Labour pooling as a source of agglomeration: An empirical investigation*

Henry G. Overman[‡]
London School of Economics

Diego Puga[§]
IMDEA and *Universidad Carlos III*

Preliminary, printed 28 November 2007

Abstract: We provide empirical evidence on the role of labour market pooling in determining the spatial concentration of UK manufacturing establishments. This role arises because large concentrations of employment iron out idiosyncratic shocks and improve establishments's ability to adapt their employment to good and bad times. We measure the likely importance of labour pooling by calculating the fluctuations in employment of individual establishments relative to their sector and averaging by sector. Our results show that sectors whose establishments experience more idiosyncratic volatility are more spatially concentrated, even after controlling for a range of other industry characteristics that include a novel measure of the importance of localized intermediate suppliers.

Key words: labour market pooling, spatial concentration.
JEL classification: R30, R12

*Thanks to seminar participants at the NBER Urban Economics Meeting in March 2007 for comments. We are extremely grateful to Roberto Picchizzolu for his help in providing us with data. This work contains statistical data from ONS which is Crown copyright and reproduced with the permission of the controller of HMSO and Queen's Printer for Scotland. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

‡Department of Geography and Environment, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom (e-mail: h.g.overman@lse.ac.uk; website: http://personal.lse.ac.uk/overman). Also affiliated with the Centre for Economic Policy Research, and the Centre for Economic Performance at the London School of Economics.

§Madrid Institute for Advanced Studies (IMDEA) Social Sciences and Universidad Carlos III de Madrid, Department of Economics, Madrid 126, 28903 Getafe, Spain (e-mail: diego.puga@imdea.org; website: http://diegopuga.org). Also affiliated with the Centre for Economic Policy Research.

1. Introduction

Spatial concentrations of establishments and workers offer great productive advantages. Many modern econometric studies have confirmed and quantified this important stylized fact. Estimates of the productivity increase from a doubling in the size of an agglomeration range between two and eight percent, depending on the sector and details of the estimation procedure (see Rosenthal and Strange, 2004, Combes, Duranton, Gobillon, and Roux, 2007).

Unfortunately, the literature has been far less successful at distinguishing between the different possible sources of urban increasing returns than at quantifying their magnitude. We have sound theoretical models providing microeconomic foundations for the economies of agglomeration, but the different mechanisms are had to distinguish empirically. The main difficulty is what Duranton and Puga (2004) call the "Marshallian equivalence" of these theories: they all predict an increase in productivity with spatial concentration, but work through mechanisms that are hard to trace.

This paper focuses on a potential source of agglomeration economies to which Alfred Marshall (1890) already devoted particular attention: labour market pooling. While there are various interpretations of labour market pooling as a source of agglomeration economies, Marshall emphasized that "a localized industry gains a great advantage from the fact that it offers a constant market for skill" (Marshall, 1890, p. 271). In a model we review in section 2, Krugman (1991) formalizes this argument by considering an industry where plants experience idiosyncratic shocks. Individual profits are convex in the plant-specific shock, since each plant responds to the shock by adjusting its levels of production and employment. However, changes in the plant's employment affect local wages, and the effect is greater the more isolated the plant is from other plants in the same sector or using similar workers. If wages are higher when the plant wants to expand production in response to a positive shock and lower when it wants to contract production in response to a negative shock, this limits the plant's ability to adapt its employment level to good and bad times. Consequently, establishments that tend to experience substantial changes in their employment relative to other plants using workers with similar skills will find it advantageous to locate in places where there is a large number of workers with such skills. As a result, the model predicts that sectors whose establishments experience more idiosyncratic volatility will be more spatially concentrated.

To assess empirically the importance of labour market pooling as a source of agglomeration economies, we use plant-level data from the United Kingdom's Annual Respondent Database (ARD), which underlies the Annual Census of Production. The data is described in section 3. We begin by constructing a plant-level measure of idiosyncratic employment shocks by calculating the difference between the percentage change in the plant's employment and the percentage change in the sector's employment. We then average this (in absolute value) across time and across establishments in the sector to obtain a sector-level measure of how much idiosyncratic volatility individual establishments in each sector face. We then check whether, consistent with the theory, sectors whose establishments experience more idiosyncratic volatility are more spatially concentrated. We find that this is indeed the case, even after controlling for a range of other industry characteristics that include a novel measure of the importance of localized intermediate suppliers.

2. The theoretical advantages of labour pooling

In this section, we review Krugman's (1991) labour pooling model.[1] The purpose of this is twofold. First, to clarify the microeconomic foundations of labour pooling as a source of agglomeration economies. Second, to derive an empirically testable prediction about how the importance of labour pooling will vary across sectors.

Consider a sector with a discrete number of production establishments distinguished by subindex i and a continuum of workers. Establishments and workers are risk neutral. Following Ellison and Fudenberg (2003), we treat location and production in this model as a two stage game. In the first stage, establishments and workers simultaneously choose their location. In the second stage, each establishment receives a productivity shock ε_i.

[1] See also Ellison and Fudenberg (2003) and Duranton and Puga (2004).

The shocks are uncorrelated across establishments and identically distributed over $[-\varepsilon,\varepsilon]$ with mean zero and variance σ_s. Establishments observe these shocks and decide how much labour to hire from the local labour pool in the sector. If establishment i chooses an employment level l_i, it has profits given by:

$$\pi_i = [\beta + \varepsilon_i]l_i - \frac{1}{2}\gamma[l_i]^2 - wl_i \ . \tag{1}$$

Following Krugman (1991), assume that each establishment takes the local wage as given. Thus, after shocks are realized, each establishment hires labour until its marginal value product equals the wage. This yields establishment i's labour demand:

$$l_i = \frac{\beta - w + \varepsilon_i}{\gamma} \ . \tag{2}$$

Substituting equation (2) into (1), this simplifies into:

$$\pi_i = \frac{[\beta - w + \varepsilon_i]^2}{2\gamma} \ . \tag{3}$$

Note that establishment profits are a convex function of the idiosyncratic productivity shock, since the establishment adjusts its production level in response to the shock. Similarly, profits are convex in the wage.

In the first stage, workers choose their location in order to maximize their expected wage $\mathrm{E}(w)$ given the location of establishments and of other workers. Denote by L the size of the local labour pool in the sector and by N the number of local establishments in the sector. Labour market clearing, together with (2), implies

$$L = \sum_{i=1}^{N} l_i = \frac{\beta - w + \sum_{i=1}^{N} \varepsilon_i}{\gamma} \ , \tag{4}$$

We can then solve for the market clearing wage from equation (4):

$$w = \beta - \gamma \frac{L}{N} + \frac{1}{N} \sum_{i=1}^{N} \varepsilon_i \ . \tag{5}$$

Taking expectations and denoting the ratio of workers to establishments by $R \equiv L/N$, yields the expected wage:[2]

$$\mathrm{E}(w) = \beta - \gamma R \ . \tag{6}$$

[2] We assume that the support of the distribution of productivity shocks is not so large that the non-negative employment constraint for some establishment might be binding under some realization of shocks. In particular, we assume that the restriction $\frac{\gamma}{\varepsilon} \geqslant \frac{2(N-1)}{L}$ holds. This follows from $l_i > 0$, and equations (2) and (5) for a case where $\varepsilon_i = -\varepsilon$ and $\varepsilon_j = \varepsilon \ \forall j \neq i$.

Since $\mathrm{E}(w)$ only depends on the ratio R of workers to establishments and is decreasing in this ratio, any individual worker would prefer to be in the location with the lowest ratio of workers to establishments. Thus, in equilibrium workers spread across cities in proportion to the number of establishments (if this was not the case, workers in locations with a higher R could profitably deviate). Since $\mathrm{E}(w)$ is decreasing in L, such an allocation is also stable with respect perturbations in the distribution of workers across space.

Turning to the first-stage location decision of establishments, they make their choice in order to maximize expected profits, given the location of workers and of other establishments. Taking expectations of the profits in equation (3) yields:

$$\mathrm{E}(\pi_i) = \frac{[\beta - \mathrm{E}(w)]^2 + \mathrm{var}[\varepsilon_i - w]}{2\gamma} \,. \tag{7}$$

Substituting (6) and $\mathrm{var}[\varepsilon_i - w] = \mathrm{var}[\varepsilon_i] + \mathrm{var}(w) - 2\mathrm{cov}[\varepsilon_i,w]$ into (7), this simplifies into:

$$\mathrm{E}(\pi_i) = \frac{\gamma}{2}R^2 + \frac{\mathrm{var}[\varepsilon_i] + \mathrm{var}(w) - 2\mathrm{cov}[\varepsilon_i,w]}{2\gamma} \,. \tag{8}$$

The first term of the right-hand side is what establishment profits would be in the absence of shocks. It increases as the ratio R of workers to establishments increases because this lowers the expected wage. The second term captures the labour pooling effect. This shows that expected profits increase with the variance of the establishment-specific productivity shock, $\mathrm{var}[\varepsilon_i]$, and with the variance of the local wage, $\mathrm{var}(w)$, because of the convexity of profits discussed above. However, they decrease with the covariance of the establishment-specific productivity shock and the local wage, $\mathrm{cov}[\varepsilon_i,w]$. The reason is that if the local wage is higher when an establishment wishes to expand production in response to a positive shock and lower when the establishment wishes to contract production in response to a negative shock, profits become less convex in the shock and fall in expectation. This is the key intuition of the model, which highlights the microeconomic foundations of labour pooling as a source of agglomeration: *establishments prefer locations where their productivity shocks get ironed out rather than heavily reflected in local wages*.

To simplify equation (8) further, we can use equations (5) and (6) to calculate $\mathrm{var}(w) = \frac{\sigma_s}{N}$ and $\mathrm{cov}[\varepsilon_i, w] = \frac{\sigma_s}{N}$. Substituting these and $\mathrm{var}[\varepsilon_i] = \sigma_s$ into (8) yields:

$$\mathrm{E}(\pi) = \frac{\gamma}{2}R^2 + \left(1 - \frac{1}{N}\right)\frac{\sigma_s}{2\gamma}\,, \tag{9}$$

where we have dropped subindex i since expected profits are equal for all establishments in the same location and sector. An equilibrium is a distribution of workers and establishments across locations such that R is the same in all locations and an individual establishment cannot increase the expected profits of equation (9) by locating elsewhere. An establishment must consider two aspects in deciding whether such a deviation is profitable. First, starting from a situation where R is constant across locations, the establishment's relocation would decrease the ratio of workers to establishments in the destination location, making the labour market tighter in expectation, and increasing the expected wage, which would reduce the establishment's profits. This labour market tightness effect is captured by the first term on the right-hand side of (9). Second, if the destination has a larger number of establishments, the establishment's productivity shocks (that get translated into employment shocks) will not affect the local wage as much, allowing the establishment to adapt better to circumstances and obtain higher expected profits. This is the labour pooling effect discussed above, summarized now by the second term on the right-hand side of (9). If we consider multiple sectors differing only in terms of the variance of productivity shocks, σ_s, the labour market tightness effect favouring establishment dispersion is equally strong across all sectors but the labour market pooling effect is stronger the higher is σ_s in the sector. Thus, the balance of agglomeration and dispersion forces tips more easily in favour of agglomeration the higher is σ_s. In particular, if a location has fewer than $\frac{\sigma_s}{2\gamma^2R^2+\sigma_s}$ times as many establishments as the largest agglomeration in the sector, all remaining establishments find it individually profitable to relocate to the largest agglomeration.[3] Thus, *the benefits of labour pooling will be greater*

[3] Stated differently, an equilibrium in this model is an allocation of workers and establishments across locations such that each location is either empty or has at least $\sigma_s/(2\gamma^2R^2 + \sigma_s)$ as many establishments as the location with most establishments and the ratio R of workers to establishments is the same in all non-empty locations as in the aggregate economy. See Ellison and Fudenberg (2003) for details.

the larger the heterogeneity of establishment-specific shocks in the sector and consequently sectors with more heterogeneous shocks are more likely to be agglomerated. We will test this prediction empirically in section 4.

3. Data

To examine the role of pooling, we will regress a measure of spatial concentration for each sector on a measure of the potential for labour pooling in the sector and a number of sectoral characteristics that are also likely to affect spatial concentration. The measure of spatial concentration and the pooling variable described below are calculated on using exhaustive establishment level data from the Annual Respondent Database (ARD) which underlies the Annual Census of Production in the United Kingdom. We use data from 1994–2004. The data set is collected by the Office for National Statistics (ONS) and covers all UK establishments (see Griffith, 1999, and Duranton and Overman, 2005, for a detailed description of this data).[4] For every establishment, we know its postcode, four-digit industrial classification, and employment. We restrict our attention to production establishments in manufacturing industries using the Standard Industrial Classification 92 (SIC 15000 to 36639) for the whole country except Northern Ireland. For the purposes of this exercise we have plant data from the ARD for 1990–2004. We observe 557,595 plants at least once. On average, we observe each plant 4.16 times.

Since the labour pooling mechanism depends on firms ability to take more or less workers from the local labour pool without difficulty, we must work with geographical units that correspond as much as possible to local labour markets. Thus, our geographical unit of analysis are the UK Travel to Work Areas (TTWA), 1998 classification. Similar to the Labor Market Areas that the Bureau of Labor Statistics defines for the United States, these TTWA are defined on the basis of commuting patterns to capture local labour markets. Specifically, the boundaries are drawn such that of the resident economically active popu-

[4] We use the terms establishment and plant interchangeably. Our description of the data is based closely on Duranton and Overman (2005).

lation, at least 75% work in the area and of everyone working in the area, at least 75% live in the area. The classification is exhaustive, with 308 covering the whole of Great Britain. UK postcodes can be uniquely mapped to TTWA so we are able to locate establishments in the ARD according to the TTWA classification. The number of plants per TTWA is rather skewed. There are 15,154 on average, while the median number is 4,545. There are 14 TTWA with less than one hundred plants, although inclusion of the very large or the very small areas does not appear to affect our results so we include the whole sample in what follows. One slight complication involves the treatment of plants that move across TTWA or change sector. We treat these as a separate observation.[5]

Our controls for other industry characteristics come mainly from the ONS Input-Output tables, available annually from 1992 to 2004.[6]. We complement these where necessary with Eurostat's Detailed Enterprise Statistics for the United Kingdom and the ARD itself. We provide more details as we introduce these controls.

4. The importance of labour pooling for industry concentration

The theoretical model of section 2 suggests that sectors whose establishments experience more heterogeneous employment shocks have greater potential benefits from labour pooling and, to exploit these, will be more spatially concentrated. In this section we consider this prediction by regressing a measure of spatial concentration for each sector on a measure of the potential for labour pooling in the sector. Of course, other characteristics of industries may also affect the extent of concentration, and we will need to control for these. That is, we estimate:

$$C_s = \alpha + \rho P_s + \boldsymbol{\phi} \boldsymbol{X_s} + \epsilon_s \ , \tag{10}$$

[5]Moves across TTWA should not actually happen as plants identifiers are supposed to designate a unique physical entity. In reality, firms sometimes report under the same plant identifier when they have actually moved plants. This justifies our decision to treat these observations separately. The issue of changing SIC is more problematic as these classifications are based in the most significant activity undertaken at a given plant so may change over time.

[6]The UK Input-Output tables use a 77 industry classification. This is compatible with NACE Rev. 1 and corresponds, roughly to NACE 3-digit. We map this to the 237 industries in the UK SIC92 by assigning the same value to all 4 digit industries under any given IO heading.

where C_s is a measure of spatial concentration for sector s, P_s is a measure of the potential for labour pooling in the sector, $\boldsymbol{X}_i$ is a vector of sector characteristics, α, ρ and $\boldsymbol{\phi}$ are parameters to be estimated, and ϵ_i is an identically and independently distributed error term.

This approach to investigate various motives for spatial concentration has been used before (see Audretsch and Feldman, 1996, and in particular Rosenthal and Strange, 2001, to which our regressions are most directly related). The main novelty of our analysis is to explicitly look at the potential for labour pooling of different sectors by measuring the heterogeneity of individual establishment's employment shocks in each sector. We will also offer an important refinement for measuring the importance of sharing intermediate input suppliers.

Measuring each sector's potential for labour pooling

The argument that labour pooling, by allowing establishments to better adapt to idiosyncratic to shocks, can be an important determinant of agglomeration is well known. However, data restrictions mean that previous studies have had to get at this effect rather indirectly by focusing, for example, on the extent to which workers in an industry are likely to have industry-specific skills. Rosenthal and Strange (2001), for example, use three measures of labour pooling: net productivity (the value of shipments less the value of purchased inputs, all divided by the number of workers in the industry), the ratio of management workers to production workers, and the percentage of an industry's workers with Doctorates, Master's Degrees, and Bachelor's Degrees. Sectors with a larger share of managers or high-skilled workers may agglomerate partly because of labour pooling, but many other reasons could also be important. For instance, agglomerations of high-skilled workers may facilitate better matching between jobs and workers Helsley and Strange (1990). Also, large markets may allow high-skilled workers to specialize in a narrower set of tasks and become more productive (Baumgardner, 1988, Becker and Murphy, 1992, Duranton, 1998) or they may facilitate solving dual-career problems for high-skilled couples (Costa and Kahn, 2000), leading sectors relying more heavily on high-skilled workers to agglomerate.

We wish to isolate labour pooling, as motivated by the theoretical argument of section 2, from other labour market considerations. The crucial point, as discussed above, is that a labour pooling advantage only arises if, whenever a plant expands employment many other plants using similar workers are contracting, and vice versa. That is, what matters is the plants' idiosyncratic need to alter employment. To capture this effect, we exploit the fact that we have a panel of plants over a long time period to construct a direct measure of the idiosyncratic nature of any given plant's employment adjustments. To measure the idiosyncratic shock to a plant in any given year, we calculate difference between the percentage change in the plant's employment and the percentage change in the industry's employment (in absolute value). This will take a high value for plant's who either expand employment when the rest of the industry is contracting, or vice versa. Differencing the plant's change from the industry's change is important because there is no labour pooling advantage if whenever the plant expands employment many other plants using similar workers also expand. We then take the average of this variable across all years and across all firms in each sector. The resulting "pooling" measure captures how much idiosyncratic volatility individual establishments in each sector face.

Measuring each sector's spatial concentration

There are a variety of statistics that can be used to measure the extent of spatial concentration. We use the index proposed by Ellison and Glaeser (1997). This measures the amount of clustering in a sector over and beyond that which we would expect to find based on randomness alone. It has the advantage of being comparable across sectors, controlling for the overall concentration of employment, and also controlling for the "lumpiness" of employment. This lumpiness arises because industrial concentration means plants are of different sizes. This is a problem for measuring spatial concentration, because even random distributions of plants across spatial units can give rise to some places having more employment than others (if they happen, by chance, to get a particularly large plant). The Ellison-Glaeser index controls for the industrial concentration of the industry and thus corrects for this problem. Let s_a be the share of sector s employment that is in area a, x_a be the share of total employment that is in area a. Then the Ellison-Glaeser index of geographical concentration is defined as:

$$C_s \equiv \frac{G_s - (1 - \sum_a x_a^2) H_s}{(1 - \sum_a x_a^2)(1 - H_s)} , \tag{11}$$

where G_s is a raw localization index equal to

$$G_s \equiv \sum_a (s_a - x_a)^2 , \tag{12}$$

and

$$H_s \equiv \sum_i z_i^2 \tag{13}$$

is the Herfindahl index of the sector's plant size distribution, with z_i denoting plant i's share of sector s's employment. Ellison and Glaeser (1997) show that if plants are randomly distributed across locations with probabilities given by x_a, then the expected value of this measure is zero. A positive value of the index indicates a level of spatial concentration over and above what one would expect by chance.

Baseline results

Although we have panel data for the Ellison-Glaeser index and some of the explanatory variables, some preliminary regressions we estimated exploiting the panel dimension of the data did not perform well. Perhaps this is unsurprising as location patterns change only slowly while some of the industry characteristics (e.g., R&D expenditure per worker) can show a considerable amount of year on year variation. Furthermore, for the labour pooling measure it seems appropriate to take into account plant-level employment shocks relative to the sector for a number of years. The simple solution to deal with the noise in the time series is to average over time. We do this by splitting the data in half and regressing the average Ellison-Glaeser index for the six years from 1998 to 2003 on the average of the industry characteristics from 1992 to 1997. This specification has a rather nice economic interpretation whereby plants are able to observe industry characteristics before making their location decisions so we would actually expect some lag from characteristics to outcomes. It also helps partially address concerns about the endogeneity of some of the industry characteristics.[7]. We now briefly consider each of the industry characteristics before turning to our results.

As we have already discussed, industries with large number of plants which face idiosyncratic shocks should be more spatially concentrated. We identify these industries through the pooling variable described above. The remaining control variables for our first specification broadly follow Rosenthal and Strange (2001). We briefly motivate all of them, but refer the reader to Rosenthal and Strange (2001) for a more detailed discussion.

The availability of natural resources may differ across regions. If natural resources are very spatially concentrated, then we would expect industries that use to them intensively to be very spatially concentrated. Of course, if natural resources are very dispersed, then the opposite effect might hold and industries which use this resource intensively may be dispersed. As we do not have independent information on the distribution of resources, we capture the effect of natural resources on spatial concentration by looking at each industry's primary inputs (from agriculture, forestry, fishing, mining and quarrying) as a share of total inputs. As the discussion makes clear we do not have a strong prior on whether the impact will be positive or negative. Industries also differ with the intensity with which they use water and energy.[8] As the price of water and energy may differ across regions the intensity with which industries use these two inputs may affect their spatial distribution. We capture reliance on water using "collection, purification and distribution of water" (1087) as a share of total inputs, from the ONS Input-Output tables. Eurostat's Detailed Enterprise Statistics provide data on the value of energy products purchased at the SIC 4 digit level, which we normalize by total inputs to provide a proxy for reliance on

[7] Although the slowly changing nature of spatial concentration means that shocks are likely to persist so that lagged values of the industry characteristics may not be ideal instruments for current values

[8] The UK water industry is comprised of a number of privatized regional monopolies who have different pricing structures. Thus, we allow for the possibility that water usage may play a role in industrial concentration, although the existence of a national regulatory is likely to restrict the importance of water in practice.

energy. We expect the coefficients on these two variable to be positive and significant if price variations across regions are large enough to affect plant location and insignificant otherwise.

Turning to agglomeration forces, we start by following Rosenthal and Strange (2001) and using the purchase of goods and services as a share of inputs to capture the importance of vertical linkages. These are calculated using the input coefficients on manufacturing (IO8-84) and non-manufacturing industries (IO107-115, 118-123), respectively, from the ONS Input-Output tables. The basic idea is that industries who buy or sell a lot from other plants may have an incentive to cluster near those plants. If the degree to which an industry buys goods and services as inputs captures this effect, then we should expect the coefficient on these two variables to be positive. As emphasized in models of new economic geography, the level of transport costs for an industry will be crucial in determining whether agglomeration forces outweigh dispersion forces leading to the spatial clustering of the industry. We use transport services (IO93-97) as a share of inputs to capture the impact of transport costs on industry spatial concentration, again using data from the ONS Input-Output tables. As Rosenthal and Strange (2001) argue, this measure is not ideal as it is most likely endogenous. Unfortunately, for the UK, alternative data are not available in the time period that we consider. Finally, we use the share of R&D expenditure in value added to capture the possible role of technological externalities and knowledge spillovers in driving the spatial concentration of high-tech industries. These are calculated on the basis of Eurostat's Detailed Enterprise Statistics for the United Kingdom.[9]

Results from a regression of the Ellison-Glaeser index (averaged over the years 1998 to 2003) on these industry characteristics (averaged over the years 1992 to 1997) are given in column 1 of table 1.

The main result of interest is the relationship between each sector's potential for labour

[9]Preliminary data for all these variables were kindly provided by Roberto Picchizolu, a PhD student in the department of Geography and Environment at the London School of Economics. The final version of our data continues to use the energy and R&D variables provided by Picchizolu, but the remaining variables are based on the authors own calculations from the ONS Input-Output tables 1992–2004.

Table 1. Regression of EG index on industry characteristics

Dep. variable: Ellison-Glaeser index	(1)	(2)
Pooling	0.046	0.05
	0.028*	0.030*
Natural resources as share of inputs	-0.157	-0.173
	0.048***	0.057***
Water as share of inputs	2.181	1.596
	3.016	2.676
Energy as share of inputs	-0.084	-0.35
	0.356	0.347
Purchase of goods as share of inputs	-0.19	-0.25
	0.077**	0.086***
Purchase of services as share of inputs	-0.538	-0.388
	0.158***	0.155**
Share of R&D expenditure in value added	-1.802	-2.084
	1.258	1.211*
Transport costs as share of inputs	-0.441	-0.437
	0.142***	0.137***
Own industry inputs as share of inputs		0.096
		0.029***
IO weighted EG index		0.505
		0.250**
Constant	0.163	0.159
	0.047***	0.049***
R^2	0.09	0.14
Observations	236	236

Errors robust to heteroscedasticity. ***, **, * significant 1%, 5%, 10% respectively.

pooling and the spatial concentration in the sector. As predicted, the role of the labour pooling variable is positive and significant. Thus, industries where, on average, plants face more idiosyncratic shocks relative to the industry are more spatially concentrated.

Turning to other determinants of spatial clustering, a high natural resource requirement actually causes industries to be less spatially concentrated than they otherwise would be. This may well reflect the fact that agricultural inputs tend to dominate for most industries where natural resource inputs are important and, at least in the UK, agricultural activity is reasonably dispersed across the country. Water and energy use have no significant effect on spatial concentration. As suggested above, this may well be because price variations are not actually that large across UK regions. Ignoring, for one moment, the role of purchases of goods and services, we see that the share of R&D expenditure in value added does not have a significant effect. The final variable, transport costs has a negative and significant effect on spatial concentration. As expected, industries with high transport costs are more dispersed.

Perhaps the biggest surprise are the negative and significant coefficients on the purchase of goods and services as a share of inputs. As we discussed above, if these variables are actually capturing vertical linkages then we would expect them to have a positive significant effect on spatial concentration. How, then to explain the negative coefficients?[10] It may be that sharing intermediate suppliers is not an important motive for agglomeration, but other evidence suggests it is.[11] The answer, it turns out, is similar to the story that allows for a negative coefficient on natural resources. When an industry buys a lot from other industries, the effect on its concentration will depend, in turn, on whether those industries are spatially concentrated or dispersed. For instance, the meat processing industry is a large buyer of inputs from farms and from the plastic film industry. However, farms are very dispersed across the country and so is the plastic film industry, since it supplies many other sectors located in different places in addition to meat processing. Hence, the meat processing industry has no reason to concentrate spatially even if it makes large intermediate purchases: it can easily find its inputs everywhere. For a sector to cluster to share intermediate suppliers, it must be the case not only that the sector makes large purchases of intermediates but also that those intermediates are supplied by industries that are themselves very spatially concentrated. Thus, to better capture the importance of vertical linkages for a particular industry, s, we calculate the input share weighted sum of the Ellison-Glaeser index across all industries from which industry s purchases intermediates. That is, we calculate

$$V_s = \sum_{j \neq s} I_{s,j} C_j \tag{14}$$

[10]Rosenthal and Strange (2001) find no significant effect for these variables.

[11]Holmes (1999) looks at variations in intermediate input purchases within sector across locations and finds a strong connection between spatial concentration and intermediate purchases.

where V_s is our new measure of vertical linkages, $I_{s,j}$ is the share of sector j in sector s's intermediate inputs from other sectors and C_s is the Ellison-Glaeser index of spatial concentration for sector j. Notice that, for obvious reasons, we exclude industry s's Ellison-Glaeser index (or left-hand side variable) from this calculation. However, we would expect, ceteris-paribus, industries that buy a large share of intermediate inputs from their own industry to be more spatially concentrated. To capture this, we can include $I_{s,s}$, the share of intermediates from own industry, in the regression in addition to the vertical linkages variable.[12].

Column (2) in table 1 shows what happens when we include these two new variables. We see that both the own industry inputs as a share of inputs and the Input-Output weighted Ellison-Glaeser index have a positive and significant impact on spatial concentration. Industries that buy a lot of intermediates from other plants in the same industry, or who buy a lot of intermediates from other industries that are spatially concentrated are, in turn, more spatially concentrated. We see that the coefficients on goods purchased and services purchased remain negative and significant. That is, purchasing large amounts of inputs per-se has a negative impact on spatial concentration. Finally, note that the coefficient on our main variable of interest, labour market pooling, remains positive and significant

5. Conclusions

Since Alfred Marshall talked about labour pooling as a source of agglomeration, it keeps being mentioned in every discussion of the topic. Yet, while existing empirical studies tend to find that labour market issues play a key role in leading industries to cluster, we have so far not had a direct test of whether ironing out plant-level shocks by drawing workers from a large local pool is, at least in part, what is behind this. In this paper, we have developed a novel measure that captures precisely this aspect: We calculate

[12] In the final version of the paper, we hope to be able to do something similar for services. However, there is a data access issue and a data availability issue (specifically, it appears that the service sector data is only available from 1997 onwards). We are working on both.

the fluctuations in employment of individual establishments relative to their sector and average these across the sector and over time. Our results show that sectors whose establishments experience more idiosyncratic volatility are more spatially concentrated, even after controlling for a range of other industry characteristics that include a novel measure of the importance of localized intermediate suppliers.

References

Audretsch, David B. and Maryann P. Feldman. 1996. R&D spillovers and the geography of innovation and production. *American Economic Review* 86(3):630–640.

Baumgardner, James R. 1988. The division of labor, local markets, and worker organization. *Journal of Political Economy* 96(3):509–527.

Becker, Gary S. and Kevin M. Murphy. 1992. The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics* 107(4):1137–1160.

Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux. 2007. Estimating agglomeration effects: Does playing with different instruments give a consistent tune? Processed, University of Toronto.

Costa, Dora L. and Matthew E. Kahn. 2000. Power couples: Changes in the locational choice of the college educated, 1940-1990. *Quarterly Journal of Economics* 115(4):1287–1315.

Duranton, Gilles. 1998. Labor specialization, transport costs, and city size. *Journal of Regional Science* 38(4):553–573.

Duranton, Gilles and Henry G. Overman. 2005. Testing for localization using micro-geographic data. *Review of Economic Studies* 72(4):1077–1106.

Duranton, Gilles and Diego Puga. 2004. Micro-foundations of urban agglomeration economies. In Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: North-Holland, 2063–2117.

Ellison, Glenn and Drew Fudenberg. 2003. Knife-edge or plateau: When do market models tip? *Quarterly Journal of Economics* 118(4):1249–1278.

Ellison, Glenn and Edward L. Glaeser. 1997. Geographic concentration in US manufacturing industries: A dartboard approach. *Journal of Political Economy* 105(5):889–927.

Griffith, Rachel. 1999. Using the ARD establishment level data: an application to estimating production functions. *Economic Journal* 109(456):F416–F442.

Helsley, Robert W. and William C. Strange. 1990. Matching and agglomeration economies in a system of cities. *Regional Science and Urban Economics* 20(2):189–212.

Holmes, Thomas J. 1999. Localization of industry and vertical disintegration. *Review of Economics and Statistics* 81(2):314–325.

Krugman, Paul R. 1991. *Geography and Trade*. Cambridge, MA: MIT Press.

Marshall, Alfred. 1890. *Principles of Economics*. London: Macmillan.

Rosenthal, Stuart S. and William Strange. 2004. Evidence on the nature and sources of agglomeration economies. In Vernon Henderson and Jacques-François Thisse (eds.) *Handbook of Regional and Urban Economics*, volume 4. Amsterdam: North-Holland, 2119–2171.

Rosenthal, Stuart S. and William C. Strange. 2001. The determinants of agglomeration. *Journal of Urban Economics* 50(2):191–229.