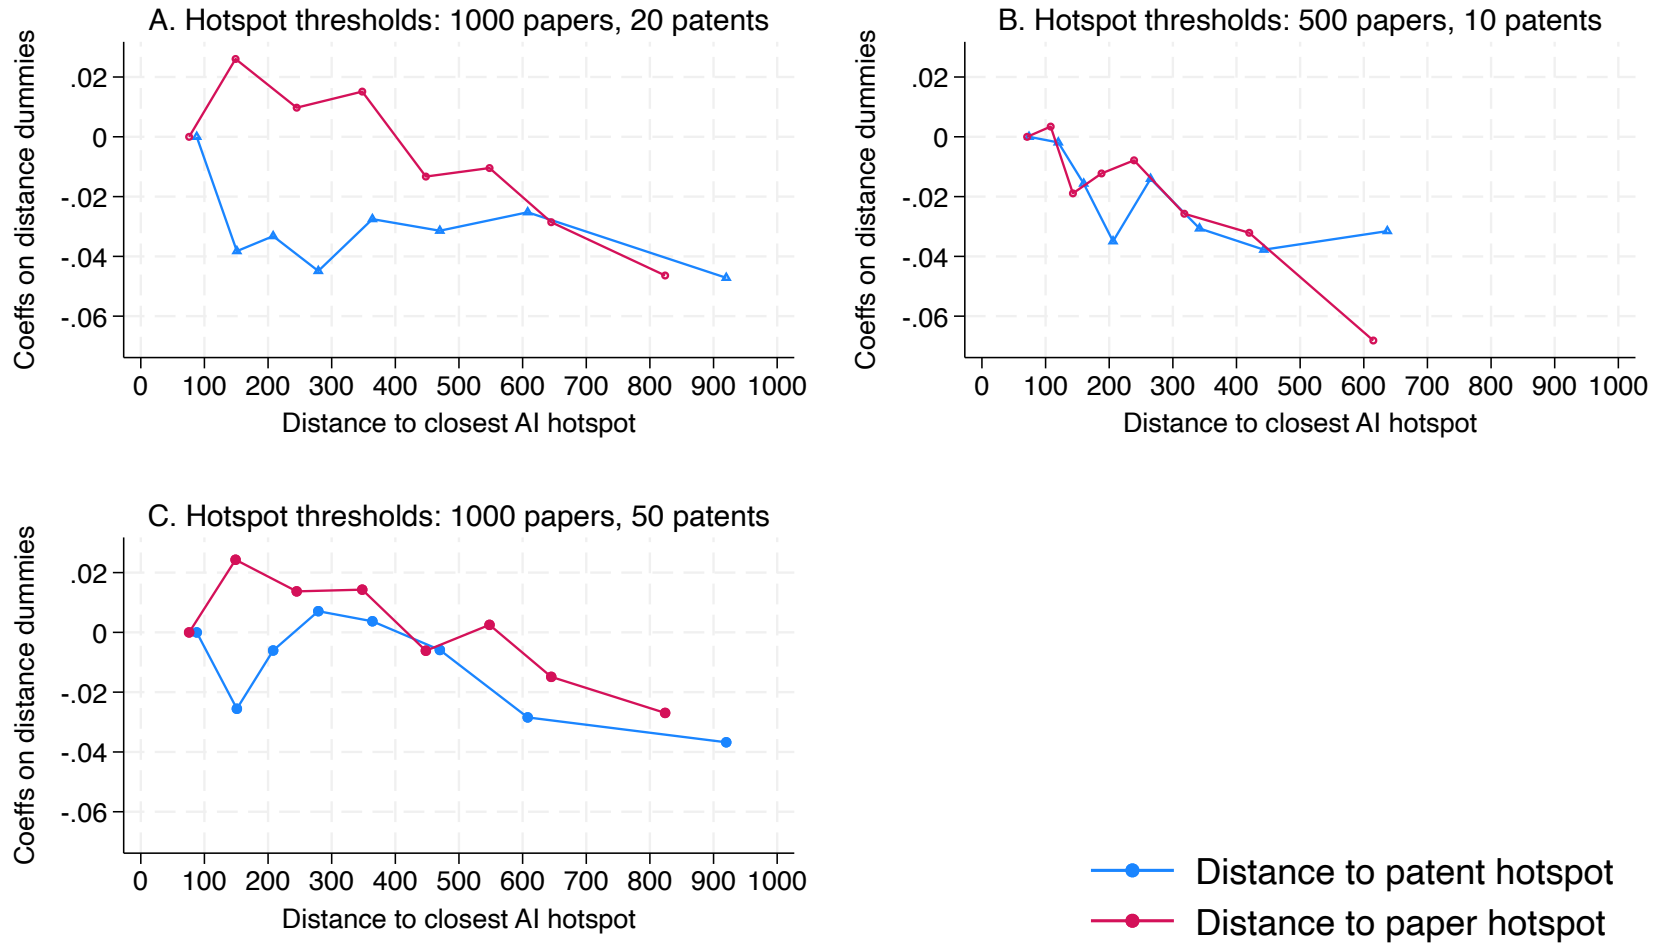


(d) 2018: 282 commuting zones with any publication



Note: Publications are the sum of patents and papers.

Figure 10: Effect of distance with flexible specification



Notes: Coefficients from dummies for distance to closest commuting zone.

Table 1: Summary statistics

	Mean	Median	Min	Max	Obs
A. Δ AI job advertisement share (%)					
$\Delta=3$	0.061	0.036	-2.46	4.70	5776
$\Delta=7$	0.138	0.094	-2.82	4.90	2888
B. $P(\Delta \text{ AI job advertisement share}) > 0$					
$\Delta=3$	0.629	1	0	1	5776
$\Delta=7$	0.758	1	0	1	2888
C. Hotspot characteristics					
Publications in 1000+ publication hotspots	1948	1743	1010	6841	722
Papers in 1000+ paper hotspots	1872	1740	1002	6680	722
Patents in 20+ patent hotspots	96	71	22	534	722
D. Initial conditions covariates					
Any AI paper prior to 2007	0.48	0	0	1	722
Any AI patent prior to 2007	0.17	0	0	1	722
AI papers prior to 2007	154	0	0	6617	722
AI patents prior to 2007	5.7	0	0	535	722
Job advertisements 2007	16,891	2613	3	696,205	722
Population in 2000 in thousands	387	106	1.19	16,393	722
IT share 2007 (%)	9.0	7.7	0	42.9	722
E. Distances (km)					
To closest 1000+ publication hotspot	356	322	40	1138	722
To closest 1000+ paper hotspot	356	322	40	1138	722
To closest 20+ patent hotspot	392	330	40	1196	722
Radius enclosing 2500 publications	357	325	40	1055	722
To closest populous commuting zone	310	280	8.7	1004	722
To other commuting zones (average)	1542	1443	1144	2673	722
To closest commuting zone	72.4	67.0	7.3	226	722
F. Differenced covariates $\Delta=3$					
AI papers	18.1	0	-101	5525	5776
AI patents	0.4	0	-40	326	5776
Log job advertisements	0.18	0.25	-3.5	2.1	5776
IT job ad share (%)	-0.3	-0.9	-45	59	5776
G. Differenced covariates $\Delta=7$					
AI papers	37.9	0	-49	6611	2888
AI patents	0.75	0	-33	419	2888
Log job advertisements	0.50	0.55	-2.75	3.26	2888
IT job ad share (%)	-1.3	-2.2	-38.5	26.1	2888
H. Mechanism covariates					
Travel time to closest hotspot (minutes)	236	231	38	588	722
Time difference of one hour?	0.17	0	0	1	722
Computer ads by firms also hiring in hotspot (%)	0.7	0.3	0	8.3	722
Immigration pc from closest hotspot (%)	0.19	0.1	0	2.3	722
CZ in same state as 1000+ publication hotspot	0.328	0	0	1	722
Any 1000+ publication hotspot in CZ's state	0.402	0	0	1	722

Notes: the distance to the closest populous commuting zone is the distance to the closest of the most populous 31 commuting zones in the United States; this is used for hotspots of at least 1000 publications and of at least 1000 papers. Immigration per capita is annual averaged over the 1990s. Alaska and Hawaii are excluded.

Table 2: Effect of distance to closest AI hotspot on change in AI jobs' share in advertisements

	(1)	(2)	(3)	(4)	(5)	(6)
A. Threshold: 1000 publications	-0.149*** (0.018)	-0.099*** (0.016)	-0.076*** (0.017)	-0.117*** (0.023)	-0.085** (0.024)	-0.079*** (0.022)
Pseudo R ²	0.04	0.16	0.16	0.16	0.18	0.23
B. Threshold: 1000 papers	-0.147*** (0.018)	-0.096*** (0.016)	-0.073*** (0.016)	-0.112*** (0.023)	-0.084*** (0.023)	-0.079*** (0.022)
Pseudo R ²	0.04	0.16	0.16	0.16	0.18	0.23
C. Threshold: 20 patents	-0.121*** (0.014)	-0.085*** (0.011)	-0.050*** (0.013)	-0.041** (0.017)	-0.050*** (0.016)	-0.047*** (0.017)
Pseudo R ²	0.04	0.15	0.15	0.15	0.18	0.23
D. Thresholds:						
1000 Papers	-0.082*** (0.027)	-0.075*** (0.023)	-0.071*** (0.022)	-0.081*** (0.030)	-0.075** (0.032)	-0.081*** (0.029)
20 Patents	-0.074*** (0.021)	-0.022 (0.018)	-0.004 (0.020)	-0.023 (0.019)	-0.027 (0.022)	-0.026 (0.018)
Pseudo R ²	0.04	0.16	0.16	0.16	0.18	0.23
CZ AI publications 2006; hotspot(s) publications 2006	--	Yes	Yes	Yes	Yes	Yes
Log job ads 2007; log population 2000; IT share in advertisements 2007	--	--	Yes	Yes	Yes	Yes
Distance to closest populous CZ(s); Population(s) of closest populous CZ(s)	--	--	--	Yes	Yes	Yes
Average distance to other CZs; distance to closest CZ; log population in closest CZ	--	--	--	--	Yes	Yes
Δ log ads, Δ IT share, Δ AI papers, Δ AI patents	--	--	--	--	--	Yes

Notes: Coefficients on distance(s) to the closest hotspot multiplied by 1000; coefficients in different panels and columns are from different median regressions. The dependent variable is the seven-year difference in AI jobs' share of all job advertisements (%). 2888 observations using data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. All regressions include year dummies. The controls for the commuting zone's AI publications through 2006 are dummies for any AI paper, any AI patent, the number of AI papers and its square and the number of AI patents. Controls for the closest hotspot or hotspots' AI publications through 2006 are publications and its square or AI papers and its square and/or AI patents, depending on how the hotspot(s) are defined. The definition of a populous commuting zone depends upon the number of commuting zones designated as hotspots, and the number of such controls equals the number of hotspots considered. The control for the population of the closest populous commuting zones is a quadratic, and the number of such quadratics equals the number of hotspots considered. Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 3: Sensitivity of distance coefficient to estimation method, sample and hotspot choices

	7-year differences			OLS (4)	3-year Median (5)
	(1)	Median regression (2)	(3)		
Distance to closest paper hotspot	-0.081*** (0.029)	-0.097*** (0.032)	-0.089** (0.034)	-0.121** (0.048)	-0.034*** (0.012)
Distance to closest patent hotspot	-0.026 (0.018)	-0.042 (0.029)	-0.047** (0.023)	0.015 (0.048)	0.004 (0.012)
Observations	2888	2888	2166	2888	5776
Pseudo R ² / R ²	0.23	0.23	0.29	0.25	0.13
Hotspot threshold (papers, patents)	1000, 20	500, 10	1000, 20	1000, 20	1000,20
Years included	All	All	Drop 2014-2007	All	All

Notes: Coefficients on distances to the closest hotspot multiplied by 1000. The dependent variable is the difference in AI jobs' share of all job advertisements (%). Data for 2007 and 2010-2019 except column 3, excluding Alaskan and Hawaiian commuting zones. Column 1 is same as Table 2 panel D column 5 and columns 1 and 3-5 have the same covariates. Certain covariates of column 2 are adapted to reflect the different hotspot definitions (see text). Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 4: Summary statistics for the 7-year change in AI job advertisement share (%), by industry

	Mean (1)	Median (2)	Share >0 (3)
All	0.138 (0.259)	0.036	0.758
Agriculture, Utilities, Mining, Construction, Manufacturing	0.083 (0.018)	0	0.342
Wholesale trade, Retail trade, Warehousing, Transportation	0.018 (0.236)	0	0.195
Information	0.212 (2.451)	0	0.232
Finance, Insurance	0.270 (0.937)	0	0.404
Real Estate, Professional and scientific services, Administration	0.245 (0.822)	0.069	0.533
Education, Health	0.049 (0.347)	0	0.432
Arts and recreation, Accommodation	0.055 (0.797)	0	0.169
Other services, Public administration	0.172 (2.844)	0	0.297
Missing industry	0.255 (0.420)	0.151	0.658

Notes: 2888 observations, excluding Alaskan and Hawaiian commuting zones. Commuting zones without job advertisements in a particular industry are assigned 0% AI job share in that industry. Each industry's statistics is based on a different underlying sample of job advertisements. Standard deviations are in parentheses.

Table 5: Impact of distance from AI hotspot on change on AI job advertisement share by industry

	7-year differences			3-year diffs
	Δ AI share	Δ AI share	P(Δ AI share>0)	P(Δ AI share>0)
	OLS (1)	Median (2)	OLS (3)	OLS (4)
All	-0.121*** (0.036)	-0.079*** (0.022)	-0.168*** (0.064)	-0.205*** (0.043)
Agriculture, Utilities, Mining, Construction, Manufacturing	0.084 (0.069)	--	-0.016 (0.068)	-0.024 (0.040)
Wholesale trade, Retail trade, Warehousing, Transportation	0.078 (0.054)	--	0.039 (0.065)	0.018 (0.038)
Information	-0.231 (0.192)	--	-0.070 (0.059)	-0.024 (0.040)
Finance, Insurance	-0.153 (0.105)	--	-0.191** (0.063)	-0.115** (0.043)
Real Estate, Professional and scientific services, Administration	-0.270** (0.114)	-0.069** (0.029)	-0.236*** (0.056)	-0.103** (0.040)
Education, Health	0.050 (0.100)	--	-0.106 (0.075)	-0.102** (0.050)
Arts and recreation, Accommodation	0.161 (0.098)	--	0.100 (0.073)	0.007 (0.052)
Other services, Public administration	0.451 (0.583)	--	-0.119 (0.077)	-0.069 (0.045)
Missing industry	-0.391*** (0.051)	-0.285*** (0.041)	-0.358*** (0.062)	-0.217*** (0.040)
Observations	2888	2888	2888	5776

Notes: Coefficients on distance to the closest hotspot (at least 1000 publications) multiplied by 1000; each coefficient is from a different regression. The dependent variables are differences in AI jobs' share of all job advertisements (%). 2888 observations using data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. All columns include covariates of Table 2 panel A column 5. The NAICS 2-digit codes are: 11, 21-23, 31-33 for Agriculture, Utilities, Mining, Construction, Manufacturing; 42, 44-45, 48-49 for Wholesale trade, Retail trade, Warehousing, Transportation; 51 for Information; 52 for Finance and insurance; 53-56 for Real Estate, Professional and scientific services, Administration; 61-62 for Education and Health; 71-72 for Arts and recreation, Accommodation; 81, 92 for Other services, Public administration. Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 6: Summary statistics for the 7-year change in AI job advertisement share (%), by occupation

	Mean (1)	Median (2)	Share >0 (3)
A. Occupations (2-digit)			
Computer and mathematical	1.219 (1.839)	0.883	0.723
Architecture and engineering	0.354 (2.224)	0	0.498
Management	0.253 (0.974)	0	0.472
Business and finance	0.148 (0.591)	0	0.348
Other	0.022 (0.169)	0.001	0.505
B. Computer/math occupations (6-digit)			
Research computer and mathematical, mathematicians, statisticians, operations research	7.320 (13.632)	2.342	0.541
Computer programmers	0.538 (4.350)	0	0.239
Software developers, applications	1.313 (2.794)	0.445	0.541
Web developers	0.710 (3.425)	0	0.247
Network and computer systems administrators	0.076 (2.395)	0	0.117
Computer user support specialists	0.088 (2.031)	0	0.142
Computer and mathematical not elsewhere classified	1.039 (3.419)	0	0.498

Notes: 2888 observations, excluding Alaskan and Hawaiian commuting zones. Commuting zones without job advertisements in a particular occupation are assigned 0% AI job share in that occupation. Each occupation's statistics is based on a different underlying sample of job advertisements. Standard deviations in parentheses.

Table 7: Impact of distance from AI hotspot on change on AI job advertisement share by occupation

	7-year differences			3-year diffs
	Δ AI share	Δ AI share	P(Δ AI share > 0)	P(Δ AI share > 0)
	OLS (1)	Median (2)	OLS (3)	OLS (4)
A. Occupations (2-digit)				
Computer and mathematical	-1.019*** (0.214)	-0.552** (0.229)	-0.237*** (0.055)	-0.226*** (0.048)
Architecture and engineering	-0.293* (0.166)	--	-0.159** (0.063)	-0.148*** (0.042)
Management	0.028 (0.086)	--	-0.073 (0.069)	-0.042 (0.048)
Business and finance	-0.083 (0.090)	--	-0.159** (0.064)	-0.072 (0.046)
Other	0.021 (0.044)	--	-0.188** (0.074)	-0.111** (0.045)
B. Computer/math occupations (6-digit)				
Research computer and mathematical, mathematicians, statisticians, operations research	-3.962** (1.893)	-2.017 (1.253)	-0.212** (0.064)	-0.159*** (0.044)
Computer programmers	-1.004 (0.629)	--	-0.096 (0.062)	-0.088** (0.039)
Software developers, applications	-1.804*** (0.452)	-0.623*** (0.176)	-0.362*** (0.058)	-0.321*** (0.044)
Web developers	-0.485 (0.489)	--	-0.056 (0.058)	-0.067 (0.040)
Network and computer systems administrators	0.541** (0.257)	--	0.128** (0.056)	0.018 (0.039)
Computer user support specialists	-0.319* (0.182)	--	-0.066 (0.050)	0.006 (0.031)
Computer and mathematical not elsewhere classified	-0.865*** (0.298)	-0.318** (0.130)	-0.284*** (0.053)	-0.159*** (0.040)

Notes: Coefficients on distance to the closest hotspot (a least 1000 publications) multiplied by 1000; each coefficient is from a different regression. Data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. All columns include covariates of Table 2 panel A column 5. Panel A 2-digit SOC codes are: Computer and mathematical (15); Architecture and engineering (17); Management (11); Business and financial operations (13). Panel B 6-digit SOC codes are: Research computer and mathematical (15-1111), mathematicians (15-2021), statisticians (15-2031), operations researchers (15-2041); Computer programmers (15-1121); Software developers, applications (15-1132); Web developers (15-1134); Network and computer systems administrators (15-1142); Computer user support specialists (15-1151); Computer and mathematical not elsewhere classified (15-1199). Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 8: Regressions testing mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Distance to closest hotspot (1000+ publications)	-0.079*** (0.022)	-0.097*** (0.037)	--	-0.066*** (0.023)	-0.043* (0.024)	-0.046** (0.023)	-0.045* (0.023)	-0.077** (0.037)
Travel time (minutes) × 1000	--	0.053 (0.069)	-0.098*** (0.039)	--	--	--	--	0.078 (0.066)
Time difference (1-hour dummy)	--	--	--	-0.008 (0.007)	--	--	-0.003 (0.007)	-0.001 (0.007)
Computer ads posted by firms also posting in hotspot, % all ads	--	--	--	0.011** (0.005)	--	--	0.014** (0.006)	0.014** (0.006)
Immigration per capita from hotspot (1990s %)	--	--	--	0.043 (0.026)	--	--	-0.003 (0.028)	0.008 (0.028)
+ square	--	--	--	-0.024* (0.013)	--	--	-0.010 (0.013)	-0.012 (0.013)
Commuting zone and closest 1000+ publication hotspot in same state	--	--	--	--	0.034** (0.008)	0.023* (0.013)	0.027* (0.014)	0.029** (0.014)
Any 1000+ publication hotspot in same state as commuting zone	--	--	--	--	--	0.012 (0.012)	0.014 (0.012)	0.011 (0.012)
Pseudo R ²	0.23	0.24	0.23	0.24	0.24	0.24	0.024	0.24
Immigration effect at mean	--	--	--	0.033 (0.022)	--	--	-0.007 (0.023)	-0.004 (0.024)

Notes: Coefficients on distance to the closest hotspot (a least 1000 publications) multiplied by 1000; median regressions. The dependent variable is the seven-year difference in AI jobs' share of all job advertisements (%). 2888 observations using data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. Column 1 is the same as Table 2 panel A column 5. The covariate in row 4 is computer and mathematical occupation advertisements posted by firms also posting computer and mathematical advertisements in the hotspot as a % of all ads. All columns include the unreported covariates of Table 2 panel A column 5. Time difference is for winter. Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 1: Skills used to designate a job advertisement as requiring Artificial Intelligence, by type

A. Unspecified Artificial intelligence and/or Machine learning only
B. Image processing Image processing
C. Other AI ChatBot, Amelia, ANTLR, Automatic Speech Recognition (ASR), Caffe Deep Learning Framework, Chatbot, Computational Linguistics, Computer Vision, Decision Trees, Deep Learning, Deeplearning4j, Google Cloud Machine Learning Platform, Gradient boosting, H2O (software), IBM Watson, Image Recognition, IPSoft, Ithink, Keras, Latent Dirichlet Allocation, Latent Semantic Analysis, Lexalytics, Lexical Acquisition, Lexical Semantics, Libsvm, Machine Translation (MT), Machine Vision, MLPACK (C++ library), MoSes, MXNet, Madlib, Mahout, Microsoft Cognitive Toolkit, Mlpy, ND4J (software), Natural Language Processing, Natural Language Toolkit (NLTK), Nearest Neighbor Algorithm, Neural Networks, Object Recognition, Object Tracking, OpenCV, OpenNLP, Pattern Recognition, Pybrain, Random Forests, Recommender Systems, Sentiment Analysis / Opinion Mining, Semantic Driven Subtractive Clustering, Semi-Supervised Learning, Sentiment Classification, Speech Recognition, Supervised Learning (Machine Learning), Support Vector Machines (SVM), TensorFlow, Text Mining, Text to Speech (TTS), Tokenization, Torch (Machine Learning), Unsupervised Learning, Virtual Agents, Vowpal, Wabbit, Word2Vec, Xgboost

Skills designated as AI by Alekseeva et al. (2021).

Appendix Table 2: Summary statistics from Burning Glass micro-data job advertisements

	Share (%)	AI required? (%)	Sample of ads requiring AI, with valid occupation Computer and math (%)	Observations
	(1)	(2)	(3)	(4)
A. By industry				
Agriculture, Utilities, Mining, Construction, Manufacturing	9.0	0.41	59.8	69,422
Wholesale trade, Retail trade, Warehousing, transportation	12.3	0.16	65.6	38,932
Information	3.0	1.10	70.1	63,147
Finance, Insurance	7.6	0.54	59.4	78,469
Real Estate, Professional, technical and scientific services, Administration	17.9	0.68	67.4	236,939
Education, Health	22.7	0.16	29.7	71,974
Arts and recreation, Accommodation	6.9	0.11	56.4	14,442
Other services, Public administration	4.5	0.22	50.7	19,187
Missing industry	16.0	0.38	74.2	121,836
All	100.0	0.37	62.6	714,348
B. By skills required				
Unspecified AI only (AI or Machine Learning)	0.14	100	68.3	264,852
Image processing	0.05	100	49.9	88,970
Other AI	0.19	100	61.6	360,526
Computer-aided design	0.76	0.36	8.1	1,465,449
Solar energy (excluding installation, sales, management)	0.07	0.33	7.3	129,134
Quantum computing	0.01	17.21	44.4	11,233

Notes: 2007-2019. 204,553,172 observations in column 1. The NAICS 2-digit codes are: 11, 21-23, 31-33 for Agriculture, Utilities, Mining, Construction, Manufacturing; 42, 44-45, 48-49 for Wholesale trade, Retail trade, Warehousing, Transportation; 51 for Information; 52 for Finance and insurance; 53-56 for Real Estate, Professional and scientific services, Administration; 61-62 for Education and Health; 71-72 for Arts and recreation, Accommodation; 81, 92 for Other services, Public administration.

Appendix Table 3: Sensitivity of distance coefficient to choice of AI hotspots

	(1)	(2)	(3)	(4)
Distance to closest paper hotspot, threshold of				
100 papers	-0.146 ^{***} (0.054)	-0.068 (0.055)	--	--
1000 papers	--	-0.081 ^{***} (0.027)	-0.068 ^{**} (0.032)	-0.054 ^{**} (0.026)
2000 papers	--	--	0.032 (0.033)	--
Distance to closest patent hotspot, threshold of				
1 patent	0.010 (0.048)	-0.006 (0.052)	--	--
20 patents	--	-0.022 (0.017)	-0.040 (0.036)	--
50 patents	--	--	--	-0.036 (0.042)
200 patents	--	--	-0.012 (0.027)	--
Pseudo R ²	0.23	0.24	0.24	0.23

Notes: Coefficients on distances to the closest hotspot multiplied by 1000; median regressions. The dependent variable is the seven-year difference in AI jobs' share of all job advertisements (%). 2888 observations using data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. The covariates in all columns correspond to those in Table 2 panel D column 5, with certain covariates adapted to reflect the different hotspot definitions (see text). Note for column 2 that the distance to the closest populous commuting zone is the same for a one-patent hotspot and a 100-paper hotspot. Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 4: Sensitivity of the coefficient on distance to the closest hotspot of 2000+ papers

	(1)	(2)	(3)	(4)	(5)
Distance to closest hotspot, threshold of					
1000 papers	--	--	-0.088*** (0.023)	-0.083*** (0.029)	-0.074*** (0.022)
2000 papers	0.029*** (0.011)	0.023** (0.010)	0.019* (0.011)	0.019 (0.031)	--
200 patents	--	--	--	0.001 (0.026)	--
Radius around commuting zone enclosing					
2500 papers	--	-0.092** (0.041)	0.061 (0.057)	0.067 (0.068)	--
200 patents	--	--	--	-0.021 (0.054)	--
Distance to closest 1000+ paper hotspot × hotspot has 2000+ papers	--	--	--	--	-0.038 (0.026)
Pseudo R ²	0.22	0.23	0.24	0.24	0.23

Notes: Coefficients on distance(s) to the closest hotspot multiplied by 1000; coefficients in different panels and columns are from different median regressions. The dependent variable is the seven-year difference in AI jobs' share of all job advertisements (%). 2888 observations using data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. All columns include covariates corresponding to those in Table 2 panel D column 5, with certain covariates adapted to reflect the different hotspot definitions (see text). Columns 2-4 also include the population(s) enclosed by the radius or radii, the number of papers within the paper-defined radius and the number of patents within the patent-defined radius. Column 5 also includes a dummy for distance to the closest populous commuting zone interacted with a dummy for a population over 3,798,017 in the populous commuting zone. Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 5: Summary statistics for the 3-year change in AI job advertisement share (%), by industry

	Mean (1)	Median (2)	Share >0 (3)
All	0.061 (0.223)	0.036	0.629
Agriculture, Utilities, Mining, Construction, Manufacturing	0.043 (0.010)	0	0.287
Wholesale trade, Retail trade, Warehousing, Transportation	0.010 (0.225)	0	0.172
Information	0.110 (2.850)	0	0.202
Finance, Insurance	0.128 (0.826)	0	0.322
Real Estate, Professional and scientific services, Administration	0.116 (0.682)	0	0.458
Education, Health	0.020 (0.301)	0	0.342
Arts and recreation, Accommodation	0.029 (0.567)	0	0.133
Other services, Public administration	-0.038 (2.894)	0	0.253
Missing industry	0.124 (0.371)	0.033	0.545

Notes: 5776 observations, excluding Alaskan and Hawaiian commuting zones. Commuting zones without job advertisements in a particular industry are assigned 0% AI job share in that industry. Each industry's statistics is based on a different underlying sample of job advertisements. Standard deviations in parentheses.

Appendix Table 6: Summary statistics for the 3-year change in AI job advertisement share (%), by occupation

	Mean (1)	Median (2)	Share >0 (3)
A. Occupations (2-digit)			
Computer and mathematical	0.589 (1.643)	0.316	0.614
Architecture and engineering	0.179 (1.777)	0	0.179
Management	0.069 (1.276)	0	0.377
Business and finance	0.066 (0.567)	0	0.275
Other	0.009 (0.151)	0	0.413
B. Computer/math occupations (6-digit)			
Research computer and mathematical, mathematicians, statisticians, operations research	3.485 (12.878)	0	0.444
Computer programmers	0.266 (4.048)	0	0.193
Software developers, applications	0.628 (3.590)	0	0.445
Web developers	0.321 (2.862)	0	0.186
Network and computer systems administrators	0.042 (2.327)	0	0.111
Computer user support specialists	0.043 (1.706)	0	0.119
Computer and mathematical not elsewhere classified	0.535 (3.155)	0	0.393

Notes: 5776 observations, excluding Alaskan and Hawaiian commuting zones. Commuting zones without job advertisements in a particular occupation are assigned 0% AI job share in that occupation. Each occupation's statistics is based on a different underlying sample of job advertisements. Standard deviations in parentheses.

Appendix Table 7: Summary statistics for the 7-year change in AI job advertisement share (%), by skill type

	Mean (1)	Median (2)	Share >0 (3)
A. AI skills (mutually exclusive)			
Unspecified (AI and/or Machine Learning only)	0.074 (0.149)	0.034	0.648
Image processing	0.003 (0.067)	0	0.350
Other AI	0.061 (0.173)	0.042	0.680
B. Non-AI skills			
Computer-aided design	0.006 (0.511)	0.013	0.511
Solar energy (excluding installation, sales, management)	0.007 (0.121)	0	0.400
Quantum computing	0.002 (0.023)	0	0.102

Notes: 2888 observations, excluding Alaskan and Hawaiian commuting zones. Data are based on the full set of job advertisements. Standard deviations in parentheses.

Appendix Table 8: Impact of distance from AI hotspot on change on AI job advertisement share by type of skills required

	7-year differences			3-year diffs
	Δ AI share	Δ AI share	P(Δ AI share>0)	P(Δ AI share>0)
	OLS (1)	Median (2)	OLS (3)	OLS (4)
A. AI skills (mutually exclusive)				
Unspecified only (AI and/or Machine Learning)	-0.088*** (0.017)	-0.045*** (0.016)	-0.353*** (0.061)	-0.274*** (0.045)
Image processing	-0.012 (0.009)	--	-0.064 (0.076)	-0.107** (0.048)
Other AI	-0.020 (0.029)	-0.046** (0.014)	-0.168** (0.067)	-0.125*** (0.041)
B. Non-AI skills				
Computer-aided design	-0.028 (0.058)	0.032 (0.028)	0.021 (0.071)	0.041 (0.049)
Solar energy (excluding installation, sales, management)	-0.006 (0.015)	--	0.034 (0.071)	0.049 (0.039)
Quantum computing	-0.006* (0.004)	--	0.003 (0.040)	-0.014 (0.025)

Notes: Coefficients on distance to the closest hotspot (a least 1000 publications) multiplied by 1000; each coefficient is from a different regression. The types of AI skill are mutually exclusive. Data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. All columns include covariates of Table 2 panel A column 5. For all regressions, data are based on the full set of job advertisements. Standard errors clustered by commuting zone are in parentheses.

*** p<0.01, ** p<0.05, * p<0.

Appendix Table 9: Summary statistics for the 3-year change in AI job advertisement share (%), by skill type

	Mean (1)	Median (2)	Share >0 (3)
A. AI skills (mutually exclusive)			
General (AI and/or Machine Learning only)	0.036 (0.107)	0.004	0.513
Image processing	0.002 (0.058)	0	0.335
Other AI	0.023 (0.171)	0.012	0.555
B. Non-AI skills			
Computer-aided design	-0.012 (0.446)	-0.002	0.482
Solar energy (excluding installation, sales, management)	0.005 (0.101)	0	0.373
Quantum computing	0.001 (0.017)	0	0.069

Notes: 5776 observations, excluding Alaskan and Hawaiian commuting zones. Data are based on the full set of job advertisements. Standard deviations in parentheses.

Appendix Table 10: Regressions testing mechanisms using OLS

	Coefficients					Gelbach
	(1)	(2)	(3)	(4)	(5)	components
Distance to closest hotspot (1000+ publications)	-0.121*** (0.036)	-0.160*** (0.053)	--	-0.134** (0.055)	-0.066 (0.039)	--
Travel time (minutes) × 1000	--	0.096 (0.103)	-0.134* (0.070)	0.179* (0.103)	--	--
Time difference (1-hour dummy)	--	--	--	-0.006 (0.014)	-0.007 (0.014)	-0.002 (0.004)
Computer ads posted by firms also posting in hotspot, % all ads	--	--	--	0.016** (0.006)	0.016** (0.006)	-0.006 (0.004)
Immigration per capita from hotspot (1990s %)	--	--	--	0.056 (0.052)	0.046 (0.052)	--
+ square	--	--	--	-0.046 (0.028)	-0.041 (0.028)	-0.006 (0.016)
Commuting zone and hotspot in same state	--	--	--	0.021 (0.018)	0.020 (0.018)	-0.025 (0.023)
Any hotspot in same state	--	--	--	0.012 (0.017)	0.012 (0.017)	-0.016 (0.024)
Pseudo R ²	0.23	0.25	0.25	0.25	0.25	--
Immigration effect at mean	--	--	--	0.038 (0.042)	0.031 (0.042)	--

Notes: Coefficients on distance to the closest hotspot (a least 1000 publications) multiplied by 1000; OLS. The dependent variable is the seven-year difference in AI jobs' share of all job advertisements (%). 2888 observations using data for 2007 and 2010-2019, excluding Alaskan and Hawaiian commuting zones. Column 1 is the same as Table 7 column 1 row 1. All columns include the unreported covariates of Table 2 panel A column 5. Standard errors clustered by commuting zone are in parentheses. Time difference is for winter.

*** p<0.01, ** p<0.05, * p<0.1

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

2037	Giuseppe Berlingieri Filippo Boeri Danial Lashkari Jonathan Vogel	Capital-skill complementarity in firms and in the aggregate economy
2036	Alessandra Fenizia Tom Kirchmaier	Not incentivized yet efficient: Working from home in the public sector
2035	Elodie Andrieu John Morrow	Can firm subsidies spread growth?
2034	Miquel-Àngel Garcia López Luz Yadira Gómez-Hernández Rosa Sanchis-Guarner	Highway traffic in Britain: The effect of road capacity changes
2033	Stephan Heblich Stephen J. Redding Yanos Zylberberg	The distributional consequences of trade: Evidence from the Grain Invasion
2032	Brian Bell Philip Johnson	Immigrant downgrading: New evidence from UK panel data
2031	Ying Chen Paul Cheshire Xiangqing Wang You-Sin Wang	Valuing consumption services as technology transforms accessibility: Evidence from Beijing
2030	Johannes Boehm Thomas Chaney	Trade and the end of antiquity
2029	Jay Euijung Lee Martina Zanella	Learning about women's competence: The dynamic response of political parties to gender quotes in South Korea
2028	Christos Genakos Blair Yuan Lyu Mario Pagliero	Asymmetric pass-through and competition

2027	Fabrice Defever Alejandro Riaño Gonzalo Varela	Evaluating the impact of export finance support on firm-level export performance: Evidence from Pakistan
2026	Andrés Barrios-Fernández Christopher Neilson Seth Zimmerman	Elite universities and the intergenerational transmission of human and social capital
2025	Agnes Norris Keiller	Brexit and investment
2024	Lidia Farré Libertad González Claudia Hupkau Jenifer Ruiz-Valenzuela	Paternity leave and child development
2023	Gabriel M. Ahlfeldt Nikodem Szumilo Jagdish Tripathy	Housing-consumption channel of mortgage demand
2022	Agnes Norris Keiller Tim Obermeier Andreas Teichgraeber John Van Reenen	When trade drives markup divergence. An application to auto markets
2021	Amanda Dahlstrand Nestor Le Nestour Guy Michaels	Online versus in-person services: Effects on patients and providers
2020	Hanwei Huang Kalina Manova Oscar Perelló Frank Pisch	Firm heterogeneity and imperfect competition in global production networks
2019	Jan David Bakker Nikhil Datta	Avenging the tenants: Regulating the middle man's rents

The Centre for Economic Performance Publications Unit

Tel: +44 (0)20 7955 7673 Email info@cep.lse.ac.uk

Website: <http://cep.lse.ac.uk> Twitter: @CEP_LSE