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

November 7, 2023

*We thank Bledi Taska and the company formerly known as Burning Glass Technologies for access to data and Erich Denk formerly of the Technology and Policy Research Initiative at Boston University for extensive data work on the Burning Glass Technologies files. We are grateful for comments from Tarek Alexander Hassan, Flavio Calvino, Aureo de Paula, Sabrina Genz, Britta Glennon, Bill Kerr, John Landon–Lane, Ruonan Xu and participants in the Australian OVERS seminar, the University of Sydney Microeconometrics and Public Policy Working Group seminar, the Technology and Policy Research Initiative seminar, the 2021 NBER Artificial Intelligence Conference and the 2023 Artificial Intelligence and the Economy conference for helpful comments on earlier versions. 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 (research papers and patents) 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 3–5% of median growth. Distance from a hotspot plays no role if a commuting zone is itself a hotspot, but distance is a greater barrier the greater a hotspot's share of publications that are patents rather than research papers. Analysis by occupation, industry and AI type suggests that the type of job posting for which distance is a barrier is jobs adapting AI for use in a new setting. We do not find convincing evidence for an effect of distance on the adoption of AI, perhaps because there is as yet little adoption.

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

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

Such considerations may seem unimportant in the face of technological progress including the telephone, modern means of transportation, email, texting, the worldwide web and video conferencing. These are likely to have reduced the role of distance in both knowledge sourcing and especially sustaining collaboration, though they may have had less impact on initiations of collaborations. Indeed, initiations of collaborations appear sensitive to even small changes in distance: Catalini (2007) finds that existing collaborations persisted after the 1997–2014 shuffling of research laboratory locations on a Paris

[1] Esposito (2023); Catalini (2017). The World Intellectual Property Organization (2019) discusses the creation of contacts and networks in an international context.

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

university campus, while collaboration between newly proximate laboratories increased greatly. Berger and Prawitz (2023) show the construction of the Swedish rail network spurred invention by spurring collaboration between newly connected rural areas. A more general empirical literature supports the hypothesis that distance is a barrier to the diffusion of inventive activity to potential inventors.[3]

A related literature examines how the adoption of technology, often across countries, is affected by the proximity of other adopters. One hypothesis is that it is advantageous for a potential adopter of a technology to be proximate to an earlier adopter because this makes adoption less risky: the later adopter could discuss adoption with the early adopter, observe the early adopter's methods and outcomes, and poach the early adopter's experienced workers. Another hypothesis is that firms could learn about distant technology through trade or their region's receiving direct investment, and distance is a barrier to trade and direct investment. The empirical adoption literature confirmed that distance is a barrier to the diffusion of adoption[4], but finds the barrier to be lower for multiestablishment or multinational firms, which presumably have internal communication channels and coordination.[5]

Our paper seeks to examine whether distance constitutes a barrier between technology production (innovation) and technology adoption or adaptation, focusing on the technology of artifical intelligence (AI). We choose to examine AI in part because the rapid growth in AI research papers and patents began only recently, allowing an examination of its geographic diffusion from early in the process. It is also of particular interest because it is potentially important for future economic growth.[6] Because AI is still immature, with few off–the–shelf applications yet available, we seek evidence for the effect of distance on

[3] 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.

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

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

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

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

Prior theoretical and empirical work suggest that an industry adopting an innovation tends to be either established in the location of the innovation, or, in the case of a mature industry, moves to the location of the innovation (Duranton 2007; Kerr (2010); Zucker, Darby and Brewer 1998). However, the role of distance has rarely been studied explicitly. 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 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 as scientific papers 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.[7]

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

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

[7] 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. Acemoğlu, Autor, Hazell and Restrepo (2021) examine the growth of AI job advertisements in the Burning Glass Technologies data and Babina et al. (forthcoming) combine resume information with Burning Glass data, but these papers do not consider geography.

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. We assume that companies are able to fill the vacancies they post, and interpret a negative effect of distance as a barrier to hiring AI workers. Distance could be a barrier to the hiring of AI workers by companies already operating in distant commuting zones, or to the establishment in distant commuting zones of companies anticipating requiring AI workers. That U.S. firms mentioning AI on their website tend to be young (51% less than five years old) suggests the latter mechanism is likely to be important.[8]

A null effect of distance 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, approximately where the effect size is largest, 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.

On the other hand, we find that distance to a hotspot with 2000 or more publications has only a small negative effect on AI job postings. The AI inventions and development in these thirteen large hotspots may be sufficiently prominent that even distant commuting zones have access to their benefits. The prominence could be due to a mixture of

[8] Dernis et al. (2023). See also Acemoğlu et al. (2022) for related statistics.

publications in more prominent journals; non–hotspot researchers, managers or owners monitoring the activity of large hotspot firms and researchers regardless of how the results are disseminated; media exposure; and large hotspots having geographically wider personal networks as the numerous students of AI researchers and developers fan out to take jobs.

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 innovation hotspot, the distance to the nearest other hotspot is almost irrelevant: with sufficient local AI, the marginal contribution of AI innovation elsewhere falls to zero. We also find that distance is a greater barrier the greater a hotspot's share of publications that are patents rather than research papers (journal articles and conference proceedings). Nevertheless, distance represents a barrier even when the hotspot has a low share of patents among the publications.

To distinguish among job vacancies reflecting innovation, adaptation and innovation, we perform analysis by type of AI and by occupation and industry. The evidence fits most closely with the hypothesis that distance is a barrier to the adaptation of AI for use in a new setting, rather than to adoption of AI. Distance is a barrier to the posting of jobs for AI computer and mathematical occupation workers, who could be engaged in innovation or adaptation, but more specifically of developers of AI software applications, likely to be engaged in adapting AI for adoption. Further, while there is suggestive evidence that distance is a barrier to searching for AI workers in finance and insurance industries, which could reflect adaptation or adoption, it is no barrier to searching for AI workers in business and insurance (or management) occupations, who would be adopting AI.

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.[9] 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

[9] 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.

aggregate papers and patents into 741 commuting zones for each year.[10] Each author is thus the source of potential spillovers, whether in the same or a different location from his or her co-authors. While we refer to the commuting zones' publications, these are really author–publication pairs.

We use these data to designate certain commuting zones as innovation hotspots, based on the cumulative number of AI papers+patents (a sum we refer to as publications) through 2006, the year before our study period. We assume that it is the total rather than per capita number of publications that matter for spillovers to other locations, and experiment with different absolute thresholds.

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 U.S. 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 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.[11]

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,

[10] 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.

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

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 [12] 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, almost 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. (2021). An apparently simple way to distinguish advertisements for an AI innovator from those for an AI adopter is to divide the specific AI skills required in the job advertisements into categories reflecting the distinction. However, this has proved difficult to do, not least because a large share of advertisements requiring AI skills simply require either "Artificial Intelligence" or "Machine Learning" (ML) skills, with no further detail specified. The detailed skills, for their part, are difficult to categorize. We therefore examine three mutually exclusive categories: unspecified AI skills only (nothing beyond AI or ML mentioned); image processing (frequently requested for health occupations), whether requested along with other AI skills

[12] Burning Glass Technologies, personal communication.

or not; and the remaining AI skills or skill combinations.[13]

We aggregate the total job advertisements, AI job advertisements and IT job advertisements to the commuting zone–year level using the county of the employer, and calculate the share of the commuting zone's total advertisements which are AI or IT advertisements in each year. Finally, we merge the data with the publication data. Our dependent variable is based on the share of advertisements that require AI skills, so that small commuting zones may experience as large an effect of distance as large commuting zones.[14]

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 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

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

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

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 whose closest innovation hotspot is commuting zone h: $\Delta^k AI^s_{cht} = \frac{AI\ job\ ads}{All\ job\ ads}_{cht} - \frac{AI\ job\ ads}{All\ job\ ads}_{ch\,t-k}$. 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 also have investigated changes in the probability of a commuting zone having any AI job advertisement. However, we find that the standard errors on the coefficient of interest in these regressions are large enough to prohibit inference, and we do not report these results.

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}:

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

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 posting of AI job vacancies. If it is zero, however, this could reflect either that distance is no barrier, or that distance is such a barrier that only innovation in the commuting zone affects a commuting zone's AI job vacancies. The conceptual randomization is the distribution of pre–2007 AI publications among commuting zones.

The first set of additional controls capturesthe commuting zone's own AI innovation prior to 2007: a quadratic in the commuting zone's own cumulative AI papers through 2006 (quadratic rather than log due to the presence of zeros), $AI\ Papers_{ct^*}$; a dummy for any such paper, $AI\ Papers > 0_{ct^*}$; a linear term in the commuting zone's own cumulative AI patents through 2006, $AI\ Patents_{ct^*}$; and a dummy for any such patent, $AI\ Patents > 0_{c,t^*}$. The next set of controls are for other initial conditions: 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 ($IT^s_{c,2007}$).

The third set of controls contains other distances that could be confounders of distance to the closest hotspot: 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. We also control for the log of the number of pre–2007 publications (papers+patents) in the closest hotspot, $log(AI\,Pubs_{ht^*}^{Hot})$, and the 2000 population in the closest hotspot, $log(Pop_{ht^*}^{Hot})$.[15]

Finally, we control for contemporaneous changes in some key variables: the number of the commuting zone's own AI papers and patents, $\Delta^k AI\,Papers_{ct}$ and $\Delta^k AI\,Patents_{ct-2}$; log job advertisements, $\Delta^k log(All\,job\,ads_{ct}))$; and the change in the IT job advertisements' share in all advertisements, $\Delta^k IT_{ct}^s$. This could constitute overcontrolling: some or all of these could be the result of growth in AI job advertisements, rather than the cause, and their inclusion could bias $\hat{\sigma}$ upward toward zero. On the other hand, 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, even if we are able to measure growth in AI job advertisements reflecting adaptation and adoption only, it is plausible that there is a positive correlation between unobserved factors affecting innovation and adaptation/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 adaptation/adoption). Controlling for changes in the number of the commuting zone's own AI papers and patents (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

[15] The last two covariates vary only by h. Were we interested in their standard errors we would have to adjust for this low variation, but since we are not, we do not.

all the covariates in the equation above.[16]

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 pre–2007 AI publications.

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 2000 population within the circle with radius $log(Pop_c^{RN})$, 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 and which allows for population to be controlled in a way free from from collinearity problems.[17]

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.[18] OLS point estimates of σ are somewhat more negative than median regression estimates, with larger standard errors. We report three–year differences, the only difference length to use all the years' data except 2007 as the first year of the difference, and as a specification check we also report the longer seven–year differences 2007–2014, 2010–2017, 2011–2018, and 2012–2019. We cluster standard errors by commuting zone.[19]

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

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

[17] We have recently become aware of the method of Sävje (2023), which overcomes problems with a spatial spillover method we rejected. We may adopt it in a future version of the paper.

[18] 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 the closest hotspot are correlated.

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

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 represents an acceleration in employment.[20] 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. With this dependent variable, all regressions using any difference length use all years of data.[21]

All these regressions establish whether distance to an AI hotspot based on papers+patents is a barrier to the growth of AI job advertisements. We would like to know whether the barrier is to spillovers from AI papers or AI papers. The high correlation between the distance to an AI paper hotspot and the distance to an AI patent hotspot makes it infeasible to define separate paper and patent hotspots and run a horse race between the distance to the closest hotspot of either type. Nor is the radius approach well–suited to distinguishing between the relative importance of the two AI publication types. Instead, we retain a hotspot definition based on thresholds of AI publications (papers + patents), and include an interaction between the distance to the closest hotspot and the share of the hotspot's publications that is patents. This has the advantage of using all the variation in the ratio of papers and patents in a hotspot and sidesteps the multicollinearity problem.

Further analysis is designed to distinguish whether the barrier is to the adaptation or adoption of AI, or merely to additional innovation in AI. For this purpose, we examine different AI types separately, and we investigate the role of distance by occupation and industry, using as the outcomes the number of AI job advertisements in a particular occupation or industry, divided by job advertisements in that occupation or industry. We use OLS for some of these regressions, where the median of the dependent variable is zero.[22]

[20] Dernis et al. (2023) make this point about Burning Glass data.

[21] However, differences involving 2007 are too short due to the missing 2008 and 2009 data.

[22] 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.

3 Descriptive statistics

The national time–series of AI job advertisements is plotted in Figure 1. The increase over time from 9000 in 2007 to 190,000 in 2018 (and 135,000 in the first six months of 2019) far outstrips the 50% increase in the total number of job advertisements online. Figure 2 shows that the AI jobs share in all advertisements rises from 0.07 percent to 0.75 percent, rising linearly with a break in the slope in 2016, when growth increases (black squares, left scale). The IT jobs share is much higher (red circles, right scale) and evolves quite differently, rising from 2007–2012, then changing non–monotonically but ending lower in 2019 than 2012. The shares of the three types of AI job advertisements in all job advertisements are shown in Figure 3: unspecified AI and "other AI" have grown equally quickly over the whole period and begin and end at the same shares, but "other AI" grew faster in the 2007–2012 period. Image processing, on the other hand, has not grown over the period. The Figure 4 maps indicating commuting zones' AI job advertisement shares show how the fraction of commuting zones with no AI job advertisement (white) 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 the mean increase in the three–year AI job advertisement increase is 0.06 percentage point, while the median increase is lower at 0.03 percentage point (first row). The minimum value of -2.46 percentage points and the maximum value of 4.70 percentage points confirm the existence of the outliers mentioned above: such large changes are caused by very small changes in the number of AI job advertisements in commuting zones with few job advertisements. The mean seven–year increase is 0.14 percentage point and the median increase is 0.09 percentage point (second row). The lower panels of Table 1 shows the means of key covariates, including those based on AI publications (panel D). The mean number of pre–2007 patents (5.5) is much lower than the mean number of pre–2007 papers (150); the median for both is zero.

The national time–series for AI papers and patents from 1950 onwards (a few publications are pre–1950) are shown in Figure 5. Papers (times the number of authors), plotted in black squares using the left scale, increased from 6 in 1950, to 11,620 in 2007, to 49,484 in 2018 and to 65,411 in the first half of 2019. Patents rose from 1 in 1950 to 469 in 2007 to 1525 in 2017; the numbers fall in 2018 and 2019, reflecting the dating using patent applications rather than patent awards. 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 papers and patents for each of the 31 commuting zones satisfying this requirement: commuting zones are ordered by papers+patents. 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 (including Pittsburgh) and California is Seattle in ninth place. 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). New York, San Jose and Seattle stand out as having a large number of patents, while Pittsburgh stands out among the top ten as having a small number of patents. The five AI "pioneer locations" designated by Bloom et al. (2021) are all in our top nine AI hotspots, though notably do not include Los Angeles or Pittsburgh.[23] The map in Figure 7 shows the distribution of cumulative pre–2007 publications, while the four maps in Figure 8 show that there is very slow diffusion of publications 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 (the sum of papers

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

and patents), and then analyze the sensitivity to changes in the hotspot threshold using the preferred covariates. Next, we test whether the distance effects are due to hotspots' patents or research papers. Third, we investigate how the role of distance varies by type of AI skill, occupation and industry, seeking to distinguish between the impact of distance on AI adoption, adaptation and innovation.

4.1 The effect of distance on AI vacancies

The effect of distance to the closest innovation hotspot on the change in AI job advertisements as a percent of all job advertisements, is presented in Table 2. In both panel A, for three–year differences, and panel B, for seven–year differences, the median regression coefficients on distance are statistically significantly negative in all columns, with the seven–year coefficient between 1.8 and 2.6 times the three–year coefficient. In the first column, the only controls are log distance to the closest hotspot, the log of its AI publications, and five controls for the commuting zone's AI papers and patents through 2006. The three–year coefficient of -0.011 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 (0.011)(0.1)=0.0011 percentage point. This is 3.1% of the median growth rate of 0.035 percentage point in Table 1, a modest effect; the equivalent number for the seven–year distance coefficient of -0.029 is also 3.1%.

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: -0.017 for three–year differences and -0.031 for seven–year differences. In column 3, we add controls for changes in log job advertisements, IT jobs' share, AI papers and AI patents, which leaves the coefficients on distance almost unchanged. The addition of the distance to the closest large commuting zone (one of the 31 most populous, since there are 31 AI publication hotspots) and the log of its population in column 4 renders the coefficient on distance slightly less negative, to -0.013 for three–year differences and -0.026 for seven–year differences. These are our preferred coefficients, corresponding to 3.7% and 2.8% of median growth respectively. The standard errors are higher in this specification than in the specification of column 3, due to the fairly high correlation of distance to closest hotspot and distance to closest large commuting zone.

We test whether our results are driven by the remote commuting zones of Alaska and Hawaii by dropping them from the estimation in column 5: this renders the coefficient of interest slightly less negative in each panel. On the other hand, using mean rather than median regression (column 6) makes the coefficients more negative, at -0.015 for three–year differences and -0.034 for seven–year differences. Results by AI type, occupation and industry in later tables which rely on OLS because the median AI job share is zero may thus slightly overstate the effect of distance.

The threshold of 1000 AI publications to designate a hotspot is arbitrary, so it is important to test the sensitivity of the distance effect to the threshold. We would expect that very low thresholds would lead to a finding of no effect of distance and we can use as a 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 hotspot publication or population–related distance control). If this is negative, a negative coefficient on distance to the closest hotspot could be picking up a spurious effect. Note that as the definition of the hotspot changes, so do three covariates in addition to the distance to the closest hotspot: the number of publications in the closest hotspot, the distance to the closest large commuting zone (potentially, as the pool of large commuting zones changes to contain the same number of commuting zones as the number of hotspots), and the population in the closest large commuting zone.

In Figure 9 panels A (three–year differences) and B (seven–year differences), we plot the point estimates for the median regression coefficients on log distance to the closest hotspot for different thresholds. With the exception of the hotspot threshold–specific covariates, the specifications behind the coefficients plotted with black circles are the same as in Table 2 columns 4–6 (i.e. full covariates), while the blue dashed curves are the same as in Table 2 column 3 (i.e. no controls related to large commuting zones). All four curves are U–shaped in the threshold range of 0–2000 publications, and equal zero at thresholds of 0 and 1 (some of the points for these two thresholds are indistinguishable). The coefficients are statistically significantly negative for thresholds in the range 300–1500 for the black curves and to a higher threshold for the blue curves. Although the coefficients along the curves based on specifications with and without large commuting zone controls are never statistically significantly different, the point estimates do differ somewhat in the threshold range 1500–2500: where the curve without the controls rises to a point estimate shy of zero at the 2000 threshold and then flattens, the curve with the controls rises to above zero at 200, then falls back below zero.

A curve sloping downward from the origin is consistent with the hypothesis that as a commuting zone's number of AI publications increases, proximity to it is increasingly beneficial. The upward–sloping part of the curve and the subsequent lack of any distance effect were not anticipated, but are consistent with the hypothesis that once hotspots become sufficiently large, their inventions and development are sufficiently prominent that even distant commuting zones have access to their benefits. The prominence could be due to a mixture of publication in more prominent journals; non–hotspot researchers, managers or owners monitoring the activity of large hotspot firms and researchers regardless of how the results are disseminated; media exposure; and large hotspots having geographically wider personal networks as the numerous students of AI researchers and developers fan out to take jobs.

In panels C and D of Figure 9, we plot the equivalent graphs for AI job publications cumulated from 2007, rather than contemporaneous AI job publications, and find similar U–shapes with similarly statistically significant coefficients for hotspot publication thresholds in the 300–1500 range. The negative point estimates in this range are smaller than in panels A and B, though not directly comparable. The results for three different specifications at the 1000 threshold are shown in Table 3 columns 1–3: coefficients are similar across the specifications.

Our second identification approach is to define the key distance covariate as the radius of the circle enclosing a certain number of pre–2007 AI publications. In columns 4–6 of Table 3, we present the median regression coefficients on the distance variable from three specifications for both three–year and seven–year differences: all are statistically significantly negative for our chosen hotspot threshold of 1000 publications. In the preferred specifications of column 6, the three–year difference coefficient of -0.016 indicates 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 0.0016 percentage point, or 4.6% of median growth; the corresponding number for the seven–year difference coefficient of -0.035 is 3.8%.

We test the sensitivity of the radius method to the hotspot threshold in Figure 10; as in Figure 9, we plot coefficients on distance to the closest hotspot with controls related to large commuting zones with black dots, and the coefficients on specifications without these controls with a dashed blue line. We observe U–shapes in both panels (for three and seven–year differences), though there are some differences compared to Figure 9. In the radius approach, there is no equivalent to the falsification test of a zero threshold, and the point estimate at a threshold of one publication is slightly negative rather than zero, though it is statistically insignificant from zero. Also, the effects at high thresholds are statistically significantly negative, rather than zero, and at high hotspot thresholds it is the blue curve based on specifications without large commuting zone controls that rises more. Considering all curves in Figures 9 and Figure 10, we conclude that once a commuting zone reaches 2000 or more pre–2007 publications, it appears to be able to influence other commuting zones with little interference from distance, but the role of distance may not disappear completely. Appendix Figure 1 shows that the results are similar when Alaska and Hawaii are dropped.

If the coefficient on the distance to the closest hotspot or on the radius encosing a given number of AI job advertisements indeed reflects barriers to geographic spillovers, it it very likely to be less negative for commuting zones which themselves have more AI publications: the more AI activity a commuting zone has, the less it needs its neighbors. Returning to the original identification strategy, we therefore include controls for the interaction of log distance to the closest hotspot with the commuting zone's own pre–2007 AI publications controls. Because there are three such interactions, we calculate the median regression effects of distance at different numbers of publications and present them in Table 4, columns 1 (three–year differences) and 2 (seven–year differences).[24]

By construction, the coefficient on the main distance to hotspot effect represents the distance effect for commuting zones with no pre–2007 AI publications of their own: it is statistically significantly negative in both columns. The three interaction coefficients are only jointly statistically significant for three–year differences, but the qualitative pattern is the same for three–year and seven–year differences, as shown in the bottom four rows: the coefficient on distance becomes less negative as a commuting zone's own AI publications rise from zero, eventually rising to a statistically insignificant estimate of close to zero by the time a commuting zone has 1000 pre–2007 AI publications of its own. The coefficient rises from -0.016 to -0.002 for three–year differences, and from -0.031 to -0.002 for seven–year differences. The results reinforce evidence that AI job posting grows faster in commuting zones that are near hotspots, which we assume means that hiring of AI workers does too. Next we turn to the question of whether this result is driven by AI papers or AI patents.

4.2 The relative effects of papers versus patents

We begin our examination of the relative importance of papers and patents by defining hotspot thresholds based on a commuting zone's number of patents only and running median regressions based on our Table 2 column 4 specification.[25] The coefficients are plotted in Figure 11 for three–year and seven–year differences, and the familiar pattern of Figure 9 emerges (if one ignores a rogue coefficient at the 110 patent threshold in panel B), along with statistically significantly negative coefficients for thresholds below 40 patents.

One difference compared to Figure 9 is that the curves slopes downward initially thanks mainly to the falsification threshold of zero patents, whose coefficient equals zero

[24] Recall that there are five pre–2007 publication main effects, but that we sum papers and patents for the purposes of interactions with distance to hotspot, due to collinearity issues.

[25] We control as before for the distance to the closest large commuting zone and the log of that commuting zone's population. As when the hotspot threshold was based on AI publications, for each AI patent threshold there are h hotspots, and we designate as large commuting zones the most populous h commuting zones.

as expected. Distance from a hotspot defined as having a single pre–2007 AI patent or more has a statistically significant negative coefficient in each panel. Because the number of AI patents is so small, we think it plausible that proximity to a commuting zone with a single pre–2007 AI patent could matter, whereas we would not find this plausible for a single pre–2007 AI paper. We do not show a figure with thresholds based on papers alone, because it is so similar to Figure 9.

This separate graph for patents cannot be the basis for disentangling papers and patents, however, since the number of AI papers and the number of AI patents in a commuting zone are highly correlated. Instead, we add as controls the share of the closest hotspot's AI publications which are patents ($\frac{AI\,patents}{AI\,patents+AI\,papers}$) and its interaction with the log distance to the closest innovation hotspot. The coefficients on this interaction are -0.095 (0.039) for three–year differences and -0.33 (0.11) for seven–year differences, indicating that the coefficient on distance from the closest hotspot is statistically significantly more negative the greater the share of patents among the hotspot's publications. An intuitive way of grasping the magnitude is to evaluate the distance effect at the 75th and 25th percentiles of the patents/publications ratio. These 75th and 25th percentile effects are -0.014 (0.003) and -0.010 (0.003) for three–year differences and -0.035 (0.007) and -0.020 (0.008) for seven–year differences: the 25th percentile effects are thus only 60–70% of the 75th percentile effects, but are still statistically significant.

We plot the 75th percentile and 25th percentile distance coefficients corresponding to various hotspot thresholds in Figure 12 panels A and B (note that as the threshold rises, the number of different values of the patents/publications ratio falls). All four curves have the U–shape of Figure 9, and both are zero at the thresholds of 1 and 2000. At the most negative part of the curves, the coefficient at the 25th percentile is as low as half that of the 75th percentile. The gap between the 25th and 75th percentiles of patents' share of publications is not generally statistically significant, but for thresholds in the range 100 to 2500 for three–year differences and 100 to 4000 for seven–year differences, the coefficient on the interaction term is statistically significantly negative (not pictured).

There is therefore a greater AI job vacancy penalty due to distance from a patent–

intensive hotspot than from a paper–intensive hotspot, but the distance penalty remains for a paper–intensive hotspot. Having established this, we turn to assessing in more detail which sorts of job advertisements lie behind distance as a barrier to the influence of an innovation hotspot, 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 research papers and patents. Acemoğlu et al. (2022) report that 2.2% of workers are employed at firms that produce AI and 12.6% of workers are employed at firms that use AI, but given likely within–firm differences across these categories in the share of workers requiring AI skills, the statistics are not very informative as to what share of our AI jobs advertisements are likely to be for workers producing AI.

4.3 The effect of distance on AI hiring by type of AI

We present the results for the effect of distance by type of AI in Table 5 and Figure 13, using median regressions except in the case of three–year differences in image processing (since the median of the dependent variable is zero). The first row of Table 5 panel A shows that the largest category of AI is the unspecified category, constituting 37% of AI advertisements (based on the advertisement micro–data), while image processing makes up 12% of AI advertisements and other AI the remaining 50%. The second row presents the share of AI advertisements in these categories advertising for a computer science or mathematics occupation, a proxy for a job likely to be in innovation or adaptation rather than adoption. The share is highest in the unspecified AI category (68%), compared to only 50% in the image processing category, with an intermediate 62% for the residual "other AI" category. This suggests that image processing could be viewed as closer to an application than the other categories.

The regression coefficients reported in Panel B (three–year differences) and panel C (seven–year differences) show that distance to the closest AI hotspot has no impact on growth the the image processing share of job advertisements (column 3), while the effects for unspecified AI (column 2) and other AI (column 4) are statistically significantly negative with similar magnitudes of about half the overall magnitude in column 1.[26]

Consistent with these results, Figure 13 shows that the coefficients for image processing fluctuate around zero and do not become gradually more negative as the hotspot threshold is raised from zero publications. Coefficients for unspecified AI and other AI, on the other hand, trace out the pattern familiar from Figure 9, with the coefficients in the threshold range approximately 200–1250 being statistically significantly negative (not pictured). Unlike the pattern for all AI types in Figure 9, however, coefficients become negative again at hotspot thresholds above 2000, sometimes statistically significantly so. Figure 13 also shows that the distance effect at the 2000 threshold is statistically significantly positive for other AI. This type of AI is therefore driving the positive (albeit statistically insigificant) point estimate at this threshold in Figure 9 panels A and B black dots; we have no explanation for this anomalous result, though we note that the point estimate is not positive without controlling for distance to the closest large commuting zone (panels A and B blue dashed line).

The results are evidence that distance from a hotspot does not hinder growth in image processing, an application, but does hinder growth in AI job advertisements for unspecified AI and for other AI. The results also suggest that for the types of AI for which distance to a hotspot is a barrier, the effect of distance does not fade away completely when the threshold for a hotspot is set very high. This is consistent with the results from the radius identification approach.

4.4 The effect of distance by occupation

A different way of testing whether distance from an AI hotspot affects growth in AI innovation or adoption is to examine the effect by occupation, creating occupation–specific microdata samples based on 2–digit SOC, and aggregating each to the commuting zone–year level. If an AI job advertisement is for a computer scientist, the job is likely to

[26] If the estimation were all OLS, the columns 2–4 coefficients would sum to the column 1 coefficient. Although all regressions except one are median regressions, the columns 2–4 coefficients nevertheless sum approximately to the column 1 coefficient.

be innovation or adaptation rather than adoption. An advertisement for an engineer is likely to be for adaptation or adoption, and an advertisement for a business, finance or sales occupation is clearly for adoption. An AI job advertisement for a manager is less clear–cut, since it could be a position managing computer scientists engaged in innovation.

In Table 6 we analyze these occupations, which have the highest AI job share levels and growth, for the hotspot AI publication threshold of 1000 (three–year differences in columns 1–3 and seven–year differences in columns 4–6). The number of observations are in columns 1 and 4: some commuting zone–years do not have job advertisements from every occupation. The means of the dependent variables are in columns 2 and 5: computer and mathematical occupations have by far the fastest growth in the share of AI job advertisements, with an average three–year growth of 0.59 percentage point and an average seven–year growth of 1.22 percentage points. Architectural and engineering occupations grow much less quickly (0.18 and 0.36 percentage point), followed by management (0.07 and 0.25 percentage point respectively) and business and finance occupations (0.07 and 0.15 percentage point respectively). Sales occupations have a higher share of AI job advertisements than all occupations pooled, but the growth is similar to that of the pooled occupations (0.01 and 0.02 percentage point for three–year and seven–year differences respectively).

The OLS coefficients on the log distance to the closest AI hotspot are reported in columns 3 and 6. The statistically significantly negative coefficients for computer and mathematical occupations (second row) are ten times larger than the coefficients for all occupations (first row), at -0.133 (three–year differences) and -0.280 (seven–year differences). These coefficients indicate percentage point effects; the percent effects for computer scientists/mathematicians, however, are similar to those for pooled occupations – a 10% increase in the distance to the closest AI hotspot reduces mean AI job advertisement growth by 2.2% (three–year differences) and 2.3% (seven–year differences). The statistically significantly negative coefficients for architecture and engineering (-0.050 and -0.112 percentage point respectively) are about three times the coefficients for all occupations, but also have comparable percent effects. In slight contrast, the statistically significantly negative (at the 10% significance level or more) coefficients for sales are small, but imply larger percent effects; the coefficient is sensitive to a few outlier commuting zones, however. Conversely, there is no negative effect of distance for management occupations, business and finance occupations or the pooled "other" occupations.

We pursue the occupations analysis by assessing the pattern of distance coefficients by hotspot threshold in Figures 14 (computer and mathematical, engineering and sales occupations) and 15 (management and business/finance occupations). The coefficients are based on median regressions for the computer and mathematical occupations (whose median is greater than zero) and on OLS regressions for engineering and sales occupations. In panel A (three–year differences, red circles), the coefficients for all computer and mathematical occupations initially become more negative as the hotspot threshold rises, before stablizing then slowly becoming less negative. The coefficients from a threshold of about 200 are statistically significantly negative. Seven–year coefficients are similar, except that the coefficient abruptly becomes zero at a threshold just before 2000.[27]

We also plot the median regression coefficients for one of the two largest detailed computer and mathematical occupations, developers of software applications (15-1134). In the micro–data, this group represents 28% of all advertisements for computer and mathematical occupations, and 35% of advertisements for computer and mathematical occupations requiring AI skills.[28] Any distance effect for this group is likely to reflect a barrier to adaptation: its practitioners are developing applications to be adopted by others. Both panels show that the pattern and magnitudes for this (yellow triangles) are fairly similar qualitatively and quantitatively to those of all computer and mathematical occupations. A difference is that in panel A, the coefficient for software applications developers gradually becomes less negative after a threshold of about 750, and becomes zero by a threshold of 2000. In this panel, the coefficient on distance is statistically significantly negative (at the 5 or 10 percent level) between the thresholds of approxmiately 400 and

[27] The jump in the coefficient occurs when the group of hotspots shrinks to exclude Austin, TX, the last hotspot in the middle of the country.

[28] The other large detailed occupation is "other" computer, which also accounts for a large share of AI job advertisements. We do not examine this as the title is too vague for the analysis to be informative.

1250. In panel B, the coefficient is statistically significantly negative between thresholds 200 and 1250; as for all computer and mathematical observations, the coefficient suddenly rises at a threshold just below 2000, and even becomes quite positive (and statistically significant at the 10% level), as seen for "other AI" in the previous figure.

The coefficients on distance for and architect/engineering occupations are similar to those for all computer and mathematical occupations and for developers of software applications, though statistically significantly negative at thresholds above 2000. The curves for sales in Figure 14 at first sight hew closely to the Figure 9 pattern, albeit with less negative coefficients. However, they are lacking the expected downward slope at low thresholds, since the apparent decline relies on the placebo hotspot threshold of zero. The haphazard patterns in Figure 15 for management and business/finance occupations confirm that distance to an AI hotspot appears to play no role in the AI job growth in management and business/finance occupations.

We have also analyzed other detailed occupations, but generally find that although there is always at least a hint of the familiar shape when we plot the coefficients on distance against threshold, that the coefficients are generally statistically insignificant. In particular, the coefficients for either occupation 15-111 (computer and information research scientists) alone, or for this occupation pooled with mathematicians, statisticians and operations researchers, which have a large share of advertisements requiring AI and are likely to be innovating in AI, the coefficients are imprecisely estimated, but possibly more negative that those reported in the graphs (results not presented).

The various patterns by occupation are consistent with the hypothesis that distance to the closest hotspot is a barrier to adapting but not adopting AI, with workers in computer and mathematical occupations (particularly developers of software applications) engaged in adaptation rather than innovation, and engineers engaged in adaptation rather than adoption. Sales may be an exception to the finding that no effect of distance is found for the clearly adoptive occupations business/finance and management.

4.5 The effect of distance by industry

Another tack for isolating job advertisements for AI innovation versus those for AI adoption or adaptation is to examine the effects by industry. To do so, we run commuting zone–year level regressions based on different underlying microdata samples, namely a sample for each of several NAICS 2 code groups. Of particular interest is the finance and insurance industry (NAICS 2 code 52), which has had the highest percentage point growth in the share of AI job advertisements, a growth which reflects adoption, or at the very least adaptation of AI, rather than production of new AI. Conversely, the information sector (NAICS 2 code 51) presumably includes much of the AI innovation performed at companies.

Results of OLS regressions for the 1000 publication threshold are reported in Table 7, based on the Table 2 column 4 specification, the left three columns for three–year differences and the right three for seven–year differences. The number of observations are in columns 1 and 4: some commuting zone–years do not have job advertisements from every industry. Columns 2 and 5 contain the means of the dependent variables: the finance and insurance industry had an average three–year increase in the AI job advertisement share of 0.131 percentage point and a seven–year increase of 0.276 percentage point, the highest of any sector. There is a similarly rapid increase in the AI share in the category of job advertisements for which the industry is missing: 0.121 percentage point and 0.249 percentage point. The third fastest growth is in the industry category grouping real estate, professional and scientific services and administration (0.115 and 0.244 percentage point respectively); while the fourth sector with rapid AI growth is unsurprisingly the information sector (0.104 and 0.208 percentage point respectively). Other sectors have much slower AI job share growth.

The OLS coefficients on the distance to the closest AI hotspot are reported in columns 3 and 6: the coefficients are negative for the four groups mentioned above, with magnitudes about double that for all industries (reported in the first row), and statistically significant except for the information sector (and one coefficient significant at the 10% level). No other industry categories yield negative or statistically significant coefficients. The most negative and most precisely estimated coefficients are for the missing industry group: -0.051 compared to -0.015 for all industries for three–year differences, and -0.109 compared to -0.034 for seven–year differences. The seven–year coefficient of -0.065 for the finance and insurance sector implies that a 10% increase in distance reduces the growth in AI's share of job advertisements by 0.0065, or 2.3% of the 0.276 percentage point average increase.

In Figure 16, we present the plots of the coefficient on distance with a range of AI hotspot thresholds for the finance and insurance sector (panel A for three–year differences and panel B for seven–year differences). These graphs display a qualitative pattern very similar to that for all industries in Figure 9, albeit with few statistically significant coefficients and point estimates for thresholds in the range 100–1000 that are much more negative (note the very different scales). Although only suggestive because of the large standard errors, the results are very interesting given the absence of any distance effect for business and finance occupations found in the previous section. This implies that distance is an obstacle to growth in AI vacancies in finance and insurance not because it is an obstacle to growth in vacancies for workers who would be adopting AI, but because it is an obstacle to growth in vacancies for workers adapting AI to this sector.

The lower two graphs of Figure 16 plot the coefficients on distance for the information sector. The pattern by hotspot threshold for the information sector is not consistent with theory, with few negative distance coefficients and no statistically significantly negative coefficients. This points to distance being no barrier to growth in AI innovation employment. We have tried to test this better by using the NAICS 6 codes to focus on narrower industries, ignoring issues related to the large share of missing values in the NAICS 6 variable. We select samples of advertisements posted by universities and by firms performing research and development, both of which presumably advertise AI jobs in innovation, rather than adoption or adaptation. Unfortunately, the standard errors on log distance to the closest hotspot are too large for conclusions to be drawn (the results are not shown).

Our industry results indicate that distance from an AI hotspot can be a modest barrier to the adaptation of AI. The lack of support for a role of distance for the information sector is some evidence against distance being a barrier for innovation, though hardly conclusive; this paper is not seeking to test that hypothesis.

The very strong role of distance from a hotspot and large increase in the share of AI jobs for advertisements with missing industry are worthy of further attention. We first confirm that when the distance coefficients are plotted against hotspot threshold, the familiar pattern of Figure 9 emerges (figure not reported). 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. 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. In results not reported, we confirm that the very strong role of distance is found in regressions based on an underlying microdata sample of job advertisements missing both industry and firm name, but not in those based on a sample of job advertisements missing a firm name but with a valid industry. This shows that the very strong role of distance is not straightforwardly associated with Burning Glass Technology's having recorded a missing firm name for advertisements posted by an employment agency. Rather, it is associated with employment agency advertisements for which Burning Glass could not establish an industry (Burning Glass Technologies codes industry based primarily on employer name; Burning Glass Technologies 2019). The occupation, skill and year mixes of the two missing firm name categories differ, probably reflecting the degree to which industry may be inferred based on this information, but not in a way that is enlightening in connection with the role of distance to closest hotspot.

4.6 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 Table 5). 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 the interaction between the count and the log distance to the closest hotspot. This is in part because the main effect of the count and the interaction term are highly correlated.

We have also investigated whether the effect of distance varies by firm size. 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. As for the analysis of firms operating in multiple locations, the analysis is necessarily based on job advertisements for which a firm name is provided. We now sum up and reconcile the results of the different sections.

5 Conclusion

Our results indicate that online vacancies for jobs requiring Artificial Intelligence (AI) skills grow more somewhat more slowly in U.S. commuting zones farther from AI innovation hotspots. We assume that companies are able to fill the vacancies they post, and equate this effect with distance being a barrier to hiring AI workers, either due to companies already operating in a distant commuting zone hiring fewer AI workers or due to companies anticipating requiring AI workers being less likely to choose to operate in a distant commuting zone.

For a hotspot definition of least 1000 AI papers or publications prior to 2007, we find that a 10% greater distance from the closest hotspot (about a standard deviation) reduces a commuting zone's growth in AI jobs' share of job advertisements by 3–5% of median growth. Our investigation of the sensitivity to the definition of a hotspot shows that if the threshold of pre–2007 AI publications is set low, proximity to a hotspot is unimportant for other commuting zones, a sign that the moniker of hotspot is inappropriate. The effect of distance to the closest hotspot becomes more negative as the threshold is raised from one publication towards around 500 publications, although the change is not always statistically significant.

However, the effect of distance to the closest hotspot becomes less negative as the threshold is raised above around 1000 publications, and even rises to zero in some analysis. This is consistent with the hypothesis that once hotspots become sufficiently large, their inventions and development are sufficiently prominent that even distant commuting zones have access to their benefits. The prominence could be due to a mixture of publication in more prominent journals; non–hotspot researchers, managers or owners monitoring the activity of large hotspot firms and researchers regardless of how the results are disseminated; media exposure; and large hotspots having geographically wider personal networks as the numerous students of AI researchers and developers fan out to take jobs.

The more pre–2007 AI innovation a commuting zone has itself, the less important is its distance from an AI hotspot: if a commuting zone is itself a hotspot as defined by a threshold of 1000 pre–2007 publications, its subsequent growth in AI vacancies is unaffected by its proximity to other hotspots. We find that distance from an innovation hotspot slows growth in AI vacancies more if the hotspot's mix of innovation is more oriented to AI patents than AI scientific papers.

Our analysis by AI type, occupation and industry leads us to conclude that much of the negative effect of distance from a hotspot is an impediment to adapting AI to new situations, rather than to adopting AI, possibly because AI is not yet mature enough for adoption to be common. We find that distance to the closest hotspot slows AI vacancy growth in the finance and insurance industry, which we interpret as the presence as a barrier to adaptation because distance does not slow AI growth for business and finance occupations (nor for managers). Even at the end of the sample period, a majority of AI job vacancies were for computer scientists, and we find that the effect of distance from the closest hotspot on growth in AI vacancies for computer scientists is ten times the level for pooled occupations (in percentage point terms; in percent terms it is similar). We interpret this as at least in part an effect on adaptation, and not merely on innovation, because we find an effect of the same magnitude for developers of software applications. We find no effect of distance on the growth in job vacancies requiring image processing skills, skills possibly reflecting AI adoption, nor on growth in job vacancies in the information sector, likely to reflect AI innovation.

References

Acemoğlu, Daron, Gary W. Anderson, David N. Beede, Cathy Buffington, Eric E. Childress, Emin Dinlersoz, Lucia S. Foster, Nathan Goldschlag, John C. Haltiwanger, Zachary Kroff, Pascual Restrepo and Nikolas Zolas. 2022. "Automation and the Workforce: A Firm–Level View From The 2019 Annual Business Survey". NBER Working Paper 30,659.

Daron Acemoğlu, David Autor, Jonathon Hazell and Pascual Restrepo. 2022. "Artificial Intelligence and Jobs: Evidence from Online Vacancies". *Journal of Labor Economics*, 40(S1): S293–S340.

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

Acemoğlu, 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. 2021. "The Demand for AI Skills in the Labor Market". *Labour Economics*, 71

Almeida, Paul and Bruce Kogut. 1999. "Localization of Knowledge and the Mobility of Engineers in Regional Networks". *Management Science*, 45(7): 905–1024.

Andersson, Roland, John M. Quigley and Mats Wilhelmsson. 2009. "Urbanization, productivity, and innovation: Evidence from investment in higher education". *Journal of Urban Economics*, 66: 2–15.

Babina, Tatian, Anastassia Fedyk, Alex He and James Hodson. Forthcoming. "Investments in Artificial Intelligence Technologies and Changes in Workforce Composition". In Susanto Basu, Lucy Eldridge, John Haltiwanger and Erich Strassner eds. *Technology, Productivity and Economic Growth*, Chicago: University of Chicago Press.

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

Berger, Thor and Erik Prawitz. 2023 "Collaboration and Connectivity: Historical Evidence from Patent Records". CEPR Discussion Paper 18,031.

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

Bernstein, Shai, Xavier Giroud and Richard R. Townsend. 2015. "The Impact of Venture Capital Monitoring". *Journal of Finance*, 71(4): 1591–1622.

Blit, Joel and Mikko Packalen. 2018. "A Machine Learning Analysis of the Geographic Localization of Knowledge Flows". 2018. University of Waterloo working paper.

Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner and Ahmed Tahoun. 2021. "The Diffusion of Disruptive Technologies". NBER Working Paper 28999.

Branstetter, Lee, Britta Glennon and J. Bradford Jensen. 2018. "Knowledge Transfer Abroad: The Role of U.S. Inventors within Global R&D Networks". NBER Working Paper 24453.

Burning Glass Technologies. 2019. "Mapping the Genome of Jobs: The Burning Glass skills taxonomy". https://www.burning-glass.com/wp-content/uploads/2019/09/Burning_Glass_Skills_Taxonomy.pdf, accessed 2 June 2021.

Carnevale, Anthony P., Tamara Jayasunder, and Dmitri Repnikov. 2014. "Understanding online job ads data". Georgetown University, Center on Education and the Workforc, Technical Report.

Catalini, Christian. 2017. "Microgeography and the Direction of Inventive Activity". *Management Science*, 64(9): 4348–4364.

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.

Dernis, Hélène, Flavio Calvino, Laurent Moussiegt, Daisuke Nawa, Lea Samek and Mariagrazia Squicciarini. "Identifying artifical intelligence actors using online data". OECD Science, Technology and Industry Working Papers 2023/01.

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

Duranton, Gilles. 2007. "Urban Evolutions: The Fast, the Slow, and the Still". *American Economic Review*, 97(1): 197–221.

Esposito, Christopher J. 2023. "The geography of breakthrough invention in the United States over the 20th centry". *Research Policy*, 104810.

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

Kerr, William R. 2010. "Breakthrough inventions and migrating clusters of innovation". *Journal of Urban Economics*, 67:46–60.

Little, Jane Sneddon and Robert K. Triest. 1996. "Technology Diffusion in U.S. Manufacturing: The Geographic Dimension". *Proceedings of the Boston Federal Reserve Bank conference on Technology and Growth*.

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.

Sävje, Fredrik. 2023. "Causal inference with misspecified exposure mappings: separating definitions and assumptions". *Biometrika*,

Thompson, Peter and Melanie Fox–Kean. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment". *American Economic Review*, 95(1): 450–460.

World Intellectual Property Organization. 2019. *World Intellectual Property Report 2019. The Geography of Innovation: Local Hotspots, Global Networks.* https://www.wipo.int/edocs/pubdocs/en/wipo_ pub_ 944_2019.pdf , accessed 29 September, 2023.

Zucker, Lynne G., Michael R. Darby and Marilynn B. Brewer. 1998. "Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises". *American Economic Review*, 88(1): 290–306.

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

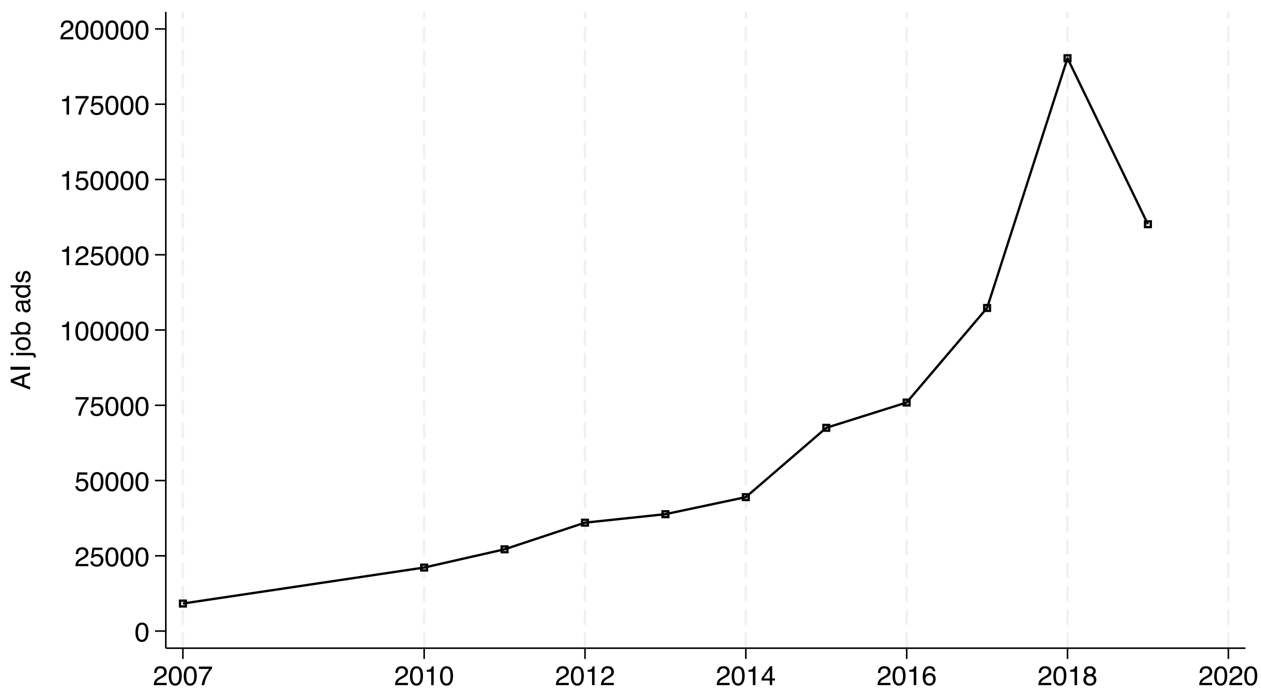

Notes: Data for 2019 are for January–July. Data for 2008 and 2009 are not available.
Source: Burning Glass Technologies.

Figure 2: AI share of job ads (%)

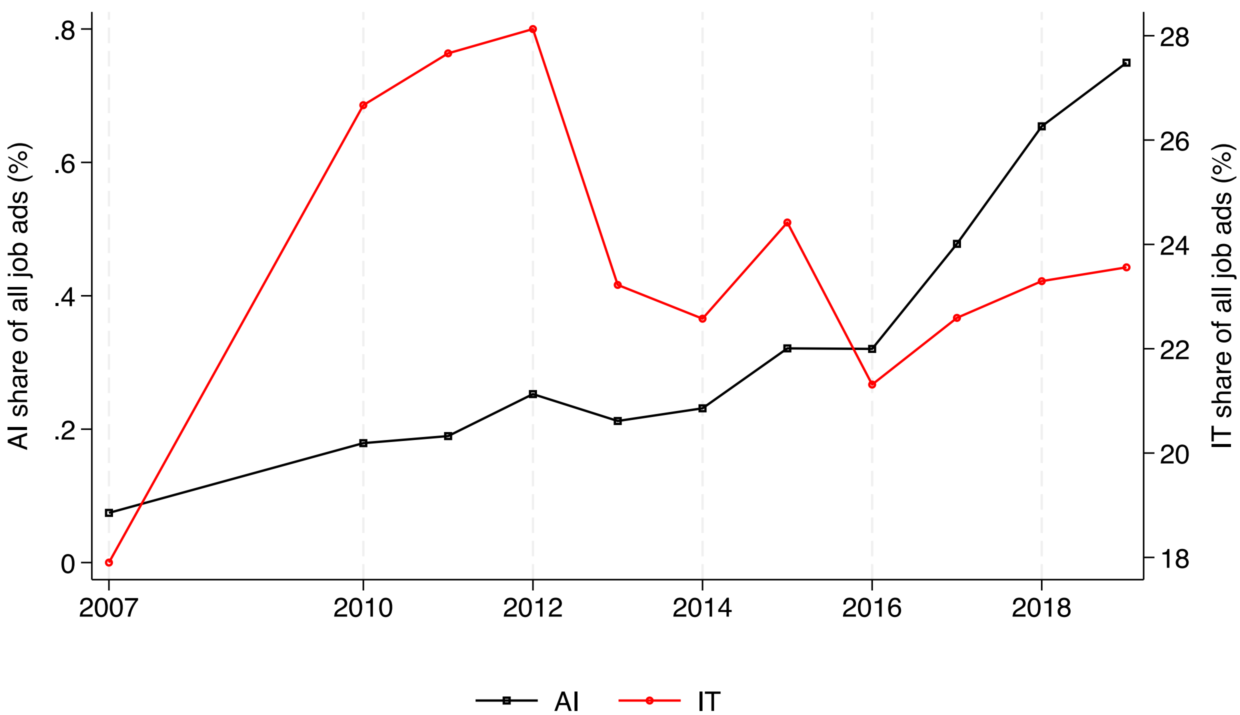

Source: Burning Glass Technologies.

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

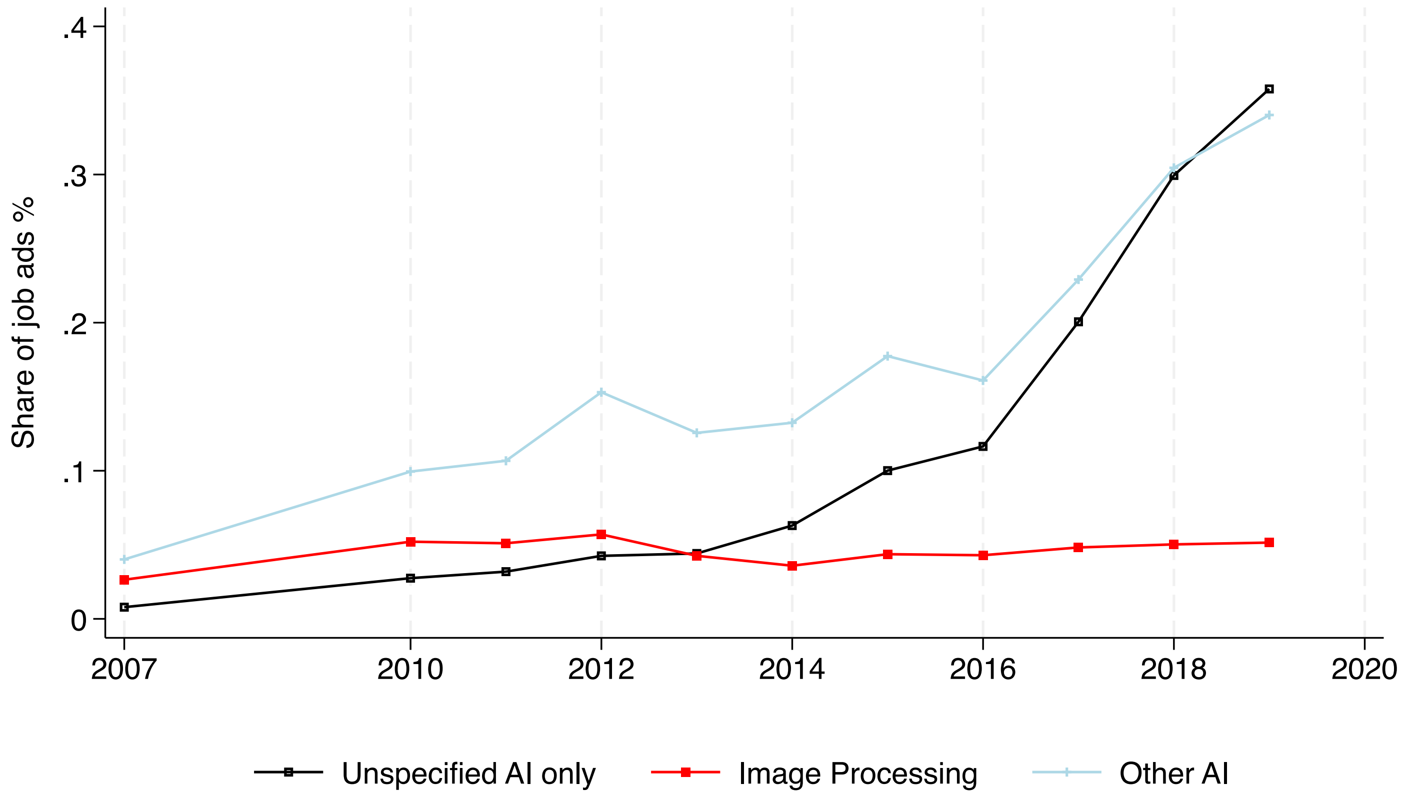

Note: Unspecified AI job advertisements require “Artificial Intelligence” and/or “Machine Learning” skills with no further detail given. Image processing AI job advertisements require image processing as one of the required skills. The “other” category is defined so the three categories are mutually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

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

(a) 2007

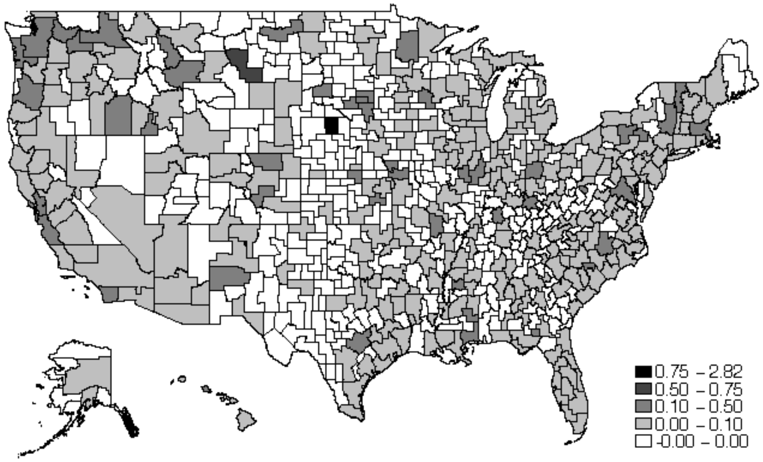

(b) 2010

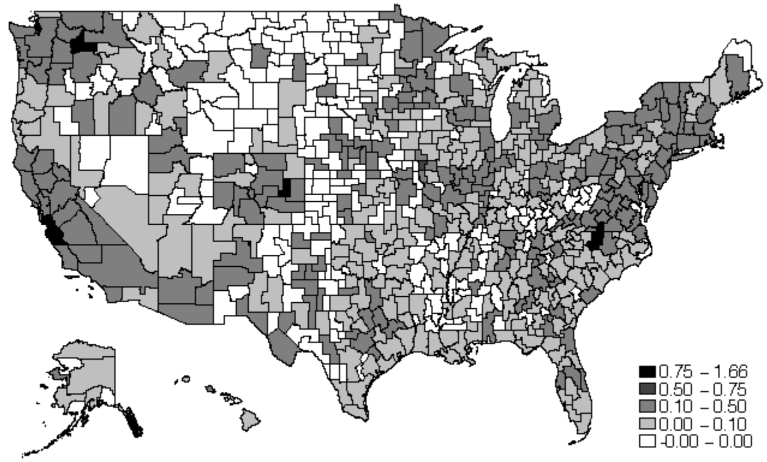

(c) 2014

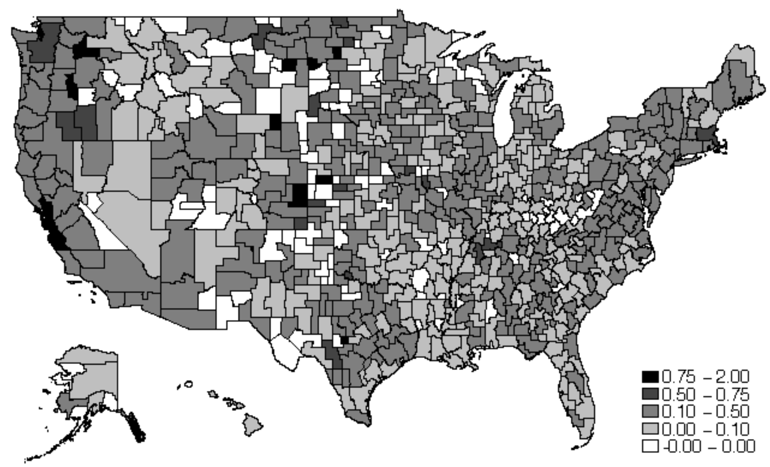

(d) 2018

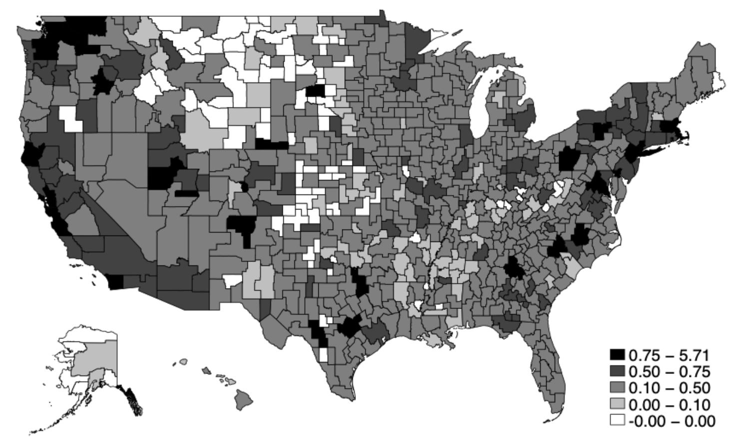

Figure 5: AI publications 1950-2019

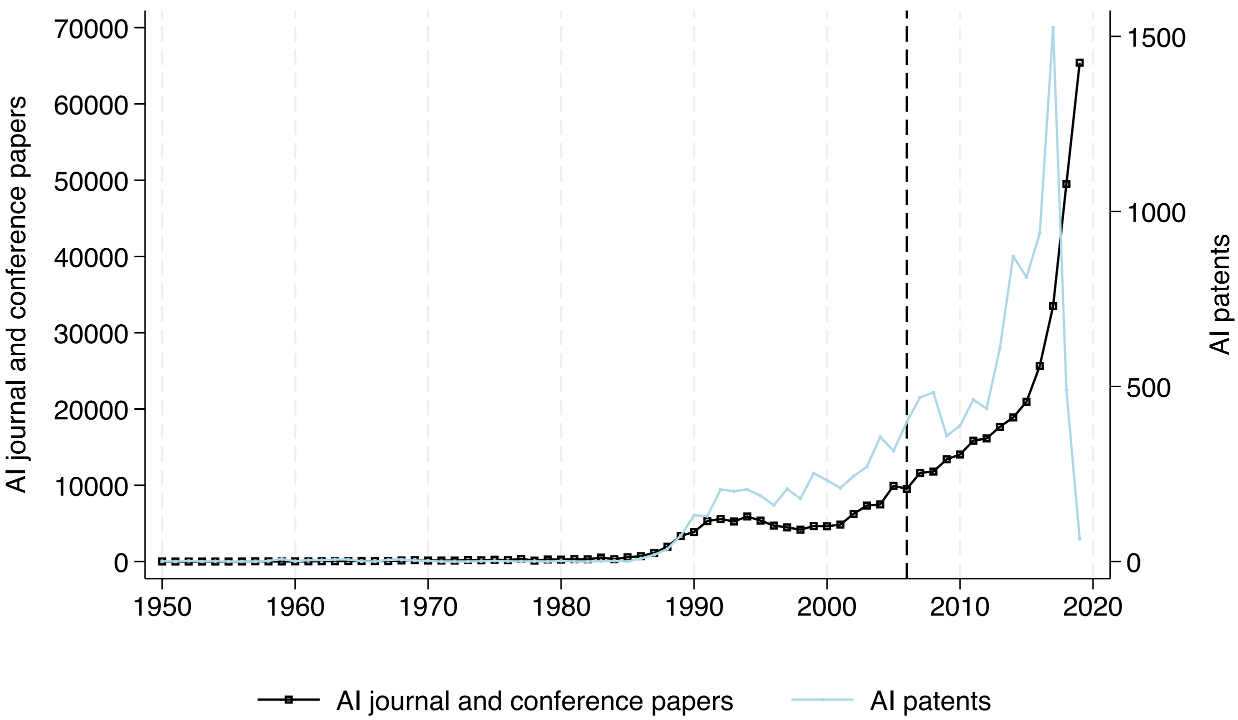

Source: Authors' dataset.

Figure 6: Innovation hotspots' AI publications through 2006

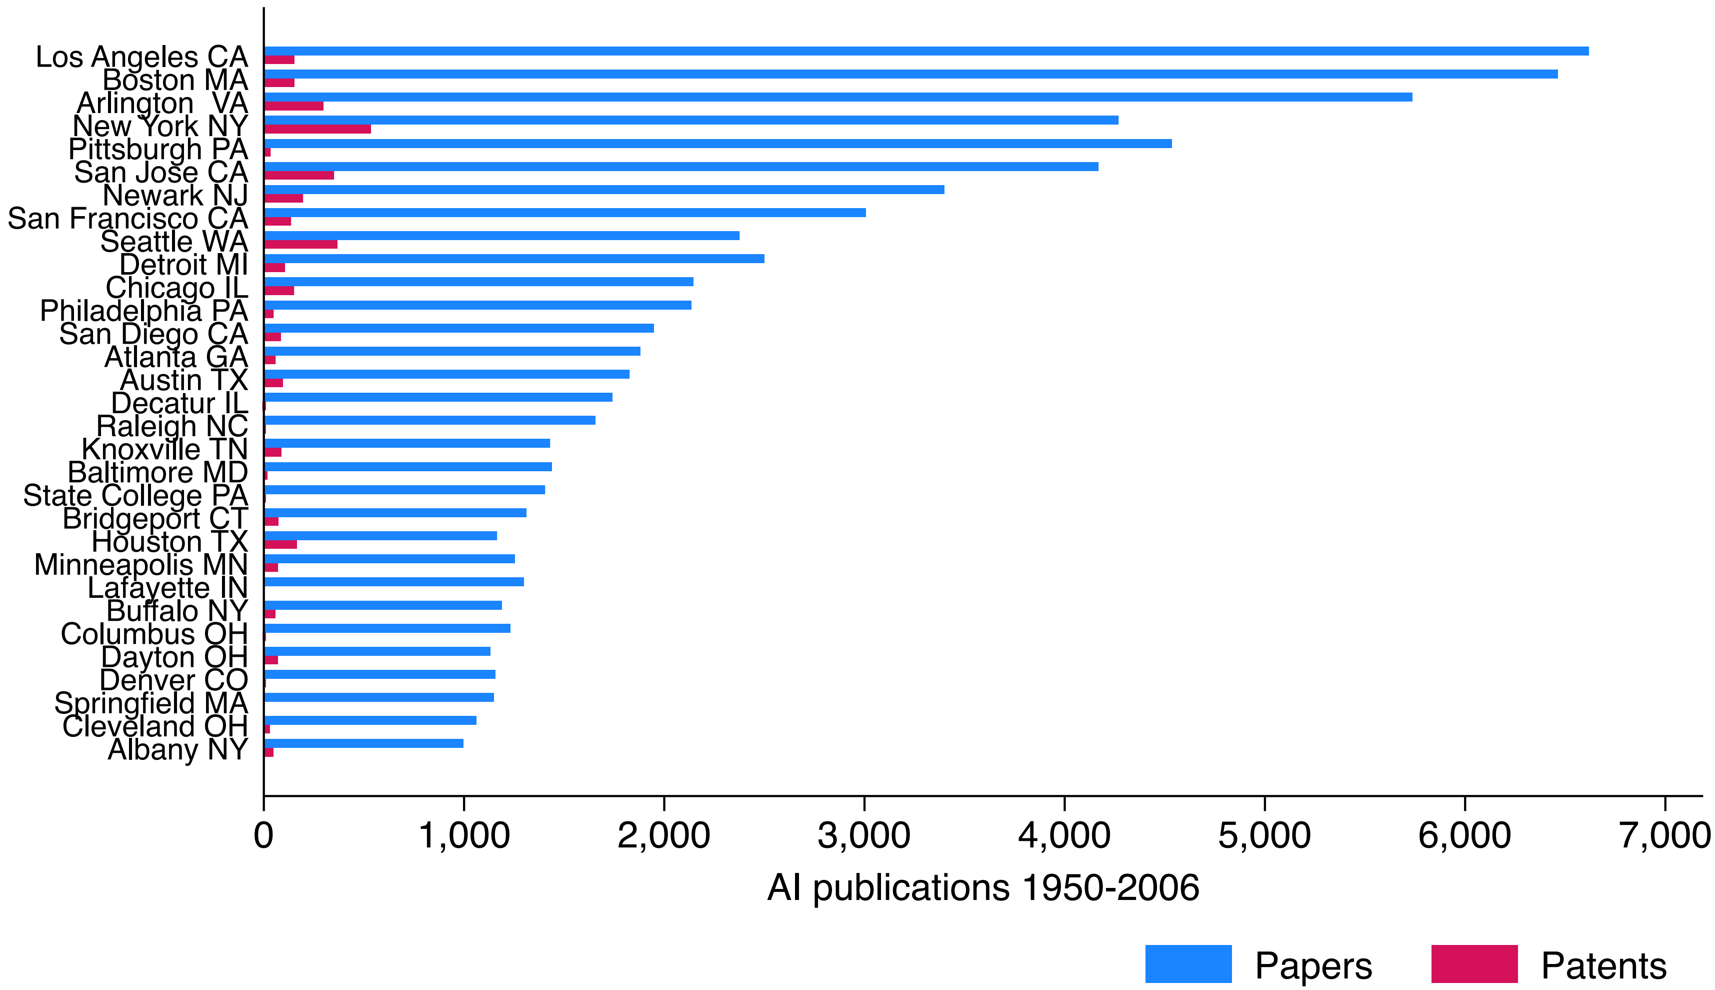

Note: The definition of a hotspot is a commuting zone with at least 1000 AI publications (papers+patents) through 2006.

Figure 7: Commuting zones' AI publications through 2006

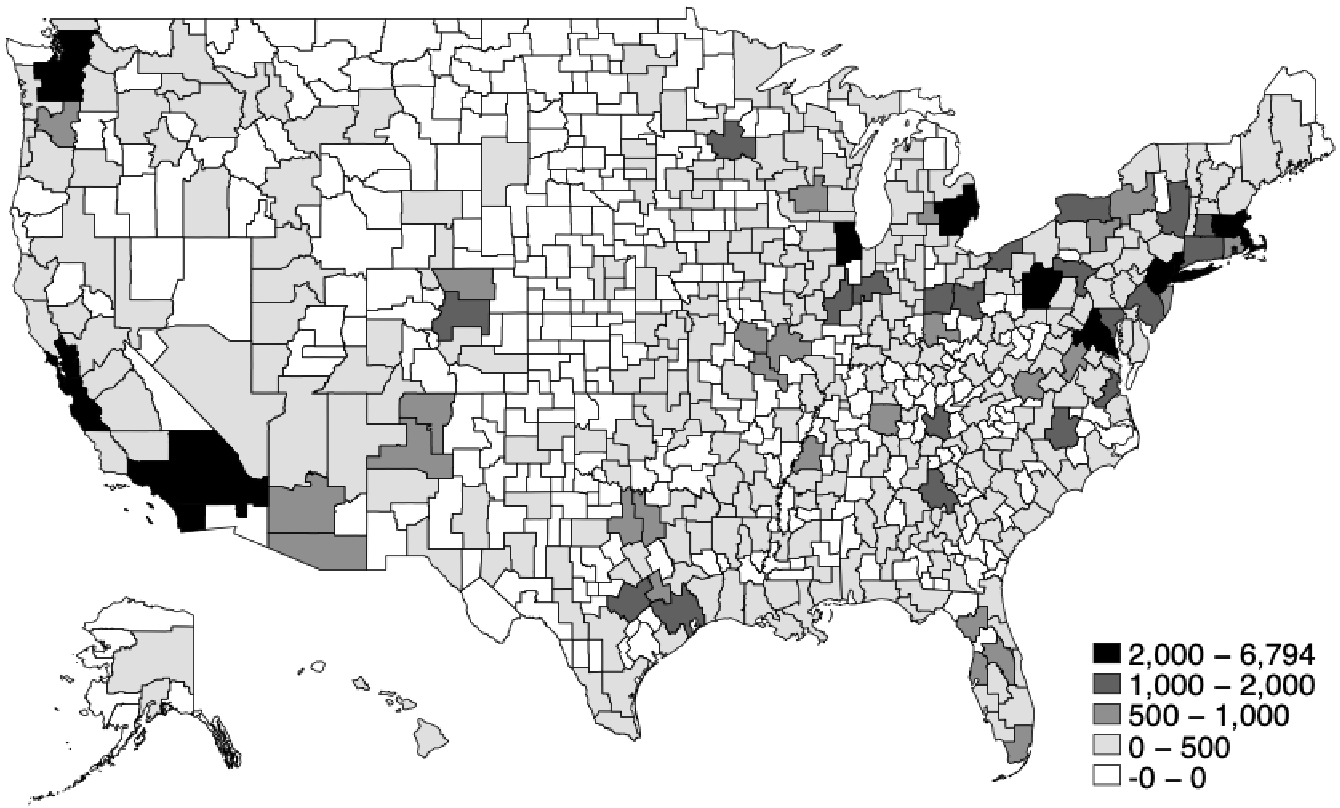

Figure 8: Commuting zones' AI publications in given year

(a) 2007: 240 commuting zones with any publication

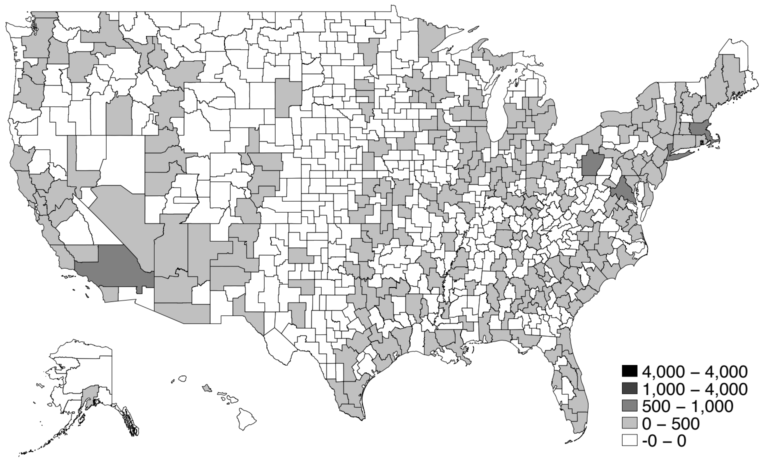

(b) 2010: 245 commuting zones with any publication

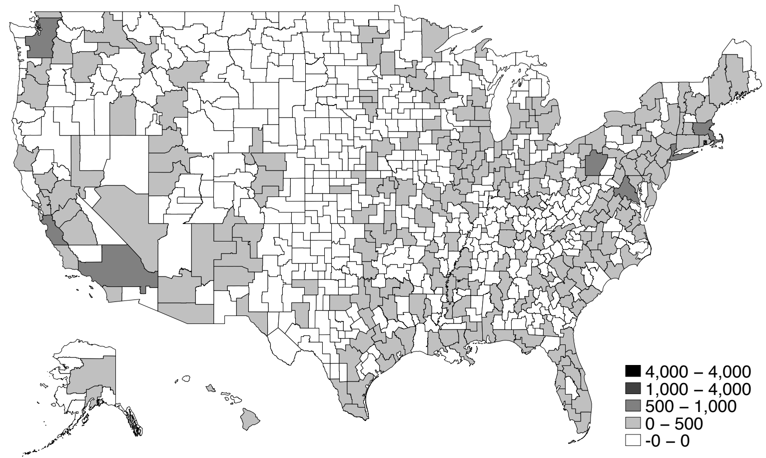

(c) 2014: 252 commuting zones with any publication

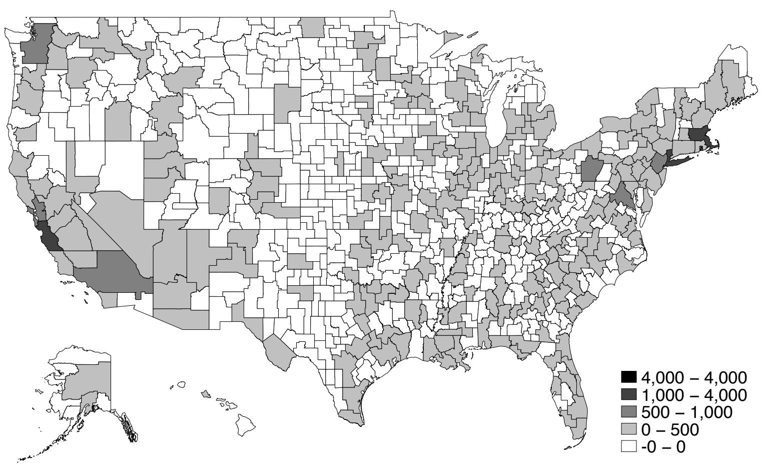

(d) 2018: 282 commuting zones with any publication

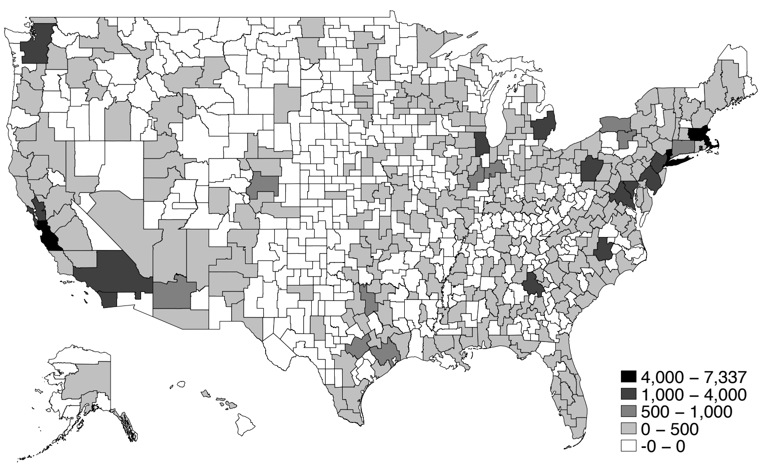

Figure 9: Effect of distance with contemporary and cumulative AI job advertisements

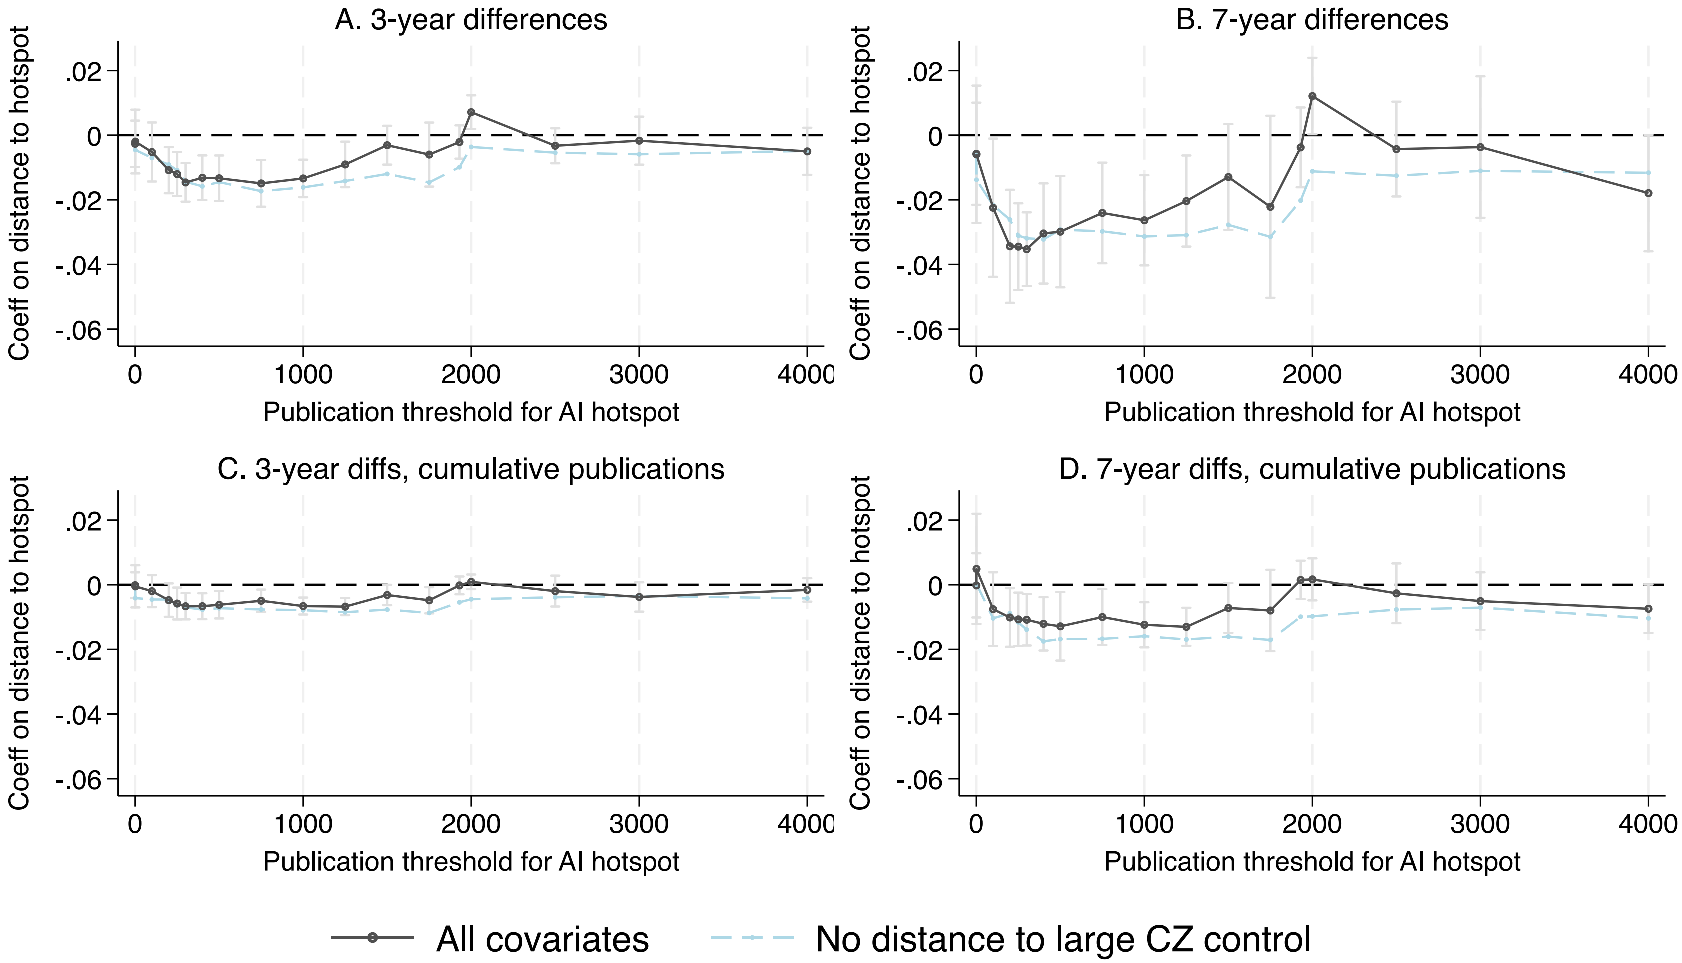

Note: The black circles are the coefficients on log distance to the nearest hotspot from median regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6; the 95% confidence intervals are also shown. The number of observations in each regression is 2964 in panels A and C and 5928 in panels B and D. The dashed blue line represents the coefficients from regressions which do not control for the log distance to the closest large commuting zone and the log of its population, but are otherwise the same. The x–axis is the number of pre–2007 AI publications designated as the threshold for a commuting zone to be an AI hotspot.

Figure 10: Effect of radius with and without population control

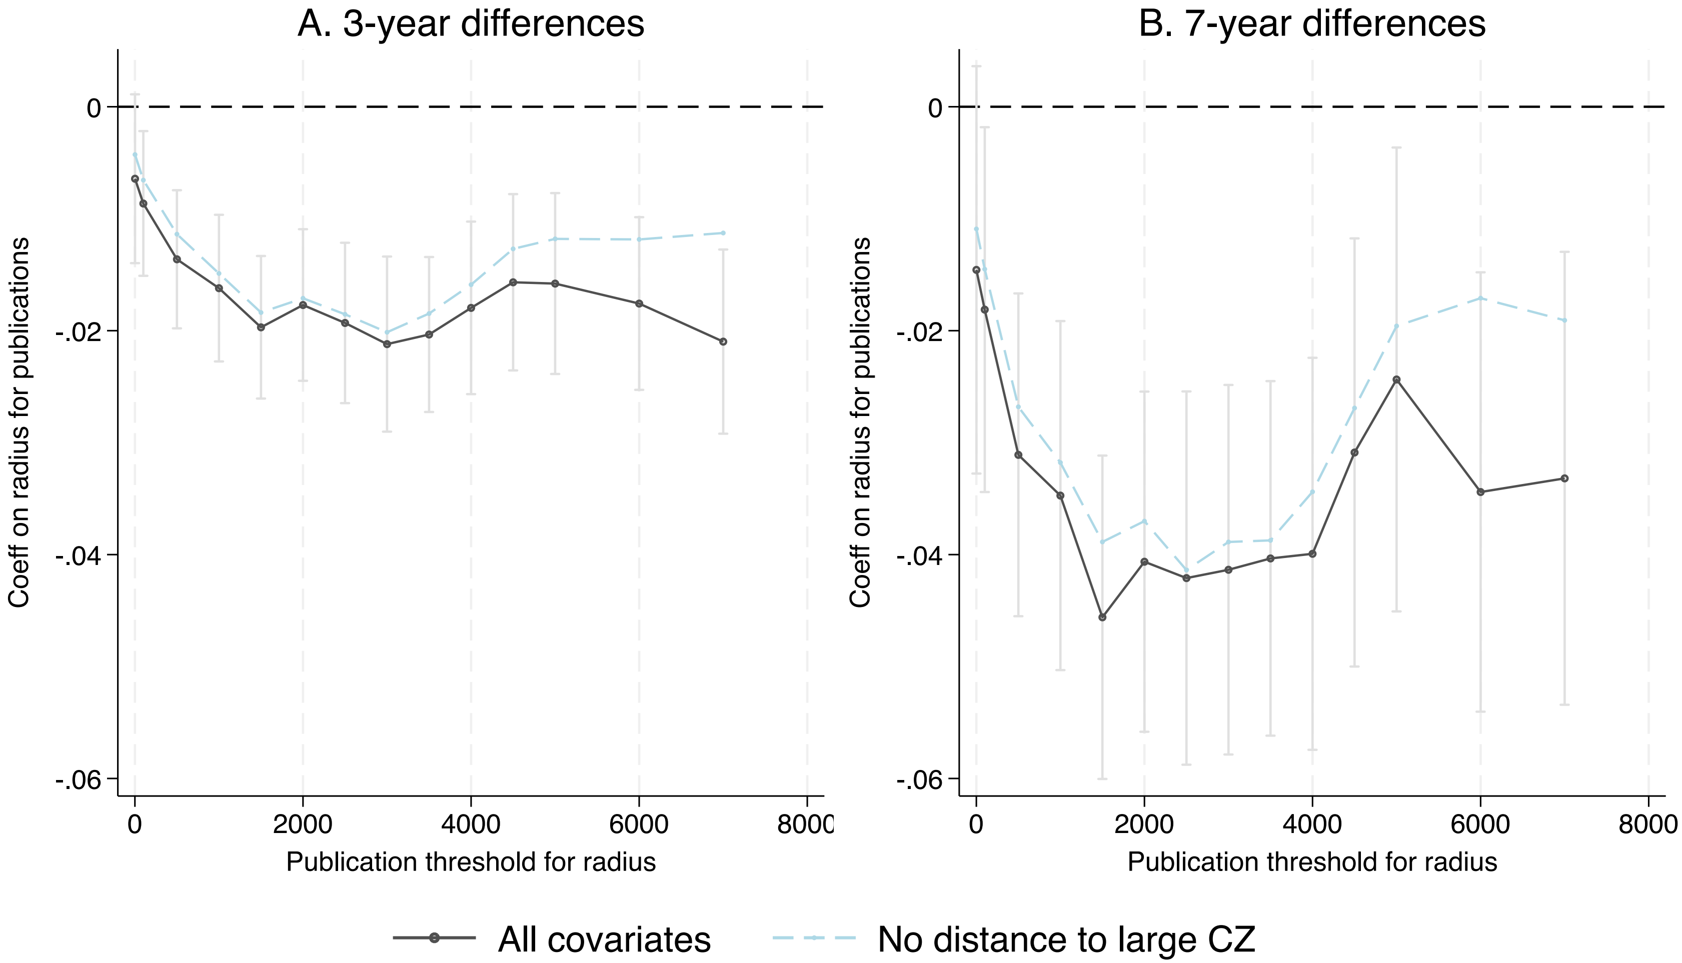

Note: The black circles are the coefficients on radius enclosing the number of pre–2007 AI publications on the x–axis from median regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 3 column 6; the 95% confidence intervals are also shown. The number of observations in each regression is 2964 in panel A and 5928 in panel B. The dashed blue line represents the coefficients from regressions which do not control for the population enclosed by the radius, but are otherwise the same.

Figure 11: Effect of distance using patents only, with population control

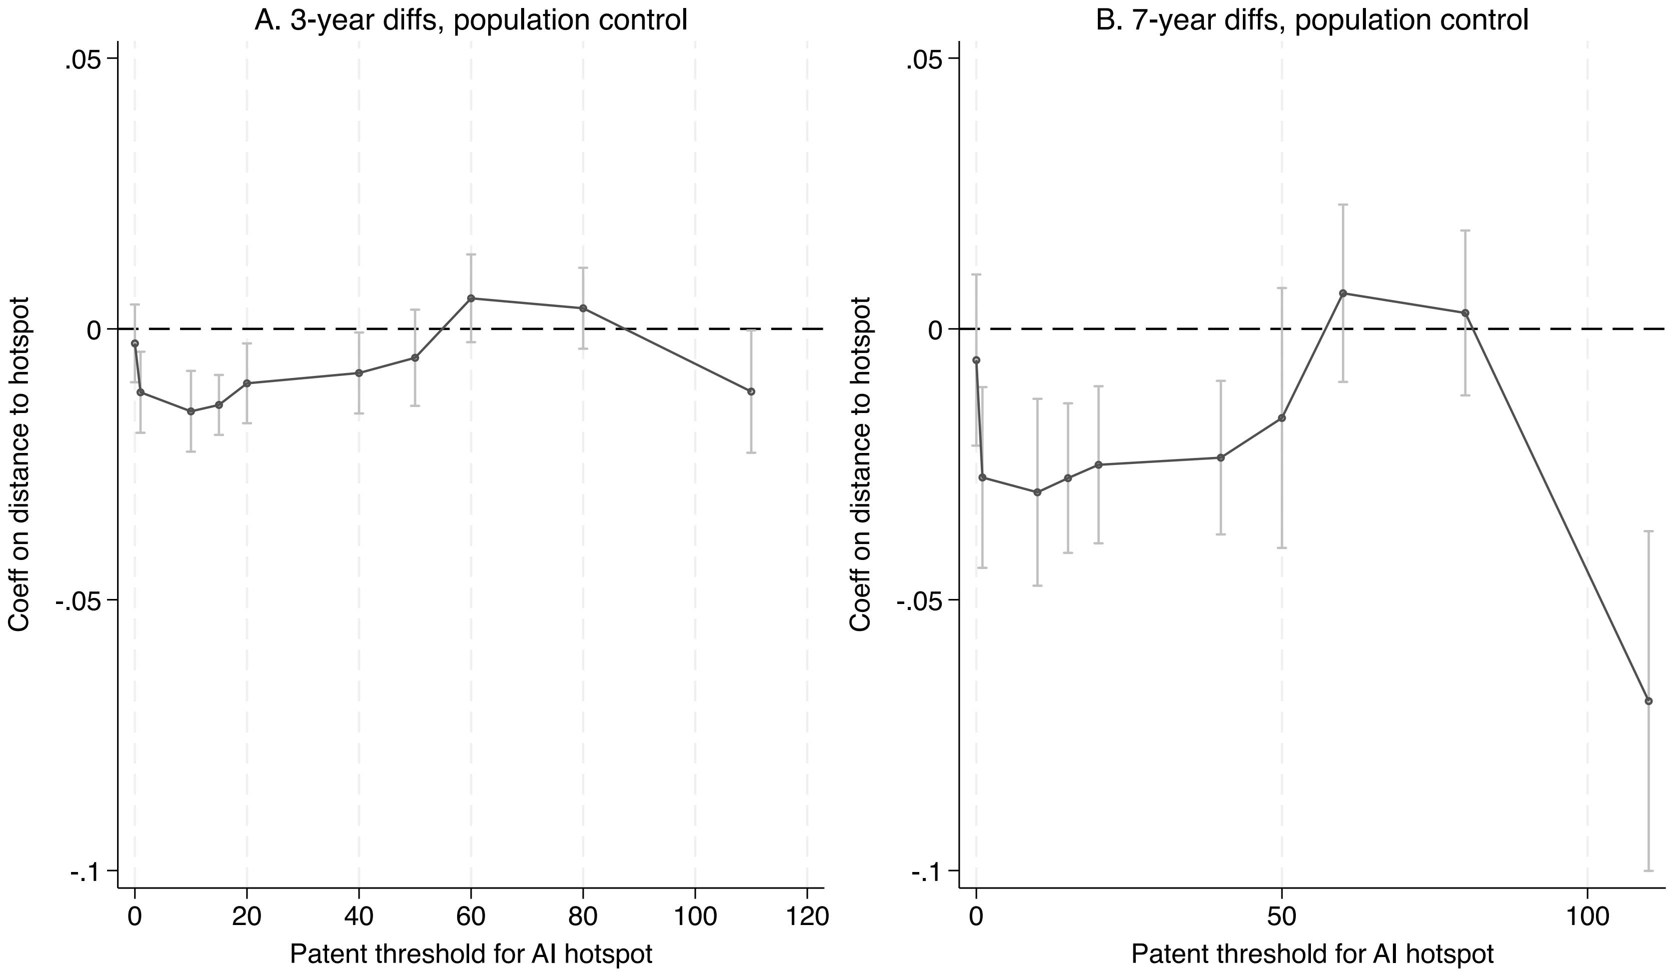

Note: The black circles are the coefficients on log distance to the nearest hotspot from median regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6; the 95% confidence intervals are also shown. The number of observations in each regression is 2964 in panel A and 5928 in panel B. The designation of AI hotspots is based on patents only; the x–axis is the number of pre–2007 AI patents designated as the threshold for a commuting zone to be an AI hotspot.

Figure 12: Compare with evaluations by percentile of hotspot patents/publications

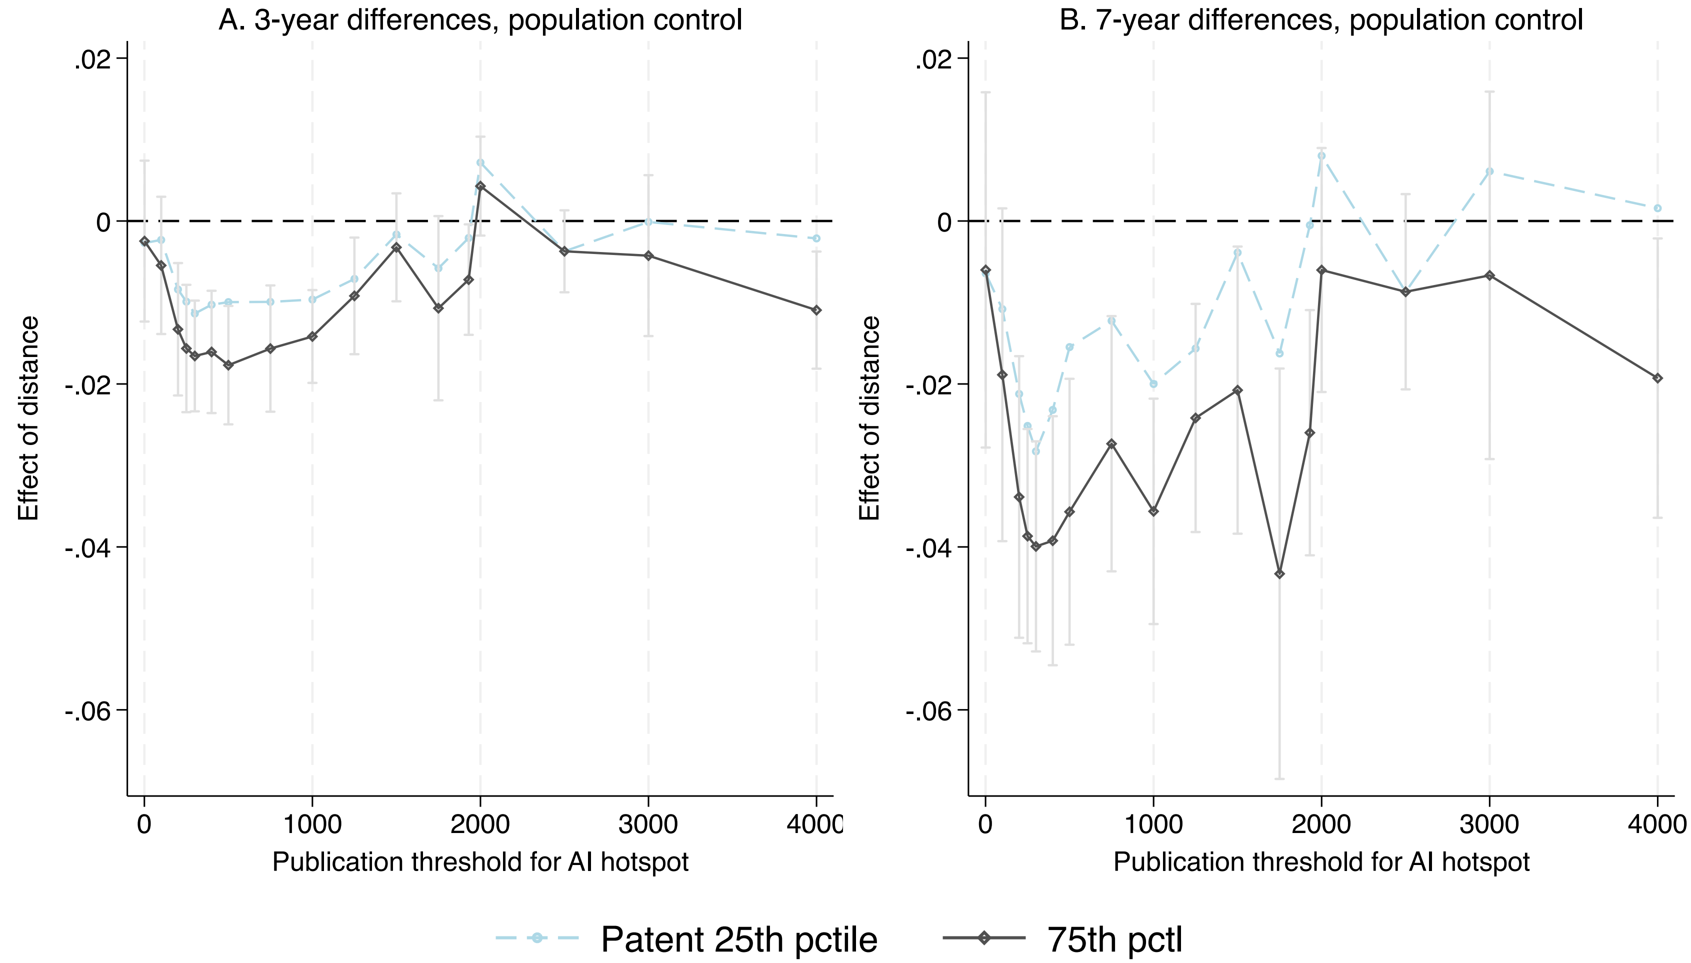

Note: The black circles are the coefficients on log distance to the nearest hotspot, evaluated at the 75th percentile of the hotspot's patents/publications ratio, from median regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6, as well as the hotspot's patents/publication ratio and its interaction with the log distance to the nearest hotspot; the 95% confidence intervals are also shown. The number of observations in each regression is 2964 in panel A and 5928 in panel B. The dashed blue line represents the 25th percentile. The x–axis is the number of pre–2007 AI publications designated as the threshold for a commuting zone to be an AI hotspot.

Figure 13: Effect of distance by type of AI skill requested in advertisements

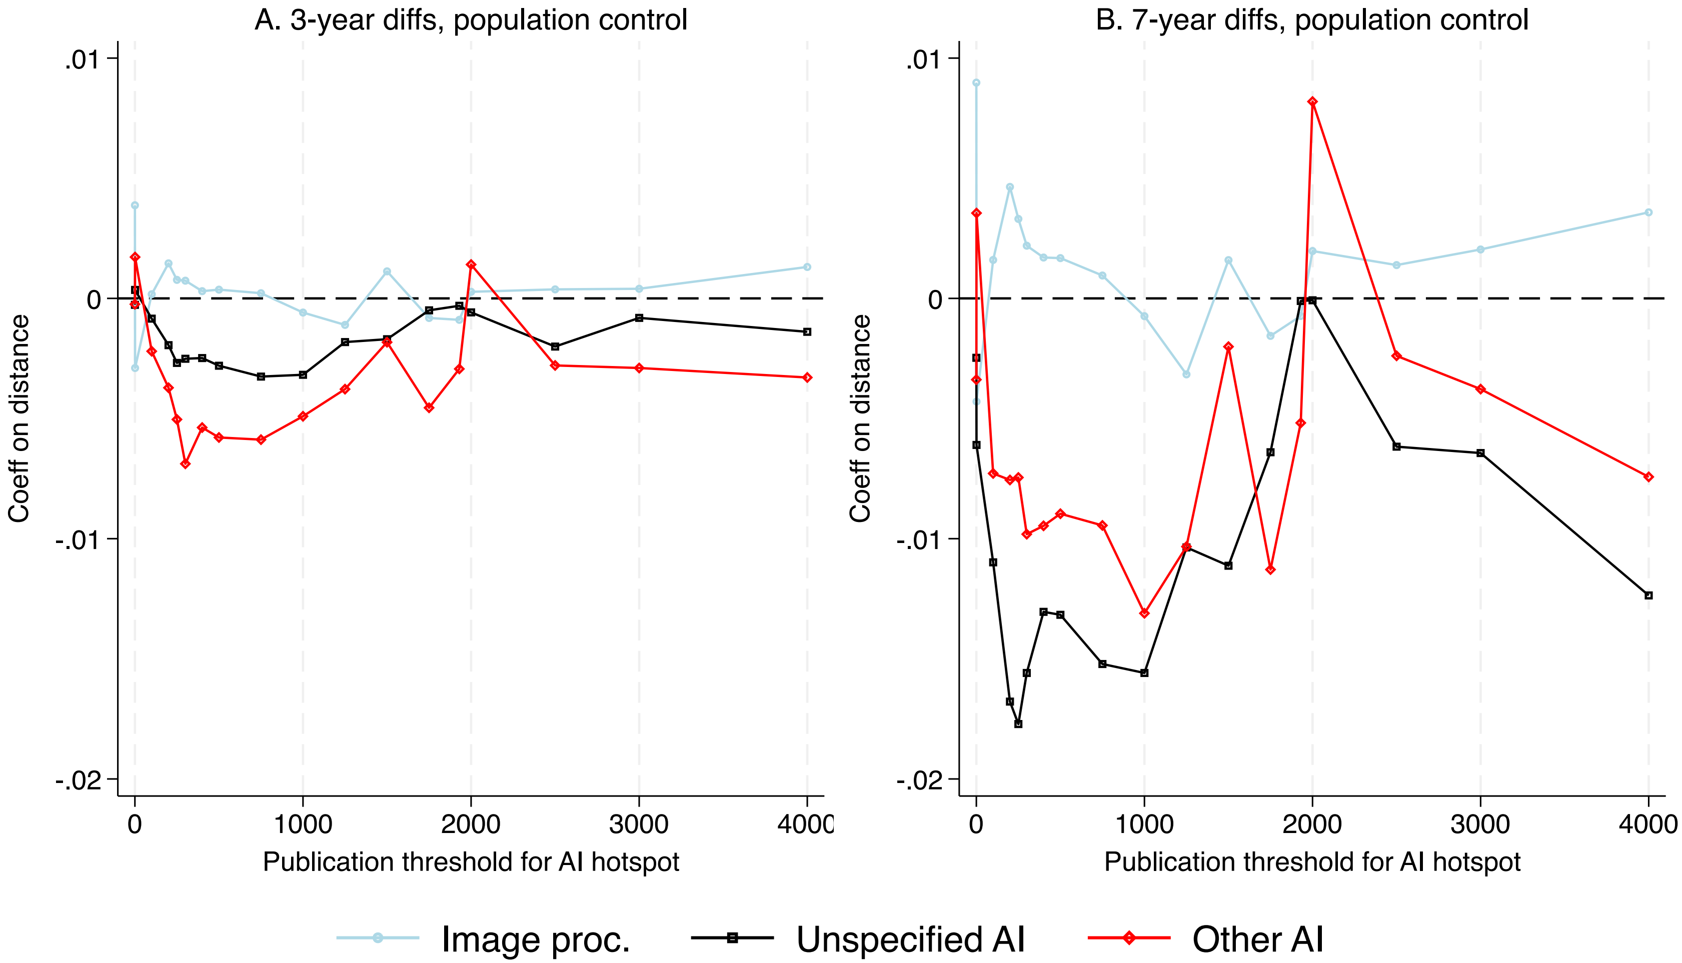

Note: The black circles are the coefficients on log distance to the nearest hotspot from median regressions of the change in the share of unspecified AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6. The number of observations in each regression is 2964 in panel A and 5928 in panel B. The other lines represent the corresponding coefficients for the two other types of AI. The x–axis is the number of pre–2007 AI publications designated as the threshold for a commuting zone to be an AI hotspot.

Figure 14: Effect of distance for computer–math, engineering and sales occupations

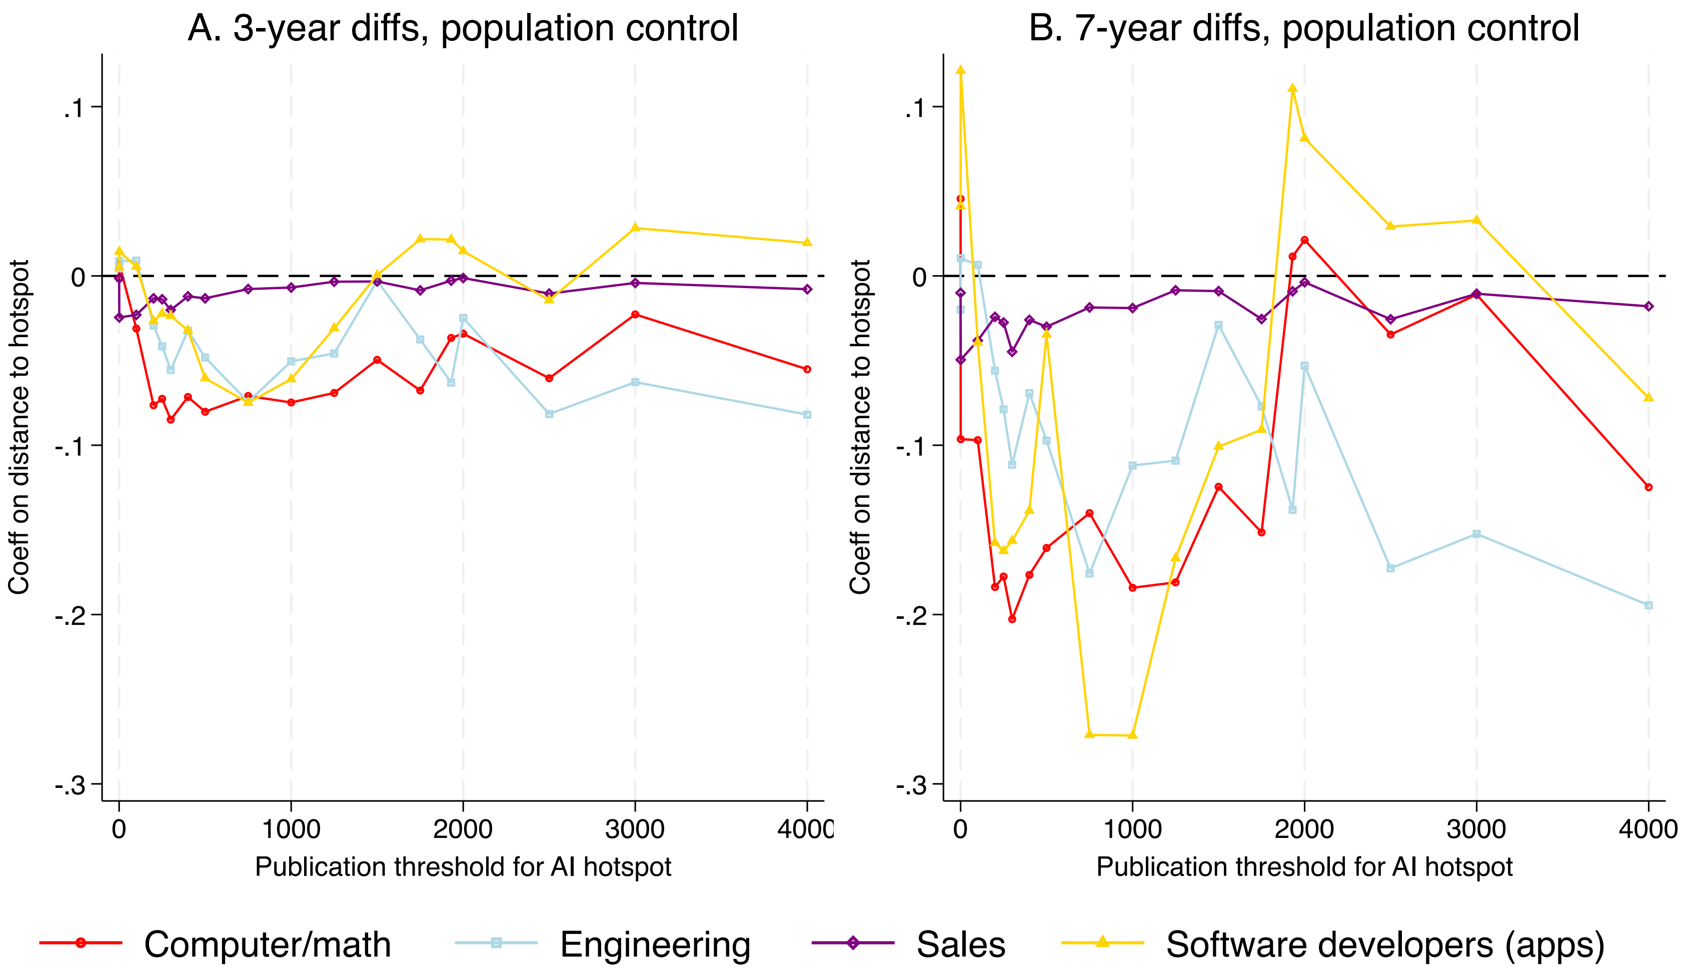

Note: The points are the coefficients on log distance to the nearest hotspot from regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6. Median regressions are used for computer/math and software developer occupations and OLS regressions for engineering and sales. The number of observations in panel A/panel B is 2934/5894 for computer–mathematical occupations; 5622/2726 for software applications developers occupations; 5867/2921 for architects and engineers and 5910/2946 for sales occupations. The underlying microdata samples are restricted in each regression to the occupation in question. The x–axis is the number of pre–2007 AI publications designated as the threshold for a commuting zone to be an AI hotspot.

Figure 15: Effect of distance for management and business–finance occupations

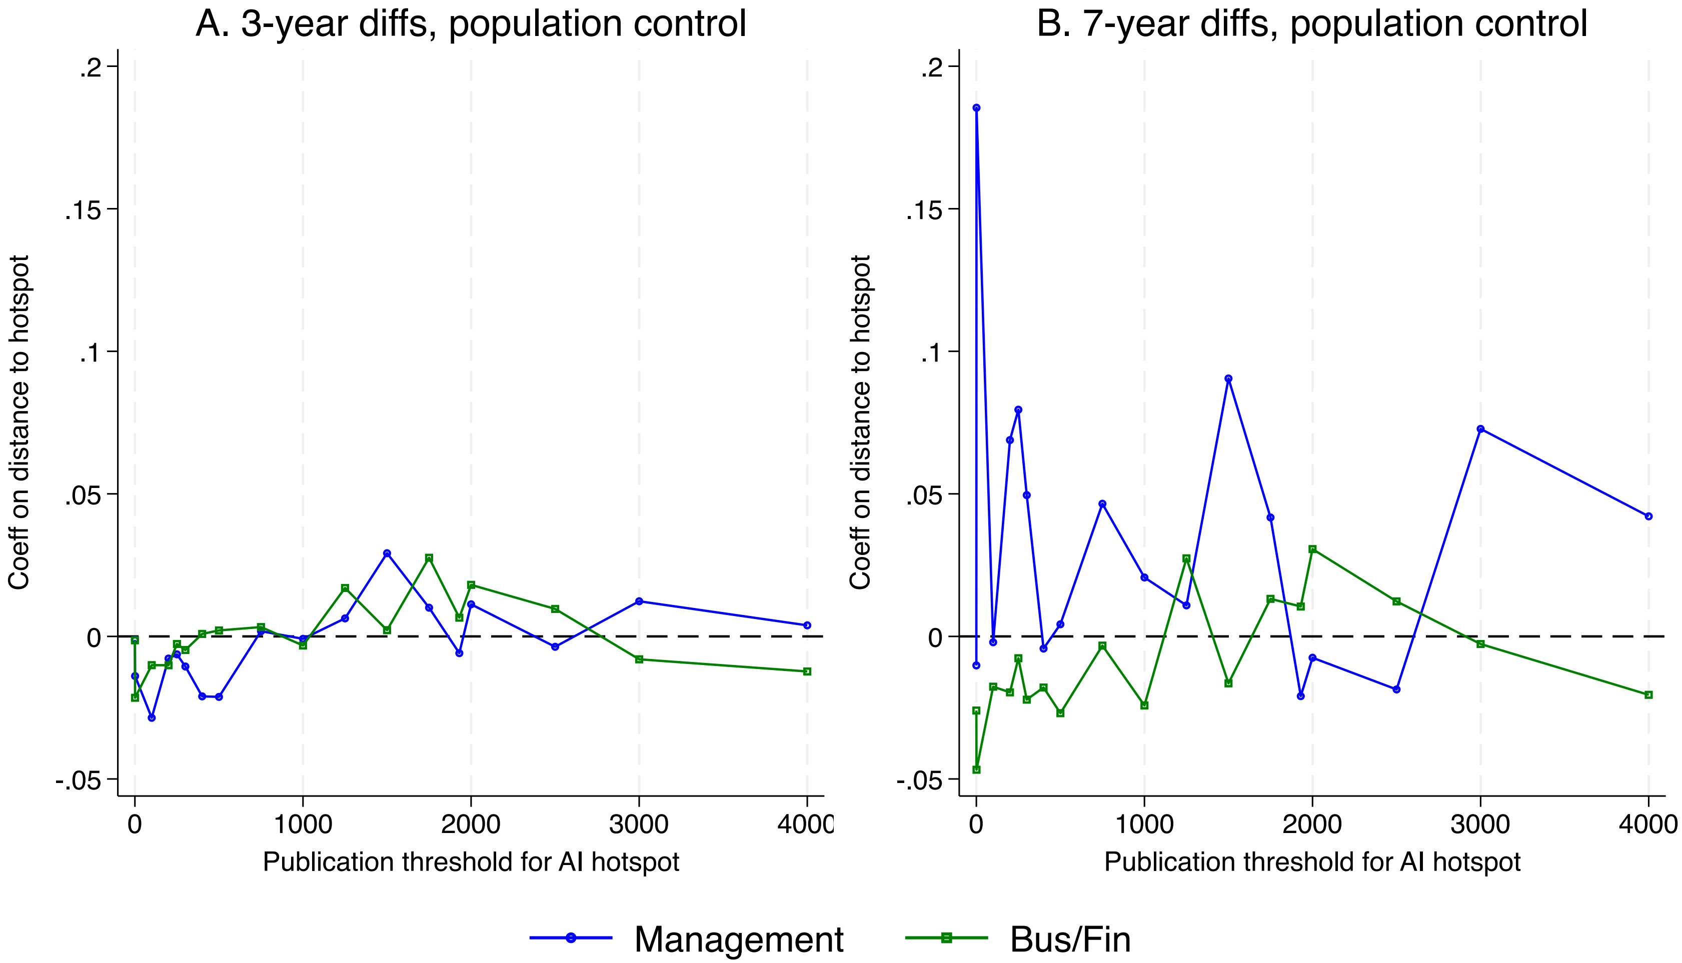

Note: The points are the coefficients on log distance to the nearest hotspot from OLS regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6. The number of observations for managers is 5925 (panel A) and 2962 (panel B); for business–finance occupations 5881 (panel A) and 2924 (panel B). The underlying microdata samples are restricted in each regression to the occupation in question. The x–axis is the number of pre–2007 AI publications designated as the threshold for a commuting zone to be an AI hotspot.

Figure 16: Effect of distance for finance/insurance and information industries

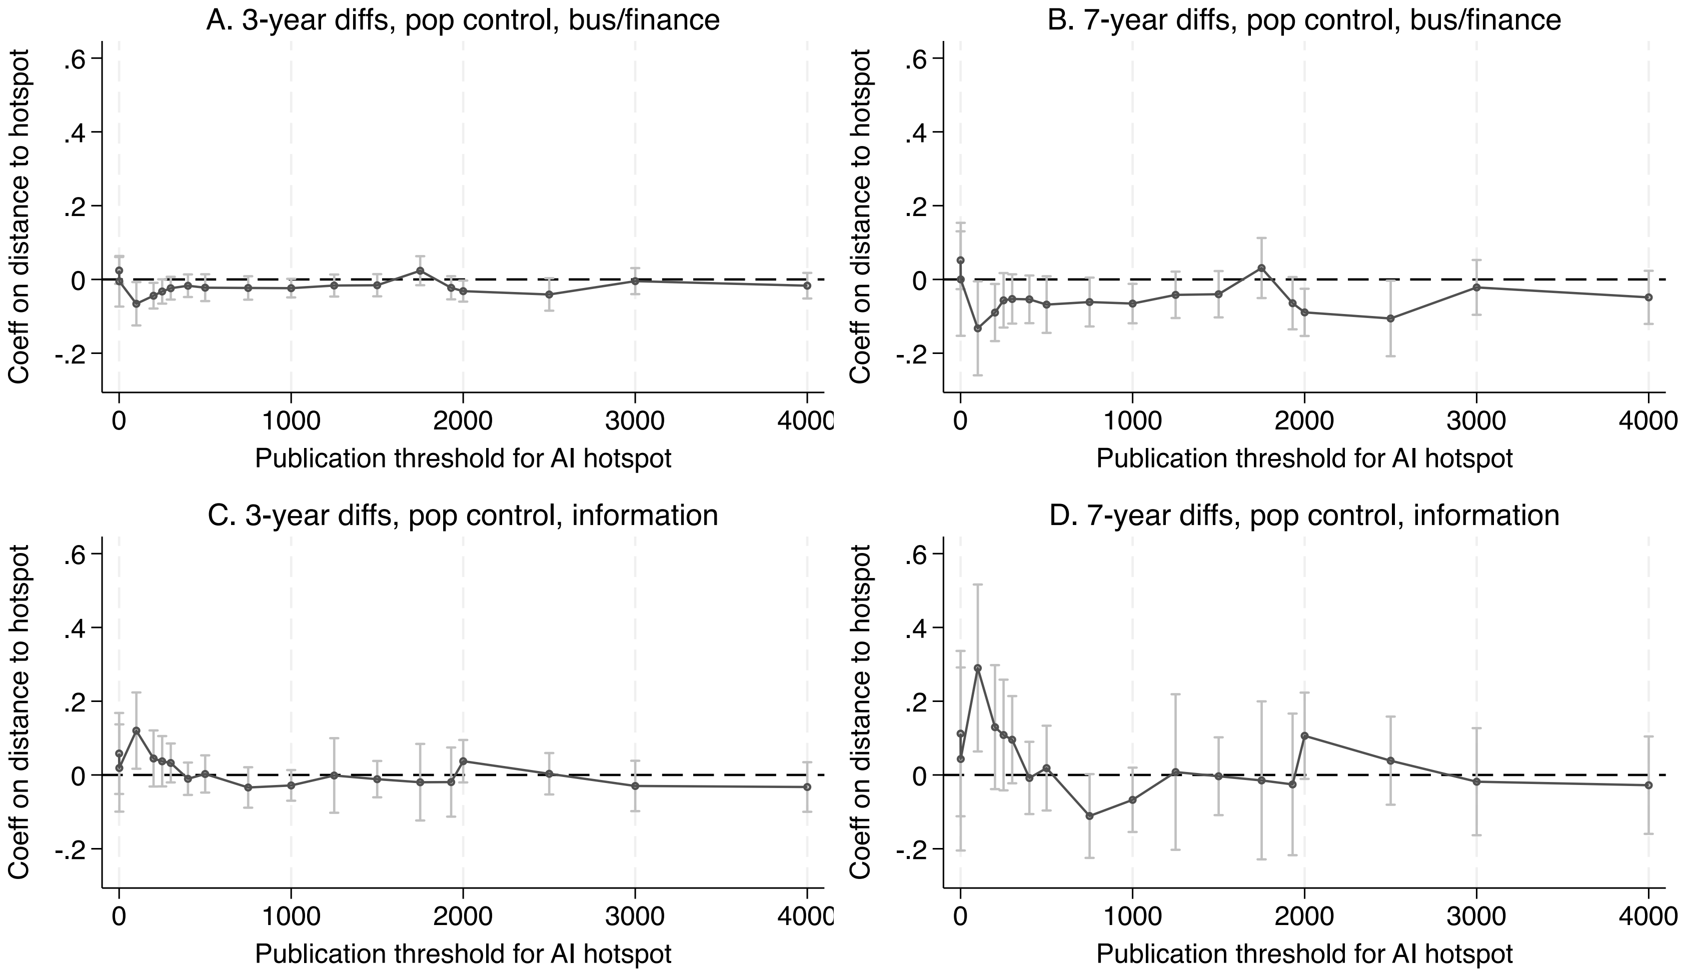

Note: The points are the coefficients on log distance to the nearest hotspot from OLS regressions of the change in the share of AI job advertisements in all advertisements, based on the specification of Table 2 columns 4–6. The number of observations is 5925 (panel A), 2961 (panel B), 5885 (panel C) and 2946 (panel D). The underlying microdata samples are restricted in each regression to the industry in question. The x–axis is the number of pre–2007 AI publications designated as the threshold for a commuting zone to be an AI hotspot.

Table 1: Summary statistics

	Mean	Median	Min	Max	Obs
A. Δ AI job advertisement share (%)					
Δ=3	0.061	0.035	-2.46	4.70	5928
Δ=7	0.136	0.093	-2.82	4.90	2964
B. Δ cumulative AI job ad share (%)					
Δ=3	0.031	0.022	-1.89	0.82	5928
Δ=7	0.074	0.062	-2.60	1.12	2964
C. Hotspot (1000+ AI pubs) characteristics					
Number of publications (papers+patents)	1938	1665	1044	6769	741
Patents*100/(papers+patents)	4.3	4.3	0	13.4	741
D. Initial conditions covariates					
Any AI paper prior to 2007	0.48	0	0	1	741
Any AI patent prior to 2007	0.17	0	0	1	741
AI papers prior to 2007	150	0	0	6617	741
AI patents prior to 2007	5.5	0	0	535	741
Job advertisements 2007	16,583	2570	3	696,205	741
Population in 2000 in thousands	380	104	1.19	16,393	741
IT share 2007 (%)	9.01	7.77	0	42.86	741
E. Distances (km)					
To closest hotspot (1000+ AI pubs)	412	324	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
F. Differenced covariates **Δ=3**					
AI papers	17.7	0	-101	5508	5928
AI patents	37.3	0	-40	326	5928
Log job advertisements	0.17	0.25	-3.5	2.1	5928
IT job ad share (%)	-0.31	-0.83	-45	46	5928
G. Differenced covariates **Δ=7**					
AI papers	37.0	0	-51	6596	2964
AI patents	0.73	0	-33	-48	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

Notes: The definition of an AI publication hotspot in the table is a commuting zone (CZ) with at least 1000 AI publications (patents+papers) 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. A paper is a journal publication or conference proceeding.

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. 3-year differences						
Log distance to closest hotspot	-0.011***	-0.017***	-0.016***	-0.013***	-0.010**	-0.015**
(1000+ AI publications)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.005)
Observations	5928	5928	5928	5928	5928	5928
Pseudo R-squared/R-squared	0.09	0.10	0.10	0.12	0.13	0.13
B. 7-year differences						
Log distance to closest hotspot	-0.029***	-0.031***	-0.031***	-0.026***	-0.022**	-0.034**
(1000+ AI publications)	(0.004)	(0.006)	(0.005)	(0.007)	(0.007)	(0.011)
Observations	2,964	2,964	2,964	2,964	2,888	2,964
R-squared/Pseudo-R-squared	0.15	0.16	0.22	0.22	0.22	0.23
AI publications through 2006, log closest hotspot publications	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, AI papers, AI patents	--	--	Yes	Yes	Yes	Yes
Log distance to closest large CZ; log closest large CZ population	--	--	--	Yes	Yes	Yes

Notes: The dependent variable is the three-year difference (panel A) or seven-year difference (panel B) in AI jobs' share of all job advertisements; the share measured in %. Data for 2007 and 2010-2019. All regressions include log number of publications (papers+patents) in closest AI hotspot; year dummies; and AI publications through 2006: dummies for any paper and for any patent, the number of papers and its square, and the number of patents. 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 clustered by commuting zone in parentheses.

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

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. 3-year differences						
Log distance to closest hotspot (1000+ AI pubs)	-0.008***	-0.008***	-0.007***	--	--	--
	(0.001)	(0.001)	(0.001)			
Log radius of circle enclosing 1000+ AI pubs				-0.014***	-0.015***	-0.016***
				(0.003)	(0.003)	(0.003)
Observations	5928					
Pseudo R-squared	0.08	0.12	0.13	0.10	0.12	0.12
B. 7-year differences						
Log distance to closest hotspot (1000+ AI pubs)	-0.017***	-0.016***	-0.012***	--	--	--
	(0.003)	(0.003)	(0.003)			
Log radius of circle enclosing 1000+ AI pubs	--	--	--	-0.029***	-0.032***	-0.035***
				(0.008)	(0.007)	(0.008)
Observations	2964					
Pseudo R-squared	0.17	0.21	0.22	0.16	0.21	0.21
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 in hotspot	--	--	Yes	--	--	--
Log population within circle	--	--	--	--	--	Yes

Note: Median regressions. The dependent variable in columns 1-3 is the three-year (panel A) or seven-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, measured in %. The dependent variable in columns 4-6 is the three-year (panel A) or seven-year difference (panel B) in AI jobs' share of all job advertisements, measured in %. Data for 2007 and 2010-2019. All regressions include year dummies; initial conditions (log number of publications in closest AI hotspot, dummies for any pre-2007 AI paper, any pre-2007 AI patent, the number of pre-2007 papers and its square, and the number of pre-2007 patents, log job advertisements 2007, log population 2000, IT share in advertisements 2007); log of average distance to other CZs and log distance to the closest CZ. 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 clustered by commuting zone in parentheses.

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

Table 4: Impacts of distance to closest AI publications hotspot and commuting zone's own publications on change in AI job advertisement share

	(1)	(2)
Log distance to closest hotspot	-0.016***	-0.031***
(1000+ AI publications)	(0.004)	(0.008)
Pre-2007 AI paper controls	Yes	Yes
(p-value joint significance)	(0.00)	(0.58)
Pre-2007 AI patent controls	Yes	Yes
(p-value joint significance)	(0.19)	(0.00)
Interactions of pre-2007 AI publications	Yes	Yes
with log distance	(0.00)	(0.35)
(p-value joint significance)		
Observations	5928	2964
R-squared/ Pseudo R-squared	0.12	0.22
Effect of distance evaluated at		
1 AI publication pre-2007	-0.015***	-0.028**
	(0.003)	(0.008)
100 AI publications pre-2007	-0.013***	-0.025**
	(0.003)	(0.008)
500 AI publications pre-2007	-0.008**	-0.013
	(0.004)	(0.011)
1000 AI publications pre-2007	-0.002	-0.002
	(0.006)	(0.018)

Notes: Median regressions. All regressions are based on the specification of Table 2 columns 4-6. Standard errors clustered by commuting zone in parentheses. Publications are the sum of papers and patents.
p<0.01, ** p<0.05, * p<0.1

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

	All AI (1)	Unspecified AI only (2)	Image Processing (3)	Other AI (4)
A. AI job ads with valid occupation				
Share AI type in AI job ads (714,348 obs)	100%	37.1%	12.5%	50.3%
Share computer scientist/mathematician	62.6%	68.3%	49.9%	61.6%
[obs]	[714,348]	[264,852]	[88,970]	[360,526]
B. 3-year differences, median regression (5928 obs)				
Log distance to closest hotspot	-0.013**	-0.004***	-0.001	-0.006***
(1000+ AI publications)	(0.003)	(0.002)	(0.001)	(0.001)
R-squared	0.12	0.16	0.02	0.04
C. 7-year differences, median regression (2964 obs)				
Log distance to closest hotspot	-0.026**	-0.016***	-0.001	-0.013**
(1000+ AI publications)	(0.007)	(0.003)	(0.002)	(0.004)
R-squared	0.22	0.25	0.02	0.09

Notes: Median regression except for column (3) panel A, which is OLS. 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. An image processing job advertisement mentions image processing among the required skills. The types of AI are mutually exclusive. All regressions are based on the specification of Table 2 columns 4-6. Standard errors clustered by commuting zone in parentheses.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

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

	3-year differences			7-year differences		
	Obs	Mean of dependent variable	OLS	Obs	Mean of dependent variable	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
All	5928	0.06	-0.015** (0.005)	2964	0.14	-0.034** (0.011)
Computer and mathematical	5894	0.59	-0.133*** (0.025)	2934	1.22	-0.280*** (0.053)
Architectural and engineering	5875	0.18	-0.050** (0.017)	2925	0.36	-0.112** (0.045)
Management	5925	0.07	-0.001 (0.008)	2962	0.25	0.021 (0.023)
Business and finance	5881	0.07	-0.003 (0.009)	2924	0.15	-0.024 (0.027)
Sales	5910	0.01	-0.007* (0.004)	2946	0.02	-0.019** (0.008)
Other occupations	5928	0.01	0.005 (0.006)	2964	0.02	0.009 (0.013)

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 based on the specification of Table 2 columns 4-6. The dependent variable is the change in the share of advertisements in the specified occupation requiring AI. Standard errors clustered by commuting zone in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

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

	3-year differences			7-year differences		
	Obs	Mean of dependent variable	OLS	Obs	Mean of dependent variable	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
All	5928	0.061	-0.015**	2964	0.136	-0.034**
			(0.005)			(0.011)
Agriculture, Utilities, Mining, Construction, Manufacturing	5926	0.044	0.010	2962	0.084	0.026
			(0.007)			(0.017)
Wholesale trade, Retail trade, Warehousing, transportation	5928	0.010	0.003	2964	0.018	0.007
			(0.003)			(0.008)
Information	5885	0.104	-0.028	2946	0.208	-0.067
			(0.021)			(0.044)
Finance, Insurance	5925	0.131	-0.024*	2961	0.276	-0.065**
			(0.013)			(0.027)
Real Estate, Professional and scientific services, Administration	5928	0.115	-0.037**	2964	0.244	-0.085**
			(0.016)			(0.032)
Education, Health	5928	0.020	0.013	2964	0.048	0.027
			(0.014)			(0.032)
Arts and recreation, Accommodation	5923	0.028	0.014	2961	0.055	0.032
			(0.014)			(0.027)
Other services, Public administration	5926	-0.038	0.073	2962	0.167	0.156
			(0.080)			(0.170)
Missing industry	5928	0.121	-0.051***	2964	0.249	-0.109***
			(0.007)			(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 based on the specification of Table 2 columns 4-6. 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 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. Other
AI ChatBot, Amelia, ANTLR, Automatic Speech Recognition (ASR), Caffe Deep Learning Framework, Chatbot, Computational Linguistics, Computer Vision, Decision Trees, Deep Learning, Deeplearning4j, Google Cloud Machine Learning Platform, Gradient boosting, H2O (software), IBM Watson, Image Recognition, IPSoft, Ithink, Keras, Latent Dirichlet Allocation, Latent Semantic Analysis, Lexalytics, Lexical Acquisition, Lexical Semantics, Libsvm, Machine Translation (MT), Machine Vision, MLPACK (C++ library), MoSes, MXNet, Madlib, Mahout, Microsoft Cognitive Tookit, Mlpy, ND4J (software), Natural Language Processing, Natural Language Toolkit (NLTK), Nearest Neighbor Algorithm, Neural Networks, Object Recognition, Object Tracking, OpenCV, OpenNLP, Pattern Recognition, Pybrain, Random Forests, Recommender Systems, Sentiment Analysis / Opinion Mining, Semantic Driven Subtractive Clustering, Semi-Supervised Learning, Sentiment Classification, Speech Recognition, Supervised Learning (Machine Learning), Support Vector Machines (SVM), TensorFlow, Text Mining, Text to Speech (TTS), Tokenization, Torch (Machine Learning), Unsupervised Learning, Virtual Agents, Vowpal, Wabbit, Word2Vec, Xgboost

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

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

	Share (%)	AI required? (%)	Sample of ads requiring AI: Occupation (%)			
			Computer and math	Management	Architects engineers	Business finance
Industry	(1)	(2)	(3)	(4)	(5)	(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.