PRIPE
7/23/07
3:00 pm

Innovation, Cities, and New Work*

Jeffrey Lin†

First version: October 2004
This version: November 2006

Abstract

I find that the supply of educated workers and local industry structure matter for the location of *new work*—that is, new types of activities that closely follow innovation. Using Census 2000 microdata, I show that regions with more college graduates and a more diverse industrial base in 1990 are more likely to attract these new activities. Across metropolitan areas, initial college share and industrial diversity account for 50% and 20%, respectively, of the variation in selection into new work unexplained by worker characteristics. I use a novel measure of innovation based on new activities identified in decennial revisions to the U.S. occupation classification system. New work follows innovation, but unlike patents, it also represents subsequent adaptations by production and labor to new knowledge. Further, workers in new activities are more skilled, consistent with skill biased technical change.

Key words: innovation, agglomeration, occupations, human capital, industrial diversity
JEL codes: J24, O33, R12, R23

---

* Preliminary; comments welcome. Gordon Hanson and Hoyt Bleakley provided invaluable guidance. Dana Bradley and Marisa Tegler at the U.S. Census Bureau provided documentation on new occupations. I thank Eli Berman, Julian Betts, Nora Gordon, Roger Gordon, Ian Savage, and seminar participants at UCSD for helpful comments. All errors are my own.
† Doctoral Candidate in Economics, University of California, San Diego. Address: 9500 Gilman Drive, La Jolla, CA 92093-0534. E-mail: jelin at ucsd dot edu. Web: http://econ.ucsd.edu/~jelin/

1 Introduction

Human capital is central to our understanding of technological change and economic growth, according to Lucas (1988). In one illustration of this role, Jacobs (1969) contends that the *creation* of new knowledge results from novel combinations of existing techniques. Often, these combinations come from knowledge spillovers between workers possessing different sector-specific experiences. Alternatively, human capital can aid in learning and understanding new ideas from others—Glaeser (1999) says that skill accelerates the *diffusion* of innovation. Finally, in turn, Schultz (1975) argues that certain abilities speed *adaptation* to disruptions from new knowledge. Here, human capital helps to better identify the changing incentives that result from innovation, allowing firms and workers to adjust their activities quickly in response to technological change.[1]

These examples all emphasize *access* to the human capital of others. For Jacobs, the relevant supply of human capital is a wide variety of industry-specific abilities that might lead to innovative combinations. Then again, industrial diversity—or more traditionally, formal education—might also supply the particular skills that firms and workers need to adapt quickly to innovation. This is a central reason why cities are important: by geographically concentrating human capital, cities provide access to the different types of skills required for the creation, diffusion and adaptation of new knowledge.[2] Certain regions, with larger or more relevant stocks of human capital, should be better at this process than others.

In this paper, I examine the characteristics that make some regions better at creating or attracting *new work*, by which I mean new types of activities.[3] These new activities follow innovation, but unlike other measures of innovation output, they also represent market acceptance and subsequent changes to the organization of production, labor demand, and labor supply. In a sense, new work is an observable, *adapted* analog to new knowledge. I focus on whether the initial supply of college-educated workers and industrial diversity matter for the location of these new activities. In a model based on Helpman (1998), scale economies and transport costs explain why some regions are better able to adapt to new technologies.[4] The central prediction is that new work appears in regions with greater initial supply of educated workers (who, locally, perform new work) and an industrially diverse base of production firms (who employ new work). In other words, differences in the ability of regions to attract new work may reflect different stocks of the particular types of human capital required to create or adapt to new knowledge.

This paper is probably most closely related to the empirical literature on city growth. For example, Glaeser and Saiz (2005) find that educated cities experience faster wage and population

[1] See also Mokyr (2002).
[2] See also Marshall (1890) and Duranton and Puga (2004).
[3] The phrase *new work* comes from Jacobs (1969).
[4] I present this model in Section 3, with details in the appendix.

growth, providing indirect evidence of the ability of these cities to adapt to innovation. However, there is controversy in this literature on the merits of industrial diversity (Glaeser et al. 1992, Henderson et al. 1995, Henderson 2003). Elsewhere, there is partial evidence on other connections between agglomeration, human capital, and innovation; for example, Rauch (1993) and Moretti (2004) find higher wages in educated cities, Jaffe et al. (1993) find that patents are geographically localized, and Feldman and Audretsch (1999) trace product inventions to industrially diverse cities. In addition, work on skill biased technical change suggests that it is college graduates who are best suited to adapt to new technologies (Berman et al. 1994, Autor et al. 1998).[5] This paper's contribution to the literature is to highlight both the supply of educated workers and local industrial structure as sources of regional advantage. Also, via new work, I am able to more specifically characterize how labor markets change in response to the creation of new knowledge.

To identify new work, I rely on Romer's (1990) definition of a unit of knowledge as a recipe—a set of instructions—for combining raw inputs into useful product. Implementations of these recipes might be labeled techniques, activities, or types of work. New recipes, requiring previously unknown actions or combinations of actions, virtually necessitate new activities.[6] This observation is at the heart of the paper. I argue that occupations are essentially measures of activities or techniques. Changes to the Census occupation classification system (OCS), the official catalog of activities, reflect the emergence of new techniques, and new knowledge. In Census 2000 microdata, I identify workers in new occupations that first appeared between the 1990 and 2000 OCS.[7] Then, I estimate a model predicting worker selection into new occupations. The main explanatory variables are 1990 measures of educated labor supply and industrial diversity across U.S. metropolitan areas. Controlling for worker characteristics, I interpret the estimated coefficients on these variables as identifying their effects on attracting new work to regions.

To preview the results, I find that regions with more college graduates and a more diverse industrial base in 1990 are more likely to attract new work in 2000. Across metropolitan areas, college share and industrial diversity account for 50% and 20%, respectively, of the variation in selection into new work unexplained by worker characteristics. I estimate that 5% of U.S. workers in 2000 participate in new work, and a change of one standard deviation in the 1990 metro college graduate share increases a worker's likelihood of selecting into new work by almost 0.6%. A similar change in 1990 industrial diversity increases selection also by 0.6%. These results highlight the importance of the supply of educated workers and industrial diversity for the

[5] Also related is Bresnahan and Trajtenberg's (1995) paper on general-purpose technologies. Certain pieces of knowledge are broadly accessible and adaptable to a wide variety of purposes; new work captures some of these properties. Bleakley and Lin (2006) find that thicker markets provide better labor market matching; thicker markets for new skills (via more skilled workers) might contribute to the quicker adoption of new work.

[6] A process of destroying old work while creating new work occurs with many innovations. Adding machine operators, paper filers, and telegraph operators may disappear as statistical analysts, database managers, and network administrators appear.

[7] This is similar to Xiang (2004), who identifies new products in revisions to industry codes. Autor et al. (2003) and Bacolod and Blum (2006) also use detailed occupational data to study skill bias and wage inequality.

location of new work. I also show that college graduates are more than four times as likely as high school dropouts to select into a new occupation, consistent with skill biased technical change. Workers in new activities also earn higher wages than observationally similar workers in older activities, consistent with higher wages in skilled cities.

I further examine possible sources of bias and discuss alternative explanations. These results do not appear to be due to unobserved region or worker characteristics. Using a coarser methodology based on more aggregate occupation codes, I also identify new occupations that appeared between 1960-1970 and 1970-1980.[8] Estimates from these earlier periods are similar to the main results.

The rest of the paper proceeds as follows. The next section describes the new work data, the characteristics of workers in new activities, and the differences between new work and patents. Section 3 models the effect of innovation (in the form of increasing varieties of a traded good) on the location of production activities. I use the central prediction to guide the reduced-form estimation described in Section 4. I discuss results and evaluate alternative explanations in Section 5. In the appendices, I provide further detail on both data collection and the theoretical model. Section 6 concludes.

2 Data and Methodology

Over time, changes in the Census occupation classification system (OCS) form an important—if accidental—record of the changing nature of work in the United States. In this section, I outline the process of collecting the new work data for the period between 1990 and 2000. The main steps are (1) understanding the OCS as a comprehensive, detailed catalog of activities, (2) comparing occupation titles in 1990 and 2000 to identify new types of work, and (3) matching new work to available labor market data. Later in this section, I describe characteristics of the population in new work and compare new work to patents, another measure of innovation output. Further details are relegated to the appendix.

2.1 The Census occupation classification system

The Census Bureau uses the OCS as a catalog of the various types of work performed in the U.S. economy. It is updated every 10 years to reflect both the changing nature of work and the changing needs of data users.[9] Each revision relies on previous versions of the OCS, field research, the *Dictionary of Occupational Titles* (produced by the Bureau of Labor Statistics), and written descriptions from Census respondents of the type of work they perform. These reviews ensure

[8] New work in earlier Census years is covered in the appendix. Because the 3-digit occupation codes in the OCS did not change between 1980-1990, it is much more difficult to identify new work in this period.
[9] See technical papers from the U.S. Census Bureau (2003) and the Bureau of Labor Statistics (2005) for good overviews of this process.

that new activities are identified every decade and that the standards for identifying them remain consistent.

Two companion volumes in each Census (the *Alphabetical* and *Classified Indexes of Industries and Occupations*) contain thousands of (5-digit) occupation titles (hereafter "titles"), which are the atomistic unit of the OCS. A title describes a small number of individual jobs that require the use of a similar set of techniques. This narrow scope means there are a large number of titles: between 1950 and 2000, the number of titles expanded from about 25,000 to 31,000. Any broad category of work may consist of hundreds of occupation titles. For example, in the 2000 OCS, there are over 500 titles that contain the word engineer. Among these are at least nine computer-related engineering occupations (e.g., computer software applications engineer, Microsoft certified systems engineer, Novell certified engineer, software requirements engineer) and at least eleven aerospace-related engineering occupations (e.g., aerospace engineer, aircraft instrument engineer, airport engineer, flight test engineer). There are also over twenty varieties of economists, with descriptions spanning specialty (e.g., econometrics, finance, labor, trade), and function (e.g., teacher, research analyst, research assistant, policy advisor).

Each 5-digit title is also assigned to a 3-digit detailed occupation code ("DOC"). Unlike titles, we can observe DOCs in public-use Census microdata. Each DOC groups together titles according to the similarity of work performed and skills required. In the 2000 Census, the median number of titles in each DOC is 33. For example, DOC 110, *Network and computer system administrators*, contains 30 occupation titles. Some of these are: certified Novell administrator, computer security information specialist, computer systems administrator, LAN administrator, UNIX systems administrator, web administrator, and Windows system administrator.

2.2 Comparing the 1990 and 2000 OCS to identify new work

I examine revisions to the OCS between 1990 and 2000, which should reflect the changing nature of work in the U.S. over time. This is somewhat complicated by changes to the OCS unrelated to innovation; in some years the taxonomic structure of the OCS shifts, due to changes in data demands. Increased interest in a specific sector may cause the Census Bureau to identify new DOCs, without any actual change in the types of work performed. For example, employment growth between 1960 and 1970 led to the separation of lawyers and judges into two separate DOCs. In addition, the 1990-2000 revision reflected significant changes related to the creation of "job families," where parts of some DOCs shifted to other DOCs (U.S. Census Bureau 2003).

Crucially, this sort of spurious identification is not a problem with the more detailed 5-digit occupation titles. A Census technical paper notes that titles found in the *Indexes,* unlike DOCs, "provide information about the intended, or 'ideal' changes from each [...] occupation code of [the 1990] classification into each [...] occupation code of [the 2000] classification" (U.S. Census Bureau 2003, p. 9). DOCs may be combined or split apart according to the needs of the Census or of a growing population, but titles remain anchored to a small number of techniques performed. In addition, the 2000 *Indexes* are derived from the 1990 *Indexes*. With some exceptions noted in internal Census documents, occupation titles carried over from 1990 to 2000 are consistent, and new titles reflect technical change. The creation of a new occupation title is based on the emergence of new types of work, using new techniques. By comparing individual occupation titles between 1990 and 2000, I can identify new types of work.

I compare electronic versions of both the 1990 and 2000 *Classified Indexes*. I first eliminate OCS 2000 titles based on exact string matches, allowing only for differences in punctuation, capitalization, and spacing. I also correct for consistent changes between the two *Indexes*. Among these are small variations in spelling (e.g., gauger/gager), abbreviation (class of worker/c.o.w.), naming convention (automobile/auto), and the elimination of gender-specific titles ("nursery man"/"nursery worker"). Finally, I manually inspect the 3,000 remaining OCS 2000 titles and compare them to the 1990 *Index*. At the same time, I consult detailed internal Census documentation on the sources of changing occupation titles. Some of these remaining titles clearly exist in 1990, but are phrased in a way that makes it difficult to match them mechanically ("fork-truck driver"/"forklift truck operator", "monorail operator"/"monorail car operator", "portable router operator"/ "portable machine cutter operator"). By eliminating these kinds of matches, I obtain about 850 new titles that appear between 1990 and 2000. I am unable to match these to *any* title appearing in the 1990 *Index*. This list of 850 titles constitutes new work in 2000.[10]

2.3 Matching new occupations to Census 2000 microdata

The final step in creating usable data involves collapsing the 5-digit titles to the 3-digit DOCs observed in Census microdata. I count the number of new titles as a share of all titles within each 2000 DOC. Table 1 contains a list of the DOCs with the highest new title shares.[11] For example, DOC 111, *network systems and data communication analysts*, contains the most new titles as a share of total titles—29 out of 30 titles (96.7%) do not match any title from the 1990 *Index*. DOCs related to information technology and medicine generally have the highest incidence of new titles. In addition to the top DOC, respondents employed as computer support specialists, network administrators, software engineers, radiation therapists, and biomedical engineers were highly likely to be engaged in new types of work.

The majority of 2000 DOCs contain no new titles at all; only a few contain as many new titles as those in Table 1. As seen in Figure 1, the distribution of new titles within DOCs is heavily skewed. Out of the 505 DOCs, 56% contain zero new titles and 75% have new title shares of less than five percent.

[10] Available at my web site (http://econ.ucsd.edu/~jelin/papers/newwork/newtitles.pdf). In Section 5.5, to show that my results are not sensitive to a particular strategy for identifying new occupations, I also use two wider definitions of new work. These definitions are described in the appendix.
[11] Tables A1 and A2 in the appendix report new DOCs between 1960-1980.

Upon inspection, new DOCs seem to reflect new labor demands that result from actual innovation. Consider again DOC 111, *network systems and data communication analysts*. Among the new occupation titles within this DOC are chat room host/monitor, computer networks consultant, network engineer, Internet developer, and web designer. According to the *Classified Index*, workers in this DOC "analyze, design, test and evaluate network systems, such as local area networks (LAN), wide area networks (WAN), Internet, intranet, and other data communications systems." Clearly, these new occupations are tied to innovations in network and computer technology that occurred in the 1980s and 1990s. That the OCS catalogs these occupations in 2000 but not in 1990 suggests a close, if slightly delayed, relationship between innovation and new work. Historically, some important innovations during this time period were the creation of the first TCP/IP wide-area network in the mid-1980s, the launch of the World Wide Web in 1991, and the development of the first graphical web browser, Mosaic, in 1993.

I match the new DOC data to 2000 Census microdata from the Integrated Public Use Microdata Series (Ruggles and Sobek et al., 2004). The data contains detailed personal characteristics and metropolitan area of residence for approximately 1% of the 2000 U.S. population. I use 1999-defined consolidated metropolitan areas[12], so that, for instance, Stamford and Newark are grouped with New York, Orange County and Riverside are with Los Angeles, and Dallas and Fort Worth are grouped together.[13] The sample (1.5 million observations) includes all workers age 20 to 65 in identified metropolitan area and occupations. Every observation includes a variable equal to the share of new occupation titles within that respondent's identified detailed occupation code. I interpret this variable as reflecting the likelihood that a worker participates in an activity that first appeared (in the OCS) between 1990 and 2000. Table A5 in the appendix contains summarizes other variables.

2.4 Characteristics of those employed in new work

I estimate that about 5% of the U.S. workforce participated in new work in 2000. (I calculate this number by multiplying new title shares by total employment for each DOC.) Table 2 describes the population in new work. In metropolitan areas, employment share in new work is 5.3%; outside metropolitan areas, new work share is only 3.7%. New work adoption appears to be an urban phenomenon. New workers are more likely to be male, and new work employment shares tend to be higher for younger age groups (peaking in the late 20s).

The map in Figure 2 compares the share of workers in new work across metropolitan areas. Darker shading corresponds to a higher local share of workers in new occupations. Greater Washington, the San Francisco Bay Area, Raleigh-Durham, Austin, and Boston had the highest shares in new work in 2000; McAllen, Texas, Fresno, Bakersfield, and Modesto had the lowest shares in new work.

---

[12] See (http://www.census.gov/population/estimates/metro-city/99mfips.txt).

[13] Section 5 reports results using unconsolidated metropolitan areas.

New occupations employ more educated workers, consistent with skill-biased technical change. In Table 3, I present sample estimates of the share of workers employed in new work for four levels of educational attainment. The second column shows that while 8.6% of college graduates are in new work, new workers comprise only 1.7% of high school dropouts. The share of workers employed in new work increases monotonically by educational attainment. The types of new work differ, too, between educational groups; while college graduates in new work are likely to be computer software engineers or computer programmers, high school dropouts in new work are more likely to participate in activities that are less skill-intensive, such as estheticians and electrolysists, both in the category of personal appearance workers.

Table 4 compares three characteristics of workers in new work to other workers not employed in new work. (I divide the sample by comparing the imputed likelihood that a worker is in a new occupation to the sample mean of 4.8%.) Those in new work (or rather, those more likely than average to be in new work) have an average educational attainment of 14.8 years, versus 13.2 years for other workers. In addition, 47% of new workers are college graduates, compared to 21% of other workers. Finally, workers in new occupations have a 30% higher hourly wage relative to other workers. Each of these differences is significant, at a 99% level of confidence, as determined by a test on the equality of means.

The wage premium for participating in new work remains even when controlling for education and experience. I perform a wage regression, with dependent variable log hourly wage and independent variables education, experience, and the new work variable. I find that employment in new work predicts a wage premium of upwards of 35% (logarithmic) over observationally similar workers, with similar levels of education and experience, who are not in new work. More detail on this regression is in the appendix.

2.5 New work compared to patents

New work has several unique strengths as a measure of new knowledge. Consider a comparison between new work and patents, a common measure of innovation. Both are outputs related the (invisible) creation of new knowledge. A patent must also pass through both patentability rules and incentives affecting the patenting decision. A new occupation reflects both market acceptance of a new idea and its subsequent effects on production, labor demand, and labor supply. These differences drive their respective strengths: patents are readily available and can identify incremental advances, while new work is broader in industrial scope and is less sensitive to firm behavior. If patents track the *birth* of an idea, then new work tracks the *effects* and the *adaptation* of an idea. The two measures might be considered complements when it comes to identifying new knowledge.

A noted weakness of patent data is the influence of firm behavior and varying sets of patentability rules. It is well known that incentives affect patenting decisions; some inventions may not be patentable at all.[14] For example, Hall (2005) notes defensive patenting is common in some industries. A new occupation title, on the other hand, is not subject to patentability rules or firm behavior. A title is recognizable as soon as a small number of people perform that activity. Occupation titles also capture activity across sectors, which is useful in identifying innovations outside manufacturing. This contrasts with more common approaches that focus solely on product innovations (e.g., Feldman and Audretsch, 1999). As Bresnahan and Trajtenberg (1992) note, new knowledge may have unexpected applications across sectors, encompassing new products and new services. Further, by matching new work data to microdata, I can identify workers engaged in new work; this data is useful for studying the effects of innovation on labor markets. Finally, new work has potential benefits for historical analyses: this strategy may be extended backwards in time, to earlier Censuses. On the other hand, new work may not include incremental innovations. For example, an increase in the clock speed of an Intel processor may merit a patent but not create a new occupation. Only over a longer period of time, with larger advances in processor technology, would new types of work emerge.

Thus, new work and patents are certainly related, though they differ in their informational content. To show this, I use accumulated utility patent counts (i.e., patents for invention) between 1990 and 1999 from the U.S. Patent and Trademark Office (2000). Patent locations are determined by the location of the first-listed patent filer. Measured across metropolitan areas, the correlation coefficient between patent filings per capita and the workforce share in new work is 0.61.

Figure 3 graphs workforce share in new work in 2000 against accumulated utility patents per capita between 1990 and 1999. Each point represents a metropolitan area. The fitted line is from a regression of new work share on patents per capita. Metropolitan areas below this line, with unexpectedly high rates of patenting, include Rochester, Boise City, and Grand Rapids. Cities above the line have higher new work shares than predicted by patents per capita. This category corresponds to many cities most associated in the popular imagination with the knowledge economy in the 1990s: Austin, Boston, San Francisco, and Seattle. This emphasizes again the differences between the creation of new knowledge and its subsequent applications. Certainly, the two processes are related, and regions with particular stocks of human capital will be better at both. This is indicated by the positive relationship between patents and new work. Variation from the fitted line, however, in part must reflect differences in the ability of regions to adapt to the creation of new knowledge.

3 Theory

I showed in the last section that certain regions attract more new work. In this section, based on Helpman's (1998) work, I formalize the effect of innovation on the distribution of production

[14] In addition, the patent office may be limited in its ability to identify truly new inventions. See Wu (2006).

activities across regions.[15] The starting point is to imagine innovation as a shock to the economy. How do workers and firms across regions adapt to this shock? Here, innovation occurs in the form of an exogenous expansion in the variety of production activities. My strategy is to solve for initial equilibria, introduce a global shock in the form of innovation, and finally solve for the new equilibria, and the new distribution of activities across regions. [16]

The main features of the model are as follows. Scale economies and transport costs encourage traded goods production to locate near consumers. Consumer love-of-variety and transport costs, along with their use as factors in production, encourage skilled households to locate near production variety. Congestion costs, in the form of increases in the price of a non-traded good (e.g., housing), provide a smooth and realistic dispersing force. A key assumption is that as the variety of traded goods expands, the aggregate expenditure share devoted to housing falls.[17]

The central prediction is that new activities concentrate in the region that, initially, contains a greater supply of skilled labor and greater production diversity. I use this prediction to guide the reduced-form estimation strategy described in Section 4. Intuitively, educated workers have the skills to work in these new activities, industrial diversity means that local firms can adopt more new activities, and, in the larger region, households consume and firms produce more new products that use new activities as inputs.[18]

3.1 Setup, preferences, and technology

There are two regions, labeled 1 and 2. Each region is endowed with a non-traded good, supplied inelastically across regions, with quantities h_1 and h_2. (Helpman calls this good housing, though it can be any non-traded good with inelastic supply.) A population of skilled labor L, mobile across regions (l_1 and l_2), supplies labor inelastically to traded goods production. They consume housing services h and differentiated varieties of the traded good x.

There are N total varieties of the traded good, each produced by a separate firm. Further, each variety is produced *using a distinct production activity*. Therefore, there is a one-to-one relationship between the number of traded goods, the number of firms, and the number of production activities. Firms are also mobile across regions, so that $n_1 + n_2 = N$.

Representative household utility U is:

(1) $$U = (h)^{1-\mu}[(\int_{j=0}^{N} x_j{}^{\alpha})^{1/\alpha}]^{\mu},$$

[15] This is one way to motivate growth in new activities on initial conditions. Krugman (1991a), Venables (1996) and Duranton and Puga (2001) present other ways in which the distribution of new activities might be modeled. I base this model on Helpman (1998) because the location of activities *used in production* is more transparent.

[16] Redding and Sturm (2006) use a similar strategy to simulate the effects of German division on the size of cities. Hanson (2005) also uses the Helpman model to examine the effect of market access on agglomeration.

[17] To avoid repetition, many details are relegated to the appendix.

[18] The last part of this intuition echoes the concept of "venturesome consumption" explained by Bhidé (2006). In that paper, he stresses that the *use* or *consumption* of innovation-related outputs matters for the development of new ideas.

where $\sigma \equiv 1 / (1 - \alpha)$ is the constant elasticity of substitution between traded goods varieties, assumed greater than 1. Let $\mu \equiv N / (N + \delta)$, $\delta > 0$, so that the expenditure share devoted to traded goods increases with the number of varieties. This is a key assumption: as N expands, the expenditure share devoted to housing ($1-\mu$) falls. Without it, growth in varieties scales production in each region, failing to deepen agglomeration. Is this assumption plausible? Bils and Klenow (2001) find that variety growth leads to lower expenditure shares on non-innovating sectors (e.g., housing). Note, too, that alternative utility specifications can also generate flexible expenditure shares—a CES aggregator over housing and traded goods, for example.

Production of each variety of traded good is subject to scale economies. This is modeled as a fixed cost f in terms of skilled labor l. Let β be the unit cost in skilled labor, then:

(2) $$l = f + \beta x,$$

where both f and β are assumed greater than zero. After production, there are iceberg transport costs. For each variety, $t > 1$ units must be shipped for 1 unit to arrive in the other region. Region 1 residents pay p_1 for every locally produced variety but tp_2 for varieties imported from region 2.[19]

3.2 Initial equilibria

Profit maximization implies that relative mill prices of the traded good (p_1 / p_2) must be equal to relative wages of skilled labor ($w \equiv w_1 / w_2$). Also, by free entry of firms, equilibrium output for each variety is constant and the same in both regions. It follows that skilled labor demand is equal across regions and varieties; therefore, $n_1 / N = l_1 / L$.

In equilibrium, I am interested in the share of production activities located in region 1. Define this value as $v \equiv n_1 / \mathrm{N}$. In the appendix, I derive two equilibrium conditions. The first relates v, the location of production activities, to w, relative prices and wages. This condition provides a *unique* solution to relative prices and wages w for each distribution of production activities v.[20] Since skilled labor is mobile across regions, a second equilibrium condition requires that household utility is equal across regions. Equilibrium is fully characterized by these two conditions, which determine two endogenous variables, v and w, in terms of parameters μ,σ, t, and h_1/h_2. Because the non-linearity of the model makes it impossible to find closed-form solutions for the equilibrium values of v and w, I solve for these values numerically. I first calculate relative utility $u \equiv u_1/u_2$ for the entire range of values of v, the share of production activities in region 1. In equilibrium, it must be that $u = 1$, or else that all activity concentrates in one region (and $u = 0$ or $u = \infty$).

[19] Unskilled labor is in the background; it is immobile, used in a constant returns to scale technology to produce another traded good (e.g., food). I assume that unskilled workers consume only food, in order to focus on the relationship between the locations of differentiated production activities and skilled labor. For evidence on the (lack of) mobility of unskilled labor, see Borjas et al. (1992) or Bound and Holzer (2000).

[20] The relationship between v and w does not depend on either f or β, which serve only to scale the number of varieties and the level of production output.

Following Helpman, the important parameters for determining the stability and uniqueness of the initial equilibria are μ ($\equiv N / (N + \delta)$), σ, and t. There are two configurations of equilibria. In the first case, a unique, stable equilibrium exists for $\sigma(1 - \mu) = \sigma\delta / (N + \delta) > 1$.[21] These conditions imply a high elasticity of substitution (households substitute easily across varieties), or large expenditure shares devoted to housing. Because households care less about variety and spend more on housing, the agglomerating forces are relatively weak, and the distribution of production activities is a function of the housing stock. I depict this relationship in Figure 4 for the case of $h_1 = 2$ and $h_2 = 1$.[22] The dashed line (labeled "before") indicates relative utility u for values of v and w. The unique, stable equilibrium is at the point of intersection between this line and the solid line indicating $u = 1$. Stability can be verified by noting that increases in the size of region 1, relative to equilibrium, lead to lower relative household utility in region 1.

Multiple stable equilibria exist if $\sigma(1 - \mu)=\sigma\delta / (N + \delta) < 1$ and $1 < \underline{t} < t < \infty$. These conditions imply low elasticity of substitution (households prefer variety) or low expenditure shares on housing, and intermediate transport costs. In this case, agglomerating forces are relatively strong, and the distribution of production activities can concentrate in one region, even conditioned on equal housing stocks. In Figure 5, again, intersections between the dashed line ("before") and the solid line ($u = 1$) mark equilibria. The initial symmetric equilibrium ($v = 0.5$) is unstable, as increases in the size of region 1 from this point raise relative household utility in region 1. The two concentrated equilibria, however, are stable. In each of these, one region initially contains more skilled labor and more production diversity, despite equal endowments of housing. Intuitively, consumers are willing to pay higher prices for housing in order to have access to a wider variety of consumption goods, and firms locate near customers and skilled workers.

3.3 Technological change and discussion of new equilibria

Having solved for configurations of initial equilibria, I now introduce a global innovation shock in the form of an exogenous expansion in the number of traded good varieties (an increase in N).[23] This is a non-unique way to formulate innovation: population growth in skilled labor or an expansion in production activities will have equivalent effects. Given new N (and hence μ), I solve again for equilibrium production share v, wages w, and relative utility u. In Figures 4 and 5, intersections between the dotted lines ("after") and the solid line ($u = 1$) indicate the new equilibria.

---

[21] Or, in a trivial case, when there are no transport costs ($t = 1$).
[22] Since in this case concentration is a function of housing endowments, setting h_1 and h_2 to different values is important for establishing differences in initial conditions. The interpretation here is that the initial equilibrium represents a distribution of skilled labor and production diversity reflecting historical processes. Note that if $h_1 = h_2$, the initial equilibrium is symmetric, with half of the skilled labor force residing in each region.
[23] Details on this simulation and parameter values are in the appendix.

In the unique equilibrium case (Figure 4), innovation deepens concentration. Region 1, which initially contained more skilled labor and production diversity, attracts more new production activities. In the multiple equilibria case (Figure 5), the effect of innovation on the location of production activities is similar. Innovation deepens concentration in the region that contained greater initial supply of skilled labor and production diversity; note that the two stable equilibria shift outwards. Unlike the previous case, however, the presence of multiple equilibria suggests that the economy may switch from a concentration of production activities in one region to a concentration in the other region. How likely is this to happen? Given an historical concentration of skilled workers and production firms in one region, no single firm or worker has an incentive to move. In other words, the concentrated equilibria are likely to be self-reinforcing. This chain of reasoning follows earlier work on regions: Saxenian (1994) argues that initial differences between regions can explain future development, and Krugman (1991b) notes that slow adjustment processes mean that factor rewards across regions persist. The distribution of production activities across regions, then, is likely to remain stable. However, dramatic changes from one possible equilibrium to another are not inconceivable: after all, the Santa Clara Valley was at some point more known for fruit than Apple (computers) and "silicon." Secular migration, region-specific innovations or population shocks each may undo previous patterns of concentration. The multiple equilibria case may be more closely aligned to the historical evidence, but it is still probable that dramatic switches between concentrated equilibria are rare.

Thus, in the two separate equilibria configurations, the model presented here predicts that new activities will appear in regions with greater initial supply of educated workers and industrial diversity. If these initial differences are products of historical processes, then they serve as initial conditions that determine how well a region does in attracting the next round of innovation and new work. It is this prediction that I use to guide the estimation strategy described in the next section.

4 Estimation

In this section, I describe an estimation strategy to assess whether new work is more likely to appear in cities that, initially, have greater supply of educated labor and industrial diversity. Define the outcome of interest, v_i, as the new title share (in all titles) within each worker i's detailed occupation code.[24] For example, workers identified as *network systems and data communication analysts* (DOC 111) have $v_i = 96.7\%$, corresponding to 29 new titles out of 30 within this DOC. I interpret this continuous variable, ranging from 0 to 1, as indicating the likelihood that each worker selects into a new activity that first appeared between 1990 and 2000.[25] For most workers, v_i is zero; all titles in their DOC can be matched to 1990 titles, and they

[24] I read v as the Greek letter "nu."

[25] Because of the aggregation from new 5-digit titles to 3-digit DOCs, it is unobserved whether each worker is actually in a new occupation. Refer again to Table 1 and Figure 1.

are unlikely to have selected into a new activity. I estimate a linear model to predict v_{ij} for each worker i living in city j:[26]

$$v_{ij} = \alpha + x_i\beta + z_j\gamma + \varepsilon_{ij}. \quad (3)$$

Because of the construction of v_i, this equation could also be estimated at the occupation level, with DOCs instead of workers being the unit of observation. However, this strategy would be unable to separately estimate the effects of worker characteristics on selection into new work.

Here, x_i is a vector of worker characteristics and z_j is a vector of initial educational attainment and industrial diversity in metropolitan area j. The focus here is on the location where workers select into new occupations. Using ordinary least squares, I regress v_i on z_j to identify the effect of initial metropolitan area education and industrial diversity on the appearance of new work. (Because z is defined over j, I cluster the standard errors at the metropolitan level.)

As v_{ij} represents occupational outcomes in 2000, z_j uses 1990 levels of education and industrial diversity. To measure initial metro education, I use the 1990 share of college graduates, in all workers, within metropolitan area j. A high value indicates that a metropolitan area has many highly skilled and educated workers. I also separately include shares of workers with some postsecondary education and those with high school diplomas to fully characterize the metro skill distribution. (The high school dropout share is the omitted category.) In contrast to including the mean educational attainment of workers within a metropolitan area, this approach emphasizes returns to higher education, and it also flexibly allows for nonlinear returns to metro education. In 1990, the college share variable ranges from 11.5% (McAllen, Texas) to 31.6% (Raleigh-Durham).

To measure 1990 metro industrial diversity, the second element of z_j, I use the number of identified 3-digit industries within metropolitan area j in 1990.[27] I then normalize this as a share of total 3-digit industries in the U.S., in 1990. A high value (near 1) indicates that a metro area contains many different industries. I classify an identified industry as one employing more than 2,000 workers within metropolitan area j in 1990.[28] In 1990, this variable ranges from 5.1% of industries (Boise, Idaho) to 92.7% (Los Angeles).

The vector x_i contains variables describing characteristics of each worker i. By using Census microdata and including controls in x_i, I can separately identify external returns from composition effects. Using flexible dummy variables, I control for individual educational attainment, sex, race, ethnicity, marital status, nativity, origin, worker class, and age. With indicator variables for high school graduation, some postsecondary education, and college

[26] The linear probability model, which provides easily interpretable estimates, is also effective in generating predicted values between 0 and 1. An alternative imputation strategy for the dependent variable, in which I code a binary variable based on whether v_i is greater than or less than the mean new title share, yields results similar to the ones presented in Section 5.
[27] Other diversity concepts yield similar results; see Section 5.1.
[28] The 1990 Census over-samples certain demographic groups and regions, so this threshold ensures a compatible measure across metropolitan areas. I experiment with different thresholds, with little difference in result.

graduation, I can see which levels of educational attainment are most likely to result in selection into new work.

Controlling for personal characteristics in x_i, I interpret γ, the coefficient on z, as identifying effects of the initial metro supply of educated labor and industrial diversity in creating and attracting new work to regions. The discussion in Section 3 suggests that the estimated effects of education and industrial diversity should be positive. The source of identification comes from variation across 1990 metropolitan areas, which I take as historically determined. In an interpretation based on the model in the last section, initial regional characteristics reflect long-running historical patterns of economic activity. Firms and households then make location decisions based on these patterns.[29]

Proper inference requires the ε_{ij} be uncorrelated to the elements of z_j. I group possible violations of this condition into three categories: (1) unobserved city characteristics, (2) sorting across metropolitan areas based on unobserved worker ability, and (3) measurement error. For the latter, I am concerned mostly with error in measuring new work and industrial diversity; Section 5.5 discusses results using alternative measures.

Unobserved city characteristics may be related to new work, city skill and industrial diversity; this will bias estimates of γ. For example, city size is positively correlated to industrial diversity, and also the appearance of new work; this will lead to a positive omitted variables bias. In all regressions, I include (log) 1990 metropolitan area population and land area. Similarly, in most specifications, I control for city-specific labor demand shocks and labor supply responses. More details on these controls can be found in the next section.

Another important source of bias is the sorting of workers across cities based on unobserved skill.[30] In this case, the estimated effect of college share on new work could be due to sorting on unobservables, rather than increasing returns. In the case of unobserved ability, (3) will overestimate the returns to initial city education and industrial diversity. I use a technique from Evans et al. (1992) to correct for this type of sorting. They argue that workers are less likely to sort on unobservable characteristics at higher levels of geographic aggregation; I therefore use characteristics of a worker's state of residence as instruments for city characteristics. Following a discussion of the main results, Section 5.4 describes these efforts in more detail.

5 Empirical Results

5.1 Main results, metro college share, industrial diversity and skill bias

I find that 1990 college graduate share and industrial diversity positively predict worker selection into new work. The main results are in Table 5; each column is a separate regression. I include

[29] One example where historical patterns determine initial regional characteristics is the Bound et al. (2004) finding that the distribution of college-educated labor is influenced by the locations of colleges, which were determined long ago.
[30] To account for sorting based on *observed* skill, I describe an approach allowing for different effects by education group in the next section.

coefficient estimates for worker educational attainment, metro education and industrial diversity, and other selected worker characteristics. Suppressed coefficients include regional population density. Estimates are reported in percentage point units; an estimate of 6.5 means that a one unit change in the independent variable increases the likelihood of selection into new work (or, specifically, a higher new title share in each worker's identified occupation) by 6.5 percentage points. Means and standard deviations of the independent variables are also reported.

Controlling for other characteristics, I find that a one standard deviation increase in 1990 metropolitan college share increases the likelihood of selection into new work by 0.4-0.6%.[31] This is calculated by multiplying the estimated coefficient (11.0 to 14.9) by the standard deviation in college share across metro areas (0.04). This change, akin to the difference between New Orleans (19.3% college graduates) and Chicago (23.4%), accounts for an increase in selection into new work comparable to the effect of graduating high school relative to dropping out. Over the entire range of observed values in college share (McAllen's 11.5% to Raleigh's 31.6%), this effect is as large as the difference between dropping out of high school and attempting some postsecondary education. The effect of metro college share is precisely estimated, and is consistent with educated cities attracting new work and adapting to innovation.

These results for college share provide a new interpretation for earlier work on the effects of citywide human capital. Rauch (1993) and Moretti (2004) find that workers have higher wages in cities with more skilled workers. Workers also appear to earn more in new occupations, both in the sample data (Table 4) and when controlling for other characteristics in a regression (Section 2.5). To the extent that new work is more productive (or, that workers in new activities enjoy rents for adapting to new technologies quickly), this result suggests that the appearance of new occupations may be an important channel for productivity spillovers observed in previous work.

In column 3, a standard deviation increase in 1990 industrial diversity, as measured by observed 3-digit industries, increases new work share by about 0.6-0.7%. This change, akin to the difference between Dayton (28.8% observed industries out of U.S. total) and San Diego (50.4%), is an increase in new work share comparable to that of metro education. Over the range of observed values in industrial diversity (Boise's 5.1% to Los Angeles's 92.8%), this effect is slighter smaller than the difference between dropping out of high school and attempting some postsecondary education. As an initial robustness check on these estimates, in column 4 I include additional metro characteristics, which I describe in Section 5.2.

Table 5's measure of industrial diversity uses the *appearance* of 3-digit industries within a region, with less emphasis on the relative *size* of each industry. In other words, the appearance of new work is driven by a greater number of observed sectors, rather than the distribution of employment across sectors. Table 6 presents estimates that show this result for different

[31] I also include shares of the city workforce with some college and high school diplomas, though coefficient estimates on these shares are not significantly different from zero. The high school dropout share of the city workforce is the omitted category.

industrial diversity concepts. Depending on the concept used to measure industrial diversity, the estimated effect ranges from 0.2% (employment share of the largest industries) to 0.7% (the observed number of industries) for standard deviation changes in industrial diversity.

Using a lower employment threshold (1,000) for counting 3-digit industries within metropolitan areas, the estimated coefficient is 2.7, implying an increase of 0.6% in new work share in response to a standard deviation increase in industrial diversity. A second diversity concept, a Herfindahl index of 1990 industry employment shares within a city, is calculated as the sum of squares of industry employment shares within each metropolitan area. A maximum value of 1 indicates that all employment within the metropolitan area is concentrated in a single industry. This measure is then inverted so larger values reflect increases in 1990 industrial diversity. Using this measure, the implied increase in new work share is 0.3% for standard deviation increases in industrial diversity.

A third industrial diversity concept is the 1990 share of the top n industries within a metropolitan area, as in Glaeser et al. (1992). Metros with high values near 1 are those with employment concentrated in a few large industries, whereas areas with lower values are those with employment scattered more evenly across industries. This measure is also inverted so larger values reflect increases in 1990 industrial diversity. For $n = 20$, the implied effect of a standard deviation change in industrial diversity is 0.2%; for $n = 50$, the implied standardized effect is 0.7%.

To illustrate the relationship between new work and the college share, I plot average residuals from the regression in column (1), which controls only for personal characteristics, against college share. This graph is Figure 6. (Individual residuals are averaged within each of the 88 metropolitan areas.) The relationship is clearly positive, with only outlier Honolulu having a very high share of college graduates and low likelihood of selection into new work. At the metropolitan level, college share alone accounts for over 50% of the variation in selection into new work left unexplained by worker characteristics (as measured by R-squared from a regression of city-averaged residuals against college share).

Workers may sort across cities based on observable skills. In particular, skilled workers may be drawn to educated cities, further increasing the likelihood that they select into new work. I allow for separate effects of metropolitan college share and industrial diversity based on workers skill by performing the same regression from Table 5, column 3, on four separate samples: college graduates, workers with some college, high school graduates, and high school dropouts. Table 7 displays estimated effects for each of the four education groups. There is evidence that metro college share and industrial diversity matter more for skilled workers; this may be due in part to sorting on observable skills. The effect of the college share is most acute for college graduates, rising about 1.0% for a standard deviation increase in the college share. Standard-deviation increases in college share predict a rise of about 0.5% in new work share for workers with some college. Estimates for industrial diversity echo the college share results; the effects of metropolitan-level variables are much smaller for high school graduates and dropouts. In part, this result can be seen as reflecting the skill bias in new activities. These estimates also support sorting of workers across cities, based on observable skill; sorting into skilled cities may be one way that skilled workers are better able to adapt to new technologies.

I perform a graphical exercise similar to Figure 6 using average residuals and observed 3-digit industries in Figure 7. The relationship is positive as well. This measure of industrial diversity alone explains approximately 20% of the variation in new work unexplained by personal characteristics. While both college share and industrial diversity are important in explaining new work across regions, a comparison of Figures 6 and 7 supports a more central role for metro education. Formal education, being more general, may be a more important form of human capital in generating regional advantage.

In Table 5, the estimated effects of individual educational attainment on selection into new work confirm the skill bias observed in the sample statistics. Controlling for other characteristics, a college graduate is 6.5% more likely to select into new work (that is, select into an occupation with a 6.5% higher new work share) than a high school dropout. Selection appears monotonic in educational attainment.

There is also an important age dimension to participation in new work. Coefficient estimates on age group dummies suggest that worker participation in new occupations peaks in ages 25-30 (0.7% higher than ages 20-25), decreasing through older age groups. For presentation purposes, I have omitted the robust standard errors, but differences in age effects (from the omitted age 20-25 group) on selection into new work are all statistically significant at the 95% level of confidence (except for the age 41-45 group). One interpretation of this result is that older workers, given investments in specific human capital tied to older types of work, are more reluctant to switch into the new types of work that appear following innovation.[32]

In sum, initial metro college share and industrial diversity are important predictors of future selection into new occupations. This is consistent with increasing returns to geographic concentration. A standard deviation change in either characteristic increases the likelihood of selection into new work by about 0.6%. The effect of college share is more precisely estimated and more central to the location of new work, accounting for about 50% of the unexplained variation in new work across metropolitan areas. The evidence so far supports the idea that stocks of human capital help regions to better adopt the new activities that follow innovation. In addition, workers with more educational attainment are more likely to select into new occupations. This is consistent with skill biased technical change.

5.2 Other region-specific characteristics

[32] See also Bleakley and Lin (2006).

These main results are robust to the inclusion of other city-specific characteristics. In most specifications, I include the Blanchard-Katz (1992) labor demand shock index as an additional regressor. This index is a weighted average of industry employment growth, where the weights are metro-specific industry employment shares.

(4) $$\hat{\eta}_j = \sum_k \xi_{jk} \eta_k .$$

The Blanchard-Katz index on the left side, η_j, is the predicted growth in employment for metropolitan area *j*, based on 1990 employment shares ξ of industries *k* in *j*, and the change in log employment η_k in industry *k* between 1990 and 2000. Some metropolitan areas may be fortunate to be specialized, for historical reasons, in particularly fast-growing industries. The Blanchard-Katz index measures idiosyncratic shocks to each metropolitan area based on historical industrial composition. For example, cities with industrial bases specialized in computer products or services in 1990 would have fared well given growth in these industries during the 1990s. High values of the Blanchard-Katz index indicate that in 1990, the city was highly specialized in industries that grew quickly during the 1990s. These shocks almost certainly are correlated with metro education and industrial diversity, and unobserved, may drive the estimates of γ.

As Bound and Holzer (2000) show, metro-specific labor demand shocks may also cause workers to migrate differentially by skill level. These responses may also confound identification of γ. Skilled workers who remain in Midwestern cities as they decline may be more likely to select into new occupations as other skilled workers leave. To control for this, I include the (log) change in metro employment between 1990 and 2000 in worker *i*'s education group.

As an alternative to the Blanchard-Katz index, in some specifications I include a modified index using industry patent activity instead of industry employment growth. Using data from the U.S. Patent and Trademark Office, I match patent counts of 2-digit industries between 1990-1999 to 1990 industry composition. For each metropolitan area, I calculate a weighted average of patent activity by 2-digit industry, using metro industry shares as weights. Formally,

(5) $$\hat{\pi}_j = \sum_k \xi_{jk} \pi_k .$$

As before, ξ_{jk} is the 1990 employment share of industry *k* in metro *j*, and π_k is log patents in industry *k* between 1990 and 1999. I use this to capture innovation shocks to each metropolitan area based on historical industrial composition. As in the case of employment shocks, including this index as an additional regressor is an attempt to mitigate identification issues caused by random historical specialization patterns.

The first two columns of Table 8 contain estimates from regressions including these variables. (Column 1 is the same as Table 5, column 4.) Both the labor demand shock index and the predicted patenting index have estimated coefficients that are positive, which is consistent with fast-growing areas attracting more new work. However, the standard errors are larger in magnitude than the estimates themselves. Also, their inclusion does not appreciably change the estimated effects of metropolitan education and industrial diversity.

There still may be a number of unobserved city characteristics related to the appearance of new work across regions that are correlated with education and industrial diversity. One approach that I take to control for remaining omitted variables is including patent activity as an additional regressor. I use the patent data described in Section 2.5 to control for additional unobserved variables related to knowledge creation. In Table 8, column 3, a higher rate of patenting predicts a higher selection into new work, as expected. A standard deviation change increases selection into new work by about 0.2%. The inclusion of actual patenting does not affect the sign or significance of the estimates for college share and industrial diversity. To the extent the patents can control for remaining unobserved factors related to knowledge creation, omitted variables do not appear to contribute significantly to the estimated effects of the college share and industrial diversity. The estimate from column 3 also illustrates the difference between patents and new work; the location that *originates* a new idea may not be the same location where that idea is *applied*.

Further, metropolitan areas may contain institutions, such as universities, that promote innovation that are also related to elements of z_j. Such an effect might lead to an overestimate of returns to metro education. I include an indicator variable for the presence of a land grant college within that region, which I interpret as measuring local infrastructure relevant to the production of skill.[33] In column 4, the presence of a land grant college does not seem to affect the location of selection into new work, nor does it provide any explanatory power beyond that of the main explanatory variables. The coefficient estimate is not significantly different from zero, with a standard error nearly as large as the point estimate. This suggests that a skilled labor force is what matters for the location of new work, rather than educational institutions themselves.

In columns 5-7 of Table 8, I include different variables measuring other aspects of 1990 metropolitan industry-occupation structures. Selection into new work is negatively related to occupational diversity, as seen in column 1. Though this is precisely estimated, the magnitude of the estimate suggests the relationship is very weak. The total effect over the entire range of observed values of occupational diversity is less than 0.06%, an order of magnitude less than that of industrial diversity. Selection into new work is also negatively related to a worker's own-industry concentration in the 1990 metropolitan area. As the 1990 own city-industry size increases, selection into new work decreases. This effect is precisely estimated and is rather large; a one standard deviation increase in own-industry share decreases the likelihood of selection by 0.6%. This is *not* consistent with within-industry externalities fostering new knowledge and new work. This also makes sense in that other workers with similar industry-specific human capital are locally available for new activities. Further, as the estimated effects of education and

[33] Moretti (2004) uses this variable as an instrument for metro college share. I am hesitant to use college location similarly, because of possible direct institutional effects on the adoption of new work.

industrial diversity are unchanged, this may be further support for the importance of industrial diversity in attracting new work to regions.

5.3 New work in traded goods industries

An alternative explanation is that the estimated values for γ may be driven by migration. The secular movement of people from cities in the northeast and Midwest to those in the southern and western U.S. may affect local demand for goods and services, and, in turn, the appearance of new work. Newer Sunbelt cities might attract high skilled workers and have higher local demand for new goods and services because of the lack of local infrastructure. This may be especially true of types of work associated with the production of non-traded goods and services. Migration trends may drive the location of new work associated with the production of non-traded goods. If this is the case, then the results from Table 5 are due to migration, not increasing returns.

In contrast, occupations associated with traded goods industries will be less tied to these movements. In other words, traded goods industries, facing a national or international market, will be less attracted to growing population centers in the south and west. Regression estimates using only a sample of workers in such industries will be more insulated from the effects of migration. Hopefully, these estimates should confirm the pattern seen in Table 5.

I therefore also identify workers employed in a new type of work, in a traded goods industry, and use this as the dependent variable.[34] I assign a full weight of 1 to the manufacturing and information industries. These industries seem most likely to produce traded goods, and are thus less sensitive to the effects of local infrastructure. I also assign half weight to segments of the following industries that also produce traded goods: wholesale trade, transportation, finance and insurance, professional services, management, higher education, and arts and entertainment.[35]

Table 9 contains three regressions that use selection into new work in traded goods as the dependent variable. Otherwise, they are identical to the baseline estimation. The first column uses new work as defined in Section 2. The next two columns use wider definitions of new work, which are described in Section 5.5 and the appendix. In all cases, the sign and significance patterns match those of the main results. In column 1, standard deviation changes in college share and industrial diversity predict increases in new work share by 0.3% and 0.7%, respectively. To the extent that traded goods production is less sensitive to the secular movement of people to the Sunbelt, the results in Table 9 suggest that migration does not drive the main results.

5.4 Correction for sorting on unobserved characteristics

Another alternative explanation for the main results is that they reflect workers sorting, based on *unobserved* ability, across cities. In this case, estimates of γ are due to omitted variables bias;

---

34 Alternatively, I calculate the national share of workers in each new occupation employed in traded goods industries, and identify respondents employed in new occupations that are mostly concentrated in traded goods industries. Results using these different measures are similar.

35 Variations to this weighting scheme do not affect the main results.

workers select into new occupations because of high unobserved ability, rather than because they live in cities that are initially skilled. This type of sorting on unobserved characteristics would cause an over-estimate of the effect of initial metro education and industrial diversity on the appearance of new work.

To correct for this type of sorting, I implement a simple technique used by Evans et al. (1992). In their application, they consider neighborhood effects on individual outcomes. They use metropolitan characteristics to instrument for neighborhood ones, arguing that the degree of bias due to geographic sorting is less severe at higher levels of aggregation. The validity of this instrument rests on the presence of moving costs from one region to another.[36] In a similar spirit, I use state-level characteristics to instrument for metropolitan characteristics in (3). The results of the instrumental variables estimation are in column 4 of Table 9. These estimates indicate that sorting on unobserved characteristics does not drive the main results.

5.5 Variations in measurement and sample

Different measures of new work in 2000 yield qualitatively similar results. The strategy described in Section 2 identifies the smallest set of occupation titles that are new—at the possible expense of excluding some new activities. A wider definition instead uses the list of 3,000 titles remaining after initial string matching. Its expanded scope may include some titles that do not represent new activities. It may also include activities that are only subtly different in 2000. The medium definition is an attempted compromise; it uses titles remaining from the wide definition list after string matching on three words.

Using these more inclusive definitions of new work, I find that the estimated effects of college share and industrial diversity on worker selection into new work are positive and significantly different from zero. Columns 5 and 6 of Table 9 contain these results for the medium and wide definitions of new work. Qualitatively, the results are similar, though somewhat smaller than the baseline estimates, using the narrow definition of new work, presented in Table 5. Estimates using the medium and wide definitions imply responses of 0.4% and 0.5% in new work share to standard deviation changes in college share and industrial diversity, respectively. The smaller estimated effect may be due in part to increased measurement error using the more liberal definitions of new work. Still, these estimates show that differences in the identification of new work do not significantly affect the main results.[37]

In addition, I narrow the geographic scope of regional effects to the metropolitan area (as opposed to the consolidated metropolitan area in most regressions.) For example, New York City and Stamford, Connecticut, are now considered as separate metropolitan areas. Results in column 7 are similar to the baseline results. In column 8, I restrict the sample to metropolitan areas that

[36] Note that this technique cannot correct for more complex error structures; see Bayer and Ross (2006).

[37] In Section 2 I noted that computer-related occupations seem to be most represented in new work. In fact, one-third of workers in new work are in some computer-related industry. However, this skewness does not drive the results; estimates separating new work in computers from new work in everything else yield similar results for both samples.

are consistently and completely identified (in terms of county composition) in both the 1990 and 2000 PUMS. With 58 metropolitan areas, results are similar to those of the full sample.

The results are not region-specific. The first two columns for Table 10 present estimates for two sub-samples, the eastern and western U.S. In both cases, the estimated effects of educational attainment on selection into new work are nearly identical, and the estimated effects college share and industrial diversity are similar. The point estimate for industrial diversity is slightly larger in the west, but it is imprecisely estimated. This is possibly due to smaller sample size in the west (37 metropolitan areas). In addition, metropolitan areas are defined as sets of counties. In the west, counties are larger, increasing measurement error.

5.6 Long-run effects of skill and industrial diversity; earlier Census years

In Section 4, I suggested that the initial distribution of educated labor and industrial diversity across regions reflects long-running historical processes. If this in fact the case, then earlier measures of regional skill and industrial diversity should also predict worker selection into new work. In the third column of Table 10, I use 1970 metro college share and industrial diversity instead of 1990 values. I find that even though these reflect regional human capital stocks several decades prior to the workers observed in 2000, they still predict the location of new work.

In addition, I perform an analysis using earlier Census data, identifying new work that emerged between 1960-1970 and 1970-1980. This data is matched to 1970 and 1980 Census microdata; as noted in the appendix, the methodology relies on matching 3-digit DOCs rather than 5-digit titles, and is therefore less precise and more unreliable.

I perform several analyses using these data. The first replicates the 1990-2000 estimation using data from both the 1970 and 1980 Census; I use occupation outcomes in 1970 and 1980 matched to metropolitan data in 1960 and 1970. These results are displayed in columns 4 and 5 of Table 10. The sign pattern is very similar; new work in 1970 and 1980 does seem to be skill biased. (In 1970 selection appears most among workers with some college; in 1980 the effect is monotonic.) The sign pattern for college share and industrial diversity is similar, though not all estimates are significantly different than zero. Differences in magnitudes are in large part due to different measurement techniques. Measurement error is also more likely in earlier Census years. However, the general pattern is the same.

A second analysis pools the 1970, 1980, and 2000 Census microdata. With multiple observations per region, I can include region fixed effects. In this way, I can control for constant unobserved metro-specific attributes, related to new work. Identification of γ comes from changes in z_j within metropolitan areas, over time. This strategy makes sense under an interpretation of a slowly evolving historical process with periodic (small) shocks to city skill and industrial diversity just large enough to identify their effects. In these specifications, survey year fixed effects are also included to account for differences in innovativeness and measurement between survey years.

In Table 11, controlling for metropolitan fixed effects does not change the flavor of the main results. (Note that the large magnitude of these estimates relative to Tables 5 and 10 is due to differences in identification strategy across Census years, and the resulting differences in variance in the dependent variable. See the appendix for details.) Controls are the same as in the Table 5 regressions. Here, college share has a similar effect as observed in the cross-section data. A one standard deviation change in college share predicts an increase of 0.8% increase in new work share. A one standard deviation increase in industrial diversity, measured either by the number of observed 3-digit industries or an inverted Herfindahl index of employment across 3-digit industries within a metropolitan area, predicts an increase of 0.3% in selection. These estimates reflect an average relationship, over time, between new work and city skill and industrial diversity. That they appear across time periods and specifications suggests that the main results hold up, even when faced with alternative explanations. I conclude that educated and industrially diverse cities do in fact attract the new activities that follow innovation.

6 Conclusions

In this paper, I find that the initial supply of educated workers and industrial diversity create advantages for regions in attracting new work—the new activities that follow innovation. The main contribution is to more specifically characterize how various forms of human capital help workers, firms, and regions better create, diffuse, or adapt to new knowledge. Further, workers who select into new occupations tend to look successful by other labor market measures—including educational attainment and wages.

New work may have further value as a way to investigate other aspects of innovation. For example, the *Dictionary of Occupational Titles* contains multidimensional characterizations of the skill content of many occupations. The attributes of older work that directly precede selection into new work may give us more insight into the innovation process. In particular, it may be possible to use this data to investigate what *kinds* of industrial diversity (that is, in skill content) matter for the creation of new activities. In addition, new work data can be matched to the Current Population Survey, in order to further understand high-frequency properties of innovation. The data may allow observations of the diffusion of new work over time, which may provide support for the firm relocation model of Duranton and Puga (2001). Finally, the wage premia experienced by workers who select into new work may be important for understanding trends in wage inequality.

Bibliography

A. del Regato, J. (1995). "One Hundred Years of Radiation Oncology," in *Current Radiation Oncology*, Vol. 2, J.S. Tobias and P.R.M. Thomas, eds. New York: Oxford University Press.

Autor, D.H., L.F. Katz and A.B. Krueger (1998). "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113 (4), 1169-1213.

Autor, D.H., F. Levy and R.J. Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration,” *Quarterly Journal of Economics* 118 (4), 1279-1333.
Bacolod, M and B.S. Blum (2006). “Two Sides of the Same Coin: U.S. ‘Residual’ Inequality and the Gender Gap,” working paper.
Bayer, P. and S.L. Ross (2006). “Identifying Individual and Group Effects in the Presence of Sorting: A Neighborhood Effects Application,” NBER working paper 12211.
Berman E., J. Bound and Z. Griliches (1994). “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers,” *Quarterly Journal of Economics* 109 (2), 367-397.
Bhidé, A. (2006). “Venturesome Consumption, Innovation, and Globalization,” working paper. www.bhide.net/publications.html
Bils, M. and P.J. Klenow (2001). “The Acceleration in Variety Growth,” *American Economic Review* 91 (2), 274-280.
Blanchard, O.J., L.F. Katz, R.E. Hall, and B. Eichengreen (1992). “Regional Evolutions,” *Brookings Papers on Economic Activity* 1992 (1), 1-75.
Bleakley, H. and J. Lin (2006). “Thick-Market Effects and Churning in the Labor Market: Evidence from U.S. Cities,” working paper.
Borjas, G.J., S.G. Bronars, and S.J. Trejo (1992). “Self-Selection and Internal Migration in the United States,” *Journal of Urban Economics* 32 (2), 159-185.
Bound, J., J. Groen, G. Kezdi and S.E. Turner. “Trade in University Training: Cross-State Variation in the Production and Stock of College-Educated Labor,” *Journal of Econometrics* 121 (1-2), 143-173.
Bound, J. and H.J. Holzer (2000). “Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s,” *Journal of Labor Economics* 18 (1), 20-54.
Bresnahan, T.F. and M. Trajtenberg (1992). “General Purpose Technologies: ‘Engines of Growth?’” *Journal of Econometrics* 65 (1), 83-108.
Bureau of Labor Statistics (2005). “Proposed Category Systems for 1990-2000 Census Occupations,” Working Paper 383, U.S. Department of Labor.
Duranton, G. and D. Puga (2001). “Nursery Cities: Urban Diversity, Process Innovation, and the Life-Cycle of Products,” *American Economic Review* 91 (5), 1454-1477.
–– (2004). “Micro-Foundations of Urban Agglomeration Economies,” in *Handbook of Regional and Urban Economics*, Vol. 4, J.V. Henderson and J.F. Thisse, eds. North-Holland.
Earle, C., C. Cao, J. Heppen, and S. Otterstrom. 1999. *The Historical United States County Boundary Files 1790 - 1999 on CD-ROM*. Geoscience Publications, Louisiana State University.
Evans, W.N., W. Oates and R. Schwab (1992). “Measuring Peer Group Effects: A Study of Teenage Behavior,” *Journal of Political Economy* 100 (5), 966-991.
Feenstra, R.C. (1994). “New Product Varieties and the Measurement of International Prices,” *American Economic Review* 84 (3), 379-399.
Feldman, M.P. and D.B. Audretsch (1999). “Innovation in Cities: Science-based Diversity, Specialization and Localized Competition,” *European Economic Review* 43, 409-429.
Glaeser, E.L. (1999). “Learning in Cities,” *Journal of Urban Economics* 46, 254-277.
Glaeser, E.L., H.D. Kallal, J.A. Scheinkman, and A. Shleifer (1992). “Growth in Cities,” *Journal of Political Economy* 100 (6), 1126-1152.
Glaeser, E.L. and A. Saiz (2005). “The Rise of the Skilled City,” NBER working paper 10191.
Hall, B. (2005). “Exploring the Patent Explosion,” *Journal of Technology Transfer* 30, 35-48.
Hall, B. H., A. B. Jaffe and M. Tratjenberg (2001). "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." NBER working paper 8498.
Hanson, G.H. (2005). “Market Potential, Increasing Returns, and Geographic Concentration,” *Journal of International Economics* 67 (1), 1-24.
Helpman, E. (1998). “The Size of Regions,” in *Topics in Public Economics*, D. Pines, E. Sadka and I. Zilcha, eds. Cambridge: Cambridge University Press.
Henderson, J.V. (2003). “Marshall’s Scale Economies,” *Journal of Urban Economics* 53, 1-28.

Henderson, J.V., A. Kuncoro, M. Turner (1995). "Industrial Development in Cities," *Journal of Political Economy* 103 (5), 1067-1090.

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

Jaffe, A.B., M. Trajtenberg and R. Henderson (1993). "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics* 108 (3), 577-598.

Krugman, P. (1991a). "Increasing Returns and Economic Geography," *Journal of Political Economy* 99 (3), 483-499.

Krugman, P. (1991b). "History Versus Expectations," *Quarterly Journal of Economics* 106 (2), 651-667.

Lucas, R.E. (1988) "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22, 3-42.

Marshall, A. (1890). *Principles of Economics*. London: Macmillan.

Moretti, E. (2004). "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data," *Journal of Econometrics* 121 (1-2), 2004.

Mokyr, J. (2002). *The Gifts of Athena: Historical Origins of the Knowledge Economy*. Princeton University Press.

Rauch, J.E. (1993). "Productivity Gains From the Geographic Concentration of Human Capital: Evidence from the Cities," *Journal of Urban Economics* 34, 380-400.

Redding, S.J. and D.M. Sturm (2006). "The Costs of Remoteness: Evidence from German Division and Reunification," working paper.

S. Ruggles, M. Sobek, T. Alexander, C.A. Fitch, R. Goeken, P.K. Hall, M. King and C. Ronnander (2004). Integrated Public Use Microdata Series: Version 3.0 [Machine-readable database]. Minneapolis, MN: Minnesota Population Center [producer and distributor]. http://www.ipums.org/

Romer, P.M. (1990). "Endogenous Technical Change," *Journal of Political Economy* 98 (5), S71-S102.

Saxenian, A. (1994). *Regional Advantage*. Cambridge: Harvard University Press.

Schultz, T.W. (1975). "The Value of the Ability to Deal with Disequilibria," *Journal of Economic Literature* 13 (3), 827-846.

U.S. Census Bureau (1950, 1960, 1970, 1980, 1990, 2000). *Alphabetic Index of Occupations and Industries*. Washington, D.C.: U.S. Department of Commerce.

––(1950, 1960, 1970, 1980, 1990, 2000). *Classified Index of Occupations and Industries*. Washington, D.C.: U.S. Department of Commerce.

–– (1968). "Technical Paper 18: Changes Between the 1950 and 1960 Occupation and Industry classifications." Washington, D.C.: U.S. Department of Commerce.

–– (1972). "Technical Paper 26: 1970 Occupation and Industry Classification Systems in Terms of Their 1960 Occupation and Industry Elements." Washington, D.C.: U.S. Department of Commerce.

–– (1989). "Technical Paper 59: The Relationship Between the 1970 and 1980 Industry and Occupation Classification Systems." Washington, D.C.: U.S. Department of Commerce.

–– (2003). "Technical Paper 65: The Relationship Between the 1990 Census and Census 2000 Industry and Occupation Classification Systems." Washington, D.C.: U.S. Department of Commerce.

U.S. Patent and Trademark Office (2000). *United States Patent Grants by State, County, and Metropolitan Area*. Washington, D.C.: U.S. Department of Commerce.

Venables, A.J. (1996). "Equilibrium Locations of Vertically Linked Industries," *International Economic Review* 37 (2), 341-359.

Xiang, C. (2005). "New goods and the Relative Demand for Skilled Labor," *Review of Economics and Statistics* 87 (2), 285-298.

Wu, T. (2006) "Weapons of Business Destruction: How a Tiny Little 'Patent Troll' got BlackBerry in a Headlock," *Slate*, February 6, 2006

Appendix A. Further discussion of new work data

This appendix includes more details about the process of identifying new occupations appearing between 1990 and 2000, the construction of new work data for 1970 and 1980, and other sources of data. I also present the result of the wage regression referenced in Section 2.5.

A.1 Identifying new occupations in 2000

As described in the main text, I identify new occupations in 2000 by comparing 5-digit occupation titles in electronic versions of the 1990 and 2000 *Classified Indexes*. An initial string match, allowing for typographic differences, is the basis for the widest or least strict new work definition. A medium definition of new work comes from matching titles from the wide list on any three common words in 1990 and 2000. The narrow definition of new work, used throughout the paper, comes from consultation with detailed internal Census documentation on 1990-2000 OCS revisions and a manual review of the remaining occupation titles.

I consider the narrow definition to be the closest in spirit and practice to actually identifying new activities that appear between 1990 and 2000. Titles in this list include "web designer," "data recovery planner," "pharmacoepidemiologist" (studies drug outcomes in large populations), "dosimetrist" (determines proper doses in radiation therapy) "AIDS counselor," and "polymerization kettle operator" ("controls reactor vessels to polymerize raw resin materials to form phenolic, acrylic, or polyester resins," according to the *Dictionary of Occupational Titles*). The complete list is available on my website (http://econ.ucsd.edu/~jelin/).

In the text I cite the appearance of DOC 111, *network systems and data communication analysts*, as evidence that new occupations indeed followed actual innovations. Another example is DOC 104, *computer support specialists*, which contains workers who provide technical assistance to users of desktop computers and database software. Desktop computers, such as the IBM PC and Apple //, and commercial database software, such as Oracle and DB2, did not widely appear until the mid 1980s. Clearly, new types of work appeared around this time to support these new innovations. Given the decennial nature of the Census, it seems reasonable that they were first cataloged for Census 2000.

Occupations related to advances in medicine and health also represent another major thread of the new work data, as illustrated by Table 1. DOC 320, *radiation therapists*, includes workers who use radiation to treat a variety of medical conditions. Though this use of radiation has been experimented with since the late 1890s, many major advances in the field have occurred in the period since 1980. These advances include the standardization of dosages, computerized dosimetry, and the use of computerized scans to target specific areas of the body (A. del Regato 1995). These examples provide intuitive verification of the kinds of changes used in this paper.

A.2 Identifying new work in 1970 and 1980

I use new work data from 1970 and 1980 to supplement the main results. Identifying new work in these earlier years is more challenging. Without electronic versions of the *Classified Indexes*, I am forced to rely on the cruder 3-digit occupation codes to generate a list of new occupations. In addition, the complications of taxonomic shifts between successive versions of the OCS are more severe. For example, between 1960 and 1970, "the occupation classification system was enlarged [...] because of requests from data users for more data" (U.S. Census Bureau 2003, p.5). The transition between 1970 and 1980 coincides with an attempt to harmonize the OCS to the Standard Occupation Classification, a multi-agency project. The changes between 1970 and 1980 are more drastic than any of the other transitions. These sorts of changes confound the identification of new occupations created by technological change. Because of the increased possibility of measurement error, I focus instead on the more reliable 2000 data, and use earlier years' results only to supplement the main findings.

I rely on Census documents to eliminate spuriously new DOCs unrelated to innovation. As in 2000, I construct three different sets of criteria for identifying new occupations, with varying strictness. The narrow list attempts to minimize the inclusion of spuriously new activities, while the wide list attempts to minimize the exclusion of actual new activities.

The primary source for identifying new DOCS that appear between 1960-1970 and 1970-1980 is a series of Technical Papers from the Census Bureau. Issued in 1972 and 1989, they detail how respondents in a preceding Census would be classified according the OCS from the subsequent Census, and vice versa. For both transitions, I rely on a table that documents how the OCS in year t would have classified workers in the previous Census, in year $t - 10$.

For a number of DOCs, the Technical Papers indicate that *no* new workers in the previous Census would have been classified in the contemporary DOC. These are the DOCs that I classify as new in the strictness, most narrow sense. In 1970, these DOCs included data processing machine repairers, marine scientists, mathematical technicians, and computer specialists. In 1980, these DOCs included marine engineers and marine life cultivation workers.

Further, the Technical Papers state that virtually all new DOCs that reflect innovation are created from previous “miscellaneous” categories. Therefore, in order to capture new occupations not measured by the narrow definition, I also examine new DOCs that are wholly from miscellaneous categories from the previous OCS. In other words, I isolate contemporary DOCs that would have been wholly classified as “miscellaneous” in the previous Census. This forms the basis of the medium and wide lists of new occupations. Further, I eliminate any DOC that, according to the Technical Paper, would have sustained a decrease in employment or would have already included a large number of workers in the previous Census. This is to discount any obviously spurious categories. The remaining miscellaneous-sourced DOCs are manually divided into two groups to form the wide and medium definitions of new work. In 1970, the list now includes computer programmers; the 1980 list includes computer science teachers, numerical control machine operators, and inhalation therapists.

In Tables A1 and A2, I present lists of new work in 1970 and 1980 under both narrow and medium definitions. (The 1980 list, which is longer, contains only selected occupations from the medium definition.) Computer-related occupations (computer programmers and systems analysts) emerge in 1970 from the miscellaneous professional categories of 1960. The 1970 list also includes types of work related to math and science (health practitioners, marine scientists, and mathematical technicians), as well as social science and policy (sociologists, political scientists, and welfare aides). A number of new occupations in the 1980 list also reflect scientific and technical advancement (agricultural and nuclear engineers, computer science teachers, communications equipment operators, and marine life cultivation workers).

After identifying new occupations, I create a variable in the 1970 and 1980 PUMS indicating whether a worker is employed in a new occupation. Table A3 displays the share of the 1970 and 1980 labor force employed in new work, for each definition. Note that the less precise identification strategy results in estimated shares significantly lower than in 2000. Also, both changes in innovativeness and taxonomy are conflated into changes in new work share over time. The bottom panel displays the share of the 1970 and 1980 labor force employed in new work, by education group. Both display similar skill bias as in 2000.

A.3 Data description

This study uses data from a number of sources, summarized in Table A4. The main body of data is the Integrated Public Use Microdata Data Series (Ruggles and Sobek et al., 2004). This data contains the person-level data used in the estimation. I use the 2000 1% sample and the 1970 and 1980 1% metro samples. The 2000 sample is nonrandom and requires the use of weights. In addition, in the 2000 1% sample, some metropolitan areas are incompletely identified. Where metropolitan areas are incompletely identified in the 1% sample but completely identified in the 5% sample, I use the 5% data, taking care to re-weight observations. Table 9 contains estimates using only a sample of completely and consistently identified metropolitan areas.

Metropolitan area data comes from a variety of sources. I define metropolitan areas using the consolidated definition created by the Office of Management and Budget in 1999. The affected consolidated metropolitan areas are Boston, Buffalo, Chicago, Cincinnati, Cleveland, Dallas, Denver, Detroit, Hartford, Houston, Los Angeles, Miami, Milwaukee, New York, Philadelphia, Pittsburgh, Portland, Providence, Raleigh-Durham, Sacramento, San Francisco, Seattle, and Washington, DC. I use land area data from the Historical U.S. County Boundary Files (Earle et al., 1999). Patent data for 1990-19999 comes from the U.S. Patent and Trademark Office (2000). Historical patent data comes from the National Bureau of Economic Research (Hall et al. 2001). Data on metropolitan college share comes from the State of the Cities Data Systems, maintained by the Department of Housing and Urban Development. This is a convenient source for metropolitan area data available from Censuses between 1970 and 2000. Measures of industrial diversity are calculated from the State of the Cities, as well as the IPUMS 1950, 1970, and 1990. Data on land grant colleges comes from Moretti (2004). Table A5 displays summary statistics for both metropolitan area worker characteristics, for most of the variables of interest.

A.4 Wage regression results

Table A6 contains wage regression results that I referred to in Section 2.5. Log hourly wage is imputed for workers based on total annual wage income divided by weeks worked and usual hours per week worked. Even when controlling for educational attainment, experience, and other demographic characteristics, the estimated coefficient on the new work variable is positive and statistically significant, implying a wage premium greater than 35%. This premium may represent both productivity advantages and rents earned by workers who more quickly adapt to new technologies.

Appendix B. Details on theory and simulation results

This appendix contains details on the theoretical model and simulation that were not presented in the main text.

B.1 Equilibrium conditions

Profit maximization for each manufacturing variety yields the price in region *i* for each locally produced variety, equal to a constant markup over marginal cost:

(B1) $p_i = \left(\frac{\sigma}{\sigma - 1}\right)\beta w_i,$ in each region i=1,2.

Profit maximization with free entry (zero profits, $p_i = (f/x + \beta)w_i$) implies that equilibrium output for each variety is constant, and the same in both regions:

(B2) $x = x_i = (f/\beta)(\sigma - 1).$

The production function (2) implies labor demand in region *i* = $(f + \beta x)n_i$, where n_i is the number of varieties of the traded good that are manufactured in region *i*. Since labor demand equals labor supply in each region, the number of varieties, and hence the variety of activities, produced in region *i* is proportional to the amount of skilled labor in region *i*.

(B3) $n_i = l_i/(f\sigma)$

Each resident of region *i* pays p_i for every locally produced traded good and tp_j (t>1) for every brand imported from region *j*. Aggregate demand for each variety produced in region 1 should equal total supply for each variety of traded good (from (1), (B2)):

(B4) $\frac{1}{\beta} f(\sigma - 1) = \frac{p_1^{-\sigma}}{n_1 p_1^{1-\sigma} + n_2 (tp_2)^{1-\sigma}} \mu e_1 + \frac{t(tp_1)^{-\sigma}}{n_1 (tp_1)^{1-\sigma} + n_2 p_2^{1-\sigma}} \mu e_2.$

Each worker/consumer spends fraction (1-μ) on housing, therefore aggregate value of housing services is (1-μ)(e_1+e_2). Aggregate income is labor income plus income from housing, $w_1 l_1 + w_2 l_2 + (1 - \mu)(e_1 + e_2)$. Assume that housing stocks are equally owned by all workers, then total spending by residents of region *i* equals

(B5) $$e_i = w_i l_i + \frac{l_i}{l_1 + l_2}\frac{1-\mu}{\mu}(w_1 l_1 + w_2 l_2).$$

Define $v \equiv n_1/N \equiv n_1/(n_1 + n_2)$ (the share of production activities located in region 1) = l_1/L (by (B3)). Define $w \equiv w_1/w_2$ (the wage in region 1 relative to the wage in region 2) = p_1/p_2 (by (B1)). By substituting (B1), (B3), and (B5) into (B4) I obtain the first equilibrium condition:

(B6) $$1 = \frac{vw^{1-\sigma}}{vw^{-\sigma} + (1-v)t^{1-\sigma}}[\mu + (1-\mu)(v + \frac{1-v}{w})] + \frac{(1-v)(tw)^{1-\sigma}}{v(tw)^{1-\sigma} + 1 - v}[\frac{\mu}{w} + (1-\mu)(v + \frac{1-v}{w})]$$

Equation (B6) relates v, share of production activities, to w, relative wages and prices. Skilled labor is mobile, so utility levels must be equal across regions in equilibrium:

(B7) $$(\frac{h_1}{l_1})^{1-\mu}\left(\frac{\mu E_1}{l_1[n_1 p_1^{1-\sigma} + n_2(tp_2)^{1-\sigma}]^{1/(1-\sigma)}}\right)^{\mu} = (\frac{h_2}{l_2})^{1-\mu}\left(\frac{\mu E_2}{l_2[n_2 p_2^{1-\sigma} + n_1(tp_1)^{1-\sigma}]^{1/(1-\sigma)}}\right)^{\mu}.$$

Substitute (B1), (B3), and (B5) into (B7) to yield relative utility $u \equiv u_1/u_2$:

(B8) $$u = 1 = \left(\frac{h_1}{h_2}\frac{1-v}{v}\right)^{1-\mu}\left(\frac{\mu w + (1-\mu)(vw + 1 - v)}{\mu + (1-\mu)(vw + 1 - v)}\right)^{\mu}\left(\frac{vw^{1-\sigma} + (1-v)t^{1-\sigma}}{v(tw)^{1-\sigma} + 1 - v}\right)^{\mu/(\sigma-1)}.$$

Equations (B6) and (B8) determine equilibrium values of v and w.

B.2 Simulation

I simulate innovation by expanding the number of production activities (increasing N). I take values of σ and μ from the literature. Following Redding and Sturm (2006), I start with Feenstra's (1994) value of $\sigma = 4$, and approximate expenditure share on housing from the Bureau of Labor Statistics' value of $(1 - \mu) = 1/3$. I set the initial N to 30, which is the approximate number of occupation titles (in thousands) in the 1990 Census OCS. $N = 30$ and $\mu = 2/3$ imply $\delta = 15$. I simulate a 10% increase in the number of activities, so that $\Delta N = +3$. This corresponds to a decrease in the expenditure share devoted to housing from 0.33 to 0.31 (μ goes from 0.67 to 0.69).

In Figure 4, the unique equilibrium configuration, I set $h1 / h2 = 2$, $\sigma = 4$, and $t = 6$. In this case, the relative housing stock is chosen so as to generate an initial concentration of production activities and skilled labor in region 1. The other parameters are set only to satisfy $\sigma(1 - \mu) > 1$ and so that the changes in N will be clearly visible on the graph. In Figure 5, the multiple equilibria configuration, I set $h1 / h2 = 1$, $\sigma = 2$, and $t = 4$. In the first case, production activity share in region 1 is 88% before technological change and 92% following the expansion of activities. In the second case, in the rightmost equilibrium, v goes from 85% to 92%.

Appendix C. Figures and Tables

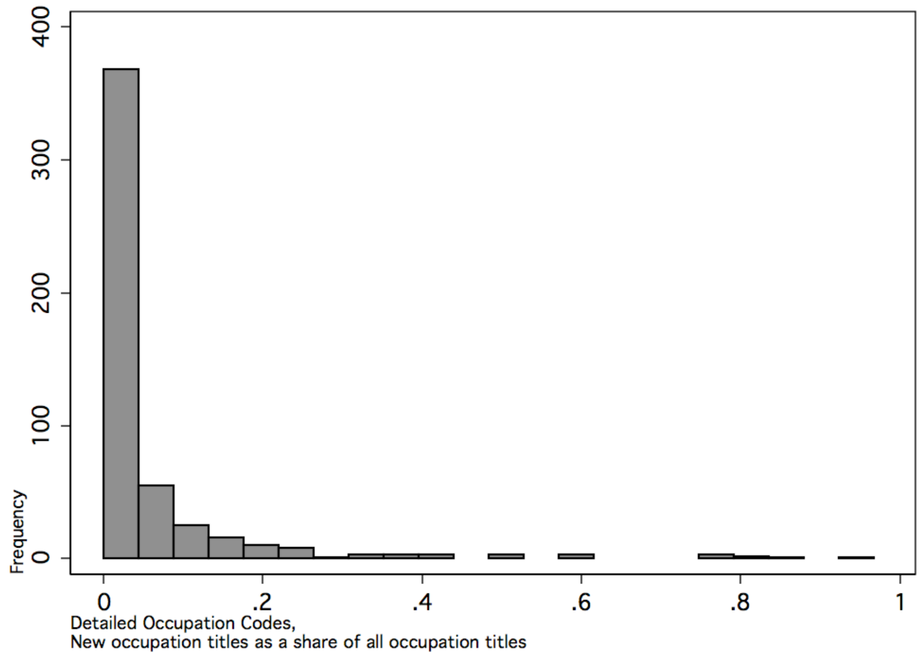

Figure 1. How "new" are (3-digit) Detailed Occupation Codes (DOCs)?
Relative frequency of DOCs by share of new titles in all titles (5-digit)

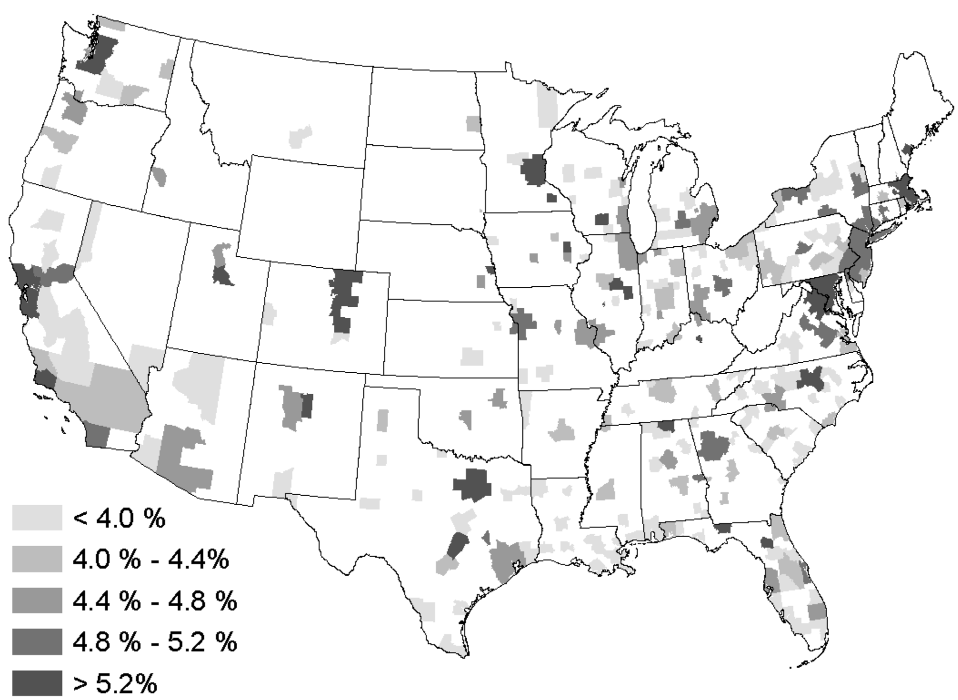

Figure 2. New work as share of local labor force, U.S. metropolitan areas, 2000

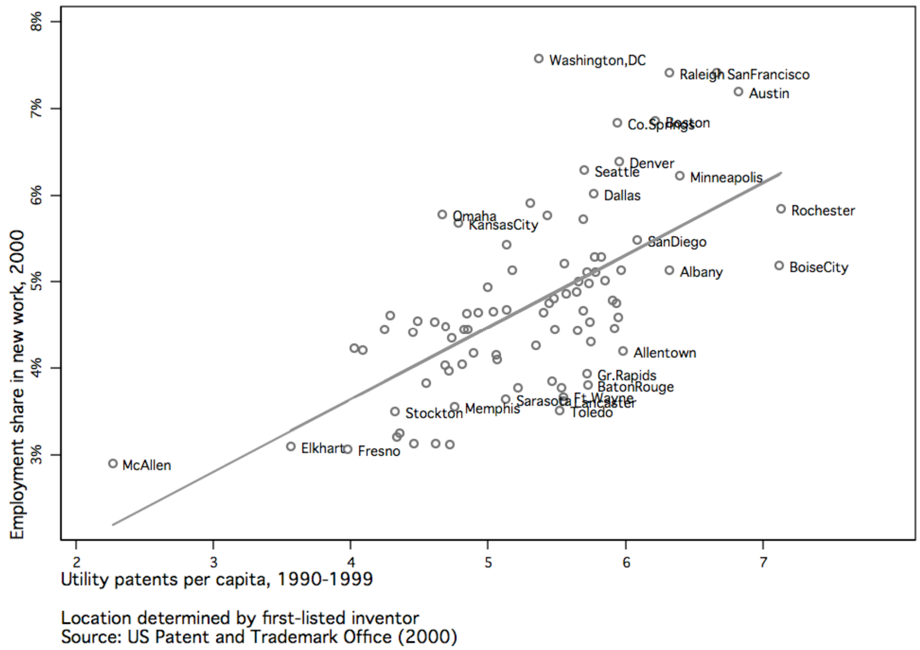

Figure 3. U.S. consolidated metropolitan areas:
Employment share in new work in 2000 and utility patents per capita, 1990-1999

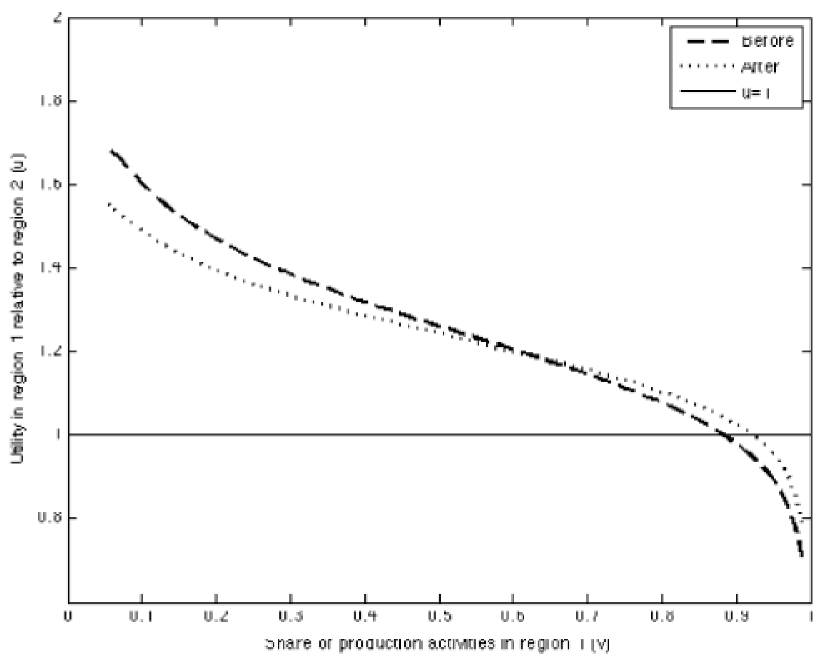

Figure 4. Innovation with unique stable equilibrium
$h_1/h_2 = 2$, $\sigma = 4$, $t = 6$, μ goes from 0.67 (before) to 0.69 (after)

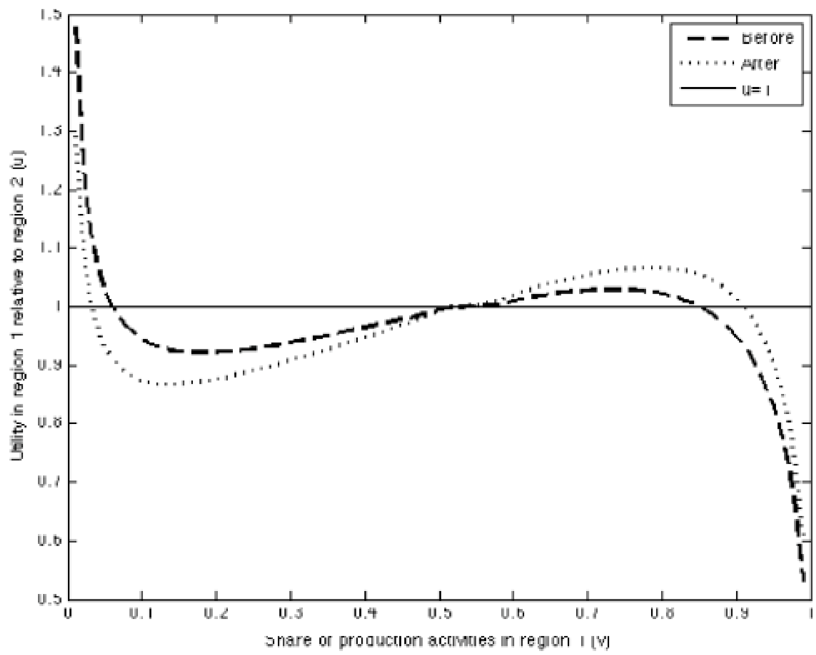

Figure 5. Innovation with multiple stable equilibria
$h_1/h_2 = 1$, $\sigma = 2$, $t = 4$, μ goes from 0.67 (before) to 0.69 (after)

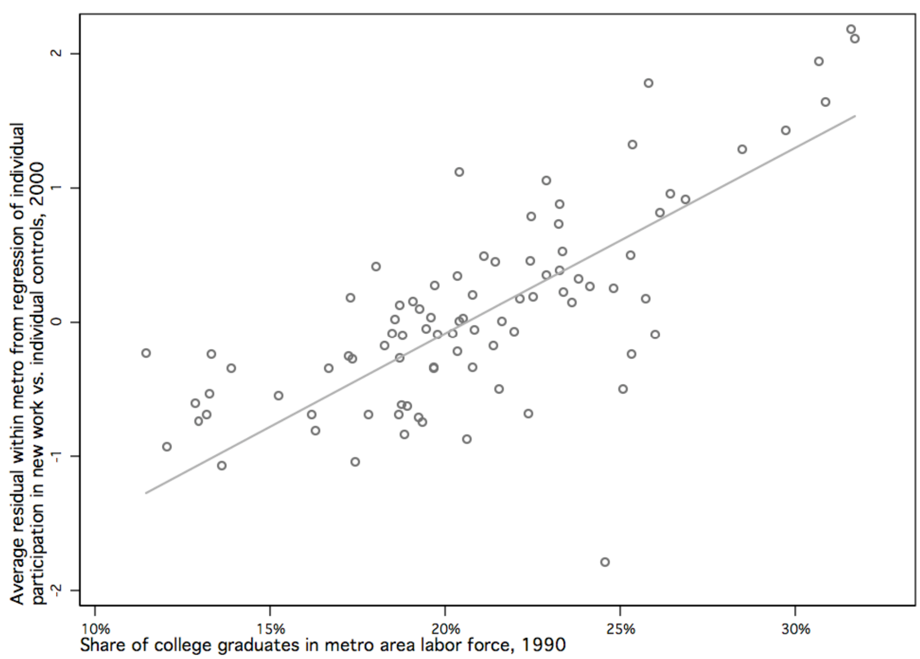

Figure 6. Selection into new work in 2000 and 1990 metro area college share

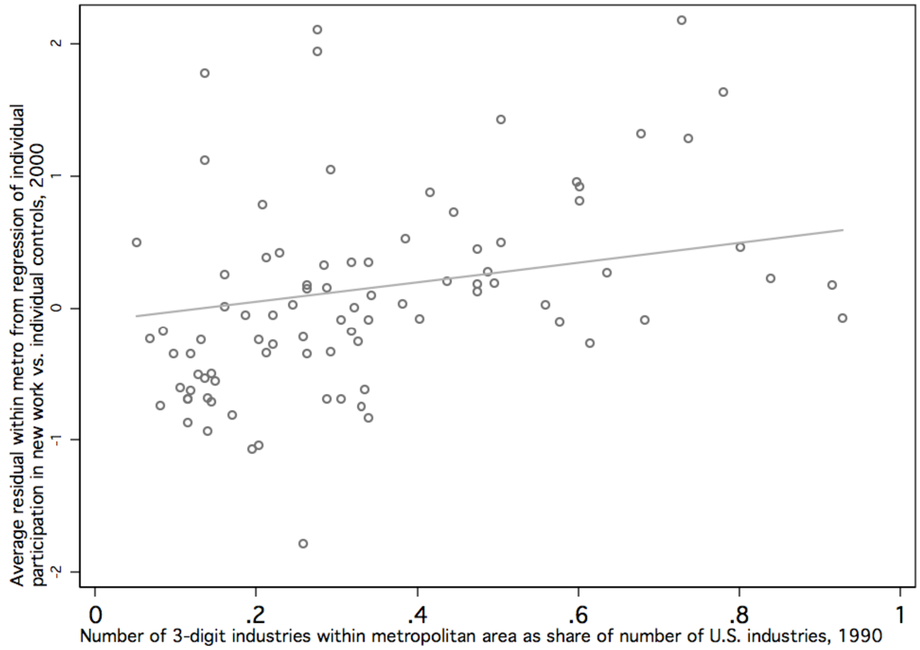

Figure 7. Selection into new work in 2000 and 1990 metro industry diversity

Table 1. Detailed occupation codes in 2000 containing the most new titles

These are the shares of new occupation titles in all titles for each DOC. New titles are identified by manual comparison of 1990 and 2000 OCS. Titles are 5-digit and DOCs are 3-digit classifications. There are 181 additional DOCs with at least one new title, and there are 285 additional DOCs with zero new titles, as in Figure 1.

DOC	*Share*	*Description*
111	0.97	Network Systems and Data Communication Analysts
104	0.86	Computer Support Specialists
110	0.83	Network and Computer Systems Administrators
102	0.80	Computer Software Engineers
106	0.77	Database Administrators
11	0.76	Computer and Information Systems Managers
320	0.75	Radiation Therapists
583	0.60	Desktop Publishers
101	0.59	Computer Programmers
134	0.57	Biomedical Engineers
194	0.50	Nuclear Technicians
70	0.50	Logisticians
140	0.50	Computer Hardware Engineers
316	0.43	Physical Therapists
201	0.41	Social Workers
142	0.39	Environmental Engineers
354	0.36	Other Healthcare Practitioners and Technical Occupations
650	0.35	Reinforcing Iron and Rebar Workers
672	0.33	Hazardous Materials
705	0.33	Electrical & Electronics Installers and Repairers, Transportation Equipment
452	0.33	Miscellaneous Personal Appearance Workers
332	0.32	Diagnostic Related Technologists and Technicians
73	0.27	Other Business Operations Specialists
182	0.26	Psychologists
430	0.25	First-Line Supervisors/Managers of Gaming Workers
230	0.25	Preschool and Kindergarten Teachers
133	0.25	Agricultural Engineers
100	0.25	Computer Scientists and Systems Analysts
291	0.24	Photographers
200	0.23	Counselors
72	0.22	Meeting and Convention Planners
712	0.20	Electronic Home Entertainment Equipment Installers and Repairers
174	0.20	Environmental Scientists and Geoscientists
322	0.20	Respiratory Therapists
90	0.20	Financial Examiners
474	0.20	Counter and Rental Clerks
711	0.19	Electronic Equipment Installers and Repairers, Motor Vehicles
84	0.18	Financial Analysts
181	0.18	Market and Survey Researchers

Table 2. Employment in new work, 2000

Share of employment (%) in new occupations. I multiply new title shares by employment in each DOC (see Table 1). Data: Census PUMS 2000, workers age 20-65 in identified occupations.

United States, entire sample	4.8%
By urban status	
Metropolitan areas	5.3
Rural	3.7
By gender	
Women	4.6
Men	4.9
By age	
< 26	4.3
26-30	5.6
31-40	5.2
41-50	4.8
51 <	4.3

Table 3. Employment in new work by educational attainment, 2000

Share of employment (%) in new occupations, by educational attainment

Educ. attainment	*Share (%)*	*Top new DOCs*
Less than high school	1.7	Misc. personal appearance workers Security guards and gaming surveillance officers Nursing, psychiatric, and home health aides
High school graduate	2.5	Computer support specialists Secretaries and administrative assistants Network systems and data communication analysts
Some college	4.8	Computer support specialists Network systems and data communication analysts Computer software engineers
College graduate	8.6	Computer software engineers Computer programmers

Table 4. Characteristics of workers in new work and others

Sample mean characteristics of workers in new occupations and others, with t-statistic from test on equality of means.

	New work	Old work	t
Average education, years	14.8	13.2	385.4 *
College graduate share	0.47	0.21	338.2 *
Log hourly wage	2.85	2.56	246.3 *

** - Statistically significant at the 99% level of confidence.*

Table 5. Main results: Estimated effects on selection into new work

	[mean] [(s.d.)]	(1)	(2)	(3)	(4)
1990 Metro area characteristics					
College share (see below)	[0.21]	-	14.9	13.8	11.2
	[(0.04)]		(1.4) *	(1.8) *	(1.7) *
No. of 3-digit industries	[0.34]	-	-	2.7	3.5
as share of U.S. total	[(0.21)]			(1.2) *	(0.8) *
Other metro characteristics		-	-	-	YES
Worker characteristics					
College graduate	[0.27]	6.0	6.0	6.0	6.0
(high school dropout omitted)		(0.3) *	(0.2) *	(0.2) *	(0.2) *
Some college	[0.31]	2.9	3.1	3.1	3.1
		(0.2) *	(0.1) *	(0.1) *	(0.1) *
High school graduate	[0.28]	0.7	0.8	0.8	0.8
		(0.0) *	(0.1) *	(0.1) *	(0.1) *
Male	[0.51]	0.6 •	1.0 *	1.0 *	1.0 *
Black	[0.12]	-0.2 •	-0.5 *	-0.5 *	-0.5 *
Asian	[0.05]	1.4 *	1.2 *	1.2 *	1.3 *
Hispanic	[0.12]	-0.3 *	-0.5 *	-0.5 *	-0.5 *
Age 26-30 (20-25 omitted)		0.6 *	0.7 *	0.6 *	0.7 *
31-35		0.6 *	0.6 *	0.6 *	0.6 *
36-40		0.4 *	0.3 *	0.3 *	0.3 *
41-45		0.1	-0.1	-0.1	-0.1
46-50		-0.3 •	-0.4 *	-0.4 *	-0.4 *
51-55		-0.4 *	-0.6 *	-0.6 *	-0.6 *
56-60		-0.6 *	-0.9 *	-0.9 *	-0.9 *
61+		-0.9 *	-1.2 *	-1.2 *	-1.2 *
R-squared		0.04	0.05	0.05	0.05

** - Statistically significant at the 99% level of confidence; • - 95% level. (Robust standard errors, clustered on metro area, in parantheses.) Each column is a separate regression, using Census weights. Dependent variable is selection into a new occupation (range 0-100, sample mean 5.3), based on new title share of identified DOC. Data: Census PUMS 2000, age 20-65, in identified occupations and metro areas. Number of observations = 2.2 million in regression 1; 1.5 million in regressions 2-4. Number of (consolidated) metro areas = 88. Omitted categories are high school dropout and 1990 metro dropout share. Additional controls for marital status, class of worker, nativity included in all regressions. Regressions 2-4 include metro some college and high school share, log population density. These coefficient estimates are not significantly different from zero. Regression 4 includes controls for labor demand shock index and others as described in text.*

Table 6. Alternative industrial diversity concepts

	[mean] [(s.d.)]	Coeff. estimate	
Measure of industrial diversity			
No. of 3-digit industries with. empl. > 2,000	[0.34] [(0.21)]	3.5 (0.8)	*
No. of 3-digit industries with empl. > 1,000	[0.49] [(0.21)]	2.7 (0.8)	*
Herf. index of empl. across 3-dig. industries (inverted, X 10)	[46.8] [(7.3)]	0.04 (0.02)	•
Share in top 20 3-digit industries (inverted)	[2.99] [(1.51)]	0.11 (0.04)	*
Share in top 50 3-digit industries (inverted)	[2.00] [(0.75)]	0.95 (0.38)	•

** - Statistically significant at the 99% level of confidence; • - 95% level. Each row is a separate regression, using Census weights, containing controls as in Table 5, column 4. Robust standard errors, clustered on metro area, in parantheses.*

Table 7. Decomposition by worker educational attainment

	Sample decomposed by worker education							
	College graduates		Some college		High school graduates		High school dropouts	
1990 Metro area characteristics								
College share	24.5 (3.1)	*	12.2 (1.5)	*	3.5 (0.8)	*	-1.8 (0.8)	*
No. of 3-digit industries as share of U.S. total	4.9 (1.3)	*	4.1 (1.3)	*	0.2 (0.5)		1.2 (0.5)	•

** - Statistically significant at the 99% level of confidence; • - 95% level. Each column is a separate regression using Census weights, containing additional controls as in Table 5, column 4. Robust standard errors, clustered on metro area, in parantheses.*

Table 8. Selection into new work, with additional metro controls

	Table 5 Column 4	(2)	(3)	(4)	(5)	(6)	(7)
1990 Metro area characteristics							
College share	11.2 (1.7) *	11.0 (1.7) *	8.9 (1.9) *	11.2 (1.6) *	12.1 (1.5) *	10.8 (1.7) *	11.2 (1.8) *
No. of 3-digit industries as share of U.S. total	3.5 (0.8) *	3.7 (0.8) *	3.0 (0.8) *	3.9 (1.0) *	3.7 (0.8) *	3.1 (0.8) *	3.2 (0.8) *
Labor demand shock	3.8 (5.2)	1.2 (4.1)	3.9 (5.4)	4.5 (5.6)	1.7 (5.3)	4.7 (5.6)	4.5 (5.3)
Predicted patent activity	0.08 (0.11)	-	-	0.09 (0.11)	0.13 (0.11)	0.08 (0.11)	0.08 (0.11)
Actual patent activity	-	-	0.27 (0.13) •	-	-	-	-
Presence of land grant college	-	-	-	-0.11 (0.10)	-	-	-
Herf. index of empl. across 3-dig. occupations (inverted)	-	-	-	-	-0.02 (0.01) •	-	-
No. of 3-digit occupations as share of U.S. total	-	-	-	-	-	0.70 (0.66)	-
Own 1990 1-digit industry concentration in metro	-	-	-	-	-	-	-2.86 (0.16) *
R-squared	0.05	0.05	0.05	0.05	0.05	0.05	0.05

** - Statistically significant at the 99% level of confidence; • - 95% level. Each column is a separate regression using Census weights, containing additional controls as in Table 5, column 4. Robust standard errors, clustered on metro area, in parantheses.*

Table 9. Robustness checks

	New work in traded goods industries				*Measurement and sample*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		Broader new work defns. Medium	Wide	Sorting correction	*Broader new work defns.* Medium	Wide	Metro areas	Consistently identified
Worker educational attainment								58 metros
College graduate	2.7	3.1	3.5	6.5	5.6	7.2	6.4	6.3
	(0.1) *	(0.2) *	(0.2) *	(0.2) *	(0.2) *	(0.2) *	(0.2) *	(0.3) *
Some college	1.4	1.7	1.9	3.2	2.7	3.6	3.1	3.0
	(0.1) *	(0.1) *	(0.1) *	(0.1) *	(0.1) *	(0.1) *	(0.1) *	(0.1) *
High school graduate	0.36	0.5	0.6	0.76	0.61	0.95	0.76	0.74
	(0.04) *	(0.1) *	(0.1) *	(0.06) *	(0.06) *	(0.06) *	(0.06) *	(0.07) *
1990 Metro area characteristics								
College share	7.4	5.9	6.5	11.0	9.4	9.5	10.9	10.3
	(1.3) *	(1.8) *	(2.0) *	(3.8) *	(1.7) *	(1.8) *	(1.1) *	(2.1) *
No. of 3-digit industries as share of U.S. total	3.5	5.4	6.0	3.6	2.6	2.4	1.7	2.7
	(0.8) *	(1.2) *	(1.4) *	(1.4) •	(0.8) *	(0.8) *	(0.7) •	(0.8) *
R-squared	0.03	0.03	0.04	0.05	0.03	0.04	0.05	0.05

** - Statistically significant at the 99% level of confidence; • - 95% level. Each column is a separate regression using Census weights, containing additional controls as in Table 5, column 4. Robust standard errors, clustered on metro area, in parantheses. New occupations defined by: (col. 3, 6) exact string match from comparison of 1990 and 2000 occupation titles, (col. 2, 5) three-word match, or (other columns) manual inspection.*

Table 10. Census regions and data from earlier Census years

	Census 2000						*Earlier Census Years*			
	East		West		Using 1970		1970		1980	
Worker educational attainment										
College graduate	6.5		6.3		6.0		0.06		0.85	
	(0.2)	*	(0.4)	*	(0.2)	*	(0.02)	*	(0.06)	*
Some college	3.3		3.0		3.1		0.11		0.33	
	(0.2)	*	(0.2)	*	(0.2)	*	(0.01)	*	(0.03)	*
High school graduate	0.8		0.8		0.8		0.05		0.10	
	(0.1)	*	(0.1)	*	(0.1)	*	(0.01)	*	(0.03)	*
Lagged metro area characteristics										
College share	13.3		11.4		21.5		0.37		0.90	
	(2.6)	*	(2.1)	*	(2.5)	*	(0.17)	*	(0.81)	
No. of 3-digit industries as share of U.S. total	2.9		5.6		1.8		0.04		0.32	
	(0.8)	*	(2.6)	•	(0.9)	•	(0.04)		(0.22)	
R-squared	0.05		0.05		0.05		0.001		0.003	

** - Statistically significant at the 99% level of confidence; • - 95% level. Each column is a separate regression, using Census weights, containing controls as in Table 5, column 4. Robust standard errors, clustered on metro area, in parantheses.*

Table 11. Pooled results including region fixed effects, 1970-2000

	(1)		(2)		(3)	
Lagged metro area characteristics						
College share	18.9		21.2		21.4	
	(3.5)	*	(3.5)	*	(3.3)	*
No. of 3-digit industries as share of U.S. total	-		1.6		-	
			(0.8)	•		
Herf. index of empl. across 3-dig. ind. (inverted)	-		-		0.03	
					(0.01)	•
Adj. R-squared	0.08		0.08		0.08	

** - Statistically significant at the 99% level of confidence; • - 95% level. Each column is a separate regression, using Census weights. Robust standard errors, clustered on metro area, in parantheses. Number of observations = 2.1 million; number of metro areas = 111. Omitted categories are high school dropout and metro dropout share. Additional controls as in Table 5, column 3, plus metropolitan and census year fixed effects.*

Table A1. New DOCs, 1960-1970

Panel A. New DOCs, narrow definition

DOC	*Description*
52	Marine Scientists
73	Health practitioners, n.e.c.
92	Political Scientists
94	Sociologists
131	Home economics teachers
156	Mathematical technicians
311	Social welfare clerical assistants
475	Data processing machine repairmen

Panel B. New DOCs, medium definition

DOC	*Description*
3	Computer programmers
4	Computer systems analyst
5	Computer specialists, n.e.c.
202	Bank officers and financial records managers
954	Welfare service aides

There are 30 additional new DOCs under the wide definition.

Table A2. New DOCs, 1970-1980

Panel A. New DOCs, narrow definition

DOC	*Description*
54	*Agricultural engineer*
353	*Communications equipment operators, n.e.c.*
483	*Marine life cultivation workers*
489	*Agricultural products inspector*
794	*Hand grinding and polishing occupations*
833	*Marine engineer*

Panel B. New DOCs, medium definition

DOC	*Description*
49	*Nuclear engineer*
98	*Inhalation Therapists*
129	*Computer Science Teachers*
158	*Special education teacher*
714	*Numerical control machine operators*

There are 18 additional new DOCs under the medium definition.
There are 26 additional new DOCs under the wide definition.

Table A3. Employment in new work across Census years

Panel A. Share of employment (%) in new occupations. Data: IPUMS, workers age 20-65 in identified occupations.

	1970	1980	2000
(1) Narrow definition	0.06	0.02	4.8
(2) Medium definition	0.89	0.72	11.2
(3) Wide definition	2.84	3.19	13.1

Panel B. Share of employment (%) in new occupations, by education group

	1970 narrow	1980 medium	2000 narrow
Less than high school	0.01	0.42	1.7
High school graduate	0.05	0.55	2.5
Some college	0.11	0.74	4.8
College graduate	0.10	1.20	8.6

Table A4. Summary of data sources and variables

Variable	*Source*
New work	Census Technical Papers; Census Alphabetical and Classified Index of Occupations and Industries
Individual and metro area data	PUMS 1970, 1980, 2000 (Ruggles et al., 2004)
Other metro area data	State of the Cities Data System (HUD)
Land area	HUSCO (Earle et al. 1999)
Land grant colleges	Moretti (2004)
Patents	NBER patent data (Hall et al. 2001), U.S. Patent and Trademark Office (2000)

Table A5. Sample summary statistics			
	1950	*1970*	*1990*
Metropolitan areas			
Number of metro areas	96	103	88
Share of labor force ... w/ college degree	0.096 (0.030)	0.134 (0.033)	0.209 (0.044)
... w/ some coll.	0.102	0.138	0.264
... w/ HS diploma	0.224	0.364	0.295
3-digit industries ... as share of U.S. total	0.181 (0.215)	0.199 (0.215)	0.337 (0.212)
... Herfindahl index (inverted)	29.7 (11.2)	50.0 (12.4)	46.8 (7.3)
	1970	*1980*	*2000*
Workers			
Number of individuals	553,555	817,240	1,541,623
Men	0.543	0.539	0.514
Blacks	0.097	0.102	0.115
Asians	0.008	0.017	0.048
Hispanics	0.039	0.057	0.115
Married	0.756	0.673	0.601
Self-employed	0.082	0.084	0.095
Foreign-born	0.053	0.067	0.157
College graduates	0.127	0.183	0.273
Some college	0.165	0.235	0.311
HS graduates	0.338	0.355	0.275

Data: IPUMS 1950, 1970, 1980, 1990, and 2000, and SOC. Metropolitan sample is aggregated from all respondents, age 20-65, in the identified metropolitan areas. Sample: all occupation-identified workers, age 20-6, in metropolitan areas identified from 1950 to 2000. Standard deviations in parantheses.

Table A6. Log hourly wages and employment in new work						
	(1)		(2)		(3)	
New work	-		0.389		0.360	
			(0.000)	*	(0.000)	*
College graduate	0.768		0.742		0.717	
	(0.000)	*	(0.000)	*	(0.000)	*
Some college	0.356		0.344		0.335	
	(0.000)	*	(0.000)	*	(0.000)	*
High school graduate	0.178		0.175		0.175	
	(0.000)	*	(0.000)	*	(0.000)	*
Potential experience	0.039		0.039		0.039	
	(0.000)	*	(0.000)	*	(0.000)	*
Potential experience squared	-0.001		-0.001		-0.001	
(*1000)	(0.000)	*	(0.000)	*	(0.000)	*
Conrols for race, age, nativity	YES		YES		YES	
Metro fixed effects?	-		-		YES	
R-squared	0.217		0.221		0.233	

Each column is a separate regression. Robust standard errors in parantheses. Dependent variable is log hourly wage. Each regression contains Census region controls. Number of observations=1.3 million.