

SPATIAL PERSISTENCE OF AGGLOMERATION IN SOFTWARE PUBLISHING*

George Deltas[†] Dakshina G. De Silva[‡] Robert P. McComb[§]

June 2, 2019

Abstract

We estimate the effects of industrial localization on the spatial persistence of employment in the software industry, using establishment data from Texas for the 2000-2006 period. Locations with an initial concentration of software employment retain an excess number of employees, beyond that expected from job turnover and job persistence at the establishment level. This is not driven by differential establishment growth or survival, but it is due to (a) the retention by establishments in a location of jobs lost by other establishments in that location, and (b) the propensity of software establishments to enter in locations with prior software establishment presence. These findings are more consistent with labor channel effects than with disembodied knowledge spillovers.

JEL Classification: R32, L86, R12.

Keywords: Agglomeration economies, labor pools, knowledge spillovers, firm growth, spatial effects.

*We would like thank Scott Adams, three anonymous referees and participants the International Industrial Organization conference, the Bank of Korea, Pusan National University, and the University of Vermont for useful comments that improved the paper. This manuscript replaces earlier research that was circulated under a similar title. For that earlier research, we want to thank Tim Dunne and seminar participants in the Dusseldorf Center for Industrial Economics (DICE), the University of Illinois, the University of Lancaster, University of Macedonia, the EARIE conference, and the Western Economic Association conference for their helpful comments. We also would like to thank Anita Schiller, and Mervin Ekanayake for their skillful research assistance, the Texas Workforce Commission for providing the data, and a number of sources with experience in the software industry.

[†]Department of Economics, University of Illinois, 1407 W. Gregory, Urbana, IL 61801, United States. (e-mail: deltas@illinois.edu), and Department of Economics, Lancaster University Management School, Lancaster University, Lancaster, LA1 4YX, United Kingdom.

[‡]Department of Economics, Lancaster University Management School, Lancaster University, Lancaster, LA1 4YX, United Kingdom. (e-mail: d.desilva@lancaster.ac.uk).

[§]Department of Economics, PO Box 41014, Texas Tech University, Lubbock, TX 79409, United States. (e-mail: robert.mccomb@ttu.edu)

1 Introduction

A voluminous literature has examined the ubiquity and sources of spatial aggregation at the industry level and investigated the effects of agglomeration on technical change, firm growth, productivity, and other outcomes. To what extent, though, is industrial agglomeration self-sustaining? What are the drivers of its persistence? These questions are the focus of this paper.

Persistent agglomeration raises policy challenges for areas and localities that offer no advantage to any industry that benefits from agglomeration economies. If persistence is shown to be significant, it would suggest that such a locality could only succeed in becoming an activity hub through a concerted effort to attract simultaneous relocation or entry of multiple enterprises. It would also suggest that economic development efforts would be more successful when targeting some specific industries rather than being applied with no specific sectors in mind. However, if agglomeration were not persistent, a laissez-faire approach would be appropriate. A policy aimed at facilitating industry clusters must identify sources of agglomeration economies which are responsible for the persistence of agglomeration. This is a challenging task because agglomeration economies operate through many channels simultaneously.¹

In this paper, we isolate the contributions of agglomeration spillovers on the spatial persistence of employment from the effects arising from other factors. We accomplish this using data from the software publishing industry in Texas, which lends itself exceptionally well to the investigation of labor market and (disembodied) knowledge spillovers. The industry's output is intangible intellectual property that sells in the national and global markets. Thus, locational factors such as access to natural resources, local demand, proximate input suppliers, and transportation costs are not particularly relevant. Capital inputs are not typically fixed (beyond the period of the current building lease) or tend to depreciate very rapidly (e.g., computer hardware and software), so past investment decisions in physical plant do not constrain firm location over the medium term. Moreover, given the dynamic and competitive nature of the industry, the rapid evolution of software development, and its heavy reliance on specialized human capital, software publishing seems well suited to benefit from labor pooling and knowledge spillovers. Software engineers and computer programmers, the primary labor input, tend to be young and mobile.² Therefore, if labor-driven localization economies are

¹These include the prevalence of specialized inputs and upstream service providers, labor pooling, and localized knowledge spillovers, as recognized since Marshall (1920). In general, all three forces can drive agglomeration, as demonstrated recently by Ellison, Glaeser, and Kerr (2010).

²According to the Bureau of Labor Statistics, 51.5% of the employees in software publishing in 2017 were classified as

present, their effects should be easy to detect. However, the absence of other plausible channels for such spillovers is a necessary but not a sufficient condition to measure the effects of labor channel localization economies. We must also control for localized economic activity that might otherwise result in spurious spatial association of software publishing. We do so using data on other business establishments at the address level to account both for local economic development and also for the presence of infrastructure and zoning that is conducive for software publishing activity.³

Geographic clusters of software establishments are highly persistent, even when measured in locations that are only two miles across. Employment at the facility level exhibits very small persistence, but employment within the locations these facilities are situated is very persistent. Most of this persistence is driven by employment by entrants who locate in different addresses but very close proximity to existing firms. This geographic clustering of entrants around incumbents takes place even though the locations with incumbent firms cover a tiny portion of the Texas area with active business enterprises. These facts are highlighted in Figure 1. Total software employment in Texas has grown slowly during our sample period. Almost all employment has been in locations with active establishments at the start of our sample, with locations defined at a very fine level of spatial dis-aggregation (1-mile radius circles; details provided in section 3.1). The preponderance of commercially active locations in Texas had no software publishers at the start of period and did not have any at the end of our period either. Such persistence could have been explained by establishment-level persistence in employment. But this is not the case: employment in establishments active at the start of our period has declined substantially. The decline is steeper if one excludes job creation at the establishment level. The gap between these two lines represents jobs lost by an establishment that were “captured” by other establishments located in the same two-mile wide location.

Motivated by these facts, we formally estimate employment persistence at various levels of aggregation and explain it on the basis of several relevant economic factors. Within a 1-mile radius, we find that the elasticity of software employment at the end of 2006 with respect to employment at the start of 2000 is approximately 0.50 – 0.75. These are very high numbers given that the majority of specific software jobs that were present at the start of 2000 were lost. In fact, at the establishment level, we find the elasticity of computer occupation professionals.

³Texas has a mix of areas with a high and a low concentration of software establishments; thus, our findings are likely to be relevant to other states. We expect the labor market channel effects we identify to have external validity to other industries that employ highly skilled labor, though in those industries other agglomeration effects, about which we have little to say, are also present.

software employment to be in the 0.30 – 0.35 range.

Our results also indicate that end-of-period employment is lower for establishments located in areas with existing software establishments, despite the absence of product-side competition. Thus, the presence of software establishments reduces rather than enhances the success and growth prospects of other incumbents. We provide evidence that many of the jobs lost by a software establishment in a location were captured by other software establishments located no more than one mile away. This, however, is not sufficient to fully explain the high elasticity at the location level, leaving the location choices of entrants as the likely source of the additional jobs. To confirm this possibility, we investigate entry rates at different locations and follow the survival of these entrants as a function of location characteristics. After controlling for localized presence in a control group of other businesses, we find that software establishments tend to enter locations within one mile of other software establishments, thus reinforcing rather than diluting any initial location advantage. We find that the attractiveness of a location diminishes rapidly with distance and almost disappears after 5 miles. There are no comparable effects on exit at any distance threshold, as we show using survival analysis.

The high spatial persistence of employment in the software industry is driven by two components: the propensity of software establishments to locate in areas with pre-existing software establishments, and the ability of co-located establishments to act as a “sponge” for jobs lost by their neighbors. Spatial persistence is not driven by differential establishment growth rates in areas with substantial software industry presence or by differential establishment survival probabilities. Spatial persistence operates primarily via the labor market. Proximity facilitates the transmission of information about employees of other firms and increases the match value for prospective employers. It also reduces switching costs, as employees can retain living and commuting arrangements when switching employers, whether incumbent or entrant. These factors reduce the set-up costs of establishments and subsequent recruiting costs. Thus, entrants tend to locate near existing establishments, and departing employees of a firm tend to locate a spin-off in the same vicinity.⁴ Our finding that the presence of proximate software establishments (marginally) decreases employment growth at the establishment level is inconsistent, all else equal, with the presence of positive knowledge externalities that are distinct from knowledge transmission through the hiring of workers. However, limited disembodied knowledge spillovers might be present and explain why the increased entry in high agglomeration locations

⁴City-wide policies aiming to attract software firms (rare in Texas to begin with) cannot drive our persistence findings given that our unit of analysis is very small.

of (some presumably more marginal) establishments does not depress average survival rates.

Though our analysis is limited to one industry over six years, it forms an important benchmark of relevance for other industries and periods. We show that spatial persistence can still be high even if the vast majority of the jobs associated with specific employers at the start of the period are no longer there. Quite certainly, the degree of persistence would decrease the longer this period is. In a less dynamic industry, that uses equally specialized labor, this level of job churn would take more years to accomplish, perhaps, say, three times as long. Other industries that are not tied geographically to a location and/or use less specialized labor would also be expected to have a smaller level spatial persistence. In contrast, industries that are partially tied geographically to a location because of location-specific factors would have higher levels of spatial persistence.

The findings of the importance of the labor pooling agglomeration effects for the spatial persistence of employment at the industry level are corroborated by studies that document the importance of labor pooling in generating agglomeration economies. Prominent among those are the contributions of Overman and Puga (2010) based on cross-industry analysis using UK data, Rosenthal and Strange (2001) who find in a cross-industry study that labor pooling effects are the dominant source of agglomeration beyond the zip code level, and Ellison, Glaeser and Kerr (2010) who measure the relative importance of various sources of agglomeration, including labor pooling, also using cross-industry data. Our study differs from these in terms of the special features of the computer industry and our focus on the link between agglomeration economies and spatial persistence of employment at the industry level.⁵

2 Conceptual Framework

We start by discussing knowledge and labor pool spillovers and illustrate their impact on the persistence of agglomeration through a stylized framework. Since professional experience is embodied in the individual, worker mobility, and direct personal and professional interaction are the primary channels of un-priced knowledge spillovers. Spillovers occurring through the hiring of workers are subsumed in our analysis under the labor pool effect (in any event, the two cannot be easily distinguished from each other). Thus, in our framework, knowledge spillovers refer to the disembodied transfer of information through personal and

⁵Agglomeration economies have been examined from the perspectives of firm location choices (Rosenthal and Strange, 2003, Woodward et al., 2006), firm exits (Staber, 2001), industry growth (Glaeser et al., 1992, Henderson et al., 1995, and Combes, 2000), and labor productivity (Ciccone and Hall, 1996).

professional contact. Both are facilitated by close geographic proximity. For example, workers who happen to know each other can easily meet (e.g., for lunch, coffee, or even car-pooling) if their places of employment are proximate. Being within walking distance reduces the barriers for such meetings. Interactions are far harder to arrange if the location of employment is even five miles away. The net knowledge spillover is expected to be positive although there may be “winners” and “losers” for any interaction. When one firm absorbs an idea from another, it might be to the detriment of the source firm if there is competition between them.⁶

The labor pool effect may have several expressions. A large pool of software engineers and programmers within commuting distance allows a firm to expand more easily by poaching the employees of other firms and reduces the set-up costs for new entrants. The overall effect on firm productivity and growth could be positive if employee-employer match quality improves by this practice; this mechanism could underpin the increased productivity of firms in close proximity to workers within commuting distance, as documented in Rice, Venables, and Patacchini (2006). The labor channel can also operate through the founding of start-up firms by employees departing their current employer. As with knowledge spillovers, these localized effects may have both winners and losers with some successful firms benefiting to the detriment of their neighbors. In fact, as one industry executive confided to us, programmers employed by different firms compare work conditions and terms of employment when in social contact and switch employment when their current employer is not competitive.⁷

Labor channel effects are not completely distinct from knowledge spillovers. In fact, such spillovers may be embodied in the workers who move from one firm to another, prompting some firms to employ “no-compete” clauses to limit their extent (see Fosfuri and Ronde, 2004 for a discussion of this and implications for agglomeration, and also Men, 2005). Consistent with the labor pooling hypothesis, Freedman (2008) finds that spatial clustering facilitates localized worker mobility in the software publishing industry. Rosenthal and Strange (2003) report a quotation from Saxenian (1994) in which a high-tech worker from Silicon Valley states, “The joke is that you can change jobs and not change parking lots.” For such localized job changes,

⁶Combes (2000) notes that a greater number of similar firms in a locality increases the likelihood of knowledge spillovers due to the greater likelihood of closer matches between firms. Spillovers might conceivably even be negative in aggregate if they lead to free-riding. Firms might reduce experimentation with new ideas hoping to piggy back on ideas developed by other firms. This might lead to a reduction of available knowledge.

⁷Some Silicon Valley firms run company buses from San Francisco to their facilities to ease commuting. Company buses also reduce, incidentally or intentionally, employee interaction across firms (relative to car pooling, public transport, or other non-company transport), suggesting that such interaction is not valued by firms.

search and transaction costs are probably negligible. The quotation has implications for the appropriate geographical area over which agglomeration externalities can be observed. Indeed, Rosenthal and Strange (2003), Wallsten (2001) and Arzaghi and Henderson (2008) note that localization externalities and knowledge spillovers attenuate rapidly within one mile.⁸

While both disembodied knowledge spillovers and labor pool effects may lead to a positive association between agglomeration and spatial persistence in employment, the effects differ subtly between these two sources of persistence. Labor market effects lead to the retention of jobs in a locale because the mere availability of labor attracts entry, even without improving the survival prospects of the entrants. Labor market effects can also improve firm growth and survival if thicker labor markets improve employee-firm matching and reduce ongoing costs of recruiting, but this effect is of second-order importance. In the absence of increased entry, increased growth and survival rates are unlikely to be driven by labor market effects.⁹ By contrast, knowledge spillovers are more likely to lead to gains for existing firms and better survival prospects for entrants (given that software firms compete in a national market). Of course, by increasing growth and survival prospects, knowledge spillovers would also increase entry. However, in the absence of any growth and survival effects, increased entry is unlikely to be driven by knowledge spillovers. Labor market effects are a more persuasive explanation in this case. A very simple stylized model is developed in the Appendix that describes a single period equilibrium of this dynamic story.¹⁰

The key insight, we believe, that leads to these disparate predictions of the two effects is that disembodied knowledge spillovers are not priced in the wage of the employees (and for an industry that is a small fraction of local economic activity, are not priced in land prices). Therefore, an entrant or incumbent in a location with a local concentration of software establishments would reap an unpriced productivity benefit that would result in higher growth and lower exit rates. Of course, this also provides incentives to enter in such areas. Knowledge embodied in human capital is largely priced into wages. A potential recruit who is offered

⁸Within urban cores, an IT cluster may emerge in part because of social agglomeration effects, i.e., satisfaction obtained from socializing with individuals of similar backgrounds (see Hutton, 2004). Though of primary importance for dense inner cities, e.g., San Francisco, Manhattan, Vancouver, they may also play a minor role in Texas.

⁹Spatial persistence may be particularly high if the source of agglomeration economies is the labor pool channel. When an establishment receives a negative shock and lays off workers, it creates a pool of local jobseekers and reduces the recruiting costs of co-located establishments, facilitating their expansion. This tends to stabilize employment at the location.

¹⁰Ideally, cost and profit data would be used to measure agglomeration economies. Survival rates are a crude proxy for profits, and only measure whether they exceed zero. Employment is correlated with profits, but imperfectly so. Unfortunately, profits are not available to us at the firm level. Though they could be obtained through use of a Census Regional Data Center, we would even then lack profits at the facility level, which is the unit of analysis in our study.

less than his market value will not sign up for a firm; an employee whose wage does not keep up with his or her marginal product will leave for another firm. Thus, following entry, establishments in areas of high agglomeration do not obtain any continuing benefits. However, potential entrants would still find it advantageous to locate there since it is harder to induce employees to relocate, not only because of upfront relocation costs, but also because these employees would forgo the option to work for other employers at a later time.

To summarize, based on this framework, we interpret possible results as follows. If we observe no excess spatial persistence in software industry employment (“excess” being higher than expected by the employment persistence at the establishment level and the economic activity in that location), then neither labor market nor knowledge spillovers are important. If we do observe excess spatial persistence, then either knowledge spillovers or labor market agglomeration economies, or both, are present. We can distinguish the presence and relative importance of each as follows. If establishments in areas of high software agglomeration exhibit faster growth and higher survival rates, and if these areas have higher entry rates, then knowledge spillovers dominate, though weaker labor market effects may also be present. If establishments in areas of high software agglomeration exhibit faster growth and lower survival rates, but there are no positive entry effects, then only knowledge spillovers are present. If establishments in areas of high software agglomeration exhibit slower growth and lower survival rates, but there are positive entry effects, then only labor agglomeration effects are present. Finally, if there are no growth and survival rate effects for establishments in areas of high software agglomeration, but there are positive entry effects, then labor pool effects dominate, but weaker knowledge spillovers may be present as well. Table A-1 provides a graphical presentation of the linkage between the possible empirical findings and theoretical implications.

3 Spatial Sampling, Estimation Framework, and Data

3.1 Data and Location Definition

The Quarterly Census of Employment and Wages (QCEW) compiled by the Texas Workforce Commission provides establishment data used in this study. The QCEW is an administrative database; establishment non-response is essentially non-compliance, and it carries administrative censure or penalties. Therefore, coverage should be (nearly) complete. Each record includes the specific location (address and latitude/longitude) of the establishment, business start-up date, and the relevant six-digit North American Industrial Classification

System (NAICS) code. Separate establishments of a firm are reported in separate records. These data contain observations from the 1st quarter of 2000 through the 4th quarter of 2006.¹¹

After restricting the analysis to the software publishing industry (NAICS code 511210), the sample contains 957 establishments corresponding to 877 unique firms (the vast majority of firms are single-establishment enterprises, and thus ‘establishments’ will be used inter-changeably with ‘firms’). Average establishment size is around 35 employees.¹² The number of software publishing establishments decreased from a high of 648 in Q4:2002 and Q1:2003 to a low of 581 in Q4:2006.¹³ However, industry employment increased from 16,600 to 21,000 over the sample years, reflecting an increase in average establishment size.

A key feature of our analysis is to identify the geographic scale at which agglomeration effects are relevant. If they dissipate rapidly over space, observing geographies at the sub-county level is critical.¹⁴ In principle, this would be accomplished by looking at county sub-divisions, such as census blocks, but these are of variable geographic extent and irregular shape. Ideally, we would like locations to be equally sized and defined in a purely spatial manner. Moreover, locations should only include plausible destinations for software establishments. At small scales, some areas do not offer basic conditions for the entry of software publishers. Including these areas is problematic in two ways. First, it is inconsistent with specifications that assume a positive expected number of software entrants. Second, it would result in a plethora of observations, only a tiny fraction of which would experience software establishment entry. In defining locations, we adopt the principle that human capital requirements can potentially be met either because required employees are already locally available or because they could be induced to move. However, the locations of firms are restricted by zoning laws, by the availability of suitable building stock, and by the presence of complementary infrastructure (e.g., roads, utilities). We posit that a software establishment cannot enter a location where these conditions are not satisfied.

Our approach to identifying potential locations is as follows. More than 90 percent of software estab-

¹¹The authors obtained these data under an agreement of confidentiality and, therefore, disclosure of the actual data is subject to certain restrictions.

¹²115 of these establishments have a single employee during our sample. Excluding these leads to results that are very similar to those in the main tables, and sometimes (marginally) more supportive of our thesis. Note that an establishment that only has one employee does not necessarily denote a completely different “type” of firm with no growth prospects. In our sample, there are 104 establishments that have one employee at some point in time, but then grow to have multiple employees (on average about 3.25).

¹³Fourteen software publishing establishments had PO Box as the official address, and their physical location could not be ascertained. These are not included in the above totals or in the subsequent analysis.

¹⁴Some of the variables are by necessity defined at the county level.

lishments share a building or address with other non-software establishments. The industries to which these other establishments belong are referred to as “control industries.” There about 700 industries on this list (at the 6-digit level) constituting most of the employment in Texas.¹⁵ We presume that any location that contains an establishment from these other control industries is one where a software establishment could potentially be situated (though we recognize in our estimation approach that the probability of doing so is not the same across all such locations). In other words, we assume that the physical infrastructure embodied is fungible across these industries and zoning laws will accommodate software publishing activities.¹⁶ Moreover, it is extremely unlikely that a software firm would locate in an area that does not contain at least one establishment on this long list of control industries.

There are 580,375 unique control industry establishment locations (including establishments that entered or exited during our sample period).¹⁷ We sort these establishments by latitude (south to north) and sequentially drop any that are closer than 2 miles to a previous establishment. We obtain a final sample of 9,299 establishments that form the center points of non-overlapping circle locations of one-mile radius.¹⁸ These locations cover about 11 percent of Texas.¹⁹ These locations are surrounded by concentric rings of 5, 10, and 25-mile radii.

When used in a panel, as in the entry analysis, the locations yield 251,073 observations (a balanced panel of 27 quarters; the first quarter is dropped because of lagging of some variables). In Table 1 Panel B, we report the summary statistics for these observations. We note that nearly all software establishments are less than a mile from a control industry establishment, and thus these control industries effectively map

¹⁵A large number of industries share facilities with software establishments because many industrial and agricultural firms have offices located separately from production facilities. These offices can share a building with software firms. A weighting scheme, described at the end of this section, mitigates issues when these industries are used to construct control variables.

¹⁶This does not imply that if a software establishment in some location shares a building with an advertising agency, then every building housing an advertising agency could potentially house a software establishment. But most such buildings would be a potential location or would be close to other buildings that would be suitable locations.

¹⁷There were 11,791 establishments with Post Office addresses. The location of almost all of those were obtained via batchgeo.com, and are thus utilized to control for localized economic activity and infrastructure.

¹⁸If we had not dropped establishments closer than 2 miles to a previously selected establishment, then the one-mile radius locations would overlap. Establishments do not need to be in the center of anything to be location centroids. This is an appealing feature. We want random locations, not locations that are central to a cluster. The set of locations is neither “sparse” nor “dense”: There is no way to identify another establishment that could be the centroid of a non-overlapping one-mile circle, but our set of non-overlapping circle locations does not maximize the number of included establishments or covers the largest area.

¹⁹The final list of establishments depends on how the initial sort is made, as the maximal set of establishment locations with mutual distances that exceed 2 miles is not unique. We have also performed the sort by longitude (from east to west); the covered area and key summary statistics differ very little.

out the large majority of potential locations.²⁰ However, after we drop locations to eliminate overlap, some spatial gaps are created, and thus the proportion of software establishments outside the final set of 1-mile circle locations rises to 28%.²¹ Figure 2, which shows the set of locations and software establishments in the Dallas area, is illustrative of this. The establishments in the “gaps” are lost from the left-hand side of the analysis, but nearly all establishments are incorporated on the right-hand side because they are within 25 miles of location centers. This is akin to explaining the behavior of a random sample of locations and establishments but, conditional on this sample, no errors-in-variables biases are introduced.²² We use two measures of software firm activity within each geographical area: (i) the number of establishments, and (ii) the number of employees in these establishments.

The control industries are not only used to define locations but also to control for the baseline propensity of software establishments to situate (or enter) there. Such control is necessary because locations with infrastructure that makes software establishment location possible differ in the extent to which they possess that infrastructure. Locations that contain a large number of control industry employees would, all things equal, be expected to contain more software establishments. Less developed locations, with a smaller number of control industry employees, are likely to contain fewer software establishments. Failure to control for this baseline might generate spurious clustering of entering establishments since a developing area attracts more entrants, including more software entrants. Thus, a positive spatial association at very small distances may be an artifact of land development patterns (or, equivalently, the abandonment of commercial land that has become economically less desirable).

We control for the baseline propensity of being situated in (or entering) a given location in two ways. The first way is to use the number of employees in the control industries as a control variable. The second way is to use the weighted number of employees as a control variable, with weights obtained from how frequently

²⁰Only 14 establishments are further than a mile from the nearest control firm facility. Most of these are one person operations, likely home-based.

²¹These establishments do not appear to be different than those within our locations, which is expected given our sampling procedure. The average establishment employment in the first quarter it is observed in our sample is 29.3 for those in our locations and 35.3 for those in the gaps between them. The corresponding standard errors are approximately 5 and 7, respectively. The standard error of the difference is larger still, and therefore, the difference in the means in average employment is far from being statistically significant; in fact, it is smaller than its standard error.

²²The only way to avoid gaps that potentially contain software establishments is to divide Texas into a fine grid, creating squares of (say) one-mile width, and dropping those that do not contain any control establishment. However, it is geometrically challenging to define neighborhoods around square locations that contain all points no more than a specified distance from the edge and center of the location, and not contain points that are further than those pre-specified distances.

establishments of each industry are co-located with software establishments. In this second approach, the employees of each industry have an industry-specific weight, which is the fraction of facilities in that industry that are co-located with software facilities. We find the second approach more appealing and report the results using the weighted control (results are robust to using the unweighted number of employees).²³

Key features of the data are provided in Table 1. Within one mile of an existing software establishment, there are, on average, approximately ten other software establishments employing 394 employees. At distances between 1 and 5 miles, there are an additional 48 software establishments employing 1,857 employees, a substantial drop off in the density of software firm activity (recall area is proportional to the square of distance). The next 5 and the following 15 miles contain 46 and 72 software establishments employing 1,810 and 2,525 employees, respectively, corresponding to ever-larger density drop-offs. Control industry employment goes from approximately 16,500 within a mile of a software establishment, to 155,000 in the next 4 miles, 314,000 in the next 5, and 303,000 in the next 15.²⁴ At small scales (up to 10 miles), the density of employment of software firms drops off proportionately faster than that of the control firms.

Considering firm characteristics, we include a measure of the establishment’s exposure to university R&D, which captures the possibility that knowledge spillovers are available from research universities.²⁵ This variable is based on the total federally funded research expenditures at the university located closest to the establishment, using the main address for the university campus. To introduce distance decay in the university R&D expenditures, we deflate total R&D expenditures by distance in miles (minimum 1 mile).²⁶ We distinguish between research universities and colleges/junior colleges. College funding is treated similarly to university funding. We include junior college funding since previous studies have found this variable to be more important than university research funding in explaining high-tech firm location (Abramovsky, 2007).

We proxy the quality of the labor pool by local recreational and cultural amenities.²⁷ As Woodward et

²³Our baseline measures are similar to those in Ellison and Glaeser (1997), after “shutting down” the channel of location-specific natural advantages. Using collocation weights to construct the baseline activity level also accounts for possible spillovers from other firms, even though we expect these synergies to be small (see Arai et al., 2004).

²⁴These averages are high because they are taken at the establishment-cross-quarter level. With most software establishments being located in the metropolitan areas (especially Dallas Metroplex and Austin), the typical urban density dominates the data.

²⁵Data on annual University R&D expenditures were obtained from the National Science Foundation (NSF). We aggregate total federal awards by all granting agencies, i.e., DoE, EPA, DoD, by each geographically distinct institution, i.e., system campuses are scored geographically separately.

²⁶University and college R&D activity is also correlated with the presence of highly educated workers and high tech activities. These are part of the “intangible assets” whose availability in a location can lead to agglomeration economies, as shown in Artis, Miguelez, and Moreno (2012).

²⁷Values of county variables are assigned based on the co-ordinates of a location’s centroid.

al. (2006) suggest, cultural and natural amenities are important in attracting and retaining skilled workers. To measure the local presence of these amenities, we compute the share of county employment in NAICS 71 (Arts, Entertainment, and Recreation), 721110 (hotels and motels), 722110 (full-service restaurants), and 722410 (drinking places, alcoholic beverages) as reported in the QCEW data set. De Silva and McComb (2012) have used a similar measure.²⁸ To account for factor costs, we use the average quarterly payroll of high-tech industries in the county.²⁹ The county unemployment rate for the final month in each quarter, as reported by the Texas Workforce Commission, proxies for economic conditions.³⁰ In the bulk of the regression analysis, we exclude establishments that have only a single employee throughout the sample period.³¹

3.2 Some Facts about Spatial Persistence of Software Industry Employment

The software industry is dynamic. There is a large turnover of establishments and substantial changes in the scale of establishments over time. For example, approximately half of the establishments that were in operation at the start of 2000 had exited by the end of our sample in 2006, and many of the establishments that did not exit experienced large drops in employment. All in all, from the 16,645 jobs that were initially in our dataset, only 7,015 (or about 41 percent) persisted in the same establishment until the end of the period.³² The persistence of jobs is not the same as worker turnover; a worker who leaves an establishment and gets replaced by another worker at that same establishment during our sample period registers as retention of that job by that establishment.

During the same period, 431 new establishments entered, almost all in different buildings than exiting establishments.³³ The jobs created by the entrants and the jobs added by growing incumbents raised total

²⁸There are other ways to account for overall labor quality (see Artis, Miguelez, and Moreno, 2012). Our findings are robust to completely dropping this variable from the regressions.

²⁹For the few counties with no high-tech industries, we used the average wage of NAICS 3133 (Manufacturing), 51 (Information), 52 (Finance and Insurance), 54 (Professional, Scientific, and Technical Services), 55 (Management of Companies and Enterprises), 61 (Educational Services), and 62 (Health Care and Social Assistance).

³⁰Land prices at the county level are a poor proxy for rental costs of software publishers, so they are not used in our analysis. Note that an area may be more “productive” for software publishers, but not in general (and may, in fact, be less productive for other firms). With software publishers being a small fraction of employment, local rents are unlikely to be significantly bid up by their presence.

³¹There are only 8 such establishments, but 5 are entrants. They are included in the entry analysis, since they may inform the location decision.

³²This percentage is equal to $(\sum_i \min\{emp_{i,1}, emp_{i,T}\})/(\sum_i emp_{i,1})$ where $emp_{i,1}$ is an establishment’s employment in the first quarter of our sample (and 0 if the establishment entered at a later date) and $emp_{i,T}$ is an establishment’s employment at the last quarter of our sample (and 0 if it exited by that quarter). Note that by construction this ratio cannot exceed unity.

³³Establishments are identified by their address. In the rare case of a firm merger, constituent establishments remain distinct.

employment in the industry to about 21,000 (an increase of 24%). This means that only one-third of the jobs at the end of the period were jobs that existed in the same establishment at the start of the period. Given that this industry does not rely on specialized infrastructure and has non-localized demand for its product, and given that entrants typically choose different addresses than incumbent or exiting establishments, there is a potential for the spatial distribution of this industry to be completely transformed in the timeframe of this study.

This turns out not to be the case. At the macro level, a quick way to assess whether spatial concentration has increased is to investigate whether the share of employment in, say, the top five counties has increased over this period. The list of the top five counties has remained the same (Dallas, Travis, Harris, Collin, and Bexar). The number of software publishing employees in those five counties has increased at approximately the same rate as in the entire state of Texas, marginally raising their combined share from about 89% at the start of the period to about 90% at the end. While these figures indicate that concentration is broadly constant at the county level, they do not provide any direct evidence about concentration at the 1-mile radius level. Some evidence at the establishment level can be provided by a pictorial examination of a few representative areas. One cannot easily plot employment into space, but can plot establishments. In Figures 3 and 4, we have plotted the software publishers and the control firms in Austin and North Dallas areas for the first and last quarter of our sample. The distribution of control firms indicates the areas where software publishers could be located. Note that software publishers are not uniformly distributed in this space, but are rather concentrated in certain areas. Moreover, areas of concentration remain stable despite entry and exit: an area with prior concentration of software publishers retains that concentration to the end of our sample period. A final observation relates to newly developed locations. Some areas were commercially undeveloped in the second quarter of 2000 and got developed by the end of 2006. But these areas did not attract a cluster of software establishments, and isolated software firm entry in them is rare (only 5% of software entrants locate more than 5 miles from an existing software facility).

Some additional evidence at the micro level is obtained by computing the number of software employees in each of the 1-mile radius locations and then seeing how that number changed over our sample period. The top ten 1-mile radius locations in terms of initial employment contained 61% of software employees in our initial quarter. The corresponding percentage at the end of our sample period is 65%, an increase in concentration despite an increase in the number of establishments and the number of employees. However,

there has been some churning of the top 10 locations. The percentage of software employees employed at the end of the sample period in the top 10 locations measured by employment at the start of the sample period is 48%, a decline from the initial level, but still a remarkable persistence in such a dynamic industry. Removing from all calculations one outlier location which jumped to first place from outside the top 10 list due to the entry of a large facility, the percentage increases to 57%, essentially unchanged from the initial level of 61% despite the dramatic changes in the industry at the establishment level.

4 Econometric Analysis and Results

4.1 Spatial Persistence in Employment and Its Components

A more formal analysis would perform the following conceptual exercise. Suppose we could exogenously increase the software employment in a particular location at the start of the period by one percent (e.g., by increasing the order flow of the establishments situated there). What would then be the percentage increase in software employment in that location at the end of the sample period?³⁴ Clearly, we cannot observe the impact of exogenously created software jobs being in a location. Rather, we observe locations that differ in the level of initial software employment and other characteristics. Using locational controls based on industries that are typically co-located, and noting that the software industry has minimal infrastructure demands and a non-physical product with a national market, we ensure that the initial variation in employment comes close to being exogenous in the statistical sense.³⁵

To better understand the content of the question we are investigating, we observe that end-of-period employment in a location is a function of jobs lost and accrued. In all of this analysis, a job in an establishment is retained regardless of employee turnover. For the moment, suppose that every job in any establishment has the same probability of being lost and that job accrual is proportional to the initial number of jobs in a location. Then the elasticity of final period employment with respect to initial employment is 1. However, such an extreme degree of persistence is unlikely. Entrants are more likely to locate in places with more existing software jobs (conditional on overall economic activity), but not proportionately more likely. Similarly, establishment job growth is not proportional to establishment size (as shown by the literature

³⁴The same hypothetical question could be expressed in terms of employment levels rather than percentages, in which case it would correspond to analysis using the number of employees rather than their log.

³⁵Initial variation in employment need not be random, but only uncorrelated with unobserved shocks to end-of-sample period employment.

on Gibrat’s law), and hence final employment in a location is not expected to be proportional to initial employment in that location. As a result, the landscape of the software industry activity would shift over time. Some areas with high concentrations would “revert to the mean” or closer to a level of software industry activity that is in proportion to overall economic activity in that location. Other areas with limited activity might exhibit higher concentrations for idiosyncratic reasons.

Spatial persistence would be higher than this benchmark if establishments located in areas with a high concentration of software industry activity grew systematically faster than establishments located in areas with a low concentration. The degree of spatial persistence of employment would balance these factors. In the absence of agglomeration economies, we would expect spatial persistence to be driven solely by the inertia of jobs at the establishment level.³⁶ With agglomeration economies, we would expect it to be higher than this value. Indeed, the observed fraction of initial jobs that remain with the same employer is so low as to be consistent with even a *negative* elasticity.³⁷ Thus, even a demonstration that the elasticity of final employment to initial employment is positive has empirical significance.

We now describe the econometric framework through which we investigate the extent of spatial persistence. Our spatial unit of econometric analysis is the one-mile radius locations described in section 3.1. For each of these locations, we estimate the end-of-period employment in the software industry as a function of initial employment and other initial conditions. We use as initial conditions the number of software establishments in the location at the start of the sample period, the number of software establishments in concentric 1-5, 5-10 and 10-25 mile rings and the number of employees of these establishments. Though the 1-mile radius locations are not overlapping, the surrounding rings are.

County effects for the five counties with major employment in this industry are included in some specifications.³⁸ For these regressions, the effect of initial conditions is identified from the within-county distribution of software publishing activity. We include the (weighted) number of employees in the set of control industries in the initial and final periods as explanatory variables in several specifications. In addition to these exogenous variables, we sometimes also include the final period number of establishments and employees in

³⁶This assumes that locations are “small” as a fraction of industry employment.

³⁷Consider two establishments (or locations), one with initial employment of 100, the other with initial employment of 200. Let the final period employment levels be flipped. Then, the number of retained jobs is 100 for each establishment, or a total of 200 (67 percent) for the industry. But the elasticity of final employment with respect to initial employment is negative!

³⁸Using country fixed effects for all counties effectively “dummies out” most locations situated in those counties, without major changes in the findings.

the concentric rings surrounding a location. These latter variables are co-determined to some extent, though the stronger influence probably goes from the larger outer rings to the center (note that concentric rings are mutually non-overlapping, i.e., outer rings do *not* include the employment of the inner rings). More formally, our most general model is given by

$$Y_{l,T} = D'_{l,t=1}\delta_0 + R'_{l,t=1}\rho_0 + C'_{c_l,t=1}\varphi_0 + D'_{l,T-1}\delta_1 + R'_{l,T-1}\rho_1 + C'_{c_l,T-1}\varphi_1 + \eta_l \quad (1)$$

where $Y_{l,T}$ is the location l 's quarterly software employment in the final quarter, T (Q4:2006), or its natural log.³⁹ The vectors $D_{l,t=1}$ and $D_{l,T-1}$ represent initial and final period (minus a lag) density of software publisher employment variables for location l (in some specifications, we also include density of software publisher establishments).⁴⁰ Similarly, we include initial and final period (minus a lag) location variables ($R_{l,t=1}$ and $R_{l,T-1}$), and variables specific to the county c_l where this location is situated ($C_{c_l,t=1}$ and $C_{c_l,T-1}$).⁴¹ Employment has a long right tail, and many of the explanatory factors are expected to act synergistically rather than in a purely additive fashion. Thus, an analysis using log-employment as the dependent variable is appealing in this context, though we report results using both specifications. Only locations with positive initial software industry employment are included in this regression. The lowest possible employment level is 1, so $Y_{l,T}$ is censored and equation (1) is estimated via maximum likelihood assuming a normal distribution for η_l and accounting for censoring.⁴² We base all inferences on robust standard errors for parameters and marginal effects.⁴³

Table 2 presents the results using log employment as the dependent variable, moving from the simplest to the more complicated specifications. Since our focus is on the marginal effect of initial employment on the

³⁹We observe no seasonal patterns in software employment, and therefore there is no gain in truncating our sample so that the initial and terminal quarter are the same. In contrast, dropping the initial or terminal quarters would have resulted in the loss of information by basing our estimates on a shorter time period.

⁴⁰These variables are counts, but since all sampling geographies are identical in shape and area (i.e., 3.14, 75.4, and 238.6 square miles at 1, 5 and 10-mile radii), they are effectively normalized densities across observations.

⁴¹Location variables include the measure of baseline economic activity (either weighted by its association with software establishment presence or un-weighted) that serves as a control for the expected presence of software establishments, and the proximity to college/university R&D. County characteristics for the county in which the location is centered include the presence of cultural and recreational amenities, the county unemployment rate and the measure of localized factor costs.

⁴²The censoring threshold in the log linear specification is 0, while in the linear specification it is 1. We distinguish between locations with one employee and censored locations with no software employment.

⁴³Our estimation strategy accounts for non-spherical disturbances and censoring, but not for possible omitted variable biases (or other sources of endogeneity). Any instrumentation strategy would be based on what would likely to be weak and coarse instruments, which would sacrifice the high spatial resolution of our data. Most other studies do not have access to such high-quality establishment data on colocation of facilities of different industries, which we use to construct our baseline index, and we believe that using this information is preferable to the alternatives.

expected value of final employment, this table (and all tables with employment as the dependent variable) report marginal effects and the associated robust standard errors. Table 3 presents the results using the same sets of regressors but with linear rather than log specifications.⁴⁴

We discuss the results of Table 2 and Table 3 together. In the simplest specification (model 1), no covariates are used except for differential intercepts for major counties. The point estimates suggest that a one percent increase in initial employment translates into a 0.60 to 0.75 percent increase in final period employment. This indicates a large spatial persistence in line with the informal evidence described earlier. The addition of software employment in close and moderate proximity and other location characteristics (model 2) slightly reduces these point estimates. Interestingly, while location characteristics are jointly significant, employment at any distance up to 25 miles is not. Adding current location conditions to the regression (model 3) shows that final employment in a location increases with initial proximate employment but decreases in final proximate employment (lagged by one quarter). Because employment in a 1-mile-radius location is typically a small fraction of the employment in the surrounding 25 radius ring, we posit any causal effects go mostly from the surrounding area to the inner circle. Thus, locations in moderate proximity to an existing software cluster grow faster, but those near a growing cluster grow slower. Knowledge and other productivity-enhancing spillovers are unlikely to yield this pattern (software activity, whether initial or final period, should increase employment in a location), but can easily be rationalized with a labor pool thesis. Establishments in a location proximate to areas with many software employees can obtain a ready supply of workers, but if those proximate areas are growing themselves, they can siphon workers away from that location.

The last reported specification (model 4) augments model 2 by adding the number of establishments at various distance thresholds as explanatory variables. In these regressions, agglomeration effects are decomposed to those arising from an increase in the number of facilities (holding employment constant) and those arising from an increase in employment (holding the number of facilities constant). Significance of the former is mixed, but the point estimates economically meaningful (significance is stronger when adding the establishment variables to model 3). The log specification, where the coefficients are both elasticities, suggests that a “replication” of the facilities in a location would double end of period employment in that

⁴⁴To make a valid comparison between the marginal effects in the two tables, those of the regressions in Table 3 are weighted by the initial number of employees in each location, i.e., they are the sample average of $\frac{Y_{l,t=0}}{\bar{Y}_{t=0}} \frac{\partial E[Y_{l,T}]}{\partial x_l}$ where $Y_{t=0}$ is the initial software employment in Texas and x_l is any regressor.

location, i.e., there is nearly full persistence.⁴⁵ Part of the explanation of why the number of initially active establishments is associated with higher terminal employment, which we have verified in our data, is that it increases the probability that some of these establishments grow substantially and reduces the possibility that all of them fail.

It is also worth noting, as indicated in both tables, that one quarter of the observations are censored. This might seem high, but about 70 percent of the locations initially have only one establishment and these locations form the bulk of those that are censored (i.e., the sole establishment in that location exited and was not replaced). Given that half of the incumbent establishments fail, a 25% censoring rate implies that in many locations with only one establishment, at least one other establishment was attracted before the exit of the incumbent. This suggests a propensity of entrants to locate close to incumbents, which is formally demonstrated in section 4.2.

In light of the earlier discussion, the magnitude of the spatial persistence of employment appears “large.” But to better assess how much larger it is compared to a benchmark of no locational advantage from any initial industry presence, we compare the results in Tables 2 and 3 with the counterpart specifications estimated at the establishment level, rather than the location level.⁴⁶

Examining models 1 and 7 of Table 4, which use the same specification as model 1 of Tables 2 and 3, the estimate of the coefficient of initial establishment employment on final period employment is much smaller than those obtained at the location level.⁴⁷ Increasing initial employment in an establishment by one percent leads to only 0.3 to 0.4 percent increase in final employment. Models 2 and 3 (for logs) and 8 and 9 (for levels) are progressively more inclusive specifications and confirm this pattern. These results corroborate the conclusion that establishment-level employment persistence is much lower than location level persistence.⁴⁸

⁴⁵Replicating the facilities results in doubling the number of establishments and doubling the total employment. The implied end-of-period effect from Table 2, column 4, is $0.483 + 0.467 = 0.95$. Adding establishment variables to model 3 (in results not reported for brevity), the corresponding figures are $0.496 + 0.606 = 1.102$.

⁴⁶Note that the intercept of the location employment and establishment employment regressions are not comparable. If the elasticity with respect to initial employment were fixed to 1, the intercepts in the former models reflect the growth rate of software employment in a typical location, while in the latter models they reflect the growth rate of a typical establishment.

⁴⁷A coefficient of 1 would imply Gibrat’s law. Most of the literature investigating Gibrat’s law has found a negative association, e.g., a mean reversion effect where large firms grow slowly while small firms grow faster (see early work by Evans, 1987, and Dunne, Roberts, and Samuelson 1989, as well as later work by Hart and Oulton, 1996, and Dunne and Hughes, 1994).

⁴⁸These estimates are not driven by small establishments that contribute little to aggregate employment; re-estimating Table 4 using initial employment weights yields very similar results. The establishments included in these regressions are those that were active at the start of our sample period, i.e., the 526 incumbents plus 2 entrants that entered at the start of 2000, and thus the number of observations is smaller than the total number of establishments reported in Table 1.

Conceptually, the difference between establishment-level and location-level employment persistence consists of three components: (i) faster growth rates of establishments in high concentration locations, (ii) the capture of jobs lost to an establishment by others in that same location, and (iii) higher entry rates of firms in locations with prior software presence and hence the creation of more jobs by entrants in high initial employment locations. The first component is evidence of positive spillovers from co-location, yielding higher employment growth among establishments present in a location with many other co-located establishments. If that were the case, then exogenously increasing the employment level of a single establishment would lead to a smaller increase in its final employment than exogenously increasing the employment levels of all establishments in a location.

We investigate this key question using establishment-level regressions by re-estimated models 1-3, and 7-8 of Table 4 after adding the initial employment by co-located establishments as a regressor (plus 1, when taking the log). The results, reported in columns 4-6 (for logs) and 10-12 (for levels), suggest that there are no such positive synergistic effects. In fact, if anything, there seems to be a negative effect from the presence of other firms (measured by their employees) on an establishment's final period employment.⁴⁹ With the addition of both number and employment of other establishments in the specification, the sign on the number of other establishments is positive but not significant. However, the coefficient on the number of employees of those establishments becomes even more negative. This means, then, that the effect of initial employment on final employment does not arise from a positive effect that initial employment has on the growth of the existing establishments. It rather arises because more of the jobs that existing establishments lose are captured by co-located establishments and because more jobs are created by entrants. This is a key piece of evidence, the implications of which we discuss later.

We next examine whether the local capture of jobs lost by an establishment is indeed a contributor to local employment persistence. This is done by re-estimating the models 1-4 in Tables 2 and 3 using as dependent variable the final employment in a location's establishments that were present at the start of the sample period (incumbents). The results, reported in columns 5-8 of these two tables, show that the coefficients on initial employment are higher than those in Table 4. This indicates that the negative effect of co-location on individual establishment growth is outweighed by the tendency of jobs to remain in areas

⁴⁹It is conceptually possible that there are positive synergies, and that these attract entrants of lower intrinsic productivity to high-density locations, resulting in this negative association. We have no direct measures of establishment productivity to test this proposition, but the null results of the duration analysis in section 4.2 suggest these effects are not material.

with larger prior employment.⁵⁰

More importantly, the coefficient on initial employment when the dependent variable is the end period incumbent establishment employment is lower than the elasticities reported when the dependent variable is the end of period total employment. For example, column 1 of Table 2 yields an elasticity of 0.605, while the elasticity for column 5 is 0.528, and it is 0.366 at the establishment level (Table 4, column 1). For the linear version of this model, the corresponding elasticities measured by the weighted marginal effect are 0.753, 0.539, 0.283, respectively, yielding substantially larger differences. Similar results are obtained when comparing the other models in these tables. The statistical significance of these differences can be tested by estimating each pair of regressions via SUR; in results we omit for brevity, the difference is always statistically significant at the 5 percent level for the linear specifications. For the log-linear specifications, the differences are not statistically significant. However, one can make a case that the linear specification is more natural, decomposing the effect of having an extra worker employed in a software establishment to increased final period by that establishment, increased employment by other incumbent establishments, and increased employment by all establishments (including entrants).⁵¹ To summarize, the employment persistence in a locality is only partially driven by incumbent establishments capturing jobs lost to other incumbents. Some of the persistence must be driven by disproportionate entry rates into locations with greater initial presence of software publishers.

4.2 Establishment Entry and Survival

In this section, we examine the entry process of establishments to confirm its impact on the spatial persistence of industry employment. Establishment entry is defined as the introduction date of a new Enterprise Identification Number (EIN) in our dataset. Our primary aim is to understand the extent to which the

⁵⁰For an illustrative example, consider four establishments, A, B, C1, and C2, operating in locations A, B, and C. Suppose establishments A and C1 have initial employment of 100 while establishments B and C2 have initial employment of 50. Let there be a simple process of mean reversion, whereby the large establishments lose 10 workers who get hired by the small ones. Let establishment C1 lose another 5 workers, so that establishment employment growth is negatively associated with location employment. Then, the elasticity of final employment with respect to initial employment is around 0.54 at the establishment level, and around 0.78 at the location level, because many jobs lost by the large establishment in the high employment location are balanced by gains in the small establishment in that same location.

⁵¹In fact, the log linear specification likely overstates the degree in persistence, because a proportional reduction in the employment of all incumbents would appear as perfect persistence in the regressions using incumbent employment. The statistical significance of the difference between the coefficients in Table 4 relative to those of Table 2 and 3 cannot be tested via SUR because Table 4 uses establishment-level data; however, the establishment level elasticities have the biggest differences from the location-based counter-parts and have the smallest standard errors, so we expect all such differences to be statistically significant.

prior presence of software publishing activity in a location influences the entry rate of software firms in that location. Since other factors are associated with entry rates, we must carefully formulate the econometric design. Even under the null hypothesis that localization economies do not influence entry probabilities, the expected number of entrants is not uniform across all one-mile radius locations. We estimate the expected number of software firms entering into any of these locations as a function of location characteristics similar to those used in the analysis of employment persistence. Of these characteristics, the number of employees in other related industries, weighted by each industry’s spatial association with software publishing firms, is the most important control variable. This is not a causal relationship, but serves as a summary statistic for the baseline propensity of software firms to enter that location.

An important difference between the employment persistence analysis and the entry analysis is that the latter takes advantage of the time variation in our data. We estimate the number of establishments entering a location in a particular quarter as a function of the location characteristics in the preceding quarter. Doing so increases the variation we can exploit in our sample. In the typical entry case, there is a single software entrant in a given location in a given quarter, while the maximum number of entering software establishments is two (except for one occurrence of 5). Given this, we estimate the entry models using an ordered probit. Our dependent variable is the number of software publishing start-ups $y_{l,t}$ in a given one-mile radius location l during a given quarter t . The basic ordered probit model is

$$Y_{l,t}^* = D'_{l,t-1}\delta + R'_{l,t-1}\rho + C'_{c_l,t-1}\varphi + \tau(t) + \alpha_{c_l} + \epsilon_{lt} \quad (2)$$

where $Y_{l,t}^*$ is a continuous latent variable with two threshold points, one delineating no entry from entry by a single establishment and the other delineating entry by one establishment from entry by two or more establishments. The independent variables can be classified into three main groups: $D_{l,t-1}$ represents software industry activity/density measures in the location, $R_{l,t-1}$ controls for other location-specific variation, $C_{c_l,t-1}$ controls for county-specific characteristics of the county c_l where location l is situated, and $\tau(t)$ is a quadratic function of time (the model cannot be meaningfully identified with time fixed effects). Most of the variables characterized by long tails are in logs. When taking the log of the number of establishments or employees, the value of 1 is added to ensure well defined regressor values for a location with no establishments. In some specifications we have included county effects α_{c_l} for the top 5 counties with the most software publishers.⁵² The random disturbance ϵ_{lt} has a standard normal distribution.

⁵²The remaining counties are pooled together as the excluded category. We use county effects rather than location effects to

The results are reported in Table 8. Localized own-industry density has a strong positive effect on entry probabilities regardless of whether the number of software establishments or the associated number of employees is used as a measure of density. The effect appears to be stronger for distances of less than one mile, somewhat important for intermediate distances (1-5 miles) and marginal or absent for distances greater than 5 miles (especially for more comprehensive specifications). The use of county fixed effects tends to weaken the association between pre-existing software establishments and subsequent entry, but only marginally, while complementing these fixed effects with time-varying location characteristics tends to have no effect. Reassuringly, our control for localized activity by firms using similar infrastructure (Other Industry Employees) is positive and strongly significant. The number of software establishments seems equally important for the location decision of potential entrants as the total number of employees of those establishments (measured by statistical significance). Moreover, the two sets of models have an approximately equally good fit as measured by the log likelihood, with the specification that uses the number of establishments as the measure of software publisher presence having a slight edge.

Finally, when both the number of establishments and the number of their employees are used in the regression, significance drops substantially, especially for the number of employees present. For this reason, we attach greater importance to the results in the specifications (1) through (6) where either the number of establishments or the number of employees of these establishments is used as a control. From the remaining variables, high tech wages and, to a lesser extent, university spillovers are associated with higher entry probabilities (the former perhaps as a proxy for demand for high tech labor) while junior college spillovers and local unemployment rates are negatively associated with entry.⁵³

Despite the small entry counts, we also estimate the entry process using the Poisson model to investigate the sensitivity of the results to the econometric specification. The estimation results are not reported for brevity, but are summarized here. Though the parameter values of the Poisson and ordered probit models are not directly comparable with each other because of differences in normalization, statistical significance and relative magnitudes can be compared. On this basis, the Poisson and ordered probit estimates for the agglomeration variables and their effect on entry probabilities are mostly similar. However, as with control for unobserved location heterogeneity, because repeat entries into any particular location are rare. Moreover, adding dummies for each county results in the loss of many observations.

⁵³We have also estimated these models using ordered logit, obtaining qualitatively similar estimates for the agglomeration variables. The only noticeable difference is in the rate at which spillover effects decay with distance, which tends to be smaller under the ordered logit than the ordered probit.

distances of up to 1-mile, localized own-industry density effects now appear slightly stronger for the 1 to 5 mile distance. But this holds only when employment levels are the measure of industry activity.

We next investigate whether establishments entering areas with high prior software firm concentration have higher survival rates. We estimate a number of different specifications of Weibull duration models, using variable sets that parallel those of entry models (though variable values correspond to the date of establishment entry). We find no statistically significant association between the number of other software establishments or their employment and the exit hazard rate of an establishment.⁵⁴ These results are corroborated by a discrete choice model of the exit decision (both sets of results are available upon request).

5 Interpretation of Results and Concluding Remarks

Software industry employment at the highly localized level (one-mile radius) exhibits excess persistence, i.e., greater persistence than implied by the employment persistence of individual establishments. This persistence is not driven by the faster employment growth of establishments in areas with a high density of software publishing activity. Indeed, incumbent establishments seem to exhibit slower growth when located near other establishments. Moreover, the persistence is only partially driven by the fact that a job “lost” by an establishment is “captured” by another in the same location. Rather, the presence of software firms increases the propensity of other software firms to enter within close proximity. But establishments entering localities with prior software firm presence do not experience differential survival rates than establishments entering localities with no prior software firm activity.

This combination of findings suggests that a prior concentration of software establishments in a locality lowers the entry costs of other software firms in that same locality, but that post-entry profitability in that locality is not higher and may even be lower (if one focuses on employment growth). Recalling the framework developed in Section 2, this pattern is most consistent with spatial effects arising from the localized labor pool, including from firms preferring to locate in a particular location because recruiting is easier.

The competing explanation, that entrants locate in an area because of synergies or direct human capital spillovers from incumbents, does not fit the facts. Our conceptual framework (see Panel A of Figure A-1) reveals that such spillovers should make the average entrant into that location more successful compared to

⁵⁴Though our sample period is rather short, turnover is sufficiently rapid that exit is observed for 176 of the 423 entering establishments in our data. Note that entry is uniformly spaced during the sample period.

entrants elsewhere and should also benefit incumbents. There is no evidence of this at the one-mile range and, therefore, human capital spillovers cannot be the only source of spatial persistence. However, some human capital spillovers might be present in conjunction with labor channel effects. These human capital spillovers would reinforce the positive effects of labor pool effects on post-entry profits in high spatial agglomeration locations. Even though lower entry costs in those locations mean that software establishments co-locating with other software establishments have on average weaker business prospects, the upward shift in the payoff function would leave average exit rates unchanged (and employment growth of incumbent establishments only marginal lower).

We find no strong relationship between end-of-period software employment in a location and software employment at intermediate distances (1 to 10 miles) from that location. However, there is still a positive, but diminished, association between entry rates and prior software publishing activity. There continues to be no association between exit rates and software publishing activity. In other words, entry costs in a location appear lower if there is a nearby labor pool, but the surrounding pool does not enhance the productivity of software establishments in that location, which supports the labor pool interpretation we suggested earlier.

Finally, employment in a location decreases in the contemporaneous employment growth at larger distances (10 to 25 miles). This is suggestive of a “pull” effect for employees from proximate high growth areas. Also, there is only a tenuous positive relationship between entry rates in the location and employment at 10 to 25 miles. This evidence is also consistent with labor pool effects.

A key feature of the identification strategy is the construction of a baseline activity against which to measure the spatial concentration of software publishers. It is important to note that many findings do *not* rely on this benchmark. These include the spatial stability of software activity, the trends summarized in Figure 1, the high employment variability at the establishment level, the concentration of new entry near existing establishments, and the observation that new establishments do not simply replace existing ones in the same address/building. These facts, though, could be interpreted differently if the explanation for them was that these one-mile locations where software employment is concentrated were the only suitable locations for software establishments. It is this last notion that our baseline activity measure dispels. Because locations differ in many dimensions, and obtaining all relevant characteristics and measuring their impact on software publisher location choices is not a feasible exercise, this baseline measure of location suitability

we construct has elements of a propensity score.

Our findings suggest that spatial persistence is primarily driven by labor pool agglomeration effects. Productivity spillovers from knowledge transmission are of subsidiary importance. Firms are attracted to locations because of the existing labor force, not because entering there gives them better prospects through spillovers. In this industry, jobs go where the workers are, rather than the other way around.

This leaves cities that want to become a magnet for an industry in a conundrum. Offering incentives to attract a single establishment (even a large one) is unlikely to succeed. What is needed is a concerted effort to bring multiple establishments in a location in a short amount of time, which is a daunting task. The alternative strategy is to attract businesses whose success does not hinge on Marshallian spillovers or which require specialized labor. A general business-friendly environment may attract a diverse set of such businesses and eventually bring about self-reinforcing Jacobian-style externalities. A city does not need specialist industries to grow; it can succeed economically by becoming an ecosystem for a diverse set of generalist businesses.

Appendix: A Stylized Model

Consider all potential establishments in a location l and rank them on the basis of their expected present value gross of entry costs. Let this present value be given by $V(p) = \theta(p) + \kappa z_l + \lambda \theta(p) z_l$, where p is the percentile ranking of an establishment, $\theta(\cdot)$ is an increasing function, z_l is a measure of value-enhancing spillovers in location l , and κ and λ are positive constants. Let the fixed costs of entry of any establishment be $FC = \gamma_0 + \gamma_1 x_l$, where x_l is a measure of the entry barriers in a location and γ_0 and γ_1 are constants.

For any location, l , the establishments with a gross present value exceeding entry costs enters and operates; the rest stay out. The fraction of potential establishments that enter, i.e., the entry threshold percentile, is given by $\hat{p}_l = \theta^{-1} \left(\frac{\gamma_0 + \gamma_1 x_l - \kappa z_l}{1 + \lambda z_l} \right)$. This is decreasing in any value-enhancing spillovers and any location characteristic that reduces entry costs. In this framework, spillovers do not necessarily increase the present value of all establishments, but they do shift the distribution of present value so that a distribution with higher first-order spillovers dominates one with lower spillovers.⁵⁵

Figure A-1, panel A, illustrates within this framework how entry and establishment performance differ across two locations, one with low and the other with high knowledge spillovers. The horizontal axis plots the percentiles of potential entrants, and the vertical axis plots the entry costs and the gross present value. The upward sloping lines are the relationship between establishment present value and establishment percentile rank for low and high spillover locations. For this figure, $\theta(\cdot)$ is assumed to be linear. The line for the

⁵⁵For simplicity, fixed costs are the same for all establishments, but an extension that relaxes this is discussed in footnote 57.

high knowledge spillovers location is above the line for the low knowledge spillover location, indicating that knowledge spillovers are value enhancing, i.e., increase the value of z_l . The high z_l line is steeper, as $\lambda > 0$, suggesting a stronger positive effect at the high end of value distribution. The fixed costs of establishment entry are the same for high and low knowledge spillover areas, consistent with the premise that knowledge spillovers do not affect entry costs. The marginal entrant in the low and high spillover location is \hat{p}^{**} and \hat{p}^* , respectively. High spillover locations attract more establishments. Moreover, the average performance (value) of these establishments is higher. With the distribution of establishments over percentiles being uniform by construction, the mean expected value of the operating establishments is increasing in z_l .⁵⁶ The continuation value and (future) employment of an establishment evolves over time, but is positively correlated with its initial expected value. Thus, locations with high values of z_l are expected to have more and faster-growing establishments. When an establishment's continuation value drops below zero, it exits. With the distribution of establishment expected present values being higher in locations with high values of z_l , exit rates in those locations should be lower.

Panel B of Figure A-1 illustrates how entry and establishment performance outcomes differ across two locations, one with low and the other with high labor market agglomeration economies. Agglomeration effects arising from the labor channel reduce the cost of entry. They can also reduce ongoing recruiting costs and thus shift up the distribution of gross present value of establishments. In our notation, labor market agglomeration economies decrease x_l and increase z_l . The second effect is expected to be smaller. Whether or not this is the case, the important point is that the labor channel affects both set-up costs and also future performance. Panel B of Figure A-1 shows that the average present value of entering establishments in locations with high labor agglomeration may be lower, even though they attract more establishments (as demonstrated in footnote 56, the expected value conditional on entry falls as x_l decreases, but rises with z_l). In fact, we believe that it is lower since the entry effect is expected to dominate, but this depends on the relative size of the shifts of the two lines and the extent to which $\theta(\cdot)$ departs from linearity (and in which way).⁵⁷ Following the discussion in the preceding paragraph, the expected value of an establishment affects its future employment level and exit probability.⁵⁸ In markets with high labor agglomeration economies, then, establishment growth rates may be lower and exit probabilities higher than in locations with low agglomeration economies.

These two figures are both static representations and describe entry in a single period for a specific level of spatial agglomeration. Spatial agglomeration persists over time, but persistence is not perfect, as random variation in establishment outcomes weakens the initial advantage of some locations and strengthens that of

⁵⁶ $E[V(p)|p > \hat{p}_l] = (1 + \lambda z_l)E[p|p > \hat{p}_l] + \kappa z_l$. Observing that $\hat{p}_l = (\gamma_0 + \gamma_1 x_l - \kappa z_l)/(1 + \lambda z_l)$, and that by definition p is distributed uniformly, we obtain $E[p|p > \hat{p}_l] = 1/2 + (\gamma_0 + \gamma_1 x_l - \kappa z_l)/(2(1 + \lambda z_l))$. Substituting in the expression for the establishment's expected value and simplifying, we get $E[V(p)|p > \hat{p}_l] = (1 + \lambda z_l + \gamma_0 + \gamma_1 x_l + \kappa z_l)/2$ which is increasing in z_l .

⁵⁷Heterogeneous entry costs do not alter any of this discussion as long as establishments with high post-entry value enter "ahead" of those with low post-entry value. A more substantial modification of this framework, under which establishments have the same post-entry performance but differ only in the entry costs (represented by a flat post-entry cardinal performance curve and upward sloping entry cost curve), results in some meaningful changes in the comparative statics, but cannot be reconciled with the empirical results.

⁵⁸We note parenthetically that growth and exit are qualitatively different in that exit is a "tail event." Thus, an increase in the variance of outcomes may not affect mean establishment employment growth, while affecting survival probabilities.

others.

References

- [1] Abramovsky, L., R. Harrison, and H. Simpson (2007). "University Research and the Location of Business R & D," *Economic Journal*, 117: C114–C141.
- [2] Arai, Y., H. Nakamura, H. Sato, T. Nakazawa, T. Musha, and K. Sugizaki (2004). "Multimedia and Internet Business Clusters in Central Tokyo," *Urban Geography*, 25: 483-500.
- [3] Artis, Michael J., Ernest Miguelez, and Rosina Moreno (2012). "Agglomeration Economies and Regional Intangible Assets: An Empirical Investigation," *Journal of Economic Geography*, vol. 12, pages 1167-1189.
- [4] Arzaghi, M. and J. V. Henderson (2008). "Networking off Madison Avenue," *Review of Economic Studies*, 75: 1011-1038.
- [5] Buenstorf, G. and S. Klepper (2009). "Heritage and Agglomeration: The Akron Tyre Cluster Revisited," *Economic Journal*, 119: 705-733.
- [6] Ciccone, A. and R. E. Hall (1996). "Productivity and the Density of Economic Activity," *American Economic Review*, 86: 54-70.
- [7] Combes, P. (2000). "Economic Structure and Local Growth: France, 1984-1993," *Journal of Urban Economics*, 47: 329-355.
- [8] De Silva, D. G. and R.P. McComb (2012). "Research Universities and Regional High-Tech Firm Start-ups and Exit," *Economic Inquiry*, 50: 112-130.
- [9] Dunne, P. and A. Hughes (1994). "Age, Size, Growth and Survival: UK Companies in the 1980s," *Journal of Industrial Economics*, 42: 115-140.
- [10] Dunne, T., M. J. Roberts, and L. Samuelson (1988). "Patterns of Firm Entry and Exit in the U.S. Manufacturing Industries," *Rand Journal of Economics*, 19: 495-515.
- [11] Ellison, Glenn, and Edward L. Glaeser (1997). "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach," *Journal of Political Economy*, vol. 105, pages 889-927.
- [12] Ellison, G., E. L. Glaeser and W. R. Kerr (2010). "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns," *American Economic Review*, 100: 1195-1213.
- [13] Evans, D. S. (1987). "Tests of Alternative Theories of Firm Growth," *Journal of Political Economy*, 95: 657-674.
- [14] Fosfuri, Andrea, and Thomas Ronde (2004). "High-Tech Clusters, Technology Spillovers, and Trade Secret Laws," *International Journal of Industrial Organization*, 22: 45-65.
- [15] Freedman, M. L. (2008). "Job Hopping, Earnings Dynamics, and industrial agglomeration in the Software Publishing Industry," *Journal of Urban Economics*, 64: 590-600.
- [16] Glaeser, E. L., H. L. Kallal, J. A. Scheinkman, A. Schleifer (1992). "Growth in Cities," *Journal of Political Economy*, 100: 1126-1152.

- [17] Hart, P. E. and N. Oulton (1996). "Growth and Size of Firms," *Economic Journal*, 106: 1242-1252.
- [18] Henderson, J. V. (1994). "Where Does an Industry Locate?," *Journal of Urban Economics*, 35: 83-104.
- [19] Henderson, V., A. Kuncoro, and M. Turner (1995). "Industrial Development in Cities," *Journal of Political Economy*, 103: 1067-1090.
- [20] Hutton, T. A. (2004). "The New Economy of the Inner City," *Cities*, 21: 89-108.
- [21] Marshall, A. (1920). *Principles of Economics*, London: MacMillan.
- [22] Men, Jarle (2005). "Is Mobility of Technical Personnel a Source of R & D Spillovers?," *Journal of Labor Economics*, 23: 81-114.
- [23] Overman, H. G. and D. Puga (2010). "Labor Pooling as a Source of Agglomeration: An Empirical Investigation," in Glaeser, E. L., editor, *Agglomeration Economics*, chapter 4, pages 133-150, The University of Chicago Press.
- [24] Rice, Patricia, Anthony J. Venables, and Eleonora Patacchini (2006). "Spatial Determinants of Productivity: Analysis for the Regions of Great Britain," *Regional Science and Urban Economics*, vol. 36, pages 727-752.
- [25] Rosenthal, S. S. and W. C. Strange (2001). "The Determinants of Agglomeration," *Journal of Urban Economics*, 50: 191-229.
- [26] Rosenthal, S. S. and W. C. Strange (2003). "Geography, Industrial Organization, and Agglomeration," *Review of Economics and Statistics*, 85: 377-393.
- [27] Saxenian, A. (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Cambridge: Harvard University Press.
- [28] Staber, U. (2001). "Spatial Proximity and Firm Survival in a Declining Industrial District: The Case of Knitwear Firms in Baden-Württemberg," *Regional Studies*, 35: 329-341.
- [29] Wallsten, S. J. (2001). "An Empirical Test of Geographic Knowledge Spillovers using Geographic Information Systems and Firm-size-level Data," *Regional Science and Urban Economics*, 31: 571-599.
- [30] Woodward, D., O. Figuerido, and P. Guimarães (2006). "Beyond the Silicon Valley: University R&D and high technology location," *Journal of Urban Economics*, 60: 15-32.