Papers in Evolutionary Economic Geography

17.01

When Buzz and Pipelines Fail

Christopher R. Esposito, David L. Rigby

Utrecht University
Urban & Regional research centre Utrecht

When Buzz and Pipelines Fail

Christopher R. Esposito and David L. Rigby

Department of Geography
UCLA
cresposito@ucla.edu

Abstract

Explanations for why some cities outperform others frequently rest on the assumed benefits of local and global interaction. Within the "buzz and pipelines" literature, the costs and returns to interaction have rarely been examined in formal settings. In this paper we extend research on knowledge sharing by modeling local and global interactions between firms distributed across city-regions. Our simulation model develops an evolutionary framework where firms explore and exploit knowledge sets that are accumulated over time by recombining technologies held by local and non-local firms. Our results make two contributions to the existing literature. First, we show why too much local interaction can induce technological lock-in and restrict cities' innovative growth. Second, we illustrate that non-local interaction entails opportunity costs that can outweigh its benefits. Together, the results unearth the conditions under which local and non-local interactions strengthen the economies of cities and when they fail to do so.

Keywords: regional economic growth, innovation, networks, computer simulation

JEL Codes: 033, R11, D83

Acknowledgements
The authors acknowledge helpful comments from Pierre-Alex Balland, Koen Frenken, and workshop participants at the 2015 Geography of Innovation conference (Toulouse, France), the 2016 Association of American Geographers annual meeting (San Francisco), and the 2016 Regional Economic Development and Innovation workshop (UCLA). The usual disclaimer applies.

Funding
This material is based upon work supported by the National Science Foundation Graduate Research Fellowship [DGE-1144087].

1. Introduction

The economic fortunes of cities and regions rest on the development of new technologies and on the spatial distribution of the resulting technological rents. Individual firms still dominate these processes, competing over the acquisition of knowledge and the translation of that knowledge into profitable technologies (Kogut and Zander, 1992; Teece and Pisano, 1994). Because the technological capabilities of firms diffuse and their competitive advantages erode, firms must continually balance their exploitation of known capabilities with their exploration for new ones (March, 1991). Recent models of invention suggest that new technologies emerge from the recombination of existing ideas (Jacobs, 1969; Kauffman, 1993; Weitzman, 1998; Fleming and Sorenson, 2001). In this respect, the size and the heterogeneity of knowledge pools are critical drivers of the pace and the direction of technological change (Rigby and Essletzbichler, 2006). Running throughout these claims are tensions between the production of ideas and their capture within the boundaries of individual firms, and the spillovers or managed flows of ideas that move between economic agents.

Within geographical settings these arguments often appear as network-based models of local and non-local interaction. Local interaction, or buzz, precipitated by dense clusters of economic agents connected within cities, is generally seen as an engine of innovation (Saxenian, 1994; Storper and Venables, 2004; Storper, 2015), though it may also be a harbinger of technological lock-in and decay (Grabher, 1993; Hassink, 2010). Non-local interaction, the global knowledge pipelines that connect economic agents located in different geographic regions, is typically regarded as a solution to technological stagnation, replenishing place-bound stocks of ideas and thus refueling economic growth (Bathelt et al., 2004).

Though broadly supportive of this general model, we are concerned with the imprecision of its current rendering. The benefits of buzz are typically imagined to rest simply on the number or the density of network partners, and the benefits of global pipelines of knowledge are generally assumed to unconditionally increase the innovativeness of their anchor regions. Too rarely are the effects of buzz and pipelines on technological heterogeneity discussed. Buzz and pipelines are believed to increase the innovativeness of regions by connecting their firms to more diverse pools of knowledge, but surprisingly little effort has been taken to see when, why, and how they might achieve this objective, and when they might fail to do so.

In this paper, we develop a simulation model to reveal when local pipelines and global buzz succeed and when they fail to improve the innovativeness of firms and their respective regions. We find that neither buzz nor pipelines are unconditionally good. Our model reveals that regions containing firms that share some knowledge outperform regions with independent firms, but that technological dynamism declines when firm boundaries largely dissolve. In other words, too much buzz can be harmful to a region’s health. We go on to explore the conditions under which knowledge flows, in the form of pipelines that connect actors located in distinct knowledge pools or cities, are beneficial. With one notable exception (Morrison et al., 2013), the story to date is that such flows yield positive returns. However, our model reveals that there are costs as well as benefits to cities or regions whose agents interact with partners elsewhere. The costs of pipelines may exceed the benefits in places with moderate local interaction. Together, these results show the conditions under which buzz and pipelines enhance innovation and when they may fail to generate the technological dynamism proponents suggest.

In the discussion below, we provide a brief sketch of an emerging literature that views technological change as driven by the interaction of multiple, competing agents. In turn, that vision is shown to extend debates on the nature of technological search, on exploitation and exploration, and on a recombinant model of invention. We review geographical extensions of these claims, largely located within the subfield of evolutionary economic geography, that highlight the spatially differentiated character of regional knowledge pools and of interaction within and across those pools. Thereafter, we change gears to discuss the structure of a simulation model developed to address the questions raised above. We first discuss the model informally and then shift to a more formal presentation. The results from the model are discussed at some length and we summarize the main findings in the conclusion.

2. Literature Review

2.1 A Recombinant Model of Technological Change

Invention, the process of knowledge creation, is increasingly imagined as recombination, of new technologies constructed by combining existing and sometimes new ideas. Such recombination might be born out of novel insight, though it may also result from a trial and error process as in Edison's search for an incandescent lamp filament or the stepwise hybridization of crop varieties (Evenson and Kislev, 1976). The vision of recombinant invention can be traced back at least as far as Isaac Newton who quipped that his insights came from "standing on the shoulders of giants." Within economics, recombination figures prominently in Schumpeter's (1942) model of creative destruction, of new technological possibilities emerging from and eliminating the old. It is made more explicit by Gilfillan (1935) who defines invention as "new combinations of prior art" in his exploration of the evolution of modern steamships from earlier forms of water-based vessels, and by Usher (1929) who likens invention to the "constructive assimilation of preexisting elements." Mokyr (1992) provides many examples of well-known technologies that combined previous ideas, from Crompton's spinning mule to the Jacquard loom. Within geographical analyses, this vision of invention extends back to Jacobs (1969), for whom the recombination of existing technologies within a city-region was key to the economic perseverance of Birmingham, England.

For Arthur (2007), technologies are combinations of components, subsets of knowledge that, once developed, can be repurposed in many different ways. This yields a mechanism to understand the evolution of technology in general. That there are limits on the possibilities of recombination is linked to the architecture of knowledge (Simon, 1962), itself controlled by the ease or complexity of coupling knowledge components. Kauffman (1993) uses these ideas to set the topology of technological landscapes, governed by the number of components that can be combined and by the interaction between those components. Processes of search across those landscapes are explored by Levinthal (1997), Rivkin (2000) and Fleming and Sorenson (2001). Recent empirical work by Ackcigit et al. (2013), Strumsky and Lobo (2015) and Youn et al. (2015) utilize technology class codes to examine the importance of recombination in U.S. patents from 1836. Weitzman (1998) formally integrates this notion of invention as the recombination of past knowledge in a model of endogenous growth. He shows that at early stages of economic development growth is bounded by the number of ideas that can be recombined. However, the recombinant knowledge stock grows much faster than the rest of the economy such that the number of idea combinations rapidly overwhelms the processing capacity of individual economic agents. Jones (1995) develops this last argument to explain why the rate of growth of productivity in the United States did not accelerate throughout the twentieth century.

2.2 Exploitation and Exploration

Alongside the metaphor of the "technology landscape," economic geographers have developed the concept of the "knowledge space" to capture the cognitive distance between different ideas and technologies (Kogler et al., 2013; Rigby, 2013). Nodes in the knowledge space represent different technologies or knowledge subsets and the distance between nodes is set by the cognitive proximity of those subsets. These ideas borrow from the product space of Hidalgo et al. (2007), from measures of the technological distance between the knowledge portfolios of firms (Jaffe, 1986; Teece and Pisano, 1994) and from attempts to distinguish and map technological fields (Engelsman and van Raan, 1994). Recombination entails assembling ideas from different parts of the knowledge space.

The growing heterogeneity of knowledge stocks and the rise of their complexity make the long-run management of technology increasingly difficult. Individual firms have limited capacity to effectively manage those stocks and so specialize in unique parts of knowledge space, organizing to exploit and extend the knowledge embodied within their workers and routines while guarding it from others (Kogut and Zander, 1992; Grant, 1996). As firms compete, so new technological combinations are discovered that rewrite the values of existing knowledge assets along with the fortunes of firms. A firm may occupy an area of the knowledge space associated with lucrative technological rents. But for Maskell and Malmberg (1999A), technological rents are eroded through time as patent protections expire, trade secrets diffuse and what was once tacit increasingly becomes codified. How do firms, and regions, maintain competitive advantage in this environment?

For March (1991), the fundamental answer to this question is the allocation of resources between exploitation and exploration. Local technological search or exploitation is typically focused on short-run improvements or refinements to existing practice that build incrementally on existing knowledge stocks and capabilities and that guide technologies along well-defined trajectories (Dosi, 1982; Stuart and Podolny, 1996). Distant search or exploration involves experimentation with new technological combinations unrelated to existing capabilities and that hold the possibility of disrupting established practice or reconfiguring the knowledge space (Christiansen, 1993; Rosenkopf and Nerkar, 2001). Thus, while exploitation is seen as enabling firms to capture the rents from knowledge production, exploration lowers the risk of competency traps and lock-in, especially around sub-optimal technologies, and offers the allure of long-run gains (David, 1985; Levitt and March, 1988; Arthur, 1989).

2.3 Interaction and Buzz

The organizational constraints of firms, including coordination and switching costs, limit the technological diversity they can internalize (Pavitt, 1999). At the same time, the growing complexity of technology means that recombination must link together subsets of knowledge that are drawn from different parts of knowledge space. Thus, technological change increasingly rests on the movement of ideas across the boundaries of individual firms. Ideas move more easily between firms in the same city than between firms in different cities (Jaffe et al., 1993; Sonn and Storper, 2008), though localized learning capabilities may be required to absorb knowledge spillovers (Lundvall, 1988; Malmberg and Maskell, 2006). More formal linkages may also be necessary to leverage the knowledge assets of network partners (Zucker and Darby, 1996; Owen-Smith and Powell, 2004). Though the precise mechanisms of knowledge flow are varied, there can be little doubt that the exchange of ideas in cities is a positive function of the number of potential interacting partners. Strong evidence for this is provided by Bettencourt et al. (2007). Of course, the quality of interaction mediates this simple message as different types of proximity influence the real density of partners (Boschma, 2005; Rutten, 2016). These ideas form the core of recent debates around buzz and pipelines.

2.4 Buzz

Storper and Venables (2004) argue that buzz, or the qualities of face-to-face exchange in dense agglomerations of interacting agents, is the reason why geography is not history in our increasingly integrated global economy. For them, face-to-face (F2F) interaction raises the efficiency of communication, especially when the information exchanged is tacit; it builds trust and thus the ease of interaction; it engenders screening and socializing functions that speed acquisition of shared values; and in performance, F2F enhances informational quality and boosts the efforts of partners. As F2F is limited to spatially bounded sets of agents, it gives rise to localized communities of practice (Brown and Duguid, 1991; Storper, 1997; Lawson and Lorenz, 1999; Gertler, 2003) that emerge alongside place-specific competences and capabilities (Maskell and Malmberg, 1999B; Rigby and Essletzbichler, 2006), often deepening the linkages between firms and regions (Schoenberger, 1999). Over time these systems of value and capabilities may lock regions into technological and institutional regimes that yield diminishing returns (Grabher, 1993), while in other cases they foster experimentation with more open knowledge architectures, labor market mechanisms and systems of entrepreneurship (Saxenian, 1994).

The presence of local buzz assumes that the boundaries of firms and other economic agents are to some degree porous. Firms must allow, consciously or not, the occasional idea to slip to their neighbors. Local buzz also assumes at least some heterogeneity in the knowledge stocks, organizational routines or institutional structures of the interacting partners. The openness of firm boundaries and the extent of regional knowledge heterogeneity are not independent of one another. If firm boundaries are very open, knowledge diffuses rapidly throughout the region and firms may adopt the ideas of their neighbors at the expense of inventing their own. Isomorphism, or overlap in the knowledge space, may reduce the heterogeneity of knowledge in the region and induce lock-in. This raises the question of whether or not there can be too much buzz, or too much openness within an economy. Insofar as openness encourages buzz, greater openness is a virtue. However, does more openness always lead to more buzz and heterogeneity and thus greater possibilities for recombination within the local/regional economy? This motivates the first question we seek to examine in this paper:

Question 1: Ceteris paribus, can local interaction be too intense, to the point that it reduces regional knowledge production?

2.5 Pipelines

An emerging effort seeks to better understand the mechanisms through which firms access knowledge (Bathelt et al., 2004; Gertler and Levitte, 2005; Maskell et al., 2006; Torre, 2008; Fitjar and Rodriquez-Pose, 2011; Morrison et al., 2013; Nomaler and Verspagen, 2016). They question the simple tacit-local and codified-global binary of knowledge flow and explore the conditions under which successful clusters source different types of knowledge within and between regions. At root, they argue that pipelines linking actors in different regions feed and reinforce local buzz, at least in part by diversifying knowledge subsets available to economic agents within regions. The knowledge types generated through pipelines are often more targeted than those available through buzz to the members of a local cluster, reflecting the cost of pipeline creation and maintenance. Empirical work supporting these claims includes Owen-Smith and Powell's (2004) study of the Boston biotechnology complex and Grabher's (2002) analysis of the London advertising sector.

Though Bathelt et al. (2004) and Morrison et al. (2013) explore some possible limits to the economies generated by both buzz and pipelines, we seek here to interrogate a more fundamental constraint. The focus of the model in Morrison et al. (2013) is to predict when and why an agent will interact with other local and non-local actors. While the results they generate are important for understanding the choice of individual agents to rationally engage in non-local knowledge search, the more fundamental question for geographers is whether local and non-local interactions are always beneficial for the regional ecosystem. Most firms are constrained by resources and thus it is reasonable to assume that the decision to interact with non-local firms will restrict local interaction (Bathelt et al., 2004). Pipelines then bear opportunity costs. Might the cost of reduced local interaction exceed the benefits of global interaction?

The answer to the above question varies across types of regions. There are three general cases of regions that have firms participating in pipelines. In the first, there is no local buzz: the boundaries of firms within regions are closed. In this case pipeline development will not benefit the regional ecosystem. Closed local firm boundaries prevent knowledge sourced by individual firms through non-local interaction from diffusing to co-located firms. While individual firms that participate in pipelines in no-buzz regions may benefit, the regional ecosystem does not, as found in Boschma and Ter Wal's (2007) empirical study.

In a second scenario, firm borders within a region are completely open, meaning that the technologies developed within each firm are available to all agents in the region. In this case, the cost to the region if one firm participates in pipeline activity will be relatively low. When a firm leaves a high-interaction region to participate in a pipeline, its home region will not miss out on much technological heterogeneity because other firms in the region have similar stocks of technologies. The opportunity cost of pipeline activity to the region is largely redundant knowledge. Instead, the pipeline will yield strong benefits: the high degree of local interaction will cause knowledge attained through the pipeline to diffuse rapidly through the cluster, enhancing its overall knowledge production. We would therefore expect regions with very strong buzz to enjoy significant returns from pipeline participation.

The third and most interesting case is when firm borders within regions are open but not perfectly so. In this situation, when a firm chooses to reduce engagement in local collaborative efforts in favor of non-local collaboration, the volume and diversity of local knowledge available for recombination within the region are decreased. The firm that engages in non-local partnership may be successful at extending its own knowledge assets, but the extent to which those assets are shared by firms across the home region and whether those assets are greater or less than those that might have been generated by purely local collaboration are unknown. The balance between the costs and benefits of pipelines are difficult to predict for cities with moderate buzz, which leads us to question 2.

Question 2: Do pipelines confer similar benefits on agents located in regions with different intensities of local interaction? For which strengths of local interaction do the benefits of pipelines exceed the costs, and for which strengths of local interaction do the costs exceed the benefits?

3. Informal Model Discussion

3.1 Introduction

We use a simulation model to examine the hypotheses and questions about buzz and pipelines in a rigorous and controlled setting. In this section of the paper we informally walk through the set-up of the simulation model, presenting a more formal explication of that framework in Section 4.

The general structure of the simulation model is relatively simple. We examine the costs and benefits of different forms of interaction that occur between groups of firms that are located in different city-regions. Firms are the only economic agents that we examine. Individual firms compete with one another by exploring for new technologies and exploiting the technologies that they develop. New profitable technologies generate monopoly rents. Firms prefer to exploit these profitable technologies rather than engage in exploration until the diffusion of knowledge increases competition and eliminates rents. This induces search for new technologies that result from the recombination of components of the firm's existing knowledge stock.

Firms interact locally and non-locally. Within each city-region, co-located firms may interact by sharing their individual knowledge stocks in a form of local buzz. This buzz does not flow outside the region. In our model, we vary the amount of buzz or interaction that occurs within each city-region and explore over a series of time-steps the total volume of knowledge produced within each. In a second model variant, one firm from each city is selected as a node in a global pipeline, modeled as a trade show, where it interacts with non-local firms and thus may access non-local knowledge subsets that it brings back to its home region.

We do not explore variations in institutional or organizational proximity, other than assuming that the borders of firms are more or less porous from one city to the next and so allow varying amounts of local knowledge sharing. More broadly, institutional and organizational factors may be regarded as fixed and thus do not influence the processes examined. We also take spatial proximity, the spatial clustering of firms into cities, as a fixed condition. Thus, the results of the simulation are driven solely by varying the propensity for firms to interact with one another, either locally or non-locally.

3.2 Firms

Firms are characterized by heterogeneous sets of knowledge that they use to produce various types of outputs. Figure 1 illustrates a firm that "knows" two types of knowledge, called Knowledge A and Knowledge B, and that has two copies of each of these two technologies. These knowledge sets serve as instruction manuals or routines that the firm follows to produce distinct outputs, in this case two units of Output A and two units of Output B.

Figure 1: Firm as a container of heterogeneous technologies

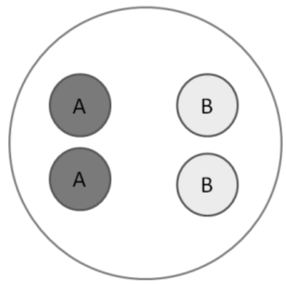

Firms from all regions sell their outputs in a single market to generate revenue. In order to focus the model on the supply-side dynamics of the economy, we assume that demand is equally robust for all types of outputs. Therefore, the prices of outputs are determined by supply conditions only. Outputs that are short in supply command high prices and outputs that are supplied in abundance command low prices. Monopolistic rents are the source of profit.

Each period, firms reinvest their revenue in search of rents. Firms that control technologies yielding high returns have large streams of revenue that they can reinvest to increase their productive capacity. Once a firm "knows" at least one A, it can attain more A through reinvestment; the firm does not need to invent more A organically. Firms that control technologies associated with low returns have minimal streams of revenue and must downsize production.

While the firms that control profitable technologies will want to continue to control, or exploit, these technologies, firms that control unprofitable technologies will not exploit their existing technological stock. Rather, these firms engage in a process of exploration, of recombining the different elements of their knowledge stocks in new combinations in the hope of discovering a new profitably technology. If the firm in Figure 1 finds that technologies A and B yield output that is less profitable than average, it recombines its existing knowledge types to produce technology AB.

In the model, a new time period begins after firms complete the adjustment of their productive capabilities, explore new technological possibilities, and produce new bundles of outputs.

3.3 Invention through Buzz

In Section 3.2 we considered how individual firms make decisions regarding the exploitation or exploration of their knowledge stocks. In our model, individual firms are distributed in even number across a set of cities where the level of local interaction or buzz varies. Therefore, where local interaction is at a maximum, firm borders are open and all firms in the city have access to exploring knowledge stocks of their local partners. Thus, when a firm in this city decides to explore with its unprofitable knowledge subsets, it may combine these knowledge types with all other knowledge variants that other firms in the same city choose to explore with. For firms located in cities with no local interaction, firm borders are closed and thus all knowledge recombination is fixed within the firm. For firms in cities with intermediate levels of interaction or buzz, some knowledge sharing occurs between the firms. No knowledge flows across city boundaries in our base, buzz-only model.

We develop the base, buzz-only version of our model to identify the optimal amount of local interaction that maximizes knowledge production, firm growth, and the technological heterogeneity within cities. Our single market contains five cities that are each endowed with five firms. In city 1 we exogenously maximize the amount of local interaction between firms by setting the degree of local interaction equal to one, allowing all firms access to the technologies of their neighbors. We endow the second city with slightly weaker interaction (local interaction strength = 0.75), where firms interact 75% of their technologies, while the remaining 25% of technologies do not spillover between firms.[1] The third city (local interaction strength = 0.5) interacts half of its technologies locally and the other half internally, and so on until the fifth city, which we endow with zero local interaction.

Figure 2 shows simplified examples of two cities with varying degrees of local interaction. Ties drawn between the technologies suggest that firm technologies are being recombined. In the first city, where the local interaction is perfectly strong, an equal number of ties are made between firms as within firms. In the second city, where local interaction is weaker, there are many more within-firm ties than between-firm ties.

Figure 2: Recombination in cities of varying degrees of local interaction

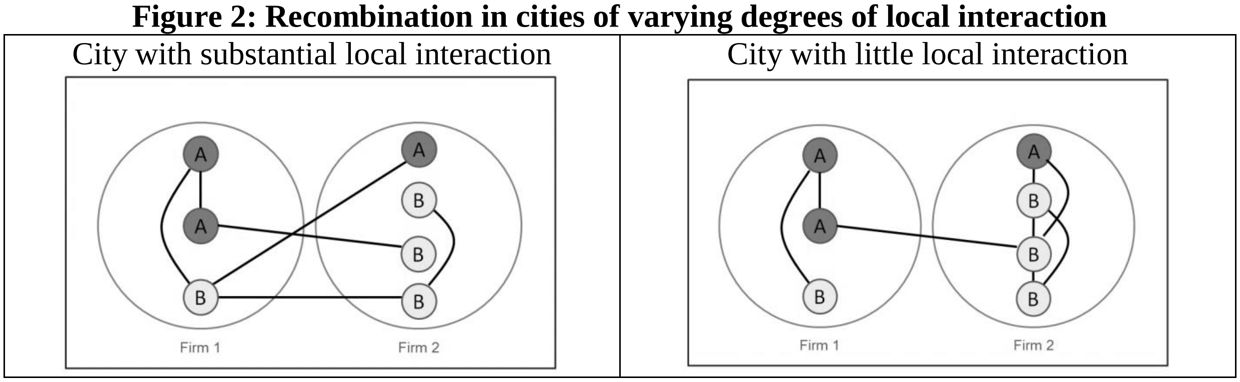

As we start the simulation model in period 1, each city is endowed with the same technologies that are distributed across their five firms. The simulation model is run multiple times for many periods and we examine the size of firms and how the firm size distribution varies across the cities with different levels of local interaction. The size of firms is given in terms of the number

[1] For example, a firm with a local interaction strength of 0.75 that is exploring with 10 A will interact 7.5 of its A locally and 2.5 of its A internally. Firms can explore with fractions of technologies.

of technologies they control. We explore how firm and city knowledge characteristics change over time.

3.4 Invention through Pipelines

We develop a second model variant in which we introduce global pipelines to explore how global interaction, or interaction between firms across cities, augments the relationship between local interaction strength and regional growth. In the pipelines model, a "tradeshow" is held late in each model run. We model pipelines as tradeshows because it is the simplest way to capture the effects of global knowledge transfer on regional knowledge heterogeneity (Maskell et al., 2006).[2] While the tradeshow is in session, one firm from each region leaves its home city to attend the tradeshow.

We assume that participation in the tradeshow occurs at the cost of local interaction such that pipeline participation is not free either to the firms directly involved or to their local partners. One fewer local partners reduce the volume of local knowledge sharing. Firms at the tradeshow share and learn from the other participants as though they were participating in a temporary economic cluster.[3] When the tradeshow ends, participants return to their home cities and resume local interaction. We depict a simplified tradeshow, with only two cities of two firms each, in Figure 3.

We run two simulation models. In the first, we study local interaction only. In the second model we add global interaction through pipelines to local interaction. Comparison of model results reveals how pipeline participation reacts to local interaction in different ways.

Figure 3: Recombination with pipelines

Before tradeshow All firms are in home city	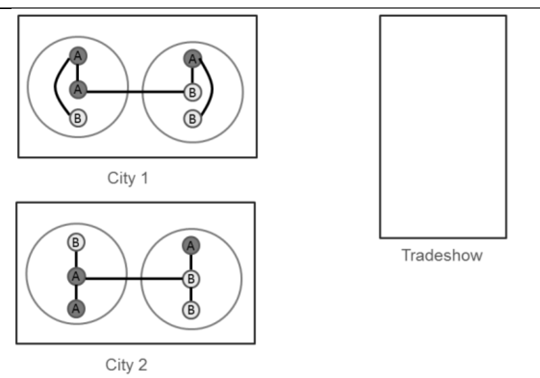

---

[2] While we acknowledge that the tradeshow representation of pipelines is limited, it is sufficient to achieve our objective of exploring the effects of global interaction on the knowledge stocks of firms and cities. Modeling global interaction as direct, peer-to-peer knowledge pipelines gives rise to complicated modeling questions about how firms choose to interact with specific partners and the extent of information they have to guide their decisions.

[3] We assign all firms participating at tradeshows medium strength local interaction, regardless of the local interaction strength of their home cities.

<table>
<tr>
<td>During tradeshow

One firm from each city joins the “temporary cluster”</td>
<td>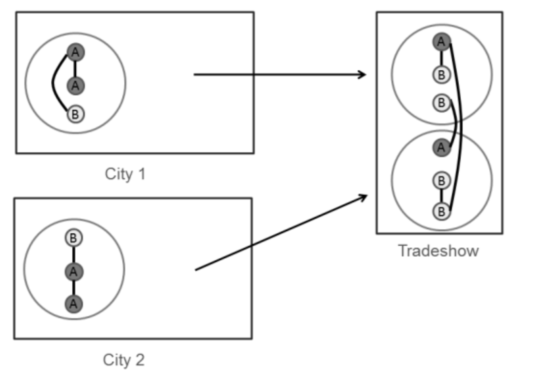
</td>
</tr>
<tr>
<td>After tradeshow

Firms return to home city and disseminate knowledge learned at tradeshow</td>
<td>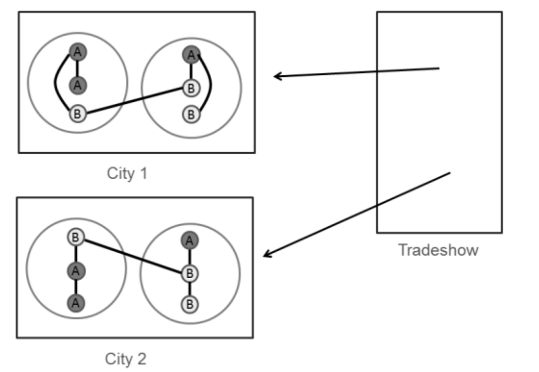
</td>
</tr>
</table>

4. Formal Model Discussion
4.1 Knowledge and Production of Outputs

We now introduce the specifics of the model. Each firm is described by a knowledge vector in which the elements record the amount of knowledge of each type that the firm controls. Knowledge types translate directly into an output. To produce output of type a, we assume the one-to-one production function so that firms compete entirely on basis of their knowledge portfolio. The production function is then given by

$$Output_{a,f} = Knowledge_{a,f}$$

Firms sell their outputs in the global market where they compete with the outputs supplied by all other firms. The price that firms receives for a given output is determined by the intersection of supply and demand for that output good. To simplify matters, we assume that demand is uniformly distributed across all outputs.[4] Therefore, the price of an output is determined by the quantity supplied by all firms.

[4] Relaxing the assumption of uniform demand requires a specification of how demand gets constituted, which requires a model of how consumer preferences evolve over time. By assuming uniform demand, we focus the model on the supply-side dynamics of the economy, as is the convention in evolutionary models of invention.

We depict the effect of supply on the prices of two outputs in Figure 4. The demand is identical in both markets, but the supply is much more abundant in the market in the right column. The abundant supply results in a relatively low market price for the output, denoted by P*.

Figure 4: Economics of ubiquity

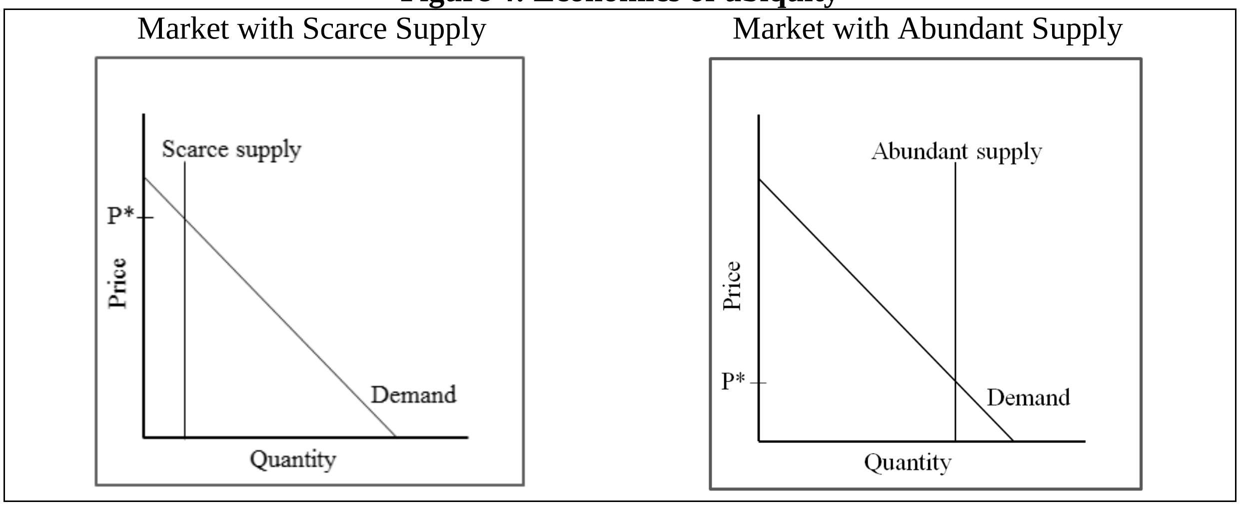

Because the supply of a given output determines its price, the price of output *a* is a decreasing function of its quantity supplied. We therefore write the price of *a* as its inverse quantity:

$$Price_{a,t} = \frac{1}{Quantity_{a,t}}$$

Firms can control more than one of each type of technology, so the total revenue generated by a given technology for a firm is given by the market price of the outputs scaled by the quantity of the technology controlled by the firm.

$$Revenue_{a,f,t} = Price_{a,t} * Quantity_{a,f,t}$$

Each period, firms reinvest their revenues in order to increase or decrease the quantity of knowledge that they hold. The scaling up of knowledge cannot be done for free. Firms must pay a cost to increase their knowledge-based productive capacity. Because we want the acquisition of novel technologies to drive competition in our model (rather than cost-based competition), we assume that the cost of production is constant across all goods, and given by the world average price level for all goods from the previous time period. This captures the idea that yesterday's outputs are today's inputs.

$$Cost_t = \frac{\sum_a^n Quantity_{a,t-1} * Price_{a,t-1}}{\sum_a^n Quantity_{a,t-1}}$$

To keep firm-level agency straightforward, we assume that firms reinvest their revenues generated by one type of technology in that same technology. Therefore, the quantity of *a* that firm f will control in the next time period is given by the revenue it generated from a divided by numeraire cost:

$$Quantity_{a,f,t+1} = {Revenue_{a,f,t}}/{Cost_t}$$

The above equation allows firms to adjust their knowledge structures based on market prices, which would lead to a convergence of prices across all outputs. We therefore introduce friction by making $Quantity_{a,f,t+1}$ a function of how much a the firm already knows. This amendment captures the fact that firms do not respond immediately to changes in the market.

$$Quantity_{a,f,t+1} = \frac{1}{2}\left(Quantity_{a,f,t} + Quantity_{a,f,t+1}\right)$$

4.2 Exploitation and Exploration

Firms continue using, or exploiting, technologies when they are profitable. Profitable technologies are priced above the numeraire cost. When technologies are no longer profitable, firms recombine them in a process of exploration or search for new hybrid technologies. Therefore, a firm will exploit Technology a if

$$Price_{a,t} \geq Cost_t$$

To invent, firms recombine their exploring technologies; a firm that explores with a and b will create the new technology, a-b. As firms continue to recombine technologies in later time periods, they can invent more complex hybrid technologies; a recombination of the technologies abb with bab, for instance, will produce the technology abb-bab. Moreover, the recombination of technologies follows a tree of knowledge that can grow infinitely tall (or deep), as shown in Figure 5.

Because recombination can result in a rapidly growing set of possible technologies, we must assume that each technology that a firm has chosen to explore with has a ρ (rho) probability of actually being exploited. While ρ can be interpreted as an exogenous, uniform degree of risk aversion, ρ's primary purpose is to curtail the amount of exploration that occurs in the model. Because the value ρ is constant across all cities and all firms, it does not bias the key findings in our results.

Figure 5: Tree of knowledge

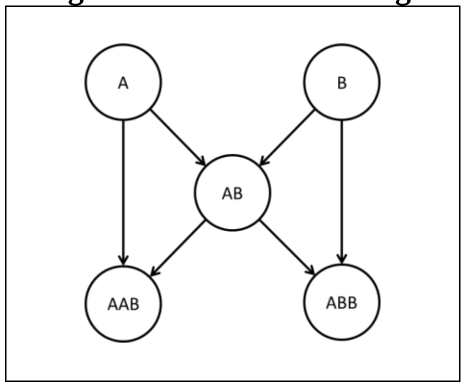

Lastly, we discuss the technical properties of how firms recombine, and the normalization methods we use to ensure that exploration does not increase the quantity of technologies. When a firm recombines, it produces all possible pairs between exploring technologies. For example, a firm exploring one a and one b explores with the Exploring Knowledge vector.

$$\text{Exploring Knowledge} = 1a;\ 1b$$

Recombining its a with b generates the Possible Pairs vector.[5]

$$\text{Possible Pairs} = 2ab;\ 1aa;\ 1bb$$

Note that the firm began with a quantity of two technologies but recombination resulted in a quantity of four technologies. We therefore normalize the Possible Pairs vector by dividing it by the number of potential partner technologies. That is, each exploring technology should produce exactly one new technology.[6] Dividing the Possible Pairs vector by the number of potential pairs technologies yields the following Realized Pairs vector:

$$\text{Realized Pairs} = 1ab;\ 0.5aa;\ 0.5bb$$

The realized Pairs vector contains two technologies (aa and bb), which have added no new variety to their base technologies, a and b. Because variety is needed to produce new technologies, aa is identical to a, and bb is identical to b. We therefore rewrite the Realized Pairs vector as

$$\text{Realized Pairs} = 1ab;\ 0.5a;\ 0.5b$$

The exploration process has concluded once firms arrive at the above Realized Pairs vector. The Realized Pairs vector represents the new technologies that the firm has added through exploration. Together with the old technologies the firm has retained through exploitation, the firm will use these technologies to produce outputs and generate rents in the next time period.

[5] We ignore the order of technologies, so B-A becomes a second A-B

[6] While each exploring technology must produce one new technology, it does not need to produce a single, whole technology. Exploring with one A, for instance, can produce 0.5 AB and 0.5 ABB.

4.3 Buzz and Pipelines

We exogenously fix the degree of local interaction in each city. The degree of local interaction determines the quantity of its technologies that firms will interact internally and locally. Local interaction strength (*LIS*) ranges from 0 to 1 and scales the quantity of technologies that firms interact internally and locally. If a firm explores with the Exploring Vector

$$\text{Exploring Knowledge} = 1a;\ 1b$$

it will divide its exploring technologies between two vectors, its Internal Exploring Knowledge vector and its Local Exploring Knowledge vector.

$$\text{Internal Exploring Knowledge} = [1a;\ 1b] * (1\text{-}LIS)$$

$$\text{Local Exploring Knowledge} = [1a;\ 1b] * (LIS)$$

The Internal Exploring Knowledge vector captures the technologies that the firm will explore with internally, following the example given in Section 4.2. The Local Exploring Knowledge Vector captures the technologies that the firm will interact with the knowledge of other co-located firms. These locally exploring technologies represent the firm's contribution to the regional buzz. A region with five firms will thus generate five Local Exploring Knowledge vectors, the pool of technologies that can be recombined locally. Local interaction otherwise occurs identically to internal interaction.

When a firm is participating in a tradeshow, *LIS* is set to the middle value, 0.5. All other aspects of the recombinatory process are preserved in tradeshow invention.

4.4 Collection of Results and Model Parameters

We run the base (without pipelines) and pipelines simulation models independently. We collect the output trace of the number of technologies in each firm after the simulation reaches the time horizon. We then sum the number of technologies across firms by city to calculate the resulting city size. We calculate the median city size across model runs to measure the average city size associated with each local interaction strength in the base and pipelines models. The model parameters and initial conditions are shown in Figure 6.

Figure 6: Model parameters and initial conditions

Parameter	Base Model	Pipelines Model
Number of Regions	5	5
Number of Firms per Region	5	5
Number of Firms in Trade Show	-	5
Initial Technologies of Firms	A, B, AB, AAB, ABB	A, B, AB, AAB, ABB

Initial Quantity of each Technology per Firm	10	10
Time Horizon	40	40
Trade Show Time Periods	-	25-30
Local Interaction Strength in Cities	[1, .75, .5, .25, 0]	[1, .75, .5, .25, 0]
Tradeshow Cluster Interaction Strength	-	0.5
ρ (Risk Aversion)	95%	95%
Number of Model Runs	5,000	5,000

5. Results

We present our model results in three sets of figures. Figure 7A contains overview results from our base model of buzz in cities; Figure 7B shows additional output from the base model that addresses identification issues. Figure 8A presents the overview results from the pipelines model, while Figures 8B and 8C address the identification issues. The final results table (Figure 9) juxtaposes output from the buzz and pipelines models to allow direct comparison between the two.

5.1 Results: Buzz without Pipelines

Figure 7A shows that average city size, an analog of the production of technologies, has an inverse-U shaped relationship with the strength of interaction that takes place within cities. Across multiple runs of our model, median city size peaks in cities where firms spill 80% of their technologies to neighboring firms and retain 20% of their technologies internally. Local interaction beyond this level is associated with less innovation and smaller cities. These results confirm Hypothesis 1: too much local interaction can impact a city negatively.

Why does a very high degree of local interaction reduce the innovativeness of cities? When the strength of local interaction is so high that it effectively removes the borders between firms, the knowledge of each firm within a region is available to all. In this situation, the knowledge stocks of local firms rapidly converge. Fruitful exploration requires the recombination of diverse knowledge types; thus when the strength of interaction is very high, the returns to exploration become increasingly limited.

Figure 7A: Buzz without pipelines

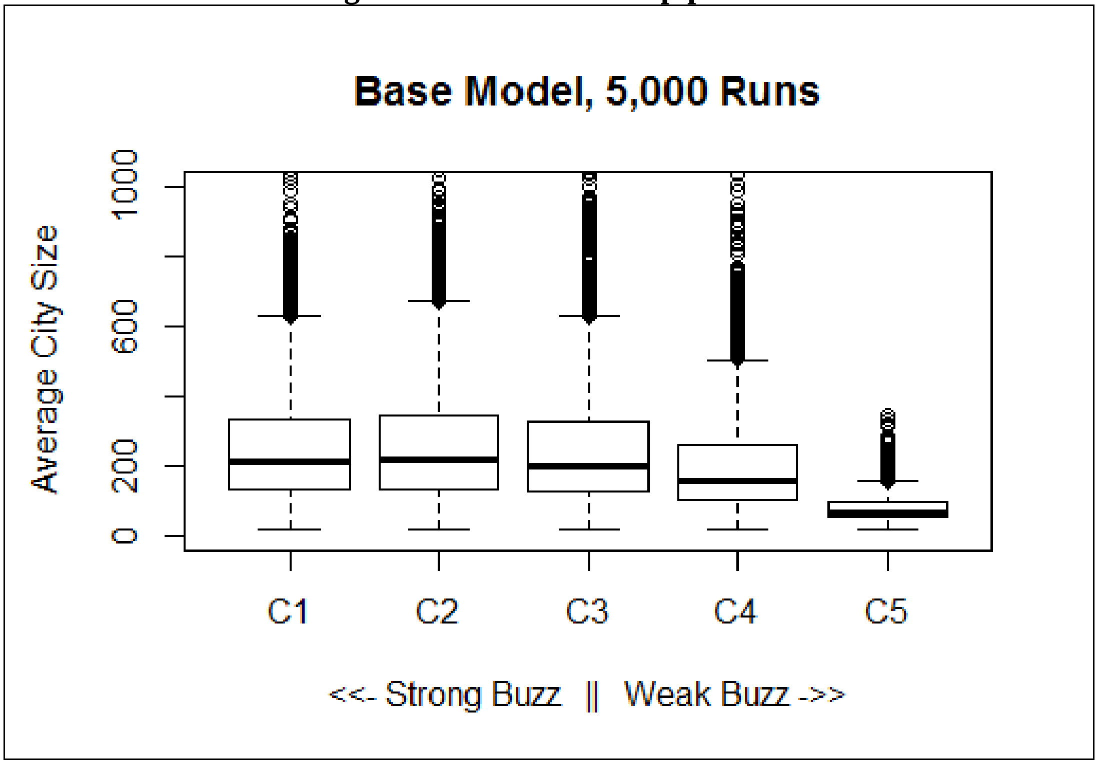

To examine this finding more closely, we develop a controlled experiment to test if very strong local interaction produces technological redundancy. We run a simplified model that contains two cities, each with two firms. The first city serves as a control and has its level of local interaction set to a moderate value (0.50) for the full duration of the model. The second city also begins with its level of local interaction set to 0.50, but we "treat" the second city by increasing the level of interaction between its firms to 1.0 after 25 time periods have passed. Each time period, we calculate the cognitive distance between the firms of each city. We are particularly interested in how the cognitive distance between the firms within each city separates following the treatment effect. Evidence of a post-treatment separation confirms our hypothesis that too intense interaction can produce knowledge redundancy.

We calculate the cognitive distance between firm i and firm j at time t in a given city using a standard distance formula, as depicted below. The distance between the firms is given by the sum of the distance between the quantity they control of each technology type, given by the range $a : n$. High distance values suggest high regional-level knowledge heterogeneity. We run this simplified model 15,000 times and average the results across the model runs for each time period.

$$Distance_{i,j,t} = \sum_a^n \sqrt{i_{a,t}^2 - j_{a,t}^2}$$

The results are presented in Figure 7B with both cities given identical technologies in period 1. The average technological distance between the firms of each region thus starts out the same. The number of technologies expands at a similar rate in each city until the treatment effect is applied to one city in period 25. The treatment, denoted by the star, changes the city growth dynamics. The treated city (bold line) shows less cognitive distance between its firms following the treatment: its firms develop a greater degree of knowledge redundancy.

Figure 7B: Controlled experiment of local interaction on heterogeneity

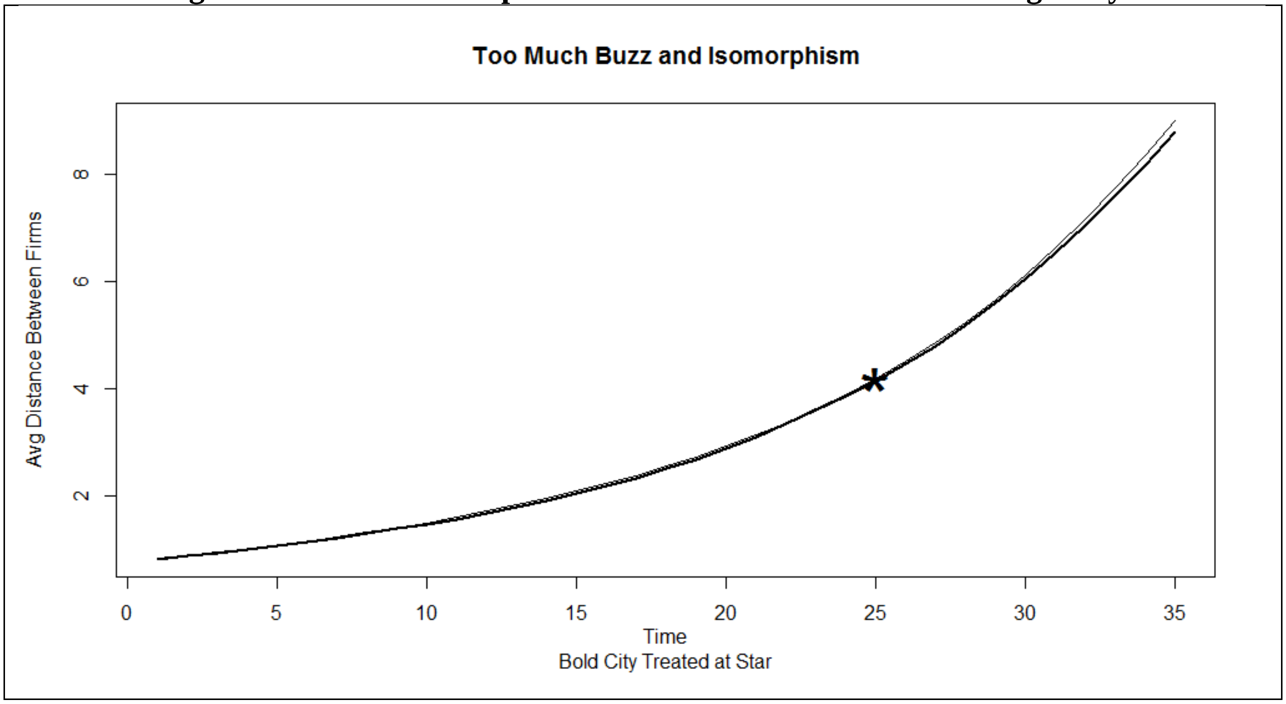

Similar findings, concerning the effects of close proximity in breeding knowledge redundancy, can be traced back to DiMaggio and Powell's (1983) work on isomorphism and have appeared in more recent and applicable research. For example, Uzzi's (1996) interviews with apparel manufacturers in New York City found that the firms with a high degree of network embeddedness were outperformed by their less embedded peers. Boschma (2005) also notes the potential for social ties that are too close to create cognitive lock-in, although to our knowledge the results presented in Figures 7A and 7B are the most systematic evidence to date that too much interaction can reduce knowledge variety.

5.2 Results: Buzz with Pipelines

We now explore the effect that global knowledge pipelines, modeled as tradeshows, has on the relationship between the strength of local interaction and regional dynamism. The results are constructed in a similar way to those of the base model. Figure 8A reveals that median city size increases continuously with the strength of local interaction when pipelines, or knowledge flows between firms located in different cities, exist. Inter-city knowledge flow enables the region with the highest level of local interaction to escape the trap of technological redundancy. The openness of the region allows it to rapidly diffuse the diverse knowledge that it sources from afar.

Figure 8A: Buzz with pipelines

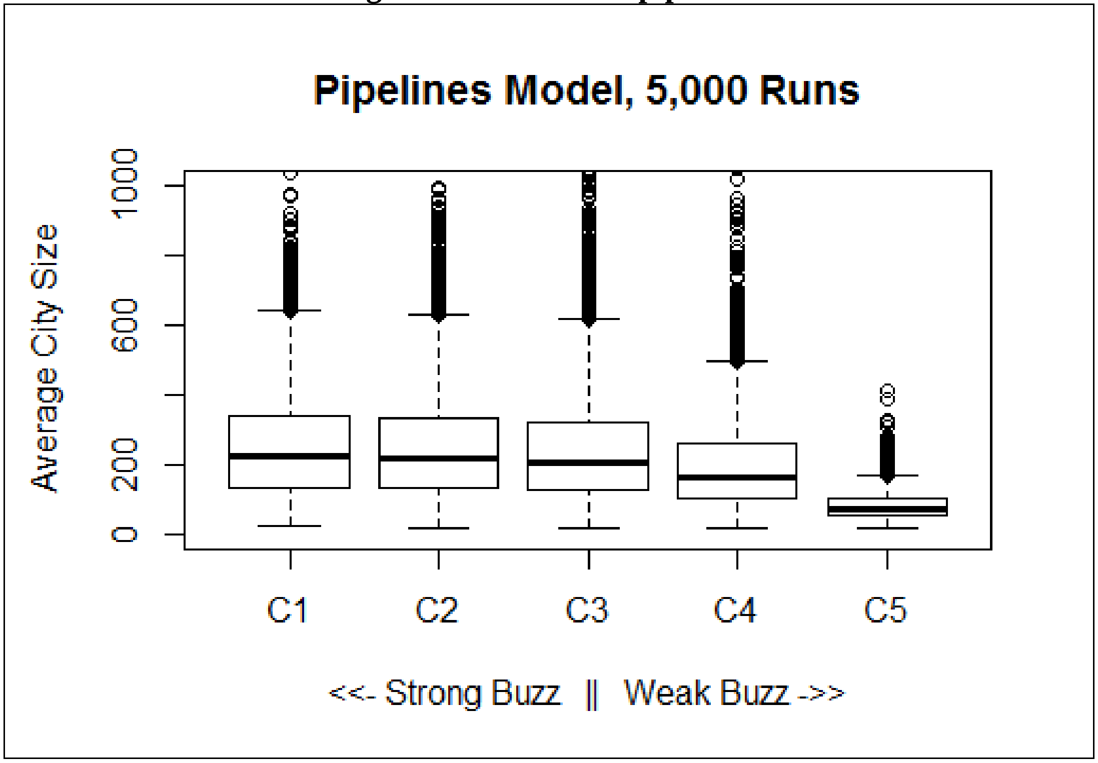

We perform a more thorough investigation of the effects of pipelines on firms and regions by looking at the size and the technological diversity of individual firms before, during, and after the tradeshow commences. Again, we develop a simplified version of our model, with fewer moving parts. The simplified model comprises three cities, each with three firms where the strength of local interaction is set at the level of 0.5. Moreover, the only parameter that varies across firms is whether or not they participate in a tradeshow. Two firms in each city do not participate in the tradeshow while the third one does.

Figure 8B reports the average changes in technological heterogeneity of tradeshow-participating firms and non-participating firms. For all firms, technological heterogeneity is measured for two time periods: between the beginning of the tradeshow (t = 24) and the conclusion of the tradeshow (t = 30), and between the conclusion of the tradeshow (t = 31) and the conclusion of the model (t = 40). Figure 8C reruns the tradeshow simulation model measuring the size of firms rather than their technological heterogeneity. The firm heterogeneity measure quickly responds to changes in the network connections of firms, while firm size responds more slowly to changes in firm heterogeneity. Comparing the results of Figures 8B and 8C helps to identify the time-dynamic influences of the tradeshow on the growth of firms.

Figure 8B: Pipeline effects on firm-level technological heterogeneity

		Firm Participated in TS? Yes	No
Time Period Being Measured	Beginning of Tradeshow to End of Tradeshow (t = 24:30)	0.17127	0.10073
	End of Tradeshow to End of Simulation (t = 31:40)	0.32113	0.32670

Figure 8C: Pipeline effects on firm-level size

		Firm Participated in TS? Yes	No
Time Period Being Measured	Beginning of Tradeshow to End of Tradeshow (t = 24:30)	1.1719	0.8305
	End of Tradeshow to End of Simulation (t = 31:40)	1.7699	1.5323

Figures 8B and 8C advance two stylized facts about how tradeshows affect the dynamics of firms. First, the results from the top row of both figures show that the non-participating firms in cities lose out when one of their firms leaves to attend a tradeshow. Firms that participate in tradeshows grow and gain heterogeneity much more rapidly than the firms that do not participate. Second, from the bottom row of Figure 8B, we see that the advantages of participating in a tradeshow end after the participating firms return to their home clusters; the tradeshow-participating firm adds 0.32113 to its technological heterogeneity while the firms that do not participate add 0.32670 to their technological heterogeneity. We believe that this result is driven by diffusion of tradeshow-sourced knowledge through the local cluster.

While the growth of technological heterogeneity equalizes across the firms after the tradeshow concludes, the size growth advantage persists for firms that participated in the tradeshow. We show this pattern in the second row of Figure 8C, where the firms that participated in the tradeshow grow by an average of 1.7699 technologies between the conclusion of the tradeshow and the end of the model simulation, while non-participants in tradeshows grow by only 1.5323 technologies. We believe that this result is driven by the time required to translate the acquisition of heterogeneity acquisition into firm size. Together, the tradeshow effects on firm heterogeneity and size indicate the costs and benefits of knowledge pipelines.

5.3 Results: Comparing Buzz to Buzz with Pipelines

Figure 9 juxtaposes output data from the base model of within city interaction with the pipelines model where interaction between cities is active. The figure explores the overall effect of tradeshows on cities with varying levels of buzz or interaction. We denote the median city size values of base model cities with circles and the median size values of pipeline cities with squares.

Figure 9: Costs and benefits of buzz and pipelines

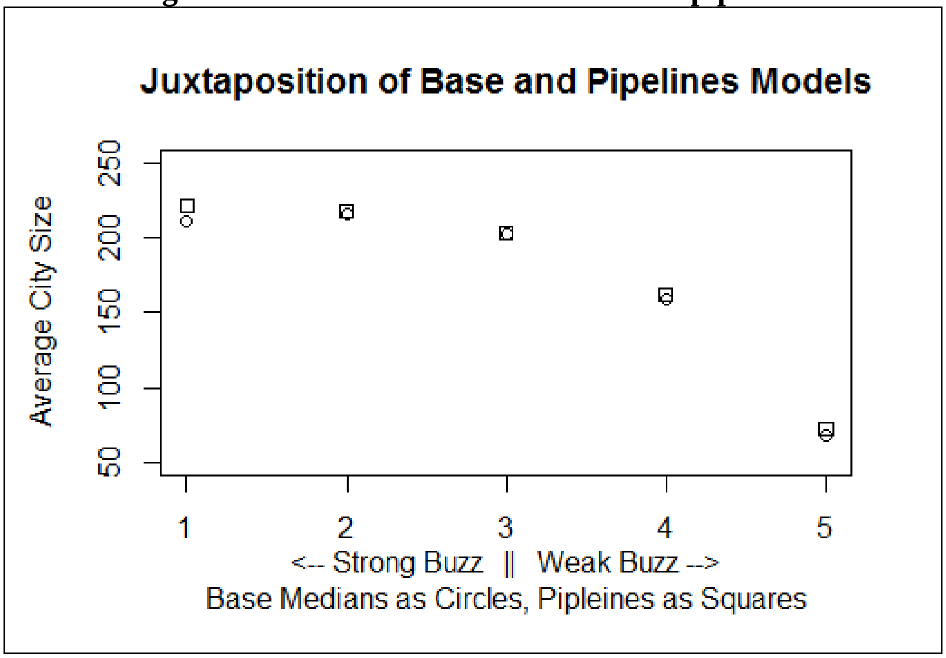

Working from the right to the left, we observe that when cities have no local interaction, they gain a slight advantage when a firm participates in a tradeshow. The interpretation of this result is straightforward. With zero local interaction, firms are completely isolated from one another and do not have the opportunity to participate in a cluster. The one firm that participates in the tradeshow will pull ahead. The square city outperforms the circle city because it is home to this fortuitous firm.

The advantages of tradeshows diminish as the strength of local interaction in cities increases. Two forces counteract the influence of tradeshows in this case. First, moderate local interaction allows the foreign technologies acquired through the tradeshow to diffuse through the local cluster. However, the tradeshow also works against the regional influence by removing one firm from the local cluster while the tradeshow is in session. For the duration of the tradeshow, the firms that remain in the local cluster have one less partner that they can interact with. At an intermediate level of local interaction, the costs of pipelines are equal to the benefits.

As the strength of local interaction increases further, the benefits of tradeshows again exceed the costs. In City 1, where knowledge flows freely across firms, tradeshow-sourced knowledge disseminates rapidly through the local cluster. City 1 loses relatively less technological variety when one of its firms leaves to participate in the trade show. Very strong local interaction and participation in pipelines are complimentary. They mitigate the costs of the other while amplifying the benefits.

6. Conclusion

Cities and regions contain the hard and soft infrastructures that enable local and non-local interaction and that drive the process of invention or knowledge production. In this paper, we provide a relatively formal rendering of invention that occurs within and across individual firms that are located in different spatial units. Firms in different regions are characterized by varying levels of local and non-local interaction. Through this model we examine the structure of costs and returns to invention at the level of the firm and the region. We find that increased local interaction and non-local interaction do not always have positive outcomes for regions. Local interaction within cities can be too great, and pipelines may generate opportunity costs that exceed the returns to non-local interaction. We confirmed the firm-level micro processes that give rise to these regional-level outcomes.

Of course, real cities do not operate according to a strict and singular logic. The particularities of places augment how buzz and pipelines operate on-the-ground, and the outcomes predicted by our model are not agnostic to its underlying parameters. One might imagine that a world with more cities, more firms, or more technologies would arrive at a different final state. The virtue of our model is not in the outcomes it predicts, but in the mechanisms it identifies. These mechanisms operate within and through cities, sometimes in the background and not always readily identifiable in empirically observable outcomes. Critical evaluation is required to assess if and how these mechanisms surface in real-world cities. We hope that the evidence presented in this paper will be useful in this effort.

Bibliography

Akcigit, U., Kerr, W., and Nicholas, T. (2013). The Mechanics of Endogenous Innovation and Growth: Evidence from Historical U.S. Patents. Working Paper, Department of Economics, University of Chicago.

Arthur, B. (1989). Competing Technologies, Increasing Returns, and Lock-In by Historical Events. *The Economic Journal* 99: 116-131.

Arthur, B. (2007). The Structure of Invention. *Research Policy* 36: 247-287.

Bathelt, H., Malmberg, A., and Maskell, P. (2004). Clusters and Knowledge: Local Buzz, Global Pipelines, and the Process of Knowledge Creation. *Progress in Human Geography* 28: 31-56.

Boschma, R. (2005). Proximity and Innovation: A Critical Assessment. *Regional Studies* 39: 61-74.

Boschma, R., and Ter Wal, A. (2007). Knowledge Networks and Innovative Performance in an Industrial District: The Case of a Footwear District in the South of Italy. *Industry and Innovation* 14: 177-199.

Brown, J., and Duguid, P. (1991). Organizational Learning and Communities-of-Practice: Toward a Unified View of Working, Learning, and Innovation. *Organization Science* 2: 40-57.

Christensen, C. (1993). The Rigid Disk Drive Industry, 1956-1990: A History of Commercial and Technological Turbulence. *Business History Review* 67: 531-588.

Cohen, W., and Levinthal, D. (1990). Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly* 35: 128-152.

David, P. (1985). Clio and the Economics of QWERTY. *The American Economic Review* 75: 332-337.

DiMaggio, P., and Powell, W. (1983). The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields. *American Sociological Review* 48: 147-160.

Dosi, G. (1982). Technological Paradigms and Technological Trajectories: A Suggested Interpretation of the Determinants and Directions of Technical Change. *Research Policy* 11: 147-162.

Engelsman, E., and van Raan, T. (1994). A Patent-Based Cartography of Technology. *Research Policy* 23: 1-26.

Evenson, R., and Kislev, Y. (1976). A Stochastic Model of Applied Research. *Journal of Political Economy* 84: 265-281.

Fitjar, R., and Rodriguez-Pose, A. (2011). When Local Interaction Does Not Suffice: Sources of Firm Innovation in Urban Norway. *Environment and Planning A* 43: 1248-1267.

Fleming, L., and Frenken, K. (2007). The Evolution of Inventor Networks in the Silicon Valley and Boston Regions. *Advances in Complex Systems* 10: 53-71.

Fleming, L., and Sorenson, O. (2001). Technology as a Complex Adaptive System: Evidence from Patent Data. *Research Policy* 30: 1019-1039.

Gertler, M. (2003). Tacit Knowledge and the Economic Geography of Context, or The Undefinable Tacitness of Being (There). *Journal of Economic Geography* 3: 75-99.

Gertler, M., and Levitte, Y. (2005). Local Nodes in Global Networks: The Geography of Knowledge Flows in Biotechnology Innovation. *Industry and Innovation* 12: 487-507.

Gilfillan, S. (1935). *Inventing the Ship*. Chicago: Follet Publishing.

Grabher, G. (1993). The Weakness of Strong Ties: The Lock-In of Regional Development in the Ruhr Area. In: *The Embedded Firm: On the Socioeconomics of Industrial Networks.* London: Routledge, 255-277.

Grabher, G. (2002). The Project Ecology of Advertising: Tasks, Talents, and Teams. *Regional Studies* 36: 245-262.

Grant, R. (1996). Toward a Knowledge-Based Theory of the Firm. *Strategic Management Journal* 17: 109-122.

Hassink, R. (2010). Locked in Decline? On the Role of Regional Lock-ins in Old Industrial Areas. In: Boschma, R., and R. Martin (eds.), *Handbook of Evolutionary Economic Geography*. Cheltenham, UK: Edward Elgar, 450-468.

Hidalgo, C., Klinger, B., Barabasi, A., and Hausmann, R. (2007). The Product Space Conditions the Development of Nations. *Science* 317: 482-487.

Jacobs, J. (1969). *The Economy of Cities*. New York: Random House.

Jaffe, A. (1986). Technological Opportunity and Spillovers from R & D: Evidence from Firms' Patents, Profits, and Market Value. *The American Economic Review* 76: 984-1001.

Jaffee, A., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics* 108: 577-598.

Jones, C. (1995). R&D-Based Models of Economic Growth. *Journal of Political Economy* 103: 759-784.

Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.

Kogler, D., Rigby, D., and Tucker, I. (2013). Mapping Knowledge Space and Technological Relatedness in US Cities. *European Planning Studies* 21: 1374-1391.

Kogut, B., and Zander, O. (1992). Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology. *Organization Science* 3: 383-397.

Lawson, C., and Lorenz, E. (1999). Collective Learning, Tacit Knowledge, and Regional Innovative Capacity. *Regional Studies* 33: 305-317.

Levinthal, D. (1997). Adaptation on Rugged Landscapes. *Management Science* 43: 934-950.

Levitt, B., and March, J. (1988). Organizational Learning. *Annual Review of Sociology* 14: 319-340.

Lundvall, B. (1988). Innovation as an Interactive Process: From User-Producer Interaction to the National Innovation Systems. In: Dosi, G., Freeman, G., Nelson, R., Silverberg, G., and Soete, L. (eds), *Technology and Economic Theory*. London: Printer Publishers.

Malmberg, A., and Maskell, P. (2006). Localized Learning Revisited. *Growth and Change* 37: 1-18.

March, J. (1991). Exploitation and Exploration in Organizational Learning. *Organization Science* 2: 71-87.

Maskell, P., Bathlet, H., and Malmberg, A. (2006). Building Global Knowledge Pipelines: The Role of Temporary Clusters. *European Planning Studies* 14: 997-1013.

Maskell, P., and Malmberg, A. (1999A). The Competitiveness of Firms and Regions: ‘Ubiquification’ and the Importance of Localized Learning. *Urban and Regional Studies* 6: 9-25.

Maskell, P., and Malmberg, A. (1999B). Localized Learning and Industrial Competitiveness. *Cambridge Journal of Economics* 23: 167-185.

Mokyr, J. (1992). *Levers of Riches: Technological Creativity and Economic Progress*. Oxford: Oxford University Press.

Morrison, A., Rabellotti, R., and Zirulia, L. (2013). When do Global Pipelines Enhance the Diffusion of Knowledge in Clusters? *Economic Geography* 89: 77-96.

Nomaler, O., and Verspagen, B. (2016). River Deep, Mountain High: of Long Run Knowledge Trajectories Within and Between Innovative Clusters. *Journal of Economic Geography*. Forthcoming.

Owen-Smith, J., and Powell, W. (2004). Knowledge Networks as Channels and Conduits: The Effect of Spillovers in the Boston Biotechnology Community. *Organization Science* 15: 5-21.

Pavitt, K. (1999). *Technology, Management and Systems of Innovation*. Cheltenham, UK: Edward Elgar.

Rigby, D. (2013). Technological Relatedness and Knowledge Space: Entry and Exit of US Cities from Patent Classes. *Regional Studies* 49: 1-16.

Rigby, D., and Essletzbichler, J. (2006). Technological Variety, Technological Change, and a Geography of Production Techniques. *Journal of Economic Geography* 6: 45-70.

Rivkin, J. (2000). Imitation of Complex Strategies. *Management Science* 46: 824-844.

Rosenkopf, L., and Nerkar, T. (2001). Beyond Local Search: Boundary-Spanning, Exploration, and Impact in the Optical Disk Industry. *Strategic Management Journal* 22: 287-306.

Saxenian, A. (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128.* Cambridge: Harvard University Press.

Schoenberger, E. (1999). The Firm in the Region and the Region in the Firm. In: Barnes, T., and Gertler, M. (eds), *The New Industrial Geography*. London: Routledge, 205-224.

Schumpeter, J. (1942). *Capitalism, Socialism, and Democracy*. New York: Harper and Row.

Simon, H. (1962). The Architecture of Complexity. *Proceedings of the American Philosophical Society* 106: 467-482.

Sonn, J., and Storper, M. (2008). The Increasing Importance of Geographical Proximity in Knowledge Production: An Analysis of US Patent Citations, 1975-1997. *Environment and Planning A* 40: 1020-1039.

Storper, M. (1997). *The Regional World: Territorial Development in a Global Economy.* New York: Guilford Press.

Storper, M. (2015). *The Rise and Fall of Urban Economies: Lessons from San Francisco and Los Angeles*. Palo Alto: Stanford University Press.

Storper, M., and Venables, A. (2004). Buzz: Face-to-face Contact and the Urban Economy. *Journal of Economic Geography* 4: 351-370.

Strumsky, D., and Lobo, J. (2015). Identifying the Sources of Technological Novelty in the Process of Invention. *Research Policy* 44: 1445-1461.

Stuart, T., and Podolny, J. (1996). Local Search and the Evolution of Technological Capabilities. *Strategic Management Journal* 17: 21-38.

Teece, D., and Pisano, G. (1994). The Dynamic Capabilities of Firms: An Introduction. *Industrial and Corporate Change* 3: 537-556.

Torre, A. (2008). On the Role Played by Temporary Geographical Proximity in Knowledge Transmission. *Regional Studies* 42: 869-889.

Usher, A. (1929). *A History of Mechanical Inventions*. New York: McGraw-Hill.

Uzzi, B. (1996). The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *American Sociological Review* 61: 674-698.

Weitzman, M. (1998). Recombinant Growth. *Quarterly Journal of Economics* 113: 331-360.

Weitzman, M. (1992). On Diversity. *Quarterly Journal of Economics* 107(2): 363-405.

Youn, H., Strumsky, D., Bettencourt, L., and Lobo, J. (2015). Invention as a Combinatorial Process: Evidence from US Patents. *Journal of the Royal Society Interface* 12: 1-8.

Zucker, L., and Darby, M. (1996). Star Scientists and Institutional Transformation: Patterns of Invention and Innovation in the Formation of the Biotechnology Industry. *Proceedings of the National Academy of Scientists* 93: 12709-12716.