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

September 19, 2021

*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 and University of Sydney Microeconometrics and Public Policy Working Group for helpful comments on more preliminary work. Hunt is also affiliated with the 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 define hotspots, we create a geocoded dataset of all AI publications (journal articles, conference proceedings and patents) through 2020, while we obtain the job vacancy information from online job advertisements scraped by Burning Glass Technologies from 2007–2019. We define hotspots based on the cumulative number of AI publications by 2006. We find that a hotspot's AI publications increasingly affect other commuting zones' AI vacancies as the hotspot publication threshold grows to 300, a threshold met by 11% of commuting zones. A 10% greater distance from such 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. The effect is almost entirely due to the one third of job advertisements posted by employment agencies, for which the industry is not known, but we do measure a small effect due to advertisements posted directly by financial firms. We also show that distance is a barrier to the posting of vacancies involving AI applications. The results suggest that for a minority of firms, distance from innovation is a moderate barrier to the adoption or adaptation of technology.

The extent to which geographic distance is a barrier to technological knowledge transfer is of interest to governments of countries distant from centers of innovation; 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 innovation, or as a prerequisite for the adoption of innovative 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 innovation.

The importance of distance for the diffusion of innovation 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 innovation. 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 innovation 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 innovation, 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 innovation between inventors.[2]

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

[1] For within–country firm to firm moves see Agrawal, Cockburn and McHale (2006); 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.

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 instead seeks to examine whether distance constitutes a barrier between innovation and 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 innovation, we create a geocoded dataset of all AI publications (journal articles, conference proceedings and patents) through 2020, while we measure 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 innovation and adoption or of the geographic diffusion of AI is by Bloom et al. (2021). 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 innovations emerging 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.[6]

We approach the question by dividing the United States into 741 commuting zones

[3] Little and Triest (1996); Comin, Dmitriev and Rossi–Hansberg (2012).

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

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

[6] Andersson, Quigley and Wilhemsson (2009) look at the impact of decentralization of Swedish universities on both prouctivity and innovation. Acemoglu, Autor, Hazell and Restrepo (2021) examine the growth of AI job advertisements in the Burning Glass Technologies data, but do not consider geography.

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 if we set a very low publication threshold for a commuting zone to be designated an innovation hotspot, distance from the nearest hotspot plays no role in the growth of AI job advertisements: presumably commuting zones designated hotspots in this way in fact have too little innovation to be influential. As the hotspot threshold grows to 300, a threshold met by 11% of commuting zones, a hotspot's AI publications increasingly affect other commuting zones' AI vacancies. A 10% greater distance from such 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. The effect is almost entirely due to the one third of job advertisements posted by employment agencies, for which the industry is not known. This suggests that for a minority of firms, distance from innovation is a moderate barrier to the adoption of technology. While this group of firms surely posts advertisements for jobs involving both innovation and adoption (and adaptation), we find that a small fraction of the effect is due to job advertisements posted directly by financial firms. We also find that distance is a barrier to the posting of vacancies involving applications of AI, rather than the development of AI. These results show that distance is a barrier to the adoption or adaptation of innovation.[7]

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

Our findings are robust to the second approach using the (log) radius of the circle enclosing a given number of AI publications: here the effect of distance grows more negative as the threshold number of publications rises to 3000. The findings are also robust to measuring job advertisements cumulatively over time instead of contemporaneously. Since distance is not a barrier for AI diffusion among vacancies posted directly by employers, it is not surprising that we find no evidence that the distance barrier between a hotspot and another commuting zone is lower when they have more advertisements placed by common employers, since this can only be measured for vacancies posted directly by employers.

1 Data

We have created our own database of AI publications and patents, and use 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 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.

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

[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.

[9] We match cities to counties using the file provided at https://simplemaps.com/data/us-cities, accessed 18 May 2021. A number of cities not in this file were hand–assigned a county using Wikipedia. 1000 job advertisements from cities not in the file remain to be matched to a county and hence commuting zone.

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 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 190 million job advertisements, we use the location, the NAICS industry code, the standard occupation classification code, classifications of keywords for required skills, and the employer name. We harmonize differing versions of employer name. A missing value for the employer name means the advertisement is posted by an employment agency: in almost all cases, Burning Glass Technologies codes the employer name as missing if the employer is an employment agency.[10] Since Burning Glass Technologies infers industry principally from the employer name, this means that most but not all vacancies with a missing firm name are also missing industry. In a few cases the employment agency name is retained and the vacancy assigned the NAICS 2 code 56 (Administrative and support services), while a few observations with missing employer name have a valid industry code different from 56.

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).[11] We designate job advertisements as being IT job advertisements if the Skill Cluster Family (most aggregate) field contains "Information Technology" as long as the Skill field does not contain names of Microsoft Office software (and the advertisement is not also an AI advertisement, though there is

[10] Burning Glass Technologies, personal communication.

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

almost no overlap). We then 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 hotspots have a closest hotspot. For use with this independent variable, we also compute the distance to the closest large commuting zone for each commuting zone, with the definition of a large 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 large commuting zone as one of the h largest 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

[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.

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: $\Delta^k AI^s_{ct} = \frac{AI\ job\ ads}{All\ job\ ads}_{c,t} - \frac{AI\ job\ ads}{All\ job\ ads}_{c,t-k}$, for several reasons. Particularly in early sample years, a large share of commuting zones have no AI job advertisements and many have only one or two, making short differences in the early years either zero or very large for many commuting zones. This also suggests avoiding fixed effects (including Poisson fixed effects), which might use such 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^{Hot}_c:

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

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 ($IT^s_{c,2007}$).

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^{min}_c. 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: 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 or adaptation, 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 $\Delta^k IT^s_{c,t}$ in all advertisements 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 vacancies 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.

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, clustering standard errors by commuting zone.[13] This also downweights the large outliers in the outcome.[14] OLS point estimates of σ are somewhat larger 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.

While it seems natural to form a panel using a dependent variable based on what we

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

[14] 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.

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 employment is more readily interpretable than a change in vacancies. 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 (or adaptation) of AI innovation, or merely to additional innovation in AI, and to shed light on the mechanism by which distance slows diffusion. For this purpose, we investigate the role of distance by industry, using as the outcomes the number of AI job advertisements in a particular industry, including missing industry, divided by total job advertisements. In other regressions, the commuting zone–year variables are calculated based on subsamples of the job advertisements e.g. subsamples 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. Occupation is not very helpful as such a large (though declining) majority of AI vacancies posted are for computer and mathematical occupations.[16]

[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. These data could be used in a future version of our paper.

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 Figure 3 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 0.31 percentage point, while the median increase was lower at 0.25 percentage point (first row). The minimum value of -2.81 percentage points and the maximum value of 3.64 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 0.15 percentage point and the median increase is 0.11 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 4. Publications (times number of authors) increased from 7 in 1950, to 12,063 in 2007, to 49,882 in 2018 and to 65,378 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 5 depicts the number of publications for each of the 32 commuting zones satisfying this requirement. The three top publishers are Los Angeles, Arlington, V.A. (the area around Washington, D.C.) and Boston, each with more than 5000 publications, followed by the trio of San Jose, New York and Pittsburgh, with more than 4000 publications each. The highest publishing commuting zone outside the Northeast and California is Seattle, W.A. in ninth place, and in the ranks after Seattle, Midwestern then Texan commuting zones begin to appear. 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 6 shows the distribution of these cumulative publications, while the succession of maps in Figure 7 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. Finally, we investigate sources of heterogeneity in the distance effect, both to distinguish among advertisements for jobs in AI innovation versus adoption, and to seek the mechanism through which distance influences AI job advertisements.

4.0.1 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). We initially 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 -0.032 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 (0.032)(0.1)=0.0032 percentage point. This is 3.0% of the median growth rate of 0.108 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, render the coefficient of interest slightly more negative at -0.035. 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. The addition of the distance to the closest large commuting zone (one of the 32 most populous, since there are 32 AI publication hotspots) in column 4 increases the coefficient on very slightly, to -0.028: 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. Thus, -0.028 is our preferred coefficient of interest, corresponding to 2.9% of median growth.

We test whether our results are affected by the remote commuting zones of Alaska and Hawaii by dropping them from the estimation in column 5: this has almost no effect on the coefficient of interest. On the other hand, using mean rather than median regression (column 6) makes the coefficient considerably more negative, at -0.054. In the first six columns, the sample includes commuting zones that are themselves hotspots. Dropping the hotspots from the sample (column 7) yields a coefficient essentially the same as in the otherwise comparable regression in column 3.

Panel B shows the corresponding coefficients from the twelve–year difference regression: 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 -0.098 (implying a 10% increase in distance reduces AI job growth by 3.2% of median growth) to -0.053 (implying only a 1.7% reduction). This may reflect the correction of a negative bias in the previous specifications.

We find similar patterns in our estimation of the impact of the same covariates on the cumulative AI share. The coefficients on the log distance to the closest hotspot, presented in columns 1–3 of Table 3, are one half to one third the magnitude of those in Table 2, though not directly comparable, and always statistically significantly negative.

Our second 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 -0.038—0.040 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 0.0038–0.0040, or 3.5–3.7% 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 0.007–0.010 percentage point, or 0.9–1.3% 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.

4.1 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 8, 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.

In all four graphs, the effect of distance when the threshold is zero or one publication is zero (in panel C the points for these two thresholds are visually indistinguishable), then becomes increasingly negative as the threshold rises to 300 publications. The effect is then similar until at least 1000 publications. Standard errors are larger when the distance to the closest large commuting zone is controlled in the right two panels, but this control has little effect on the point estimates.[17]

In all four graphs, the effect of distance seems to weaken after the stable set of coefficients between 300 and 1000 publications , in the case of the right hand graphs to zero. All patterns are similar in Figure 9, which contains the corresponding graphs for cumulative AI job advertisements. Especially in the right hand graphs of the two figures, the weakening of the distance effect at high thresholds involves a jump when Austin, TX (1922 AI publications in 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 (that technology diffuses immediately from the largest hotspots, for example).

The corresponding graphs using the radius of the circle enclosing a given number of AI publications are shown in Figure 10. Using this identification strategy, the effect of

[17] 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 correspondingly large.

distance does not go to zero in any graph at high thresholds, in the threshold range we have plotted.[18] In Figure 11, we plot the corresponding graphs for the hotspot identification strategy and the probability of having any AI job advertisement. It is less evident here that the effect of distance increases with threshold at low thresholds, but we again see the sensitivity to Austin.

4.2 Disaggregating the effect of distance

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, both to shed light on the mechanism and 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.

An obvious step is to determine the industries through which the distance effect operates.[19] We show in Table 4 that the insight to be obtained from industry is less than might be hoped. Column 1 reproduces the preferred specification from Table 2 (column 3) for seven–year differences (panel A) and twelve–year differences (panel B). Columns 2 and 3 represent the same regressions, but with the values for each commuting zone–year observation calculated using job advertisements with valid industry (column 2) and missing industry (column 3); 37% of advertisements have missing industry (see Appendix Table 2).[20] The effect of distance is four times more negative when based on advertisements with missing than with valid industry (-0.053 compared to -0.014 in panel A, and a similar ratio in panel B), though all effects are statistically significant.

Since Burning Glass Technologies codes industry based primarily on firm name (Burn-

[18] The radii for thresholds in the range 10,000–40,000 all have a correlation of at least 0.9 begin to enclose significant proportions of the national AI publications as of 2006 (about one third at 40,000), which is why we plot only to 10,000. The population controls over this range are not highly correlated, however, and the coefficient on distance does increase over this range and eventually become positive in specifications including population.

[19] Examining the precise type of AI skill required in the job advertisement (if specified) would provide further evidence; we defer this to the next version of the paper.

[20] Missing industry means missing NAICS 3. However, some job advertisements have a valid NAICS 2 but missing NAICS 3. The next version of the paper will distinguish based on NAICS 2.

ing Glass Technologies 2019), we investigate whether the strong effect of distance among job advertisements with missing industry fundamentally reflects missing firm name, and therefore the effect of vacancies posted by employment agencies. Columns 3 and 4 show that indeed, the effect of distance for job advertisements with a valid firm name is small and statistically insignificant (coefficient -0.006 in panel A and -0.019 in panel B), while the effect for those with a missing firm name (one third of job advertisements) is an order of magnitude higher (-0.077 and -0.170 respectively).

Descriptive statistics in Appendix Table 2 suggest that compared to all job advertisements, advertisements without a firm name are less likely to require AI; more likely to be (explicitly) in administrative and support services; and more likely to advertise in computer and mathematical occupations. Because advertisements due to (named) firms with few total advertisements are the most similar in terms of these characteristics, we analyze the effect of distance using an underlying job advertisement sample of vacancies at firms posting 100 or fewer per year (Table 3 column 6, using OLS). The implications of the seven–year difference coefficient (a statistically significant -0.028 in panel A) and the twelve–year difference coefficient (a small and statistically insignificant -0.012 in panel B) differ: while the latter suggests the influence of advertisements by unnamed firms is not due to their being small, the intermediate size of the former leaves open the possibility that the unnamed firms may be somewhat disproportionately small.[21]

Although Table 4 suggests that distance is not a barrier to job advertisements with a valid industry as a whole (column 2), distance could nevertheless constitute a barrier for some particular industries. Accordingly, we present in Table 5 the results of regressions based on the full set of job advertisements, with as dependent variable AI job advertisements in a particular industry group (including missing industry) as a share of all job advertisements. Dividing AI job advertisements in this way leads to a median value of zero for the dependent variable in some regressions, making OLS regression more appropriate. The coefficients from the industry regressions sum to the coefficient for all industries, and

[21] The twelve–year difference coefficient is not sensitive to the inclusion of changes in job advertisements, IT share and AI publications.

the relative sizes of the effects reflect both the effect for the industry and the size of the industry.

The first row in columns 1 (seven–year differences) and 2 (twelve–year differences) reproduce the OLS coefficients for all industries already shown in Table 2 column 6 (-0.053 and -0.079, respectively). The last row shows that almost all (column 1) or even more than 100% (column 2) of the total distance effect is due to AI in advertisements with missing industry, as expected. The only other industry group with a consistently statistically significant effect is administrative and support services, with much smaller distance coefficients of -0.0021 (column 1) and -0.0054 (column 2). However, the point estimates for the industry group containing information, finance, insurance, real estate and management are considerably more negative (-0.0064 and statistically insignificant in column 1; -0.015 and statistically significant in column 2).

Although there is considerable overlap between job advertisements with missing industry and missing firm name, some observations with missing firm name have a valid industry. If the industry distribution were similar among advertisements with and without a firm name (admittedly unlikely), we could learn from analyzing industry based on the sample of advertisements with a firm name, as in columns 3 and 4 of Table 3. The first row shows the relatively strong distance effect for all industries (-0.11 and -0.18), while the last row shows these are fully accounted for by the missing industry group. As in column 1 and 2, the coefficients corresponding to other industry groups are much smaller, but statistically significant for adminstrative and support services, and the industry group comprising information, finance, insurance, real estate and management. Unlike in the first two columns, the estimates for the education and health group are also statistically significantly negative, and the most negative point estimates are for administrative and support services.

In unreported regressions, we have further disaggregated industries, which allows more clarity as to whether distance is retarding innovation or adoption. We find statistically significant distance effects for firms in the finance and software industries, the former result suggesting that some of the effect measured does indeed represent adoption or adaption of AI, rather than further innovation.

Another 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. A large share of advertisements requiring AI simply require either "Artificial Intelligence" or "Machine Learning" skills, with no further detail specified. These unspecified AI skills comprise our first category. The remaining categories are not mutually exclusive. Image processing, which seems less than tightly linked to AI, is its own 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). The aggregate shares of these AI categories in all job advertisements are shown in Figure 12, while the exact definitions are in Appendix Table 1.

We present the results for the effect of distance on these AI categories' shares in Table 6. Panel A's first row shows that the largest category of AI is the unspecified category, constituting 37% of AI ads, (based on the micro–data), with only small minorities of advertisements mentioning AI tools (9.1%) or applications (19.1%). The second row presents the share of AI advertisements in these categories advertising for a computer science or mathematics occupation, as a proxy for a job closer to innovation than adoption. The share is highest in the AI tool category (80.0%), compared to only 53% in the AI application category, suggesting the category distinctions are meaningful. The lowest share is in image processing (50%).

Panel B presents the results of seven–year difference median regressions, and panel C the results of twelve–year difference median regressions. The effect of distance from an AI hotspot (with at least 1000 publications) is statistically significantly negative for all categories except image processing (and AI applications in the seven–year difference specification). The results suggest that distance from innovation is a barrier to both innovation and adoption of AI, and that image processing is possibly inappropriately included as AI.[22]

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 62% of AI advertisements (see Appendix Table 2). We hypothesize that when more such ties exist, the effect of distance to the closest hotspot will be smaller. 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 firm name, and our distance effect is found to operate among advertisements with no firm name (employment agencies).

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. The effect works almost entirely through the one third of advertisements placed by employment agencies, indicating that firms likely to hire through employment agencies are also likely to experience distance as a moderately sized barrier to AI job growth. Although the industry of the job is not available for most advertisements placed by employment agencies, who presumably post jobs on behalf of clients in a mix of industries, the facts that distance

[22] Median regressions do not converge in the case of image processing, so we present OLS results.

is a barrier for vacancies posted directly by financial firms and for vacancies using AI applications indicate that distance from AI innovation hinders adoption and adaptation of AI, rather than simply hindering further AI innovation.

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.

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.

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.

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

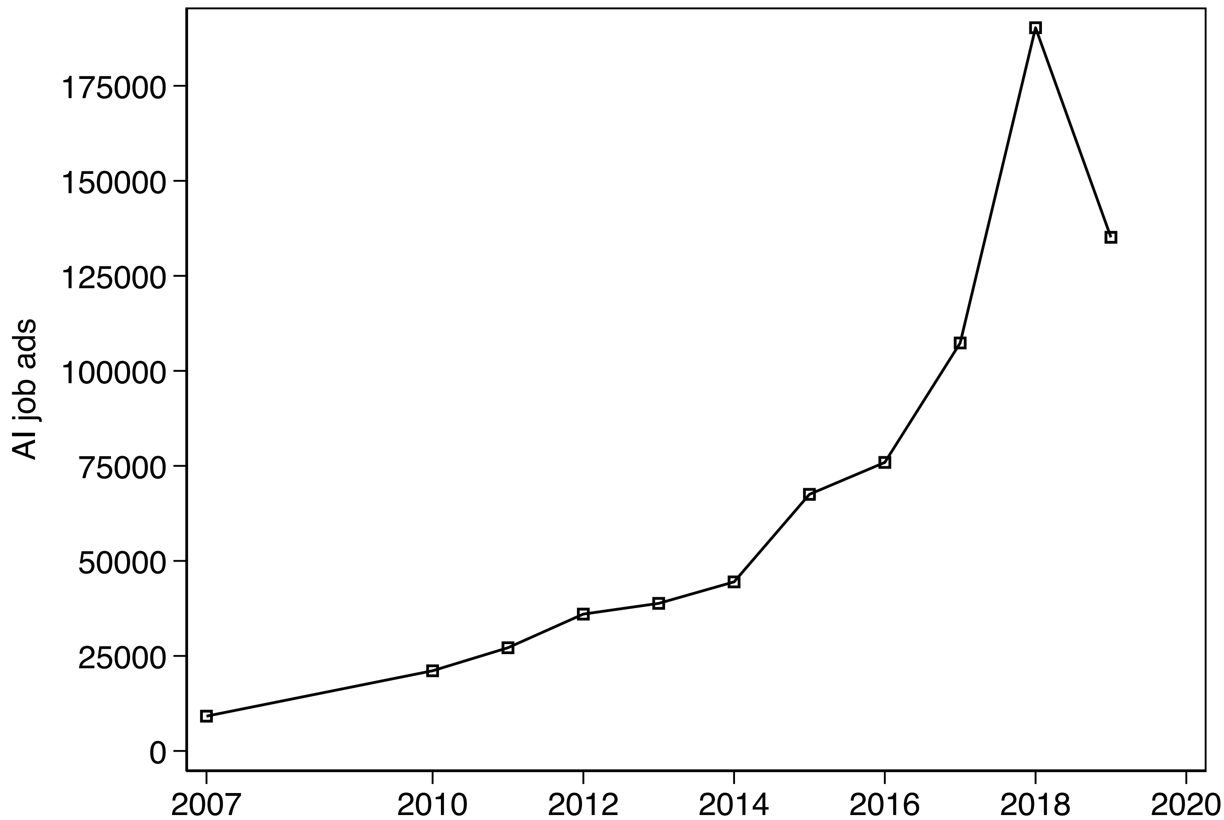

Notes: Data for 2019 are for January–July. Data for 2008 and 2009 are not available.
Source: Burning Glass Technologies.

Figure 2: AI share of job ads (%)

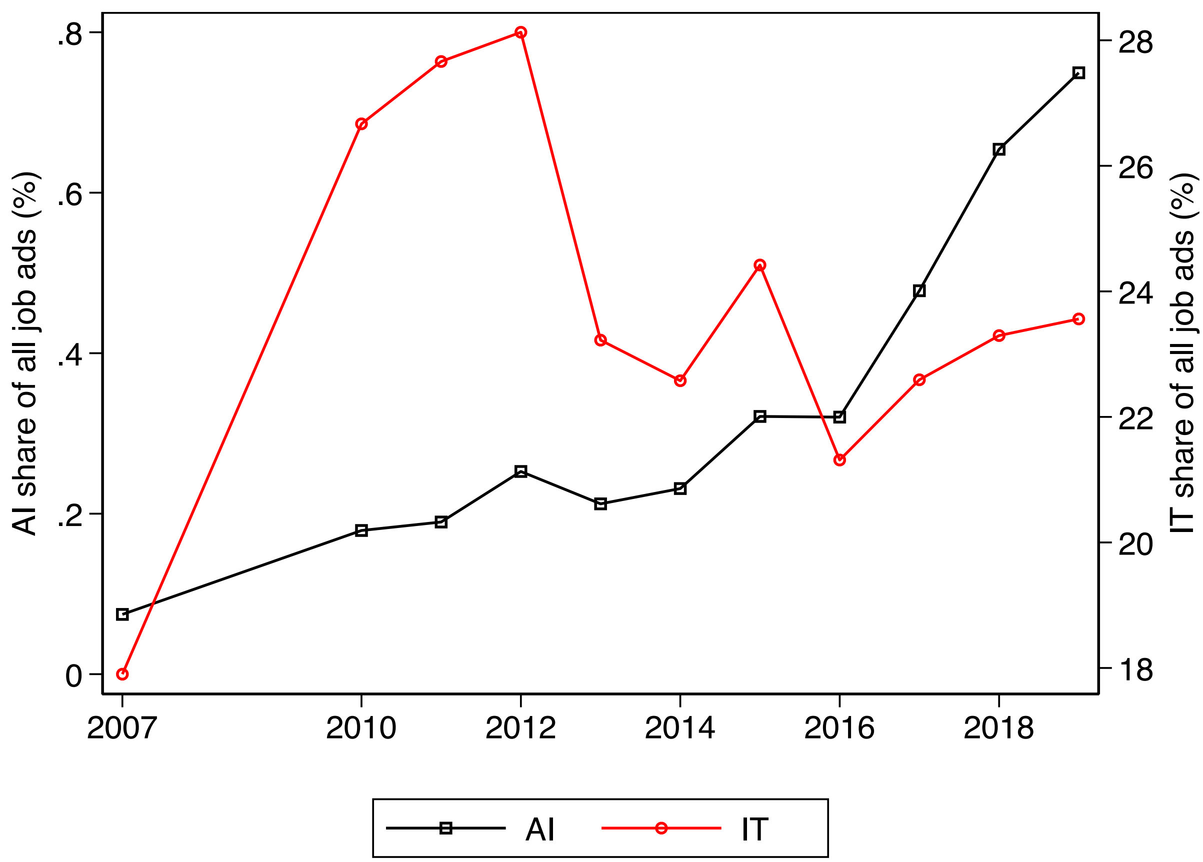

Source: Burning Glass Technologies.

Figure 3: AI job advertisements as percent of jobs advertisements in given year

(a) 2007

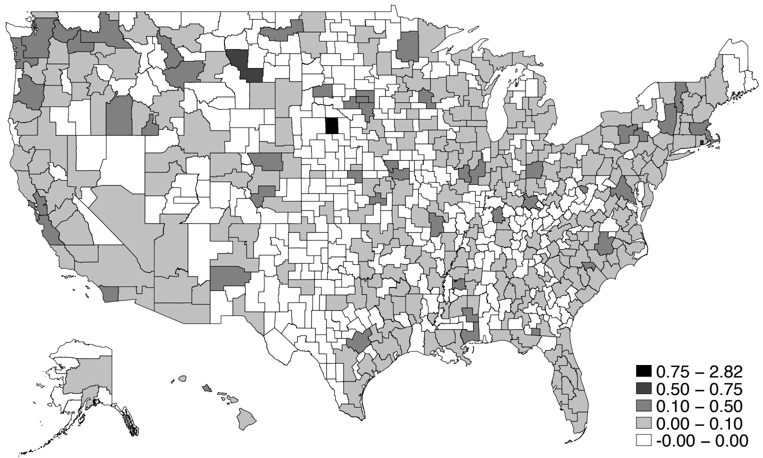

(b) 2010

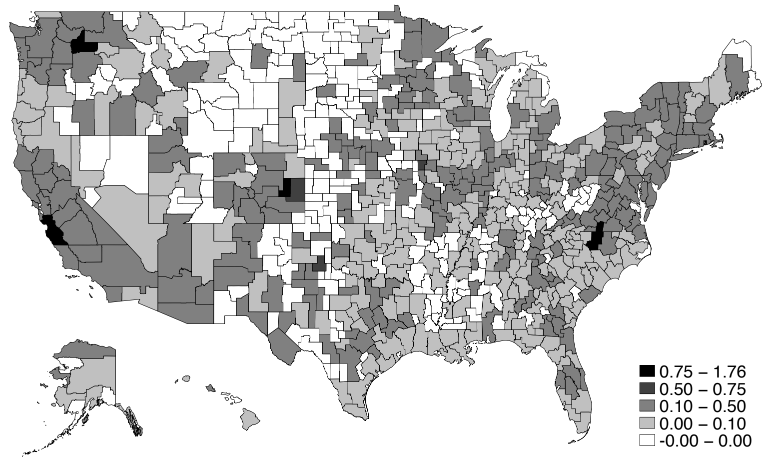

(c) 2014

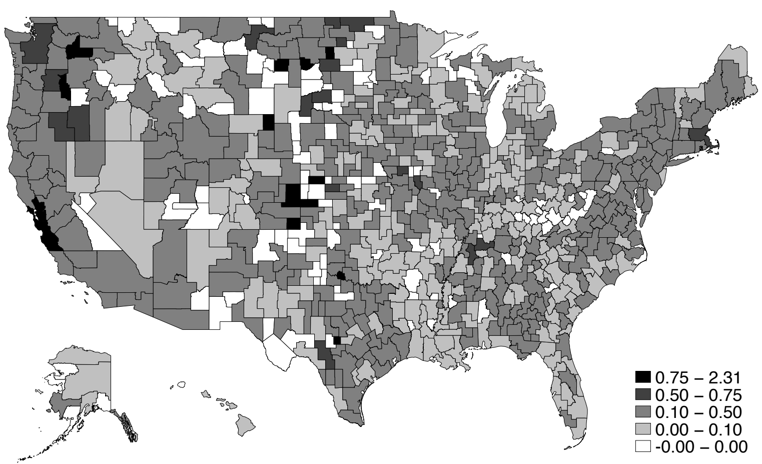

(d) 2018

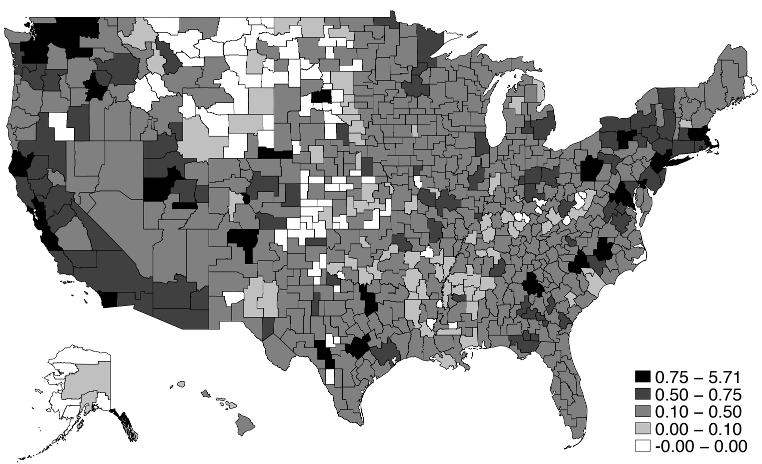

Figure 4: AI publications 1950-2019

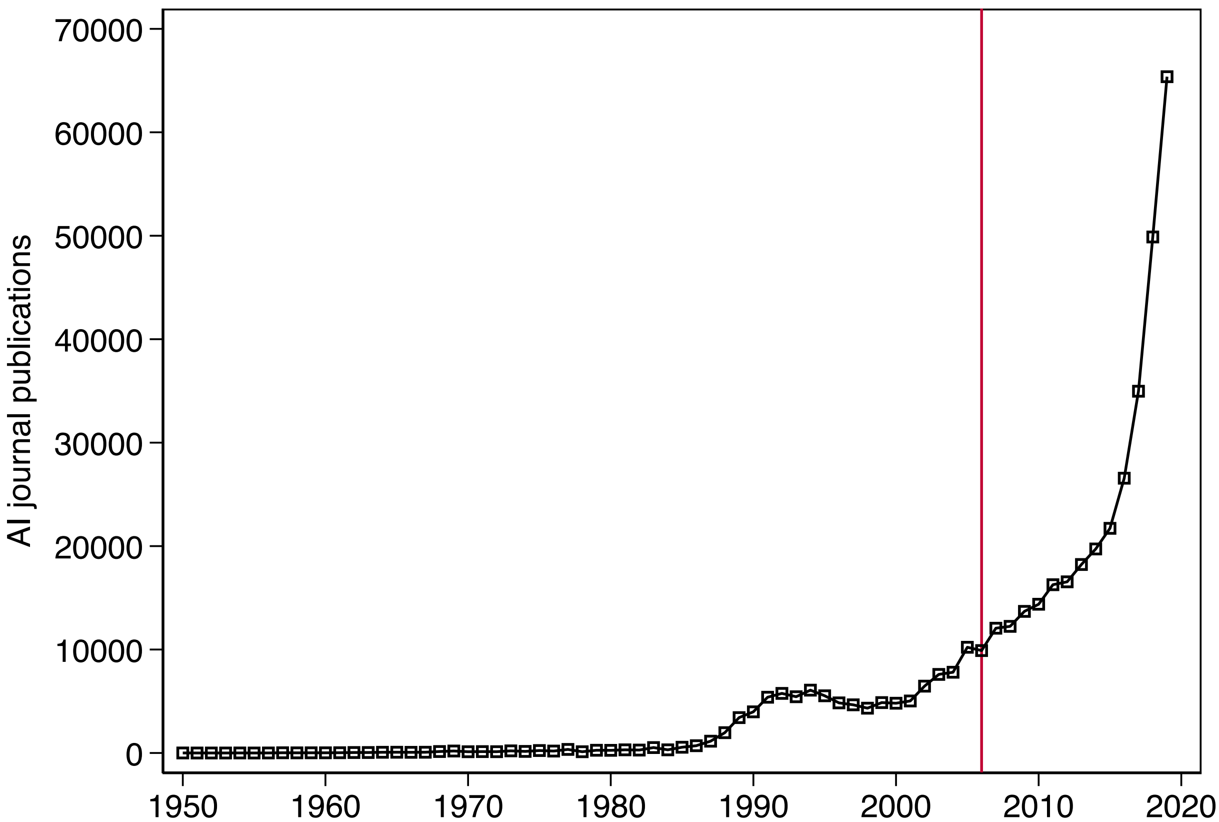

Source: Authors' dataset.

Figure 5: Innovation hotspots' AI publications through 2006

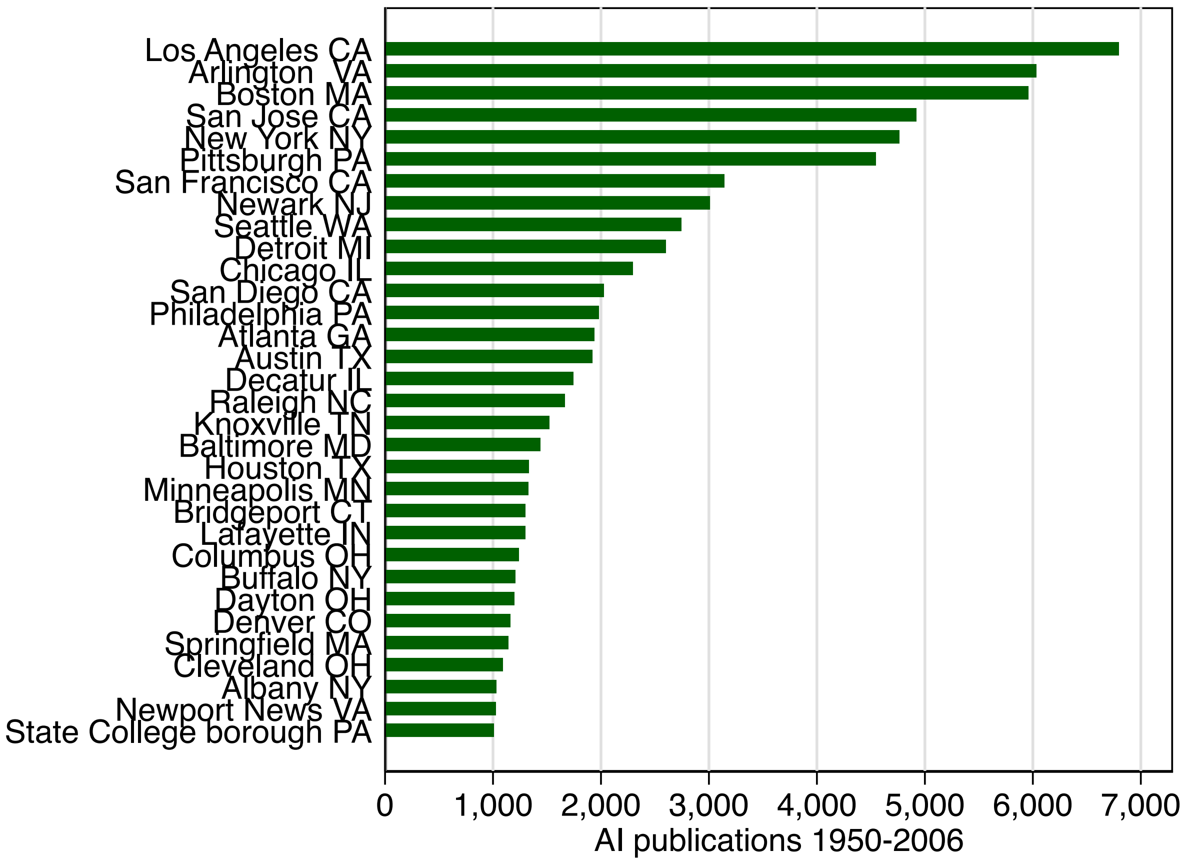

Note: The definition of a hotspot here is a commuting zone with at least 1000 AI publications through 2006.

Figure 6: Commuting zones' AI publications through 2006

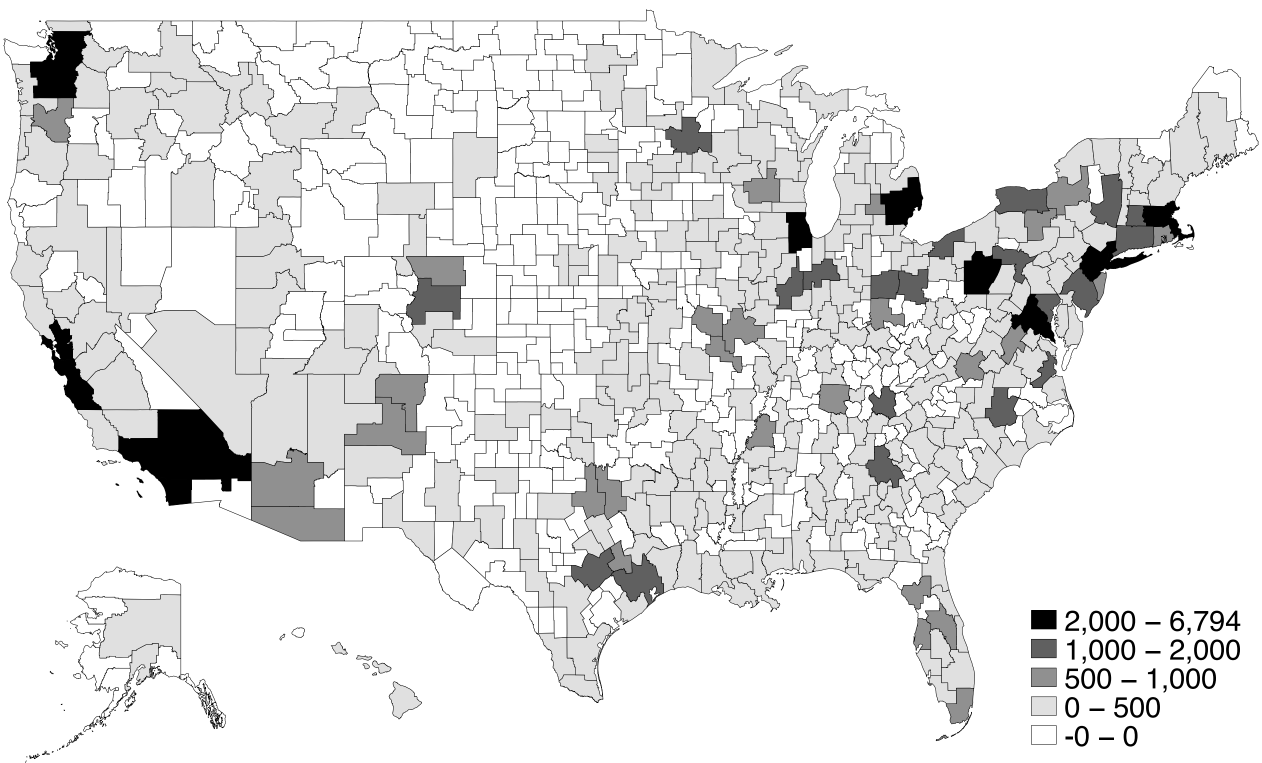

Figure 7: Commuting zones' AI publications in given year

(a) 2007: 240 CZs with any publication

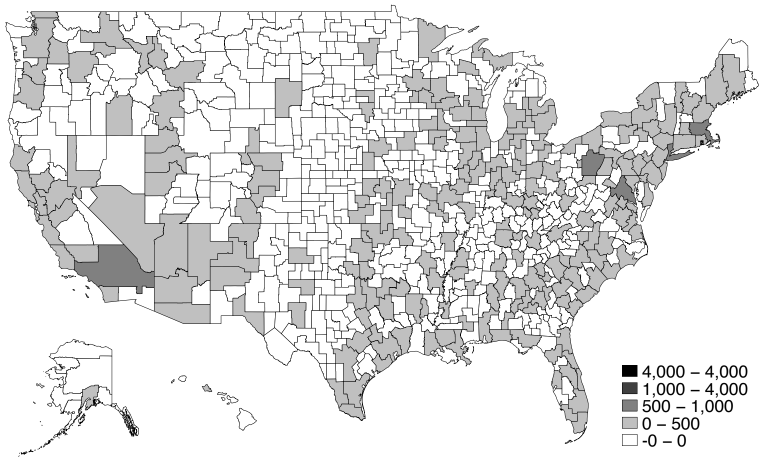

(b) 2010: 246 CZs with any publication

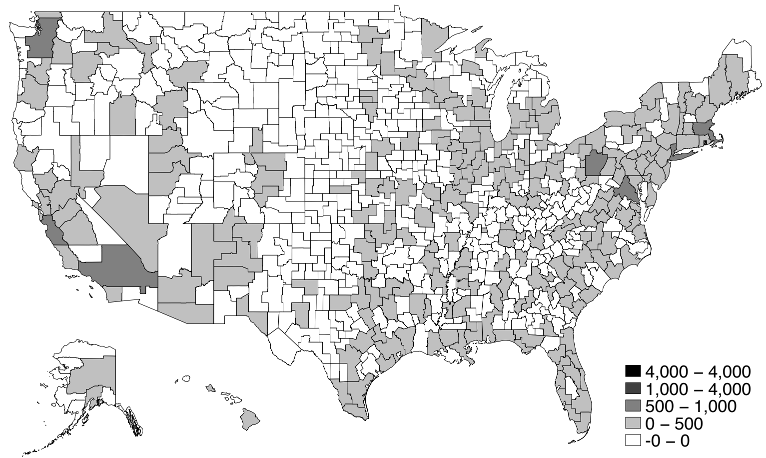

(c) 2014: 254 CZs with any publication

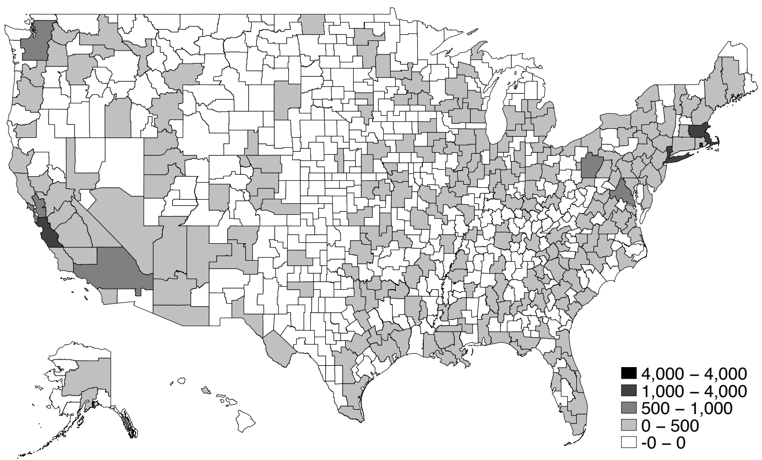

(d) 2018: 292 CZs with any publication

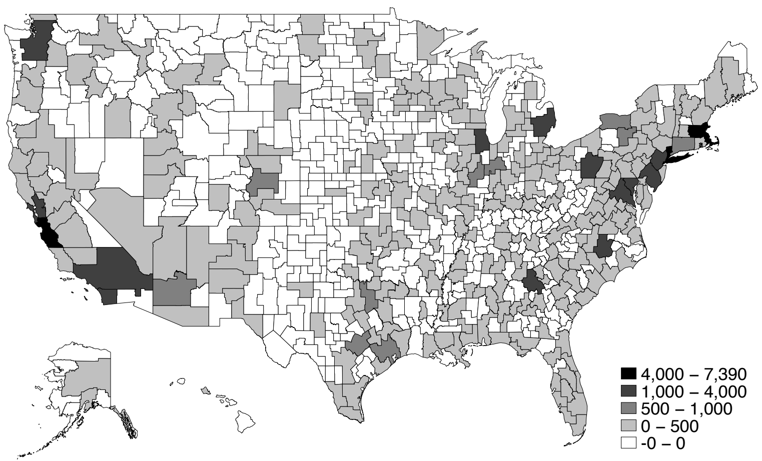

Figure 8: Coefficients on distance to closest hotspot

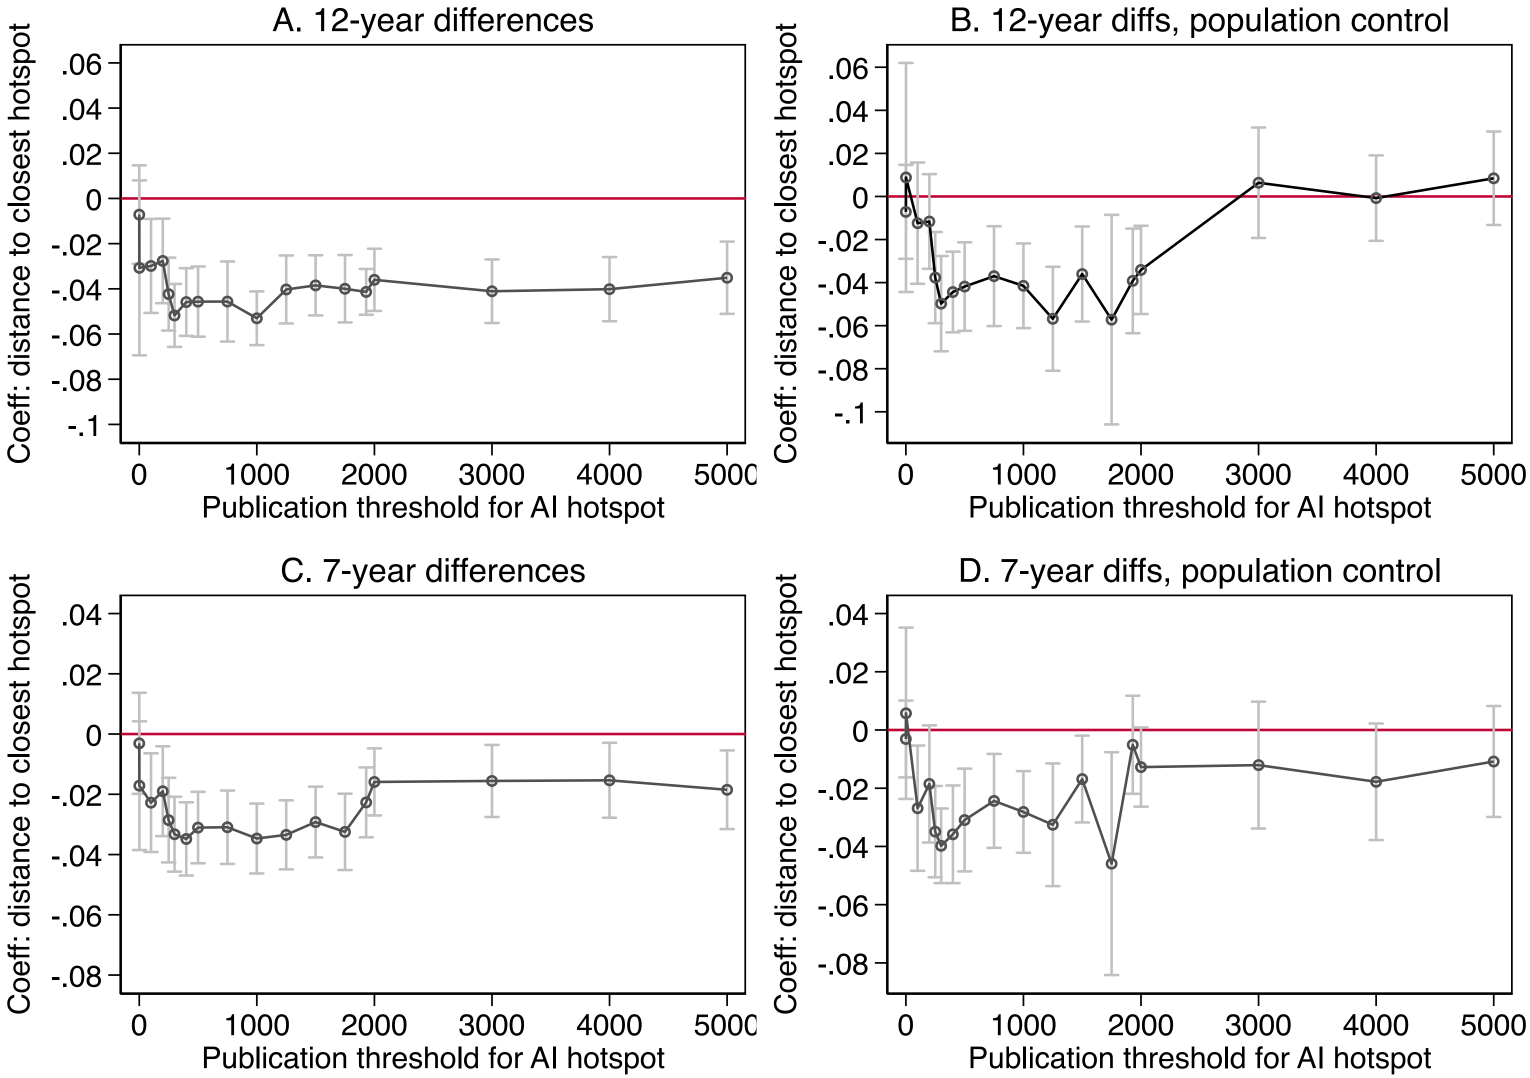

Note: The dependent variable is the change in AI jobs' share in total job advertisements. Population control refers to the distance to the closest large commuting zone.

Figure 9: Coefficients on distance to closest hotspot: Cumulative AI job advertisements

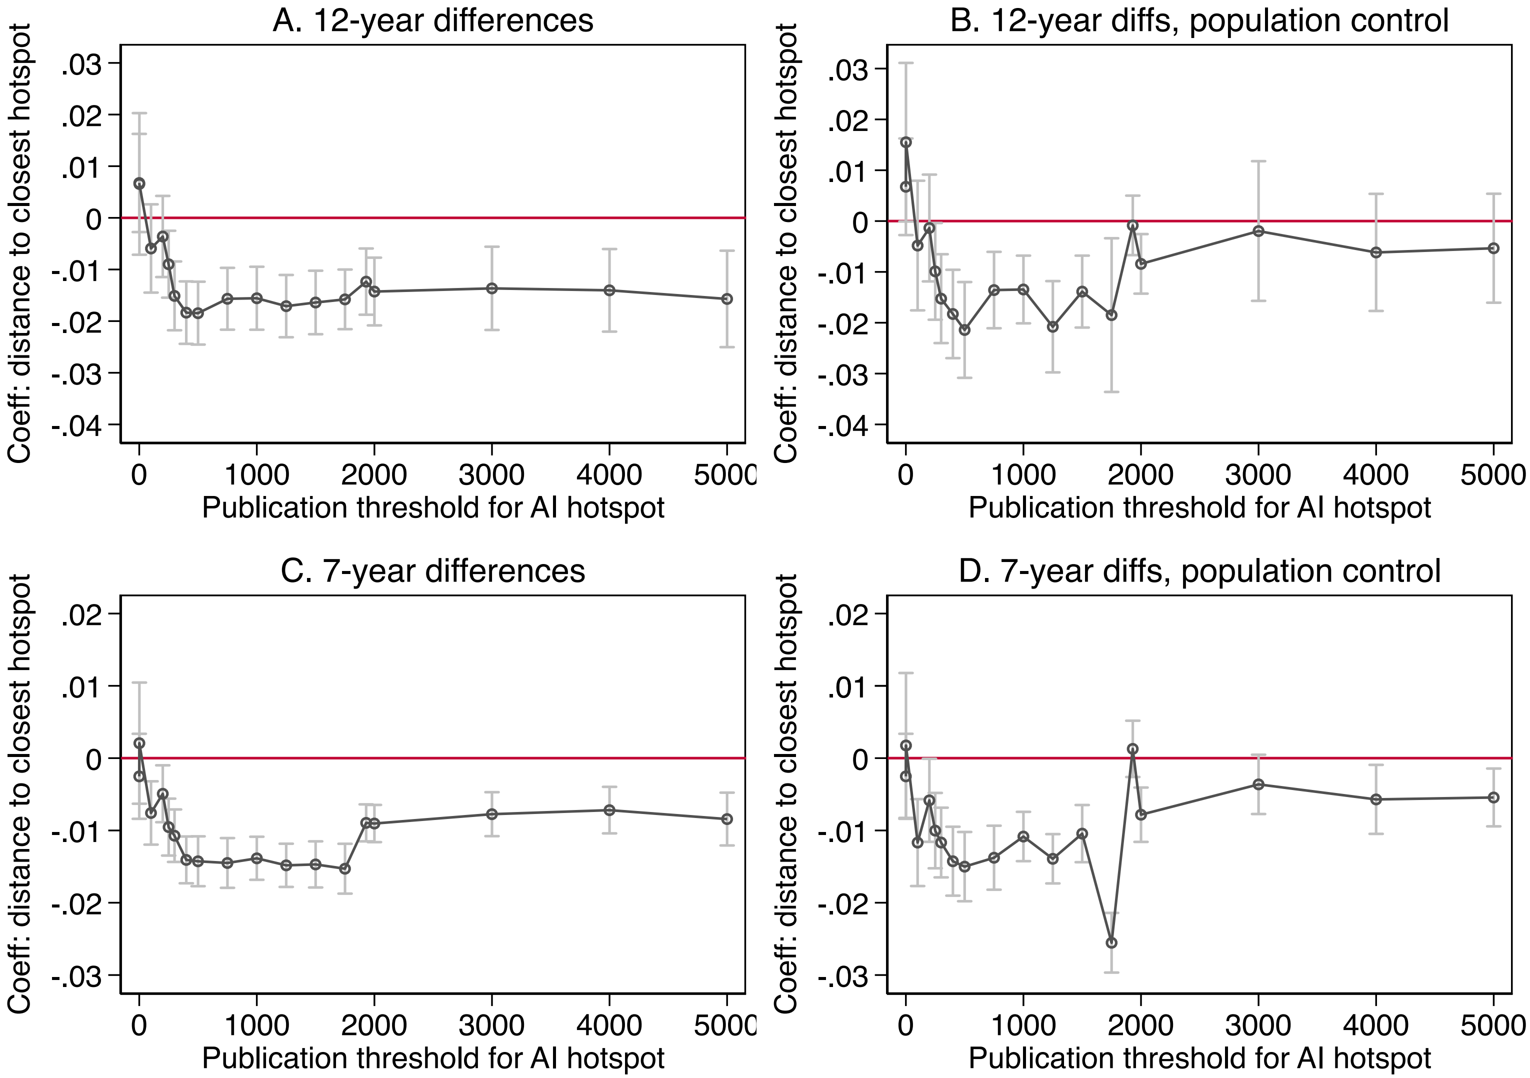

Note: The dependent variable is the change in cumulative AI jobs' share in total job advertisements. Population control refers to the distance to the closest large commuting zone.

Figure 10: Coefficients on log radius including threshold number of publications

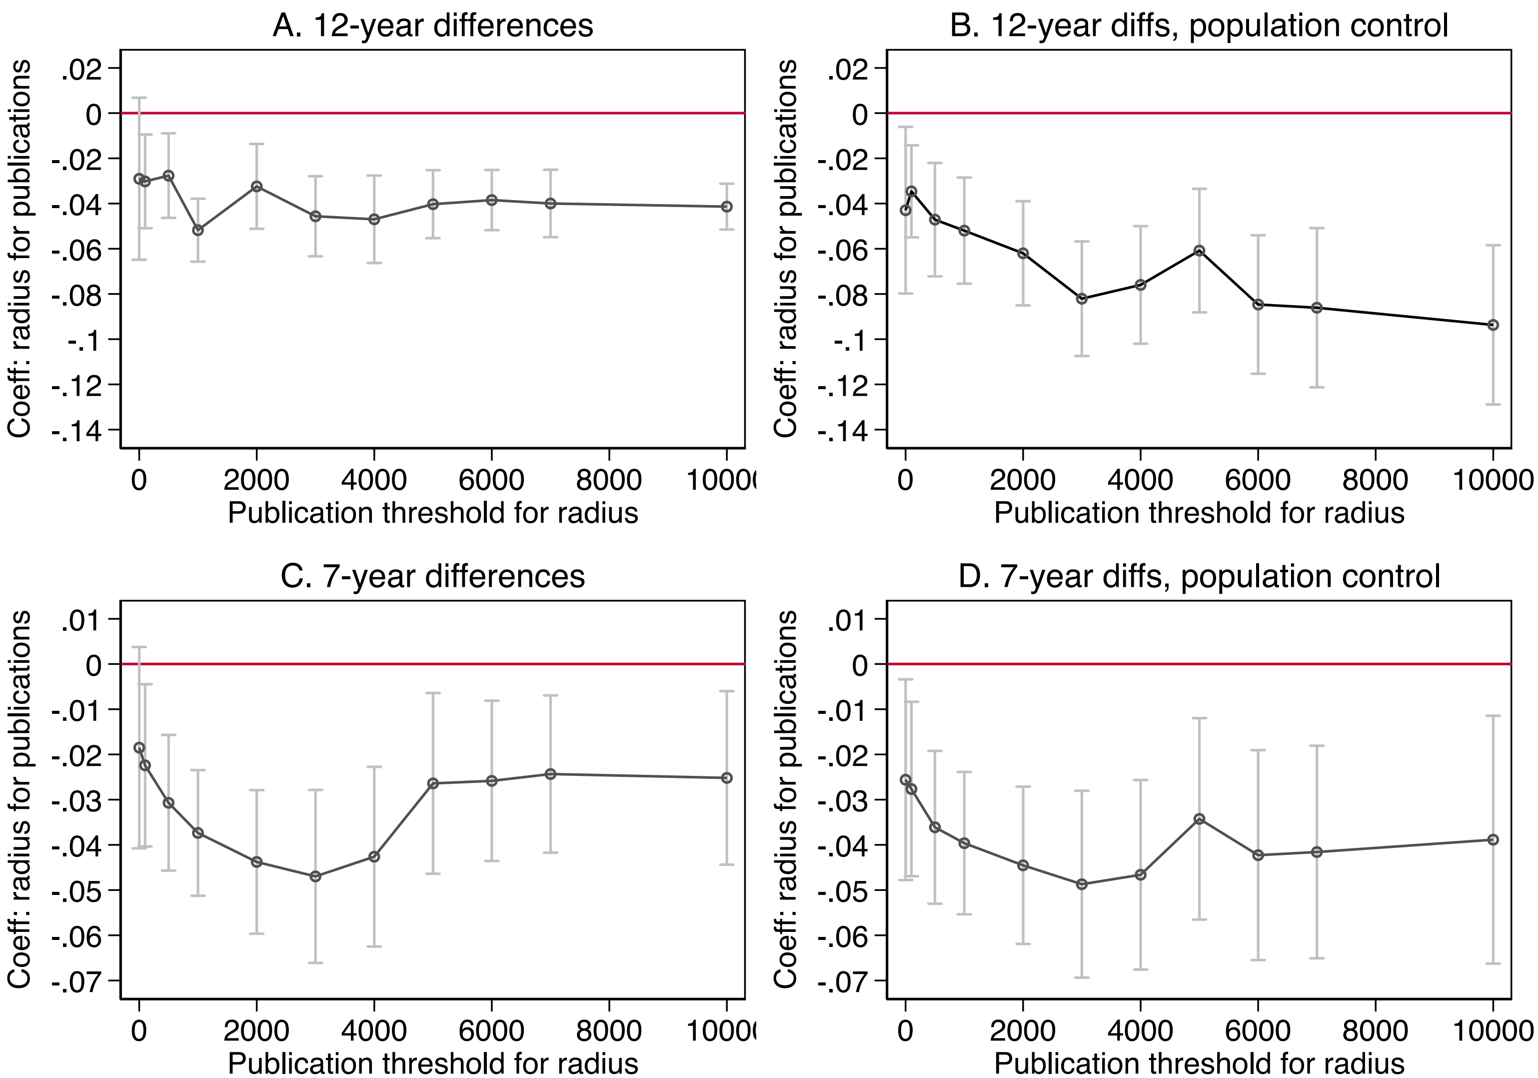

Note: The dependent variable is the change in AI jobs' share in total job advertisements. The coefficients plotted are on the log of the radius around the commuting zone which includes the threshold number of publications. Population control refers to the log population included within the radius.

Figure 11: Coefficient on distance to closest hotspot: extensive margin

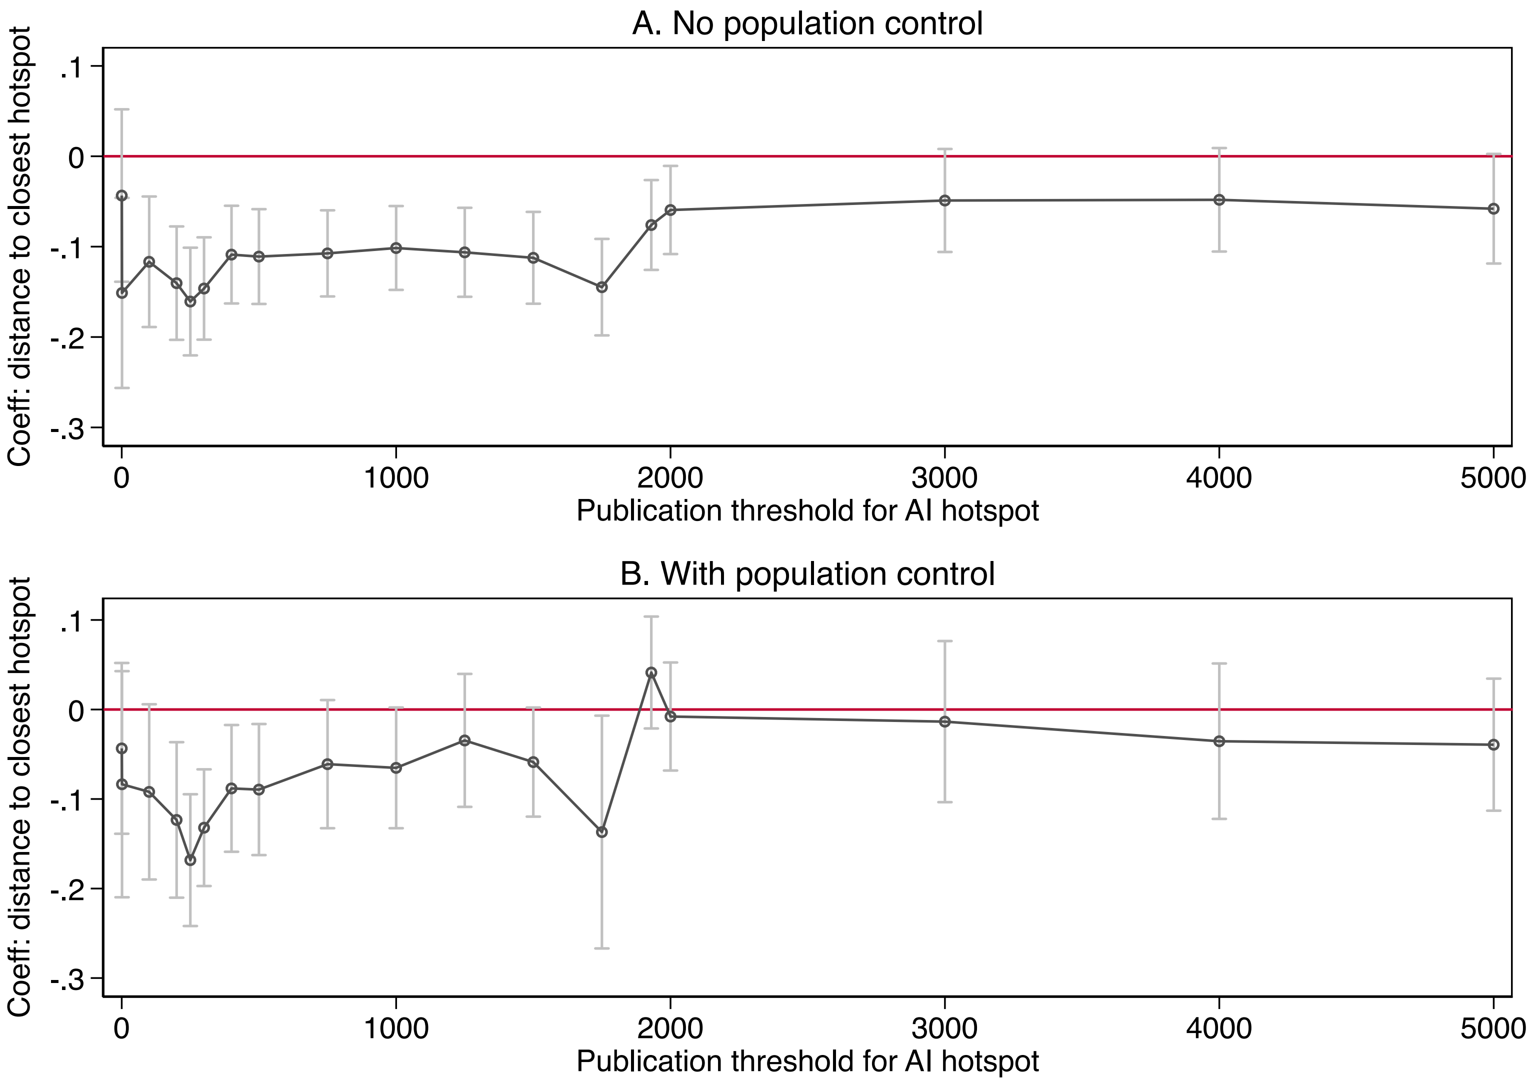

Note: The dependent variable is the probability of having any AI job advertisement in 2018 conditional on having none in 2007. Population control refers to the distance to the closest large commuting zone.

Figure 12: Growth in share of job advertisements accounted for by different types of AI (%)

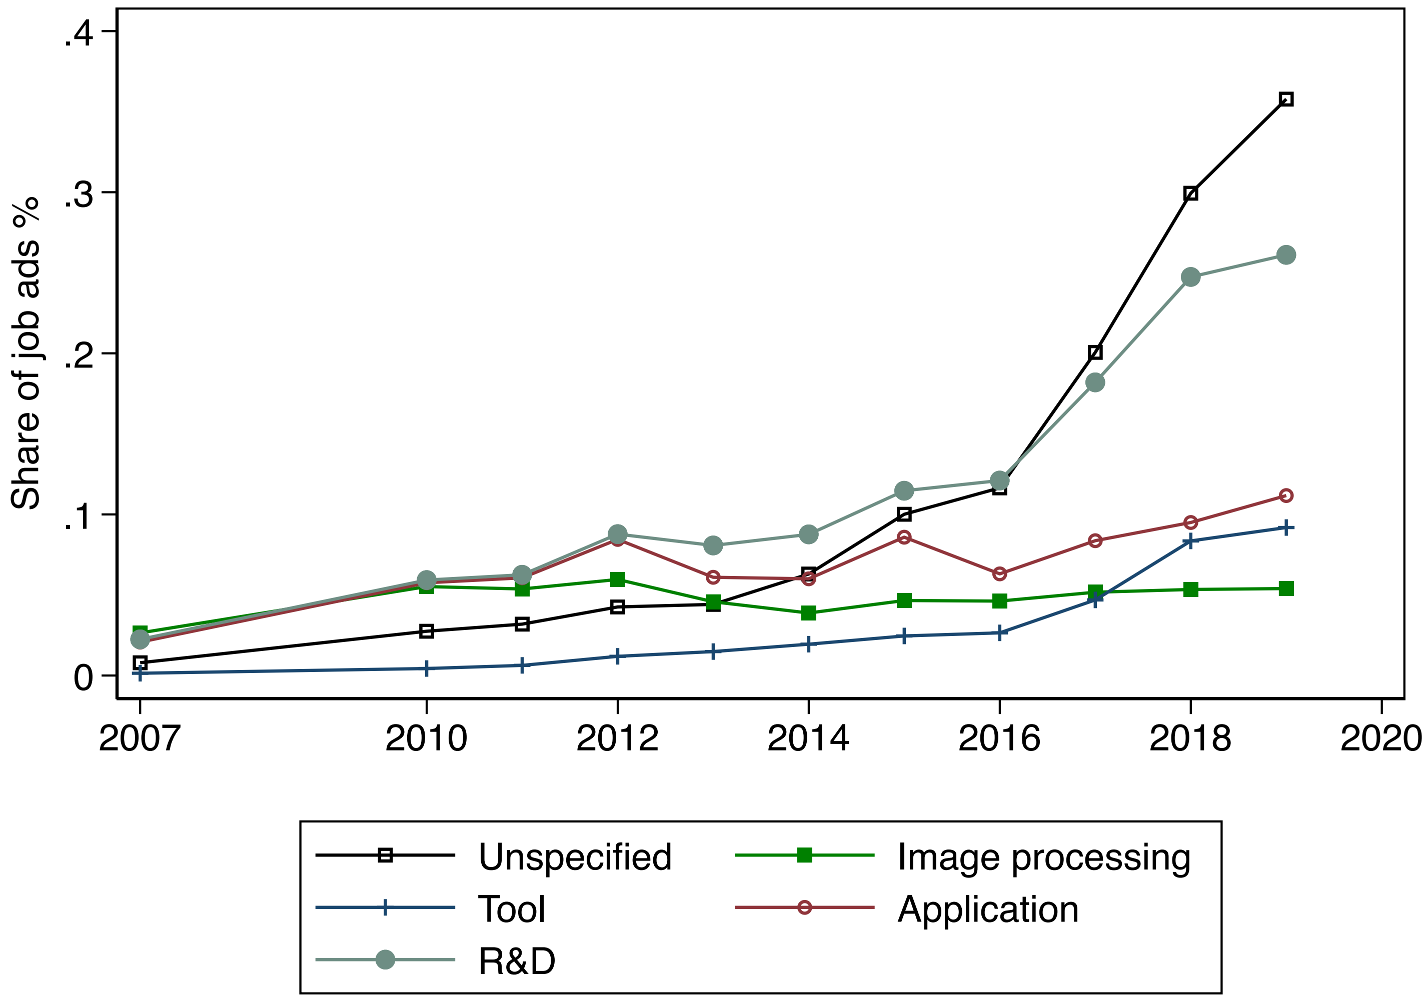

Note: Unspecified AI job advertisements require “Artificial Intelligence” and/or “Machine Learning” skills with no further detail given. The other categories are not mutually exclusive.

Table 1: Summary statistics

	Mean	Median	Min	Max	Obs
A. Δ AI job advertisement share (%)					
2007-2019 **Δ=12**	0.307	0.249	-2.81	3.64	741
2007-2019 **Δ=7**	0.155	0.108	-2.81	5.51	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	285	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 thru 2006	0.48	0	0	1	741
AI publications thru 2006	155	0	0	6794	741
Job advertisements 2007	15,221	2105	3	654,605	741
Population 2000 in thousands	380	104	1.19	16,393	741
IT share 2007 (%)	9.28	7.86	0	42.86	741
E. Differenced covariates **Δ=12**					
AI publications	72.0	0	-15	6919	741
Log job advertisements	0.48	0.48	-2.15	3.08	741
IT job ad share (%)	8.53	8.34	-23.50	31.12	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 (32 CZs); the distance to the closest large CZ is the distance to the closest of the most populous 32 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					Mean	Median
	All commuting zones				No AK/HI	All	<1000 pubs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Seven-year differences							
Log distance to closest hotspot	-0.032***	-0.035***	-0.035***	-0.028***	-0.026***	-0.054***	-0.037***
(1000+ AI publications)	(0.005)	(0.006)	(0.006)	(0.007)	(0.006)	(0.008)	(0.006)
Observations	2,964	2,964	2,964	2,964	2,888	2,964	2,836
R-squared/Pseudo-R-squared	0.12	0.13	0.19	0.19	0.20	0.20	0.07
B. Twelve-year difference							
Log distance to closest hotspot	-0.081***	-0.098***	-0.053***	-0.041***	-0.053***	-0.079***	-0.058***
(1000+ AI publications)	(0.011)	(0.010)	(0.006)	(0.010)	(0.008)	(0.015)	(0.005)
Observations	741	741	741	741	722	741	612
R-squared/Pseudo-R-squared	0.17	0.21	0.30	0.30	0.31	0.37	0.21
AI publications through 2006 (any, level, square)	Yes	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	Yes
Change in log ads, IT share, log AI pubs	--	--	Yes	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 (32 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 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)	-0.015***	-0.014***	-0.011***	--	--	--
	(0.002)	(0.001)	(0.002)			
Log radius of circle enclosing 1000+ AI pubs	--	--	--	-0.038***	-0.037***	-0.040***
				(0.007)	(0.007)	(0.008)
Observations						
Pseudo R-squared	0.10	0.13	0.13	0.13	0.19	0.19
B. 12-year differences				2964		
Log distance to closest hotspot (1000+ AI pubs)	-0.032***	-0.016***	-0.013***	--	--	--
	(0.004)	(0.003)	(0.002)			
Log radius of circle enclosing 1000+ AI pubs	--	--	--	-0.103***	-0.047***	-0.052***
				(0.015)	(0.010)	(0.012)
Observations				741		
Pseudo R-squared	0.26	0.29	0.29	0.20	0.30	
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	--	Yes	Yes	--	Yes	Yes
Log distance closest large CZ	--	--	Yes	--	--	Yes
Log population within circle	--	--	--	--	--	Yes

Note: 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 (32 CZs); the distance to the closest large CZ is the distance to the closest of the most populous 32 CZs.

Table 4: Impact of distance to closest AI publications hotspot in different job advertisement samples

Job advertisements in underlying micro sample:	All	Valid industry	Missing industry	Valid firm name	Missing firm name	Ads placed by small firms
	(1)	(2)	(3)	(4)	(5)	(6)
A. 7-year differences						
Log distance to closest hotspot (1000+ AI publications)	-0.035***	-0.014***	-0.053***	-0.006*	-0.077***	-0.028***
	(0.006)	(0.004)	(0.008)	(0.003)	(0.009)	(0.009)
Observations	2,964	2,963	2,963	2,960	2,964	2906
R-squared	0.19	0.12	0.13	0.05	0.19	0.06
B. 12-year differences						
Log distance to closest hotspot (1000+ AI publications)	-0.053***	-0.027***	-0.099***	-0.019*	-0.170***	-0.012
	(0.006)	(0.0010)	(0.013)	(0.010)	(0.017)	(0.019)
Observations	741	740	740	739	741	697
R-squared	0.30	0.22	0.29	0.22	0.27	0.12

Notes: Median regressions for columns 1-5; OLS for column 6. Each column's dependent variable is a dummy for a job advertisement requiring AI, based on different underlying samples of job advertisements. Missing industry refers to missing NAICS 3. A small firm is one which posts 1-100 job vacancies in a given year (a valid firm name is required for the calculation). 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 AI job advertisement share by industry

	All job advertisements		Job advertisements with missing firm name	
	7-year differences (1)	12-year differences (2)	7-year differences (3)	12-year differences (4)
All	-0.0536***	-0.0785***	-0.1090***	-0.1813***
	(0.0082)	(0.0154)	(0.0100)	(0.0196)
Agriculture, Utilities, Mining, Construction, Manufacturing	-0.0012	-0.0004	-0.0000	0.0011
	(0.0021)	(0.0074)	(0.0004)	(0.0008)
Wholesale trade, Retail trade, Warehousing, transportation	0.0017	0.0019	-0.0005	-0.0000
	(0.0019)	(0.0026)	(0.0005)	(0.0002)
Information, Finance, Insurance, Real estate, Management	-0.0064	-0.0150**	-0.0002	-0.0075***
	(0.0036)	(0.0072)	(0.0023)	(0.0013)
Administrative and support services (incl. employment agencies)	-0.0021***	-0.0054***	-0.0061***	-0.0144***
	(0.0008)	(0.0015)	(0.0016)	(0.0030)
Education, Health	-0.0015	-0.0005	-0.0040***	-0.0043***
	(0.0032)	(0.0037)	(0.0010)	(0.0009)
Arts and recreation, Accommodation	-0.0000	-0.0003	0.0001	-0.0002**
	(0.0007)	(0.0006)	(0.0006)	(0.0001)
Other services, Public administration	0.0009	-0.0028	-0.0012	-0.0046
	(0.0022)	(0.0033)	(0.0009)	(0.0044)
Missing industry	-0.0449***	-0.0878***	-0.0969***	-0.2118***
	(0.0054)	(0.0086)	(0.0083)	(0.0179)

Notes: Each cell contains the coefficient on log distance to an AI publication hotspot (at least 1000 publications) from a different OLS regression with 2964 observations and full covariates (including year dummies in columns 1 and 3). The commuting zone-year observations are based on the full sample of job advertisements in columns 1 and 2, and on the subsample with missing firm name in columns 3 and 4. The dependent variable is a dummy for a job advertisement requiring AI in the specified industry. The NAICS 2 codes for each row are a) 11, 21-23, 31-33; b) 42, 44-45, 48-49; c) 51-55; d) 56; e) 61-62; f) 71-72; g) 81, 92. In columns 1 and 3, standard errors are clustered by commuting zone.

Table 6: Impact of distance from AI hotspot on AI job advertisement share by AI type

	All AI	Unspecified AI only	Image Processing	AI Tool	AI Application	AI R&D
	(1)	(2)	(3)	(4)	(5)	(6)
A. AI job ads with valid occupation (680,652 obs)						
Share AI type in AI job ads	100%	37.1%	12.4%	9.1%	19.1%	34.6%
Share computer scientist/mathematician	62.5%	70.0%	49.6%	80.0%	52.7%	66.3%
B. 7-year differences, median regression (2964 obs)						
Log distance to closest hotspot	-0.035***	-0.017***	0.004*	-0.003***	-0.002	-0.007***
(1000+ AI publications)	(0.006)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
R-squared	0.19	0.24	0.01	0.12	0.03	0.20
Median of dependent variable (pct points)	0.108	0.033	0.004	0	0.005	0.020
Effect of 10% greater distance as % of median	-3.2%	-5.2%	8.9%	--	-3.0%	-3.7%
C. 12-year diffs, median regression (741 obs)						
Log distance to closest hotspot	-0.053***	-0.024***	-0.001	-0.004**	-0.012***	-0.015***
(1000+ AI publications)	(0.006)	(0.003)	(0.004)	(0.002)	(0.003)	(0.003)
R-squared	0.30	0.34	0.01	0.22	0.14	0.24
Median of dependent variable (pct points)	0.249	0.090	0.025	0.016	0.024	0.062
Effect of 10% greater distance as % of median	-2.1%	-2.7%	-0.5%	-2.5%	-5.0%	-2.4%

Notes: Median regression except column 3, which is OLS. Each column's depending 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 general AI versus other types). Standard errors in parentheses; clustered by commuting zone in panel A, robust in panel B. *** 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

	Share (%)	AI required? (%)	Admin and support services? (%)	Sample of ads requiring AI: Occupation			
				Computer and math	Management	Architects engineers	Business finance
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Industry							
Agriculture, Utilities, Mining, Construction, Manufacturing	6.0	0.55	0	61.7	9.7	14.5	2.9
Wholesale trade, Retail trade, Warehousing, transportation	11.9	0.23	0	66.3	12.9	4.7	4.1
Information, Finance, Insurance, Real estate, Management	16.6	0.79	0	62.3	12.3	5.5	7.8
Administrative and support services (incl. employment, temp agencies)	3.5	0.25	100	67.1	6.1	5.3	3.3
Education, Health	15.4	0.21	0	26.1	9.2	2.6	1.7
Arts and recreation, Accommodation	6.1	0.06	0	46.0	11.7	2.2	3.8
Other services, Public administration	3.8	0.21	0	47.4	18.9	7.3	3.0
Missing industry	36.7	0.46	0	71.3	8.1	6.7	3.7
All	100.0	0.42	3.5	62.5	10.6	6.4	4.8
B. Size of firm posting vacancy							
Firm has 1-100 ads per year	12.5	0.38	3.3	62.9	9.8	5.9	4.3
Firm has 101-2000 ads per year	22.3	0.48	2.8	59.8	10.4	6.2	4.3
Firm has 2001-10,000 ads per year	16.0	0.41	2.4	56.4	12.4	7.7	4.5
Firm has more than 10,000 ads per year	16.5	0.52	2.2	59.5	14.9	4.6	7.6
Missing firm name	32.7	0.33	5.2	72.4	5.8	7.6	3.4
All	100.0	0.42	3.5	62.5	10.6	6.4	4.8

Notes: 2010-2019. 192,265,310 observations in columns 1-3; 680,652 observations in columns 4 -7 (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) 51-55; d) 56; e) 61-62; f) 71-72; g) 81, 92.

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

	Median regression					Mean	Median
	All commuting zones				No AK/HI	All	<1000 pubs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log distance to closest hotspot	-0.032***	-0.035***	-0.035***	-0.028***	-0.026***	-0.054***	-0.037***
(1000+ pubs)	(0.005)	(0.006)	(0.006)	(0.007)	(0.006)	(0.008)	(0.006)
Log distance to closest large CZ	--	--	--	-0.010	--	--	--
				(0.008)			
Any AI publication through 2006	0.029***	-0.004	0.000	-0.003	0.002	-0.005	0.001
	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.014)	(0.009)
AI publications through 2006/1000	0.188***	0.134***	0.048**	0.048**	0.031	0.096***	0.041
	(0.023)	(0.027)	(0.021)	(0.023)	(0.024)	(0.029)	(0.090)
AI publications through 2006/1000^2	-0.019***	-0.013***	-0.011***	-0.012***	-0.009**	-0.019***	-0.051
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.143)
Log job advertisements 2007	--	-0.002	-0.001	-0.002	-0.001	-0.008	-0.003
		(0.005)	(0.005)	(0.005)	(0.005)	(0.007)	(0.005)
Log population 2000	--	0.013***	0.009	0.010*	0.007	-0.002	0.011**
		(0.005)	(0.005)	(0.005)	(0.006)	(0.009)	(0.005)
IT share in advertisements 2007 (%)	--	0.004***	0.005***	0.005***	0.006***	0.007***	0.005***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Log distance to other CZs (average)	--	0.061***	0.056***	0.065***	0.115***	0.109***	0.054***
		(0.019)	(0.016)	(0.016)	(0.020)	(0.025)	(0.015)
Log distance to closest CZ	--	0.007	0.006	0.004	0.003	0.002	0.007
		(0.009)	(0.009)	(0.009)	(0.008)	(0.010)	(0.008)
Change in log advertisements (7-year)	--	--	0.014**	0.014**	0.014**	0.008	0.009
			(0.007)	(0.006)	(0.007)	(0.013)	(0.006)
Change in IT share (%) x 1000 (7-year)	--	--	0.004***	0.004***	0.004***	0.008***	0.004***
			(0.001)	(0.001)	(0.001)	(0.002)	(0.001)
Change in AI publications (7-year)	--	--	0.337***	0.355***	0.347***	0.361***	0.783***
			(0.007)	(0.044)	(0.046)	(0.040)	(0.192)
Observations	2,964	2,964	2,964	2,964	2,888	2,964	2,836
R-squared/Pseudo-R-squared	0.12	0.13	0.19	0.19	0.20	0.20	0.07

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 (32 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. *** p<0.01, ** p<0.05, * p<0.1

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

	Median regression					Mean	Median
	All commuting zones				No AK/HI	All	<1000 pubs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log distance to closest hotspot	-0.081***	-0.098***	-0.053***	-0.041***	-0.053***	-0.079***	-0.058***
(1000+ AI publications)	(0.011)	(0.010)	(0.006)	(0.010)	(0.008)	(0.015)	(0.005)
Log distance to closest large CZ	--	--	--	-0.011	--	--	--
				(0.011)			
Any AI publication through 2006	0.083***	0.017	0.003	-0.000	0.000	-0.011	-0.016
	(0.018)	(0.018)	(0.014)	(0.016)	(0.016)	(0.028)	(0.018)
AI publications through 2006/1000	0.255***	0.168***	0.047	0.051	0.053	0.057	0.864
	(0.041)	(0.020)	(0.049)	(0.052)	(0.048)	(0.047)	(1.381)
AI publications through 2006/1000^2	-0.017***	-0.008***	-0.019	-0.020	-0.019	-0.020**	-11.907
	(0.006)	(0.003)	(0.018)	(0.019)	(0.018)	(0.008)	(19.130)
Log job advertisements 2007	--	0.012	0.005	0.007	0.006	-0.011	0.007
		(0.010)	(0.016)	(0.017)	(0.016)	(0.027)	(0.017)
Log population 2000	--	0.018	0.006	0.003	0.005	0.014	0.013
		(0.012)	(0.017)	(0.019)	(0.019)	(0.020)	(0.019)
IT share in job advertisements 2007	--	0.004*	0.021***	0.020***	0.021***	0.021***	0.020***
		(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)
Log distance to other CZs (average)	--	0.148***	0.048**	0.051**	0.018	0.131**	0.083***
		(0.034)	(0.022)	(0.025)	(0.035)	(0.057)	(0.023)
Log distance to closest CZ	--	0.015	0.019*	0.015	0.003	0.012	0.004
		(0.018)	(0.011)	(0.011)	(0.012)	(0.018)	(0.015)
Change in log job advertisements 2007-2019	--	--	0.005	0.010	0.003	-0.014	0.010
			(0.017)	(0.019)	(0.018)	(0.028)	(0.020)
Change in IT share x 1000, 2007-2019	--	--	0.023***	0.024***	0.025***	0.022***	0.022***
			(0.002)	(0.002)	(0.002)	(0.003)	(0.002)
Change in AI publications, 2007-2019	--	--	0.286**	0.286**	0.276**	0.320***	0.499
			(0.122)	(0.125)	(0.123)	(0.049)	(0.715)
Observations	741	741	741	741	722	741	612
R-squared/Pseudo-R-squared	0.17	0.21	0.30	0.30	0.31	0.37	0.21

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 (32 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, eleven-year differences

	All commuting zones				No AK/HI	<1000 pubs
	(1)	(2)	(3)	(4)	(5)	(6)
Log distance to closest hotspot	-0.193***	-0.096***	-0.101***	-0.065*	-0.084***	-0.101***
(1000+ AI publications)	(0.023)	(0.024)	(0.023)	(0.034)	(0.024)	(0.023)
Log distance to closest large CZ	--	--	--	-0.057	--	--
				(0.044)		
Any AI publication through 2006	0.179***	0.033	0.017	0.012	0.015	0.017
	(0.028)	(0.031)	(0.030)	(0.031)	(0.030)	(0.030)
AI publications through 2006/1000	0.499*	0.174	0.369	0.413	0.237	0.369
	(0.271)	(0.314)	(0.275)	(0.268)	(0.284)	(0.275)
AI publications through 2006/1000^2	-0.579	-0.282	-1.003	-1.224*	-0.660	-1.003
	(0.494)	(0.534)	(0.651)	(0.651)	(0.668)	(0.651)
Log job advertisements 2007	--	0.067**	0.298***	0.302***	0.292***	0.298***
		(0.029)	(0.062)	(0.061)	(0.063)	(0.062)
Log population 2000	--	0.150***	-0.050	-0.056	-0.041	-0.050
		(0.031)	(0.056)	(0.056)	(0.057)	(0.056)
IT advertisements 2007/1000	--	-0.828***	-0.756***	-0.765***	-0.736***	-0.756***
		(0.213)	(0.213)	(0.217)	(0.211)	(0.213)
IT advertisements 2007^2	--	35.166**	39.459***	38.785***	39.125***	39.459***
		(14.635)	(14.667)	(14.799)	(14.590)	(14.667)
Log distance to other CZs (average)	--	-0.072	-0.019	0.017	0.086	-0.019
		(0.083)	(0.081)	(0.087)	(0.103)	(0.081)
Log distance to closest CZ	--	0.037	0.007	0.008	0.016	0.007
		(0.050)	(0.049)	(0.048)	(0.051)	(0.049)
Change in log job advertisements, 2007-2018	--	--	0.307***	0.312***	0.296***	0.307***
			(0.070)	(0.070)	(0.074)	(0.070)
Change in IT advertisements x 1000, 2007-2018	--	--	-0.070***	-0.067***	-0.071***	-0.070***
			(0.022)	(0.022)	(0.023)	(0.022)
Change in AI publications, 2007-2018	--	--	1.006	1.116	0.933	1.006
			(0.734)	(0.767)	(0.668)	(0.734)
Observations	401	401	401	401	389	401
R-squared	0.18	0.38	0.41	0.42	0.38	0.41

Notes: The dependent variable is the eleven-year difference (2007-2018) in whether a commuting zone (CZ) had any AI job advertisement. Estimation is with linear probability. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1