Is Distance from Innovation a Barrier to the Adoption of Artificial Intelligence? *

James Bessen †
Boston University

Iain Cockburn ‡
Boston University and NBER

Jennifer Hunt §
Rutgers University and NBER

August 10, 2022

*We thank Bledi Taska and Burning Glass Technologies for access to data; Erich Denk of the Technology and Policy Research Initiative at Boston University for extensive data work on the Burning Glass Technologies files; and Aureo de Paula 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 and the 2021 NBER Artificial Intelligence seminar for helpful comments on more preliminary work. Hunt is also affiliated with the University of Sydney, IZA (Bonn), CEPR (London) and DIW (Berlin).

†jbessen@bu.edu. Technology and Policy Research Initiative, Boston University School of Law, 765 Commonwealth Avenue, Boston, M.A. 02215.

‡cockburn@bu.edu. Strategy and Innovation Department, Questrom School of Business, Boston University, 595 Commonwealth Avenue, Boston, M.A. 02215

§jennifer.hunt@rutgers.edu. Department of Economics, Rutgers University, 75 Hamilton Street, New Brunswick, N.J. 08901.

Abstract

We investigate whether online vacancies for jobs requiring Artificial Intelligence (AI) skills grow more slowly in U.S. locations farther from AI "innovation hotspots." To do so, we create a dataset of AI publications and define hotspots based on locations' cumulative number of AI publications by 2006. The source for job vacancies is online job advertisements scraped by Burning Glass Technologies from 2007–2019. With a hotspot defined as a commuting zone with at least 1000 AI publications, a 10% greater distance from a hotspot (about a standard deviation) reduces a commuting zone's growth in AI jobs' share of job advertisements by 2-3% of median growth. Analysis by industry shows distance is a barrier for sectors likely to be adopting AI, like finance and insurance, and not merely for sectors likely to be innovating in AI, like the information sector. Analysis by type of AI skill also suggests distance is a barrier for both innovation and adoption. Further results suggest distance from an AI innovation hotspot is a barrier to the adoption of AI because the specialized workers necessary for adoption are unwilling to move long distances from hotspots to non–hotspots with unskilled labor forces.

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 development; 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 new aspect of the latter, by examining the geography of U.S. firms' adoption of Artificial Intelligence (AI) in response to AI research.

The importance of distance for the diffusion of inventive activity has received considerable attention. Despite the longstanding availability of the telephone and modern means of transportation, personal contact is hypothesized to be important for stimulating invention. This could take the form of an inventor interacting with a potential inventor at another firm or university, which is more likely to happen if the two people work or live in physical proximity. A more indirect effect of distance is easier to test: that inventors diffuse inventive activity by moving themselves, either to another firm in the same location or to another location. Distance is then a barrier because distance is a barrier to migration. If physical mobility is central to the diffusion of inventive activity, distance could persist as a barrier even as widespread email and video conferencing have reduced the cost of communication. The importance of inventors' moving has been demonstrated empirically[1], and while there has been some debate, the empirical evidence overall supports the hypothesis that distance is a barrier to the diffusion of inventive activity to potential inventors.[2]

A related literature examines how the adoption of technology, often across countries,

[1] For within–country firm to firm moves see 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).

[2] 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. Singh and Marx (2013) find political borders, including those within countries, to be larger barriers than distance itself. For analysis of country R&D as a proxy for innovation, see Keller (2002) and papers in Keller's (2004) survey.

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. This adoption literature has also found distance to be a barrier[3], but finds the barrier to be lower for multiestablishment or multinational firms, which presumably have internal communication channels and coordination.[4]

Our paper seeks instead to examine whether distance constitutes a barrier between technology research and technology adoption. We choose to examine AI because there are data on its use beginning when relatively few firms had adopted it; because adoption has since spread rapidly; and because this spread is potentially important for future economic growth.[5] To measure what we will term 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 adoption using job vacancy information from U.S. online job advertisements scraped by Burning Glass Technologies from 2007–2019.

The only existing analysis of geographic links between knowledge creation or technology development and technology adoption 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 Burning Glass job advertisements involving the technology group. Bloom et al. do not, however, consider explicitly the link between distance from a pioneer location and the growth of the technologies, nor do they consider new knowledge emerging

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

[4] Branstetter, Blennon and Jensen (2018).

[5] Aghion, Jones and Jones (2017); Goldfarb, Taska and Teodoridis (2019).

as scientific publications rather than patents. Thus, the contributions of our paper are a new question, its application to a new technology, and new data linking AI publications and job advertisements.[6]

We approach the question by dividing the United States into 741 commuting zones and using them as a panel. Our first approach involves designating as innovation hotspots those commuting zones whose cumulative AI publications before our study period were 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 (log) distance to the closest innovation hotspot. A negative effect of distance means that distance is a barrier, while a null effect could either mean that the barrier is so high that commuting zones have no effect on one another, or that there is no barrier. Our second identification strategy defines the key distance covariate as the (log) radius of the circle around the commuting zone which encloses more than a certain threshold of cumulative AI publications before our study period (exclusive of the commuting zone's own publications). This is essentially a variant of the first identification strategy incorporating more AI publication information.

We find that as the hotspot threshold surpasses 300 publications, a threshold met by 11% of commuting zones, a hotspot's AI publications affect other commuting zones' AI vacancies. At a threshold of 1000 publications, where the effect size is no longer sensitive to the threshold, a 10% greater distance from a hotspot (about a standard deviation) reduces a commuting zone's growth in AI jobs' share of job advertisements by 2–3% of median growth. Our findings are robust to the second approach using the (log) radius of the circle enclosing a given number of AI publications and to measuring job advertisements cumulatively over time instead of contemporaneously.

To distinguish job vacancies reflecting adoption and those reflecting innovation, we perform analysis by type of AI and by industry. Analysis by type of AI skills required

[6] Other related papers are by Andersson, Quigley and Wilhemsson (2009),and Dittmar and Meisenzahl (2022), who look at the impact of universities on local innovation, and Acemoglu, Autor, Hazell and Restrepo (2021), who examine the growth of AI job advertisements in the Burning Glass Technologies data, but do not consider geography.

shows that distance from a hotspot reduces growth in job advertisements requiring the use of AI applications, not merely advertisements requiring more technical AI skills. Analysis by industry shows distance is a barrier for sectors likely to be adopting AI, like finance and insurance, and not merely for sectors likely to be innovating in AI, like the information sector. These results show that distance from innovation is a barrier to the adoption of innovation, and not just to innovation itself.[7]

Further analysis suggests that distance is a barrier to the adoption of AI because it is a barrier to migration of the necessary specialized workers. Workers at firms or institutions innovating in AI may form an important hiring pool for firms seeking to adopt AI. Adopting AI is not yet as simple as buying off–the–shelf software, but rather requires adapting existing software: even in job advertisements requiring the use of AI applications, 53% specify computer and mathematical occupations, while 59% of AI job advertisements in finance and real estate specify these occupations. If a firm's commuting zone (or prospective commuting zone) has an unskilled labor force or is geographically isolated, or is the site of AI innovation irrelevant for the firm's desired application, the firm must search further afield for workers (or locate elsewhere). In this case, greater distance from an AI hotspot is an impediment to hiring, apparently because specialized workers cannot be enticed to move long distances (except to other hotspots). This mechanism echoes knowledge diffusion through movement of inventors themselves, studied in the literature mentioned above, but the workers in question need not themselves have published or patented.

The evidence for this mechanism comes indirectly from distinguishing between job advertisements posted by an employment agency and those posted by the ultimate employer. Vacancy–level analysis indicates that employment agencies are used disproportionately by firms which have difficulty hiring locally: a job advertisement is more likely to be posted by an agency if the commuting zone is geographically isolated or has few college workers or if the vacancy is for computer and mathematical occupations or workers with IT or AI skills. Further, commuting–zone level analysis differs when based on an underlying

[7] A future version of the paper will distinguish between AI journal articles and AI patents.

sample of advertisements placed by the ultimate employer versus those placed by employment agencies: in the former case, there is no effect of distance to an AI hotspot on the share of AI advertisements, while in the latter case, a 10% increase in distance reduces the AI share by 5.5–6.5% of median growth. This suggests that it is firms that cannot hire locally which are less likely to attempt to adopt and hire in AI if they are far from an AI hotspot, or that firms anticipating the need for AI skills do not set up in commuting zones far from hotspots. In regressions whose underlying micro sample is all vacancies, we find imprecisely estimated support for a more negative effect of distance in commuting zones with a higher share of agency advertisements.

Regression coefficients other than the coefficient on distance to a hotspot lend support to the hypothesized mechanism. When the micro sample is job advertisements posted by the ultimate employer, the growth in the commuting zone's AI share is faster when it has a large number of pre–2007 AI publications, while when the micro sample is job advertisements posted by an agency there is no such effect. This suggests that firms hiring non–locally are in commuting zones whose AI innovation is not of a type relevant to their potential AI applications and does not provide a pool of potential hires. Reinforcing this, pre–2007 AI publications have no effect on AI job growth in finance and insurance.

The effect of distance to the closest hotspot is greatly reduced when the number of local pre–2007 AI publications becomes large: for firms in a commuting zone that is itself an AI hotspot, the distance to the nearest other hotspot is almost irrelevant. Though this could have several explanations, one is that such firms have a wide selection of local workers and can also attract AI workers from far away who would be unwilling to move long distances to non–hotspots; this drawing power is part of the reason for the high skill of the hotspot labor force. In sum, distance from an AI innovation hotspot is a barrier to the adoption of AI because the specialized workers necessary for adoption are unwilling to move long distances from hotspots to non–hotspots with unskilled labor forces.

1 Data

We have created our own database of AI publications and patents, and merge it with Burning Glass Technologies job advertisement information.

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 citation linkages. 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. The authors of these publications work at firms and research institutes as well as universities.

Where possible, the location of each author was carefully geo-coded using information on their organizational affiliation at the time of publication. 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.[8] 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. Of the 442,563 publication-author pairs which we identified as having

[8] 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 affiliationid 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 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 publications into 741 commuting zones for each year.[9] 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 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. Suppemental data are the county (and hence commuting zone) unemployment rates from the Bureau of Labor Statistics, and the share of college graduates in each commuting zone's 2000 population from opportunityatlas.com.

1.2 Burning Glass Technologies job advertisements

Burning Glass Technologies is an employment analytics and labor market information firm which since 2007 has daily scraped the web's online job postings and produces files with duplicates eliminated standardized information for each advertisement. Its 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. Unfortunately, there are no data for 2008 and 2009, which influences our estimation strategy, so our sample period is February–December 2007, all years and

[9] 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.

months from 2010–2018, and January–July 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.

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. 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. We harmonize differing versions of employer name. 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 missing value for the employer name means the advertisement is posted by an employment agency [10] Burning Glass apparently also 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. Rather, since Burning Glass Technologies infers employer industry principally from the employer name, about half of vacancies (44%) with a missing firm name are also missing industry. Some employment agency names do appear, presumably because employment agencies do hire some workers.

We designate a job advertisement as being an AI job advertisement if the required skills include the general 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. (2019). This table also shows how we classify AI skills by type, seeking to distinguish skills used for adoption versus innovation. 37% of advertisements requiring AI simply require either "Artificial Intelligence" or "Machine Learning" skills, with no further detail specified. These unspecified AI skills comprise

[10] Burning Glass Technologies, personal communication.

our first category. The remaining categories are not mutually exclusive, though there is little overlap. Image processing, which seems less than tightly linked to AI, is the second category. The third category comprises skills in AI software that is a tool for more work in AI; the fourth category comprises skills in AI software that is an application of AI to be used by non–specialists (such as IBM Watson and recommender software); while the fifth, which we denote "R&D", comprises a set of detailed AI terms that imply knowledge of the underlying details of AI (such as supervised learning) as well as more general terms that constitute a field of research (such as computer vision).[11]

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, 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 are AI, so that small commuting zones may experience as large an effect of distance as large commuting zones.[12]

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.

To construct the independent variable we emphasize, we combine the distances with the hotspot information to compute the distance to the closest innovation hotspot for each commuting zone. Unless there is only one hotspot (a case we do not consider), even

[11] See Burning Glass Technologies (2019) for a description of how required skills are codified.

[12] For a small proportion of postings, the county is missing, but as state is never missing, missing counties are assigned randomly within the state.

hotspots have a closest hotspot. For use with this independent variable, 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 publication 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 present results using a different independent variable that does not use the concept of a hotspot. For each commuting zone, we calculate the radius of the circle around it which encompasses a given number of AI publications; we calculate this at the commuting zone level.

2 Methods

We choose as our primary dependent variable long differences (length k) in AI jobs' share of job advertisements in commuting zone c: $\Delta^k AI^s_{ct} = \frac{AI\ job\ ads}{All\ job\ ads}_{c,t} - \frac{AI\ job\ ads}{All\ job\ ads}_{c,t-k}$. We use shares to avoid having the variation in AI reflect variation in commuting zone population. We use long differences because at a given point in time, a large share of commuting zones have no AI job advertisements and many have only one or two, meaning short differences in shares would often reflect very small absolute changes in the number of AI advertisements. Because the number of job advertisements is often also small in such commuting zones, there is a considerable number of outliers in the change in AI share, a problem mitigated with longer differences. We avoid using 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. We therefore estimate this equation in our first identification approach, with our key dependent variable, distance to the nearest innovation hotspot, defined D_c^{Hot}:

$$\begin{aligned}\Delta^k AI_{ct}^s =& \alpha + \sigma log(D_c^{Hot}) \\ &+ \beta_1 AI\ Pub > 0_{c,t^*} + \beta_2 AI\ Pubs_{c,t^*} + \beta_3 (AI\ Pubs_{c,t^*})^2 \\ &+ \gamma_1 log(All\ job\ ads_{c,t^*}) + \gamma_2 log(Pop_{c,t^*}) \\ &+ \nu IT_{c,t^*}^s \\ &+ \phi_1 log(\bar{D}_c) + \phi_2 log(D_c^{Pop}) + \phi_3 log(D_c^{min}) \\ &+ \rho_1 \Delta^k AI\ Pubs_{c,t} + \rho_2 \Delta^k log(All\ job\ ads_{c,t}) + \rho_3 \Delta^k IT_{c,t}^s \\ &+ \eta_t + \Delta^k \epsilon_{ct},\end{aligned}$$

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 covariate of interest is σ. If σ is negative, distance constitutes a barrier to the adoption of innovation. 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 adoption.

The first set of additional controls captures initial conditions. A quadratic in the commuting zone's own cumulative AI publications through 2006 (quadratic rather than log due to the presence of zeros), $AI\ Pubs_{c,t^*}$, and a dummy for any such publication $AI\ Pub > 0_{c,t^*}$, capture the own–effect counterpart to the spillovers from innovation hotspots. We control for the initial number of job advertisements of all types, $log(All\ job\ ads_{c,2007})$, and the population in the most recent pre–study period census, $log(Pop_{c,2000})$, despite the fact that the dependent variable is scaled, to control for variation in the size of online job boards relative to population. To avoid the AI publication covariates picking up variation in non–AI IT, we control for IT's share of job advertisements in 2007 ($IT_{c,2007}^s$).

We also control for 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}. To ensure that σ is not capturing any general disadvantage due to distance from a large commuting zone as well as the disadvantage due to distance from an innovation hotspot, we control for the distance to the nearest large commuting zone ($\bar{D}_c^{Pop}$): these two distances are very positively correlated.

The last set of covariates is intertwined with the question of the conditions under which $\hat{\sigma}$ is unbiased. If AI publications cleanly measure innovation, and AI job advertisements cleanly measure adoption, and unobservable variables affecting commuting zones' propensity to adopt or adapt AI do not affect their propensity to innovate in AI, $\hat{\sigma}$ will be unbiased in regressions with the covariates described thus far. In this case, controlling for changes in the number of the commuting zone's own AI publications $\Delta^k AI\ Pubs_{c,t}$, the change in log job advertisements $\Delta^k log(All\ job\ ads_{c,t})$ and the change in the IT job advertisements' share in all advertisements $\Delta^k IT^s_{c,t}$ is likely to constitute overcontrolling: some or all of these could be the result of growth in AI adoption, rather than the cause, and their inclusion could bias $\hat{\sigma}$ upward toward zero.

However, some of the AI job vacancy growth reflects innovation, so for regressions considering all types of AI but omitting these covariates, $\hat{\sigma}$ will be biased down (the classic spatial spillover problem described in Gibbons and Overman 2012). Furthermore, it is plausible that there is a positive correlation between unobserved factors affecting innovation and adoption, a further reason $\hat{\sigma}$ is likely to be biased down in such specifications (distance to innovation will be negatively correlated with the error term including unobserved influences on adoption). Controlling for changes in the number of the commuting zone's own AI publications (for example) could reduce the downward bias stemming from both issues: this would control for the part of the growth in the dependent variable due to growth in innovation, and would proxy for unobserved determinants of growth in adoption. Our preferred specification is therefore the one including all the covariates in the equation above.

It seems likely that distance to AI publication hotspots is irrelevant for commuting zones that themselves generate a large number of AI publications, but important for commuting zones which themselves have few AI publications. To test this hypothesis, in some specifications we include controls for the interactions of log distance to a hotspot with all the covariates involving AI publications. To test the importance of hiring non–locally, as proxied by the use of employment agencies, in some specifications we also control for the share of job advertisements in the commuting zone for which the employer name was missing (indicating a posting by an employment agency) in the initial year of the difference, and its interaction with log distance to the closest hotspot.

Our second approach involves replacing $log(D_c^{Hot})$ with the radius of the circle enclosing N or more pre–2007 AI publications $log(R_c^N)$, exclusive of the commuting zone's own AI publications. We also replace the population control $log(Pop_{c,t^*})$ with the log of the population within the circle with radius $log(R_c^N)$, exclusive of the commuting zone's own population. In addition, we control for the (log) number of AI publications within the circle, since this varies due to the lumpy geographic nature of AI publications at the commuting zone level. This approach is not so much a different identification strategy as a specification using the pre–2007 AI publications data more fully. We considered a large number of other specifications, and explain in the Methodological Appendix why we did not pursue them.

Due to the significant number of zeros in the dependent variable despite the focus on long differences, we estimate the equation using median regression. This also downweights the large outliers in the outcome.[13] OLS point estimates of σ are somewhat more negative than median regression estimates, with larger standard errors. We use both the single twelve–year difference 2007–2019, which has the advantage of capturing long term effects with fewer outliers, but does not use most of the data, and the pooled seven–year differences 2007–2014, 2010–2017, 2011–2018, and 2012–2019, thus using data for all available years except 2013. For seven–year differences, we cluster standard errors by commuting zone, while we calculate robust standard errors for twelve–year differences.[14]

While it seems natural to form a panel using a dependent variable based on what we obtain directly from the data, job vacancies, it would be more desirable to base the dependent variable on AI employment rather than vacancies, since a change in vacancies

[13] 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; also, total job advertisements and distance to a hotspot are correlated.

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

represents an acceleration in employment. In the absence of this information, we rerun the estimation using cumulative AI vacancies, which would equal employment if those jobs were never destroyed or vacant again. In this regression, even the twelve–year difference uses data from all years.[15]

We also investigate the probability of a commuting zone having any AI job advertisement in 2018, conditional on having none in 2007. In these regressions, IT advertisements are expressed as numbers rather than shares, and IT in 2007 is captured with a quadratic in the number of advertisements. We focus on the longest difference, which is 2007–2018, since 2019 is only a partial year.

All these regressions establish whether distance is a barrier to the growth of AI job advertisements. Further analysis is designed to distinguish whether the barrier is to the adoption of AI, or merely to additional innovation in AI, and to shed light on the mechanism by which distance slows diffusion. For this purpose, we examine different AI types separately, and we investigate the role of distance by industry, using as the outcomes the number of AI job advertisements in a particular industry, divided by job advertisements in that industry.[16] Occupation is not very informative as a majority of AI vacancies posted are for computer and mathematical occupations (even for individual AI types and in 2019). In other regressions to establish the mechanism by which distance affects AI adoption, the commuting zone–year variables are calculated based on subsamples of the job advertisements with valid or missing industry and subsamples with valid or missing employer name. For a small number of commuting zones in some years, some of these subsamples have no observations.

We also run linear probability regressions at the job advertisement–level to estimate the probability of an advertisement being posted by an employment agency (have no associated employer name):

$$P(\textit{Posted by agency}_{ioc}) = \delta_0 + \delta_1 S_{ioct} + \delta_2 X_{ct} + \zeta_o + \eta_t + \nu_{ioc}, \tag{1}$$

[15] However, differences involving 2007 will be too small due to the missing 2008 and 2009 data.

[16] 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.

where i indexes job advertisement and o indexes occupation, S_{ioct} are controls for skills, X_{ct} are characteristics of the commuting zone (c) and ζ_o are 23 occupation dummies (including a dummy for missing occupation). We do not control for commuting–zone dummies since otherwise unexplained geographic variation is not of interest.

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 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, and that the IT jobs share is much higher (see the right scale) and evolves quite differently. The aggregate shares of these AI categories in all job advertisements are shown in Figure 3: the fastest growth is in the unspecified AI/Machine learning category. The Figure 4 maps indicating commuting zones' AI job advertisement shares show how the fraction of commuting zones with no AI job advertisement (white) shrank with time, and how the non–zero shares rose with time (as represented with darker shading) to a maximum of 4.0% in San Jose in 2019 (and one other small commuting zone).

In panel A of Table 1 we show that over the whole twelve–year study period 2007–2019, the mean AI job advertisement share increased by 0.28 percentage point, while the median increase was lower at 0.22 percentage point (first row). The minimum value of -2.82 percentage points and the maximum value of 3.34 percentage points confirm the existence of the outliers mentioned above: such large changes are caused by very small changes in the number of AI job advertisements in commuting zones with few job advertisements. When all seven–year changes are pooled in the fourth row, the mean increase is 0.14 percentage point and the median increase is 0.09 percentage point.

The lower panels of Table 1 shows the means of key covariates, including those based on AI publications. The national time–series for AI publications from 1950 onwards (a few publications are pre–1950) is shown in Figure 5. Publications (times number of authors) increased from 7 in 1950, to 12,089 in 2007, to 49,974 in 2018 and to 65,476 in 2019. The 2007–2018 increase is therefore much smaller in both absolute and percentage terms than the rise in AI job advertisements. Appendix Table 2 shows summary statistics based on the underlying vacancy micro–data.

One definition of an innovation hotspot we use is having at least 1000 cumulative publications by 2006, and Figure 6 depicts the number of publications for each of the 31 commuting zones satisfying this requirement. The three top publishers are Los Angeles, 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 and California is Seattle, W.A. in ninth place, and the ranks after Seattle include some Midwestern commuting zones as well as the Texan commuting zones of Austin and Houston. Some of the hotspots are recognizable as technology and university centers, others as university towns, and others as centers of military activity (Los Angeles is all three). The map in Figure 7 shows the distribution of these cumulative publications, while the succession of maps in Figure 8 shows that there is very slow diffusion of publishing through 2014, but faster diffusion afterwards.

4 Regression analysis

We begin by presenting various specifications of regressions in which the definition of an innovation hotspot is having at least 1000 pre–2007 publications, and another set in which the key distance variable is the radius of the circle enclosing 1000 pre–2007 publications. We then choose a preferred specification, and analyze the sensitivity to changes in the distance thresholds. Next, we investigate sources of heterogeneity in the distance effect, to distinguish among advertisements for jobs in AI innovation versus adoption. Finally, seeking the mechanism through which distance influences AI job advertisements, we investigate the characteristics of job advertisements posted by employment agencies.

4.1 Results for AI hotspot publications threshold of 1000

The effect of distance to the closest innovation hotspot on the change in AI job advertisements ($\times 100$) as a share of all job advertisements, is presented in Table 2 (the full sets of coefficients are presented in Appendix Tables 3 and 4). Consider panel A, containing results from seven–year differences: the coefficients on distance are always statistically significantly negative. In the first column, the only controls are distance to the closest hotspot and three controls for the commuting zone's AI publications through 2006. The coefficient of -0.029 implies that a 10 percent greater distance, which is approximately the standard deviation of the distance, reduces the median growth rate of AI jobs' share by (0.029)(0.1)=0.0029 percentage point. This is 3.1% of the median growth rate of 0.093 percentage points in Table 1, a modest effect.

In column 2, the addition of other initial conditions, average distance to other commuting zones and distance to the nearest commuting zone, leave the coefficient of interest little changed at -0.031. In column 3, we add seven–year differences in log job advertisements, IT jobs' share and AI publications, which leaves the coefficient on distance unchanged at -0.031.[17] The addition of the distance to the closest large commuting zone (one of the 31 most populous, since there are 31 AI publication hotspots) in column 4 increases the coefficient on slightly, to -0.024: for this definition of hotspot, the distance to the closest hotspot and to the closest large commuting zone are not excessively correlated, and the latter has an unreported small and statistically insignificant coefficient. In this specification, a 10% increase in distance to the closest hotspot reduces AI share growth by 2.6% of median growth.

We test in column 5 whether our results are affected by the remote commuting zones of Alaska and Hawaii, by dropping them from the column 3 specification: this increases the coefficient of interest to -0.022. On the other hand, using mean rather than median regression (column 6) makes the coefficient considerably more negative, at -0.049.

Panel B shows the corresponding coefficients from the twelve–year difference regres-

[17] The coefficient is -0.042 if the 2007–2014 difference is dropped.

sion: as in panel A, all coefficients on distance are statistically significantly negative. In columns 1 and 2, the coefficients are somewhat more than double the size of the panel A coefficients, corresponding to a time period that is almost twice as long. However, unlike in panel A, the addition to the covariates of changes in job advertisements, IT share in job advertisements, and AI publications as a share of job advertisements (column 3) considerably weakens the effect of distance, raising the coefficient from -0.089 (implying a 10% increase in distance reduces AI job growth by 4.1% of median growth) to -0.041 (implying only a 1.8% reduction). This may reflect the correction of a negative bias in the previous specifications.

Examining changes in job vacancies is somewhat similar to examining an acceleration in employment, and we therefore also run regressions with the change in cumulative AI share, an approximation to change in employment, as the outcome. The patterns are qualitatively similar: the coefficients on the log distance to the closest hotspot, presented in columns 1–3 of Table 3, are always statistically significantly negative. The coefficients are one half to one third the magnitude of those in Table 2, though not directly comparable, implying that a 10% increase in distance reduces growth in AI share by 2.3% of median growth (coefficient of -0.014 in column 2 panel A, median growth of 0.06) or 1.6% (coefficient of -0.018 in column 2 panel B, median growth of 0.11).

Our second identification approach is to define the key distance covariate as the radius of the circle enclosing a certain number of cumulative AI publications as of 2006. In columns 4–6 of Table 3, we present the coefficients on this variable from three specifications reflecting our preferred sets of covariates. The general patterns, including an always statistically significantly negative effect of distance, continue to be similar. The seven–year difference coefficients of -0.031—0.034 indicate that a 10% increase in the radius of the circle enclosing at least 1000 AI publications (about three–quarters of a standard deviation) reduces AI job advertisement growth by 0.0031–0.0034, or 3.3–3.7% of median growth.

In Appendix Table 5, we present results for the effect of distance to the closest AI hotspot on the (linear) probability of a commuting zone's having any AI job advertisement in 2018 if it had none in 2007. The magnitude of the preferred coefficients implies that a 10 percent increase in distance reduces the probability of having any AI job advertisement by 0.006–0.010 percentage point, or 0.8–1.3% of the 0.79 mean percentage point growth in Table 1. The effect of distance to a hotspot thus operates more strongly through the intensive rather than extensive margin.

4.2 Sensitivity to choice of AI publication threshold

Thus far, the analysis has used the apparently arbitrary hotspot and radius threshold of 1000 AI publications through 2006. We now turn to testing the sensitivity of the results to the threshold. We would expect that very low thresholds would lead to a finding of no effect of distance and indeed, this could be considered a falsification test. In the case of the distance to the closest hotspot approach we can use as the falsification test the coefficient on the distance to the closest commuting zone with at least zero AI publications by 2006 i.e. the distance to the closest commuting zone (without any other publication distance control). In the radius approach there is no equivalent to this, and the closest test is using the threshold of one publication. If there is a genuine effect of distance, it should emerge as the threshold is increased.

In Figure 9, we plot the point estimates and 95 percent confidence intervals for the coefficients on distance to hotspot with different thresholds, using the conservative specification from column 3 Table 2 for the left graphs, and the column 4 specification adding distance to the closest large commuting zone (appropriately adjusted for the hotspot threshold) in the right graphs. The upper two graphs are based on twelve–year differences, while the lower two are based on seven–year differences. Interesting patterns emerge, though it is worth noting that there are no statistically significant differences between the plotted coefficients.

In none of the four graphs is there an effect of distance when the AI publication threshold is zero, satisfying the falsification test, and in panels B–D there is no effect when the threshold is one publication. In panels B–D the effect becomes increasingly negative (and statistically significant) as the threshold rises to 300 publications, then remains similar until at least 1250 publications, while in panel A this pattern is less apparent, with a statistically significant negative effect somewhat similar for all positive thresholds. In panels B–D, the effect of distance unexpectedly seems to weaken as the threshold approaches 2000, reaching zero at high enough thresholds in the case of the right hand graphs. Standard errors are larger when the distance to the closest large commuting zone is controlled in the right two panels (C, D), but this control has little effect on the point estimates for thresholds in the range 300–1250.[18]

Patterns are similar in Figure 10, which contains the corresponding graphs for cumulative AI job advertisements, although here panel A is similar to the other panels in having the effect of distance become more negative as the hotspot publication threshold rises to 300. More so than in the previous figure, the weakening of the distance effect at high thresholds in panels B–D involves a jump when Austin, TX (1922 AI publications by 2006) no longer meets the threshold for an innovation hotspot, leaving no hotspots in the middle of the country. The sensitivity to Austin suggests that the small distance effects at high thresholds may not have an economic interpretation. However, a possible economic explanation is that technology diffuses immediately from the largest hotspots, perhaps because new developments in such places receives nationwide publicity.

The corresponding graphs using the radius of the circle enclosing a given number of AI publications are shown in Figure 11. Using this identification strategy, the effect of distance does not go to zero in any graph at high thresholds, in the threshold range we have plotted, though for seven–year differences (lower graphs) the locus of points has a U–shape as in most previous graphs.[19] In Figure 12, we plot the corresponding graphs for the hotspot identification strategy and the probability of having any AI job advertisement.

[18] When the threshold is 1750 publications, the correlation between the distance to the closest hotspot and the distance to the closest large commuting zone is 0.94 and the standard error is correspondingly large.

[19] The radii for thresholds in the range 10,000–40,000 have correlations with one another of at least 0.9 and begin to enclose significant proportions of the national cumulative AI publications by 2006 (about one third at 40,000), which is why we plot only to 10,000. The coefficient on distance does become less negative over the 10,000–40,000 range, however, eventually becoming positive in specifications controlling for distance to the closest large commuting zone.

With no control for distance to the closest large commuting zone (panel A), the effect of distance is negative and statistically significant for publication thresholds from one to 2000. With the addition of the control for distance to the closest large commuting zone (panel B), the coefficients become less negative and mostly lose their statistical significant. Compared to analysis of the AI share, it is less evident that the effect of distance increases with threshold at low thresholds, but we again see the sensitivity to Austin.

4.3 Distinguishing innovation and adoption

We now turn to assessing in more detail which sorts of job advertisements lie behind distance's constituting a barrier to the influence of AI innovation, to ascertain whether the AI jobs influenced by distance are in fact AI adoption or adaptation jobs rather than simply jobs that will lead to more AI innovation and publications.

4.3.1 The effect of distance by type of AI

One approach to distinguishing between innovation and adoption is to divide the specific AI skills required in the job advertisements into categories reflecting the distinction, rather than grouping them all together as in the analysis until now. We present the results for the effect of distance on the shares of different AI types in Table 4. Panel A's first row shows that the largest category of AI is the unspecified category ("Artificial Intelligence" or "Machine Learning" only), constituting 37% of AI ads (based on the micro–data), closely followed by the category of skills we have termed "R&D" (e.g. Computational Linguistics or Neural Networks) at 34%. Only small minorities of advertisements mentioning AI applications (19%) or tools (9.1%), with image processing accounting for 12%. The second row presents the share of AI advertisements in these categories seeking computer or mathematical workers, which could be a proxy for a job closer to innovation than adoption. The lowest share is in image processing (50%). The share is highest in the AI tool category (80%), compared to only 52% in the AI application category, suggesting the category distinctions are meaningful. However, the low share of AI advertisements requiring use of AI applications and the frequent requirement for computer and mathematical occupations even in this category suggests the use of AI by non–specialists is as yet limited.[20]

Panel B presents the results of seven–year difference median regressions, and panel C the results of twelve–year difference median regressions; column 1 reproduces the results of Table 2 column 3 for comparison purposes. The coefficients of most interest are those for AI applications in column 2: they are negative for both seven and twelve–year differences, and statistically significant for the twelve–year differences. The latter coefficient of -0.010 is large compared to the all–AI effect of -0.043 in column 1; a 10% increase in distance to the closest hotspot reduces median growth by 5.0%, compared to 1.9% for all AI. The point estimate of -0.0013 for seven–year differences is similar, when scaled by median growth, to the all–AI effect (-3.3% in both cases). The effect of distance is statistically significantly negative in all other regressions in the table except those for image processing, which may be inappropriately designated as AI.[21] The results are consistent with distance from innovation being a barrier to both innovation and adoption of AI.

4.3.2 The effect of distance by industry

The next step is to determine which industries are most affected by distance. We categorize industries using two–digit NAICS codes, forming groupings of various levels of aggregation. The most disaggregated groups are Information (NAICS 52), much of whose AI is likely to be innovation, and Finance and insurance (NAICS 53), which has rapid AI growth that can be categorized as adoption. The fast–growing AI of the group encompassing Real Estate (NAICS 53), Professional, Scientific and Technical Services (NAICS 54), Management of Companies and Enterprises (NAICS 55) and Administrative and Support and Waste Management and Remediation Services (NAICS 56), is likely to be a mixture

[20] Note that the computer and mathematical occupations category includes actuaries, who if they use AI tend to use AI applications.

[21] Median regressions do not converge in the case of image processing, so we present OLS results. Columns 2–6 do not sum to column 1 since non–linear regression is being used and because the categories are not mutually exclusive.

of innovation and adoption.

Accordingly, we present in Table 5 (seven–year differences) and Appendix Table 6 (twelve–year differences) the results of regressions based on samples of job advertisements in particular industries, with as dependent variable AI job advertisements in a particular industry as a share of that industry's job advertisements. Dividing AI job advertisements in this way leads to a median value of zero for the dependent variable for all industries except one and we therefore present OLS regressions. For those industries whose 75th percentile AI share growth is not zero, we also present quantile regression at the 75th percentile. For some industries there are no job advertisements in certain commuting zones in certain years, so the number of observations in the regressions varies slightly by industry.

The means of the dependent variables in columns 2 show that the fastest average growth in AI occurred in Finance and insurance (0.28 percentage point), job advertisements with missing industry (0.25 percentage point) and Real estate, professional services, management and administration (0.24 percentage point), ahead of Information (0.21 percentage point). Despite its fast growth (and high level), Information's AI growth at the 75th percentile is zero (column 4), indicating high geographic concentration (resulting in 18 commuting zone–years having no job advertisements in this industry).

The regression results in Table 5 suggest that distance from a hotspot is a barrier to both the adoption of AI and further innovation in AI. The OLS effect of distance (column 3), whether measured by the coefficient on distance or its size relative to mean growth, is most negative for job advertisements with missing industry, followed by the two sectors Information and Real estate, professional services and administration. The coefficient of -0.061 for Finance and insurance is more negative than for all industries pooled (-0.049), but the effect for Finance and insurance is less negative than for all industries relative to mean growth (-2.2% versus -3.5% in response to a 10% increase in distance). The coefficients are statistically significantly negative for all these sectors, but statistically insignificant for the other sectors (and small, except for the oppositely signed coefficient of 0.096 for Other services and public administration). The coefficients on distance for OLS and quantile regressions are broadly similar (when both are available), though they are measuring different things. Similar patterns emerge for twelve–year differences in Appendix Table 6 (though the coefficient for the Information sector is statistically insignificant).

4.4 The role of employment agencies and a possible mechanism

We next follow up on the finding that distance is the greatest barrier for the "sector" of missing industry. Thanks to our imputation of many industry values based on employer name, only six percent of job advertisements with missing industry have a valid employer name. The large effect of distance for job advertisements with missing industry is therefore likely to reflect a peculiarity of job advertisements with missing employer, and therefore posted by an employment agency. Since Burning Glass Technologies codes industry based primarily on employer name (Burning Glass Technologies 2019) and many employment agencies are not associated with a particular industry, many job advertisements placed by employment agencies do not have a valid industry. Altogether, job advertisements with a missing industry and a valid employer name represent 1% of all job advertisements, compared to 15% missing both employer name and industry and 19% with valid industry but missing employer name. The occupation and skill mixes of the latter two categories differ, probably reflecting the degree to which industry may be inferred based on this information.

In Table 6, we use median regression to compare the effect of distance based on all job advertisements (column 1, reproducing Table 2 column 3), based on job advertisements with valid (column 2) and missing (column 3, the median regression counterpart of Table 5 columns 2 and 5) industry, and based on job advertisements with valid (column 4) and missing (column 5) employer name. All regressions are at the commuting zone–year level, but based on different underlying job advertisement samples. Seven–year differences are presented in panel A and twelve–year differences in panel B.

As expected based on Table 5, the distance coefficient is much more negative when based on advertisements with missing than with valid industry (-0.094 compared to -0.012 in panel A, and -0.126 compared to -0.029 in panel B), though all effects are statistically significant. A large gap remains when the coefficients are scaled by median growth in AI share, which is larger for job advertisements with no employer name: the effect of a 10% increase in distance is -2.0 or -1.2% of median growth for the job advertisements with valid industry (column 2 panels A and B), and -6.5 or -5.5% for job advertisements with missing industry (column 3 panels A and B).

In column 4 panels A and B, we can see that the coefficients on distance are small and statistically insignificant for regressions based on job advertisements with a valid employer name, while in column 5 we see that the coefficients are much more negative in regressions based on job advertisements with a missing employer name. The effect of a 10% increase in distance is -1.1 or -1.4% of median growth for the job advertisements with valid employer name (column 4 panels A and B), and -6.5 or -5.5% for job advertisements with missing employer name (column 5 panels A and B).

Since the determinant of having a valid industry is how much information is available to Burning Glass Technologies, while the determinant of having a valid employer name is an economically meaningful firm type, we turn our focus to the presence or absence of an employer name in the job advertisement. In Table 7, we present results of linear probability regressions using the micro job advertisement data. We regress a dummy for the absence of an employer name (34% of the advertisements) on characteristics of the job advertisement and its commuting zone, whose means are given in column 1. Column 2 shows that a one percentage point higher share of college graduates in the commuting zone labor force reduces the probability of using an employment agency by a statistically significant 0.73 percentage point. It also shows that a vacancy in a commuting zone that is 10% farther from its nearest neighbor is a statistically insignificant 0.13 percentage point more likely to be posted by an employment agency, and that a vacancy in a commuting zone 10% farther from all other commuting zones is a statistically significant 0.6 percentage point more likely to be posted by an employment agency. Advertisement–level dummies for whether a (non–AI, non–Microsoft Office) IT skill is required and whether any of the types of AI skill is required have positive coefficients of at least 0.03, except for the dummy for image processing, which has a negative sign. These results point to employers turning to employment agencies when skilled workers (of the type typically advertised in Burning Glass) cannot be found close at hand, though the low R–squared of 0.10 indicates that many unobserved factors are important.

Column 2 coefficients also indicate that the greater the number of skills specified in the job advertisement, the less likely the job advertisement is posted by an employment agency (coefficient of -0.012): this seems unintuitive, but may simply reflect advertising practices of employment agencies. The commuting zone unemployment rate has the expected negative sign, but is very imprecisely estimated. We also check for a relation between using an employment agency and distance from the closest AI hotspot, to understand any selection into the subsamples used above, but find none.

We also control for the pre–2007 cumulative number of AI publications (divided by 1000): the coefficient indicates that an additional 1000 publications increase the probability by 0.4 percentage point that advertisements in that commuting zone are posted by an employment agency, a small reponse to such a large change in publications. The final set of covariates are dummies for occupation, including missing occupation (3.7% of advertisements); we report only a subset of the occupation coefficients. With the exception of legal occupations (coefficient not reported), jobs in the omitted category of computer and mathematical occupations are most likely to be advertised by an employment agency. For example, advertisements for architects and engineers are 3.7 percentage points less likely to be posted by an employment agency. Results for the sample of finance and insurance jobs (means in column 3, regressions results in column 4) and for the sample of jobs requiring AI application skills (columns 5 and 6) are qualitatively similar.

The commuting zone–year level regressions provide further support for the hypothesis that employers turn to employment agencies when they cannot hire locally. Although not shown in Table 6, the determinants other than the role of distance to a hotspot also differ between the samples based on valid versus missing employer name. We show this in Table 8 for seven–year differences (and in Appendix Table 7 for twelve–year differences).

The first two columns of Table 8 represent the same regressions as in columns 4 and 5 of Table 6, but show the effects of AI publications in addition to the effect of distance to a hotspot. The point estimate of an increase in AI publications is twice as high in the valid–name regression (0.380, column 1), which has lower growth in the AI share, as in the missing–name regression (0.187,column 2); both are statistically significant. The reported coefficients are multiplied by 1000, so the coefficient of 0.380 implies that if the number of AI publications rises by 1000, the AI share in job advertisements rises by 0.38 percentage point.

More importantly (since the covariates are more clearly exogenous), the covariates capturing the number of pre–2007 AI publications have jointly statistically significant coefficients only in the valid–name regression (column 1). The effect of an additional pre–2007 AI publication is positive and statistically significant in these regressions for moving from no publications to one publication, at 100 publications, and at 1000 publications (reported in the third panel). An increase from no publications to one publication, given that the coefficient of 0.093 is multiplied by 1000, increases the AI share in job advertisements by 0.00009 percent point. By contrast, the point estimates are negative and statistically insignificant for regressions based on missing–name micro samples in column 2.

Thus, employers hiring through employment agencies are little affected by innovation in their own commuting zone but suffer from being distant from innovation hotspots. Combined with the results above, this suggests that firms hiring through employment agencies are seeking skills not closely connected to skills required for the innovation taking place in their commuting zone. Firms hiring directly are strongly affected by innovation in their own commuting zone but unaffected by distance from hotspots: this could mean either that they do not need to hire outside their commuting zone, or that they can easily attract workers regardless of their distance (perhaps because the firms are themselves in a hotspot).

If this mechanism is correct, the effect of distance should decline with the amount of innovation prior to 2007. We investigate this hypothesis for job samples of valid and missing employer names in columns 3 and 4, by interacting log distance to the closest hotspot with the three covariates for AI publications prior to 2007. The key numbers are in the bottom panel, where we evaluate the effect of distance at various levels of pre–2007 AI publications. For the valid employer name sample (column 3), the effect of distance is small and statistically insignificant at all levels of pre–2007 innovation, while for the missing employer name sample (column 4), the initially statistically significantly negative effect of distance increases towards zero with the level of pre–2007 innovation. For a commuting zone with 1000 AI publications prior to 2007, the coefficient on distance is a statistically insignificant -0.022, compared to about -0.06 at low numbers of publications close to the mean and median numbers.[22]

It should be possible to detect the mechanism in a regression based on the full sample of job advertisements. We use a covariate measuring the share of employer names missing in a commuting zone in the first year of the pair of years being differenced and control for its main effect, and its interaction with the log distance to the closest hotspot, expecting to find that commuting zones with a higher share of missing employer names have a more negative effect of distance. The unreported point estimates support this: the coefficient on the main distance effect goes to zero, and the coefficient on the interaction term is negative (-0.048). However, the interaction coefficient is not close to statistically significant. [23]

The mechanism also implies that in samples of job advertisements reflecting adoption, the effect of pre–2007 AI publications should be smaller than in the full sample, which also reflects innovation. Column 5 confirms there is no effect of these publications in a sample of finance and insurance advertisements. On the other hand, in the column 6 sample of advertisements requiring AI application skills, the pre–2007 AI publication coefficients are jointly statistically significant and the calculated effects at various levels are positive, though statistically significiant only at the 10% level (the coefficients are small, but the

[22] We do not include the interaction of the *change* in AI publications and log distance to the closest hotspot because it is so strongly correlated with the main effect of the change in pre–2007 AI publications.

[23] The regressions of Table 8 are repeated for twelve–year differences in Appendix Table 7, yielding similar patterns except for this interaction result. This may be due to the necessary reliance on the 2007 share of job advertisements missing a employer name: the share is much higher in that year than in later years, possibly reflecting either a different selection of job advertisements into online job boards at this time, or a different treatment of the data by Burning Glass Technologies.

growth in AI share is also small, as shown in Table 4).

Although the evidence is not conclusive, we summarize our mechanism as follows. Certain commuting zones are geographically isolated, have a poorly educated population, innovate little or are specialized in a narrow type of innovation. Firms (or prospective firms) in these commuting zones must cast a wide net geographically (with the help of employment agencies) to find the specialized workers needed to adopt AI, which is not yet available as off–the–shelf software. These specialized workers are particularly numerous in innovation hotspots, but cannot be attracted to move long distances to non–hotspots with unskilled labor forces. Reluctance to move long distances therefore slows the geographic diffusion of AI and presumably other new technologies.

4.4.1 Firms operating in multiple locations and firm size

Previous papers have shown that firms operating in multiple locations speed the transfer of technology. We examine this hypothesis in our context by creating a variable measuring the number of 2007 job advertisements in a commuting zone placed by firms which also post in the closest AI hotspot in 2007. It is irrelevant for our purposes that certain types of firms such as supermarkets have locations spread across commuting zones including AI hotspots. Therefore, we base our counts on job advertisements in computer and mathematical occupations: such advertisements account for 63% of AI advertisements (see Appendix Table 2). If the mechanism through which distance deters AI adoption is as a barrier to information about technology, then when more such ties exist, the effect of distance to the closest hotspot will be smaller. If the mechanism is related to the physical mobility of AI specialists, the effect of distance could be smaller if the difficulty is identifying appropriate workers at a distance, rather than attracting them.

In unreported results, we find no role for this interaction term. This is in part because the main effect of the count and the interaction term are highly correlated. However, the results are not surprising given that the counts are necessarily based on job advertisements with a valid employer name, and our distance effect is found to operate principally among advertisements with no employer name.

We have also investigated whether the effect of distance varies by firm size, something we again must measure using job advertisements with a employer name. We count the number of job advertisements by each firm advertising in a given year, and create thresholds adjusted for the total number of job advertisements in the year. Unreported results show no pattern by firm size.

5 Conclusion

Our results indicate that online vacancies for jobs requiring Artificial Intelligence (AI) skills grow more slowly in U.S. locations farther from AI innovation hotspots. A 10% greater distance from a hotspot (about a standard deviation) reduces a commuting zone's growth in AI jobs' share of job advertisements by 2-3% of median growth. Analysis by industry shows distance is a barrier for sectors likely to be adapting and adopting AI, like finance and insurance, and not merely for sectors likely to be innovating in AI, like Information. Further, analysis by type of AI skill required shows distance from innovation is a barrier to growth in job advertisements requiring the use of AI applications, and not merely job advertisements for workers innovating in AI. Distance from innovation thus slows growth in adoption.

Although the evidence is not conclusive, we propose the following mechanism in the light of these and other results. Certain commuting zones are geographically isolated, have a poorly educated population, innovate little or are specialized in a narrow type of innovation. Firms (or prospective firms) in these commuting zones must cast a wide net geographically (with the help of employment agencies) to find the specialized workers needed to adopt AI, which is not yet available as off–the–shelf software. These specialized workers are particularly numerous in innovation hotspots, but cannot be attracted to move long distances to non–hotspots with unskilled labor forces. Reluctance to move long distances therefore slows the geographic diffusion of AI and presumably other new technologies.

References

Aghion, Philippe, Benjamin F. Jones and Charles I. Jones. 2017. "Artificial Intelligence and Economic Growth". NBER Working Paper 23928.

Acemoglu, Daron, David Autor, Jonathon Hazell and Pascual Restrepo. 2021. "AI and Jobs: Evidence from Online Vacancies". NBER Working Paper 28257.

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. 2021. "International Trade and Innovation". NBER Working Paper 29611.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila and Bledi Taska. 2019. "The Demand for AI Skills in the Labor Market". University of Navarra working paper.

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.

Bahar, Dany, Prithwiraj Choudhury and Hillel Rapoport. 2020. "Migrant Inventors and the Technological Advantage of Nations". IZA Discussion Paper 12994.

Bernard, Andrew B, Andreas Moxnes and Yukiko U. Saito. 2020. "The Geography of Knowledge Production: Connecting Islands and Ideas". Dartmouth College working paper.

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 Workforc, Technical Report.

Comin, Diego, Mikhail Dmitriev and Esteban Rossi–Hansberg. 2012. "The Spatial Diffusion of Technology". NBER Working Paper 18534.

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.

Dittmar, Jeremiah and Ralf Meisenzahl. 2022. "The Research University, Invention and Industry: Evidence from Germany History". CEPR Discussion Paper 17,383.

Ganguli, Ina, Jeffrey Lin and Nicholas Reynolds. 2019. "The Paper Trail of Knowledge Spillovers: Evidence from Patent Interferences". Federal Reserve Bank of Philadelphia Working Paper 17–44.

Gibbons, Stephen and Henry G. Overman. 2012. 'Mostly Pointless Spatial Econometrics?" *Journal of Regional Science*, 52(2): 172–191.

Goldfarb, Avi, Bledi Taska and Florenta Teodoridis. 2019. "Could Machine Learning Be a General–Purpose Technology? Evidence from Online Job Postings". University of Toronto working paper.

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.

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. 2008. "Ethnic Scientific Communities and International Technology Diffusion". *Review of Economics and Statistics*, 90(3): 518–537.

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*.

Parente, Paulo and João Santos Silva. 2016. "Quantile Regression with Clustered Data". *Journal of Econometric Methods*, 5(1):1–15.

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.

Methodological Appendix: Specifications not used

We considered and rejected several ways of capturing the effect of AI publications in other commuting zones.

A.1 Controlling for neighbor commuting zone publications

We could have made a definition of a neighboring commuting zone and controlled for the average or total number of pre–2007 AI publications in neighboring commuting zones. We considered this limited by the need to select neighbors and by the assumption of no effects of non–neighbors.

A.2 AI Publications weighted by distance from commuting zone

An ostensibly more appealing approach is to control for a weighted average of AI publications (*AI Pubs*) in all other commuting zones, with the weights a function of distance $d_{j \neq c}$ from commuting zone c. Common functions used in the spatial econometrics literature are the reciprocal of distance or the exponential of negative distance, leading to specifications such as:

$$\Delta^k AI_c^s = \delta_0 + \delta_1 AI\ Pubs_c + \delta_2 \sum_{j \neq c} \frac{AI\ Pubs_j}{d_j^{\rho}} + \Delta\nu_c,$$

or

$$\Delta^k AI_c^s = \phi_0 + \phi_1 AI\ Pubs_c + \phi_2 \sum_{j \neq c} AI\ Pubs_j e^{-\rho d_j} + \Delta\eta_c.$$

As written, the spillover coefficients (δ_2 and ϕ_2) depend on the units chosen for distance, and while the weights can be normalized to fix this in the reciprocal specification, this cannot be done in the exponential specification. Both specifications require testing robustness to a parameter (ρ), though this drawback is shared with our preferred approach. However, it is not possible to test the hypothesis that $\rho = 0$, implying that spillovers exist but are independent of distance, since in this case the spillovers are not identified (the two terms in *AI Pubs* sum to a constant, the total *AI Pubs*). In our preferred approach this case is econometrically identified, though not distinguishable from the no–spillover case.

The biggest drawback of this approach arises due to the need to control for the effect of population in addition to the effect of AI publications: although commuting zone population and AI publications are only moderately highly correlated, once they are weighted by a function of distance there is almost perfect collinearity.

A.3 Controlling for distances to more than one AI hotspot

It would be desirable to be able to judge from a single specification how hotspots defined with different thresholds affect AI job advertisements. For example, the covariates could be distance to the closest commuting zone with 500–999 AI publications and the distance to the closest commuting zone with more than 1000 AI publications, and their interaction (since the effects are unlikely to be additive). This might give an idea of whether 500–999 publications constitute as influential a hotspot as more than 1000 publications (though the distance at which to evaluate the partial effects is not obvious), but getting a precise idea would be difficult as the specification would include many main highly correlated main effects along with the interaction terms.

A.5 Interacting distance to closest hotspot with AI publications in hotspot

This approach is a hybrid of thinking there is a genuine threshold above which a commuting zones causes spillovers and thinking the actual threshold is unknown. We prefer to vary the AI publication threshold for a hotspot.

A.4 Defining relative rather than absolute hotspots

It is possible that spillovers from location A to location B depends on the relative number of AI publications rather than A's absolute number of AI publications. Using the ratio of publications raises the obvious problem of locations with zero publications, however. Furthermore, it seems unlikely that spillovers from A to B would be the same in the case where A has 2000 publications and B 1000, and in the case where A has two publications and B one. This makes introducing an absolute threshold tempting, yet any sizeable threshold yields a hotspot measure highly correlated with a purely absolute hotspot measure. Using the difference between the publications between A and B does not seem intuitive.

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

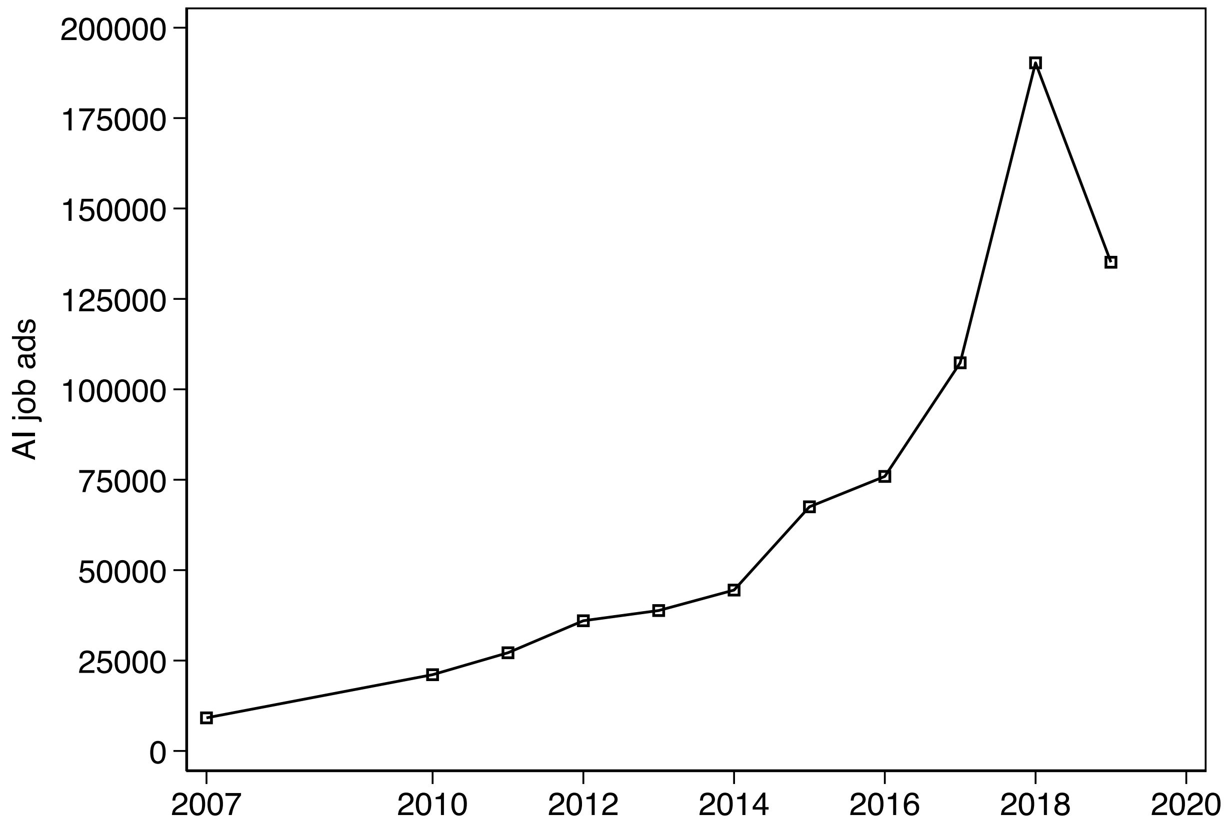

Notes: Data for 2019 are for January–July. Data for 2008 and 2009 are not available.
Source: Burning Glass Technologies.

Figure 2: AI share of job ads (%)

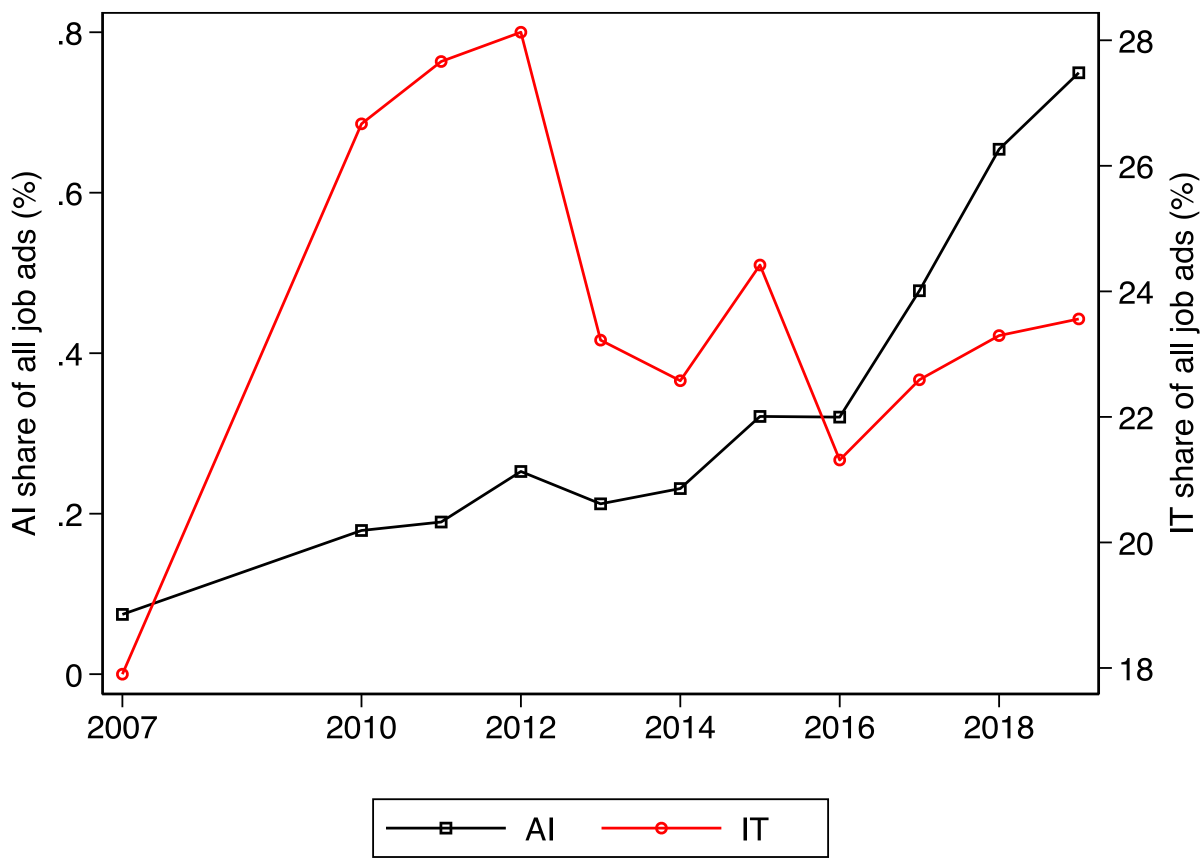

Source: Burning Glass Technologies.

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

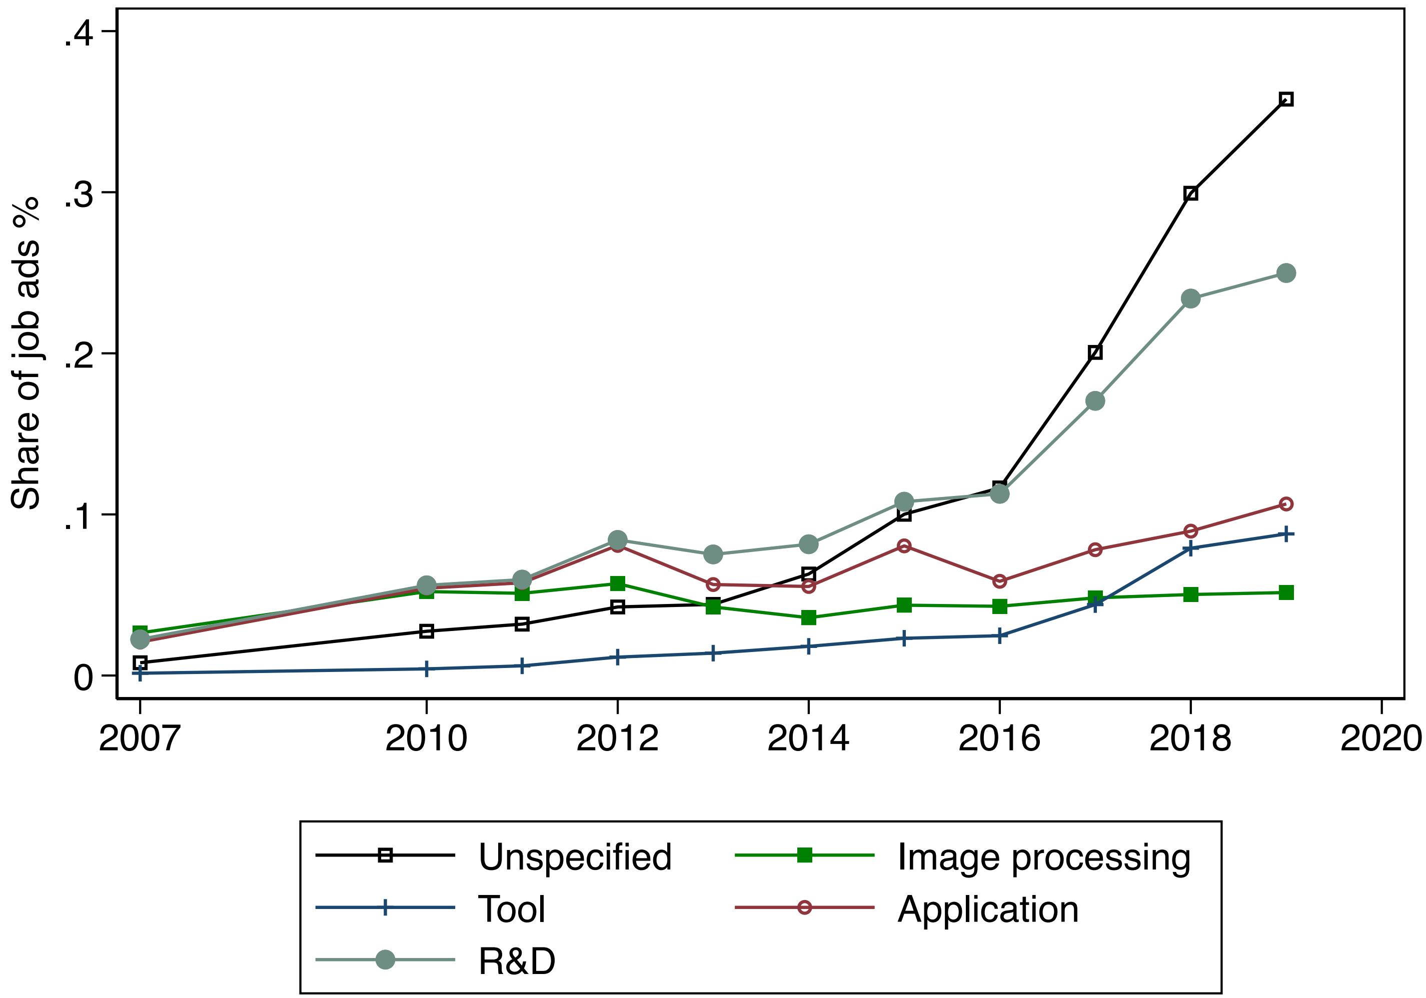

Note: Unspecified AI job advertisements require "Artificial Intelligence" and/or "Machine Learning" skills with no further detail given. The other categories are not mutually exclusive.

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

(a) 2007

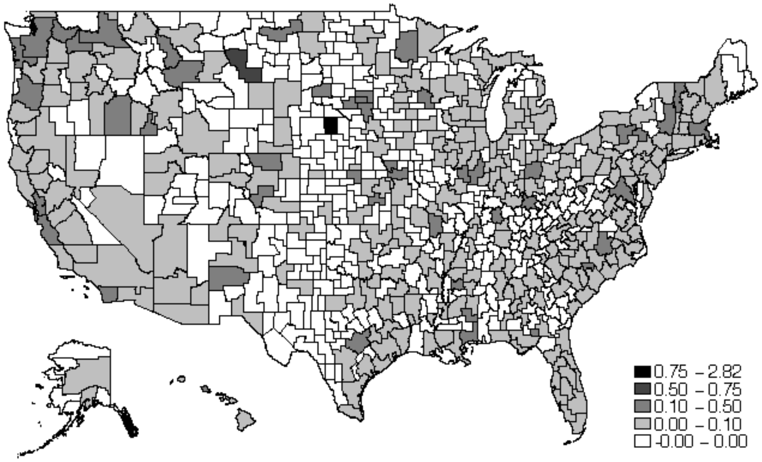

(b) 2010

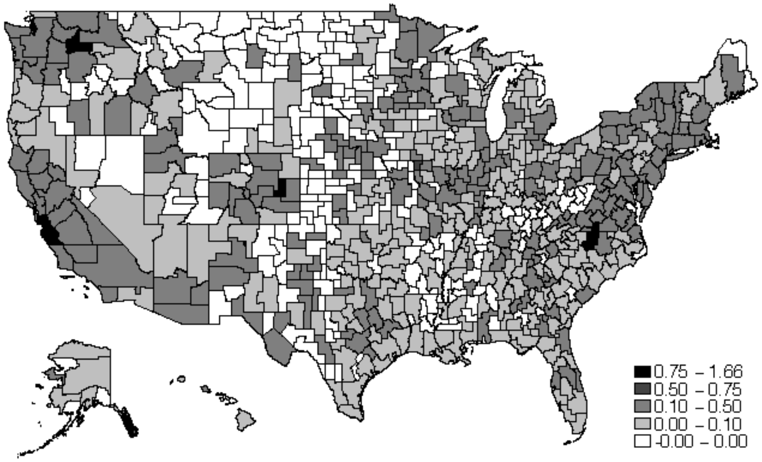

(c) 2014

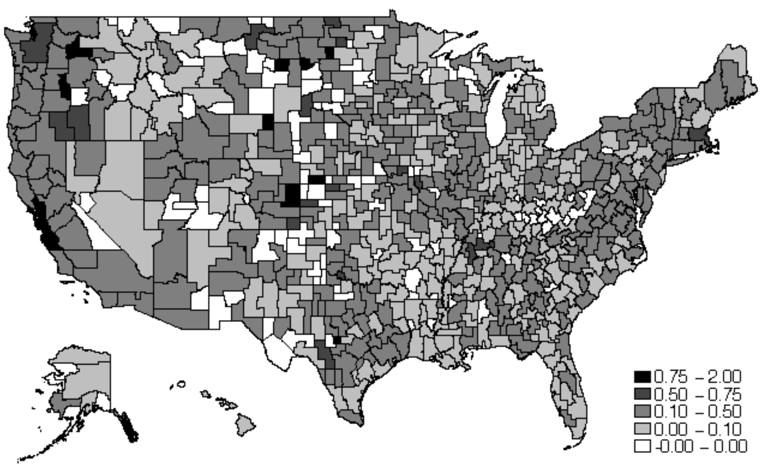

(d) 2018

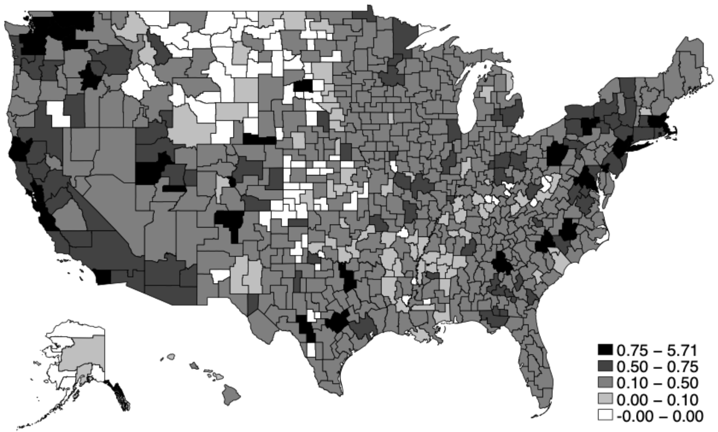

Figure 5: AI publications 1950-2019

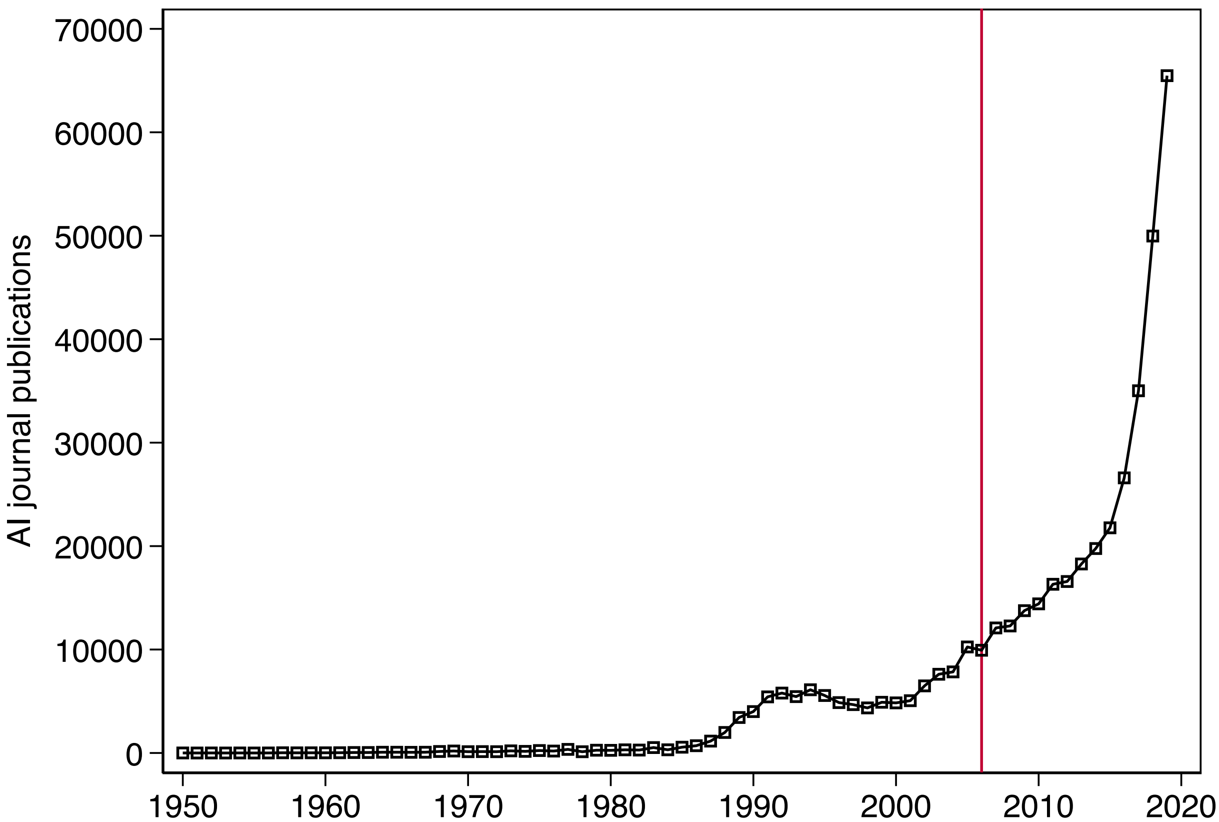

Source: Authors' dataset.

Figure 6: Innovation hotspots' AI publications through 2006

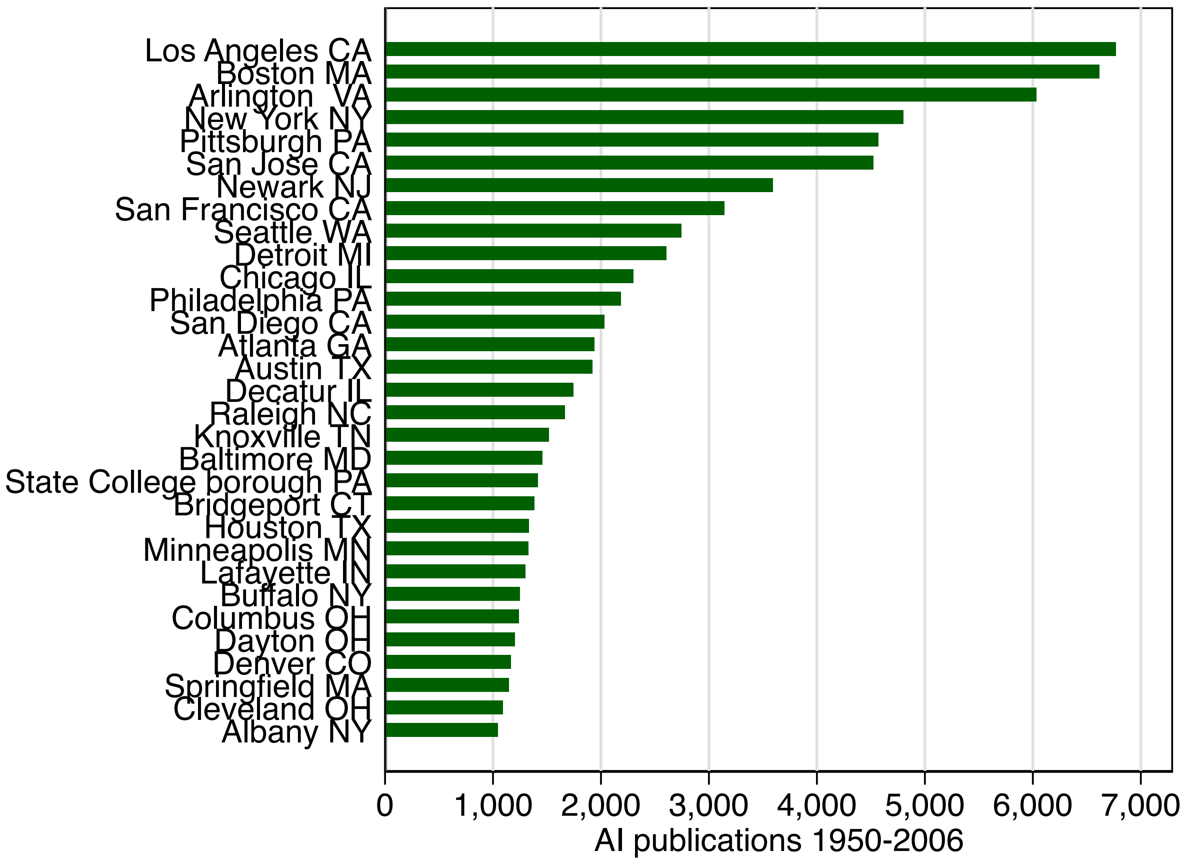

Note: The definition of a hotspot here is a commuting zone with at least 1000 AI publications through 2006.

Figure 7: Commuting zones' AI publications through 2006

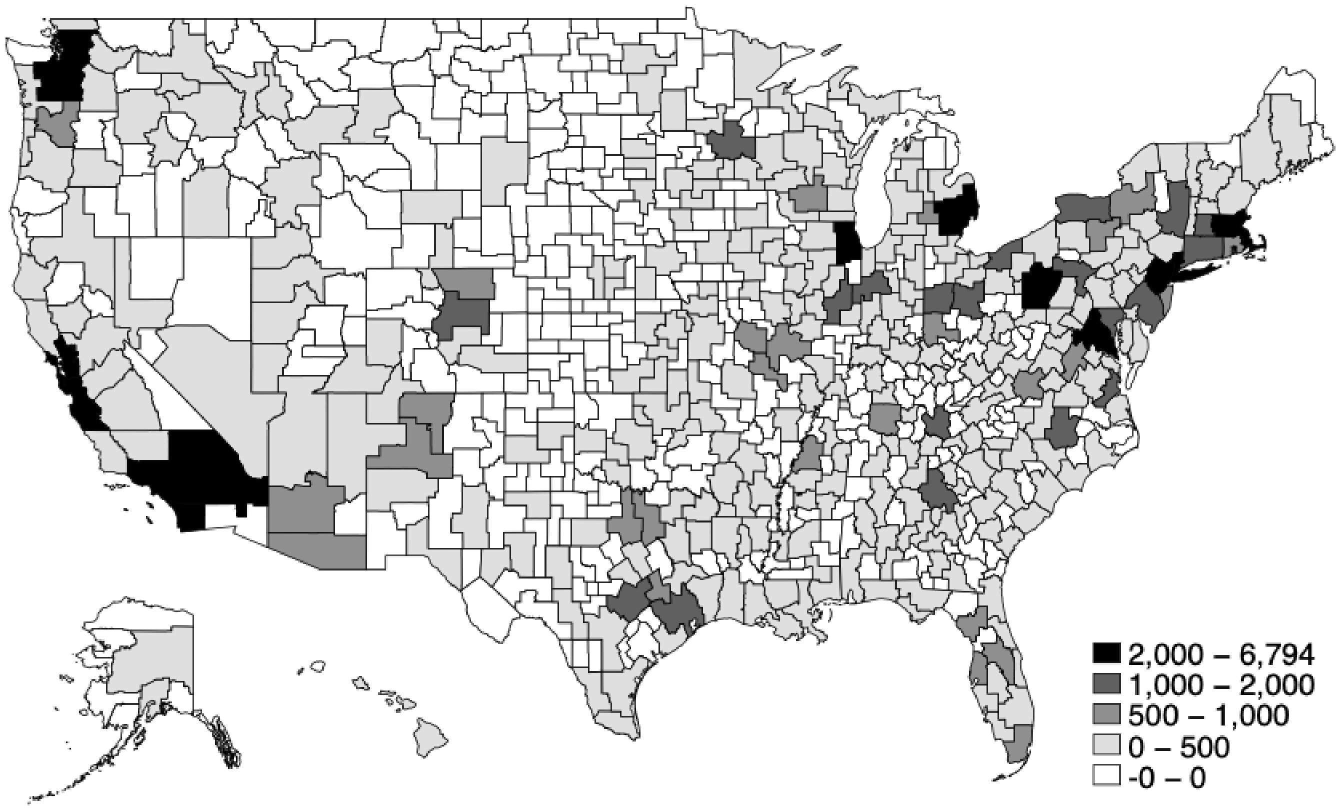

Figure 8: Commuting zones' AI publications in given year

(a) 2007: 240 CZs with any publication

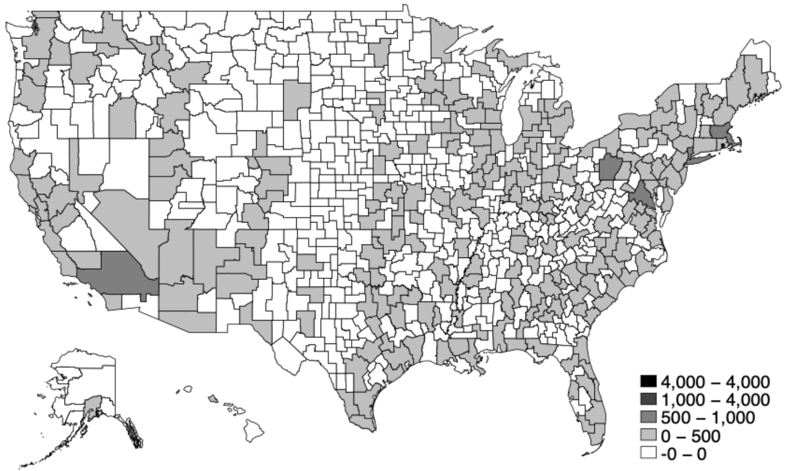

(b) 2010: 245 CZs with any publication

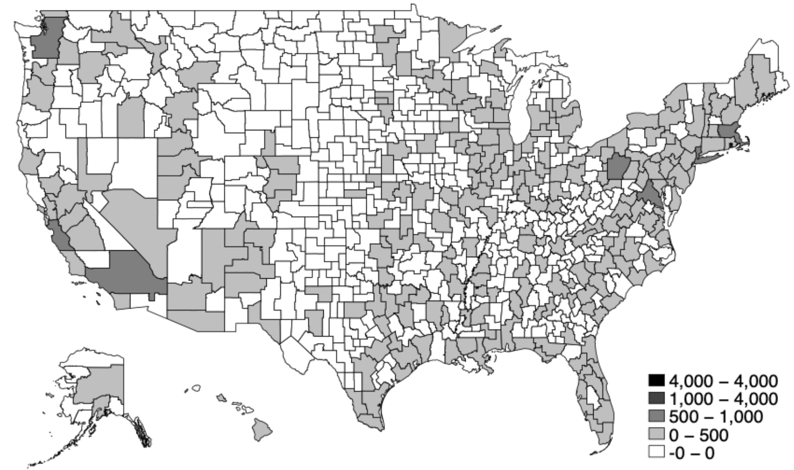

(c) 2014: 252 CZs with any publication

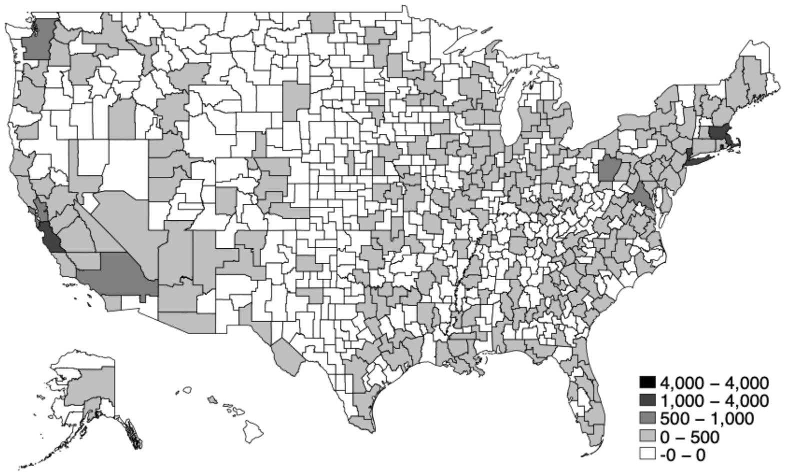

(d) 2018: 282 CZs with any publication

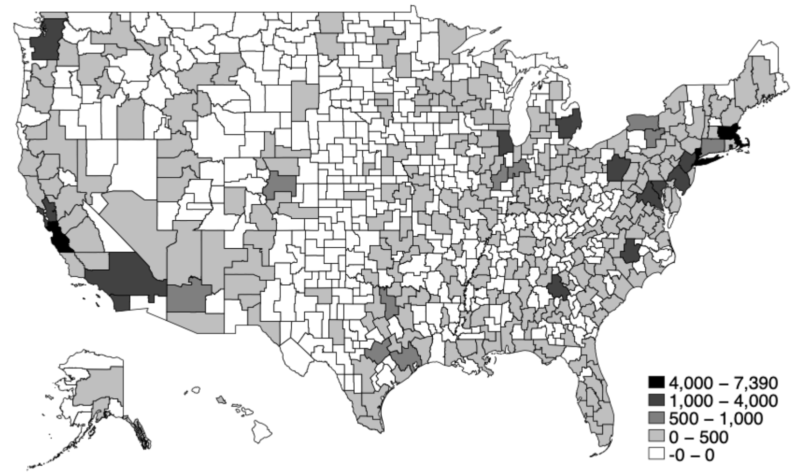

Figure 9: Coefficients on distance to closest hotspot

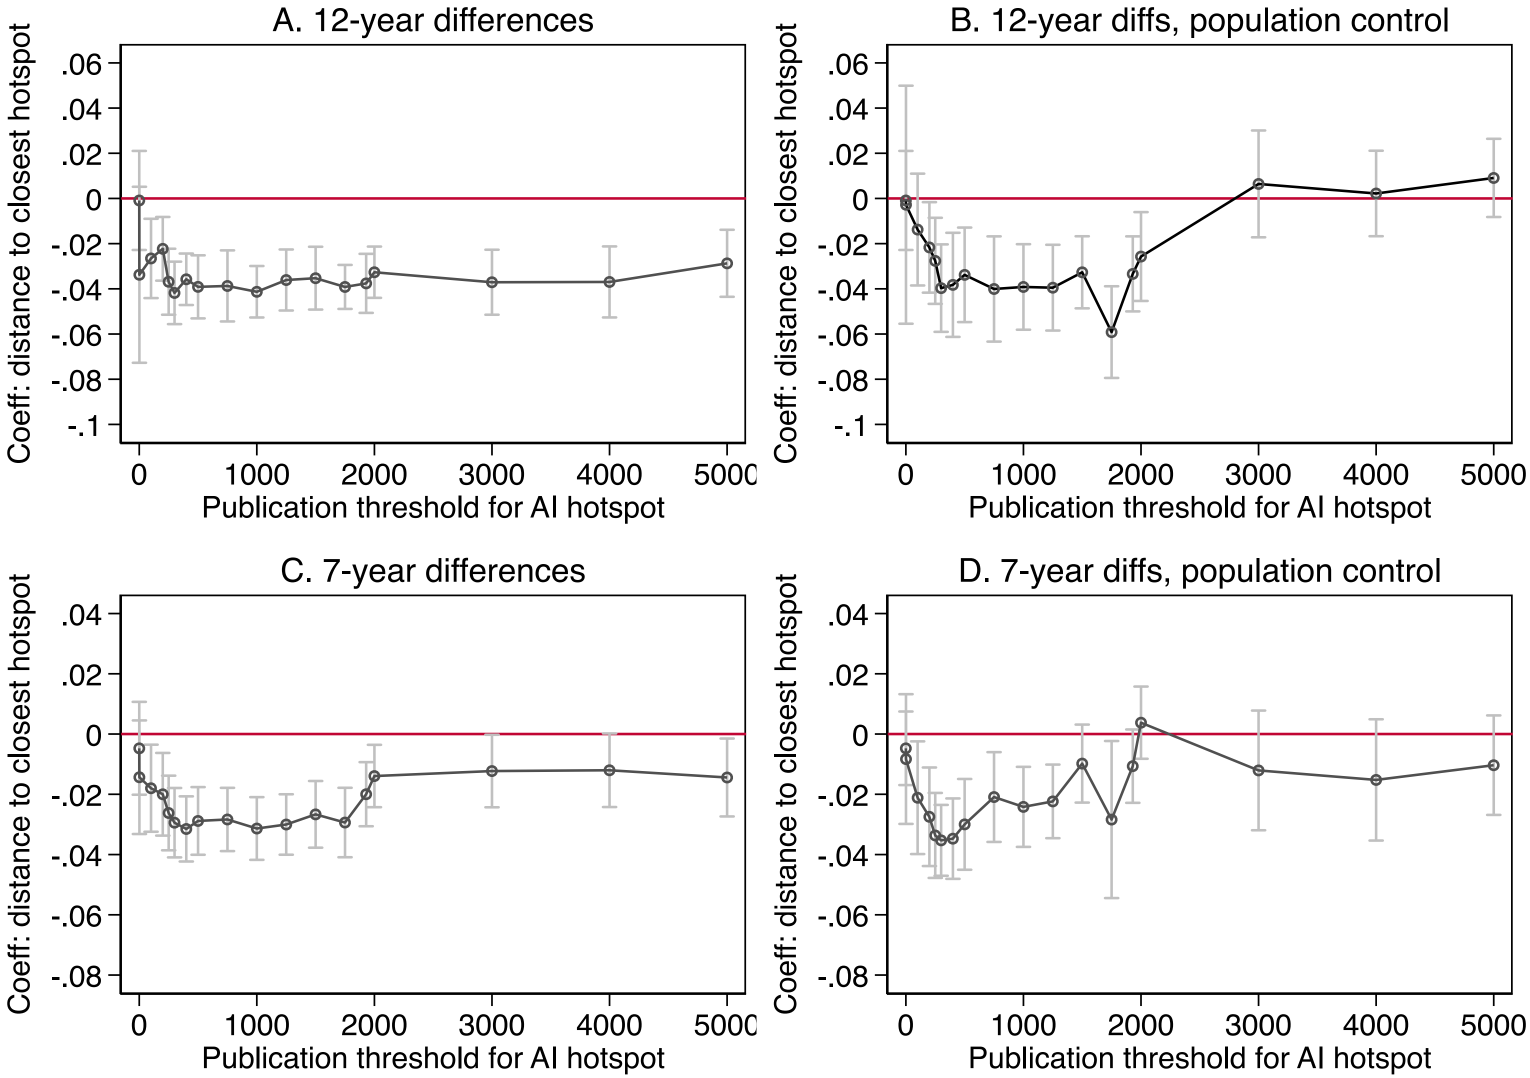

Note: The dependent variable is the change in AI jobs' share in total job advertisements. Population control refers to the distance to the closest large commuting zone.

Figure 10: Coefficients on distance to closest hotspot: Cumulative AI job advertisements

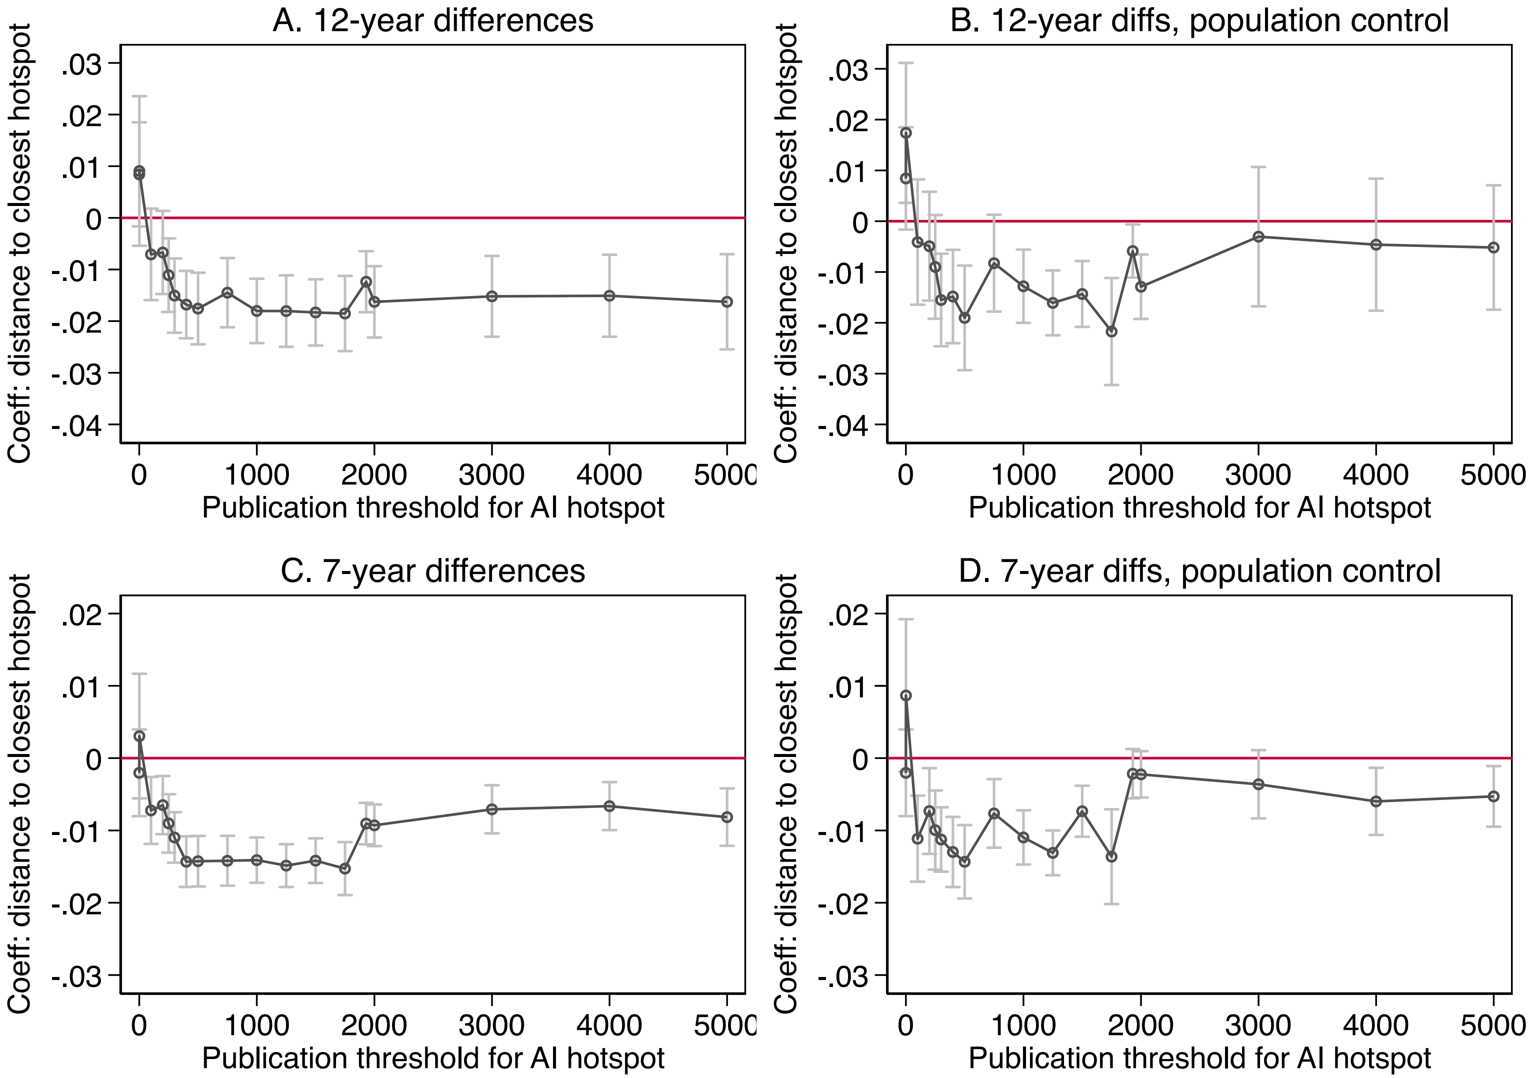

Note: The dependent variable is the change in cumulative AI jobs' share in total job advertisements. Population control refers to the distance to the closest large commuting zone.

Figure 11: Coefficients on log radius including threshold number of publications

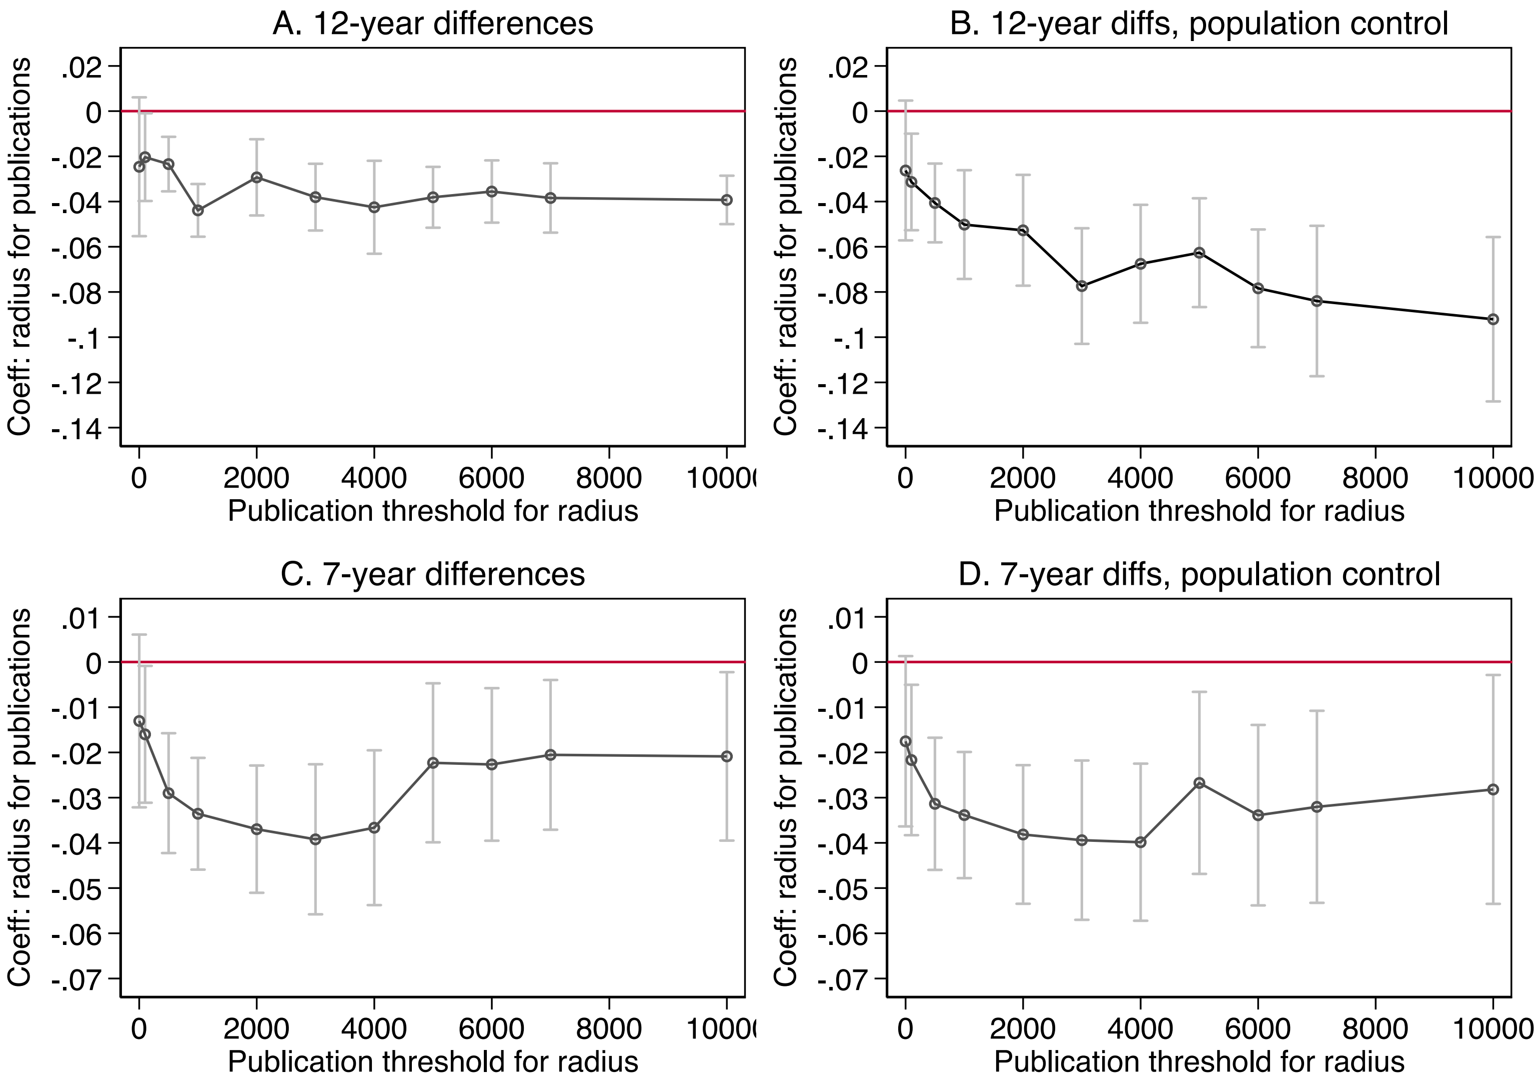

Note: The dependent variable is the change in AI jobs' share in total job advertisements. The coefficients plotted are on the log of the radius around the commuting zone which includes the threshold number of publications. Population control refers to the log population included within the radius.

Figure 12: Coefficient on distance to closest hotspot: extensive margin

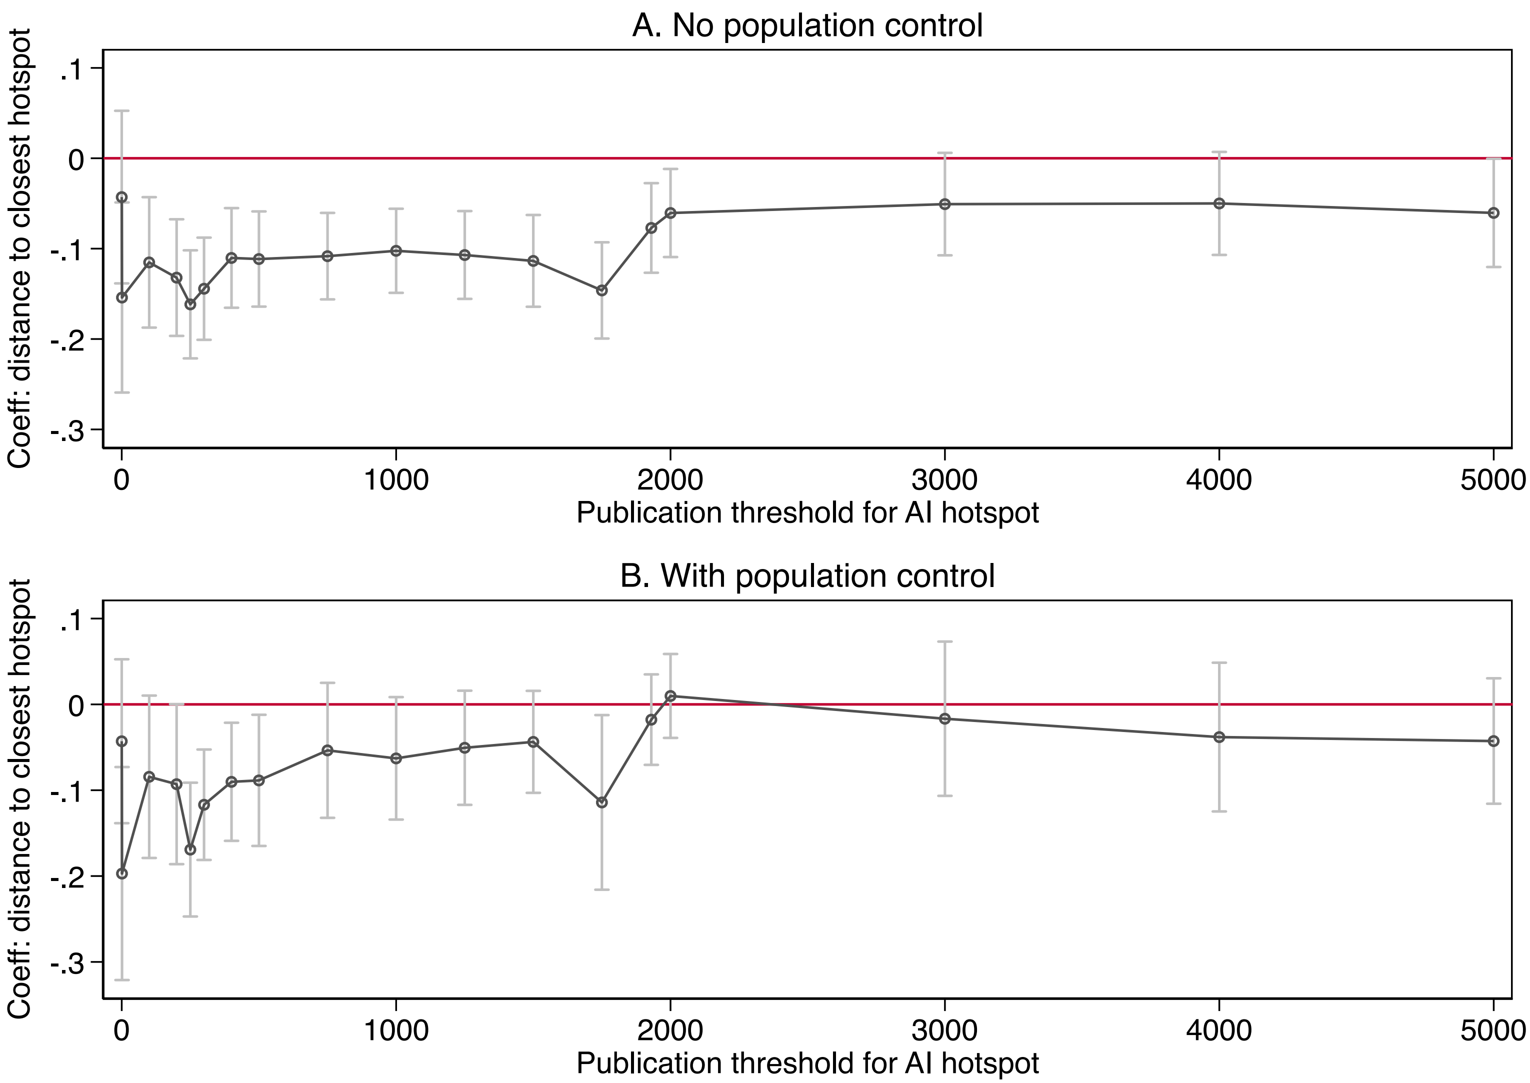

Note: The dependent variable is the probability of having any AI job advertisement in 2018 conditional on having none in 2007. Population control refers to the distance to the closest large commuting zone.

Table 1: Summary statistics

	Mean	Median	Min	Max	Obs
A. Δ AI job advertisement share (%)					
2007-2019 **Δ=12**	0.279	0.222	-2.82	3.34	741
2007-2019 **Δ=7**	0.136	0.093	-2.82	4.90	2964
B. Δ Any AI job advertisement					
2007-2018 **Δ=11**	0.791	1	0	1	401
C. Distances (km)					
To closest hotspot (1000+ AI pubs)	412	323	40	3946	741
Radius of circle with 1000+ AI pubs	324	228	40	3946	741
To closest large CZ	372	286	8.75	3946	741
To other CZs (average)	1630	1451	1144	6385	741
To closest CZ	76.5	67.7	7.2	540	741
D. Initial conditions covariates					
Any AI publication prior to 2007	0.49	0	0	1	741
AI publications prior to 2007	156	0	0	6769	741
Job advertisements 2007	16,583	2570	3	696,205	741
Population 2000 in thousands	380	104	1.19	16,393	741
IT share 2007 (%)	9.01	7.77	0	42.86	741
E. Differenced covariates **Δ=7**					
AI publications	37.4	0	-48	6543	2964
Log job advertisements	0.49	0.54	-2.75	3.26	2964
IT job ad share (%)	-1.24	-2.11	-28.53	26.09	2964
F. Differenced covariates **Δ=12**					
AI publications	72.0	0	-15	6899	741
Log job advertisements	0.34	0.42	-3.85	3.08	741
IT job ad share (%)	7.02	6.99	-27.07	28.66	741

Notes: The definition of an AI publication hotspot in the table is a commuting zone (CZ) with at least 1000 AI publications by 2006 (31 CZs); the distance to the closest large CZ is the distance to the closest of the most populous 31 CZs. The location of a CZ is based on the locations of job advertisements, so the distance between adjacent CZs is positive.

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

	Median regression					OLS
	All commuting zones				No AK/HI	All
	(1)	(2)	(3)	(4)	(5)	(6)
A. Seven-year differences						
Log distance to closest hotspot	-0.029***	-0.031***	-0.031***	-0.024***	-0.022***	-0.049***
(1000+ AI publications)	(0.005)	(0.006)	(0.005)	(0.007)	(0.005)	(0.008)
Observations	2,964	2,964	2,964	2,964	2,888	2,964
R-squared/Pseudo-R-squared	0.13	0.15	0.21	0.21	0.22	0.23
B. Twelve-year difference						
Log distance to closest hotspot	-0.074***	-0.089***	-0.041***	-0.039***	-0.045***	-0.071***
(1000+ AI publications)	(0.010)	(0.012)	(0.006)	(0.009)	(0.006)	(0.014)
Observations	741	741	741	741	722	741
R-squared/Pseudo-R-squared	0.17	0.22	0.31	0.31	0.32	0.39
AI publications through 2006 (any, level, square)	Yes	Yes	Yes	Yes	Yes	Yes
Log job ads 2007; log population 2000; IT share in advertisements 2007; Log average distance to other CZs; log distance to closest CZ	--	Yes	Yes	Yes	Yes	Yes
Change in log ads, IT share, log AI pubs	--	--	Yes	Yes	Yes	Yes
Log distance to closest large CZ	--	--	--	Yes	--	--

Notes: The dependent variable is the seven-year difference (panel A) or twelve-year difference (panel B) in AI jobs' share of all job advertisements; the share measured in %. Panel A regressions include year dummies and are based on 2014-2007, 2017-2010, 2018-2011, 2019-2012. Panel B is based on 2019-2007. "Change" refers to seven-year differences in panel A and twelve-year difference in panel B. The definition of an AI publication hotspot is a commuting zone (CZ) with at least 1000 AI publications by 2006 (31 CZs); the distance to the closest large CZ is the distance to the closest of the 32 most populous CZs. Standard errors clustered by CZ in parentheses (panel A) or robust (panel B).
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3: Effect of distance to closest hotspot on change in AI jobs' share calculated cumulatively; effect of radius of circle enclosing 1000 AI publications

	AI share calculated cumulatively			AI share calculated contemporaneously		
	(1)	(2)	(3)	(4)	(5)	(6)
A. 7-year differences						
Log distance to closest hotspot (1000+ AI pubs)	-0.014***	-0.014***	-0.011***	--	--	--
	(0.002)	(0.002)	(0.002)			
Log radius of circle enclosing 1000+ AI pubs	--	--	--	-0.031***	-0.034***	-0.034***
				(0.007)	(0.006)	(0.007)
Observations	2964					
Pseudo R-squared	0.10	0.12	0.13	0.14	0.21	0.21
B. 12-year differences						
Log distance to closest hotspot (1000+ AI pubs)	-0.024***	-0.018***	-0.013***	--	--	--
	(0.004)	(0.003)	(0.004)			
Log radius of circle enclosing 1000+ AI pubs	--	--	--	-0.081***	-0.043***	-0.050***
				(0.016)	(0.010)	(0.012)
Observations	741					
Pseudo R-squared	0.26	0.29	0.29	0.21	0.31	0.31
Initial conditions covariates	Yes	Yes	Yes	Yes	Yes	Yes
Log av. distance other CZs, log distance closest CZ	Yes	Yes	Yes	Yes	Yes	Yes
Log AI pubs in circle	--	--	--	Yes	Yes	Yes
Change in log ads, IT share, log AI pubs, log pop	--	Yes	Yes	--	Yes	Yes
Log distance closest large CZ	--	--	Yes	--	--	Yes
Log population within circle	--	--	--	--	--	Yes

Note: Median regressions. The dependent variable in columns 1-3 is the seven-year (panel A) or twelve-year (panel B) change in the cumulative number of AI job advertisements since 2007 divided by the cumulative number of all job advertisements since 2007, multiplied by 100. The dependent variable is the seven-year (panel A) or twelve-year difference (panel B) in AI jobs' share of all job advertisements; the share measured in %. Panel A regressions include year dummies and are based on 2014-2007, 2017-2010, 2018-2011, 2019-2012. Panel B is based on 2019-2007. Initial conditions covariates are AI publications through 2006 (a dummy for any, the number and its square), log job advertisements 2007, log population 2000, IT share in advertisements 2007. "Change" refers to seven-year differences in panel A and twelve-year difference in panel B. The definition of an AI publication hotspot is a commuting zone (CZ) with at least 1000 AI publications by 2006 (31 CZs); the distance to the closest large CZ is the distance to the closest of the most populous 31 CZs. Standard errors in parentheses; clustered by commuting zone in panel A, robust in panel B.

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

Table 4: Impact of distance from AI hotspot on change in AI job advertisement share by AI type

	Median regression					OLS
	All AI	AI Application	Unspecified AI only	AI Tool	AI R&D	Image Processing
	(1)	(2)	(3)	(4)	(5)	(6)
A. AI job ads with valid occupation						
Share AI type in AI job ads (714,348 obs)	100%	19.2%	37.1%	9.1%	34.4%	12.5%
Share computer scientist/mathematician	62.6%	52.2%	68.3%	80.2%	66.5%	49.9%
[obs]	[714,348]	[136,995]	[264,852]	[65,356]	[245,868]	[88,970]
B. 7-year differences, median regression (2964 obs)						
Log distance to closest hotspot	-0.031***	-0.0013	-0.017***	-0.0025***	-0.0069***	0.0030*
(1000+ AI publications)	(0.005)	(0.0010)	(0.002)	(0.0004)	(0.0016)	(0.0017)
R-squared	0.21	0.03	0.24	0.15	0.22	0.01
Median of dependent variable (pct points)	0.093	0.004	0.033	0.000	0.017	0.000
Effect of 10% greater distance as % of median	-3.3%	-3.3%	-5.2%	--	-4.1%	--
C. 12-year diffs, median regression (741 obs)						
Log distance to closest hotspot	-0.043***	-0.010***	-0.025***	-0.0030*	-0.012***	-0.0013
(1000+ AI publications)	(0.007)	(0.002)	(0.003)	(0.0016)	(0.003)	(0.0033)
R-squared	0.31	0.14	0.34	0.23	0.28	0.01
Median of dependent variable (pct points)	0.222	0.020	0.090	0.015	0.056	0.000
Effect of 10% greater distance as % of median	-1.9%	-5.0%	-2.8%	-2.0%	-2.1%	--

Notes: Each column's dependent variable is the share of that type of AI job advertisement in all job advertisements (in %). An AI job advertisement with unspecified AI requires "Artificial Intelligence" or "Machine Learning" skills but no more specific AI skills. The types of AI are not mutually exclusive (except for unspecified AI versus other types). Controls are those of Table 2 column 3. Standard errors in parentheses; clustered by commuting zone in panel A, robust in panel B.
*** p<0.01, ** p<0.05, * p<0.1

Table 5: Impact of distance from AI hotspot on change on AI job advertisement share by industry, seven-year differences

	Obs (1)	Mean of dependent variable (2)	OLS (3)	75^{th} percentile of dependent variable (4)	Quantile regression, 75^{th} percentile (5)
All	2964	0.14	-0.049^{***} (0.007)	0.21	-0.041^{***} (0.008)
Agriculture, Utilities, Mining, Construction, Manufacturing	2962	0.08	0.002 (0.014)	0.11	-0.007^{*} (0.004)
Wholesale trade, Retail trade, Warehousing, transportation	2964	0.02	0.004 (0.008)	0	--
Information	2946	0.21	-0.079^{*} (0.042)	0	--
Finance, Insurance	2961	0.28	-0.061^{**} (0.027)	0.35	-0.052^{**} (0.020)
Real Estate, Professional and scientific services, Administration	2964	0.24	-0.090^{**} (0.030)	0.36	-0.076^{***} (0.018)
Education, Health	2964	0.05	-0.000 (0.013)	0.09	-0.008^{*} (0.005)
Arts and recreation, Accommodation	2961	0.05	-0.001 (0.020)	0	--
Other services, Public administration	2962	0.17	0.096 (0.110)	0.07	--
Missing industry	2964	0.25	-0.136^{***} (0.013)	0.39	-0.125^{***} (0.016)

Notes: Each cell in columns 3, and 5 contains the coefficient on log distance to an AI publication hotspot (at least 1000 publications) from a different regression with full covariates (those of Table 2 column 3). The dependent variable is the change in the share of advertisements in the specified industry requiring AI. The NAICS 2 codes for each row are a) 11, 21-23, 31-33; b) 42, 44-45, 48-49; c) 51; d) 52; e) 53-56; f) 61-62; g) 71-72; h) 81, 92. Standard errors in parentheses, clustered by commuting zone.
*** p<0.01, ** p<0.05, * p<0.1

Table 6: Impact of distance to closest AI publications hotspot on change in AI job advertisement share in different job advertisement samples

Job advertisements in underlying micro sample:	All (1)	Valid industry (2)	Missing industry (3)	Valid employer (4)	Missing employer (5)
A. 7-year differences					
Log distance to closest	-0.031***	-0.012***	-0.094***	-0.005	-0.067***
hotspot (1000+ AI publications)	(0.005)	(0.004)	(0.011)	(0.003)	(0.008)
Observations	2,964	2,963	2,962	2,961	2,964
R-squared	0.21	0.15	0.17	0.0	0.16
Median dependent var (percentage points)	0.093	0.060	0.144	0.047	0.103
Effect of 10% distance as % of median growth	-3.3%	-2.0%	-6.5%	-1.1%	-6.5%
B. 12-year differences					
Log distance to closest	-0.041***	-0.029***	-0.126***	-0.018*	-0.131***
hotspot (1000+ AI publications)	(0.007)	(0.010)	(0.013)	(0.011)	(0.014)
Observations	741	740	739	739	741
R-squared	0.31	0.23	0.31	0.22	0.26
Median dependent var (percentage points)	0.222	0.154	0.249	0.137	0.240
Effect of 10% distance % of median growth	-1.8%	-1.2%	-5.1%	-1.4%	-5.5%

Notes: Median regressions. Each column's dependent variable is the change in AI share based on different underlying samples of job advertisements. Missing industry refers to missing NAICS 2. Controls are those of Table 2 column 3. Standard errors in parentheses, clustered by commuting zone in panel A, robust in panel B.

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

Table 7: Job advertisement-level determinants of missing employer name (using employment agency)

	All industries, skills		Finance, insurance		AI application skills	
	Mean	OLS	Mean	OLS	Mean	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Commuting zone characteristics						
College educated share of labor force 2000	0.261	-0.727***	0.269	-0.666***	0.284	-1.596***
	(0.068)	(0.067)	(0.062)	(0.082)	(0.071)	(0.233)
Log distance to closest commuting zone	4.2	0.013*	4.2	0.018	4.2	0.051*
	(0.5)	(0.008)	(0.5)	(0.012)	(0.4)	(0.029)
Log average distance to other commuting zones	7.5	0.064***	7.5	0.048*	7.5	0.051
	(0.3)	(0.015)	(0.3)	(0.019)	(0.3)	(0.054)
Log distance to closest hotspot (1000+ AI publications)	5.3	-0.012**	5.2	-0.006	5.1	-0.008
	(0.9)	(0.005)	(0.9)	(0.007)	(0.9)	(0.022)
Log distance to closest large commuting zone	5.2	0.005	5.1	-0.001	5.0	-0.010
	(0.8)	(0.005)	(0.9)	(0.008)	(0.8)	(0.023)
Unemployment rate	0.056	-0.040	0.056	0.051	0.054	-0.801
	(0.023)	(0.209)	(0.022)	(0.257)	(0.022)	(0.531)
AI publications pre-2007 (/1000)	1.7	0.004**	1.9	0.002	2.5	0.005
	(2.1)	(0.002)	(2.1)	(0.003)	(2.3)	(0.009)
Job characteristics						
Number of skills required	8.4	-0.0122***	10.8	-0.0061***	19.0	-0.0019***
	(7.5)	(0.0003)	(7.5)	(0.0003)	(11.7)	(0.0004)
Management	0.121	-0.130***	0.147	-0.011**	0.099	-0.157***
		(0.008)		(0.005)		(0.016)
Business and finance	0.069	-0.051***	0.207	0.095***	0.044	-0092***
		(0.007)		(0.007)		(0.019)
Computer science and math	0.119	--	0.137	--	0.499	--
Architecture and engineering	0.033	-0.037***	0.009	0.010*	0.082	0.002
		(0.008)		(0.006)		(0.017)
IT skill required	0.236	0.053***	0.303	0.051***	0	--
		(0.005)		(0.005)		
AI skills required (x 100)						
AI application	0.070	0.030**	0.078	0.043***	1	--
		(0.012)		(0.011)		
AI unspecified only	0.137	0.044*	0.217	0.041**	0	--
		(0.023)		(0.012)		
AI tool	0.034	0.103***	0.045	0.037***	0.076	--
		(0.020)		(0.009)		
AI R&D	0.127	0.037**	0.225	0.040***	0.225	--
		(0.018)		(0.008)		
Image processing	0.046	-0.034**	0.025	0.017	0.025	--
		(0.013)		(0.014)		
Other occupation dummies	--	Yes	--	Yes	--	Yes
R-squared	--	0.10	--	0.07	--	0.11
Observations	204,469,684		15,592,591		143,250	

Notes: Dependent variable is a dummy for whether the job advertisement is missing the employer name (mean is 0.34 for all industries, 0.19 for finance and insurance and 0.26 for job advertisement requiring AI application skills). Covariates include year dummies. There are 24 occupational categories including missing occupation (3.7% of occupations). Standard errors in parentheses, clustered by commuting zone.
p<0.01, ** p<0.05, * p<0.1

Table 8: Impacts of distance to closest AI publications hotspot and commuting zone's own publications on change in AI job advertisement share, seven-year differences

Sub-sample of micro sample:	Employer name is:				Industry:	Skills:
	Valid	Missing	Valid	Missing	Finance, insurance	AI application
	(1)	(2)	(3)	(4)	(5)	(6)
Log distance to closest hotspot	-0.005	-0.067***	-0.003	-0.087***	-0.061**	-0.0013
(1000+ AI publications)	(0.003)	(0.008)	(0.003)	(0.011)	(0.027)	(0.0010)
Change in AI publications (x1000)	0.380***	0.187**	0.378***	0.173	0.299***	0.024**
	(0.006)	(0.082)	(0.012)	(0.158)	(0.045)	(0.009)
Pre-2007 AI publication controls	Yes	Yes	Yes	Yes	Yes	Yes
(p-value joint significance)	(0.00)	(0.10)	(0.00)	(0.00)	(0.60)	(0.00)
Interactions of pre-2007 AI controls with	--	--	Yes	Yes	--	--
log distance (p-value joint significance)			(0.01)	(0.01)		
Observations	2961	2964	2961	2964	2961	2964
Pseudo R-squared	0.06	0.16	0.06	0.17	0.09	0.03
Effect of pre-2007 AI publications (x1000) evaluated at						
Change 0 to 1 AI publication	0.093***	-0.044**	--	--	-0.049	0.007*
	(0.020)	(0.021)			(0.080)	(0.004)
100 AI publications	0.098***	-0.033	--	--	0.016	0.008*
	(0.022)	(0.020)			(0.065)	(0.004)
1000 AI publications	0.065***	-0.033	--	--	0.006	0.005
	(0.015)	(0.016)			(0.047)	(0.004)
Effect of log distance evaluated at						
1 AI publication pre-2007	--	--	-0.009	-0.063**	--	--
			(0.006)	(0.013)		
100 AI publications pre-2007	--	--	-0.007	-0.058***	--	--
			(0.006)	(0.012)		
500 AI publications pre-2007	--	--	-0.001	-0.040***	--	--
			(0.013)	(0.010)		
1000 AI publications pre-2007	--	--	0.002	-0.021	--	--
			(0.022)	(0.016)		

Notes: Median regressions, except OLS in column 5. Each column's dependent variable is the change in AI share based on different underlying samples of job advertisements. Controls include those of Table 2 column 3, including a dummy for any pre-2007 AI publication, the number of pre-2007 AI publications and its square. The coefficients for pre-2007 AI publications are multiplied by 1000. Standard errors clustered by commuting zone 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. Tools
ANTLR, Automatic Speech Recognition (ASR), Caffe Deep Learning Framework, Deeplearning4j, Google Cloud Machine Learning Platform, H2O (software), Ithink, Keras, Libsvm, MLPACK (C++ library), MXNet, Madlib, Mahout, Microsoft Cognitive Tookit, Mlpy, ND4J (software), Natural Language Toolkit (NLTK), OpenCV, OpenNLP, Pybrain, TensorFlow, Torch (Machine Learning), Vowpal, Wabbit, Xgboost

D. Applications
AI ChatBot, Chatbot, IBM Watson, IPSoft Amelia, Lexalytics, Machine Translation (MT), Machine Vision, MoSes, Object Recognition, Recommender Systems, Sentiment Analysis / Opinion Mining, Sentiment Classification, Speech Recognition, Text Mining, Text to Speech (TTS), Virtual Agents, Word2Vec

E. R&D
Computational Linguistics, Computer Vision, Decision Trees, Deep Learning, Gradient boosting, Image Recognition, Latent Dirichlet Allocation, Latent Semantic Analysis, Lexical Acquisition, Lexical Semantics, Natural Language Processing, Nearest Neighbor Algorithm, Neural Networks, Object Tracking, Pattern Recognition, Random Forests, Semantic Driven Subtractive Clustering, Semi-Supervised Learning, Supervised Learning (Machine Learning), Support Vector Machines (SVM), Tokenization, Unsupervised Learning

Note: In analysis by AI type, unspecified is mutually exclusive of the other categories. Image processing, tools, applications and R&D are not mutually exclusive of one another.

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

Industry	Share (%) (1)	AI required? (%) (2)	Sample of ads requiring AI: Occupation (%) Computer and math (3)	Management (4)	Architects engineers (5)	Business finance (6)
All	100.0	0.37	62.6	10.6	6.4	4.8
Agriculture, Utilities, Mining, Construction, Manufacturing	9.0	0.41	59.8	9.5	16.3	3.0
Wholesale trade, Retail trade, Warehousing, transportation	12.3	0.16	65.6	12.3	4.9	4.4
Information	3.0	1.10	70.1	13.1	5.4	3.5
Finance, Insurance	7.6	0.54	59.4	15.5	2.7	12.1
Real Estate, Professional, technical and scientific services, Administration	17.9	0.68	67.4	9.9	6.3	5.2
Education, Health	22.7	0.16	29.7	9.5	2.9	1.9
Arts and recreation, Accommodation	6.9	0.11	56.4	11.9	2.8	3.4
Other services, Public administration	4.5	0.22	50.7	16.4	7.6	3.4
Missing industry	16.0	0.38	74.2	6.1	6.9	3.0

Notes: 2007-2019. 204,553,172 observations in columns 1-3; 714,348 observations in columns 3-6 (means for occupations are calculated based on advertisements requiring AI and with a valid occupation only). The NAICS 2 codes for each row are a) 11, 21-23, 31-33; b) 42, 44-45, 48-49; c) 51; d) 52; e) 53-56; f) 61-62; g) 71-72; h) 81, 92.

Appendix Table 3: Determinants of change in AI jobs' share in advertisements, seven-year differences

	Median regression					Mean
	All commuting zones				No AK/HI	All
	(1)	(2)	(3)	(4)	(5)	(6)
Log distance to closest hotspot	-0.029***	-0.031***	-0.031***	-0.024***	-0.022***	-0.049***
(1000+ pubs)	(0.005)	(0.006)	(0.005)	(0.007)	(0.005)	(0.008)
Log distance to closest large CZ	--	--	--	-0.012*	--	--
				(0.007)		
Any AI publication through 2006	0.028***	-0.006	-0.004	-0.005	-0.000	-0.011
	(0.007)	(0.009)	(0.008)	(0.008)	(0.008)	(0.012)
AI publications through 2006/1000	0.189***	0.136***	0.050*	0.050*	0.037	0.095***
	(0.026)	(0.024)	(0.028)	(0.027)	(0.024)	(0.027)
AI publications through 2006/1000^2	-0.019***	-0.014***	-0.010***	-0.010***	-0.009**	-0.020***
	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)
Log job advertisements 2007	--	-0.002	-0.001	0.001	-0.000	-0.007
		(0.004)	(0.005)	(0.005)	(0.005)	(0.007)
Log population 2000	--	0.012***	0.008	0.008	0.006	-0.002
		(0.004)	(0.005)	(0.005)	(0.005)	(0.008)
IT share in advertisements 2007 (%)	--	0.003***	0.005***	0.005***	0.005***	0.006***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Log distance to other CZs (average)	--	0.056***	0.054***	0.062***	0.103***	0.101***
		(0.017)	(0.016)	(0.016)	(0.017)	(0.023)
Log distance to closest CZ	--	0.009	0.003	0.003	0.003	0.001
		(0.008)	(0.008)	(0.008)	(0.008)	(0.009)
Change in log advertisements (7-year)	--	--	0.012**	0.012**	0.013**	0.008
			(0.006)	(0.006)	(0.006)	(0.012)
Change in IT share (%) x 1000 (7-year)	--	--	0.004***	0.004***	0.004***	0.008***
			(0.001)	(0.001)	(0.001)	(0.001)
Change in AI publications (7-year)	--	--	0.334***	0.338***	0.332***	0.353***
			(0.050)	(0.050)	(0.048)	(0.037)
Observations	2,964	2,964	2,964	2,964	2,888	2,964
R-squared/Pseudo-R-squared	0.13	0.15	0.21	0.21	0.22	0.23

Notes: The dependent variable is the seven-year difference in AI jobs' share of all job advertisements; the share measured in %. Differences included are 2014-2007, 2017-2010, 2018-2011, 2019-2012. All regressions also include year dummies. The definition of an AI publication hotspot is a commuting zone (CZ) with at least 1000 AI publications by 2006 (31 CZs); the distance to the closest large CZ is the distance to the closest of the 31 most populous CZs. Standard errors in parentheses, clustered by commuting zone. *** p<0.01, ** p<0.05, * p<0.1

Appendix Table 4: Determinants of change in AI jobs' share in advertisements, twelve-year differences

	Median regression					Mean
	All commuting zones				No AK/HI	All
	(1)	(2)	(3)	(4)	(5)	(6)
Log distance to closest hotspot	-0.074***	-0.089***	-0.041***	-0.039***	-0.045***	-0.071***
(1000+ AI publications)	(0.010)	(0.012)	(0.006)	(0.009)	(0.006)	(0.014)
Log distance to closest large CZ	--	--	--	-0.003	--	--
				(0.007)		
Any AI publication through 2006	0.069***	0.004	-0.006	-0.007	0.001	-0.021
	(0.016)	(0.018)	(0.013)	(0.014)	(0.014)	(0.025)
AI publications through 2006/1000	0.287***	0.158***	0.053	0.052	0.048	0.059
	(0.043)	(0.042)	(0.039)	(0.042)	(0.037)	(0.042)
AI publications through 2006/1000^2	-0.022***	-0.007	-0.021***	-0.020***	-0.019**	-0.021***
	(0.007)	(0.008)	(0.009)	(0.008)	(0.008)	(0.007)
Log job advertisements 2007	--	0.010	-0.019	-0.018	-0.011	-0.016
		(0.010)	(0.015)	(0.014)	(0.015)	(0.025)
Log population 2000	--	0.020	0.029*	0.028*	0.021	0.018
		(0.012)	(0.017)	(0.016)	(0.017)	(0.027)
IT share in job advertisements 2007	--	0.005*	0.023***	0.023***	0.022***	0.021***
		(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Log distance to other CZs (average)	--	0.166***	0.030	0.033	0.012	0.116**
		(0.035)	(0.022)	(0.022)	(0.031)	(0.053)
Log distance to closest CZ	--	-0.004	0.014	0.014	0.005	0.010
		(0.018)	(0.011)	(0.011)	(0.012)	(0.017)
Change in log job advertisements 2007-2019	--	--	-0.018	-0.016	-0.015	-0.020
			(0.015)	(0.015)	(0.015)	(0.026)
Change in IT share x 1000, 2007-2019	--	--	0.026***	0.025***	0.026***	0.023***
			(0.002)	(0.002)	(0.002)	(0.003)
Change in AI publications, 2007-2019	--	--	0.292***	0.287***	0.280**	0.312***
			(0.098)	(0.098)	(0.095)	(0.039)
Observations	741	741	741	741	722	741
R-squared/Pseudo-R-squared	0.17	0.22	0.31	0.31	0.32	0.39

Notes: The dependent variable is the twelve-year difference (2007-2019) in AI jobs' share of all job advertisements; the share measured in %. The definition of an AI publication hotspot is a commuting zone (CZ) with at least 1000 AI publications by 2006 (31 CZs); the distance to the closest large CZ is the distance to the closest of the 32 most populous CZs. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix Table 5: Determinants of having any AI job advertisement in 2018 ads if none in 2007

	All commuting zones				No AK/HI
	(1)	(2)	(3)	(4)	(5)
Log distance to closest hotspot	-0.194***	-0.096***	-0.102***	-0.063*	-0.085***
(1000+ AI publications)	(0.024)	(0.024)	(0.023)	(0.036)	(0.024)
Log distance to closest large CZ	--	--	--	-0.062	--
				(0.047)	
Any AI publication through 2006	0.170***	0.016	-0.005	-0.008	-0.008
	(0.031)	(0.035)	(0.035)	(0.036)	(0.036)
AI publications through 2006/1000	0.632**	0.269	0.457	0.485*	0.336
	(0.298)	(0.330)	(0.302)	(0.293)	(0.309)
AI publications through 2006/1000^2	-0.818	-0.457	-1.163	-1.359*	-0.856
	(0.532)	(0.570)	(0.718)	(0.717)	(0.722)
Log job advertisements 2007	--	0.067**	0.298***	0.304***	0.292***
		(0.030)	(0.062)	(0.061)	(0.063)
Log population 2000	--	0.151***	-0.049	-0.057	-0.041
		(0.031)	(0.056)	(0.056)	(0.057)
IT advertisements 2007/1000	--	-0.797***	-0.721***	-0.724***	-0.697***
		(0.218)	(0.219)	(0.223)	(0.217)
IT advertisements 2007^2	--	35.3**	37.0**	36.8**	36.4**
		(14.9)	(15.0)	(15.2)	(15.0)
Log distance to other CZs (average)	--	-0.072	-0.019	0.021	0.087
		(0.082)	(0.081)	(0.088)	(0.102)
Log distance to closest CZ	--	0.038	0.007	0.007	0.017
		(0.050)	(0.049)	(0.048)	(0.051)
Change in log job advertisements, 2007-2018	--	--	0.308***	0.314***	0.297***
			(0.070)	(0.070)	(0.073)
Change in IT advertisements x 1000, 2007-2018	--	--	-0.069***	-0.068***	-0.070***
			(0.022)	(0.023)	(0.023)
Change in AI publications, 2007-2018	--	--	0.992	1.113	0.953
			(0.749)	(0.773)	(0.678)
Observations	401	401	401	401	389
R-squared	0.18	0.38	0.41	0.42	0.38

Notes: Estimation is with linear probability. Robust standard errors in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Appendix Table 6: Impact of distance from AI hotspot on change in AI job advertisement share by industry, twelve-year differences

	Obs (1)	Mean of dependent variable (2)	OLS (3)	75^{th} percentile of dependent variable (4)	Quantile regression, 75^{th} percentile (5)
All	741	0.28	-0.071^{***} (0.014)	0.39	-0.098^{***} (0.015)
Agriculture, Utilities, Mining, Construction, Manufacturing	739	0.07	0.002 (0.030)	0	--
Wholesale trade, Retail trade, Warehousing, transportation	741	0.06	0.016^{*} (0.008)	0	--
Information	735	0.39	-0.085 (0.074)	0	--
Finance, Insurance	738	0.60	-0.151^{**} (0.075)	0.86	-0.223^{**} (0.065)
Real Estate, Professional and scientific services, Administration	741	0.38	-0.149^{**} (0.071)	0.51	-0.144^{***} (0.039)
Education, Health	741	0.16	-0.008 (0.015)	0.18	--
Arts and recreation, Accommodation	739	0.14	-0.007 (0.034)	0	--
Other services, Public administration	740	0.12	-0.008 (0.034)	0	--
Missing industry	741	0.32	-0.190^{***} (0.018)	0.48	-0.196^{***} (0.023)

Notes: Each cell in columns 3, and 5 contains the coefficient on log distance to an AI publication hotspot (at least 1000 publications) from a different regression with full covariates (those of Table 2 column 3). The dependent variable is the change in the share of advertisements in the specified industry requiring AI. The NAICS 2 codes for each row are a) 11, 21-23, 31-33; b) 42, 44-45, 48-49; c) 51; d) 52; e) 53-56; f) 61-62; g) 71-72; h) 81, 92. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Appendix Table 7: Impacts of distance to closest AI publications hotspot and commuting zone's own publications on change in AI job advertisement share, twelve-year differences

Employer names in micro sample:	Valid	Missing	Valid	Missing	Finance, insurance	AI applications skills
	(1)	(2)	(3)	(4)	(5)	(6)
Log distance to closest hotspot	-0.018*	-0.132***	0.003	-0.146***	-0.151**	-0.010***
(1000+ AI publications)	(0.011)	(0.014)	(0.013)	(0.024)	(0.075)	(0.002)
Change in AI publications	0.439***	0.078	0.486*	0.089	0.317**	0.024
	(0.101)	(0.112)	(0.268)	(0.058)	(0.117)	(0.014)
2006 AI publication controls	Yes	Yes	Yes	Yes	Yes	Yes
(p-value joint significance)	(0.00)	(0.63)	(0.06)	(0.34)	(0.57)	(0.00)
Interactions of 2006 AI controls with	--	--	Yes	Yes	--	--
log distance (p-value joint significance)			(0.16)	(0.70)		
Observations	739	741	739	741	738	741
Pseudo R-squared	0.22	0.26	0.23	0.26	0.05	0.14
Change 0 to 1 AI publication	0.183***	-0.036	--	--	-0.096	0.020***
	(0.056)	(0.056)			(0.182)	(0.004)
100 AI publications	0.161**	-0.045	--	--	0.073	0.012***
	(0.061)	(0.056)			(0.136)	(0.003)
1000 AI publications	0.073	-0.050	--	--	0.047	0.008*
	(0.050)	(0.050)			(0.108)	(0.005)
Log distance evaluated at						
1 AI publication in 2006	--	--	-0.037**	-0.131***	--	--
			(0.018)	(0.030)		
100 AI publications in 2006	--	--	-0.031**	-0.128***	--	--
			(0.015)	(0.019)		
500 AI publications in 2006	--	--	-0.012	-0.116***	--	--
			(0.041)	(0.075)		
1000 AI publications in 2006	--	--	0.005	0.043	--	--
			(0.073)	(0.052)		

Notes: Median regressions, except OLS in column 5. Each column's dependent variable is the change in AI share based on different underlying samples of job advertisements. Controls include those of Table 2 column 3, including a dummy for any pre-2007 AI publication, the number of pre-2007 AI publications and its square. The coefficients for pre-2007 AI publications are multiplied by 1000. Robust standard errors in parentheses.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$