Groupe de Recherche en
Économie Théorique et Appliquée

Spreading Big Ideas? The effect of Top Inventing Companies on Local Inventors

Carlo MENON

OECD
PARIS

&

Spatial Economics Research Center

London

Cahiers du GREThA
n° 2014-11
July

GRETHA UMR CNRS 5113
Université Montesquieu Bordeaux IV
Avenue Léon Duguit - 33608 PESSAC - FRANCE
Tel : +33 (0)5.56.84.25.75 - Fax : +33 (0)5.56.84.86.47 - www.gretha.fr

La propagation des grandes idées? L'impact de l'activité de brevet des firmes leader sur les inventeurs locaux.

Résumé

Le document examine si l'activité de brevetage des entreprises les plus inventives affecte le nombre de brevets délivrés à d'autres inventeurs locaux dans la même zone métropolitaine aux États-Unis. La théorie économique prédit que les effets positifs des économies d'agglomération peuvent être contrebalancés par une pression à la hausse sur les salaires, laquelle est plus prononcée à court terme et au sein de chaque classe de technologie. L'analyse empirique exploite la structure en panel des données pour inclure différents effets fixes, et adopte une approche avec variable instrumentale pour démontrer la causalité. Les résultats montrent que l'effet est globalement positif, qu'il est plus marqué avec un décalage dans le temps, et qu'il ne se limite pas à la même catégorie de technologie, ce qui suggère que la diffusion de connaissances à la Jacob entre secteurs domine les autre sources d'économies d'agglomération intra-sectorielles, y compris les mécanismes de partage et d'appariement. Les implications pour la politique de développement local sont discutées.

Mots-clés : diffusion locale de connaissances, brevets, innovation.

Spreading Big Ideas? The effect of Top Inventing Companies on Local Inventors

Abstract

The paper investigates whether the patenting activity of the most inventive companies has any causal effect on the number of patents granted to other local inventors in the same metropolitan area in the United States. Economic theory predicts that positive agglomeration economies may be counterbalanced by upward pressure on wages, which are stronger within technological classes and in the short term. The empirical analysis exploits the panel dimension of the dataset to account for various fixed effects, and adopts an instrumental variable approach to prove causality. The results show that the effect is overall positive, it is stronger with a time lag, and it is not bounded within narrow technological categories, suggesting that Jacob-type knowledge spillovers across sectors tend to prevail over other source of agglomeration economies within sectors, including sharing and matching mechanisms. The implications for local development policy are discussed.

Keywords: localized knowledge spillovers, patents, innovation.

JEL: R10; O31

Reference to this paper: MENON Carlo (2014) Spreading Big Ideas? The effect of Top Inventing Companies on Local Inventors, *Cahiers du GREThA*, n°2014-11. Travail réalisé dans le cadre de la Chaire régionale d'Economie de l'Innovation portée par Francesco Lissoni (GREThA, UMR CNRS 5113, Université de Bordeaux).

http://ideas.repec.org/p/grt/wpegrt/2014-11.html.

GRETHA UMR CNRS 5113
Université Montesquieu Bordeaux IV
Avenue Léon Duguit - 33608 PESSAC - FRANCE
Tel : +33 (0)5.56.84.25.75 - Fax : +33 (0)5.56.84.86.47 - www.gretha.fr

1 Introduction[1]

In a recent paper, Greenstone, Hornbeck and Moretti (2010, henceforth GHM) discuss the importance of understanding whether large plant openings raise the productivity of local incumbents. They find a substantial effect: five years later the opening of a new large plant, the productivity of incumbent plants located in the same US county is 12% higher. The authors argue that the findings are extremely relevant both for economic theory, since they provide evidence on the mechanisms underlying the agglomeration of economic activities, and for local development policies, which often subsidize the location of large industrial investments.

At least since Marshall (1890), we know that knowledge spillovers are one of the leading mechanisms of agglomeration economics. This paper focuses primarily on this channel, expanding GHMs' analysis. Exploiting a rich patent database for the United States, I identify *Top Inventing Companies (*henceforth *TICs)* by proxying the company size with the stock of owned patents. I then assess whether the aggregate number of patents developed by inventors working for TICs has any causal effect on the number of patents granted to other companies (non-TIC) located in the same Metropolitan Statistical Area (MSA). A priori, the effect is not necessarily positive: in a general equilibrium framework, positive agglomeration externalities (knowledge spillovers, sharing and learning mechanisms) may be counterbalanced by upward pressures on nominal wages. Depending on the relative strength of the two mechanisms, the net effect could also be null or negative.

The results show that positive effects tend to prevail over negative ones at city level, and are stronger with a temporal lag. According to the regression results, a 10% increase in the number of TIC patents leads to an increase of about 2% in the number of non-TIC patents over the following 4-8 years. However, within narrowly defined technological sub-categories, where negative wage effects tend to be stronger due to a higher skill substitutability, the net effect is zero. One interpretation of these findings, also supported by additional tests based on citation data, is that that Jacob-type knowledge spillovers, which are not confined within technological categories and may require more time to produce effects, tend to prevail over other more sector-specific source of agglomeration economies, including sharing and matching mechanisms.

Using the NBER/USPTO patent database, I estimate a model where the number of non-TIC patents produced in a given city, time period, and technological category is a function of the

[1]**Acknowledgements**: I thank Kristian Behrens, Chiara Criscuolo, Alfonso Gambardella, Bronwyn Hall, Andrea Lamorgese, Henry Overman, Olmo Silva, John Van Reenen, two anonymous referees, and participants to the Patents Statistics for Decision Makers conference (Paris, November 2012) and to the Zvi Griliches Seminar on Economics of Innovation (Barcelona, July 2011) for useful comments on previous versions of the paper. The views expressed in the article are those of the author and do not involve the responsibility of the OECD and the Bank of Italy. Usual disclaimers apply.

number of TIC patents developed in the same city, period, and category. I exploit the panel dimension of the dataset to account for time, city and technology fixed effects. Causality is inferred with two-stage least squares (2SLS) estimations. The instrumental variable for TIC patents is based on the interaction of the historical presence of TICs in a given MSA with on contemporaneous variation in innovation activity of other plants of the same company. Consistently with economic theory, results show that positive effects prevail with a broad sectoral classification and a time lag, while negative effects are stronger in the short run and within narrow technological sectors. This evidence is suggestive of the relative importance of cross-industry knowledge spillovers as a source of agglomeration externalities.

This paper fills in two important gaps in the related economic literature. First, by providing empirical evidence on the effect of the interaction of heterogeneous actors in cities, it improves our understanding of urban agglomeration. As Duranton and Puga (2004) forcefully argued, heterogeneity of workers and firms is at the root of most of the theories of urban agglomeration, as interactions within an "army of clones" would not theoretically motivate the existence of modern cities. It is therefore important to empirically investigate which level of heterogeneity really matters. Second, the study is related to the wide literature on the economics of innovation and patenting, by exploring under-investigated aspects of patent data, i.e., the skewness of the distribution of patents across inventors and companies.

The paper also offers interesting insights for policy. It is well known that innovation activity is highly concentrated in a small number of cities and regions; these spatial disparities have pushed a number of policies aimed at enhancing local innovation (Agrawal et al., 2012), often based on subsidizing the location of R&D labs of large companies. Very little is known, however, about the effectiveness of these policies, i.e., whether they produce any additional effects on the innovation performance of local firms, or, rather, are just a windfall for large companies, at the taxpayer's cost. The findings of this paper are on line with the general conclusions of many studies assessing the effect of cluster policies based on an "agglomeration economies" rationale.[2] All in all, although there is a positive effect of TIC patents on non-TIC patents, this is likely to materialize in different sectors than those possibly targeted by the policy, and with some delay. Furthermore, the fact that agglomeration economies are in place does not automatically imply that those can be easily triggered by ad-hoc policies.

The next paragraph reviews the relevant literature on patents and innovation; the third one introduces and discusses the definition of TICs; the fourth one describes the empirical methodology and the fifth presents the results; the sixth paragraph concludes.

[2] E.g.: Accetturo and de Blasio, 2012; Duranton, 2011; Martin, Mayer and Mayneris, 2011; Duranton, Mayer and Mayneris, 2010.

2 Patents, localized knowledge spillovers, and the size of innovation

Patent data have become extremely popular in the economic literature in the last two decades, as they represent an easy and accessible way to proxy for innovation activity, which is otherwise generally very hard to measure. Furthermore, the availability of citation linkages further spurred more interest in patent data: for the first time, researchers had a tool to "trace" knowledge spillovers, which previously had been considered one of the most difficult variable to define empirically. A popular book by Jaffe and Trajtenberg (2005), and the free availability of the USTPO dataset from the NBER website, also contributed to multiply the empirical applications based on patent data.

A significant part of this literature has focused on the geographic component of innovation, with a particular interest in the spatial decay of knowledge spillovers. A seminal contribution by Jaffe et al. (1993) shows that a cited-citing patent couple is twice as likely to be located in the same US metropolitan area as a couple of technologically similar patents with no citation links.[3] Similarly, Peri (2005) examines the flows of citations among 147 European and US regions to find that "only 20% of average knowledge is learned outside the average region of origin", and Jaffe (1989) demonstrates that academic research has large effects on the number of private patents developed in the same US state. Finally, Carlino et al. (2007) use patent data for a cross-section of US metropolitan areas to investigate the relationship between urban density and innovation intensity (as measured by patents per capita) finding a positive and robust association. All these contributions (and many similar ones which I omit for brevity) highlight that knowledge spillovers have a geographically limited distance decay.

The nature and causes of knowledge spillovers are still debated. For instance, Breschi and Lissoni (2009), building on previous contributions by Breschi and Lissoni (2001), Zucker et al (1998), and Almeida and Kogut (1999), highlight that defining localized knowledge spillovers as an *externality* can be misleading, as most of the knowledge diffusion may take place through market interactions - namely the spatially-bounded mobility of inventors among workplaces - rather than through informal contacts. Using data on US inventors' applications to the European Patent Office, they were able to show that after controlling for inventors' labour mobility and the related professional network, the role of proximity in explaining knowledge diffusion is greatly

[3] These findings have been strongly criticized by Thompson and Fox-Kean (2005), who argue that the methodology underlying the construction of the control group is seriously flawed. With a more robust approach, based on a finer technological classification of patents, the main results of the paper disappear. However, Murata et al. (2013) perform a similar analysis adopting a continuous definition of space (i.e., abstracting from MSA boundaries) to find that, even when controlling for the fine technological classification, the results of Jaffe et al. (1993) are confirmed.

reduced.

These issues are related to the growing interest in peer effects in science and in the spillovers originating from star scientists. Azoulay et al. (2010) exploit the exogenous variation in the number of "superstar scientists" in US universities due to the sudden death of these individuals to estimate the loss in productivity of their collaborators. They find an average 5-10% decline in their average publication rates, starting 3-4 years after the superstars' death and enduring over time, but no differential effect for co-located collaborators. Waldinger (2010) estimates the effect of the dismissal of scientists from Germany universities during Nazism. Similarly to Azoulay et al., he finds a strong effect on coauthors (13-18%), but no significant effects at department level. Therefore, both studies challenge the existence of localized positive spillovers originating from stars in academic environments.

Similarly, the advocates of the "death of distance" theory argue for a decreasing importance of the role of spatial proximity following the progress of communication technologies (*e.g.*, Friedman, 2005; Quah, 1999; Cairncross, 1997). On the other side, a few studies maintain that technological progress has actually increased the scope for proximity for innovative activities due to the greater importance of face-to-face contacts and agglomeration externalities (*e.g.*, Coyle, 1999). The few empirical assessments of the issue seem to support the "death of distance" hypothesis (Griffith et al., 2011; Ioannides et al., 2008), indeed suggesting that localized knowledge spillovers are fading over time.

The relationship between highly inventive companies and other inventors have been much less explored: to the best of my knowledge, contributions on the subject are confined to the role played by academic star scientists on other researchers (*e.g.*, Azoulay et al., 2010; Oettl, 2012), while industrial patenting is not considered, the only exceptions being Fons-Rosen (2010) and Agrawal et al. (Forthcoming). Fons-Rosen (2010) uses data on the entry of foreign firms into Central and Eastern Europe during the 1990s to analyze the effect on knowledge flows on local incumbent inventors; he compares the MNEs which won the privatization bids with the control group of those which also applied to the bid but lost, finding that winners receive 20% more citations by local inventors, on average, than losers. Differently from this paper, its analysis is at national level and is limited to patent citations. Agrawal et al. (Forthcoming) explore the spatial distribution of large and small (patenting) labs across US MSAs, finding that the birth rate of new start-ups (defined using patents filed for by inventors who were previously employed by large labs) is higher in metropolitan areas which are more diverse, i.e., where large and small labs coexist.

Non-TIC patents, however, can be very important for economic growth. Balasubramanian and Sivadasan (2011) link patent records to Census firm data for the US, in order to assess the impact of patents on firm performance. They focus in particular on firms that patent for the first time, and find a significant and large effect of the first patent on firm growth (but, interestingly, little change in factor productivity).

For local development policies, patents by smaller innovators are probably more relevant than patents filed for by TICs. To the extent that the latter are the outcome of formal R&D activity of large companies, they may have weaker implications on the local economy. Since patenting firms are generally large (Balasubramanian and Sivadasan, 2011), they are often multilocated, and the productivity gains of these inventions are spread across the different plants (and localities).

3 Top Inventing Companies (TICs)

The analysis is based on the NBER/USPTO database, which lists all the patents granted in the United states from 1975 to 1999.[4] For each patent, the database contains the name and city of residence of the inventor(s), the name of the applicant(s),[5] an unique applicant identifier added by the NBER working group on patents (based on the standardization of the name of company and ancillary information), the application and grant year, and the number of citation received. Patents are classified according to the synthetic technological classification developed by Hall et al. (2001) who define five technological categories: Chemical (excluding Drugs); Computers and Communications (C&C); Drugs and Medical (D&M); Electrical and Electronics (E&E); Mechanical.[6] Following a common practice in the patent literature, the geographical location of the patent is derived from the city reported in the first inventor's address field. More details on the data, the geographical assignment, and the geocoding process are reported in Appendix A.[7]

At first glance, the abundance of data makes a micro analysis at inventor or applicant level the most appealing alternative. A deeper view, however, reveals the complexity of such an approach, because the dataset is about patents, not inventors or applicants, implying that when an inventor or applicant is not patenting, her location and activity status are unknown. Furthermore, there is not an unique inventor identifier in the original dataset, and the only information available is full name and city of residence. Spelling errors are frequent. As a consequence, the longitudinal tracking of inventors would require a fuzzy matching of names and cities of residence, with inevitable errors which can easily be non random (*e.g.*, more frequent in cities where duplicate

[4] The dataset is described in details in Hall et al., 2001.

[5] The applicant is the legal entity - either a company or an individual - which owns the right to exploit the invention. In the large majority of cases, the applicant is the employer of the inventor.

[6] The sixth technological category, called "other", is a residual classification and is excluded.

[7] The NBER database has recently been extended until the 2006 or 2008, depending on the version. However, inventors data are not publicly available yet; without information on the city of residence of inventors,it is not possible to correctly geolocate the patents, therefore the date cannot be used here.

surnames are more common, or with a higher rate of inventors with foreign origins). The problem would be perhaps negligible if the focus was only on very productive inventors or applicants; but given that I am interested also in occasional inventors, the issue is crucial.

The analysis is therefore run at city level, focusing on the number of *patents* produced by two groups of applicants: *TICs* and other applicants. The classification of patents into the two groups is based on the total number of patents granted to the applicant over the whole period of analysis (1980-1999): TIC patents are assigned to the most inventive companies in a given technological category. There are many different ways in which the "most inventive companies" can be defined: among those, one can adopt a relative cut-off (e.g., the top 1%), a fixed patent threshold (e.g., more than 1,000 patents), or a ranking rule. I opt for the last method: the most inventive companies are defined as those ranked among the top 50 in their technological category; this corresponds to a lower threshold of 247 patents and a maximum of 701 patents across the five categories. The fixed threshold is preferred to using a relative cut-off as the latter may make the definition of TICs endogenous across categories: stronger positive effects of TIC patents on non-TIC patents increase the number of TIC companies and therefore inflate the denominator of the relative cut-off, increasing the TIC threshold and lowering the number of TIC patents. This would lead to overestimating of the effect of TICs. Another reason to prefer a rank-based definition is that the total number of active firms is not a very meaningful figure in the patent database, as the same firm may be counted more than once.[8] Nevertheless, in the following of the paper I report a robustness test based on a relative cut-off which identifies a roughly similar number of TICs, being set at the top 5‰ of patenting companies. The 5‰ threshold corresponds to ranks from 48th to 133th across the five categories. Results are similar to those obtained from estimations based on the top 50 threshold. Results based on alternative definitions, e.g. limiting the TIC definition to the top 25 or 75 companies, or to companies cumulatively owning 50% of patents in the category, do not affect the main conclusions of the paper either.

Also, since the patent literature offers many examples of large companies filing for patents for reasons unrelated to new inventions (*e.g.* patent thickets), and considering that generally such non-inventive patents are not cited by other patents, I exclude from the analysis all TIC patents which do not receive any citations. In the robustness section I replicate the analysis with forward citations weighted patents obtaining comparable results.[9]

The following step is the definition of the temporal dimension of the analysis. Patent data

[8] The unique identifier for small companies is derived from name harmonization and it might not be fully reliable for smaller companies due to spelling errors, homonymy, and changes of name across time; the identifier for large companies is somehow more reliable, due to their smaller number and the notoriety of their different denominations.

[9] A drawback of using citation-weighted patent counts is the censoring of forward citation information for the later years in the sample.

are rather imprecise in the time dimension: usually a patent is granted[10] 2-3 years after the first application; and it is not possible to know how long an inventor has been working on a patent before applying for it. Timing when local knowledge spillovers may have effect is also difficult: it could be while both the source and destination inventors are working on their respective patents, but it could also happen a few years after the TIC inventor has applied for (or has been granted) it. By inspecting the data I found that the median and mean value of the citation lag of patents in the same MSA is four years, and I therefore choose to adopt periods of the same length (Kerr, 2008, also adopt a period of the four years).[11] This is a reasonable choice in order to "average out" some of the measurement error in the temporal dimension. Six time periods of four years each are therefore defined, spanning from 1976 to 1999. Econometric analysis is generally limited to the last three periods (from 1988 to 1999), as MSA controls are unavailable for the first three periods. I define six periods, however, as the first is used to build the instrumental variables and lagged variables, and the second and the third are used to calculate lags of the patent variables. Table 2 lists the first 10 TICs for each technological category. As it is possible to see, only few companies (e.g., General Electric and General Motors) appear in more than one list. Table 1 reports some basic descriptive statistics for TICs and non-TICs: there are only 164 TICs defined, but they account for around 30% of all patents in the sample, and for around 20% of the inventors. Patents granted to TICs also receive more citations than patents granted to non-TICs, on average (5.11 vs. 3.74).

Being a residual category, the non-TICs are a heterogeneous group of companies. Since the distribution of patents across companies is very skewed - with many firms owning one or few patents, and a small number of companies owning many of them - the majority of firms in the non-TICs category is represented by occasional inventors: the median firm owns just one patent, and the average firm 8.9 .(see Table 1). The fact that their patent stock is small does not obviously imply that these firms are small in term of employment or turnover; on the contrary, available evidence suggests that patentees are larger businesses than non patentees (Balasubramanian and Sivadasan, 2011; Andrew, Criscuolo and Menon, 2014). A detailed characterization of these firms is not possible with available data. However useful insights on the difference between patenting and non patenting firms can be found in papers using patent data matched with US business register records (Balasubramanian and Sivadasan, 2011) or with the Survey of Industrial R&D

[10]The reason why I use the grant year, rather than the application year, is to avoid the bias given by data truncation. More precisely, using the application year would automatically exclude all the patents not granted (but applied for) before 1999, as they are not included in the dataset. This subsample could easily be non-random, *e.g.* better patents may take longer to be examined, etc. However, robustness tests (available upon request) based on the application year produce almost identical results.

[11]I restricted the calculation to patent couples with a maximum citation lag of ten years, as longer lags are unlikely to be related to knowledge spillovers. The citation lag is calculated as the difference between the grant year of the citing and cited patents.

conducted by the Census Bureau and National Science Foundation (Kerr and Fu, 2008).

3.1 Why should TIC patents affect the number of patents produced by other local inventors?

An increase in the number of TIC patents, due to an increase in the productivity or in the number of TIC inventors, may have both positive and negative effects on the number of other patents in the same city. Positive effects can be derived from theories of agglomeration externalities in cities; negative effects originate from a general equilibrium approach to local labour markets (Moretti, 2011). This section describes some mechanisms that could be in place. In the empirical analysis, I estimate a reduced-form model, taking into account only the aggregate net effects of all those mechanisms. Exploring the effect of individual channels would be much more demanding - also in terms of data - and it is above the scope of the present paper. The empirical challenge of disentangling the different - but observationally equivalent - sources of agglomeration economies is well known in urban economics under the name of "Marshallian equivalence".

It is also worth noticing that some of the mechanisms described below may, in theory, work also in the opposite direction (from non-TICs local patents to TICs); to easy the exposition this is not explicitly addressed in the discussion below, however the empirical methodology is designed to be robust to reverse causality.

3.1.1 Positive effects

According to the insightful taxonomy introduced by Duranton and Puga (2004), positive effects may occur through learning, sharing, or matching mechanisms.

Learning generally involves interactions with others and many of these interactions have a face-to-face nature. Cities are therefore a fertile environment for learning: the idea that cities foster the diffusion of knowledge goes back to Marshall (1890) and it is the backbone of endogenous growth theory (Lucas, 1988). However, there is still limited evidence on the channels through which knowledge spillovers take place (Feldman and Avnimelech, 2011). In the specific context of patenting in cities, it is possible to think about at least five different mechanisms:

a) Informal (or tacit) knowledge spillovers: TIC inventors and non-TIC inventors develop informal (personal) contacts due to residential proximity or other kind of face-to-face interactions. Thanks to frequent direct contacts with the TICs inventor, the local non-TIC inventor absorbs ideas for her projects. This channel is nick-named as "ideas over beers" by GHM and is more formally defined as ”diffusion of information” or ”social learning” by Duranton and Puga (2004).

b) Formal knowledge spillovers: TIC inventors transfer their expertise to non-TIC inventors in more formal ways, *e.g.* during seminars or conferences.

c) Workplace contacts: (future) local non-TIC inventors may have the opportunity to work in a TIC, without necessarily being inventors themselves (they may be employed in different duties, or they may leave the institution at an early stage of their career).

d) Workplace mobility and spin-off: active TIC inventors leave their company to start their own business, or they are hired by a local non-TIC. As correctly pointed out by Breschi and Lissoni (2009) and Almeida and Kogut (1999), the previous work experience may be fully priced into the inventor's wage, in which case the spillover is not an externality.

e) Display/attraction effects: the presence of many labs of TICs may attract other inventors to the same city, as they may expect to enjoy the effects of points a, b, and c. This is therefore an indirect form of positive knowledge spillover.

All the five mechanisms may require some time to become effective, thus they may be found in the data with a time lag.

The sharing of public or private goods is also a source of agglomeration economies. To the extent that an increase in the activity of TICs in a city generates or attract the provision of expensive, indivisible goods which are not found in cities where TICs are less active and that also non-TIC inventors benefit of, sharing can also explain a positive effect of an increase of TIC patents on non-TIC patents. An example of sharing could be a display effect similar to that presented above as an indirect knowledge spillover: TICs may contribute to make a city a notorious innovative hotspot, which in turn may attract non-TICs. For instance, for a start-up the location in an innovative city may be a positive signal to potential investors. Such "city identity asset" can be considered as a shared public good. Other more common examples of shared goods originating from the location of TICs and benefiting also non-TIC could be specialized universities, testing facilities and laboratories, or specialized patent attorneys or intermediares.

Finally, positive effects may also arise from better matching of firms and employees. An increase in the number of TICs leads to higher concentration of scientists and employers and therefore generates a thicker labour market for inventors, with a more efficient skill matching. This in turn raises inventors' expected wage and reduces their unemployment risks, eventually attracting more non-TICs to the city, and helping local firms innovate and patent more.

It is worth noticing, however, that both sharing and matching mechanisms can be generally categorized into the group of location (or specialization) economies, i.e., the positive externalities originating from the local proximity of activities in the same industry or sector; while learning mechanisms are usually classified into the group of Jacobian externalities or economies of diversity, i.e., positive externalities arising from the interaction with a wide spectrum of different economic activities, in line with the theory of "cross-fertilization of ideas" developed by Jacobs (1969), later formalized and empirically validated by Glaeser et al. (1992). The specific context under scrutiny in this paper is not an exception. In particular, sharing and matching mechanisms should be stronger within technological categories, rather than between. For matching mechanisms, this should be true almost by definition. For the sharing channels, it reasonable to assume that the kind of goods or facilities that inventors share, e.g. laboratories or equipment, are sector-specific. Therefore, evidence that the positive effect is stronger across technologies, rather than within, would be suggestive that Jacob-type knowledge spillovers tend to dominate over sharing and matching mechanisms.

3.1.2 Negative effects

Potential negative effects may be derived in a general equilibrium approach to local labour markets (Moretti, 2011), and they may mainly occur through an increase in nominal wages. Indeed, a raise in innovation activity in a local TIC plant corresponds to an upward shift in the demand for local scientists, which in turn raises local nominal wages in the sector, at least in the short run (in the longer run, workers may migrate in from other cities, but the inflow is limited by the local supply of housing which affects real wages). Both mechanisms affect negatively the number of non-TIC patents, since local scientists become more costly, without a corresponding increase in productivity (assuming zero positive effects). The actual impact of these effects depends on the skill substitutability among TICs and other inventors, and on the elasticity of supply of labour (also through migration). Higher the skill substitutability, larger would be the increase in local nominal wages; a more elastic supply of labour, instead, would compress the wage growth. Since labour supply is likely to be rigid in the short run, the negative effects are expected to be stronger in the short term, and then to fade over time. Also, negative effects are expected to be stronger within narrowly-defined technological sectors, since skill substitutability of workers is higher, and correspondingly the wage effect is larger.

4 Analysis

This section investigates whether the production of TIC patents in a city affects the production of non-TIC patents in the same city and period, and quantifies this effect. The model also includes a one period lag in the spillover effects of TICs. The following panel with fixed effects is estimated:

$$NonTICs_t^{ik} = \beta_1 \cdot TICs_t^{ik} + \beta_2 \cdot TICs_{t-1}^{ik} + \theta_1 \cdot \sum_{q \neq k} TICs_t^{iq} + \\ + \theta_2 \cdot \sum_{q \neq k} TICs_{t-1}^{iq} + \gamma_1 \cdot Totemp_t^{ik} + \gamma_2 \cdot HH_t^i + \phi^i \delta^k + \delta^k \tau_t + \varepsilon_t^{ik} \quad (1)$$

where i, k, and t index MSAs, categories, and time periods, respectively; TICs and nonTICs are the number of patents in the respective groups, *Totemp* and *HH* are additional time-variant controls, and δ, τ, ϕ are category, time, and MSA fixed effects.[12] The analysis is limited to periods 3-4-5, as MSA controls are not available for previous periods. All variables are expressed in logarithmic form.

The model is estimated at three alternative levels of technological classification, and consequently the index k refers to the three different technological classification levels the data are aggregated at: the MSA level (i.e., no sectoral decomposition, k is constant); the five technological categories level; and the 27 subcategories level. The category and sub-category classification is described in Hall et al., 2001. The five technological categories are the following: Chemical (excluding Drugs); Computers and Communications (C&C); Drugs and Medical (D&M); Electrical and Electronics (E&E); Mechanical; the residual category “Other" is excluded from the sample. The 27 subcategories are more detailed classifications nested inside the five categories, e.g., the category Computers and Communications is further classified into the following subcategories: Communications; Computer Hardware & Software; Computer Peripherals; Information Storage. In addition, the model is also estimated on aggregated data at the MSA level. The three different level of technological aggregation (MSA, category, and sub-category) may give interesting insights on the technological boundaries of knowledge spillovers.

In order to check the consistency of the results across different specifications, regressions are based on five different estimations of model 1. The first is an OLS estimation including all the controls. Regressions two to five are estimated with 2-stages least squares (2SLS): the second and the third include the contemporaneous and lagged TIC patents variable, respectively, excluding all other continuous controls and including all fixed effects; the fourth includes both the contemporaneous and lagged TICs patent variable jointly; finally, the fifth estimation also add all the controls. Details on the instrumental variable strategy are reported in the next section. I total, therefore, I present results from five different estimations at three different aggregation levels.

[12] The sixth technological category, called "other", is a residual classification and is excluded. This does not affect the coefficients but increase precision of the estimates.

The first control variable included in the first and fifth estimations is the total number of TIC patents in technological categories different from i. It is worth noting that this variable might be endogenous: non-TIC inventors might produce knowledge spillovers benefiting TICs in the other technological categories. However, the inclusion of this variable has a limited effect on the main coefficients, especially with the 2SLS estimator. As the latter is robust to omitted variables bias, the estimate of the coefficient for the variable of interest (the number of TIC patents) is consistent even excluding the (endogenous) control.[13]

The other two control variables included are total employment at MSA, category or sub-category level (*Totemp*), depending on the specification, to control for time-variant agglomeration economies and size effects;[14] and the Herfindahl–Hirschman index of technological diversity (*HH*), calculated as the sum of the square of the shares of 4-digit International Patent Classification (IPC) classes by MSA and time period:

$$HH_{it} = \sum_{j} (share_{jit}^{2}) \tag{2}$$

where i indexes MSAs, j IPC classes, and t time periods.

Finally, as mentioned above, all regressions include a wide set of fixed effects. When data are aggregated at MSA level, all regressions includes a MSA fixed effect, as well as period fixed effects. When regressions are aggregated at MSA and (sub)category level, all regressions include a (sub)category-MSA fixed effects, and (sub)category-period fixed effects.

4.1 The choice of the MSA as areal unit

Ideally, the spatial unit at which individual observations are aggregated should match the spatial decay of both knowledge spillovers and labour market clearing forces. Since both boundaries are indefinable entities, the spatial definition inevitably entails a substantial degree of approximation; furthermore, data limitation are particularly stringent at a detailed geographical level. With respect to labour market analysis, the choice of commuting-defined areas, like the MSAs in US, is now widely considered to be a viable option.[15] The definition of the spatial decay of knowledge spillovers is more debated: while several studies have adopted spatial areas as large as US States (*e.g.* Jaffe, 1989; Peri, 2005), available evidence suggests that the effect of knowledge

[13] Attempts to instrument the variable with the sum of the instrument in the other categories provide similar results, but estimates were less precise, due to the large number of endogenous variables and instruments.

[14] The employment variable is sourced from the County Business Patterns (CBP) database maintained by the US Census Bureau. In order to obtain aggregates at technological category and sub-category level, the SIC industry classification is converted into the US patent classification (USPC) using the concordance table provided by the USTPO; the USPC classification is then converted into the NBER technological class and subclass definition adopted in this paper.

[15] See Menon (2012) for a discussion of the statistical properties of MSAs.

spillovers may be stronger within cities, rather than between (Jofre-Monseny et al., 2011), or even rapidly fade out within less than a miles (Arzaghi and Henderson, 2008). It is worth stressing, however, that the findings from the latter study are specific to a single and peculiar industry (advertising), and the authors carefully acknowledge that the mechanisms that they isolate are "proximity benefits for information exchange and networking" (p. 1012), i.e., a specific kind of proximity benefits (those listed at point *a* in section 3.2.1). Using patent data to explore very localized informal knowledge spillovers looks promising, however it also faces important data limitations: only a very small fraction of patent data reports detailed information on inventors' address which would allow to explore knowledge spillovers at a fine spatial scale. An intermediate alternative could be using postcode information to allocate patents to US postcode area, but again the coverage is critically limited - and potentially selected - as in the patent sample used in this paper only 15% of patents report the postcode. Furthermore, US postcode areas are extremely heterogeneous in size and population, and their borders do not have any political nor economic significance, which implies that a postcode-level analysis could be seriously flawed by a Modifiable Areal Unit Problem (MAUP) bias.[16] Furthermore, they would jeopardize the estimation of labour market effects, as they are generally much smaller than local labour market areas.

Therefore, since the reduced-form effect that I estimate is supposedly a mix of labour market and knowledge spillovers mechanisms - which definition in this case is broader than benefits for information exchange and networking - the Metropolitan Statistical Area is the most sensible spatial unit of analysis, among the limited number of available options. However, the effect of short-decay knowledge spillovers may be underestimated. It is therefore appropriate to specify that the analysis takes into account only MSA-level knowledge spillovers, which may not fully reflect other short-decay spillovers.

4.2 The instrumental variable for the number of TIC patents

Estimates of equation 1 can be inconsistent due to reverse causality or omitted variable biases, especially for the main variable of interest (the number of TIC patents). For instance, non-TIC inventors may affect the productivity of TICs, and a dynamic university (or public subsidies) may attract a large number of TICs and non-TIC inventors to the same city. I therefore create an instrumental variable for the number of TIC patents in order to allow a causal interpretation of the results.

The intuition for the instrument builds on the fact that most TICs are located in several MSAs

[16] See Openshaw (1983) for an introduction to the MAUP.

and in different US states. Table 6 lists the top 25 assignees in the period under examination (1980-1999), reporting the number of different MSAs and states where at least 100 patents are authored, and the highest share of patents authored in an individual MSA: only one company is located in one MSA (Ford Motor), while all the remaining assignees are located in several different cities and states. Smaller TICs show a similar pattern. Therefore, an exogenous variation in the productivity of TIC inventors in a given MSA and period may arise from the interaction of two factors: i) an historical presence of inventors working for a given company in that MSA, and ii) a US-wide increase in the productivity or the market share of this company in the given period. To the extent that the first factor is path-dependent and exhibits some inertia over time, it is exogenous to contemporaneous MSA-specific factors, conditional on MSA fixed effects. At the same time, the productivity of TIC inventors working for the same companies in different cities is likely to be correlated, due to the sharing of similar strategies and resources, competition pressure, market demand, etc. Therefore, a US-wide productivity shift in a given company translates into MSA-specific productivity shocks in proportion to the number of inventors working for that company in the given MSA.

The IV strategy is close in spirit to the approach of Bartik (1991) and Blanchard and Katz (1992), among others, who instrument regional economic growth interacting the lagged sectoral structure of a region with the contemporaneous national sectoral trend. In the next section the construction of the instruments is explained in detail.

4.2.1 Instrumental variable construction

The instrumental variable is calculated as follows:

a) For the period 0, each MSA, and each TIC, I calculate the share of active inventors over all TIC inventors in the given MSA. In the case of TIC inventors with multiple MSAs or assignees in the same period, the modal one is chosen.

b) For each period, each TIC, and each MSA, I calculate the average number of patents produced in the whole US, excluding the patents authored by inventors located in the given MSA.

c) For each MSA, period, and assignee, I multiply the share of inventors in the period 0 calculated at point a by the average number of patents produced by TIC inventors sharing the same assignee in period t calculated in b. Subsequently, I sum the outcome by MSA, period, and technological category (if an inventor has patented in different categories in the same period, the modal one is chosen). The result is the instrumental variable for total number of TIC patents in period t, by MSA and category.

Formally, it can be summarized by the following equation:

$$IV_{ikt} = \Sigma_a (TIC\,Inv_{ika0} \cdot AvPat_{iat})/Pat_{ik0} \tag{3}$$

where i indexes MSAs, t periods, k technological categories, and a the assignees. In the few cases in which the value of point b is missing (because there are not other TIC inventors with the same assignee in other MSAs), it is replaced with the contemporaneous US-wide average productivity of TIC inventors in the same technological category.

4.2.2 Instrument's validity and falsification tests

The validity of the IV relies on an exclusion restriction related to point a, i.e., once MSA fixed effects are controlled for, the number of TIC inventors working for a given assignee in the first period has no independent effect on the number of non-TIC patents developed in period n in the same MSA/category; and on an assumption of exogeneity related to point b, i.e., given that TICs and other inventors have different assignees, it is assumed that the average productivity of an assignee in the whole US (calculated excluding the given MSA) has no independent effect on the productivity of non-TIC inventors of that MSA.

A concern related to point b is that the address of residence of a few inventors does not truly reflect the location of their workplace while working on the patent; this can be due to errors in the data or geocoding process, in the city name spelling, or to a subsequent change in the inventor's address. This would threaten the exogeneity of the IV, as the productivity of those inventors might not be exogenous to local unobservables, especially in the case in which their real MSA and the MSA the instrument is built for coincide. Therefore, to be on the safe side, when creating the IV all the company-MSA pairs with less than 100 patents are dropped; i.e., patents assigned to MSAs which are unlikely to host a R&D lab, and which may misreport the real inventor's location, are excluded. This should reduce the "noise" in the geographical assignment of the R&D labs of TICs, minimizing the risk of endogeneity of the IV.

A second concern relates to the plausibility of the exclusion restrictions for assignees' shares in period 0. The historical presence of one or more TICs in a MSA might depend on a persistent trend over time, which in turn might also correlate with non-TIC patents in the following periods. For instance, an dynamic university created in the early '70s may have attracted a productive TIC to a city; over the '90s the same expanding university may affect positively both the number of TIC and non-TIC patents in that city, challenging the exclusion restrictions on the instrumental variable.

I design a falsification test to address this and related concerns. The test is based on creating a "placebo IV" which is calculated following formula 3, but inventors' shares at time 0 are multiplied by the average productivity at time t of a randomly picked TIC within the same MSA and period, rather than to the correct one. Therefore, the placebo instrument absorbs only the (supposedly) time invariant component of the correlation with the correspondent number of TIC patents in the same MSA and category. If period zero shares were endogenous due to an unobserved MSA-specific persistent trend which correlates with both the number of TIC and non-TIC inventors, the placebo IV would correlate with the number of TIC patents; if the correlation were strong enough, the placebo IV would be significant in the first stage regression. The table added to the paper shows that this is never the case: all the coefficients are very far from being statistically different from zero (table 9).

5 Results

The regression results are reported in table 5-6-7.[17] As noted in the previous section, the estimations are based on five different specifications - OLS with controls (col. 1), 2SLS with contemporaneous TIC patents only (col. 2), 2SLS with lagged TIC patents only (col. 3), 2SLS with both contemporaneous and lagged TIC patents (col. 4), and 2SLS with all the variables (col. 5) - at three aggregation levels: MSAs (table 5), MSAs and five technological categories (table 6), and the MSAs and 27 technological subcategories (table 7). All columns with IV estimation also report the Angrist and Pischke (2009, pp. 217-18) first-stage F statistics for tests of weak identification when there is more than one endogenous regressor (AP). When just one variable is considered to be endogenous, the test is equivalent to the traditional first stage F-statistic.[18] In most cases, results from first-stage regressions confirm that the instrument is strong, especially at category and subcategory level.

Overall, the results suggest that there is a positive effect of TIC patents on non-TIC patents at MSA level. The effect is stronger with a one-period time lag, and it is not confined within technological categories. Within technological sub-cateogories, the estimations fail to find any

[17]Standard errors are clustered at MSA level. Alternative estimates based on clustering at the state-year pairwise combination give almost identical standard errors. Since the distribution of total patents across MSAs shows a large variance, all regressions are (analytically) weighted by the total number of patents over the period of analysis (see Angrist and Pischke, 2008, for a detailed discussion on the suitability of weighted regressions when the sample is composed by grouped individual observations). I also dropped all the MSA-Category pairs with less then 10 patents over the whole period of analysis. Unweighted regression results and full sample results are qualitatively similar but less precise. Logarithmic transformation is applied to the patent count augmented by one unity, in order to keep observations with zero patents in the sample. Robustness tests including a zero TIC patent dummy or limited to the subsample with strictly positive patents suggest that the procedure does not introduce any significant bias.

[18]The Angrist-Pischke (AP) first-stage F statistic is calculated for each individual endogenous regressors by "partialling-out" linear projections of the other endogenous regressors. The AP test will fail to reject if a particular endogenous regressor is unidentified. Values of the AP first-stage F can be compared to the Stock-Yogo (2002, 2005) critical values for the Cragg-Donald F statistic with K1=1.

significant positive effects, once endogeneity is dealt with. This implies that across technologies and with a time lag, positive agglomeration effects prevail; within narrowly defined technologies and the same time period, negative wage effects fully counterbalance the positive ones. This may happen because negative effects are stronger within sub-categories, and/or because positive effects are weaker. Both the justifications are consistent with the theoretical predictions. On the one hand, wage effects, which push the coefficient downward, are stronger within narrow technological categories due to higher skill substitutability - and in the short run, before workers' relocation takes place. On the other hand, the positive agglomeration effects arising from learning mechanims take a few years before being effective, and are technologically complementary. Without knowing the exact magnitude of the two effects, it is hard to say which of the two factors contributes most to explaining the absence of positive effects at sub-category level. However, the results are at least suggestive that most of the positive effects arise from Jacob-type, cross-technology knowledge spillovers, rather than from other sector-specific externalities, like sharing or matching mechanisms. In the following, I will present an additional test based on citation data which further substantiate this interpretation of the results.

The results at MSA level are reported in table 5. The first column report the results of the OLS estimation: both the contemporaneous and the lagged level of TIC patents are significant, and both the coefficients are equal to 0.09. Column 2 to 5 report different specifications using the instrumental variable approach. When included alone, the contemporaneous and the lagged number of TIC patents are both significant, with a coefficient of 0.15 and 0.22, respectively. Once the two TIC variables are jointly included in the same specification, only the lagged one is significant, with a value ranging from 0.20 to 0.22, depending on whether the other controls are included or not. Given that the specifications are log-linear, coefficients can approximately be interpreted as elasticities: a 10% increase in the number of TIC patents in a given city corresponds to around a 2% increase in the number of non-TIC patents, *ceteris paribus*.

When the sample is disaggregated to technological categories (table 6), the results are qualitatively similar to those obtained at MSA level, but the coefficients' magnitude is slightly reduced, and their significance is weaker. The contemporaneous coefficient is now never significant across the five specifications. The lagged coefficient, instead, is significant and equal to 0.06 in the OLS regression (col. 1), while ranges from 0.15 to 0.23 in the IV regression, but loosing progressively significance once further controls are introduced.

When the sample is further disaggregated to 27 technological subcategories (table 7), the contemporaneous TIC coefficient is never significant, and the lagged one is significant only in the OLS regressions (col 1). Across all IV specifications (col. 2-5), the contemporaneous coefficient turns even negative, and the lagged one is very close to zero. Therefore, as anticipated, at the sub-category level the regression analysis fails to find any causal effect of TIC patents on non-TIC patents.

Why MSA- and category-level OLS estimates are downward biased? There are at least three plausible explanations for that: negative selection, measurement error, and local average treatment effect (LATE). Negative selection may arise because, in general, those TIC inventors that are more "exposed" to non-TIC inventors might produce less knowledge spillovers than the average TICs inventor. In other words, TIC inventors localized in "non-TICs cities" may be less productive than TIC inventors localized in "TICs cities". As this lower quality is unobserved, it introduces a (downward) bias in the OLS estimates. Another plausible explanation for the downward bias could be a measurement error in the TICs variable: the intensity of activity of TIC inventors in a locality is approximated by the number of patents they produce, but the measure is clearly noisy, as patents are heterogeneous in quality. To the extent that the measurement error of the instrumental variable is independent from the one in the endogenous variable, IV estimates may eliminate the "attenuation bias" of the OLS coefficient. The independence of the two errors is plausible as the variables are measured using patents in different localities (in the specific city and in the whole US excluding that city, respectively). Finally, to the extent that the elasticity of the endogenous regressor with respect to changes in the instrumental variables is not constant across groups, 2SLS estimates may correspond to a local treatment effect, rather than to an average treatment effect (ATE) (Imbens and Angrist, 1994). In this specific context, it is likely that the elasticity of the endogenous variable to the instrument is higher for incumbent plants, since one of the component of the instrument is the historical presence of TIC inventors in the MSA. Incumbent plant inventors may have a stronger effect, since they are more connected with local other non-TIC inventors; this may explain an higher local treatment effect.

5.1 Testing for the knowledge spillover channel

As discussed in section 3.2, positive effects of TIC on non-TIC patents may occur through learning, sharing, and matching mechanisms. I have argued that the results suggest that the learning channels tend to prevail over sharing and matching mechanisms. In what follows I present additional evidence supporting this view.

Citations are often used in patent analysis as a way to “trace” knowledge spillovers (at least since Jaffe et al., 2003). Therefore, is it plausible to assume that most of locally cited patents are among those that generate knowledge spillovers to non-TIC patents, although the full group of patents generating spillovers might be much larger, as not every knowledge spillover results in a citation. However, if regressions based on such a tiny share of patents still pointed to the existence of positive effects of TIC patents on non-TIC patents, this would further support the view that knowledge spillovers play an important role in explaining the positive effect of TIC patents over non-TIC patents. The Table 8 reports MSA-level regressions in which the number of TIC patents have been restricted to those patents that are cited by non-TIC patents in the same MSA. Those patents represent only a tiny minority - only 5% on average - of all TIC patents produced in a given MSA. In all other respects, the estimations are identical to those reported in 5.

The results reported in table 8 indeed support the relevance of learning mechanisms, as OLS coefficients are quite similar to those presented in the baseline regressions (col 1). The results can however simply be similar because the two group of TIC patents - those cited by local non-TIC and the others - might be strongly correlated. Cols. 2 and 3 therefore also include the number of other TIC patents on the right hand side of the regression. As it is possible to see, while some of the contemporaneous effect is clearly due to the correlation of the two variables, the lagged effect is almost entirely attributable to the TIC patents cited locally. This result is in line with previous findings, and are a further evidence that knowledge spillovers need a few years before materializing. IV results are also in line with those discussed above, although coefficients loose some significance, probably because of a much weaker first stage, as highlighted by the lower AP statistics. All in all, the test is thus supportive of the prominent role played by knowledge spillovers in explaining the positive effect of TIC patents on non-TIC patents at MSA level.

5.2 Robustness

The core results are subject to a variety of robustness tests - regarding the econometric specification, the TIC definition, the geographic assignment of patents, etc. which leave the main conclusions unaffected.

A first robustness test consists in adding a lagged dependent variable to the right hand side of the regression equation, as its omission might lead to un omitted variable bias in the estimation of the coefficient on the lagged TIC patent variable, and this might explain why the latter tends to be larger than the contemporaneous one. The model is estimated with a GMM Arellano-Bond dynamic panel; the lagged levels (at *t-2* and *t-3*) of the dependent variable are set as instruments for the lagged dependent variables in the first-differences equation, while the endogenous TIC patents variables are instrumented with the exogenous IV. The results, available upon request, show that the lagged non-TIC variable is positive and significant at MSA level, however all other coefficients are aligned with the findings of the 2SLS regressions.

A further robustness test challenges the choice of limiting the definition of TICs to the top 50 companies within a given technological category. I therefore replicate the analysis using a relative cut-off, set at the 5‰ (per mil) level; in such a way, the TIC definition involves a number of companies which is only slightly larger than the top 50 threshold. The results (available upon request), are very similar to those obtained from the main specifications. I also replicated the analysis using two different ranking thresholds, equal to 25 and 75, respectively. Although estimates are less precise in a few cases, the values of the main coefficients are again close to those presented above, and my general conclusions are unaffected by the change in the threshold. Finally, I repeat the empirical analysis defining as TICs those companies cumulatively owning 50% of patents in the category: this also does not affect the main conclusions of the paper (the results of these two last tests are available upon request).

Another potential source of concern may be the choice of considering only the first author of the patent in the geolocalization process.[19] This is based on the assumption that the first author is the leading scientist, but it would introduce a bias if authors are listed in alphabetical order. Therefore I check whether authors whose surname begins with one of the first letters of the alphabet are more likely to be reported as first authors, compared with second or third authors, finding that differences in probability are very low and fade out after the first five-six letters (table available upon request). This evidence therefore suggests that the first author should be the project leader. However, to be on the safe side, I also replicate all the estimations limiting the sample only to single-authored patents (54% of the sample), or using the city of the second inventor for the geographical assignment of patents. In both cases, the results (available upon request) are similar to the baseline estimates.

The choice of patent count as a measure of productivity of TIC inventors may also be questioned, since patents are very heterogeneous in quality and value. As a consequence, patent counts can be a very noisy proxy. Although I exclude from the TICs group all patents which do not receive any citations, this might not be enough. If the patent value heterogeneity behaves as a classic measurement error and if the measurement error in the instrumental variable is independent from those in the endogenous variable, the 2SLS results are still consistent. However, given that both the endogenous and IV variables are based on information from the same companies, the assumption of the independence of the measurement error may not hold. Another possible solution is weighting TIC patents by the number of forward citations, since the latter has been shown to be a reasonably good proxy for patent value (Hall *et al.*, 2005). I thus replicate the analysis using the quality-corrected measure of TIC patents. The results (available upon request)

[19] In the patent literature, using only the first author is probably the most common option, although some researchers also use fractional count or multiple allocation.

confirm that the coefficient values are reassuringly close to those of the main specifications.

6 Conclusions

This paper assesses whether the number of patents developed by inventors working for the top inventive companies (TICs) has any causal effect on the number of patents granted to other inventors located in the same MSA. TICs are defined as those companies which are ranked among the top 50 in the number of owned patents within their technological category. Causality is inferred through instrumental variable estimation.

Economic theory predicts that an increase in innovation activity of TICs affects the production of non-TIC patents positively through knowledge spillovers and other agglomeration externalities, and negatively through increased local wages. The empirical findings are coherent with the theoretical framework: results show that positive effects prevail; they are stronger with a time lag and are not necessarily bounded within sectors, providing support for the relevance of economies of diversity. A 10% increase in the number of TIC patents leads to an increase of about 2% in the number of non-TIC patents in the same MSA over the following 4-8 years. Within narrowly defined technological categories, however, the positive effects completely disappear; this may happen partly because negative effects are expected to be stronger within sectors, and partly because positive knowledge spillovers are expected to be stronger across sectors. Results survive to a number of potentially demanding robustness tests. Taken together, and including also an additional test based on citation data, the results suggests that Jacob-type knowledge spillovers, which are not confined within technological categories, tend to prevail over other sector-specific source of agglomeration economies, including sharing and matching mechanisms. These findings are in line with a substantial stream of research proving the economic relevance of localized knowledge spillovers.

As Duranton and Puga (2004) forcefully argue, most of the urban agglomeration mechanisms arise from the interactions of heterogeneous agents; empirically, it is important to understand which level of heterogeneity really matters. This paper provides some evidence on that, by showing that technological and temporal heterogeneity is needed for the interaction of TIC companies with other inventing firms to produce positive effects.

The findings also bring in relevant implications for local development policies. As discussed by GHM, policy makers are increasingly keen in subsidizing the local investments of large companies, with the idea that these may generate agglomeration spillovers and benefit local firms. Do this paper's findings provide ground to these policies? Given the positive effect of TIC patents on the number of other local patents, the attraction of TICs to a city may have a positive effect on the local economic environment: in the medium run, TICs positively affect local patenting activity, which in turn might foster the birth of new plants, the innovation output of local businesses, and the generation of new employment. Thus, even though R&D labs of big corporations may have only a limited direct effect on the local economy, as most the of the employment and value added is located elsewhere, they might still be beneficial through a number of indirect channels. However, the attraction of TICs may impact sectors and time periods which are not those directly affected by the policy intervention, making difficult for policy makers to target specific sectors and to grasp the benefits in the short term.

References

Accetturo A., de Blasio G., 2012. Policies for local development: An evaluation of Italy's "Patti Territoriali". Regional Science and Urban Economics, vol 42:1-2, pp. 15-26.

Agrawal, A. K., Cockburn, I. M., Galasso, A., & Oettl, A., Forthcoming2014. Why are some regions more innovative than others? The role of firm size diversity. Journal of Urban Economics.

Almeida, P., and B. Kogut, 1999. Localization of knowledge and the mobility of engineers in regional networks. *Management Science* 45.

Andrews D., C. Criscuolo and C. Menon, 2014. Do resources flow to Patenting Firms? Cross-Country Evidence from Microdata. Working paper, OECD.

Angrist J. D. and J-S. Pischke, 2008. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press.

Arzaghi M. and J. V. Henderson, 2008. Networking off Madison Avenue, *Review of Economic Studies*, Oxford University Press, vol. 75(4), pages 1011-1038.

Azoulay, P., J. S. Graff Zivin, and J. Wang, 2010. Superstar Extinction. *Quaterly Journal of Economics* 125:2, 549-589.

Balasubramanian N, J Sivadasan, 2011. What Happens when Firms Patent? New Evidence from US Economic Census Data, *The Review of Economics and Statistics*, vol. 93;1, pp 126-146.

Bartik T.J., 1991. *Who Benefits from state and Local Economic Development Policies?* W.E. Upjohn Institute.

Bessen J, 2008. The value of U.S. patents by owner and patent characteristics. *Research Policy* 37:5.

Blanchard O., L. F. Katz, 1992. Regional Evolutions, *Brookings Papers on Economic Activity.* 23:1. pp 1-76.

Breschi, S. and Lissoni, F.,2001. Knowledge spillovers and local innovation systems: a critical survey. *Industrial and Corporate Change.* 10: 975–1005.

Breschi S., and F. Lissoni, 2009. Mobility of skilled workers and co-invention networks: an anatomy of localized knowledge flows.*Journal of Economic Geography* 9 pp. 439–468.

Cairncross F., 1997. *The death of distance: how the communications revolution will change our lives.* Harvard Business School Press. Boston.

Carlino G.A., S. Chatterjee, R.M. Hunt, 2007. Urban density and the rate of invention. *Journal of Urban Economics*, 61: 389-419.

Coyle, D., 1999. *The Weightless World: Strategies for Managing the Digital Economy.* MIT Press.

Duranton, G., 2011. California Dreamin': The feeble case for cluster policies. *Review of Economic Analysis*, 3(1), 3-45

Duranton, G., P. Martin, T. Mayer and F. Mayneris, 2010. *The Economics of Clusters: Lessons from the French Experience.* OUP Catalogue, Oxford University Press.

Duranton, G., D. Puga, 2004. Micro-foundations of urban agglomeration economies, in: J. V. Henderson and J. F. Thisse (eds.), *Handbook of Regional and Urban Economics*, edition 1, volume 4, chapter 48, pages 2063-2117 Elsevier.

Ellison, G.D., E.L. Glaeser, and W.R. Kerr, 2010. What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns. *The American Economic Review*, 100:3, pp 1195-1213

Feldman MP, G. Avnimelech, 2001. Knowledge spillovers and the geography of innovation revisited: a 20 years' perspective on the field on geography of innovation. In D. B. Audretsch, O. Falck, S. Heblich,A. Lederer (eds), *Handbook of Research on Innovation and Entrepreneurship*, Edward Elgar.

Fons-Rosen, C., 2010. Knowledge Flows Through FDI: the Case of Privatisations in Central and Eastern Europe. Working paper.

Friedman, T., 2005. The World is Flat. New York: Farrar, Strauss and Giroux.

Glaeser, E.L., Kallal, H.D., Scheinkman, J.A., and Shleifer, A., 1992. Growth in cities. *Journal of Political Economy* 100, 1126-1152..

Greenstone M., R. Hornbeck, E. Moretti, 2010. Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings. *Journal of Political Economy*, 118(3).

Griffith, R., S. Lee, and J. Van Reenen, 2011. Is distance dying at last? Falling home bias in fixed-effects models of patent citations. Quantitative Economics, 2: 211–249. doi: 10.3982/QE59

Griliches Z, 1990. Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature*, 28: 4.

Hall, B.H., A. B. Jaffe, M. Trajtenberg, 2001. The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools. *NBER Working Paper* 8498.

Hall, B.H., A. B. Jaffe, M. Trajtenberg, 2005. Market Value and Patent Citations. *RAND Journal of Economics*, vol. 36, issue 1, pages 16-38

Imbens G. W., J. D. Angrist, 1994. Identification and Estimation of Local Average Treatment Effects. *Econometrica,* Vol. 62, No. 2 (Mar., 1994), pp. 467-475.

Ioannides Y., H.G. Overman, E. Rossi-Hansberg, and K. Schmidheiny, 2008. The effect of information and communication technologies on urban structure. *Economic Policy*, 23, pp 201-242.

Kerr, W.R., 2008. Ethnic Scientific Communities and International Technology Diffusion. *The Review of Economics and Statistics*, 90, (3), 518-537

Kerr, W.R., and S. Fu, 2008. The survey of industrial R&D—patent database link project. *The Journal of Technology Transfer,* 33.2 (2008): 173-186.

Jacob, J., 1969. The Economy of Cities. New York: Vintage.

Jaffe A.B., 1989. Real Effects of Academic Research. *The American Economic Review*, 79:5, pp. 957-970.

Jaffe A.B., M. Trajtenberg, 2005. *Patents, Citations, and Innovations: A Window on the Knowledge Economy.* The MIT Press.

Jaffe A.B.. M. Trajtenberg, and R. Henderson, 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics* 10, 577-598.

Jofre-Monseny J, R. Marín-Lópeza 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*, 70:2-3, pp. 61-74

Marshall, A., 1890. *Principles of Economics.* Macmillan, London.

Martin, P., T. Mayer and F. Mayneris, 2011. Public support to clusters: A firm level study of French Local Productive Systems. *Regional Science and Urban Economics*, Elsevier, vol. 41(2), pages 108-123, March.

Menon, C., 2012, The bright side of MAUP: Defining new measures of industrial agglomeration. *Papers in Regional Science*, 91: 3–28

Moretti E., 2011, Local Labor Markets. in O. Ashenfelter, D. Card (eds.), *Handbook of Labor Economics*, vol. 5.

Murata, Y. R. Nakajima, R. Okamoto, and R. Tamura, (2013). Localized Knowledge Spillovers and Patent Citations: A Distance-Based Approach, *Review of Economics and Statistics* (forthcoming).

Oettl, A., 2012. Reconceptualizing Stars: Scientist Helpfulness and Peer Performance, *Management Science*, 58(6), 1122-1140.

Openshaw, S., 1983. *The modifiable areal unit problem.* Norwick: Geo Books.

Peri G, 2005. Determinants of Knowledge Flows and Their Effect on Innovation.*The Review of Economics and Statistics*, vol. 87, issue 2, pages 308-322.

Silverberg G, B Verspagen, 2007. The size distribution of innovations revisited: An application of extreme value statistics to citation and value measures of patent significance. *Journal of Econometrics*, 139:2.

Quah, D., 1999. The Weightless Economy in Economic Development. *CEP Discussion Paper* No. 417.

Trajtenberg M., G. Shiff, R. Melamed, 2006. The "Names Game": Harnessing Inventors' Patent Data for Economic Research. *NBER Working Paper* No. 12479.

Waldinger F., 2010, Quality Matters: The Expulsion of Professors and the Consequences for PhD Student Outcomes in Nazi Germany. *Journal of Political Economy*, vol. 118, no. 4

Zucker, L. G. and M.R. Darby, 2007. Star Scientists, Innovation and Regional and National Immigration *NBER Working Paper* no. 13547.

Zucker, L. G., Darby, M. R., Armstrong, J.,1998. Geographically localized knowledge:spillovers or markets? *Economic Inquiry*. 36: 65–86.

Table 1: Summary statistics of TIC and non-TIC patents and companies

Statistics	TIC	Non-TIC
No. of inventors	68,258	270,903
No. of companies	164	54,284
Total no. of patents	209,198	483,551
Average patents per company	1275.59	8.90
Median patents per company	740	1
Average citations per patent	5.11	3.74
Median citations per patent	2	1

Note: the table reports descriptive statistics comparing Top Inventing Companies (TICs) with other companies (non-TICs), as well as the patents assigned to the two categories of companies. The statistics are limited to periods 3, 4 and 5, i.e., to the 1988-1999 interval. See paragraph 3 for the definition of TIC and non-TIC companies.

Table 2: The top 10 inventing companies (TIC) by technological category

Chemical (excluding Drugs)	no. of patents	Computers and Communications	no. of patents	Drugs and Medical	no. of patents
E I DU PONT DE NEMOURS	6417	INT BUSINESS MACHINES	10787	MERCK	2204
DOW CHEM	5255	MOTOROLA	5751	ELI LILLY	1501
EASTMAN KODAK	4923	AT & T	4622	BRISTOL MYERS SQUIBB	1238
GEN ELECTRIC	4790	HEWLETT PACKARD	2746	PROCTER & GAMBLE	1221
MOBIL OIL	3586	XEROX	2261	AMERICAN CYANAMID	1031
PHILLIPS PETROLEUM	3543	TEXAS INSTR	2259	UNIV OF CALIFORNIA	1015
EXXON RES & ENG	2798	INTEL	2119	MEDTRONIC	999
UOP	2746	UNISYS	2010	ABBOTT LAB	985
MINNESOTA MINING & MFG	2524	LUCENT TECH	1660	PFIZER	881
MONSANTO	2340	EASTMAN KODAK	1572	WARNER LAMBERT	876

Electrical and Electronics	no. of patents	Mechanical	no. of patents
GEN ELECTRIC	8144	GEN MOTORS	4243
MOTOROLA	4910	XEROX	3542
INT BUSINESS MACHINES	4801	EASTMAN KODAK	3114
WESTINGHOUSE ELECTRIC	4492	GEN ELECTRIC	2713
TEXAS INSTR	3471	CATERPILLAR	2220
RCA	3244	FORD MOTOR	1959
AT & T	3123	INT BUSINESS MACHINES	1717
GEN MOTORS	1802	BOEING	1315
AMP	1731	WESTINGHOUSE ELECTRIC	1291
HUGHES AIRCRAFT	1635	UNITED TECH	1283

Note: the table ranks the companies according to the total number of assigned granted patents in each technological category by the USTPO, over the period 1975-1999. Patents authored by non-US inventors are excluded. Source: author's elaboration on NBER Patent database.

Table 3: The top 10 MSAs in the number of TIC and non-TIC patents

No. of TIC patents	Metropolitan Statistical Area
10,366	Rochester, NY MSA
9,182	New York-Northern New Jersey-Long Island, NY-NJ-CT-PA CMSA
8,416	San Francisco-Oakland-San Jose, CA CMSA
6,028	Philadelphia-Wilmington-Atlantic City, PA-NJ-DE-MD CMSA
4,621	Chicago-Gary-Kenosha, IL-IN-WI CMSA
4,230	Austin-San Marcos, TX MSA
3,712	Dallas-Fort Worth, TX CMSA
3,656	Minneapolis-St. Paul, MN-WI MSA
3,601	Albany-Schenectady-Troy, NY MSA
3,083	Detroit-Ann Arbor-Flint, MI CMSA
No. of non-TIC patents	**Metropolitan Statistical Area**
43,798	San Francisco-Oakland-San Jose, CA CMSA
40,044	New York-Northern New Jersey-Long Island, NY-NJ-CT-PA CMSA
36,599	Los Angeles-Riverside-Orange County, CA CMSA
24,263	Chicago-Gary-Kenosha, IL-IN-WI CMSA
20,763	Boston--Worcester--Lawrence--Lowell--Brockton, MA--NH NECMA
18,200	Philadelphia-Wilmington-Atlantic City, PA-NJ-DE-MD CMSA
17,930	Detroit-Ann Arbor-Flint, MI CMSA
12,985	Minneapolis-St. Paul, MN-WI MSA
12,037	Houston-Galveston-Brazoria, TX CMSA
11,266	Dallas-Fort Worth, TX CMSA

Note: the table ranks tMetropolitan Statistical Areas (MSAs) according to the total number of patents assigned to Top-Inventing Companies (TICs) and to other companies (non-TIC) , respectively, over the period 1988-1999. See paragraph 3 for the definition of TIC and non-TIC companies.

Table 4: The location of top inventing companies

Company	Nr of patents	Nr of MSAs	Nr of States	Share of the 1st MSA
GEN ELECTRIC	12892	20	14	0,49
INT BUSINESS MACHINES	12281	16	16	0,18
EASTMAN KODAK	8828	4	4	0,86
MOTOROLA	8383	7	5	0,35
AT & T	7010	10	10	0,62
E I DU PONT DE NEMOURS	5991	5	5	0,89
XEROX	5918	3	2	0,78
GEN MOTORS	5330	10	4	0,51
DOW CHEM	5197	5	5	0,49
MINNESOTA MINING & MFG	5064	3	3	0,93
MOBIL OIL	4830	4	4	0,56
TEXAS INSTR	4617	5	2	0,74
WESTINGHOUSE ELECTRIC	3663	7	7	0,46
RCA	3548	4	4	0,44
HUGHES AIRCRAFT	3377	3	2	0,91
FORD MOTOR	3135	1	1	1,00
ALLIED SIGNAL	2969	8	8	0,45
HEWLETT PACKARD	2963	7	6	0,42
UNITED TECH	2907	4	2	0,63
UNISYS	2664	11	10	0,16
EXXON RES & ENG	2599	3	3	0,71
ROCKWELL INT	2459	5	5	0,59
AMERICAN CYANAMID	2119	3	3	0,60
MONSANTO	2087	4	4	0,70
CATERPILLAR	1986	2	1	0,81

Note: the first column reports the number of patents owned by the company, the second (third) column report the number of different MSAs (US States) in which at least 100 patents have been authored by local inventors, and the fourth column reports the share of patents authored in the MSA with the largest number of authored patents.

Table 5: The effect of TIC patents on other patents, level of aggregation: MSA

	(1)	(2)	(3)	(4)	(5)
Dep. var.			Non-TIC patents		
Method	OLS		IV - 2SLS		
TICs pats. (t)	0.090**	0.157*		0.016	-0.019
	(0.045)	(0.087)		(0.119)	(0.125)
TICs pats. (t-1)	0.090***		0.228***	0.220**	0.199**
	(0.020)		(0.066)	(0.101)	(0.100)
Total MSA empl.	0.196				0.185
	(0.122)				(0.141)
HH index	0.162***				0.204***
	(0.036)				(0.073)
FIRST STAGE REGRESSION					
AP TICs pats. (t)		13.99		7.692	7.158
AP TICs pats. (t-1)			15.27	11.40	10.11
Period f.e.	YES	YES	YES	YES	YES
MSA f.e.	YES	YES	YES	YES	YES
Observations	840	840	840	840	840

Note: robust standard errors, clustered at MSA level, in parentheses. The dependent variable is the number of non-TIC patents. The unit of observation is the MSA-period. The time interval is 1988-1999, divided in three times period of four years each. The endogenous variable are TICs pats. (t) and TICs pats. (t-1). The excluded instruments are IV (t) (col. 2); IV(t-1) (cols. 3); IV(t) and IV(t-1) (col. 4-5). All variables are expressed in logarithmic form. All regressions are (analytically) weighted by the total number of patents over the period of analysis. The first-stage pane reports the Angrist and Pischke (2009) first-stage F statistics for tests of weak identification. *** p<0.01, ** p<0.05, * p<0.1

Table 6: The effect of TIC patents on other patents, level of aggregation: MSA-Category

	(1)	(2)	(3)	(4)	(5)
Dep. var.			Non-TIC patents		
Method	OLS		IV - 2SLS		
TICs pats. (t)	-0.017	0.124		-0.048	-0.096
	(0.025)	(0.151)		(0.168)	(0.186)
TICs pats. (t-1)	0.062***		0.223**	0.234*	0.149
	(0.022)		(0.101)	(0.122)	(0.110)
TICs other cats. (t)	0.076**				0.097*
	(0.036)				(0.058)
TICs other cats.(t-1)	0.053**				0.024
	(0.023)				(0.045)
Total empl.	0.586***				0.600***
	(0.154)				(0.178)
HH index	0.144***				0.152***
	(0.053)				(0.051)
FIRST STAGE REGRESSION					
AP TICs pats. (t)		20.98		11.28	10.82
AP TICs pats. (t-1)			17.84	20.72	14.86
MSA*cat f.e.	YES	YES	YES	YES	YES
Cat.*Period f.e.	YES	YES	YES	YES	YES
Observations	3,521	3,521	3,521	3,521	3,521

Note: robust standard errors, clustered at MSA level, in parentheses. The dependent variable is the number of non-TIC patents. The unit of observation is the MSA-category-period combination. The time interval is 1988-1999, divided in three times period of four years each. The endogenous variable are TICs pats. (t) and TICs pats. (t-1). The excluded instruments are IV (t) (col. 2); IV(t-1) (cols. 3); IV(t) and IV(t-1) (col. 4-5). All variables are expressed in logarithmic form. All regressions are (analytically) weighted by the total number of patents over the period of analysis. The first-stage pane reports the Angrist and Pischke (2009) first-stage F statistics for tests of weak identification. *** p<0.01, ** p<0.05, * p<0.1

Table 7: The effect of TIC patents on other patents, level of aggregation: MSA-Subcategory

	(1)	(2)	(3)	(4)	(5)
Dep. var.			Non-TIC patents		
Method	OLS		IV - 2SLS		
TICs pats. (t)	-0.011	-0.088		-0.162	-0.297
	(0.018)	(0.180)		(0.220)	(0.305)
TICs pats. (t-1)	0.041***		0.047	0.091	0.059
	(0.011)		(0.119)	(0.151)	(0.181)
TICs other cats. (t)	0.061**				0.135*
	(0.029)				(0.074)
TICs other cats.(t-1)	0.049**				0.064
	(0.020)				(0.051)
Total empl.	0.754***				0.898***
	(0.253)				(0.303)
HH index	0.089*				0.117*
	(0.050)				(0.062)
FIRST STAGE REGRESSION					
AP TICs pats. (t)		13.93		8.128	5.567
AP TICs pats. (t-1)			19.01	11.75	14.10
MSA*subcat f.e.	YES	YES	YES	YES	YES
Subcat.*Period f.e.	YES	YES	YES	YES	YES
Observations	15,023	15,023	15,023	15,023	15,023

Note: robust standard errors, clustered at MSA level, in parentheses. The dependent variable is the number of non-TIC patents. The unit of observation is the MSA-subcategory-period combination. The time interval is 1988-1999, divided in three times period of four years each. The endogenous variable are TICs pats. (t) and TICs pats. (t-1). The excluded instruments are IV (t) (col. 2); IV(t-1) (cols. 3); IV(t) and IV(t-1) (col. 4-5). All variables are expressed in logarithmic form. All regressions are (analytically) weighted by the total number of patents over the period of analysis. The first-stage pane reports the Angrist and Pischke (2009) first-stage F statistics for tests of weak identification. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 8: The effect of TIC patents cited by local non-TIC

	(1)			(2)	(3)
Dep. var.			Non-TIC patents		
Method	OLS			IV - 2SLS	
TICs pats loc. cited . (t)	0.045***	0.011		0.119*	
	(0.016)	(0.020)		(0.067)	
TICs pats.loc. cited (t-1)	0.115***	0.093***	0.093***		0.198***
	(0.027)	(0.023)	(0.026)		(0.051)
Other TICs pats (t)		0.095*			
		(0.051)			
Other TICs pats. (t-1)		0.020	0.051*		
		(0.020)	(0.027)		
Other controls	YES			YES	YES
FIRST STAGE REGRESSION					
AP TICs pats. (t)				10.64	
AP TICs pats. (t-1)					11.02
Period f.e.	YES	YES	YES	YES	YES
MSA f.e.	YES	YES	YES	YES	YES
Observations	840	840	840	840	840

Note: robust standard errors, clustered at MSA level, in parentheses. The TIC patents variables include only patents which are cited by non-TIC patents authored in the same MSA. The time interval is 1988-1999, divided in three times period of four years each. The other controls are those listed in tab. 5.

Table 9: First-stage regression and falsification test

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.			TIC patents			
Level of aggregation	MSA		MSA-CAT		MSA-SUBCAT	
Estimation method	OLS	OLS	OLS	OLS	OLS	OLS
IV	0.973***		1.247***		0.933***	
	(0.289)		(0.320)		(0.305)	
Placebo IV		0.490		-0.092		-0.021
		(0.419)		(0.264)		(0.064)
Other controls	YES	YES	YES	YES	YES	YES
Period f.e.	YES	YES	YES	YES	YES	YES
MSA f.e.	YES	YES	NO	NO	NO	NO
MSA-Cat. f.e.	NO	NO	YES	YES	NO	NO
MSA-Subcat. f.e.	NO	NO	NO	NO	YES	YES
Cat.*Period f.e.	NO	NO	YES	YES	YES	YES
Subcat.*Period f.e	NO	NO	NO	NO	YES	YES
Observations	840	840	3,521	3,521	15,023	15,023

Note: robust standard errors, clustered at MSA level, in parentheses. The time interval is 1988-1999, divided in three times period of four years each. The placebo IV is defined in section 5.1. The other controls are those listed in tab. 5, 6, and 7.

Appendix A: Patent data and geographical assignment

Patent data come from the United states Patent and Trademark Office (USTPO) database as processed by the National Bureau of Economic Research (NBER), described in Hall et al., 2001. To the original dataset I add the inventors' unique identifier developed by Trajtenberg et al (2006) and the standardized assignee name available in Prof. Bronwyn H. Hall's website.[20] The latter, however, is not fully reliable as i) the complex ownership structure of companies may imply that differently named assignees correspond, in fact, to the same company, and ii) the same company name can be spelled in different ways (and the standardization routines cannot completely solve the problem).

I eliminate patents granted to inventors residing outside US and geolocated all the cities of residence of inventors through the ArcGis geolocator tool (based on the 2000 gazzetter of US places from US Census) and the Yahoo! Maps Ib Services. In the database both the address of the inventor(s) and the applicant(s) are reported, but the former is considered to be a more precise indication of the inventors' workplace, as often applicants' address refer to the location of the headquarter, which may not necessarily coincide with the location of the R&D lab. In the case where several authors are listed for the same patents and they live in different cities, the city of residence of the first author is chosen; this is a standard procedure in patent literature, and Carlino et al. (2007) show that the approximation is substantially innocuous, also considering that the majority of the patents are single-authored. A robustness test based only on single-authored patents, i.e. the part of the sample which should not be affected by a measurement error arising from an imprecise geographical assignment, produces results which are very similar to the baseline estimates. The geocoding operation was successful for 1,161,650 patents, which correspond to 97% of the database. I then assigned cities to counties using the ArcGis spatial join tool, and subsequently counties into MSAs (1993 definition).

[20] http://elsa.berkeley.edu/~bhhall/

Cahiers du GREThA
Working papers of GREThA

GREThA UMR CNRS 5113

Université Montesquieu Bordeaux IV
Avenue Léon Duguit
33608 PESSAC - FRANCE
Tel : +33 (0)5.56.84.25.75
Fax : +33 (0)5.56.84.86.47

http://gretha.u-bordeaux4.fr/

Cahiers du GREThA (derniers numéros – last issues)

2013-25 : LUSTOSA J. Maria Cecilia, FAURE Yves-André, *Changements institutionnels et politiques publiques au Brésil. Aux sources de la réduction des inégalités sociales, de la pauvreté et des disparités régionales*

2014-01 : BLANCHETON Bertrand, PASTUREAU Guillaume, *Le Mont-de-Piété à Bordeaux, les raisons d'un succès (1802-1913)*

2014-02 : FRIGANT Vincent, MIOLLAN Stéphane, *La restructuration de la géographie de l'industrie automobile en Europe durant les années 2000*

2014-03 : BLANCHETON Bertrand, *Les improvisations financière de la guerre de 1914-1918 en France. Les enjeux de la liquidité*.

2014-04 : ARNAUD Brice, *Extended Producer Responsibility and Green Marketing: an Application to Packaging*

2014-05 : CARAYOL Nicolas, DELILLE Rémi, VANNETELBOSCH Vincent, *Allocating value among farsighted players in network formation*

2014-06 : ARIFOVIC Jasmina, YILDIZOGLU Murat, *Learning the Ramsey outcome in a Kydland & Prescott economy*

2014-07 : BECUWE Stéphane, BLANCHETON Bertrand, *Relations internationales et discriminations tarifaires : le cas de la France (1850-1913)*

2014-08 : BECUWE Stéphane, BLANCHETON Bertrand, *La politique commerciale de la France et les filières sucrières de ses vieilles colonies sous le Second Empire*

2014-09 : FRIGANT Vincent, ZUMPE Martin, *Are automotive Global Production Networks becoming more global? Comparison of regional and global integration processes based on auto parts trade data*

2014-10 : BEN OTHMEN Asma, *De la mise à contribution des bénéficiaires au financement de la préservation des espaces naturels : tarification de l'accès ou augmentation de taxe?*

2014-11 : MENON Carlo, *Spreading Big Ideas? The effect of Top Inventing Companies on Local Inventors*

La coordination scientifique des Cahiers du GREThA est assurée par Emmanuel PETIT. La mise en page est assurée par Anne-Laure MERLETTE.