

Heterogeneous Agglomeration*

Giulia Faggio
Spatial Economics Research Centre,
London School of Economics,
Houghton Street, WC2A 2AE, London, UK.
Tel.: +44 (0)20 7852 3555
G.Faggio1@lse.ac.uk

Olmo Silva
Department of Geography and Environment,
London School of Economics,
Houghton Street, WC2A 2AE, London, UK.
Tel.: +44 (0)20 7955 6036
O.Silva@lse.ac.uk

William C. Strange
Rotman School of Management
University of Toronto
105 St. George St.
Toronto, ON M5S 3E6, Canada
wstrange@rotman.utoronto.ca
Tel.: +1 (416)978-1949

December 10, 2013

PRELIMINARY: COMMENTS ARE WELCOME

*We thank the Social Sciences and Humanities Research Council of Canada and STICERD-LSE for financial support. We also thank Simona Iammarino, William Kerr and Gianluca Tarasconi for sharing data. Finally, we thank Steve Gibbons, Stephan Heblich, Antonio Miscio, Henry Overman, Jacques Thisse and seminar participants at LSE, Syracuse University, and 2013 NARSC Annual Meeting for helpful comments. We are responsible for any errors or omissions.

Heterogeneous Agglomeration*

Abstract*

Many prior treatments of agglomeration either explicitly or implicitly suppose that all industries agglomerate for the same reasons, with traditional Marshallian (1890) factors affecting all industries similarly. An important instance of this approach is the extrapolation of the agglomeration experience of one key sector or cluster to the larger economy. Another is the pooling of data to look at common tendencies in agglomeration. This paper uses UK establishment level data on coagglomeration to document heterogeneity across industries in the microfoundations of agglomeration economies. The pattern of heterogeneity that we document is consistent with both traditional Marshallian theories and with alternative approaches that emphasize the adaptive and organizational aspects of agglomeration.

*Disclaimer: This work was based on data from the Business Structure Database and the Quarterly UK Labour Force Survey, produced by the Office for National Statistics (ONS) and supplied by the Secure Data Service at the UK Data Archive. The data are Crown Copyright and reproduced with the permission of the controller of HMSO and Queen's Printer for Scotland. The use of the data in this work does not imply the endorsement of ONS or the Secure Data Service at the UK Data Archive in relation to the interpretation or analysis of the data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

I. Introduction

This paper considers heterogeneity across industries in the microfoundations of agglomeration economies. Marshall (1890) notes the existence of three sources of agglomeration economies: labor market pooling, input sharing, and knowledge spillovers. Many subsequent treatments of agglomeration either explicitly or implicitly suppose that all industries agglomerate for the same reasons, with the three traditional Marshallian factors affecting all industries similarly. An important instance of this approach is the practice of extrapolating the agglomeration experience of one key sector or cluster to the larger economy. Another is the pooling of data to look at common tendencies in agglomeration.

The paper shows that these approaches miss important and systematic differences in how agglomeration economies impact different industries. As we elaborate below, we document clear heterogeneity in the role played by the various Marshallian forces in the agglomeration of different industries. For instance, in some industries knowledge spillovers are key. Other industries are instead driven by labor market pooling or input sharing. In short, different industries agglomerate for different reasons. The pattern of heterogeneity, in addition to being important in its own right, also allows us to develop new ways to test some fundamental theories that emphasize the adaptive and organizational aspects of agglomeration.

The paper builds on the theoretical and empirical literatures on agglomeration. Both of these lines of research begin, of course, with Marshall.¹ Marshall offers insightful but informal analysis of specific industries, such as textiles and cutlery, describing how external increasing returns are generated. The theoretical and empirical agglomeration literatures also build on Jacobs (1969), who stresses the unplanned nature of the creation of new work in cities, and Vernon (1960), who discusses how cities help manage the instability involved in certain production processes. Like Marshall, Jacobs focuses on specific industries and is informal in her analysis. Vernon, while also tending to extensive discussions of specific industries, moves beyond his verbal and case-based predecessors by bringing data from a range of industries to bear on the issue of understanding the sources of agglomeration economies. Chinitz's (1961) analysis of the positive role of small firms in the generation of agglomeration economies is another example of this sort of research. So is Porter's (1990) analysis of industry clusters. The data in his research are largely presented as descriptions of highly illuminating cases, rather than being used to test the predictions generated by theory.

The econometric literature on agglomeration has many such tests.² A number of papers have separately considered Marshall's agglomeration forces. See, among others, Fallick et al. (2006) and Almazan et al. (2007) on job hopping, Holmes (1999) on input sharing, and Jaffe et al. (1993), Arzaghi and Henderson (2008) and Lin (2012) on patents, networking and learning, and the creation of new work. This body of work

¹ Marshall is usually the oldest reference in an agglomeration paper's bibliography, but some of the same issues were discussed previously in von Thünen (1826, 1966), as noted by Fujita (2011).

² See Rosenthal and Strange (2004).

presents persuasive evidence that the three Marshallian forces are present. A related body of work runs what might be called “horse races” looking at the relative importance of these forces. For example, Audretsch and Feldman (1996) and Rosenthal and Strange (2001) regress levels of agglomeration on proxies for the presence of labor market pooling, input sharing, and knowledge spillovers. More recently, Ellison et al. (2010) regress coagglomeration of industry pairs – the tendency of industries to locate together – on proxies for Marshall’s forces. Another recent approach is Jofre-Monseny et al. (2011), who estimate count models of new firms as functions of proxies for Marshallian forces.

While the agglomeration literature has much to say about how agglomeration economies are generated, it has less to say about heterogeneity in microfoundations. The case studies are, of course, particular, and they respect industry heterogeneity. They do not, however, offer much guidance in how the specific cases they examine can be extrapolated to a more general phenomenon. The same can be said about econometric analyses of one industry. Conversely, the aggregation in the horse-race papers removes heterogeneity from consideration by treating industries as observations. As for the theory literature, little is said about heterogeneity.

The empirical literature also has less to say about microfoundations beyond assessing Marshall’s three forces. Duranton and Puga (2001) is an exception. Although it is largely a theoretical exercise, it also presents empirical evidence on location decisions over an industry’s life cycle that is consistent with the model of cities as nurseries that tend to young industries. Strange et al. (2006) is another exception, showing a systematic tendency for industries facing more uncertainty in Marshallian dimensions to agglomerate. In both of these papers, agglomeration is fundamentally adaptive. Glaeser and Kerr (2009) and Rosenthal and Strange (2010) consider the idea that agglomeration economies are stronger when there are many small firms. In these papers, agglomeration is organizational. There are also papers that ask whether Jacobs or Marshall is right by looking at relative performance in growth or productivity of diverse and specialized cities. See, for instance, Glaeser et al. (1992) and Henderson et al. (1995).

Our paper contributes to this literature by analyzing heterogeneity in agglomeration and its implications for microfoundations using a unified framework that captures a range of agglomeration theories. To do this, we make use of confidential establishment-level data from the UK’s Business Structure Database (BSD) covering the years 1997-2008. Unlike the vast majority of the papers in the agglomeration literature, the paper focuses on the implications of coagglomeration for the nature of agglomeration economies. We initially calculate the relationship between measures of industry links and coagglomeration across all manufacturing industries, as in Ellison et al. (2010). These models use industry pairs as observations. The pooled regressions should be interpreted as benchmarks that allow for no heterogeneity of effects between industries.

We then consider heterogeneity in ways that allow us to assess alternative approaches to agglomeration. We focus primarily on the Jacobs and Vernon notions of adaptation, and Chinitz's organizational approach. We also examine more recent related ideas about human capital (e.g., Rauch , 1993, Glaeser and Saiz, 2004, and Berry and Glaeser, 2005). Using coagglomeration to look at these adaptive and organizational aspects of agglomeration is unique in the literature. To investigate unplanned interactions as in Jacobs, we compare industry pairs that coagglomerate extensively with those that do not. To further consider adaptation, we compare industry pairs with substantial entry with those without much entry. In addition, we compare industries by age. We also contrast high-technology industry pairs with low-technology pairs, and industries with highly educated workforces against those with less educated workers. Finally, to analyze organizational issues, we compare industries by the sizes of incumbents and entrants. In all of these approaches, we examine the interaction between Marshallian forces and other elements of agglomeration rather than looking at Marshall as a rival to Jacobs and others, as is common in the literature. In this sense, the paper is an attempt to create a détente between Marshall and Jacobs.

The empirical analysis leaves no doubt that agglomeration works differently for different industries. The key empirical results are as follows. First, in a great variety of coagglomeration models, we show the robust predictive power of Marshall's agglomeration forces. This confirms prior work and supports our pursuit of looking at heterogeneity and at interactions between Marshallian and non-Marshallian approaches. Second, a quantile regression that differentiates pairs by their tendency to coagglomerate provides results that are consistent with Jacobs' analysis of unplanned interactions. Third, there is robust evidence of an adaptive element to agglomeration, consistent with Vernon. This manifests itself more strongly in labor pooling and in knowledge spillovers than in input sharing. Fourth, differentiation by the sector's technology orientation and workforce education shows that agglomeration is not just a high-technology phenomenon. However, high-technology sectors show stronger evidence of knowledge spillovers, while lower-technology industries instead show stronger evidence of input sharing and labor pooling. Last, agglomeration effects tend to be stronger when firms are smaller, consistent with Chinitz. The paper's approach of considering these non-Marshallian effects using heterogeneity in coagglomeration patterns is new to the literature.

Taken as a whole, our results argue for caution in extrapolation from individual cases of agglomeration. This is important because extrapolation from cases is a central part of the justification for cluster policy.³ Unfortunately, as satisfying as it is to draw conclusions from interesting and highly salient examples of agglomeration such as the Silicon Valley and computers or Detroit and cars, the findings show clearly that different industries respond differently to agglomerative forces. Similarly, the results suggest that one should interpret horse-race models on the relative strength of agglomeration effects (i.e., Audretsch and

³ See, for instance, Porter (1990) and the critique in Duranton (2011).

Feldman, 1996, Rosenthal and Strange, 2001, and Ellison et al., 2010) with care since these specifications do not allow for heterogeneous effects across industries.

To further consider the pattern of heterogeneity, we estimate single-industry models for several salient industries, including the cutlery, textiles, computers and motor vehicle industries. In these models, we identify microfoundations from the coagglomeration of the industry in question with all other manufacturing industries with which it is potentially linked. This analysis largely confirms the results obtained using the organizational and adaptation breakdown discussed above. However, we find that there is some heterogeneity even within industries that look similar in these features. This is further evidence of heterogeneity, and it is even more reason to extrapolate from case-based evidence only with the greatest care.

The remainder of the paper is organized as follows. Section II discusses our data and our measures of coagglomeration and of the three Marshallian forces. Section III makes the case for considering non-Marshallian forces and discusses how our data can capture them. Section IV presents the Marshallian analysis. Section V considers Jacobs' idea of unplanned interactions, while Section VI presents the key analysis where industries are heterogeneous along other non-Marshallian dimensions. Section VII concludes by presenting single-industry models.

II. Coagglomeration and agglomeration

A. Measuring coagglomeration

Our analysis of microfoundations is based on the tendency of industries to co-locate across metropolitan areas. We use the Ellison and Glaeser (1997) measure of coagglomeration. Let N_m denote total employment in metropolitan area m , and n_{im} denote the employment in industry i in metropolitan area m . Let $s_{im} = n_{im} / N_m$ denote the share of a given industry i 's employment in metropolitan area m , and let x_m denote the metropolitan area's share of national employment. For two industries, the Ellison-Glaeser measure of coagglomeration between industries i and j can be written as (Ellison et al., 2010):

$$\gamma_{ij}^c = \frac{\sum_{m=1}^M (s_{mi} - x_m)(s_{mj} - x_m)}{1 - \sum_{m=1}^M (x_m)^2}. \quad (1)$$

This measure is related to the covariance of industries across metropolitan areas.

The key property of this measure for our purposes is that it is related to the strength of agglomeration economies. Ellison et al. (2010, Mathematical Appendix) prove this in the context of a specific model of agglomeration. In their model, industries are partitioned into groups that must co-locate in order to have positive profit. With sequential location choices in this all-or-nothing agglomeration model, industries that

benefit from coagglomeration will coagglomerate. One can imagine weaker versions of the same sort of property from other models of agglomeration. When it is more profitable for firms to coagglomerate, they will tend to do so. It is worth noting, however, that there is a fundamental coordination problem in the determination of city composition (Helsley and Strange, 2013), and it is possible that coagglomeration fails to occur even when it would be mutually beneficial or that coagglomeration does occur when it is not. Even in this case, however, it is possible that the selection among multiple equilibria will be skewed at least somewhat in favor of coagglomeration occurring when it is beneficial.

B. Coagglomeration in UK manufacturing

In order to construct measures of pair-wise coagglomeration of UK manufacturing industries we use data from Business Structure Database (BSD). The data cover the period 1997 to 2008. The basic dataset is an annual snapshot (taken in April at the closing of the fiscal year) of the Inter-Departmental Business Register (IDBR), which consists of constantly-updated administrative data collected for taxation purposes.

In this paper, we use BSD data at the local unit (i.e., plant or establishment) level including both single- and multi-plant enterprises. For each local unit, information is available on employment, industrial activity, year of birth (start-up date) and death (termination date), and postcodes. We use this detail to assign each local unit active in England, Wales and Scotland to a Travel-to-Work Area (TTWA, see below). We neglect Northern Ireland because of poor data coverage. The raw data include approximately three million local units every year. In order to prepare the data for analysis, we carry out a series of checks and drop a number of units. Details are provided in the Data Appendix. After applying these restrictions, our dataset still comprises of more than two million plants annually over 12 years.

To quantify coagglomeration we focus on three-digit industries of the UK Standard Industry Classification (SIC) 1992 and restrict our attention to manufacturing (SIC151-SIC372)⁴. After excluding and recombining a number of sectors which present a limited or erratic evolution in the number of plants and/or their employment, we are left with a final sample consisting of 94 manufacturing 3-digit sectors. This gives a total of 4,371 unique pairs a year for twelve years (1997-2008) for a total of 52,452 observations.⁵

⁴ The UK SIC is a system for classifying industries by a four-digit code. Before the introduction of the six-digit North American Industry Classification System (NAICS code) in 1997, the US used its own four-digit SIC code. Ellison et al. (2010) construct industry pairs using three-digit industries classified according to the US SIC 1987.

⁵ We exclude Tobacco (SIC160) and combine the following industries: Leather (SIC181) and Fur clothes (SIC183) with Other wearing apparel (SIC182); Vegetal and animal oils and fats (SIC154) with Other food products (SIC158); Reproduction of recorded media (SIC223) with Printing (SIC222); Coke oven products (SIC231) and Processing of nuclear fuel (SIC233) with Refined petroleum products (SIC232); Man-made fabrics (SIC246) with Other chemical products (SIC247); Articles of concrete, plaster and cement (SIC266) with Manufacture of cement, lime and plaster (SIC265).

The level of geographical aggregation that we use is the Travel-to-Work Area (TTWA). TTWAs are geographical entities defined so that at least 75% of the resident population works in the area and at the same time 75% of the people working in the area reside there. TTWAs were devised by UK government agencies to delineate areas that can be considered as self-contained labor markets and economically relevant aggregates. As from 2007, there were 243 TTWAs within the United Kingdom. In our analysis, we focus on England, Wales and Scotland, and we select urban TTWAs only. In splitting urban and rural TTWAs, we follow Gibbons et al. (2010) and re-aggregate some areas so that our final partition splits Great Britain into 158 local economic areas, of which 84 are single urban TTWAs (with population in excess of 100,000 residents). More details are provided in the Data Appendix. In some extensions to our analysis, we also consider rural TTWAs.

To measure coagglomeration, we follow Ellison et al. (2010) and compute their metric based on the total employment shares of the selected 94 three-digit industries contained in the 84 urban TTWAs. We label this measure γ^C . Descriptive statistics are presented in the first row of Table 1. The mean and the median values of γ^C are centered at zero. It has a standard deviation of 0.005, with a minimum of -0.028 and a maximum of 0.107. Relative to Ellison et al. (2010) our coagglomeration measure displays less dispersion, although it is similarly skewed towards displaying co-location – i.e. towards positive values of γ^C . For comparison, we also calculate the Ellison and Glaeser (1997) measure of spatial concentration for our three-digit sectors. This has an average value of 0.027 and a median 0.009, compared with mean of 0.051 and a median 0.026 for the four-digit US classification reported in Ellison and Glaeser (1997). This suggests that – similarly to the US – ‘own coagglomeration’ (more familiarly, localization) is more important than coagglomeration. In general, manufacturing is less spatially concentrated in the UK than in the US.⁶

More information is provided in Table A1 in the Appendix, where we list the fifteen most coagglomerated sectors alongside the three top TTWAs where most of the coagglomeration takes place. Textiles, products of clay and ceramic, and basic metals are the most recurrent sector groupings. Some clear geographic patterns emerge. Textile-related industries tend to co-locate in Bradford, Leicester, Manchester and Nottingham. Basic-metal activities cluster around Birmingham and Sheffield. Clay- and ceramic-product manufacturers group around Stoke-on-Trent. The publishing and printing sectors tend instead to locate in London. The two most negative values of γ^C are found for “Knitted and crocheted articles (SIC177)” and “Publishing (SIC221)” with $\gamma^C = -0.027$ and for “Publishing (SIC221)” and “Ceramic goods other than construction purposes (SIC262)” with $\gamma^C = -0.025$.

⁶ UK figures calculated using a four-digit industrial classification do not affect this conclusion. For example, the Ellison-Glaeser (1997) metric used in Overman and Puga (2010) has a mean of 0.030 and a median of 0.010.

In extensions to our core analysis, we use variants of the γ^C measure. In particular, we calculate: (i) a measure of coagglomeration constructed using the number of plants rather than employment; (ii) a version of γ^C that excludes London; (iii) a measure that includes single-plant companies only; (iv) a version that includes both urban and rural areas; and (v) a measure that excludes publishing (SIC221) and printing and reproduction of media (SIC222). Descriptive statistics of these alternative measures are very similar to those presented in Table 1 for γ^C . Furthermore, their correlation with our main measure is always high, between 0.76 (when only including single-plants firms) and 0.99 (when considering both urban and rural areas).

C. Marshallian agglomeration forces

Marshall attributed the spatial concentration of industry to three forces: labor market pooling, input sharing, and knowledge spillovers. To the extent that these forces operate across industry boundaries, as they surely must to some degree, then businesses in a given industry will benefit from co-locating with businesses in other industries. The magnitude of this benefit will depend on the strength of the Marshallian links between industries. Some industries can readily share workers, but others cannot. Likewise, some industries can interact on factor markets or can share knowledge, but others cannot. In this section, we discuss the proxies that we use in order to measure the flow of goods, people and ideas across industrial pairs and that we relate to our measure γ^C in order to understand the microfoundation of coagglomeration. One empirical threat to our investigation is that firms will also tend to locate in the same metropolitan area if they use the same set of resources and infrastructures that are available locally. This might give rise to patterns of co-location that are not related to the Marshallian forces. In order to mitigate this problem, we use information on industry's use of primary resources and other non-manufacturing inputs, and control for industry-pair similarity in these respects. We discuss the details of this approach after presenting our proxies for the Marshallian agglomerative forces.

1. Labor market pooling

Firm clusters are often driven by the desire to take advantage of groups of workers with similar skills and expertise. In order to assess the importance of labor pooling as a microfoundation of coagglomeration, we use the UK Labour Force Survey (LFS) data between 1995 and 1999 and investigate whether sectors that use the same types of workers co-agglomerate with one another. The LFS is a quarterly representative survey of households living in the UK and is conducted by the ONS to track labor market dynamics (more details are provided in the Data Appendix). The data contain information on both workers' industry of occupation – which we use at the three-digit level (3-digit SIC) – and their Standard Occupation Classification (SOC 1990). The UK SOC is a standard job classification that categorizes occupations on the basis of their skill level and skill content at a detailed level.

We use the 331 occupation groups defined by the three-digit SOC classification in conjunction with our 94 manufacturing industries defined at the three-digit SIC level to calculate $Share_{io}$ and $Share_{jo}$. These measure the shares of employees of occupation o in the total employment in industry i and j , respectively. Using this information, we then measure the similarity of employment in industries i and j by computing the correlation between $Share_{io}$ and $Share_{jo}$. Descriptive statistics are presented in Table 1. The mean value is 0.237 with a standard deviation of 0.188. The maximum value of 0.968 is found for the “Manufacture of electronic valves, tubes & electronic components” and “Manufacture of TV/radio transmitters & telephones/telegraphs” pair. The minimum (at -0.022) is displayed by the combination of “Manufacture of veneer sheets, plywood, laminboard & other panels/board” with “Manufacture of Transport Equipment Not Elsewhere Classified”.

2. Input sharing

An alternative reason why firms locate near one another is to reduce the costs of obtaining intermediate inputs and shipping goods to downstream customers.⁷ To assess the importance of this factor, we use the ONS Input-Output Analytical Tables (I-O henceforth) for 1995 to 1999 to measure the extent to which industries buy and sell intermediate inputs from one another. In particular, we calculate the shares of inputs that each industry within a pair buys from each other as fractions of their total intermediate inputs, and the shares of outputs that they sell to each other as fractions of their total output (excluding sales directly to consumers). The sector classification used in the ONS I-O Tables is more aggregated than the three-digit SIC classification used so far and only includes 77 manufacturing industries. In order to assign input-output shares to SIC three-digit sectors belonging to the same I-O sector code we use an apportioning procedure based on their employment share within the group (averaged between 1995 and 1999).

We find that the largest input-sharing link is between “Manufacture of pesticides & agro-chemical products” buying from “Manufacture of basic chemicals” (0.547 of its total input), and between “Manufacture of jewelry & related articles” buying from “Manufacture of basic precious & non-ferrous metals” (0.468 of the total input) when we focus on industries that do not belong to the same two-digit industrial grouping. In terms of output, the two most closely linked sectors are “Textile weaving” selling 0.546 of its output to “Manufacturing of apparel & accessories”.

Using these data, we construct three different proxies for the extent of input-output links between industrial pairs. To begin with, we consider the maximum between the share of inputs that sector i is buying from sector j , and vice versa. This is a proxy for upstream linkages. Next, we recover the maximum between

⁷We refer to this as “input sharing” in line with prior usage even though there are both upstream and downstream elements to our measure, as in Krugman (1991).

the share of output that sectors i is selling to sector j , and vice versa. This measure captures downstream linkages. Finally, we consider the maximum of either of these two proxies as a synthetic measure of the linkages between sector pairs. This approach follows Ellison et al. (2010). Descriptive statistics are presented in Table 1. The mean value for all three proxies is small and close to zero, suggesting that most industries do not share intermediate goods to an important degree. Indeed, 30% of the sector pairs do not share any input or output, while 75% of the pairs share less than 0.005.

3. Knowledge spillovers

A third reason why firms co-locate is to speed the flow of ideas across industrial sectors. Indeed, since the original writings of Marshall (1890), agglomerations foster learning and knowledge sharing since clustering makes “the mysteries of the trade become no mysteries; but are as it were in the air, and [workers] learn many of them unconsciously.” Despite the appeal and simplicity of this idea, measuring flows of knowledge is highly challenging.

In order to construct a proxy for knowledge spillovers, we track patent citation flows using information on UK inventors contained in the European Patent Office (EPO) data. In particular, we make use of the EPO-CESPRI data provided by Bocconi University which provides clean and consistent information on patents applied for at the EPO between 1977 and 2009. Approximately 144,000 patents were filed by UK inventors over this period generating more than 77,000 citations of patents associated to UK based inventors. More information on this data is provided in our Data Appendix.

Using this data, we measure the extent to which patents associated with industry i cite patents associated with industry j and vice-versa. The main difficulty lies with creating a mapping between sectors and patents – which are categorized using technological classes rather than a standard industrial classification. Following the literature, we adopt two approaches and use: (1) a probabilistic mapping based on the Industry of Manufacture (IOM); and (2) an alternative probabilistic mapping based on the Sector of Use (SOU). This correspondence was developed by Silverman (2002) and researchers at Statistics Canada by studying approximately 150,000 patents filed at the Canadian Patent Office between 1990 and 1993.⁸

Once this mapping has been constructed, we are able to investigate the number of citations that a patent in sector i is receiving from patents in sector j , and the number of patents in sector j that a patent in sector i is citing. These measures are analogous to the input sharing proxies described above. Our two final indicators consider the maximum patent-citation flow between sector i and sector j – normalized by total citations in that industry – using either the IOM or the SOU probabilistic mapping. Descriptive statistics are

⁸ More information is available from Silverman’s website:
http://www-2.rotman.utoronto.ca/~silverman/ipcsic/documentation_IPC-SIC_concordance.htm

presented in Table 1. The average knowledge spillover shares are between 0.012 (SOU) and 0.016 (IOM). Both distributions are highly skewed with median values in the order of 0.003/0.004 and 75% of the industries having citation flows below 0.011/0.013. “Manufacture of pesticides & agro-chemical products” and “Manufacture of pharmaceuticals, medicinal chemicals & botanical products”, “Manufacture of motor vehicles” and “Manufacture of railway/tramway locomotives/rolling stock” and “Manufacture of plastic products” and “Manufacture of wooden containers” are among the most connected industries using these two metrics. In most of our analysis we use the proxy based on industry of manufacture (IOM), since the probabilistic mapping based on sector of use (SOU) is partly related to purchases of intermediate goods containing the relevant patents, and thus incorporates some elements of input-output sharing. However, we present a number of robustness checks using this second metric.

4. Natural resources

A well-recognized determinant of firm clustering is the spatial distribution of natural resources and infrastructure. If transport costs are not trivial, industries that rely on inputs that are unevenly distributed across space or that need to source/send goods from/to far afield will show a tendency to locate close to where these inputs can be found and near transport infrastructures. In order to address this problem, Ellison et al. (2010) use the spatial distribution of natural resources, transport costs and labor inputs to predict coagglomeration that would only stem from differences in costs. This approach cannot be replicated in the UK for two reasons. First, the geographical scale of the country makes the spatial distribution of resources and “natural infrastructure” – e.g. access to the sea – much more homogeneous. Second, differences in the cost of resources – such as gas, oil, water and electricity – are negligible due to regulatory constraints.

In order to deal with this issue, we capture the effect of natural resources on geographic concentration by looking at industries’ use of primary inputs as a share of total inputs. Using again the ONS 1995-1999 I-O Tables, we start by building a measure of the share of inputs that an industry is purchasing from the seven I-O primary industries, namely agriculture, forestry and fishing (industries 1, 2 and 3), and mining and quarrying (industries 4, 5, 6 and 7). We refer to this group as “Natural Resources.” We also control for usage of water and energy by separately considering the share of inputs bought from water-related service companies (IO industry 87 “Collection, purification and distribution of water”) and from energy-related industries (IO codes 85 “Electricity production & distribution” and 86 “Gas distribution”). Further, in order to control for the importance of transport costs in an industry’s production, we consider the share of inputs bought from transport-related sectors (using IO codes 93 to 96 and corresponding to railways, air, water and other land transport). Finally, following Overman and Puga (2010)’s analysis of labor market

pooling, we create a proxy for access to business services by considering the share of inputs bought from industries in this compartment.⁹

Using these shares we construct proxies for the dissimilarity of industry pairs by measuring (one half of) the absolute value of the difference in the shares of these various inputs used by the pair. Descriptive statistics are presented in the bottom panel of Table 1. The largest mean value is found for the dissimilarity in the use of natural resources (at 0.041), while the smallest relates to use of water (0.001). The other three measures have similar mean values at around 0.015.

III. Beyond Marshall: Adaptive and organizational agglomeration economies

A. Jacobs: Unplanned interactions

While Marshall's (1890) three forces have dominated the agglomeration literature, there have been other important approaches to agglomeration. The approach that has received the most attention is Jacobs (1969), who focuses on the unplanned nature of the creation of "new work". It is common to treat Jacobs as proposing an alternative to Marshall, as in the Glaeser et al. (1992) and Henderson et al. (1995) papers on urban growth. There is a natural sense in which this is true. Marshall sees his increasing returns forces as promoting the spatial concentration of industry. Jacobs, in contrast, focuses primarily on knowledge spillovers, and sees the creation of new work as a force that is enhanced by local diversity and that ultimately produces a diverse local economy.

There is another sense, however, in which Jacobs and Marshall ought not be presented as polar opposites. The Jacobs analysis of knowledge is certainly in the spirit of Marshall, and she also mentions labor- and input-market aspects of the creation of new work. For instance, she notes that Henry Ford's first assembly line made use of locally outsourced parts. In this spirit, this paper examines complementarities between Marshall and Jacobs.

To do so, we depart from standard practice in the agglomeration literature. This approach, which can be found in numerous papers, is to regress productivity, growth, wage, or rent on measures of local specialization or diversity. While the specialization measure is fairly tightly tied to Marshall, the diversity variable is more loosely linked to Jacobs. The idea is that a diverse city offers opportunities for interactions between different industries, leading to more frequent instances of new work creation. Diversity is typically measured using a Herfindahl index of industrial concentration. While this is a logical approach to formalizing Jacobs verbal analysis, it is certainly not the only approach that might be taken.

⁹ This group stretches over IO codes 107-115 and 118-123 and includes, among others, computer services, R&D activities, legal consulting, accounting services, market research and management consulting, and advertising. Results do not depend on the inclusion of this variable.

This paper offers a new approach. The basic idea is that some agglomeration is deliberate and other agglomeration is unplanned. Both kinds tend to be encouraged by the Marshallian forces. The unplanned kind can arise, per Jacobs, from a wide range of potential interactions that might take place in a city. Looking at coagglomeration, there are some industry pairs that coagglomerate to a high degree. Their tendency to co-locate is likely to represent, at least to some extent, deliberate agglomeration arising from strategic migration choice or from differences in entrepreneurial survival. Other industry pairs do not coagglomerate so much. Jacobs' agglomeration economies are to be found in these industry pairs. We therefore investigate heterogeneity in the response to Marshallian forces between more- and less-coagglomerated industry pairs to identify Jacobs-type agglomeration economies.

B. Other non-Marshallian approaches

Jacobs is not, of course, Marshall's only important successor in the study of agglomeration. In Chinitz (1961), New York is different from Pittsburgh because it is more diverse and because its industry is organized in a less-concentrated fashion, making it a friendlier environment for startups and innovation. Porter (1990) argues that competition is healthy for a business cluster, a point that echoes Chinitz. Saxenian (1994) extends this analysis further, arguing that the differences are not simply matters of industrial organization, but are instead deeper matters of business culture. Vernon (1960) writes about the importance of "instability" for increasing-returns industries, arguing that newer industries with more entry are the ones that benefit more from the options that are present in a large city. Others have emphasized the importance of human capital (e.g., Rauch, 1993, and Glaeser and Saiz 2004) and creativity (Florida, 2003), both of which are obviously related to a city's adaptive capacity. Like for Jacobs, we believe that these approaches to agglomeration should be seen as complements to Marshall rather than as substitutes or alternative explanations. This intuition informs our empirical work.

As discussed in the Introduction, there is a long record of empirical research that has considered Marshall's three forces. The literature on non-Marshallian forces is much less developed. On the theory side, Duranton and Puga (2001) show how cities can serve as "nurseries" for infant firms and industries, while Strange et al. (2006) focus on the broad issue of uncertainty in the sense of Vernon. Helsley and Strange's (2001) model of input sharing captures Chinitz and Porter effects. On the empirical side, Glaeser and Kerr (2009) and Rosenthal and Strange (2010) find evidence consistent with Chinitz by considering entry patterns. Rosenthal and Strange (2003) also consider the role of firm size, in this case by considering growth.

In order to test the non-Marshallian mechanisms discussed above, we examine the pattern of heterogeneity in Marshallian agglomeration effects using a theory-driven sectoral breakdown. To begin, we use information collected by the OECD in 1997 to classify our sectors as high- or low-technology (see the Data Appendix for more details). Next, we gather data on the share of college graduates in each industry

using the LFS over the period 1995 to 1999 and split sectors according to whether this share is above or below the median (at 0.078). Finally, we use information gathered from within the BSD to split our sample along the following dimensions: (a) sectors where the first year of opening of currently operating plants is above or below the median across all years and sectors;¹⁰ (b) sectors where the share of entrants – i.e. the incidence of new firms at time t in the total number of firms in that year – is above or below the median across all years and industries; (c) sectors where the average size of the entrants – i.e. firms operating at time t that did not exist at time $t-1$ – is above or below the median size across all years and sectors; and (d) sectors where the average size of the incumbents – i.e. firms operating both at time t and $t-1$ – is above or below the median size across all years and sectors. Given that the level of observation in our dataset is the industry pair, we use this information to classify combinations where both sectors belong to one group – for example both high-technology or both low-technology – and mixed pairs where the two sectors belong to different groups – for example one high- and one low-technology.¹¹

These groupings of industry pairs and some related descriptive statistics are presented in Table 2. In the top two panels, we look at new vs. old pairs and dynamic vs. steady combinations. We find that all three Marshallian forces increase as we move towards less turbulent sectors with older companies and little entry of new firms. The intensity of input-output sharing and knowledge flows is 2-2.5 times larger among old and steady pairs than for new/dynamic combinations. Similarly, the extent of labor pooling increases by 25%-30% when going from turbulent to steady pairs. We also find that new/dynamic pairs are more localized than mix and old/steady pairs, with localization measured by Ellison and Glaeser's (1997) γ . As might be expected, the average share of entrants across the two sectors decreases as we move from dynamic pairs (at 13.1%) to combinations with fewer new firms (at 8.2%). As for the first year of establishment of currently operating plants, this is on average 1972 for young sector combinations and almost 30 years older (1945) for long-established sectors.¹²

In the two central panels we consider high- vs. low-technology and high- vs. low-education pairs. Many more sector combinations are either low or mix-tech (20,532 and 24,780, respectively) than high-tech (7,140). The intensity of labor pooling and knowledge flows is more pronounced among high-tech pairs and declines as we move towards industries in the low-tech group. This tendency is not so clear when looking at input-output sharing. Finally, high-technology pairs are less localized than mix and low-technology pairs. A

¹⁰ We rank industries by the age of the oldest *currently operating* plant, not the age of the industry itself. We believe that this captures the degree to which an industry's operations are settled.

¹¹ Table 2 reports the exact values of the medians.

¹² Note that the number of observations/pairs in these partitions is balanced between top and bottom, and approximately twice as large in the mixed groups. This pattern is mechanical and due to the fact that we split by using the median. This remark applies to all remaining groupings, except for the technology partitioning which is based on the OECD (1997) classification.

similar pattern emerges when we rank sectors on the basis of the share of college graduates. The incidence of college graduates (on average across the two sectors) is expectedly higher in the top group (at 14.8%) than in the other two groups (at 9.7% and 4.7%).

Lastly, in the bottom two panels we rank pairs by the size of entrants and incumbents. We do not detect very strong patterns in terms of the extent of the three different Marshallian forces. The intensity of labor pooling seems roughly constant across all groups, while knowledge flows seem more prevalent between sectors with small incumbents and large entrants. Conversely, we find that input-output sharing increases as we move from small to big sector pairs. As for the localization of these industries, pairs with large entrants and incumbents are more geographically concentrated than mixed and small pairs. Expectedly, the average size of entrants (at 5.7 employees) and incumbents (at 10.5) is approximately twice as small among small pairs than for sectors in the large grouping (the corresponding figures are 14.5 and 33.7).

IV. Coagglomeration and Marshallian microfoundations: UK Evidence

A. Univariate and multivariate OLS regression analysis

In this section we study the microfoundation of agglomeration economies by linking the proxies for the three Marshallian forces discussed above to a measure of industrial-pair coagglomeration. Our results come from regressions of the following kind:

$$\gamma_{ijt}^C = \alpha + \beta_{LP} LP_{ij} + \beta_{IO} IO_{ij} + \beta_{KS} KS_{ij} + \sum_{k=1}^5 \lambda_k \text{Diss}_{ij}^k + \varepsilon_{ijt}, \quad (2)$$

where γ_{ijt}^C is the Ellison et al. (2010) measure of coagglomeration between sectors i and j at time t ; LP_{ij} , IO_{ij} and KS_{ij} are proxies for labor-pooling, input sharing and knowledge spillovers between sectors i and j averaged over the relevant years (see Section III for details); and Diss_{ij}^k is one of the five measures of dissimilarity between sectors i and j in terms of use of primary resources and non-manufacturing inputs. Finally, ε_{ijt} is an error term uncorrelated with all other variables. Note that we allow for an arbitrary degree of correlation in the shocks of sector pairs over the years and cluster standard errors at this level. As already discussed, the dataset consists of 4,371 unique combinations of 94 manufacturing sectors over 12 years, giving a total number of 52,452 observations. Throughout the regression analysis, we standardize our variables to have unitary standard deviation. This eases comparison of the relative strength of different agglomerative forces.

The first set of results is presented in Table 3. Columns (1) and (2) tabulate results from univariate regressions where we consider only one Marshallian force at the time (and include dissimilarity controls in Column 2). The results show that labor pooling has the largest and most significant association with coagglomeration. A one standard deviation increase in this measure corresponds to 19% of a standard

deviation increase in coagglomeration. The R-squared associated to this regressor is 0.037. We find that a one standard deviation change in input-output sharing has a 14% of a standard deviation impact on coagglomeration and an associated R-squared of 0.033. A smaller effect is instead associated with knowledge spillovers where a one standard deviation change in our proxy is linked to around 10% of a standard deviation change in γ^C with a related R-squared of 0.011. Note that the joint R-squared of the five dissimilarity measures is approximately 0.039, suggesting that natural advantages explain approximately as much of the variation in coagglomeration as labor pooling and input-output sharing, but three times as much as knowledge spillovers. This pattern is consistent with the results by Ellison et al. (2010), who also document weaker agglomerative effects from knowledge spillovers than from worker and input sharing. Interestingly, controlling for the five dissimilarity proxies does not change in any meaningful way the effects of the three Marshallian forces suggesting that access to natural resources and local non-manufacturing industries do not bias the results in the simple models without the additional controls.

Columns (3) and (4) present regression coefficients from the multivariate regressions. We still find that labor pooling has the strongest relation with coagglomeration with an estimated effect of approximately 0.16 of a standard deviation. On the other hand, the impacts of input sharing and knowledge spillovers decline to 0.082 and 0.024-0.031 respectively – though both retain statistical significance. A test for the equality between the effects of LP_{ij} and either IO_{ij} or KS_{ij} rejects the null with p-values of 0.022 and 0.015, respectively. Conversely, an equality test on the effects of IO_{ij} and KS_{ij} accepts the null with a p-value of 0.896. The joint R-squared of the three Marshallian proxies (from the specification in Column 3) is 0.045 – slightly higher than the joint R-squared of the dissimilarity indices. The combined effect of agglomerative forces and natural advantages (Column 4) is approximately 0.06. Relative to the results of Ellison et al. (2010), we find that labor pooling matters more than input-output sharing, and that the effect of knowledge spillovers is more attenuated when the three agglomerative forces are included together. However, all in all our findings are comparable to theirs as they show that all three of Marshall’s forces matter.¹³

To further corroborate our results, we carry out a number of robustness checks and extensions. Initially, we study whether upstream linkages are more important than downstream connections. Column (1) of Table A2 in the Appendix shows that the effect of input sharing is twice as large as the effect of output sharing, but this distinction is not significant (p-value of an equality test: 0.533) and does not affect the other coefficients. Further, we investigate whether focusing on a specific year in our sample changes the overall picture. Columns (2) to (4) of Table A2 present results for 1997, 2002 and 2008 separately. Although we

¹³ Note also that we follow Ellison et al. (2010) and do not correct γ^C for differences in the variance of the area-industry employment shares. Their Mathematical Appendix explains why this is not necessary. Nevertheless, we assess the robustness of our findings against this issue by including in our specification industry i and industry j dummies. When we do this, we find very similar labour-pooling effects; slightly smaller, but still significant input-output effects; and larger and more precisely estimated knowledge-spillover effects.

find a slight attenuation in the effect of LP_{ij} , IO_{ij} and KS_{ij} as we move towards more recent years, the differences are not substantial. This pattern can be partly explained by the fact that our proxies are measured at the beginning of the observation period (1995 to 1999). Using proxies calculated over longer time periods yields more stable and more significant coefficients across the years.¹⁴

We also investigate whether using a proxy for knowledge spillovers based on the sector-of-use (SOU) probabilistic mapping affects the findings. Results are presented in Table A3 in the Appendix. The broad conclusion reached so far still holds: all three Marshallian forces matter, though the effect of labor pooling seems somewhat stronger. We also find that the impact of knowledge spillovers is stronger when using this proxy and very close to the effect of input-output sharing, which is instead weakened. As discussed, this alternative mapping is partly based on the technology (and the related patents) contained in goods that are bought and sold as intermediates across industrial sectors. Therefore, it incorporates some of the linkages stemming from input-output sharing and attenuates the effect of IO_{ij} . Given this issue, our preferred proxy is the one based on the industry-of-manufacture (IOM) probabilistic mapping which we will use throughout the rest of the paper.¹⁵

Finally, we check that our results are not affected if we change our measure of coagglomeration to: (a) be based on number of plants as opposed to total employment; (b) be based on local units belonging to single-plant enterprises only; (c) be based on both urban and rural areas. We also experiment with excluding publishing (SIC221) and printing and related activities (SIC222) since these sectors are classified among services in the US industrial classification. None of these robustness checks affect our findings.¹⁶

B. Addressing endogeneity concerns

The literature on the microfoundations of agglomeration economies has put forward two sources of possible biases in OLS estimation, namely: (i) reverse causation; and (ii) sorting. In this section, we discuss these issues from a theoretical point of view and provide a set of robustness checks and Instrumental Variable (IV) estimates to address them empirically.

¹⁴ If we collapse γ_{ijt}^C to its average across all years in our sample and run regressions that exploit variation over 4,371 observations only, we find identical results to those presented in Table 3.

¹⁵ One related concern is that input-output linkages might partly capture knowledge spillovers because our patent-based proxy measures the latter imprecisely. To investigate this issue, we run specifications where we further include two-way and three-way interactions of the Marshallian forces. The only significant interaction is the one between input sharing and knowledge spillovers, which carries a negative and significant coefficient of -0.009. In this specification, the effects of IO_{ij} and LP_{ij} remain very similar, while the effect of KS_{ij} rises to around 0.100 (significant). This suggests that the input-sharing and knowledge proxies do *not* capture similar effects. Moreover, this pattern is consistent with the sectoral heterogeneity presented later in the paper, where we show many instances in which pairs that significantly respond to knowledge spillovers are less affected by input-output links.

¹⁶ Results are not tabulated for space reasons, but are available upon request.

The reverse causation argument is laid out by Ellison et al. (2010). Firms in industries with strong Marshallian links could choose to locate together in order to benefit from those links. Alternatively, firms that locate together for other reasons could forge Marshallian links. In contrast to Ellison et al. (2010), we see the reverse phenomenon of coagglomeration leading to productive links as being itself a type of agglomeration economy. For instance, if two firms realize after choosing locations that they can hire from the same labor market, then they benefit from labor market pooling. Similarly, if two firms learn from each other *ex post*, then the resulting technological improvement is an instance of knowledge spillovers. These agglomeration economies are in fact in the spirit of Jacobs (1969), who gives numerous examples of accidental agglomeration economies. Even so, we describe below a strategy to address this issue and thus arrive at estimates that capture the effect of Marshallian links on agglomeration, rather than the reverse.

As for the sorting issue, the main concern is that agglomeration – which increases productivity in possibly unobservable ways – might be correlated with coagglomeration. To clarify matters, consider two industries, e.g. apparel and printing/publishing, which are agglomerated for historical reasons in London. Assume that both industries are highly productive either because of some advantages connected to the location or because of selection at work in this competitive environment. Further assume that more productive industries use a wider range of inputs. For example, they are more likely to use a variety of workers because their human resources divisions are better at spotting the right type of workers in a large agglomerated market. Conversely, think of two other industries, e.g. wood laminate and manufacturing of furniture, that operate in a non-agglomerated region, are low productivity and not very efficient at sharing workers. Estimating the effect of labor pooling on coagglomeration by comparing these pairs would bias the results by conflating the true effect of LP_{ij} with a productive advantage arising because of urbanization economies enjoyed by firms locating in more agglomerated places. Although this is a logically correct argument, the unobservables that would give rise to these patterns would need to have a very particular structure and imply that agglomeration is correlated to both coagglomeration *and* the linkages between sectors that we measure with our proxies. One of the advantages of the method developed by Ellison et al. (2010) is that, by studying the relation between co-location and industry-pair links, this approach deals with a number of unobservables that are not easily related to these pair-specific linkages. Therefore, we do not believe the arguments brought forward in the productivity literature in terms of sorting-related endogeneity are very compelling and sufficiently strong to undermine our findings. Nevertheless, we next provide a number of additional results that lend further support to our conclusions.

A first set of regression results is presented in Table A4 in the Appendix. Column (1) mitigates reverse causation by staggering our regressions and considering the effect of the three Marshallian forces measured up to 1999 on coagglomeration γ_{ijt}^C for the years 2000-2008. This check confirms our previous results. Columns (2) to (5) investigate whether any correlation between agglomeration and coagglomeration

has the potential to bias our findings. To begin with, we exclude London – the biggest agglomeration in the UK – from the calculations of γ_{ijt}^C and re-estimate our empirical models. Although we find that the effect of KS_{ij} is attenuated and the impact of IO_{ij} is marginally stronger, our broad conclusions are essentially unaffected. In Columns (3) and (4) of the table we include proxies for the extent of agglomeration of the areas where the two sectors in the pair are operating. In particular, we include: (i) the mean population density of all the TTWAs in which the sectors are operating, averaged across the pair (Column 3); and (ii) the mean employment density of all the areas where the sectors are operating, averaged across the pair (Column 4). Employment density is calculated as total employment across all sectors in a TTWA divided by the area size expressed in square kilometers, so this proxy captures general urbanization economies – much as population density – stemming from operating in a larger market.¹⁷ Note that we calculate simple average densities across TTWAs. Weighted averages that take into account the size of the TTWAs where the sectors are more heavily represented yield similar results. As shown in Columns (3) and (4), appending these controls to our regressions has very little effect on our estimates. Finally, in Column (5) of the table we add to our specification the average Herfindahl index across the sector pair to check whether industrial concentration (as opposed to urbanization economies) affects our results. Once again, we find no evidence that our results are sensitive to these considerations and confirm our previous conclusions.

To conclude this section, we discuss a number of IV regressions where we instrument the three Marshallian forces using proxies constructed using US data. This approach follows Ellison et al. (2010). We instrument LP_{ij} using a measure of the correlation between sector pairs in their use of different types of workers as categorized by the National Industrial-Occupation Employment Matrix (NIOEM) published by the Bureau of Labor Statistics. We instrument IO_{ij} with an identical measure obtained using the 1987 Benchmark Input-Output Accounts published by the Bureau of Economic Analysis (BEA). Finally, we instrument the flow of patent citation among UK inventors as recorded by the EPO using the flows of citations among US inventors as tracked by the NBER Patent Database. In both cases we line up technological classes to industries using an IOM probabilistic mapping. More information on the data construction is provided in the Data Appendix.¹⁸ The validity of this approach relies on having thoroughly controlled for co-location that is driven by natural advantages and shared use of non-manufacturing resources. Hence in all our IV specifications we control for the five proxies for sector dissimilarities.

Results are presented in Table 4. Following Ellison et al. (2010), we exclude from the sample sector pairs where the two three-digit industries fall in the same two-digit group and a number of sectors which were aggregated in the data construction process. The note to the table provides more details. Columns (1)

¹⁷ The correlation between the two urbanization proxies and coagglomeration is small and negative at -0.148 for population density and -0.082 for employment densities. These numbers shrink to zero and 0.007 if we exclude London.

¹⁸ We are extremely grateful to William Kerr for sharing his data and codes with us.

and (3) show that OLS results do not change as a result of these exclusions. Column (2) presents IV regressions where we include and instrument one Marshallian force at the time. The IV coefficients are very close to their OLS counterparts in Column (1). Column (4) presents multivariate IV regressions where we enter and instrument all three Marshallian forces simultaneously. We find positive and significant effects for LP_{ij} and IO_{ij} . The size of the associations here is similar to the OLS counterparts (see Column 3). However, we find that KS_{ij} loses its significance and turns slightly negative. We believe this can be attributed to some collinearity between measures that makes instrumented knowledge spillovers hard to disentangle from labor pooling and input-output sharing. A similar argument is put forward by Ellison et al. (2010) who report in their appendix weak results when instrumenting KS_{ij} . To partly address this issue, in Column (5) to (7) we enter the proxies for Marshallian forces two at the time. In Column (5), we only include LP_{ij} and IO_{ij} and confirm that both have a positive and significant association with γ_{ij}^C . In Column (6), we only consider IO_{ij} and KS_{ij} . We now find that both measures are positively and significantly associated with coagglomeration and that the effect of knowledge-spillovers is very similar to the one documented using OLS regressions. Finally, in Column (7) we include and instrument LP_{ij} and KS_{ij} and find that both are positively associated with coagglomeration. Although only the effect of labor pooling is significant at conventional levels, the impact of knowledge spillovers points in the right direction and is reasonably sizeable at about half of its OLS counterpart. All in all, we take the evidence in Table 4 as confirming our previous findings. Moreover, it supports our claim that endogeneity is not likely *a priori* to significantly bias OLS results.

V. Heterogeneous agglomeration: Jacobs meets Marshall

This section begins the presentation of results that allow for heterogeneity across industries. Specifically, it takes a new approach towards examining Jacobs' (1969) analysis of how new work is created by exploring complementarities between Jacobs and Marshall. The approach has at its core a simple idea: the coagglomeration of industries that only rarely co-locate is different than the coagglomeration of industries that are often found together. It is the former that captures the sorts of unplanned interactions that Jacobs has in mind. We estimate Marshallian models of the sort discussed in the previous section without constraining random and planned coagglomeration to have the same effect in order to assess the presence of these sorts of interactions. More precisely, we estimate equation (2) in a way that allows the effects to vary between the most and least coagglomerated pairs. The estimation is carried out using quantile regressions that simultaneously include all three Marshallian forces as well as controls for natural advantages. Figures 1a–1c

present the results for labor pooling, input-output sharing, and knowledge spillovers, respectively. The confidence intervals on the figures come from bootstrapped standard errors clustered on industry pairs.¹⁹

It is immediately clear that the pattern presented in Table 3 conceals considerable variation between industry pairs. Table 3 showed all three forces are at work, with somewhat weaker evidence for knowledge spillovers. Figures 1a–1c show that the agglomeration forces have effects that vary substantially across groupings of industry pairs.

Figure 1a presents results for labor pooling. There is clear heterogeneity across pairs according to their coagglomeration. While labor market pooling has a positive and statistically significant contribution to industry-pair coagglomeration across the board, the effect is much larger for the less coagglomerated pairs. Specifically, labor pooling has an impact of around 0.22 and 0.16 (both significant) for industry pairs in the two bottom deciles declining to around 0.06-0.08 (significant) in the top half of the coagglomeration distribution. Figure 1b shows a pattern for input-output sharing that is exactly opposite. For this force, the effect is larger for the most coagglomerated pairs. The input-sharing impact increases from approximately 0.03 (insignificant) for the bottom decile to 0.15 and 0.23 (both significant) for the top two deciles. Figure 1c shows yet another pattern, with the marginal effect of knowledge spillovers positive and significant up to the 60th percentile. The coefficient becomes smaller over the range, declining from 0.03 in the bottom decile to 0.02 at the median, and the estimation becomes increasingly imprecise for the most coagglomerated industries, for which the effect is no longer significantly different from zero.

More importantly for our purposes, the pattern of heterogeneity has interesting implications for Jacobs' ideas regarding unplanned interactions. The input sharing results in Figure 1b are contrary to Jacobs. They suggest that input sharing primarily impacts the collocation of pairs that coagglomerate extensively. There is little to be gained from links between industries that are not very coagglomerated. In other words, there is strong evidence that interactions that are most likely to be planned have the largest marginal effect on coagglomeration. It is useful to consider an example from Marshall. He writes that:

“Many cutlery firms [in Sheffield] for instance put out grinding and other parts of their work, at piece-work prices, to working men who rent the steam power which they require, either from the firm from whom they take their contract or from someone else.” (Marshall, 1890, 8th ed., p. 172).

Our data show that Sheffield continues to be a center of cutlery production today. Moreover, cutlery and manufacturing of basic iron and steel are one of the most highly coagglomerated sectors. It is entirely understandable that a cutlery maker would deliberately plan its location in a way that secures its metal input

¹⁹ Note that we performed a similar analysis where we investigate heterogeneity in the effect of the Marshallian forces by stratifying our sample *directly* along the quantiles of γ^C , i.e. by considering on the quantiles of the unconditional coagglomeration distribution. This approach yields similar results.

supply. See Table A1 in the Appendix for some other highly coagglomerated industry pairs, such as spinning of textiles and textile weaving – also discussed by Marshall in a similar fashion.

On the other hand, the results on labor market pooling and on knowledge spillovers are much more in the spirit of Jacobs. Regarding knowledge, the effects are not even significant for highly coagglomerated pairs. As for labor market pooling, the effects diminish drastically as coagglomeration increases. In other words, both of these sorts of interactions between industries have a larger marginal effect when the industries co-locate less frequently and interactions are likely to be of an unplanned nature.

There is an interesting parallel here to the Duranton and Puga (2001) nursery-city phenomenon. Certain interactions seem to have greater effects with less frequent co-location. Others have greater effects with more frequent co-location. This pattern points to agglomeration effects that depend on the maturity of the industry, as in Duranton and Puga and also Vernon (1960). We will return to this issue below.

These initial Jacobs-type results fairly strongly suggest that there is heterogeneity in the impacts of Marshallian forces. The next section looks more systematically at this heterogeneity by considering adaptation and industrial organization.

VI. Heterogeneous agglomeration: Adaptation and organization

A. Non-Marshallian approaches

This section presents empirical results on the adaptation and organizational aspects of agglomeration economies. We extend the traditional Marshallian approach by examining how the patterns of coagglomeration depend on the interaction between Marshallian forces and the nature of the industry in question. We thus have a unified approach that nests a range of microfoundations in the same model. As above, Marshall's microfoundations are allowed to be complementary to other explanations.

As a preliminary step, we look directly at the relationship between the industry differences and coagglomeration by including additional controls in equation (2). Table 5 presents the results from these models. The first important conclusion from the table is that the coefficients of the Marshallian forces are fairly constant as different controls are introduced. The coefficients are also similar in magnitude to the corresponding effects in Column (4) of Table 3 showing that estimates of the Marshallian forces are robust.

Turning to the non-Marshallian variables, an interesting pattern emerges. We find a positive and significant coefficient on year-of-opening (the inverse of age) in Column (1) and an insignificant coefficient on entry share in Column (2). The former result is consistent with the nursery city/unstable industry/unplanned interactions ideas discussed above. The latter is weakly supportive. The dummies for high- and mixed-technology pairs in Column (3) are instead both negative and significant. Controlling for Marshallian forces, we see more coagglomeration of low-technology industries. This result is the opposite of what one might expect to find based on the predictions of a nursery-city model. In Column (4), the

coefficient of average college share is significant and negative. Given the strength of human capital effects in other models (e.g., Rauch, 1993, or Rosenthal and Strange, 2008), this is somewhat unexpected.²⁰ Finally, entrant size has a positive and marginally significant coefficient in Column (5), as does incumbent size in Column (6). Controlling for Marshallian forces, we do not find much of a small-firm effect.

In sum, simply including controls for non-Marshallian forces in a coagglomeration/microfoundation model, while failing to allow for heterogeneity, generates a weak and sometimes puzzling pattern of results. We now therefore turn to less restrictive models that let the Marshallian effects differ across industries.

B. New industries and entry

We begin by looking at adaptation and nursery city ideas by estimating models based on two different partitions of industry pairs. The first focuses on industry age. As noted above, this is captured in the data by the age of the oldest active plant. The second focuses on the share of new firm entry in the industry. This approach generates the partitions of pairs detailed in Section III. In the case of age, the new industries group includes pairs where both sectors are younger than the median age across industries. In the case of entry, the dynamic industries group includes pairs where both industries have an entrant share that is above the median. In both cases, we include additional corresponding controls for, respectively, age and entry share averaged across the two sectors in the pair.

Results are reported in Table 6. For industry age, we find the largest agglomeration effects for new industry pairs. This is true for all three Marshallian forces. For labor market pooling, the effects are smaller for the mixed and the old industry pairs (at 0.153 and 0.081, respectively) than for new pairs (at 0.310). However, all coefficients are highly significant and the extent of variation in these effects is more muted than for the other two Marshallian forces. The knowledge spillover results are very much in the nursery city/unstable industry spirit discussed above. They show knowledge effects that are five to ten times stronger for young pairs, at 0.236 (significant), than for mixed and old pairs, at 0.040 (significant) and 0.026 (insignificant), respectively. The same is true for the results on input sharing. The coefficients move from 0.270 (significant) for new industry pairs to 0.049 (insignificant) for the mixed group, and finally to 0.041 (significant) for old pairs. While the pattern of effects on input sharing is not consistent with a nursery city model, the heterogeneity in the coefficients for knowledge spillovers clearly supports this theory.²¹

The results have a relatively similar pattern for industry dynamism. Labor market pooling is always significant with the coefficients for the three groups fairly constant and ranging between 0.181 and 0.144.

²⁰ These result hold if we exclude London. Conversely, dropping the three Marshallian proxies from the specifications, yields insignificant estimates of the effect of either human capital or technology on coagglomeration.

²¹ We further investigate whether new/old pairs respond differently to Marshallian linkages when they are measured closer/further in time relative to coagglomeration. To do this, we run separate regressions for 1997, 2002 and 2008. The patterns presented in Table 6 were confirmed with no evidence of additional significant heterogeneity.

We still find that the largest result for knowledge spillovers occurs for the dynamic industries at 0.181 (significant). This shrinks to 0.033 (significant) and -0.020 (insignificant) for mixed and steady industry pairs. Input-output linkages are closer to a nursery pattern in these dynamic-industry models than in the previous age grouping. Input sharing displays significant coefficients for the mixed and the steady pairs, at 0.103 and 0.052 respectively, and has no significant effect for dynamic industries.

C. High-technology and high-education

We now turn to the related issue of how the relationship between Marshallian forces and coagglomeration depends on the technological status of the industry in question. As noted above, we characterize an industry's technological status (high-tech or not) according to the OECD (1997) classification. This generates three types of industry pairs: both high-technology, both low-technology, or mixed. We estimate equation (2) for each type.

Results are reported in Table 7. We find that labor pooling is significant in all three groups, but its effect on coagglomeration is much larger for the low-technology industry group (at 0.332) than for the high-technology group (at 0.046). Input sharing also has the largest coefficient in the low-technology group, at 0.091. While this Marshallian force has positive coefficients for all three groups, the high-technology coefficient is small and insignificant. Knowledge spillovers display the opposite pattern. The largest coefficient is found for high-technology (significant at 0.053), while the effect becomes small and insignificant for low-technology (at 0.039).

These results clearly show agglomeration economies are not simply a high-technology phenomenon. Labor pooling has a stronger effect in the low-technology group, while input sharing is stronger in the mixed- and low-technology groups of industry pairs. Knowledge spillovers, reassuringly, is different with the largest effect for high-technology industries. This suggests that some of the weaker results for knowledge spillovers reported above in Tables 3 and 4, and also presented in Ellison et al. (2010), arise because the sample includes low-technology industries (as well as old and steady industries) where knowledge spillovers are not important.

Table 7 also tabulates results of a similar exercise where industries are partitioned according to the education levels of their workers. The one different element here is that we further control for the average share of college graduates across the pair in order to control for direct effects of this variable within groups.²² The pattern of results is similar to the high-technology vs. low-technology heterogeneity discussed above.

²² While it is nearly universal to equate education and skills (e.g., Glaeser and Mare, 2001), Bacolod et al. (2009a, 2009b) argue that education and skills are not identical, with the former being only one of many inputs into the latter. In our setting, their procedure for identifying skill from occupation does not seem natural, since we are working at the aggregate rather than at the individual level.

Knowledge spillovers have significant effects in high-education (at 0.048) and mixed education (at 0.050) industry pairs, but not in low-education pairs (insignificant at 0.030). Conversely, input sharing and labor pooling have the largest and most significant effects in low-education pairs, at 0.123 and 0.391 respectively. These shrink to 0.007 (insignificant) and 0.046 (borderline significant) for the high-education pairs.

Taken as a group, these findings are broadly consistent with learning playing an important role in the agglomeration process. Jacobs (1969) calls this phenomenon “the creation of new work”. Vernon (1960) instead discusses the process by which new products reach stability. The evidence is also consistent with Duranton and Puga’s (2001) nursery-city phenomenon, where new products are created in diverse cities and move to specialized cities upon reaching maturity. The authors provide evidence of firm migration following this pattern in France to support their conclusions. To the best of our knowledge, our paper is the first one to have examined coagglomeration in this light. The observation that only high-technology/high-education pairs are found to have their coagglomeration encouraged by knowledge links between the industries is consistent with the nursery-city idea. So is the finding that low-technology/low-education pairs have coagglomeration encouraged by the somewhat more routine labor and input links.

D. Industrial organization

The final set of results deals with industry structure. We present two groups of estimates. The first uses a partition based on the size of entrants. The second uses a partition based on the size of incumbents. Both splits correspond to Chinitz (1961), who argues that the presence of small firms allows entry by other small firms.²³ Naturally, there are other treatments of agglomeration that similarly depend on industry structure, and they will be discussed as well.

Results are presented in Table 8. As above, the table reports results for the three Marshallian forces estimated over three groups of industry pairs. The models include controls for, respectively, entrant size and incumbent size averaged across the pair, alongside the usual controls for natural advantages. The results are consistent with a small-firm effect, with the largest coefficient found on input sharing for the small entrants and small incumbents models (significant at 0.193 and 0.159, respectively). It is worth noting that input sharing is significant for all groups in both models. However, the effects are substantially smaller for industries characterized by mixed- and large-entrants and incumbents, and ranging between 0.068 and 0.082. The labor pooling coefficients are all significant and more comparable in sizes in the three entrant- and the three incumbent-size models. The mixed-incumbent industries display the largest labor pooling effects at

²³ Strictly speaking, Chinitz’s argument is about heterogeneity in the industrial organization of cities – not of sectors. However, we find that approximately 96% of the variation in the size of entrants and 99% of the variation in the size of incumbents is within-TTWAs across sectors, with the remaining part being within-sectors across TTWAs. This suggests that our analysis that treats industrial composition as fixed across cities and only focuses on sectoral heterogeneity capture the most relevant variation.

0.223. However, the coefficients for this agglomerative force are above 0.10 for all but one grouping and mainly in the 0.15-0.20 range. Finally, as in all the regressions so far, knowledge is not a statistically significant predictor of coagglomeration in the universe of models. In this case, we find that knowledge flows are only significant and sizeable in the small-entrant sample (at 0.051) and in the mixed-entrant group (at 0.028). They are instead insignificant and small, but always positive and between 0.020 and 0.040 for all three groups based on the size of incumbents.

Chinitz focused largely on input sharing among small firms being a driver of agglomeration. Our results clearly confirm this prediction. Conversely, Vernon and Jacobs offer anecdotes of knowledge spillovers generated by large and by small firms alike. Our results on coagglomeration seem to suggest, however, that the effect of knowledge flows is more systematic for small firms. Regarding labor pooling, neither Vernon nor Jacobs directly engages with the implications of firm size for this agglomerative force. While the coefficients on labor pooling are significant for all sizes of entrants and incumbents, we find the smallest coefficients for small entrants and small incumbents. One factor that could potentially come into play is labor poaching, where firms hire away each other's skilled workers (Combes and Duranton, 2006). If small firms are threatened to a greater degree by the possibility of poaching, they might be less likely to co-locate with firms hiring from the same labor pool.

E. Robustness

As discussed in Section IV, it is possible that the OLS results might be affected by endogeneity issues. We address these concerns here similarly to our approach in Section IV.B. To check whether any correlation between coagglomeration and agglomeration biases our findings, we estimate the models reported in Tables 6-8 excluding London. We also estimate specifications that control for density measured either as employment or population per land area. Furthermore, we estimate IV models which deal with both reverse causation and omitted variables.²⁴ Results are available on request.

The key patterns of results continue to hold. For industries where adaptation is crucial, we find stronger evidence of knowledge spillovers. This is again in the spirit of Jacobs, Vernon, and Duranton and Puga. For industries dominated by small firms, we find a stronger relationship between input sharing and coagglomeration, supporting the theories of Chinitz. Finally, labor pooling continues to have a consistently significant and sizable relationship with coagglomeration across all groups.

One further issue is that the pattern documented above might be related to the extent of localization of the industry pairs in the different groups. To consider this possibility, we perform all the regressions in

²⁴ We estimate both univariate and multivariate IV models that yield similar results.

Tables 6 to 8 adding a control for the average localization index (Ellison-Glaeser, 1997, γ) across the pair. This does not affect the results in any significant way.²⁵

VI. Conclusion: Single-industry models

This paper has considered the microfoundations of agglomeration economies using patterns of coagglomeration of UK industries. The analysis reached several key conclusions. First, there is notable heterogeneity between industries. Different industries are impacted differently by Marshallian factors. Second, this heterogeneity is consistent with Jacobs' (1969) ideas about unplanned interactions being an important aspect of the agglomeration process for labor pooling and knowledge spillovers, but not for input sharing. Taken together, these two points suggest that Jacobs and Marshall should be seen as complementary rather than as alternatives. Third, the pattern also provides support for the idea of nursery cities (Duranton and Puga, 2001) in particular, and adaptive agglomeration economies more in general (Vernon, 1960). Finally, the pattern of heterogeneity is also consistent with Chinitz's (1961) ideas that agglomeration effects have an organizational dimension in the sense of being stronger in industries dominated by small firms.

One reason why it is important to recognize this heterogeneity is that there are numerous instances in the agglomeration literature where the circumstances of individual industries and clusters are presented as having broad relevance across industries. While the detailed analysis of individual cases is often quite informative – as attested by the influence of this kind of extrapolation – it is important not to simply accept this generalization without further investigation.

We conclude the paper by considering a few salient instances of agglomeration. This is done in Table 9, which presents two sorts of information. First, it places the chosen industries in our scheme of sectoral division, describing, for instance, whether the industry is high-technology or not. Second, it presents the results of coagglomeration models estimated for individual industries. This approach exploits the variation between how one industry co-locates with the other 93 sectors over the 12 years of data that we use. While it is natural to estimate such models since they allow for maximal heterogeneity, it is not surprising that the results are somewhat noisy given the limits imposed on this approach by the data.

We have already discussed one of the salient industries we use, namely cutlery (SIC286 in Table W1), which was also considered by Marshall (1890). As Table 9 shows, this sector has both small incumbents and entrants. It is a steady industry with relatively moderate entry rates, and relatively young plants. Finally, it is a low technology and low education industry. This industrial organization suggests that one ought to find Chinitz's small firm effects. The regression results in Table 9 show exactly this. The

²⁵ Note also that the correlation between co-agglomeration and average localization of the industry pairs is small and negative at -0.028, so localization features of industries cannot explain the patterns of Figures 1a-1c.

coefficient on input linkages is large and significant (at 0.719). It is twice as large as the coefficient on labor market pooling (at 0.345), which is also significant. Knowledge spillovers are not significant. Although plants in this industry are new, entry rates are low and the industry is both low technology and low education. This evidence suggests that the development of the cutlery industry is likely to be more responsive to policies that facilitate input-output sharing than policies that target knowledge. Since for this industry, as with most others, we also find a significant labor pooling effect, labor market policies also have potential for the cutlery sector.

The picture is different when considering textiles. This sector has been of historical importance for the development of manufacturing in the UK (Landes, 1969). In considering this industry, we estimate a stacked model using data from seven sectors (SIC171-SIC177; see Table W1 for more details). Regarding our sectoral breakdown, there is heterogeneity within this group. Both entrant size and incumbent size are mixed, with roughly half of the industries being characterised by large firms and half by small ones. Regarding adaptation, these industries are largely dynamic, with relatively high entry rates and with young plants. However, the industry's workers are low education in nearly all the sectors, and all the industries are low technology. It is thus more difficult than for the cutlery sector to anticipate what the regressions might show. Table 9 presents large and significant labor pooling effects (at 0.711). We also find significant but smaller knowledge spillovers (at 0.171). Conversely, the coefficient of input sharing is small (at 0.070) and insignificant. These results show that one would not want to generalize from cutlery to textiles, and illustrate more generally the perils of extrapolation.

Without doubt, the computer industry is the salient industry in the agglomeration literature. For example, Saxenian (1994) offers an important and often quoted analysis of the Silicon Valley. In our framework, the computer industry (SIC300) has small entrants and incumbents, dynamic entry and old plants. It is also a high-technology industry, and its workers are highly educated. It is no surprise that the regression results in Table 9 show a very large and significant coefficient on knowledge spillovers at 0.378. The input sharing coefficient is also significant (at the 10% level) but substantially smaller (at 0.033). This industry is somewhat unusual in displaying an insignificant coefficient on labor market pooling (at -0.042). It is worth noting that the latter result does not mean that there is no labor market pooling in this industry. There could be significant labor pooling taking place within the computer sector itself. Note that although the computer industry is similar to the cutlery industry in having small entrants and incumbents, the organizational/Chinitz effects are dominated here by the adaptive/knowledge effects. It is thus manifestly problematic to extrapolate from the cutlery industry to the computer industry. Furthermore, however appealing it might be to use the computer industry to illustrate agglomeration economies, the logic of extrapolating from the computer industry is similarly strained.

The car industry is also highly salient in the agglomeration literature. In the US, this industry's declining cluster centered around Detroit is often contrasted to the prosperous computer cluster in Great San Jose. A very informative discussion along these lines can be found in Glaeser (2011). Somewhat surprisingly, our sectoral breakdown uncovers many similarities between these two sectors in the UK. Looking at Table 9, the car industry (SIC341) is a high technology industry with highly educated workers. It is also marked by small entrants in a dynamic environment. The only two differences with respect to the computer industry are that incumbents in the car sector are large and plants are young. Despite the similarities, the pattern of the regression coefficients differs. Labor pooling has a large and significant effect (at 0.439), while knowledge spillovers have a smaller but still significant coefficient (at 0.176). Input sharing has a small and insignificant impact (at -0.003). This evidence shows that even within sectors that are similar organizationally and with regard to adaptation, the microfoundations of agglomeration can be quite different. Agglomeration is, thus, very heterogeneous.

References

- Almazan, A., A. De Motta, and S. Titman (2007), "Firm location and the creation and utilization of human capital," *Review of Economic Studies* 74, 1305-1327.
- Arzaghi, M. and J.V. Henderson (2008), "Networking off Madison Avenue," *Review of Economic Studies* 75, 1011-1038.
- Audretsch, D.B. and M.P. Feldman (1996), "R&D spillovers and the geography of innovation and production", *American Economic Review* 86, 630-640.
- Bacolod, M., B. Blum, and W.C. Strange (2009a), "Skills in the City", *Journal of Urban Economics* 65(2), 127-135.
- Bacolod, M., B. Blum, and W.C. Strange (2009b), "Urban Interactions: Soft Skills vs. Specialization", *Journal of Economic Geography* 9(2), 227-262.
- Berry, C.R. and E.L. Glaeser (2005), "The divergence of human capital levels across cities", *Papers in Regional Science* 84, 407-444.
- Breschi, S. and F. Lissoni (2004): "Knowledge networks from patent data: methodological issues and research targets", CESPRI Working Paper 150.
- Chinitz, B. (1961), "Contrasts in agglomeration: New York and Pittsburgh", *American Economic Review* 51, 279-289.
- Combes, P.-P., and G. Duranton (2006), "Labour Pooling, Labour Poaching and Spatial Clustering", *Regional Science and Urban Economics* 36, 1-28.
- Duranton, G. (2011), "California dreamin': the feeble case for cluster policies," *Review of Economic Analysis* 3, 3-45.
- Duranton, G. and D. Puga (2001), "Nursery cities", *American Economic Review* 91(5), 1454-1477.
- Ellison, G. and E.L. Glaeser (1997), "Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach", *Journal of Political Economy* 105, 889-927.
- Ellison, G., E. L. Glaeser, and W. R. Kerr (2010), "What causes industry agglomeration? Evidence from coagglomeration patterns", *American Economic Review* 100, 1195-1213.
- Fallick, B., C. Fleischman and J. B. Rebitzer (2006) "Job hopping in the Silicon Valley: The micro-foundations of an industry cluster", *Review of Economics and Statistics* 88, 472-481.
- Florida, R. (2003), *Rise of the Creative Class: And How It's Transforming Work, Leisure, Community and Everyday Life*, New York: Basic Books.
- Fujita, M. (2011), "Thünen and the New Economic Geography", Research Institute of Economy, Trade and Industry Discussion Papers 11074.
- Gibbons, S., H.G. Overman and P. Pelkonen (2010), "Wage Disparities: Places or People?", SERC Discussion Paper 60.
- Glaeser, E.L. (2011), *The Triumph of the City, How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier and Happier*, New York, NY: Penguin US.
- Glaeser, E.L., H. Kallal, J. Scheinkman, and A. Schleifer (1992), "Growth in cities", *Journal of Political Economy*, 100, 889-927.
- Glaeser, E.L., and D.C. Mare (2001), "Cities and Skills", *Journal of Labor Economics* 19(2), 316-342.
- Glaeser, E.L. and W.R. Kerr (2009), "Local industrial conditions and entrepreneurship: How much of the spatial distribution can we explain?", *Journal of Economics and Management Strategy* 18, 623-663.
- Glaeser, E.L. and A. Saiz (2004), "The rise of the skilled city", *Brookings-Wharton Papers on Urban Affairs*, 47-105.
- Hall, B., A. Jaffe and M. Trajtenberg (2000), "Market Value and Patent Citations: A First Look," NBER Working Paper 7441.

- Hall, B., A. Jaffe and M. Trajtenberg (2001), "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tool", NBER Working Paper 8498.
- Helsley, R.W. and W.C. Strange (2001), "Innovation and input sharing", *Journal of Urban Economics* 51, 24-45.
- Helsley, R.W. and W.C. Strange (2013), "Coagglomeration", working paper, mimeo, The Rotman School of Management.
- Henderson, J.V., A. Kuncoro and M. Turner (1995), "Industrial development in cities", *Journal of Political Economy*, 103, 1067-1085.
- Holmes, T.J. (1999), "Localization of industry and vertical disintegration", *Review of Economics and Statistics* 81, 314-25.
- Jacobs, J. (1969), *The Economy of Cities*. New York, NY: Vintage.
- Jaffe, A.B., M. Trajtenberg and R. Henderson (1993), "Geographic localization of knowledge spillovers as evidenced by patent citations", *Quarterly Journal of Economics*, 108(3), 577-598.
- Jofre-Monseny, J., R. Marín-López and E. Viladecans-Marsal (2011), "The mechanisms of agglomeration: Evidence from the effect of inter-industry relations on the location of new firms", *Journal of Urban Economics* 70, 61-74.
- Krugman, P. (1991), "Increasing returns and economic geography", *Journal of Political Economy* 99, 483-499.
- Landes, D.S. (1969), *The Unbound Prometheus, Technological Change and Industrial Development in Western Europe from 1750 to the Present*, Cambridge, UK: Cambridge University Press.
- Lin, J. (2012), "Technological adaptation, cities, and new work," *Review of Economics and Statistics* 93, 554-574.
- Marshall, A. (1890), *Principles of Economics*, MacMillan: London.
- OECD (1997): "Revision of High-Technology Sector and Product Classification", STI Working Paper 1997/2.
- Office for National Statistics, *Business Structure Database, 1997-2011: Secure Access* [computer file]. 3rd Edition. Colchester, Essex: UK Data Archive [distributor], October 2012. SN: 6697.
- Office for National Statistics. Social Survey Division, and Northern Ireland Statistics and Research Agency. Central Survey Unit, *Quarterly Labour Force Survey, 1992-2011: Secure Access* [computer file]. 3rd Edition. Colchester, Essex: UK Data Archive [distributor], May 2013. SN: 6727.
- Overman, H.G. and D. Puga (2010), "Labor pooling as a source of agglomeration: An empirical investigation", in E. L. Glaeser (Ed.), *Agglomeration Economics - NBER Collections*, Chicago, IL: The University of Chicago Press, 133-150.
- Porter, M. (1990), *The Competitive Advantage of Nations*, New York, NY: The Free Press.
- Rauch, J. (1993), "Productivity gains from geographic concentration of human capital: Evidence from the cities", *Journal of Urban Economics* 34, 380-400.
- Rosenthal, S. and W.C. Strange (2001), "The determinants of agglomeration", *Journal of Urban Economics* 50, 191-229.
- Rosenthal, S. and W.C. Strange (2003), "Geography, industrial organization, and agglomeration", *Review of Economics and Statistics* 85, 377-393.
- Rosenthal, S. and W.C. Strange (2004), "Evidence on the Nature and Sources of Agglomeration Economies", in Henderson, J.V. and Thisse, J.-F. (Eds.), *Handbook of Urban and Regional Economics, Volume 4*, Amsterdam: Elsevier, 2119-2172.
- Rosenthal, S., and W.C. Strange (2008), "The attenuation of human capital spillovers", *Journal of Urban Economics* 64, 373-389.

- Rosenthal, S. and W.C. Strange (2010), “Big Establishment/Small Effects: Agglomeration, Industrial Organization, and Entrepreneurship”, Chapter 10 in E. L. Glaeser (Ed.), *Agglomeration Economics - NBER Collections*, Chicago, IL: The University of Chicago Press, 277-302.
- Saxenian, A.-L. (1994), *Regional Advantage: Cultural Competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.
- Silverman, B.S. (2002), *Technological Resources and the Logic of Corporate Diversification*, London, UK: Routledge.
- Strange, W.C., W. Hejazi and J. Tang (2006), “The Uncertain City: Competitive Instability, Skills, Innovation, and the Strategy of Agglomeration”, *Journal of Urban Economics*, 59(3), 331-351.
- Vernon, R. (1960), *Metropolis 1985*, Cambridge, MA: Harvard University Press.
- von Thünen, Johann Heinrich (1966), *Isolated state: an English edition of Der isolierte Staat*. New York, NY: Pergamon Press.

Table 1: Descriptive statistics

| | <i>Mean</i> | <i>Std. Dev.</i> | <i>Min</i> | <i>Max</i> |
|---|-------------|------------------|------------|------------|
| <i>Coagglomeration measures and Marshallian forces</i> | | | | |
| TTWA total employment coagglomeration (γ^C) | 0.000 | 0.005 | -0.028 | 0.107 |
| Labor pooling (correlation) | 0.237 | 0.188 | -0.022 | 0.968 |
| Input-output sharing (maximum) | 0.009 | 0.033 | 0.000 | 0.547 |
| Input sharing (maximum) | 0.007 | 0.029 | 0.000 | 0.547 |
| Output sharing (maximum) | 0.005 | 0.021 | 0.000 | 0.546 |
| Knowledge spillovers – prob. mapping, industry of manufacture (IOM, maximum of inward/outward citation) | 0.016 | 0.037 | 0.000 | 0.413 |
| Knowledge spillovers – prob. mapping, sector of use (SOU, maximum of inward/outward citation) | 0.012 | 0.026 | 0.000 | 0.540 |
| <i>Additional Controls</i> | | | | |
| Energy dissimilarity index | 0.013 | 0.016 | 0.000 | 0.097 |
| Water dissimilarity index | 0.001 | 0.001 | 0.000 | 0.006 |
| Transport dissimilarity index | 0.014 | 0.018 | 0.000 | 0.084 |
| Natural Resources dissimilarity index | 0.041 | 0.076 | 0.000 | 0.369 |
| Services dissimilarity index | 0.018 | 0.016 | 0.000 | 0.082 |

Note: All pairwise combinations of manufacturing SIC1992 3-digit industries are included except Manufacture of tobacco (SIC160). In addition, we combined Manufacture of leather clothes (SIC181) and Dressing and dyeing of fur (SIC183) with Manufacture of wearing apparel (SIC182); Manufacture of coke oven products (SIC231) and Processing of nuclear fuel (SIC233) with Refined petroleum products (SIC232). We also combined the following sectors: Manufacture of vegetable and animal oils and fats (SIC154) with Manufacture of other food products (SIC158); Manufacture of man-made fibers (SIC247) with Manufacture of other chemical products (SIC246); Manufacture of cement, lime and plaster (SIC265) with Manufacture of articles of concrete, plaster and cement (SIC266); Reproduction of recorded media (SIC223) with Printing (SIC222). Our final sample consists of 94 manufacturing 3-digit sectors for a total of 4,371 unique pairwise correlations a year for twelve years (1997-2008). The complete dataset contains 52,452 observations. Labor correlation indices are computed from the UK Labour Force Survey 1995-1999. Input-Output measures are calculated ONS UK Input-Output Tables for 1995-1999. Knowledge spillover measures are calculated using the UK data retrieved from the KITES-PATSTAT dataset kindly made available to us by Bocconi University. Cited patents sampled for the years 1978 to 1997. Citing patents sampled for the years 1981 to 2000. Additional control measures are calculated using the UK Input-Output tables for 1995-1999.

Table 2: Characterizing sectoral breakdown

| Sector pair is: | N. of Obs./Pairs | Mean γ^C | Mean LP | Mean IO | Mean KS (IOM) | Mean of Cut-off Variable | Mean E-G Localization Index |
|------------------|------------------|-----------------|---------|---------|---------------|--------------------------|-----------------------------|
| New | 12972/1081 | 0.000 | 0.203 | 0.006 | 0.007 | 1972 | 0.039 |
| Mixed | 26508/2209 | 0.000 | 0.234 | 0.008 | 0.016 | 1958 | 0.027 |
| Old | 12972/1081 | 0.000 | 0.277 | 0.014 | 0.022 | 1945 | 0.014 |
| Dynamic | 12972/1081 | 0.000 | 0.219 | 0.007 | 0.012 | 0.131 | 0.032 |
| Mixed | 26508/2209 | 0.000 | 0.234 | 0.008 | 0.015 | 0.106 | 0.027 |
| Steady | 12972/1081 | 0.000 | 0.262 | 0.012 | 0.020 | 0.082 | 0.021 |
| High tech. | 7140/595 | 0.000 | 0.412 | 0.014 | 0.029 | -- | 0.009 |
| Mix tech. | 24780/2065 | 0.000 | 0.221 | 0.007 | 0.016 | -- | 0.023 |
| Low tech. | 20532/1711 | 0.000 | 0.195 | 0.010 | 0.011 | -- | 0.037 |
| High education | 12972/1081 | 0.000 | 0.328 | 0.009 | 0.022 | 0.148 | 0.015 |
| Mix education | 26508/2209 | 0.000 | 0.219 | 0.008 | 0.014 | 0.097 | 0.027 |
| Low education | 12972/1081 | 0.000 | 0.184 | 0.011 | 0.012 | 0.047 | 0.038 |
| Small entrants | 12972/1081 | 0.000 | 0.248 | 0.005 | 0.016 | 5.676 | 0.018 |
| Mix entrants | 26508/2209 | 0.000 | 0.233 | 0.010 | 0.014 | 10.08 | 0.027 |
| Large entrants | 12972/1081 | 0.000 | 0.234 | 0.012 | 0.018 | 14.48 | 0.035 |
| Small incumbents | 12972/1081 | 0.000 | 0.240 | 0.005 | 0.019 | 10.54 | 0.017 |
| Mix incumbents | 26508/2209 | 0.000 | 0.233 | 0.009 | 0.014 | 22.11 | 0.027 |
| Large entrants | 12972/1081 | 0.000 | 0.243 | 0.012 | 0.016 | 33.69 | 0.036 |

Note: Number of pairs refers to unique (non-repeated) sector combinations. High-tech and low-tech industries are categorized according to the OECD classification (1997). High-education and low-education industries are classified according to the share of college graduates above and below the median across all years (the median is 0.0783). Education level calculated using the UK LFS 1995-1999 data. New/old industry pairs consist of industries where the first year of opening is above/below the median across all years (the median is 1967). Dynamic/steady industry pairs consist of industries where the share of entrants is above/below the median across all years (the median is 0.100). Small/large entrants refer to industry pairs where the average size of entrants is below/above the median across all years (the median is 8.59). Small/large incumbents refer to industry pairs where the average size of incumbents is below/above the median across all years (the median is 18.95). Mixed pairs consist of one old/big entrants/big incumbents/steady industry and one new/small entrants/small incumbents/dynamic industry. Variables in the penultimate column refer to the cut-off variables (first year of opening, entry share, size of entrants and size of incumbents) averaged across the sector pairs. Mean Ellison-Glaeser localization index across all sectors: 0.027 (std. dev.: 0.048).

Table 3: The relationship between coagglomeration γ^C and Marshallian forces

| | (1) | (2) | (3) | (4) |
|------------------------------------|---------------------|---------------------|---------------------|---------------------|
| Specification details: | OLS – Univar. | OLS – Univar. | OLS – Multivar. | OLS – Multivar. |
| Labor pooling (LP) | 0.191 (0.018)*** | 0.198 (0.018)*** | 0.156 (0.019)*** | 0.165 (0.020)*** |
| Input-output sharing (IO) | 0.138 (0.026)*** | 0.137 (0.027)*** | 0.083 (0.025)*** | 0.082 (0.025)*** |
| Knowledge spill. – IOM (KS) | 0.106 (0.015)*** | 0.099 (0.014)*** | 0.031 (0.013)** | 0.024 (0.013)* |
| <i>Resource use diss. Controls</i> | <i>No</i> | <i>Yes</i> | <i>No</i> | <i>Yes</i> |

Note: See note to Table 1 for details on variable definitions. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses. All regressions consider the period 1997-2008.

Table 4: Instrumental variable regressions

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|------------------------------------|---------------------|---------------------|---------------------|---------------------|------------------------------|------------------------------|------------------------------|
| | OLS – Univar. | IV – Univar. | OLS – Multivar. | IV – Multivar. | IV – Multivar. (no KS) | IV – Multivar. (no LP) | IV – Multivar. (no IO) |
| Labor pooling (LP) | 0.161 (0.017)*** | 0.113 (0.020)*** | 0.133 (0.018)*** | 0.116 (0.032)*** | 0.100 (0.024)*** | -- | 0.116 (0.030)*** |
| Output sharing (IO) | 0.105 (0.017)*** | 0.127 (0.026)*** | 0.061 (0.016)*** | 0.083 (0.024)*** | 0.082 (0.024)*** | 0.121 (0.028)*** | -- |
| Knowledge spillovers – IOM (KS) | 0.078 (0.013)*** | 0.088 (0.016)*** | 0.033 (0.013)** | -0.021 (0.023) | -- | 0.031 (0.019)* | 0.017 (0.022) |
| <i>First-stage statistics</i> | | | | | | | |
| <i>t-stat on LP</i> | -- | 19.93 | -- | 16.91 | 18.78 | -- | 17.04 |
| <i>t-stat on IO</i> | -- | 6.39 | -- | 6.22 | 6.32 | 5.95 | -- |
| <i>t-stat on KS</i> | -- | 6.83 | -- | 5.94 | -- | 6.45 | 6.56 |

Note: See note to Table 1 for details on variable definitions. All regressions control for dissimilarity in use of resources. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses. Sample excludes 3-digit SIC sectors within the same 2-digit SIC sector and the following sectors which were aggregated in the data construction: Manufacture of wearing apparel and accessories (SIC182); Manufacture of refined petroleum products (SIC232); Manufacture of other chemical products (SIC246); Manufacture of other food products (SIC158); Manufacture of articles of concrete, plaster and cement (SIC266); and Printing and service activities related to printing (SIC222). Number of observations 43,644 (3,637 industry pairs). Instrumental variable regressions use labor correlation, input-output linkages and patent citations flows calculated using US data. See Data Construction Appendix for more details. Cells in Column (1) and (2) come from separate regressions. Cells in Columns (3)-(7) come from regressions that simultaneously enter Marshallian forces as detailed in the headings. Knowledge spillovers measure is based on probabilistic mapping – Industry of manufacturing (Knowledge Spillovers – IOM).

Table 5: The relationship between coagglomeration γ^C , Marshallian forces and non-Marshallian mechanisms

| | (1) | (2) | (3) | (4) | (5) | (6) |
|-----------------------------|---------------------|---------------------|----------------------|----------------------|---------------------|---------------------|
| Additional Control details: | Year of Opening | Entry share | Tech. | Education | Size of entrants | Size of Incumbents |
| Labor pooling (LP) | 0.166 (0.020)*** | 0.165 (0.020)*** | 0.188 (0.022)*** | 0.176 (0.021)*** | 0.164 (0.020)*** | 0.161 (0.019)*** |
| Input-output sharing (IO) | 0.085 (0.026)*** | 0.082 (0.025)*** | 0.075 (0.025)*** | 0.079 (0.025)*** | 0.080 (0.025)*** | 0.081 (0.025)*** |
| Knowledge spill. – IOM (KS) | 0.028 (0.013)** | 0.024 (0.014)* | 0.027 (0.014)** | 0.029 (0.014)** | 0.023 (0.014)* | 0.024 (0.014)* |
| Year of opening | 0.046 (0.014)*** | | | | | |
| Entry share | | 0.003 (0.012) | | | | |
| High tech | | | -0.227 (0.055)*** | | | |
| Mix tech | | | -0.087 (0.035)*** | | | |
| Share college graduates | | | | -0.059 (0.018)*** | | |
| Size of entrants | | | | | 0.028 (0.014)* | |
| Size of incumbents | | | | | | 0.032 (0.016)* |

Note: See note to Table 1 for details on variable definitions. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses. All regressions consider the period 1997-2008, control for dissimilarity in use of resources and include all Marshallian forces simultaneously. Estimates obtained using OLS. Year of opening refers to the first year of openings across all plants in a given industry averaged across the sector pairs. Share of entry refers to the fraction of new firms averaged across sector pairs. High-tech and low-tech industries are categorized according to the OECD classification (1997). The mixed-tech industry pairs consist of one high-tech industry and one low-tech industry. Share of college graduates refers to the average share across the sector pairs. Education level calculated using the UK LFS 1995-1999 data. Size of entrants and size of incumbents refer to the number of employees of new and existing plants averaged across the sector pairs.

Table 6: The heterogeneous relationship between coagglomeration γ^C and Marshallian forces –
Adaptation

| | (1) | (2) | (3) |
|---------------------------------|---------------------|---------------------|---------------------|
| | New | Mixed | Old |
| Labor pooling (LP) | 0.310 (0.058)*** | 0.153 (0.026)*** | 0.081 (0.021)*** |
| Input-output sharing (IO) | 0.270 (0.083)*** | 0.049 (0.030) | 0.041 (0.018)** |
| Knowledge spillovers – IOM (KS) | 0.236 (0.121)** | 0.040 (0.017)** | 0.026 (0.019) |
| N of. Observations/Pairs | 12972/1081 | 26508/2209 | 12972/1081 |
| | Dynamic | Mixed | Steady |
| Labor pooling (LP) | 0.181 (0.059)*** | 0.144 (0.026)*** | 0.180 (0.028)*** |
| Input-output sharing (IO) | 0.113 (0.095) | 0.103 (0.031)*** | 0.052 (0.021)** |
| Knowledge spillovers – IOM (KS) | 0.181 (0.074)** | 0.033 (0.018)* | -0.020 (0.015) |
| N of. Observations/Pairs | 12972/1081 | 26508/2209 | 12972/1081 |

Note: See note to Table 1 and Table 2 for details on variable definitions. Number of pairs refers to unique (non-repeated) sector combinations. All regressions control for dissimilarity in use of resources. Regressions further control for the following variables averaged across the sector pairs: first year of opening (top panel); entry share (bottom panel). Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses.

Table 7: The heterogeneous relationship between coagglomeration γ^C and Marshallian forces –
Technology and education

| | (1) | (2) | (3) |
|---------------------------------|---------------------|---------------------|---------------------|
| | High-tech | Mixed-tech | Low-tech |
| Labor pooling (LP) | 0.046 (0.017)*** | 0.110 (0.019)*** | 0.332 (0.049)*** |
| Input-output sharing (IO) | 0.020 (0.012) | 0.064 (0.020)*** | 0.091 (0.045)** |
| Knowledge spillovers – IOM (KS) | 0.053 (0.024)** | 0.031 (0.016)* | 0.039 (0.041) |
| N of. Observations/Pairs | 7140/595 | 24780/2065 | 20532/1711 |
| | High-edu. | Mixed-edu. | Low-edu. |
| Labor pooling (LP) | 0.046 (0.023)* | 0.167 (0.031)*** | 0.391 (0.061)*** |
| Input-output sharing (IO) | 0.007 (0.013) | 0.066 (0.029)** | 0.123 (0.057)** |
| Knowledge spillovers – IOM (KS) | 0.048 (0.020)** | 0.050 (0.021)** | 0.030 (0.040) |
| N of. Observations/Pairs | 12972/1081 | 26508/2209 | 12972/1081 |

Note: See note to Table 1 and Table 2 for details on variable definitions. Number of pairs refers to unique (non-repeated) sector combinations. All regressions control for dissimilarity in use of resources. Regressions in the bottom panel further control for the share of college graduates averaged across the sector pairs. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses.

Table 8: The heterogeneous relation between coagglomeration γ^C and Marshallian forces – Organization

| | (1) | (2) | (3) |
|----------------------------|---------------------|---------------------|---------------------|
| | Small entrants | Mixed entrants | Large entrants |
| Labor pooling | 0.113 (0.027)*** | 0.181 (0.027)*** | 0.223 (0.056)*** |
| Input-output sharing | 0.193 (0.112)* | 0.068 (0.034)** | 0.082 (0.039)** |
| Knowledge spillovers – IOM | 0.051 (0.030)* | 0.028 (0.016)* | -0.022 (0.030) |
| N of. Observations/Pairs | 12972/1081 | 26508/2209 | 12972/1081 |
| | Small incumbents | Mixed incumbents | Large incumbents |
| Labor pooling | 0.065 (0.026)** | 0.223 (0.031)*** | 0.149 (0.047)*** |
| Input-output sharing | 0.159 (0.076)** | 0.071 (0.035)** | 0.068 (0.039)* |
| Knowledge spillovers – IOM | 0.034 (0.024) | 0.020 (0.018) | 0.040 (0.031) |
| N of. Observations/Pairs | 12972/1081 | 26508/2209 | 12972/1081 |


Note: See note to Table 1 and Table 2 for details on variable definitions. Number of pairs refers to unique (non-repeated) sector combinations. All regressions control for dissimilarity in use of resources. Regressions further control for the following variables averaged across the sector pairs: size of entrants (top panel); size of incumbents (bottom panel). Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses.

Table 9: The heterogeneous relation between coagglomeration γ^C and Marshallian forces – Single-industry models

| Sector Description | SIC Code | Sector characteristics | Effect of LP | Effect of IO | Effect of KS | E-G Index |
|--|----------|---|---------------------|--------------------|---------------------|-----------|
| Mfg. of cutlery, tools & general hardware | 286 | New, steady, low tech, low education, small entrants, small incumbents | 0.345 (0.131)** | 0.719 (0.310)** | 0.003 (0.091) | 0.034 |
| Mfg., preparation, weaving & finishing of textiles | 171-177 | New, dynamic, low tech, low education, mix entrants, mix incumbents | 0.711 (0.200)*** | 0.070 (0.066) | 0.171 (0.066)*** | 0.077 |
| Mfg. of office machinery & computers | 300 | Old, dynamic, high tech, high education, small entrants, small incumbents | -0.042 (0.068) | 0.033 (0.020)* | 0.378 (0.072)*** | 0.009 |
| Mfg. of cars, engines & bodies for vehicles | 341 | New, dynamic, high tech, high education, small entrants, large incumbents | 0.439 (0.090)*** | -0.003 (0.101) | 0.176 (0.102)* | 0.030 |


Note: Sectors are classified based on the taxonomy presented in Table 2. ‘Mixed’ categories refer to group of sectors where a roughly equal number of 3-digit industries fall in either category of the breakdown (e.g. ‘mixed entrants’ means that a roughly equal number of sectors have big and small entrants). Regression coefficients come from single-industry regressions that exploit the variation in the coagglomeration of the industry in question with other industries (mutually exclusive pairs only) over twelve years. N. of observations: 1116 for SIC286, SIC300 and SIC341; 7560 for SIC171-177. Standard errors clustered at the industry pairs. Average Ellison-Glaeser index across all sectors: 0.027 (std. dev.: 0.048).

Figure 1a: The effect of Marshallian forces at different quantiles of γ^C : Labor pooling




Note: See note to Table 1 for details on variable definitions. Variables are transformed to have unit standard deviation for interpretation. The figure plots regression coefficients and 95% confidence intervals from quantile regressions that simultaneously include all three Marshallian forces. Confidence intervals from bootstrapped standard errors clustered on industry pairs. All regressions control for dissimilarity in use of resources.

Figure 1b: The effect of Marshallian forces at different quantiles of γ^C : Input-output sharing



Note: See note to Table 1 for details on variable definitions. Variables are transformed to have unit standard deviation for interpretation. The figure plots regression coefficients and 95% confidence intervals from quantile regressions that simultaneously include all three Marshallian forces. Confidence intervals from bootstrapped standard errors clustered on industry pairs. All regressions control for dissimilarity in use of resources.

Figure 1c: The effect of Marshallian forces at different quantiles of γ^C : Knowledge Spillovers – IOM



Note: See note to Table 1 for details on variable definitions. Variables are transformed to have unit standard deviation for interpretation. The figure plots regression coefficients and 95% confidence intervals from quantile regressions that simultaneously include all three Marshallian forces. Confidence intervals from bootstrapped standard errors clustered on industry pairs. All regressions control for dissimilarity in use of resources. Knowledge spillovers measure is based on probabilistic mapping – Industry of manufacturing (Knowledge Spillovers – IOM).

Data Appendix

The Business Structure Database (BSD)

Our measures of coagglomeration of UK manufacturing sectors are constructed aggregating micro-level data from the Business Structure Database (BSD) covering the period 1997 to 2008. The data is an annual snapshot (taken in April at the closing of the fiscal year) of the Inter-Departmental Business Register (IDBR), which consists of constantly-updated administrative data collected for taxation purposes. Any business liable for value-added taxation (VAT) and/or with at least one employee registered for tax collection appears on the IDBR. Estimates produced by the Office for National Statistics (ONS) in 2004 show that the businesses listed on the IDBR account for almost 99 per cent of economic activity in the UK.

The data are structured into enterprises and local units. An enterprise is the overall business organization. The local unit can be thought of as a plant or establishment. In the majority of cases (70 per cent), enterprises only have one local unit, while the remaining 30 per cent of the cases represent enterprises with multiple local units. In our work, we make use of data at the local unit level including plants belonging both to single- and multi-plant enterprises and located in England, Wales and Scotland. We neglect Northern Ireland because of poor data coverage.

The initial raw data includes approximately three million local units every year. However, before using the data for our analysis, we carry out a series of checks and drop a number of units. In particular, we investigate the consistency of opening and closing dates of BSD units with their actual existence in the dataset and drop a number of anomalous cases where we identify establishments opening/closing in a specific year, disappearing/reappearing in a subsequent year only to open/close again in a subsequent wave. Stated differently, we only count firms' birth and death once. Secondly, we check the consistency of units' postcodes and sectors of activity over the years, and drop cases with missing or anomalous information.¹ For example, when we observe two or more plants operating in the same 3-digit industry, sharing the same postcode and being part of the same enterprise, we believe this being a reporting error and drop them. Similarly, we observe a non-trivial number of same-postcode same-three-digit industry combinations representing anomalous concentration of identical activities at a single address. We believe this is another coding error and drop the plants that belong to the top 5% of distribution of the number of plants sharing same three-digit industry and the same postcode. Finally, we drop active units with zero employment since this figure includes the owners/managers of the establishment, so it cannot be zero for an active unit, as well

¹ A UK postcode usually corresponds to a very limited number of addresses or a single large delivery point. While it might not always be a geographically accurate description of where a property is located, it is generally a good approximation. For instance, a building which contains several flats or businesses, but only one external door will only have the external door listed as a delivery point. This example shows that UK postcodes are geographically accurate up to the level of a front door in a particular street.

as units with an unusually large size (i.e., total employment above the 99th percentile of the distribution for each three-digit industry sector). After applying these restrictions, our dataset still comprises of more than two million plants annually over 12 years (1997-2008).

In terms of industrial classification, we focus on manufacturing and adopt the three-digit Standard Industry Classification (SIC) 1992. We also apply a number of restrictions and re-combine a number of sectors to avoid having a limited or erratic evolution in the number of plants and employment during the sample period. Specifically, we exclude Tobacco (SIC160) because the number of plants in this sector tends to be small throughout the sample period (e.g. 43 in 1997). In addition, we combine Leather (SIC181) and Fur Clothes (SIC183) with Other Wearing Apparel (SIC182) to avoid small sample size problems in SIC181 and SIC183. For similar reasons, we also combine the following industries: Manufacture of Vegetal and Animal Oils and Fats (SIC154) with Other Food Products (SIC158); Reproduction of Recorded Media (SIC223) with Printing (SIC222); Coke Oven Products (SIC231) and Processing of Nuclear Fuel (SIC233) with Refined Petroleum Products (SIC232); Man-made Fabrics (SIC246) with Other Chemical Products (SIC247); Articles of Concrete, Plaster and Cement (SIC266) with Manufacture of Cement, Lime and Plaster (SIC265). Our final sample consists of 94 manufacturing 3-digit sectors for a total of 4,371 unique pairwise correlations a year for twelve years (1997-2008). The complete dataset contains 52,452 observations.

In terms of geography, our unit of aggregation is the Travel-to-Work Area (TTWA). These are entities constructed to guarantee that at least 75% of the resident population works in the area and that 75% of the people working in the area are resident there. TTWAs were devised by UK government agencies to delineate areas that can be considered as self-contained labor markets and economically relevant aggregates. In 2007, there were 243 TTWAs within the United Kingdom. In our analysis, we exclude Northern Ireland and only consider urban TTWAs. In splitting urban and rural TTWAs, we follow Gibbons et al. (2010) and re-aggregate some areas so that our final partition includes 158 local economic areas, of which 84 are single urban TTWAs with population in excess of 100,000 residents. This definition is slightly different from the one in Gibbons et al. (2010) since we aggregate the individual areas of Clacton, Colchester, Lincoln, Grantham, Torquay, and Paignton-Totes into the following urban TTWAs: (1) Clacton & Colchester; (2) Lincoln & Grantham; (3) Torquay & Paignton-Totes. Even before the aggregation, the areas of Colchester, Lincoln and Torquay each had a population above the 100,000 threshold. In addition, we reclassify the areas of Chesterfield, Eastbourne, and Lancaster-Morecambe as urban TTWAs because each has a population approaching our threshold.

UK Labour Force Survey (LFS)

The UK Labour Force Survey (LFS) is a quarterly representative survey of households living at private addresses in the UK and is conducted by the Office for National Statistics (ONS) to collect information

about individuals' labor market experiences. In our analysis, we use the years between 1995 and 1999 which allow for a consistent coding of the industrial and occupational classification of workers' jobs.

Each quarter of the LFS contains between 64,000 (earlier years) and 52,000 (later years) households, equivalent to about 120,000-150,000 individuals. In our analysis, we focus on 16-59 aged women and 16-64 aged men, and on individuals either working as employees or as self-employed. Excluding self-employed individuals does not affect our analysis.

In order to assign each individual to a TTWA, we retain workers living in England, Scotland and Wales (LFS data for Northern Ireland has poor coverage), and with a valid geographical identifier, namely the ward of residence (roughly equivalent to a US census tract). Additionally, we select individuals with non-missing information on: (i) gender and age; (ii) educational qualifications; (iii) industry and occupation. We exclude people working for the armed forces.

These restrictions leave us with a set of approximately 200,000 individuals each year for a total of 1.03m, of which 820,000 and 210,000 live in urban and rural areas, respectively. Next, we select individuals living in urban areas and working in manufacturing only. The final sample consists of about 35,000 workers a year for a total of 166,000 individuals. We use the UK Standard Industrial Classification (SIC) 1992 and the UK Standard Occupational Classification (SOC) 1990 at the three-digit level for these individuals' jobs to construct a proxy for the extent of labor pooling occurring between manufacturing sectors.

The EPO-CESPRI Dataset

The main data source for our analysis of patent citation flows is the EPO-CESPRI data provided by Bocconi University. This database provides cleaned and consistently coded information extracted from the European Patent Office (EPO) data for the period 1977 and 2009. Approximately 144,000 patents were filed by 160,000 UK inventors (multiple-inventors can be recorded for each patent). These generate a stream of more than 77,000 citations of UK patents over the observed time-window.

In order to construct our knowledge spillover measures we impose a number of restrictions. First, we exclude self-citations from the same inventor or the applying company at which he/she is based. Second, we exclude citing patents filed after 2000 and before 1981, and cited patents filed after 1997. The aim of these restrictions is twofold: (a) we want to guarantee that on average citing patents are at least three years older than cited ones; (b) we want to guarantee that our knowledge-spillover measures are constructed for the initial years of our sample (i.e. up to 2000) so that they are measured at a similar time as the labor-pooling and input-output sharing metrics. Expanding the sample to include all available years does not affect the results.

Finally, it should be noted that while the US Patent Office (USPTO) requires patent applicants to declare all relevant references and citations, the EPO does not apply this rule and all citations come directly from the patent examiners. As a result, the average number of EPO patent citations is much smaller than the

corresponding figure for USPTO patents, and EPO numbers do not suffer from UPSTO-type “citation inflation” (see Hall et al. 2000). According to Breschi and Lissoni (2004), USPTO patents cited approximately 13 other patents and received on average 10.2 citations. The corresponding numbers for EPO patents are much lower at 4 and 2.8, respectively.

The OECD Technology Classification

Based on the intensity of both direct R&D (i.e. R&D expenditure) and indirect R&D (i.e. embodied technology flows) in the output of manufacturing sectors across 10 OECD countries over the period 1980 to 1996, the OECD classifies as high tech or medium-high tech the following manufacturing industries: SIC241 “Manufacturing of basic chemicals” to SIC246 “Manufacturing of other chemicals & man-made fibres”; SIC291 “Manufacturing of other machinery for production/use of mechanical power N.E.C.” to SIC297 “Manufacturing of domestic appliances”; SIC300 “Manufacturing of office machinery & computers”; SIC311 “Manufacturing of electric motors, generators & transformers” to SIC316 “Manufacturing of electrical equipment N.E.C.”; SIC231 “Manufacturing of electronic valves, tubes & electronic components” to SIC323 “Manufacturing of TV/radio receivers & sound/video recording/reproducing”; SIC331 “Manufacturing of medical, surgical & orthopedic equipment” to SIC335 “Manufacture of watches & clocks; SIC341 “Manufacturing of motor vehicles” to SIC343 “Manufacturing of parts & accessories for vehicles/engines”; and SIC352 “Manufacturing of railway/tramway locomotives/rolling stock” to SIC355 “Manufacturing of other transport equipment N.E.C.”. See Table W1 and OECD (1997) for more details.

US Data for Instrumental Variables

In order to address potential endogeneity and reverse causality issues, we follow Ellison et al. (2010) and instrument our UK-based proxies for the three Marshallian forces using (almost) identical measures obtained from US data.

Starting with labor pooling, we create a measure of the similarity in the occupational inputs of two industries using the National Industrial-Occupation Employment Matrix (NIOEM) published by the Bureau of Labor Statistics. Following the approach we have taken for the UK data, we construct the shares of different types of workers used in each manufacturing sector and then correlate the percentage of different types of occupations across industry pairs to obtain a proxy for labor sharing. In order to link this proxy to our data, we map US NIOEM industry codes to UK SIC codes. Since the US manufacturing classification is less detailed than the one that we adopt (79 vs. 94 sectors), we attribute the same US industry-occupation shares to multiple UK sectors. Note also that we construct the US labor correlation measure using all available data spanning the period 1983 to 1998. Restricting the calculations of this instrument to the period 1995-1998 does not affect our IV results.

We construct an instrument for input-output sharing following a similar approach. To begin with, we use the 1987 Benchmark Input-Output Accounts published by the Bureau of Economic Analysis (BEA) to measure the flows of intermediate inputs exchanged between US industries at the same level of aggregation as used in Ellison et al. (2010) and map these values from the 140 US manufacturing sectors to the 94 UK industries. In our regression analysis we focus on the maximum between the inputs and outputs that two industries are sharing irrespective of the direction of the flow (given that our data treats industry pairs symmetrically and considers them only once). Consistently, we use US data to construct a proxy for the maximum of the input-output linkages between industries and use this as an instrument.

Finally, we construct our instrument for knowledge spillovers using the NBER Patent Data initially assembled by Hall et al. (2001). The data cover patents granted by the US patent and Trademark Office (USPTO) between 1975 and 1999 and record citation flows across patents. Following our main approach, we use a probabilistic mapping based on the industry of manufacture (IOM) to map technology to industrial classes and convert citations across US sectors to our UK classification based on 94 industries. Note that this instrument is different from the one adopted by Ellison et al. (2010) who used UK patents registered at the USPTO to instrument for US patents registered at the same office. Conversely, we use information coming from the USPTO about flows of citations among US patents to instrument for citations among UK patents registered at the European Patent Office (EPO). Part of the mechanical problems discussed by Ellison et al. (2010) in relation to this instrument is thus by-passed.

Appendix tables

Table A1: Fifteen most co-agglomerated industry pairs – based on coagglomeration measure γ^c in 1997

| Rank | Industry 1 | Industry 2 | Coagglomeration | 1st TTWA | 2nd TTWA | 3rd TTWA |
|------|--|---|-----------------|----------------|--------------|---------------|
| 1 | Ceramic goods other than construction (262) | Ceramic tiles and flags (263) | 0.105 | Stoke-on-Trent | Exeter | London |
| 2 | Knitted and crocheted fabrics (176) | Knitted and crocheted articles (177) | 0.086 | Leicester | Nottingham | Derby |
| 3 | Publishing (221) | Jewellery and related articles (362) | 0.054 | London | Birmingham | Sheffield |
| 4 | Spinning of textiles (171) | Textile weaving (172) | 0.054 | Bradford | Huddersfield | Leeds |
| 5 | Publishing (221) | Printing and reproduction of recorded media (222) | 0.037 | London | Manchester | Birmingham |
| 6 | Finishing of textiles (173) | Knitted and crocheted articles (177) | 0.037 | Leicester | Manchester | Nottingham |
| 7 | Finishing of textiles (173) | Knitted and crocheted fabrics (176) | 0.035 | Leicester | Nottingham | Manchester |
| 8 | Ceramic goods other than construction (262) | Construction products in baked clay (264) | 0.033 | Stoke-on-Trent | Crawley | Peterborough |
| 9 | Basic iron and steel and ferro-alloys (271) | Cutlery, tools and general hardware (286) | 0.033 | Sheffield | Birmingham | Wolverhampton |
| 10 | Basic iron and steel and ferro-alloys (271) | Other first processing of iron and steel (273) | 0.033 | Sheffield | Dudley | Swansea |
| 11 | Other first processing of iron and steel (273) | Forging, pressing, stamping and roll forming of metal (284) | 0.031 | Dudley | Birmingham | Sheffield |
| 12 | Tanning and dressing of leather (191) | Footwear (193) | 0.031 | Northampton | Hull | Glasgow |
| 13 | Knitted and crocheted articles (177) | Footwear (193) | 0.030 | Leicester | Northampton | Blackburn |
| 14 | Iron and Steel tubes (272) | Other first processing of iron and steel (273) | 0.029 | Dudley | Birmingham | Sheffield |
| 15 | Spinning of textiles (171) | Finishing of textiles (173) | 0.028 | Bradford | Manchester | Huddersfield |

Note: See note to Table 1 for details on variable definitions.

Table A2: Additional regressions of coagglomeration measure γ^C on Marshallian forces

| Dependent variable/ Timing is: | (1) | (2) | (3) | (4) |
|-----------------------------------|-------------------------|---------------------|---------------------|---------------------|
| | γ^C 1997-2008 | γ^C 1997 | γ^C 2002 | γ^C 2008 |
| Labor pooling (LP) | 0.166 (0.020)*** | 0.188 (0.024)*** | 0.166 (0.022)*** | 0.146 (0.020)*** |
| Input-output sharing (IO) | -- | 0.083 (0.026)*** | 0.099 (0.028)*** | 0.071 (0.025)*** |
| Knowledge spill. – IOM (KS) | 0.026 (0.013)* | 0.028 (0.014)** | 0.029 (0.015)* | 0.022 (0.013) |
| Input sharing | 0.057 (0.028)** | -- | -- | -- |
| Output sharing | 0.025 (0.031) | -- | -- | -- |

Note: See note to Table 1 for details on variable definitions and samples. All regressions include all Marshallian forces at the same time and control for dissimilarity in use of resources. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses.

Table A3: Main results with Knowledge Spillover – Sector of Use (SOM)

| | (1) | (2) | (3) | (4) | (5) |
|-----------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Specification | 1997-2008 | 1997-2008 | 1997 | 2002 | 2008 |
| details: | Mutivar. | Mutivar. | Mutivar. | Mutivar. | Mutivar. |
| Labor pooling (LP) | 0.152 (0.019)*** | 0.154 (0.019)*** | 0.170 (0.023)*** | 0.153 (0.021)*** | 0.138 (0.019)*** |
| Input-output sharing (IO) | 0.071 (0.024)*** | -- | 0.067 (0.024)*** | 0.088 (0.027)*** | 0.063 (0.024)*** |
| Knowledge spill. – SOU (KS) | 0.075 (0.023)*** | 0.077 (0.022)*** | 0.101 (0.026)*** | 0.080 (0.025)*** | 0.057 (0.022)*** |
| Input sharing | | 0.049 (0.027)* | | | |
| Output sharing | | 0.021 (0.030) | | | |

Note: See note to Table 1 for details on variable definitions. All regressions include all Marshallian forces at the same time and control for dissimilarity in use of resources. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses.

Table A4: Further robustness checks and extensions

| | (1) | (2) | (3) | (4) | (5) |
|----------------------------|------------------------------------|-----------------------------------|----------------------------------|---------------------------------|----------------------------------|
| | Staggered Marshallian forces | γ^C excluding London | Control for popul. density | Control for empl. density | Control for Herfind. Index |
| Labor pooling | 0.160 (0.019)*** | 0.139 (0.022)*** | 0.142 (0.018)*** | 0.158 (0.020)*** | 0.165 (0.020)*** |
| Output sharing | 0.081 (0.026)*** | 0.119 (0.033)*** | 0.091 (0.026)*** | 0.088 (0.026)*** | 0.083 (0.026)*** |
| Knowledge spillovers – IOM | 0.023 (0.014)* | 0.017 (0.014) | 0.026 (0.014)* | 0.023 (0.014)* | 0.026 (0.014)* |

Note: See note to Table 1 for details on variable definitions. All regressions control for dissimilarity in use of resources. Variables are standardized to have zero mean and unit standard deviation. Robust standard errors clustered on industry pairs are reported in parentheses. Knowledge spillovers measure is based on probabilistic mapping – Industry of manufacturing (knowledge spillovers – IOM). Column (1) considers γ^C for years from 2000 and Marshallian forces calculated up to 1999. Column (2) excludes London from the calculations of γ^C . Column (3) controls for the average (un-weighted) population density of the TTWAs in which the two sectors are operating, averaged within the pair. Column (4) controls for the average (un-weighted) employment density of the TTWAs in which the two sectors are operating, average within the pair. Column (5) controls for the Herfindahl index of the two sectors, averaged within the pair.