

Is Specialization Good for Regional Economic Development?

Thomas Kemeny^{*} and Michael Storper[†]

Accepted to *Regional Studies*.

Please cite:

Kemeny, T., & Storper, M. (2014). Is Specialization Good for Regional Economic Development?. Regional Studies, 1-16. DOI:10.1080/00343404.2014.899691

Abstract: Debates about urban growth and change often center on specialization. However, arguments linking specialization to metropolitan economic development contain diverse, and sometimes conflicting, claims. Is it better to be highly specialized or diversified? Does specialization refer to the absolute or relative scale of an activity in a region? Does specialization have static or evolutionary effects? This paper investigates these questions in theoretical and empirical terms. By analyzing local agglomerations over time, we find that growing absolute specialization is positively linked to wages, while changes in relative concentration are not significantly associated with wage dynamics.

Keywords : regional economic development; specialization; agglomeration
JEL Codes: R11; R12; O21; O18

^{*} Corresponding author; Department of Geography & Environment, London School of Economics, Houghton Street, London, WC2A 2AE, UK; email: t.e.kemeny@lse.ac.uk; phone: 020 7955 6152

[†] Department of Geography & Environment, London School of Economics, Houghton Street, London, WC2A 2AE, UK; email: m.storper@lse.ac.uk; phone: 020 7955 6550

1. Introduction: The Fascination with Specialization

Discussions of urban growth and change often center on specialization. Urban planners, economic development authorities, consultancies and private businesses want to know about the prospects of metropolitan economies, and a principal way they do this is by assigning some kind of causality to patterns of industrial activity in the region. We often hear cities and metropolitan areas described in terms of their iconic activities, such as finance, high tech, logistics, services, or labor-intensive manufacturing. And such labels carry implicit value judgments. In recent years, the world's richest cities have been those whose economies contain concentrations of employment in information technology and finance. In developed countries, big manufacturing regions are in decline, in terms of their income rank and often in their population, while in the developing world, hubs of export-oriented labor-intensive manufacturing, such as Guangzhou, are said to have the secret to growth. Specialization is a principal way, then, that urban economies are viewed, labeled and classified by practitioners and policymakers, and it defines the public imagination about specific cities.

Specialization also features prominently in academic debates over economic development. Specialization and its flip side, diversification, are notions that apply to the tradable part of any economy. Although the majority of any economy – regional or even national – consists of the production of non-tradable goods and services, what the economy does in the tradable sector has strong effects on the overall level of regional employment and income. The tradable sector generates income that is spent on non-tradables in its “home market,” influencing wages in local-serving firms and industries in a variety of ways. The level of regional income is strongly influenced by specialization because a regional economy's external terms of tradeⁱ are set by its

tradable sector, and its overall level of output is influenced by tradables because demand for them is not limited by the producing region's income. A favorable specialization pattern (terms of trade and growth of external demand) is clearly good for the economy of the region. Evidence for the U.S. is suggestive: the bulk of national income growth between 1994 and 2000 was driven by large gains in just five of the country's 3,141 counties; these counties feature iconic clusters of tradable activity in information technology and financial services: Santa Clara, CA; San Mateo, CA, San Francisco, CA; King, WA; and Manhattan, NY (GALBRAITH and HALE, 2004).

In economic development circles, it has long been debated whether it is better for an economy to be diversified or highly specialized (HOOVER, 1948; RICHARDSON, 1968; QUIGLEY, 1998; BEAUDRY and SCHIFFAUEROVA, 2009). For our purposes, let us define a diversified region as one that contains a wide array of unrelated sectors in its economic base, with no specific sector dominating. As we shall see, translating such conceptual notions into precise empirical guidelines is challenging, but for the moment let us stick to the conceptual level.

Three justifications have been advanced for the virtues of diversification. The most common, for economic development professionals and some academics, is that diversification spreads the risk from economic fluctuations; this is the virtue of not putting all one's eggs in the same basket. Just as diversifying an individual's investment portfolio buffers against the volatility inherent in any single company's performance, so does the diversification of regional economic activity hedge against ups and downs in individual sectors (ATTARAN, 1986; KOREN and TENEREYRO, 2003). This argument is intuitively appealing, but since it is principally addressed to offsetting negative shocks, it does not consider whether diversification has opportunity costs, depriving an economy of benefits that could come from specialization.

A second, subtler argument for diversification holds that urbanization economies supply general inputs at efficient scales that are useful to many activities in a region. Therefore, a big metropolitan economy has reason to be diversified, and this will be reflected in its relatively high average total productivity. The major problem with this argument is obvious: diversification would be an outcome of being big; moreover, a city might have become big in the first place by being specialized. Another doubt comes from the nature of factor services supplied by urbanization economies: roads, infrastructure, and such, are the most general types of input to a modern economy. Beyond them, sectors need different and specific inputs (capital, labor, knowledge, supply chains). By definition, urbanization economies do not provide specialized resources dedicated to particular outputs; localization economies do, and localization economies are a force not for diversification, but for specialization.

A third argument for diversification concerns the dynamics of the regional economy. The idea here would be that a modern economy is a vast and very complex social division of labor. For an economy to move into, or capture, new activities, it needs to be able to draw quickly and easily from a shifting set of inputs and factors. This is a kind of “mix and match” view of the dynamics of economic development. A diversified economy might be able to do this better than a highly specialized one.

Table 1 provides an entry point into these complex relationships; it shows how levels of diversification vary with per capita income and employment. Taking data from the U.S. Census Bureau’s *County Business Patterns*, we calculate Herfindahl indices of concentration for Metropolitan and Combined Statistical Areas, where values approaching zero indicate more highly diversified regional economies, while a value of one indicates complete specialization in a single sector. The most detailed industrial data available is used: four-digit Standard Industrial Classification (SIC) codes for 1970, and six-digit North American Industrial Classification

System (NAICS) codes for 2009.ⁱⁱ These are combined with data on per capita personal income and employment data from the Bureau of Economic Affairs.

[TABLE 1 ABOUT HERE]

Table 1 shows that specialization levels for U.S. metro areas in 1970 are distributed in a fairly narrow arc, both in major cities as well as the average of all U.S. consolidated statistical areas. The largest regional economies are, of course, more diversified than the overall distribution of cities, but there is scant variation among large cities. Differences are even narrower in 2009. And yet the economies of these regions varied widely in terms of income levels, and growth of population and income. To take one example, Atlanta was the most diversified of the selected cities, while Los Angeles was the second most diversified. Los Angeles was nearly a quarter richer than Atlanta in 1970; since that time, Atlanta has nearly caught up to Los Angeles in terms of income levels, and its employment growth has dramatically outstripped that of Los Angeles. Meanwhile, San Francisco was much more highly specialized in 1970; its income grew considerably faster than both economies, while its employment base grew slower than both. And yet diversification levels in San Francisco, Atlanta and Los Angeles converge to quite similar levels by 2009. We therefore need to dig deeper, and we will now do so by thinking more about definitions of specialization, and then by testing some relationships empirically.

2. Relative or absolute specialization: Shares or size?

When making claims about specialization such as “New York is highly specialized in financial services,” or “Austin is ranked as the fourth most specialized U.S. metropolitan area in information technology,” the vast majority of reports and media buzz are referring to an industry’s employment share in the metropolitan economy. This is “relative” specialization. But specialization can also be thought of in absolute terms: having a particular activity may be the source of many jobs, or a high level of output, or large number of firms.

Absolute and relative concepts of specialization provide very different images of the economy. A small metropolitan area whose local employment base is dominated by work in a particular activity would rank higher in specialization than a large metropolitan area with a low share but a much higher absolute level of employment or output; the same is true in reverse.

Table 2 ranks U.S. metropolitan areas according to their relative and absolute specializations in a particular set of activities.ⁱⁱⁱ For exposition, we focus on information technology, but any tradable sector would do. The left column of Table 2 ranks regional economies based on the relative importance of the sum of employment in a set of 43 six-digit sectors that cover information technology activities, as defined by trade groups like Joint Venture Silicon Valley, as well as academic experts like Saxenian (1994).^{iv} The right column ranks them according to their absolute specialization in these same sectors, that is, on the basis of the actual number of workers they employ. The rankings are broadly different and there is only partial overlap of the two lists. In other words, one can generate very different images of “strong” and “weak” regions in different activities, just through this maneuver; and this is exactly what happens in the wide use of such rankings by academics, policymakers, and consultants.

[TABLE 2 ABOUT HERE]

The case of Los Angeles is instructive. Southern California hosts a large agglomeration of information technology, centered on Orange County. It is one of the nation's largest in absolute terms. Yet Los Angeles appears nowhere in the higher echelons of relative specialization (it ranks 31st among all metropolitan areas on this basis), and its location quotient is low. Although it is the fourth largest agglomeration in the U.S. – making it larger than those of celebrated clusters in Boston and Seattle – the hub of information technology concentrated in the Los Angeles region is rarely mentioned in discussions of U.S. high technology centers. Public (as well as much scholarly) debate, implicitly centered on relative, not absolute specialization, obscures this complex reality.

Of these two measures, the clearest theoretical case exists for specialization based on absolute size of the activity in the region. Increasing the size of a localized activity should positively affect productivity through the three main mechanisms specified by models of the New Economic Geography: sharing of input suppliers; matching of specialized labor demand and labor supply, especially in a context of high-turnover industries; and technological learning or spillovers, especially where innovation involves many different types of actors spread across different organizations (DURANTON and PUGA, 2004; ROSENTHAL and STRANGE, 2004).

By contrast, there is less consensus around whether having a high share of an activity would improve economic performance. Over the years, three principal notions have been developed that suggest that growing relative specialization will produce economic benefits. The first concerns competition between sectors for resources in the regional economy. Consider a regional economy with a sector that has a high share of regional employment and output. Due to

this footprint, the agglomeration will exercise a dominant role in regional demand for labor, land, infrastructure and other resources. This is descriptively plausible. Firms in any given industry might prefer not to have competition from other sectors in the local labor market. But the region might very well benefit from including other activities, even if they raise competition for factors and resources and even if they thereby ultimately drive out the dominant sector. Such diversification might stimulate movement up the ladder of technological sophistication and productivity and this would be better for regional development than remaining locked into its previous specialization. There is no general model that explains how relative specialization, by minimizing resource competition, would be systematically good or bad for regional economic development. Thus, upon closer examination, it does not provide much justification for the benefits of a narrow regional economic base.

There is a second, institutional, version of this argument. CHINITZ (1961) once proposed that dominant industries command the political attention of the region in which they are located. Contrasting New York and Pittsburgh, Chinitz suggested that the outcomes of this could be favorable if the industry is a promising or dynamic one, while it can be negative if it is not. Subsequently, MANCUR OLSON (1965) developed a more general theory of how interest groups capture attention, leading to “institutional sclerosis,” whereby the ability of institutions to reallocate resources to new domains of activity and functioning is diminished. Thus, if we borrow from Chinitz’s positive example, it follows that some forms of relative specialization could be helpful to a regional economy, because of the way they create dynamic industry groups. But if we borrow from his less positive example or more generally from the Olson hypothesis, relative specialization leads to elite capture and sclerosis.

These are obviously interesting and plausible theoretical notions. In political science, they have been tested in a number of policy-making areas, and are a major theme in large-scale

institutional theory as applied to long-term processes of national economic development (PERSSON and TABELLINI, 2002; GROSSMAN and HELPMAN, 2002; ACEMOGLU et al, 2001; ACEMOGLU and ROBINSON, 2008). To our knowledge, however, there has been no large-scale test of whether high levels of relative specialization at the regional scale lead to these mechanisms, and in turn whether they positively or negatively shape long-term adjustment of regional economies.

A third version of the relative specialization hypothesis can be drawn from recent debates in economic geography and what is known as the “new regionalism.” These discussions draw on theories of agglomeration. They explore the idea that an agglomeration of producers is simultaneously an interacting supply system; a local labor market matching system; and a context for knowledge exchange and spillover. But it is more than the sum of these parts: it is also a functioning ecosystem, tied together by many kinds of specialized economic agents, such as “dealmakers;” supportive local governments and associations; habits and soft conventions; and supportive inputs such as finance and R&D (STORPER 1997; MORGAN, 1997; FELDMAN and ZOLLER, 2012). It stands to reason that there is just so much room for these ecosystems in any given region, even in very big ones. This third hypothesis about relative specialization would then be that if a region wants to have these highly-performing ecosystems, it cannot simultaneously accommodate too many major ones.

No discussion of relative specialization would be complete without mentioning a commonly used applied version of it: the idea that a region is relatively specialized when an industry has a higher share in the regional economy than it does in the national economy. This concept, canonized in the location quotient, is an indicator in search of a theory. The strongest theory one can adduce in its support is the notion that there is a fixed external (national or international) demand for the output of a sector, so that if a region is specialized in a sector with

external demand that increases faster than the regional demand, then the specialization will be favorable to regional growth. But it can readily be seen that it offers no general predictions about whether a high location quotient will be good or bad for regional income or employment; that depends entirely on whether one specializes in a sector with high external growth or not. Evidently, this could go either way.

The academic literature exploring whether development is associated with either specialization or diversification presents academic perspectives on whether specialization should be understood in absolute or relative terms. DE GROOT et al (2009) survey this field, examining more than 25 peer-reviewed publications that present approximately 200 regressions using data drawn from 15 countries. However defined, the authors observe that specialization is very inconsistently associated with productivity, employment and innovation, with studies finding a wealth of positive, negative as well as nonexistent relationships. We explored how specialization was operationalized in the individual studies surveyed, and found that only six of 26 papers measured specialization in absolute terms; following GLAESER et al (1992) the majority proxied for specialization using location quotients, and secondarily other forms of relative specialization. None of the papers considered how relative and absolute conceptualizations of specialization might operate differently in relation to their chosen outcomes.

[TABLE 3 ABOUT HERE]

Table 3 summarizes the foregoing discussion of the various notions of specialization, and provides an overview of theoretical arguments as well as the evidentiary basis for each.

3. What goes together as a specialization? Relatedness in the economy

The central dilemma in understanding specialization is how to define a set of activities that “go together” so that we can consider them to be part of a specialization; and inversely, where to draw the boundaries between activities that will then be labeled “diverse” or “different.” This is a thorny conceptual and empirical matter that goes to the heart of work on the subject.

JANE JACOBS’ (1961, 1969) made what became canonical pronouncements about the virtues of diversification, but she did so without any precise definitions that would allow us to see whether she was thinking about serendipitous contact among similar (specialized) activities or diverse (different) ones. This blurriness has been picked up in recent literatures, where researchers argue for the virtues of economic “complexity” (HIDALGO and HAUSMANN, 2009), while others see cities as “nurseries” (DURANTON and PUGA, 2001), where firms can experiment with ideas and inputs from other activities, possibly recombining them to produce innovations that in turn spur regional development. If economies really do develop better over time through recombination (WEITZMAN, 1998), are they actually recombining inputs from sectors that are related, or at least close neighbors in terms of technology and underlying knowledge base, or are they recombining truly different, unrelated things, and hence benefiting from diversification? Everything depends on what we mean by different and diverse versus similar and specialized.

Along these lines, FRENKEN et al (2007) distinguish between what they call “related” and “unrelated” forms of diversity. They argue that a region’s long-run economic prospects for novelty are best when its industrial structure spans many distinct, but related product spaces. A variety of closely-related activities offer seedbeds for interaction, leading to gains in productivity

and innovation. This, in their view, ought to be better than having activities that are too distant from one another, because this excessive diversity inhibits recombination and “filling in the missing” product spaces. But these notions are highly sensitive to the theoretical language used to describe them. Specifically, FRENKEN et al (2007) choose to label a set of “related” activities a “related variety” (hence evoking diversity), and a set of “unrelated” activities an “unrelated variety” (a different form of diversity). Notice there is no term for specialization in these two configurations. Yet a group of activities that we define as being highly related should, by any logical extension, constitute a specialization of the regional economy. If they are related, it would have to be in ways that link their productivity, labor sharing, technological spillovers, and some kind of co-development dynamics through recombination and problem-solving to fill in the gaps in the regional supply structure. We have come full circle.

Moreover, whatever the definitions of specialization we use, a dynamic approach needs to address the question identified in the previous section: whether the current virtuous mix of sectors is a cause or an outcome of being previously diversified. The idea that specialization leads to a more complex industrial structure was suggested by GUNNAR MYRDAL (1956), and it has been revived in the New Economic Geography’s core-periphery model, which demonstrates how an economy that starts with successful specialization gets big and diversifies as a result of its economies of scale in consumption (its home market). Instrumental variables estimates produced using small- T panels, such as HARTOG ET AL (2012) are unhelpful in this regard – we need to look at a longer historical process of development in order to tease out how a virtuous complex specialization patterns evolves.

The ambiguities are not only conceptual, but empirical as well. The standard statistical categories for capturing specialization are supposed to group together activities that have similar outputs, and by virtue of this, would be based on similar production techniques and factor inputs.

In the United States, this is the idea behind the Standard Industrial Classification (SIC), and more recently, the North American Industrial Classification System (NAICS); it is also the logic that shapes the International Standard Industrial Classification (ISIC). With each system, different levels of similarity will be captured by the scale of aggregation used to perform the empirics of specialization, ranging from the highly-aggregated one-digit level that distinguishes manufacturing from wholesale activities and so on, to far more detailed six-digit industries.

Many academic articles, and most consulting reports, characterize specialization patterns using two- or three-digit industry codes. But the choice of aggregation or “granularity” is vitally important. For instance, consider two regions, A and B, each with large quantities of employment in apparel manufacturing (a three-digit NAICS sector). Examination at a more disaggregate scale might reveal that output in Region A is focused chiefly on low-cost T-shirts, while Region B does high fashion. When high levels of industrial aggregation are used, they appear to have comparable specializations, but in reality they are apples and oranges, in terms of labor demand, skills and wages, and unit prices.

This heterogeneity is suggested in the top part of Table 4, which compares wage levels in “Professional, Scientific and Technical Services” (NAICS 541) in two regions: Los Angeles and San Francisco. On average, workers in this industry in Los Angeles earn two-thirds the income of their colleagues in the Bay Area. But this category of “industry” contains graphic design, tax preparation and the design of computer systems – activities that evidently differ in important ways. So, one sensible reading of this comparison is that San Francisco and Los Angeles are specialized in different detailed tasks and subsectors within this broad activity area.

[TABLE 4 ABOUT HERE]

More detailed data is a logical solution, but this turns out to be not entirely the case. The lower part of Table 4 compares wages across Los Angeles and San Francisco within individual, six-digit information technology sectors - the most detailed industrial data commonly available. To ensure we are not examining small outliers, we confine our results to sectors in which both regions employ large numbers of workers. Interestingly, interregional wage gaps remain large at this more detailed level, and they are actually larger in “Computer Equipment and Software Merchant Wholesalers” (423430). Such wage variation could reflect differences in productivity within a subsector, but it is not implausible that Los Angeles produces outputs that can be meaningfully differentiated from those in San Francisco, using different techniques and factor inputs. Indeed, in studies on international trade and technological upgrading, researchers find considerable international variation in sophistication even using finely-grained 10-digit product-level data (SCHOTT, 2005; KEMENY, 2011).

And there may also be such a thing as too much disaggregation. To take an example, it seems sensible to jointly consider changes in specialization in such six-digit NAICS sectors as “Custom Computer Programming Services” (541511) and “Computer Systems Design” (541512). As we mention above, specialist industry groups like Joint Venture Silicon Valley do consider them to play parts within a singular coherent specialization. But if we address the issue of internal heterogeneity by defining industries using the greatest industrial detail, we arrive at another problem: we have now considered that each six-digit sector ought to exist within an entirely isolated silo, with no relationships to other six-digit industries.

It seems then, that an improved approach would seek to combine very detailed sectoral data into larger groupings reflecting substantive interconnections – our assemblage of six-digit

“information technology” sectors in Table 2 is an artisanal example of this idea, combining such sectors as “Semiconductor and Related Device Manufacturing,” (334413) and “Computer System Design Services” (541512) into something that better resembles our understanding of specialization in a set of related activities in information technology, despite the fact that, on the basis of their location in the classification system, these would be listed as industries with a great distance between them. For a large-scale application of this logic, however, we need an algorithmic method of capturing groups of industries that are strongly related through sharing, matching and learning. Some economic geographers and urban economists have experimented with approaches that address this issue, whether described via ‘industrial distance’ ‘product spaces’ and ‘related variety’ (ELLISON and GLAESER, 1997; FRENKEN et al., 2007; BOSCHMA and IAMMARINO, 2009, NEFFKE et al, 2011). Yet we lack a widely agreed upon method for distinguishing related from unrelated segments of the economy. The operationalization pursued in the related variety literature, unfortunately, mostly assumes the problem away by accepting the boundaries of three-digit sectors as demarcating “unrelated” activities, an assumption which we have shown to largely beg the central question of relatedness and hence specialization. Given this state of affairs, statements about specialization – descriptive, statistical, academic and nonacademic – should be interpreted with prudence; “league table” or rankings of hot spots should be taken with an even larger grain of salt.

4. Exploring the Relationship between Industry Specialization and Productivity

In a single paper, we cannot hope to empirically address all of the issues and questions that we have raised. But we can contribute to the process of empirical assessment by asking

whether absolute or relative specialization enhance productivity and wages, thus evaluating theories II-B and III in Table 3. Like most of the literature, we measure productivity using data on wages since they are the best available gauge of worker productivity (FELDSTEIN, 2008). And in the context of cities, evidence suggests that rising worker productivity is expressed in higher wage levels (COMBES et al, 2005).^v

A standard approach in the agglomeration literature links productivity to the relative or absolute size of a sector (and often a city). This approach predicts the wages of individual workers, as follows,

$$w_{ijk} = \alpha + \beta_1 S_{jk} + \beta_2 X'_i + \beta_3 C'_k + \varepsilon_i \quad (1)$$

where w represents wages for individual i in industry j and city k ; S indicates some index of industry specialization or agglomeration; X' describes a vector of individual characteristics, such as educational attainment, experience, gender etc.; C' is a vector of city-specific characteristics; and ε is an error term satisfying classical regression properties. Estimates of Equation 1 commonly use ordinary least squares (OLS) on large cross-sectional data like public-use samples of the Decennial Census of Population and Housing (for some prominent examples, SEE WHEATON and LEWIS, 2002; and GLAESER and MARÉ, 2001). This method offers some advantages, not least that such data cover large numbers of individuals.

However, this approach suffers from at least two major issues. The largest and most widely discussed problem is that of bias due to unobserved heterogeneity. While the available large, individual-level datasets commonly include a variety of wage covariates, they do not cover the full breadth of worker differences. Bias from this source could be large; for instance,

YANKOW (2006) finds that two-thirds of the city-size wage premium is due to unobserved worker differences.^{vi} Variation in wages could be due to specialization or they could instead reflect unobserved differences in worker ability or effort.

A second issue arises from the dearth of data on individuals over time that could be used in order to track the co-movement of specialization and wages. At its heart, any theory about the links between specialization and economic outcomes is about how changes in specialization patterns might produce changed economic circumstances. Unfortunately, such rich linked time-series data are, at best, extremely scarce.^{vii} Cross-sectional worker data simply do not allow us to shed light on dynamics. One sensible compromise is to use data offering repeated measures on industries in regional economies. Following this more feasible approach, we adopt the following model,

$$\bar{w}_{jkt} = \beta_1 \bar{w}_{jkt-1} + \beta_2 AS_{jkt} + \beta_3 RS_{jkt} + \beta_4 N'_{jk} + \beta_5 C'_k + \mu_{jk} + \eta_t + v_{jt} \quad (2)$$

where \bar{w} is the average wage for workers in industry j in city k at time t ; AS measures the level of absolute specialization for an agglomeration(*industry x city*); RS is the level of relative specialization for a given *industry x city*; N' is a vector of time-varying *industry x city* characteristics; C' is a vector of dynamic city-level characteristics; μ represents an individual *industry x city* fixed effect; η represents a year fixed effect, and v is the standard error term. Equation 2 also adds a one-period lag of the average wages in an agglomeration, since workers' wage levels are not set anew each year, but are instead anchored by the wages earned in the previous period. Just as an individual's wage is not annually renegotiated from a blank slate, average *industry x city* wages in the current year should be related to average wage levels from the prior year.^{viii}

Equation 2 explores how productivity levels in an agglomeration respond to changes in its relative and absolute levels of specialization. Taking a concrete example, our approach seeks to identify how the wages of workers in New York City’s financial services sector are influenced by changes in this agglomeration’s absolute size and relative footprint in the region. The *industry x city* fixed effect absorbs all stationary heterogeneity across agglomerations. Thus, it solves the serious analytical problem we described above that would plague cross-sectional studies, in which identification of a specialization effect depends upon a comparison between two regions’ agglomerations in industry X, ignoring relevant, if unobservable differences. Meanwhile, the year dummy variable accounts for unobserved time-specific shocks that exert uniform impacts across all *industry x city* units, such as business cycles. Equation 2 therefore offers a number of advantages over estimates of the impact of specialization on wages produced using the more common specification shown in Equation 1. First, Equation 2 accounts for a wide array of sources of spurious correlation, not least the problem of comparing industrial apples and oranges. It also exploits temporal dimensions of the data. Moreover, by confining the studied relationship to within-sector effects, we avoid having to consider an almost-unlimited number of other possible causes of inter-sectoral wage spillover effects. For these reasons, it ought to reliably gauge the relationship between specialization and productivity. Any result will be robust and provide conservative estimates of this relationship.

4.1 Data

Our primary data source is the U.S. Census Bureau’s *County Business Patterns*. *County Business Patterns* provides annual information about industries in individual counties. The data offer a number of attractive features. First, they are comprehensive: they provide details of every

industry in each county in the U.S. Second, because they are an annual series, they can be assembled and analyzed as a panel dataset. Third, they offer detailed industrial granularity, with industries defined at the 6-digit North American Industrial Classification System (NAICS) level. Fourth, they are released in a relatively timely manner, such that our analytical data run from the incorporation of the NAICS system in 1998, all the way up to 2010.

The data are not, however, without their own issues. They describe a small range of characteristics of regional agglomerations, chiefly payroll, employment and information about the number and size distribution of firms.^{ix} Moreover, their high degree of geographic and industrial detail means that it is difficult to supplement the minimal data with other information from external sources, since these supplementary data can not match their granularity. Such a small range of variables would be highly problematic in cross-sectional studies. However, using fixed effects, any stationary differences among industrial clusters are irrelevant to the analysis. This approach may not suit all research questions, but it is apt for an investigation into the responsiveness of productivity to changes in specialization.

The ‘regions’ to be studied are Metropolitan Areas, as defined by the Office of Management and Budget (OMB). OMB defines metropolitan areas to reflect functional social and economic integration as determined by commuting ties. *County Business Patterns* includes information on 292 metropolitan areas. The dependent variable in the forgoing analysis is the average annual wage income for workers in each *industry x city* agglomeration, derived by dividing total annual payroll in an agglomeration by the number of its employees. We measure absolute specialization as the number of employees in a local agglomeration. We calculate relative specialization as the share of employment in a local agglomeration in total metropolitan employment. As controls, we include total metropolitan employment. This indicates the breadth of overall agglomeration economies, which may be related to wages and productivity. Prior

research also suggests that its absence may bias estimates using measures of relative specialization (COMBES, 2000). Because of evidence indicating that industry productivity is partly a function of the distribution of the sizes of its constituent firms (ACS et al, 1999, PAGANO and SCHIVARDI, 2003), we also include an indicator of average industry firm size.

As we discuss above, defining the boundaries of a specialization is tricky for several reasons. First, there is the problem of granularity: if the boundaries of an industry are defined too narrowly, then we will miss changes in specialization that involve related sectors. Conversely, if industry definitions are too broad, then changes in employment will include many unrelated activities. Second, industrial classes such as NAICS are defined on the basis of output, ‘adjacent’ industrial classes are not always functionally integrated or more involved in sharing, matching and learning than are sectors classified as distant. Addressing this second issue lies beyond the scope of this paper, though clearly more work needs to be done to deal with this problem. We do try to address the problem of granularity through sensitivity analysis. Though we present ‘main’ estimates using industries defined at the four-digit level, we complement our findings at this scale with results produced defining sectors at 2-, 3- 5- and 6-digit levels.

Rather than estimating the impact of changes in specialization in the full range of sectors in the economy, we focus on tradable sectors for the reasons discussed in previous sections. Following JENSEN AND KLETZER (2006), we identify tradable industries as those that are not geographically ubiquitous, and by contrast spatially ubiquitous sectors are non-tradables. The following Herfindahl index of geographical concentration is constructed for each four-digit sector:^x

$$Conc_j = \sum_{k=1}^K \left(\frac{e_{jk}}{E_j} \right)^2 \quad (3)$$

where e measures employment in industry j and city k ; and E is total employment across all cities in industry j . Industries with Herfindahl values near zero will be those that exhibit a uniform distribution over space, while Herfindahl values closer to one indicate sectors where activity is highly concentrated in only a few locations.

As with Jensen and Kletzer, we must choose a cutoff point in the distribution of concentration values at which tradable activities are distinguished from non-tradables. There is no clear theoretical guidance on such a cutoff. By closely examining the data, we settle on a cutoff point of 0.036. Industries with Herfindahl values below 0.036 conform to our expectations regarding industries that ought to be non-tradable: retail stores, death care services, car repair, warehousing, architectural services, machine shops and other general purpose machinery manufacturing. Meanwhile, industries with index values above 0.036 seem likely to be tradable, such as motor vehicle parts manufacturing, software publishing, electric lighting equipment manufacturing, and pipeline transportation of crude oil. While the precise location of this cutoff is not theoretically-derived, in practice it sensibly differentiates non-tradable from tradable sectors.

4.2 Results

Results reported in the first three columns of Table 5 are estimated using pooled ordinary least squares, for exposition purposes. The final model uses a different estimation technique and represents our best estimate of the relationships of interest. Year fixed effects are included in all

models in order to account for economy-wide time-specific shocks; coefficients for these dummy variables are not reported.^{xixii}

Model 1 estimates a simplified version of Equation 2 in which relative specialization is the sole specialization measure; Model 2 does the same using only absolute specialization. Relative and absolute specialization are related by construction, though they are only moderately correlated ($\text{corr}=0.34, p=0.000$). This is because metropolitan employment, which is the denominator of the relative specialization measure, is influenced by a host of factors unrelated to the dynamics of individual industrial clusters. Diagnostics performed on OLS estimates, such as the variance inflation factor (VIF) test, indicate no problems of multicollinearity among these or other variables. Nonetheless, our initial two models focus on each specialization measure separately. In pooled cross-sectional models, both measures are positively and significantly related to average wages when they alone indicate specialization. We can interpret Model 1 as indicating that industries that occupy larger shares of their regional economy also pay higher wages, while Model 2 shows that urban industries that employ larger numbers of workers tend to pay higher wages. In Model 3 we include both aspects of specialization at once. Though magnitudes of the coefficients for each specialization measure decline somewhat, both remain positively and significantly related to average *industry x city* wages. Hence an initial interpretation of these results would say that New York's finance workers earn more than their counterparts in Los Angeles both because Wall Street employs more workers, and because it agglomeration occupies a larger share of overall employment in New York than the same industry does in Los Angeles.

However, these preliminary results ignore four important econometric considerations. First, as we discuss above, for the purposes of identification, it makes sense to utilize repeated observations on *industry x city* units. The OLS models pool together all *industry x city x time*

observations, but do not recognize the temporal relationships within *industry x city* units. By exploiting the time dimension, we can incorporate dynamics while permitting fixed effects estimation that shifts the examined relationship to one occurring within groups. Taking a fixed effects approach, we can model how wages in a particular local agglomeration change in relation to changes in specialization over time in that unit.

Second, given the likelihood that average wages depend on previous wages, it is desirable to include a lagged iteration of average wages on the right side of the equation. In the context of the OLS models (1) – (3), we opted not to do so given well-documented issues of bias in that context (ACHEN, 2000; KEELE and KELLY, 2006). Even in a panel setup, dynamic panel bias is a widely discussed problem. The standard solution is to apply some form of the Generalized Method of Moments (GMM) estimator (BOND, 2002; ARELLANO and HONORE, 2001). In addition to being apt in the presence of an autoregressive dependent variable, this class of model is also suitable for large- N , small- T panels such as the one at hand. For this reason, rather than applying the standard fixed effects estimator to equation (2), we estimate the model using two-step GMM-FE.

[TABLE 5 ABOUT HERE]

Endogeneity, and specifically bias from reverse causation represents a third potential estimation issue. While theory predicts a causal relationship running from specialization to productivity, rising wages and productivity could also stimulate changes in specialization. Employment in sectors with rising wages may grow in absolute and relative terms as workers shift from other locations, as well as from other industries in the same city. Both indicators of specialization are potentially endogenous in this regard. Lacking ready access to randomized control trials, we address the problem of bias due to endogeneity using instrumental variables

techniques. The GMM estimator is useful in this respect, as it provide methods of incorporating lagged regressors as instruments. We also add a ‘substantive’ instrument for absolute specialization, adapting a shift-share approach that CARD (2001) applies in the context of the economic effects of immigration. We calculate the ‘predicted’ size of employment in a region’s industry in time t on the basis of its size in period $t-1$ and the overall national industry growth rate between $t-1$ and t . Industry-specific national historical employment growth rates are given by:

$$g_{jt-(t-1)} = \left[\left(\frac{e_j}{E} \right)_t - \left(\frac{e_j}{E} \right)_{t-1} \right] / \left(\frac{e_j}{E} \right)_{t-1} \quad (4)$$

where g_j is the growth rate in employment e for industry j in the national economy with a total employment of E between t and $t-1$. Given these growth rates, the shift-share ‘predicted absolute specialization’ index $\overline{\overline{AS}}$ is constructed as follows:

$$\overline{\overline{AS}} = e_{jkt-1} \left[1 + (g_j)_{t-(t-1)} \right] \quad (4)$$

Since current wages can determine neither prior levels of employment in a local agglomeration, nor historical national industry employment growth, this index is a potentially useful exogenous source of variation. We discuss its effectiveness in the current context below.

Serial autocorrelation represents a fourth and final estimation problem, one which could bias standard errors. We detect the presence of serial autocorrelation in the panel data using a test created by WOOLDRIDGE (2002).^{xiii} We apply the standard Newey-West approach that uses the Bartlett kernel to produce heteroscedasticity- and autocorrelation-consistent (HAC) estimates. In initial work, we explored bandwidths from 2 to 5 and found consistent results in each case. For brevity, we present findings estimated with a bandwidth of 2.

Model 4 addresses these four econometric concerns; it is a fixed effects model with lagged as well as substantive instruments for potentially endogenous regressors, estimated using two-step GMM with HAC covariance estimation with a bandwidth of 2. Together, these methodological choices ought to produce efficient estimates of the coefficients and standard errors, while strengthening confidence on the direction of causality in the observed relationship, while also accounting for dynamic panel bias and serial autocorrelation. We estimate the model on over 20,000 local *industry x city* agglomerations. Due to the shift in estimation strategy from OLS to FE, the magnitudes of coefficients in Model 4 are substantially different from those obtained in Models 1–3.

Model 4 shows that absolute specialization is positively and significantly related to wages. The coefficient on this variable suggests that, as employment in a local agglomeration grows by 100 workers, average annual wages in that cluster will rise by around \$29. This seems fairly modest, but it is worth considering that this effect is larger than the overall urban agglomeration effect: with a coefficient of 4.32, a similar increase in urban population will augment wages by only \$0.43. Interestingly, after accounting for the temporal dimension of the data, relative specialization is not significantly related to wages (and its coefficient changes sign). In fact, over a very wide variety of fixed-effects estimates, ranging from those with no instruments and lagged dependent variables to fuller models with all of the characteristics accounted for in Model 4 absolute specialization is uniformly positive and significant, while relative specialization is insignificant.^{xiv} The striking differences between cross-sectional and panel results points to the need to revisit the findings of prior studies that do not explore temporal dynamics.

The lower panel of Table 5 displays diagnostics of the instrumental variables. Specifically, the first-stage F statistic is far above the threshold value of 13.43, suggesting that we can conclude that our instrument set is not weak. The Hansen J value indicates that at least one of our

instruments can be treated as endogenous. These results increase the confidence that the direction of the observed relationship goes from specialization to wages and not the other way around.

[TABLE 6 ABOUT HERE]

We explore the sensitivity of the results in several ways. To boost confidence that the four-digit level provides a reasonable basis for making claims about specialization, we estimate Equation (2) using two-step GMM-FE for 2-, 3-, 5- and 6-digit NAICS industries. This necessitates re-examination of the distinctions between tradable and non-tradable sectors at each level of industrial granularity, which we determine again by exploring how different thresholds produce more and less plausible groups of tradable and non-tradable industries.^{xv} Table 6 presents estimates of Equation (2) at these different levels of granularity. For estimates produced using 3- and 5-digit sectors, absolute specialization is positive and significantly related to average wages; relative specialization is significant at the three-digit level. In estimates produced with 2- and 5-digit industries, neither measures of specialization are significantly related to wages. We take this as evidence that results cohere around the 4-digit level. This is not to say that 4-digit is the intrinsically correct scale at which to measure specialization, but rather that changes in the absolute scale of moderately detailed sectoral classes appears consistently associated with rising wages, whether we somewhat loosen or tighten what constitutes ‘moderate.’

[TABLE 7 ABOUT HERE]

We also explore how our results may be sensitive to the range of cities included in the analytical sample. Our baseline sample of 281 Metropolitan Areas covers most of the population

of 366, and includes all cities of a reasonable size, and most smaller ones. However, the effects of specialization could work differently in different parts of the urban hierarchy. To explore whether this may be true, we re-estimates Equation (2) for the 100-, 150- and 200-largest cities by population, as well as for the 200-smallest cities. Table 7 displays the results, which suggest that the positive link between growing industry employment and rising wages applies not just for the entire distribution, but in a similar fashion for the largest and smallest cities.^{xvi}

5. Conclusion: specialization and the dynamics of economic development

Consistent with theories of agglomeration under which the scale of an industry augments productivity through the mechanisms of sharing, matching and learning, we find a robustly significant positive relationship between absolute specialization and wages. In careful dynamic estimates, the relative footprint of an industrial specialization in broader regional employment, i.e. relative specialization, is not significantly associated with wages. This insignificant relationship stands in contrast to results obtained using cross-sectional, between-industry approaches, perhaps because our method eliminated a lot of the noise (unobserved heterogeneity) inherent in those approaches.

Our empirical exercise leaves unexplored many other potential dimensions of the relationship between specialization and regional economic development. One such dimension is the link between incomes and the type, rather than the level, of specialization. New Yorkers might be richer on a per capita basis than Angelenos because NY has high relative and absolute specialization in finance and business services, which are higher-wage specializations than

entertainment. We have only confirmed that as finance grows bigger in absolute terms, New Yorkers working in that sector will see their wages rise. Research at the international scale confirms that countries with tradable sectors positioned near the top of the global ladder of product sophistication and quality do indeed have higher incomes than those oriented toward activities in the lower rungs (KEMENY, 2011, HAUSMANN et al, 2007). Applied to metropolitan regions, this reasoning suggests that specialization is related to development not so much through a general effect of overall levels of specialization, whether absolute or relative, as through the ‘what’ of specialization. It is good to do a lot of something, but even better to do a lot of something good.

Of course, in smaller regional economies, it follows that devoting greater effort to a more sophisticated activity will enhance the favorable effect of that specialization on the regional economy. This will mechanically raise levels of absolute and relative specialization in the favorable sector, and unleash the productivity effect that we detect above. The combined effects of ‘doing the right thing’ and doing so at a larger absolute scale, will move wages and incomes in the same positive direction. Inversely, an economy positioned far down on quality and innovation ladders is unlikely to resolve its income level problem by simply by increasing the scale – relative or absolute – of its agglomeration.

The most significant dimension of specialization, then, is the classical meaning of the term, i.e. concerning not the scale but the ‘what.’ This issue is dealt with in development theory through the notion of comparative advantage; in economic geography it features in theories that account for the locational sorting of tradable activities between regions, combined with agglomeration economies.

In the background of any consideration of the dynamics of specialization in an open global economy is the issue of the complex relationship between forces for regional convergence

and divergence. Why do some city-regions fall down the income rankings (Cleveland, Detroit), while others climb up (Houston, Dallas), and still others manage to maintain their positions at the top while transitioning their tradable sectors (San Francisco, Boston), and still others climb up a bit and then stagnate in the middle of the ladder (Las Vegas, Phoenix)? This evidently, though not entirely, has to do with the shifting industrial makeup of these places. In that process, change in specialization is not an entirely exogenous cause – it is partly an outcome – but it plays an important role.

Along these lines, some of the relative specialization hypotheses we discussed in section 3, but which we did not test in this paper, make claims about possible favorable effects of good relative specialization at t leading to good (or better) specialization at $t+n$. Notice that these hypotheses are not about maintaining or growing the same favorable specialization over time, but about a process of succession by which specializations dynamically affect one another over time and space. There is little in the empirical literature that tests this rigorously.^{xvii} The treatment of this issue remains largely qualitative and anecdotal. It reframes the specialization debate as one about development, but we are far from having the theory or measurement techniques adequate to this task. This debate raises the bar for evolutionary theories of the benefits of relatedness and for institutional theories of adjustment.

Practitioners' and policymakers' concern with specializing in the right thing lies behind the popular rankings of regional economies on the basis of their focus on finance, information technology, biotechnology, green technology, corporate headquarters and so on. These actors are rightly concerned with identifying successful places by virtue of the 'what' of specialization. But we have shown that, in many cases, their rankings are based on dubious measures; more careful approaches are needed. This observation applies to more syncretic academic concepts of specialization as well, of which we cite two very popular ones in recent years: "global cities" and

“creative cities” (SASSEN, 2001; FLORIDA, 2002). These concepts are at base making claims that regional economic performance is a function of having a regional economic base that is specialized in activities that are, respectively, ‘global’ or ‘creative’; each has spawned cottage industries in which cities are evaluated and ranked along these lines. Both are about specialization, but both suffer from many definitional problems. The concepts of globalness or creativity (the independent variables) mix sectors, labor force characteristics, and sometimes regional environmental features (such as “tolerance”). Moreover, neither has a clear dependent variable, opting for composite notions of “economic performance” (FLORIDA, 2002) or globalness (SASSEN, 2001). The most global cities – New York, London and Tokyo, and many of the rest of the top ten – are not the metropolitan areas with the highest per capita incomes. These wealthiest cities are actually mostly B-level globalization centers such as San Francisco, Oslo, Zurich, and Vancouver. The most “creative” metro areas are generally very high income regions, but we cannot tell whether this is because of their specialization in certain activities, their concentration of certain types of labor, or their environmental characteristics, nor how these different factors interact in any putative causal sequence (STORPER and SCOTT, 2009). One could obtain almost identical results to the “creative city” ranking by throwing out the labor force and environmental variables, and just ranking on the basis of specialization and wages in the tradable sectors; one could equally reverse it and obtain the ratings by using just the occupational composition (reflecting specialization, of course). In other words, neither of these analyses seem to add anything that is not done more crisply by simply analyzing the specialization of these region’s tradable economies.

Finally, we can return to the practical issues of using rankings in economic development practice and policymaking. As long as practitioners continue to believe that by shaping regional specialization patterns, they can improve economic development, then rankings such as location

quotients or other common measures will continue to exist, no matter that they remain fairly far away from more academic notions of specialization and its dynamics.

But even on their own terms, such ranking practices could be vastly improved. Rankings and classifications need to artfully mix concepts of relative and absolute specialization when they consider a particular set of industries or industry (e.g. finance, high tech, or ‘high wage’ or ‘high skill’ industries), or perhaps include both. A second lesson is that such rankings are basically uninformative at high levels of aggregation, at which there will be little or no relation to income effects. And issues of granularity are just one of several major issues around measuring specializations. Remaining is the problem of industrial relatedness or similarity that requires that researchers get closer to theorized mechanisms that ought to determine the boundaries of an industry.

A third and final lesson has to do with the relationship between specialization and quantitative growth prospects of regional economies. As noted, the principal practical and academic tool for attempting to estimate these effects is through relative specialization measures, in particular the location quotient. Such measures do poorly at their stated objectives because they cannot capture dynamics in the locational structure of the industry in question. A rise in external demand will not automatically benefit a regional economy if the industry’s locational structure is changing, rendering it highly contestable across locations. A good contemporary example of this is the logistics industry in Southern California. The region has a high level of absolute and relative specialization in this sector, and a high national location quotient. But this cannot be used to predict anything about quantitative employment changes in the region if the sector’s overall economic geography is shifting (new Panama Canal) or if capital is rapidly being substituted for labor (e.g. bigger ships, containers and trucks). Shift-share analysis can only capture this retrospectively, and – cruelly – even when it captures a favorable shift-in-share, it

cannot simultaneously include the absolute size of the industry at national scale, nor the industry's national employment density and quality.

This brings us back, once again, to the multidimensional nature of measuring specialization and the need to triangulate among the several facets of specialization – absolute, relative, share, and quality – to have any value to applied regional analysis. Both the academics and economic development professionals are in general far from such a high standard. This paper is an attempt to move us one step forward.

Acknowledgements:

The authors would like to thank the participants of the Evolutionary Economic Geography sessions at the 2012 Association of American Geographers, as well as two anonymous reviewers for helpful suggestions.

References

- ACEMOGLU, D., JOHNSON, S. and ROBINSON, J.A. (2001) The Colonial Origins of Comparative Development: An Empirical Investigation. *American Economic Review* 91, p.1369-1401.
- ACEMOGLU, D., and ROBINSON, J.A. (2008) Persistence of powers, elites and institutions. *American Economic Review* 9, p.267-293.
- ACHEN, Christopher H. 2000. Why Lagged Dependent Variables Can Suppress the Explanatory Power of Other Independent Variables. Presented at the Annual Meeting of Political Methodology, Los Angeles

- ACS, Z. and MORCK, R. and YEUNG, B. (1999) Productivity Growth and Firm Size Distribution. In: *Entrepreneurship, Small and Medium-sized Enterprises, and the Macroeconomy*, Acs, Z, Carlsson, B and Karlsson, C, Eds. Cambridge: Cambridge University Press.
- ARELLANO, M and HONORE, B. (2001) Panel Data Models: Some recent developments. In: *Handbook of Econometrics*, Vol. 5, Heckman, J.J and Leamer, E.E., eds. Amsterdam: North-Holland
- ATTARAN, M (1986) Industrial diversity and economic performance in U.S. areas. *Annals of Regional Science*, 20(2), p.44-54.
- BEAUDRY, C. and SCHIFFAUEROVA, A. (2009) Who's right, Marshall or Jacobs? The localization versus urbanization debate. *Research Policy*, 38(2), p.318-337.
- BOND, S.R. (2002) Dynamic Panel Data Models: A Guide to Micro Data Methods and Practice. *Portuguese Economic Journal*, 1(2) p.141-162.
- BOSCHMA, R. and IAMMARINO, S. (2009) Related variety, trade linkages, and regional growth in Italy. *Economic Geography* 85(3), p.289-311
- CICCONE A. and HALL, R. (1996) Productivity and the density of economic activity. *American Economic Review*, 86, p.54-70
- CHINITZ, B. (1961) Contrasts in agglomeration: New York and Pittsburgh. *American Economic Review*, 51(2), p.279-289.
- COMBES, P., DURANTON, G., and OVERMAN, H. (2005). Agglomeration and the adjustment

- of the spatial economy. *Papers in Regional Science*, 84(3), p.311–349.
- DE LA ROCA, J, and PUGA, D. (2013) Learning by working in big cities. CEPR Discussion Paper 9243.
- DURANTON, G. and PUGA, D. (2004) Micro-foundations of urban agglomeration economies. In: *Handbook of Regional and Urban Economics, Volume 4*, J.V. Henderson and J.V. Thisse, Eds., p.2063-2117. Elsevier BV.
- DURANTON, G., & PUGA, D. (2001). Nursery cities: Urban diversity, process innovation, and the life cycle of products. *American Economic Review*, 1454-1477
- ELLISON, G., GLAESER, E. (1997). Geographic concentration in U.S. manufacturing industries: A dartboard approach. *Journal of Political Economy* 105, 889-927.
- FELDMAN, M and ZOLLER, T. (2012) Dealmakers in place: social capital connections in regional entrepreneurial economies. *Regional Studies*, 46(1), p.23-37.
- FELDSTEIN, M. (2008). Did wages reflect growth in productivity? *Journal of Policy Modeling*, 30(4): p.591–594.
- FLORIDA, R. (2002) *The rise of the creative class*. New York: Basic Books
- FRENKEN, K; VAN OORT, F; VERBURG, T (2007).“Related variety, unrelated variety, and regional economic growth.” *Regional Studies* 41, p.685-697.
- GALBRAITH, J.K. and HALE, T. (2004) Income distribution and the information technology bubble. University of Texas Inequality Project, Working Paper 27.

GLAESER, E. and MARÉ, D.C. (2001) Cities and skills. *Journal of Labor Economics*, 19(2), p.316-342.

GLAESER, E. L., KALLAL, H. D., SCHEINKMAN, J. A., and SHLEIFER, A. (1992). Growth in Cities. *Journal of Political Economy*, 1126-1152.

GROSSMAN, G. and HELPMAN, E. (2002). *Special Interest Politics*. Cambridge, MA: MIT Press.

HARTOG, M., BOSCHMA, R., and SOTARAUTA, M. (2012). The impact of related variety on regional employment growth in Finland 1993–2006: high-tech versus medium/low-tech. *Industry and Innovation*, 19(6), 459-476.

HAUSMANN, R. HWANG, J. and RODRIK, D. (2007) What you export matters, *Journal of Economic Growth*, 12(1), p.1-25.

HIDALGO, C., KLINGER, B., BARABASI, A.-L., and HAUSMANN, R. (2007) The product space conditions the development of nations. *Science*, 317(5387), p.482-487.

HIDALGO, C., and HAUSMANN, R. (2009) The building blocks of economic complexity. *Proceedings of the National Academy of Sciences*, 106(26), p.10570-10575

HOOVER, E. (1948) *The Location of Economic Activity*. New York: McGraw-Hill.

ISSERMAN, A.M., and WESTERVELT, J. (2006) 1.5 million missing numbers: Overcoming employment suppression in County Business Patterns data, *International Regional Science Review* 29(3), p.311-335.

JENSEN, J. B. and KLETZER, L. (2006) Tradable Services: Understanding the Scope and Impact of Services Offshoring. *Offshoring White-Collar Work – Issues and Implications*, Washington, DC: Brookings Trade Forum, Lael Brainard and Susan M. Collins, eds.

KEELE, L.J., and KELLY, N.J. (2006) Dynamic Models for Dynamic Theories: The Ins and Outs of Lagged Dependent Variables. *Political Analysis* 14, p.186–205.

KEMENY, T. (2011) Are international technology gaps growing or shrinking in the age of globalization? *Journal of Economic Geography*, 11(1), p.1-35.

KOREN, M, and TENREYRO, S. (2003) Diversification and Development. Federal Reserve Bank of Boston Working Paper 3-3.

MORETTI, E. (2010) Local multipliers. *American Economic Review*, 100(2), p.373–377.

MORGAN, K. (1997) The learning region: institutions, innovation and regional renewal, *Regional Studies*, 31(5), p.491-503

MYRDAL, G. (1957). *Rich Lands and Poor Lands*. New York: Harper Brothers.

- OLSEN, M (1965) *The logic of collective action: Public goods and the theory of groups*.
Cambridge, MA: Harvard University Press.
- PAGANO, P. and SCHIVARDI, F. (2003) Firm Size Distribution and Growth. *The Scandinavian Journal of Economics*, 105(2), p.255-274.
- PERSSON, T. and TABELLINI, G. (2002). *Political Economics: Explaining Economic Policy*.
Cambridge, MA: MIT Press.
- QUIGLEY, J. M. (1998) Urban diversity and economic growth. *Journal of Economic Literature*, 12(2), p.7-138.
- RICHARDSON, H. W. (1969) *Regional Economics*. New York: Praeger.
- ROODMAN, D. (2005). XTABOND2: Stata module to extend xtabond dynamic panel data estimator. Center for Global Development, Washington.
- ROSENTHAL, S., and STRANGE, W. (2004) Evidence on the Nature and Source of Agglomeration Economies. In: *Handbook of Regional and Urban Economics, Volume 4*, J.V. Henderson and J.V. Thisse, Eds. Elsevier BV.
- SASSEN, S. (1991) *The global city: New York, London, Tokyo*. Princeton University Press.
- SAXENIAN, A. (1994). *Regional Advantage*. Cambridge, MA: Harvard University Press.

SCHOTT, P. (2007) The relative sophistication of Chinese exports. *Economic Policy*, 23(53), p.5-49

JOINT VENTURE SILICON VALLEY NETWORK (2008) Index of Silicon Valley. Retrieved August 21, 2012 from:
<http://jointventure.org/images/stories/pdf/The%20Index%20of%20Silicon%20Valley%202008.pdf>.

STORPER, M., SCOTT, A.J. (2009) Rethinking human capital, creativity and urban growth. *Journal of Economic Geography*, 9(2), p.147-167.

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

WEITZMAN, M.L. (1998). Recombinant Growth. *Quarterly Journal of Economics* 113(2), p.331-360.

WHEATON, W.C. and LEWIS, M.J. (2002) Urban wages and labor market agglomeration. *Journal of Urban Economics*, 51(3), p.542-562.

WOOLDRIDGE, J. M. (2002) *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA.

WOLFSON, M.C. (1997) Divergent inequalities: theory and empirical results. *Review of Income and Wealth* 43(4), p.401-421.

YANKOW, J. (2006). Why do cities pay more? An empirical examination of some competing theories of the urban wage premium. *Journal of Urban Economics*, 60(2), p.139–161.

Tables

Table 1: Regional Specialization and Selected Development Indicators for Major Combined Statistical Areas

	<u>1970</u>		<u>2009</u>		<u>1970-2009</u>	
	Specialization (Herfindahl)	Per Capita Income	Specialization (Herfindahl)	Per Capita Income	Income CAGR	Employment CAGR
				\$37,10		
Atlanta	0.01	\$3,932	0.015	1	5.92	4.26
Boston	0.008	4,430	0.015	48,831	6.35	0.89
Chicago	0.004	4,861	0.013	43,047	5.75	0.66
Dallas	0.009	4,167	0.014	39,811	5.96	2.67
Houston	0.009	4,131	0.015	42,523	6.16	2.78
Los Angeles	0.003	4,857	0.012	39,301	5.51	1.52
New York	0.006	5,212	0.013	52,354	6.09	0.47
Philadelphia	0.008	4,458	0.014	44,905	6.1	0.63
San Francisco	0.008	5,265	0.015	54,062	6.15	1.44
Washington DC	0.011	4,802	0.016	52,646	6.33	1.62
				35,7		
U.S. Average	0.027	3,711	0.022	63	5.992	1.532
				5,31		
U.S. Std. Dev	0.03	616	0.009	1	0.322	0.980

Note: Herfindahl indices produced using County Business Patterns. Larger numbers indicate that sectoral employment patterns deviate from a uniform distribution. Results are not directly comparable across years due to the switch in classification schemes in 1997 from SIC (4-digit) to NAICS (6-digit). SD indicates standard deviation for all U.S. metropolitan areas. Selected development indicators from the Bureau of Economic Affairs Regional Economic Accounts. CAGR stands for compound annual growth rate. Income figures are presented in nominal U.S. dollars. Employment figures exclude proprietors.

Table 2. Relative and Absolute Specialization in Employment in Information Technology among U.S. Metropolitan (and Combined Statistical) Areas, 2010

Metro Area	Relative		Absolute
San Francisco, CA	10%	San Francisco, CA	255,334
Washington DC-MD-VA-WV	8	Washington DC-MD-VA-WV	240,721
Seattle, WA	7	New York, NY-NJ-CT-PA	184,917
Austin, TX	7	Los Angeles, CA	153,524
Boston, MA-NH-ME-CT	5	Boston, MA-NH-ME-CT	122,474
Atlanta, GA	5	Seattle, WA	90,511
San Diego, CA	5	Dallas, TX	85,989
Dallas, TX	4	Chicago, IL-IN-WI	82,549
Portland, OR-WA	4	Atlanta, GA	74,566
Denver, CO	4	Philadelphia, PA-NJ-DE-MD	52,871

Note: Authors' calculations using employment data from County Business Patterns. To filter out small metropolitan areas, we present results for cities with an employment base over 500,000.

Table 3: Typology of theories of the development effects of specialization

SPECIALIZATION Type	ARGUMENT	SOLID ARGUMENT?	EVIDENCE
IA. Overall level of specialization / diversification	Spreads risk from external shocks	<ul style="list-style-type: none"> - Addresses shocks, not opportunities - Urbanization economies do not enhance diversification - Main benefit is from size not diversification per se 	No hard evidence that diversification raises long-run regional employment levels or quality
IB. Overall level	Dynamic version: relatedness through diversification helps evolution	Is it diversity or complex “related” specialization?	Evidentiary claims extremely sensitive to definition of “related” . No consensus about this.
IIA. Relative (share) specialization	Reduces competition for factors/congestion costs	Not clear why would be good for regional economy as whole	
IIB. Relative (share) specialization)	Focuses political-elite attention	Chinitz hypothesis supported by institutionalist literature	<ul style="list-style-type: none"> -Difficult to test at any scale -No large-sample tests at regional scale
IIC. Relative (share) specialization	New regionalism	Not just industries, but their supporting environments, ecosystems	Case studies suggest this, but lots of conceptual imprecision. No large-scale tests.
III. Absolute specialization (size of cluster)	Scale leads to greater productivity	<ul style="list-style-type: none"> -Theory on sharing, matching, learning = at least the first two strongly scale dependent; third should have positive scale effect through specialization and diversity of knowledge community 	Some confirming evidence in urban economics

Table 4: Average Wages in Information Technology Sectors, 2010

Sectors	Average Wages: Los Angeles	Average Wages: San Francisco
<i>Three-digit Sector</i>		
Professional, Scientific and Technical Services (541)	\$66,736	\$100,834
<i>Selected Individual Six-digit Sectors</i>		
Software publishers (511210)	\$128,583	\$169,432
Custom Computer Programming Services (541511)	89,295	111,648
Computer System Design Services (541512)	90,874	111,312
Computer Equipment and Software Merchant Wholesalers (423430)	80,416	155,961

Note: Authors' calculations based on data from County Business Patterns. Wages are averages expressed in nominal 2010 dollars.

*Table 5. Estimates of dynamic relationship between specialization and wages, 1998-2010
4-digit NAICS Industries*

Dependent Variable: Average Industry x Region Annual Wage				
	(1) OLS	(2) OLS	(3) OLS	(4) GMM-FE IV BW(2)
VARIABLES				
Relative Specialization	3,839*** (109.8)		1,953*** (126.5)	-265.5 (649.6)
Absolute Specialization		0.597*** (0.009)	0.486*** (0.010)	0.279*** (0.081)
Lagged Average Wages				0.233*** (0.033)
Metro employment (000s)	2.025*** (0.028)	1.272*** (0.031)	1.425*** (0.032)	4.48*** (0.783)
Avg. Employees per Firm			0.683 (0.797)	-28.48*** (4.968)
Constant	27,499*** (151.2)	28,248*** (203.7)	27,764*** (150.8)	
Observations	114,155	114,155	114,155	72,923
Groups				17,160
Year Dummies	Yes	Yes	Yes	Yes
R-squared				0.17
First-stage <i>F</i> -statistic				91.14
Hanson <i>J</i> Statistic				1.044
(Chi-square p-value)				(0.307)

Note: Asterisks indicate significance levels:*** p<0.01, ** p<0.05, * p<0.1. Models 1-3 estimated with heteroscedasticity-robust standard errors. Model 4 estimated using two-step robust GMM with HAC standard errors produced with a 2 year bandwidth.

Table 6. Estimates of dynamic relationship between specialization and wages, 1998-2010, at varying levels of industrial aggregation

	(1) 2-digit	(2) three	(3) 5-digit	(4) 6-digit
Relative Specialization	-35.61 (92.80)	1,152** (511.4)	281.1 (2,147)	5,553 (7,830)
Absolute Specialization	0.007 (0.009)	0.0971** (0.0407)	0.830*** (0.231)	-0.292 (0.461)
Lagged Average Wages	-0.019 (0.062)	0.0336 (0.168)	0.244*** (0.0426)	0.00233 (0.025)
Metro Employment (000s)	6.107*** (1.777)	4.423 (2.777)	5.630*** (0.894)	5.019*** (0.743)
Avg. Employees per Firm	-115.7*** (20.89)	-214.9*** (65.77)	-42.69*** (6.921)	-44.57*** (9.622)
Observations	39,552	39,115	92,954	180,611
Groups	4,710	7,575	23,862	53,282
Year Dummies	Yes	Yes	Yes	Yes
R-squared	0.080	0.030	0.139	0.025
First-stage <i>F</i> -statistic	28.32	16.31	25.87	9.593
Hanson <i>J</i> Statistic	2.412	0.123	1.059	0.894
(Chi-square p-value)	(0.1204)	(0.7255)	(0.303)	(0.344)

Note: Asterisks indicate significance levels:*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All results presented in this table produced using two-step robust GMM-FE with HAC standard errors produced with a 2-year bandwidth.

*Table 7. Estimates of dynamic relationship between specialization and wages, 1998-2010
4-digit NAICS Industries, Varying city groups by total employment*

	(1) 100 largest Metros	(2) 150 largest Metros	(3) 200 largest Metros	(4) 200 smallest Metros
Relative Specialization	1,375 (3,780)	1,014 (2,366)	492.1 (1,471)	492.1 (1,471)
Absolute Specialization	0.225** (0.0978)	0.237*** (0.0852)	0.258*** (0.0838)	0.258*** (0.0838)
Lagged Average Wages	0.272*** (0.0424)	0.271*** (0.0343)	0.241*** (0.0380)	0.241*** (0.0380)
Total Metropolitan Employment	3.722*** (0.860)	3.971*** (0.799)	4.283*** (0.800)	4.283*** (0.800)
Average Firm Size	-17.78*** (6.614)	-24.12*** (6.328)	-25.79*** (5.519)	-25.79*** (5.519)
Observations	35,181	47,796	58,369	58,369
R-squared	0.199	0.192	0.176	0.176
Groups	8,547	11,393	13,803	13,803

Note: Asterisks indicate significance levels: *** p<0.01, ** p<0.05, * p<0.1. All results presented here produced using two-step robust GMM with HAC standard errors produced with a 2-year bandwidth.

ⁱ For the present purposes, “terms of trade” refers to the relative prices of the region’s output compared to the prices of the goods and services it imports. If the region’s output enjoys increasing ratios of its unit prices relative to what it imports, then its terms of trade are said to be improving.

ⁱⁱ Acknowledging all the limitations of the industrial data that we discuss in more detail below

ⁱⁱⁱ To minimize the importance of smaller metropolitan areas, we present results only for metropolitan and combined statistical areas with a total employment base over 500,000.

^{iv} Relative rankings correspond to those that would be produced using location quotients.

^v Wage data, as compared with output data from the Census of Manufactures, is also less likely to introduce bias due to mis-measurement (CICCONI and HALL, 1996).

^{vi} Though contrasting evidence exists, for instance De la Roca and Puga (2013)

^{vii} The Census Bureau’s Longitudinal Employer-Household Dynamics (LEHD) is the closest data of this kind for the U.S., though it offers very scant establishment information. Access to such data are also somewhat out of reach: access is restricted to approved researchers, with approval often taking very lengthy periods.

^{viii} Including lagged dependent variables as predictors can be a tricky procedure, with the possibility that such variables will (incorrectly) swamp the effects of other predictors of interest. We discuss this and methods of correcting for such problems further in the results section, but this problem does not afflict the results of this empirical inquiry.

^{ix} There are also some issues with employment data that is suppressed due to reasons of confidentiality (ISSERMAN and WESTEVELDT, 2006), though this may not be true in more recent samples.

^x Though JENSEN and KLETZER use locational Gini coefficients, the Herfindahl index made more sense to us, because it is explicitly about concentration – another way to say specialization. See WOLFSON (1997) for a comparison of the two measures. We explored the sensitivity of results to the choice of alternate years, including 2000 and 2005. Results did not materially vary.

^{xi} In initial exploration, city and industry dummy variables were also included. These would account for the effect of any stationary city-wide or industry-wide shocks. Since these did not materially change the results for the variables of interest, we do not report these here. These dummies also got unwieldy in the more complex approaches that follow.

^{xii} While it is common for researchers to log transform some variables, especially wages, we opt against this approach, choosing to leave variables in their natural scale. We do so mainly because of the size of our dataset. While non-normality of predictors can indicate potential problems of non-normality of the residuals, this issue is not likely to bias estimates produced using a dataset with so many observations. In most cases, logging did not materially affect results.

^{xiii} We conduct Wooldridge’s test using the Stata command ‘xtserial.’

^{xiv} This is also true for estimates produced using system GMM, which is ideal for short panels with lagged outcome variables included as predictors. What distinguishes the results presented from those produced with system-GMM is that the latter produced a very large number of instruments (by definition, all lags of all instruments), which can cause efficiency problems in panels deeper than 8 (see for instance, Baum 2013). In this case, either the instrument matrix did not satisfy diagnostics, or, when limiting lags, AR(2) behavior was significant. Given that results for coefficients were consistent, we opted to present results from the two-step GMM-FE procedure..

^{xv} Results are not particularly sensitive to moderate changes in these thresholds.

^{xvi} Using GMM, we additionally find consistent results when we estimate the relationship for the mix of Metropolitan and, where available, Consolidated Statistical Areas, though questions remained about instrument validity here.

^{xvii} HIDALGO et al (2007) and NEFFKE et al (2011) are notable exceptions.