

Innovation, Cities, and New Work*

Jeffrey Lin[†]

This version: October 22, 2007

Abstract

Where does adaptation to innovation take place? The supply of educated workers and local industry structure matter for the subsequent location of *new work*—that is, new types of labor-market 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 output 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 technologies. Further, workers in new activities are more skilled, consistent with skill-biased technical change.

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

*I thank Gordon Hanson and Hoyt Bleakley for invaluable guidance. Dana Bradley and Marisa Tegler at the U.S. Census Bureau kindly provided documentation on new occupations. I also thank Eli Berman, Julian Betts, Roger Bohn, Jerry Carlino, Steve Erie, Bob Hunt, Nora Gordon, Roger Gordon, Leslie McGranahan, Leonard Nakamura, Jim Rauch, Albert Saiz, Ian Savage, and audiences at UCSD, the Kiel Institute, the NBER Summer Institute (PRIPE), and the Fed System Regional for helpful comments and discussion. Enrico Moretti generously shared data on land grant colleges. Evan Taylor provided exceptional research assistance. All errors are my own.

[†]Economist, Research Department, Federal Reserve Bank of Philadelphia. Postal address: Ten Independence Mall, Philadelphia, PA, 19106. Electronic mail: jeff dot lin at phil dot frb dot org. Web: <http://jeffr.lin.googlepages.com/>. The views expressed here are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at <http://www.philadelphiafed.org/econ/wps/>.

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, 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 experiences that might lead to innovative combinations. Then again, industrial diversity, or more general human capital from formal education, may 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 aggregate stocks of skills, 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*.³ By new types of work what I mean are the new varieties of activities or techniques that emerge in the labor market in response to new ideas or technologies. These new activities follow innovation, but unlike other measures of innovation output, the concept of new work captures both the creation of new knowledge as well as market acceptance and subsequent changes to the organization of production, labor demand, and labor supply. Through new work, we can observe more closely how workers and firms *adapt* to technological change. I focus on whether the initial supply of college-educated workers and industrial diversity matter for the subsequent location of these new activities. Scale economies and transport costs can explain why some regions are better able to adapt to new technologies.⁴ A central prediction is that new

¹Mokyr (2002) discusses relationships between human capital and technological change at length.

²Agglomeration economies in learning are usually sourced to Marshall (1890); Duranton and Puga (2004) provide an excellent survey. Rosenthal and Strange (2004) survey some recent empirical evidence on knowledge spillovers.

³I borrow the phrase *new work* from Jacobs (1969).

⁴I present a model based on Helpman (1998) in Section 3.

work appears in regions with greater initial supply of educated workers (who 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.

To identify new work in data, 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, necessitate new activities.⁵ 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, provide information about 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.⁶ Then, I estimate a model predicting worker selection into new occupations.⁷ 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 the effects of aggregate skill 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 0.5%. A similar difference in 1990 metro industrial diversity increases selection by about 0.6%. These results highlight the importance of initial aggregate skill for the subsequent 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

⁵A process of destroying old work while creating new work occurs with many innovations. For example, adding machine operators, paper filers, and telegraph operators may disappear as statistical analysts, database managers, and network administrators appear.

⁶Construction of the new work data is covered at length in Section 2 and Appendix A.

⁷See Section 4.

also earn higher wages than observationally similar workers in older activities, consistent with earlier work identifying higher wages in skilled cities. To supplement these main results, I also find that new work outperforms patents in predicting a measure of regional productivity growth.

I further examine alternative hypotheses of unobserved regional characteristics and idiosyncratic shocks related to regional comparative advantage, adaptability, or productivity growth. The main results are robust to the inclusion of additional worker- and region-level controls. Results are consistent using various measures of new work and industrial diversity. Using a coarser methodology based on more aggregate occupation codes, I identify new occupations that appeared between 1960-1970 and 1970-1980.⁸ Estimates from these earlier periods reflect similar patterns.

This paper is most closely related to the empirical literature on human capital and regional growth. Glaeser and Saiz (2003) find that educated cities experience faster wage and population growth, providing indirect evidence of the ability of these cities to adapt to innovation. Glaeser et al. (1992) examine the relationship between diversity in industry-specific experiences and growth across regions.⁹ There is evidence on other connections between agglomeration, human capital, innovation, and adaptation. For example, Rauch (1993) and Moretti (2004) find higher wages in educated cities, Bacolod, Blum, and Strange (2007) decompose the urban wage premium by occupation-level skills, Beaudry, Doms and Lewis (2006) relate PC adoption to differences in the supply of skilled labor across regions, and Carlino, Chatterjee, and Hunt (2007) find that patent intensity is positively related to employment density.¹⁰ In the literature on skill-biased technical change, Berman, Bound, and Griliches (1994) and Autor, Katz, and Krueger (1998) suggest that college-educated workers are best suited to adapt to new technologies.

Methodologically, the approach here is similar to that in several recent papers. Xiang (2005) identifies new products in revisions to industry codes. Autor, Levy, and Murnane (2003) and Bacolod and Blum (2005)

⁸New work in earlier census years is covered in Appendix A.2. I use documentation on changes in 3-digit occupation codes to obtain information about new activities in these decades. Because the 3-digit occupation codes in the OCS did not change between 1980-1990, it is much more difficult to identify new work. Occupation titles, available electronically in 1990 and 2000, are only available in hard copy for 1980. I thus did not attempt to identify new occupations for this period.

⁹See also Henderson, Kuncoro, and Turner (1995) and Henderson (2003).

¹⁰See also Jaffe, Trajtenberg, and Henderson (1993), who find that patents are geographically localized, and Feldman and Audretsch (1999), who trace product inventions to industrially diverse cities. Another related paper is Bresnahan and Trajtenberg's (1992) work on general-purpose technologies. Certain pieces of knowledge are broadly accessible and adaptable to a wide variety of purposes; by picking up changes in labor market structure across sectors, new work captures some of these properties. Also, Bleakley and Lin (2007) find that thick markets provide better labor market matching; thicker markets for newer vintages of skill (via more skilled workers) might also contribute to the quicker adoption of new work.

examine the skill content of detailed occupational codes to study technical change and wage inequality. The main contribution of this paper is to use changes in occupation classifications to empirically identify new production activities in the labor market. New work can be thought of as more specifically characterizing how technical change affects the structure of production and labor. I am then able to provide more direct evidence on how skills, at both the worker and aggregate levels, help firms and workers adapt to new ideas.

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, which contain information on workers’ activities and locations. Later in this section, I describe characteristics of the population in new work and compare new work to other measures of innovation-related output. Further details are relegated to Appendix A.

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.¹¹ 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 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 Index 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.

¹¹Technical papers from the U.S. Census Bureau (Scopp, 2003) and the Bureau of Labor Statistics (Meyer and Osborne, 2005) provide good overviews of this process.

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 with 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 (Scopp 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” (Scopp 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*.¹² 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*. 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 840 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 about 840 titles constitutes new work in 2000.¹³ Note that the main results are not sensitive to alternative (and orthogonal) algorithms for identifying new work. These algorithms are described in further detail in Appendix A.1, and Section 5.2 discusses results using these different measures.

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 lists DOCs with the highest new title shares.¹⁴ For example, DOC 111, *network systems and data communication*

¹²These are available upon request.

¹³Available at my web site (<http://jeffr.lin.googlepages.com/>).

¹⁴Appendix Table A reports new DOCs in the period 1960-1980.

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.¹⁵

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 5%. The skewness of new titles across DOCs mitigates some concern over an uneven distribution of employment within DOCs, across titles. Instead, results are driven by a small fraction of DOCs with new title shares close to 1.¹⁶

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 contain detailed personal characteristics and metropolitan area of residence for approximately 1% of the 2000 U.S. population. I use 1999-defined consolidated metropolitan areas.¹⁷ The sample (1.5 million observations) includes all workers age 20 to 65 in identified metropolitan

¹⁵A few DOCs listed in Table 1 do not reflect new activities in obvious ways. Some of this is due to measurement error; see Section 5.2 for results using alternative measures of new work. However, other DOCs may reflect implementation of new technologies in non-obvious ways. For example, DOC 452 contains estheticians and other types of skin care specialists classified as new, and DOC 200 contains AIDS and HIV counselors.

¹⁶I expand on this point and discuss alternative imputation strategies in Section 5.2.

¹⁷For instance, Stamford and Newark are grouped with New York, Orange County and Riverside are with Los Angeles, and

areas 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. Appendix Table C 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.)¹⁸ Table 2 describes the population in new work. In metropolitan areas, employment share in new work is 5.1%; 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 age distribution of new workers suggests a possible selection bias toward more recent vintages of skills. These patterns are also repeated in regression estimates presented in Section 5.1.

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, Fresno, Bakersfield, and Modesto had the lowest shares in new work.

New occupations employ more educated workers, consistent with skill-biased technical change. In the lower part of Table 2, 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 7.9% 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 electrolysis, both in the category of personal appearance workers. While on average new workers are more educated, the presence of new work at all levels of the skill

Dallas and Fort Worth are grouped together.

¹⁸A maintained assumption in this section is that the distribution of employment across new occupation titles (with DOCs) is equal across different groupings. Part of the basis for this assumption is the skewed distribution of new titles across DOCs as discussed in the previous section.

distribution highlights the interpretation of new work as reflecting *changes in the organization of production*, rather than exclusive and specialized activities requiring the most cutting-edge skills.

Table 3, Panel A, 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. Using the 2000 PUMS, I perform a regression of log hourly wage on education, experience, and the new work variable. Column (4) in Table 3, Panel B, contains the most complete set of regressors, with fixed effects for workers' metropolitan area of residence and 3-digit industry.¹⁹ Using this specification, I find that employment in new work predicts a wage premium of upwards of 21% (logarithmic) over observationally similar workers who are not in new work.

One interpretation of the wage premium is that new workers are more productive and, at the same time, especially well-suited for these new activities. A second interpretation is that entry into one of these new occupations is inherently risky—who knows if this activity is an effective way to use a new technology?—and the wage premium reflects this risk. Future research may explore different explanations for the wage premium, but for now, I only emphasize that workers performing new activities appear different from other workers, performing existing work, across many dimensions.

2.5 New work compared to other measures of innovation and adaptation

New work has several unique strengths as a measure of adaptation to new knowledge. Consider a comparison between new work and patents, a common (and complementary) measure of innovation. Both are outputs related to the (invisible) creation of new knowledge. A patent must also pass through both patentability rules and incentives affecting the patenting decision. A new occupation instead reflects both market acceptance

¹⁹There are a total of 20,241 fixed effects in this regression.

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 strategy. If patents track the birth of an idea, then new work tracks the effects and the adaptation of an idea.

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.²⁰ 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 strategy. 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. As Bresnahan and Trajtenberg (1992) note, new knowledge may have unexpected applications across sectors, encompassing new products and new services. Importantly, new work can capture innovations in processes, the organization of production. Adam Smith's pin factory example is commonly used to illustrate the division of labor. I argue that this creation of new specialized activities can also be thought of as reflecting responses in the labor market to the underlying new ideas related to re-engineering the production process.

By matching new work data to microdata, I can identify workers engaged in new work; these data are useful for studying the effects of innovation on labor markets. New work also has potential benefits for historical analyses: this strategy may be extended backward 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, but they differ in their informational content. To show this, I use accumulated patent counts²¹ 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.

²⁰In addition, the patent office may be limited in its ability to identify truly new inventions. See Wu (2006).

²¹Utility patents, or patents for invention.

Figure 3 graphs workforce share in new work in 2000 against accumulated 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.

Finally, I compare the geographic distributions of new work and productivity. Lacking any established sources of estimates for metropolitan area total factor productivity, I instead impute it using estimates from manufacturing data.²² This procedure yields an imputed measure of (log) TFP growth by metropolitan area. This measure is positively correlated to new work share, as seen in Figure 4. I then regress this measure against the share of each metro area's manufacturing workforce in new work. Estimates from this regression are in Table 3, Panel C. New work share in manufacturing is a strong predictor of imputed metro area productivity growth, outperforms the measure of log patents per capita, and is robust to the inclusion of other metropolitan controls.

3 Theory

I showed in the last section that certain regions attract more new work. In this section, based on Helpman (1998), I formalize the effect of innovation on the distribution of production activities across regions.²³ The starting point is to imagine innovation as a shock to the economy. How do workers and firms across

²²I take estimated TFP in 1990 and 1996, the latest year available in the data, for 2-digit SIC industries from the NBER-CES database (Bartelsman, Becker and Gray, 2000). Then, by metropolitan area, I take a weighted average of industry TFP, where the weights are 1990 industry employment shares within each metro area.

²³This 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.

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.²⁴

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 .²⁵ 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:

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

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? Bills 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

²⁴Redding and Sturm (2006) use a similar strategy to simulate the effects of German division on the size of cities. Hanson (2005) uses the Helpman model to examine the effect of market access on agglomeration.

²⁵Helpman calls this good housing, though it can be any non-traded good with inelastic supply.

f in terms of skilled labor l . Let β be the unit cost in skilled labor, then

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

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 2.²⁶

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 $v \equiv n_1/N$. In Appendix B.1, 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 .²⁷ 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 . I solve for equilibrium values of v and w 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 $u = \infty$).

Following Helpman, the important parameters for determining the stability and uniqueness of the initial equilibria are μ ($\equiv N/(N + \sigma)$), σ , and t . There are two configurations of equilibria. In the first case, a unique, stable equilibrium exists for $\sigma(1 - \mu) = \sigma\delta/(N + \delta) > 1$.²⁸ These conditions imply a high

²⁶Unskilled 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).

²⁷The relationship between v and w does not depend on either f nor β , which only scale the number of varieties and the level of production output.

²⁸Or, in a trivial case, when there are no transport costs ($t = 1$).

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 5, Case A, for the case of $h_1 = 2$ and $h_2 = 1$.²⁹ 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 $1 < \tilde{t} < t < \infty$. 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, Case B, 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 subsequent 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).³⁰ Given new N (and hence μ), I solve again for equilibrium production share v , wages w , and relative utility u . In Figure 5, intersections between the dotted lines (“after”) and the solid line ($u = 1$) indicate the new equilibria.

In the unique equilibrium case (Case A), innovation deepens concentration. Region 1, which initially contained more skilled labor and production diversity, attracts more new production activities. In the mul-

²⁹Since 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.

³⁰Details on this simulation and parameter values are in the appendix. Note that this is a non-unique way to formulate innovation: population growth in skilled labor or an expansion in production activities will have equivalent effects.

multiple equilibria case (Case B), 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 outward.³¹ 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 a 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 a greater supply of educated labor and industrial diversity. Define the outcome of interest, ν_i ³², as the new title share (in all titles) within each worker i 's detailed occupation code. For

³¹Intuitively, 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. This intuition echoes the concept of “venturesome consumption” explained by Bhidé (2006). He argues that the use or consumption of innovation-related outputs matters for the development of new ideas.

³²i.e., “nu”

example, workers identified as *network systems and data communication analysts* (DOC 111) have $\nu_i = 96.7\%$, 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 2000.³³ For most workers, ν_i is zero; all titles in their DOC can be matched to 1990 titles, and they are unlikely to have selected into a new activity.³⁴ I estimate a linear model to predict ν_{ij} for each worker i living in city j :³⁵

$$\nu_{ij} = \alpha + X_i\beta + Z_j\gamma + \varepsilon_{ij} \quad (3)$$

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 ν_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 ν_{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 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.³⁶ In 1990, the college share variable ranges from 11.5% (McAllen, Texas) to 31.6% (Raleigh-Durham).

³³Because of the aggregation from new 5-digit titles to 3-digit DOCs, it is unobserved whether each worker is actually in a new occupation. See Section 2.3.

³⁴See Figure 1. Because of the construction of ν_{ij} , this equation could also be estimated at the occupation level, with DOCs instead of workers being the unit of observations. However, I would then be unable to separately estimate the effects of worker characteristics on selection into new work.

³⁵The 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 ν_i is greater than or less than a particular threshold, yields similar results. See Section 5.2. Probit estimates also yield similar results; see Appendix Table D.

³⁶See Acemoglu and Angrist (2000) and Moretti (2004).

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 1990.³⁷ 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 1990.³⁸ 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, industry, and age. With indicator variables for high school graduation, some postsecondary education, and college 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 $\hat{\gamma}$, the estimated coefficient on Z_j , 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 3.3 suggests that the estimated effects of these aggregate skills should be positive. The source of identification comes from variation across 1990 metropolitan areas, which I take as historically determined. In light of the model in Section 3, the initial 1990 equilibrium represents a distribution of skilled labor and production diversity reflecting long-running historical patterns of economic activity.³⁹ These initial differences then serve as initial conditions for firms and households to make location decisions, ultimately determining how well a region attracts the next round of innovation and new work. A model-specific interpretation is that within regions containing better initial stocks of aggregate skills, households consume and firms produce more new products that use new activities as inputs.

Proper inference requires that ε_{ij} be conditionally uncorrelated to the elements of Z_j . Section 5 addresses a number of possible violations of this condition. In particular, unobserved regional characteristics or shocks to regions that are related to regional adaptiveness may bias estimates of γ . For example, city size

³⁷Results are not sensitive to this measure of industrial diversity. In Table 6, I present estimates using alternative measures of industrial diversity.

³⁸The 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. Alternatively, the Herfindahl index used in Table 6 uses aggregate census data from State of the Cities, avoiding this issue altogether.

³⁹One 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.

is positively correlated to industrial diversity and also the appearance of new work; this will lead to omitted variables bias. In all regressions, I include (log) 1990 metropolitan area population and land area. Similarly, in most specifications, I control for region-specific labor demand shocks and labor supply responses. In Section 5.3, I discuss these measures and a host of other region-specific controls that may confound identification. Other checks in Section 5 include estimates using a variety of new work and industrial diversity measures, decompositions by industry, geography, and educational attainment, an analysis of possible sorting of workers across cities based on unobserved skills, and new work using historical (pre-1990) census data.

5 Results

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

I find that 1990 college graduate share and industrial diversity positively predict worker selection into new work. The main results are in Table 4; each column is a separate regression. I include coefficient estimates for worker educational attainment, metro education and industrial diversity, and other worker characteristics. Suppressed coefficients include metro log population and land area.⁴⁰ Estimates are reported in percentage point units; an estimate of 4.7 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 4.7 percentage points. Means and standard deviations of the independent variables are also reported in brackets.

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 0.4-0.5%.⁴¹ This is calculated by multiplying the estimated coefficient (10.3 or 12.5) 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

⁴⁰In most specifications, the coefficient estimates on log population and log land area are small and statistically insignificant. Once I control for measures of aggregate skills, overall urban scale has little explanatory power for the location of new work.

⁴¹I 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.

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 3) and when controlling for other characteristics in a regression (Section 2.4). To the extent that new work is more productive (or that workers in new activities enjoy rents for selecting into new work quickly), this result suggests that the appearance of new occupations may be an important channel for productivity spillovers observed in previous work.

In columns (3) and (4), 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 comparable to 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 5.3.

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).

I perform a similar graphical exercise 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 the two figures supports a

more central role for metro education. Formal education, characterizing more general human capital, may be a more important form of skill in generating regional advantage.

In Table 4, 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 4.7% more likely to select into new work (that is, select into an occupation with a 4.7% higher new work share) than a high school dropout. Selection appears monotonic in educational attainment. At sample means, the coefficient estimates imply that college graduates are more than four times as likely as high school dropouts to select into new work.

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 (0.6% 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. A second interpretation is that changes in the structure of production may be biased toward more recent vintages of human capital, which would favor selection by younger workers. In related work, Bleakley and Lin (2007) provide evidence of more occupational churning among young workers in large metropolitan areas. Other worker-level characteristics that predict selection into new work resemble those that predict labor market outcomes in other contexts. The negative coefficient estimate on self-employment must reflect the number of workers who operate businesses in *existing* categories of work.

In sum, initial metro college share and industrial diversity are important predictors of future selection into new occupations. 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 initial regional stocks of skill help workers 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 Alternative measures of new work and industrial diversity

These results are robust to different algorithms for identifying new work and different measures of industrial diversity. Table 4'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 explained by a greater number of observed sectors, rather than the distribution of employment across sectors. This is consistent with the model in Section 3; since firm-industries are symmetric, relative size plays no role. In this section I consider alternative measures of industrial diversity.

The first alternative is a modified Herfindahl-Hirschman index of 1990 industry employment shares within a city, calculated as $Herf_j = 1/|\sum_k (s_{kj} - s_k)|$, where s_{kj} is employment share for industry k in city j . This index is similar to the standard sum-of-squared-shares Herfindahl-Hirschman index, except that it is inverted and corrects for differences in industry employment shares at the national level (s_k).⁴² A second alternative industrial diversity measure is the 1990 share of the top m industries within a metropolitan area, following Glaeser et al. (1992). I choose $m = 20$, but other values of m yield essentially similar results. This measure is also inverted so that it increases with industrial diversity.

In addition, I use alternative algorithms for identifying new occupation titles. The *baseline* measure, as described in Section 2.2, relies on a manual review of occupation titles across the 1990 and 2000 OCS. A second (and orthogonal) algorithm (*census rules*) uses detailed internal census documentation to decide whether a 2000 OCS title is new work. Details on this algorithm are in Appendix A.1. Four other algorithms are based on variations of the census rules or baseline definitions of new work. *Baseline & census rules* includes occupation titles in the intersection of the previous two sets. The *wide* definition of new work is based on a string matching algorithm between 1990 and 2000 occupation titles, as described in Section 2.2. The *baseline binary 1* definition of new work equals 1 if the baseline new title share is above the 95th percentile across DOCs, zero otherwise. (This strategy yields a dependent variable mean approximately equal to the baseline measure.) The *baseline binary 2* definition of new work equals 1 if the baseline new title share is above 90%, zero otherwise. This measure exploits the skewness of the distribution of new titles across DOCs to try to mitigate concerns that variation in within-DOC employment shares are driving the main results. Here, I classify new work only if the new title share is extremely high.

⁴²See Duranton and Puga (2000), p. 535, for discussion.

Table 6 presents coefficient estimates on the 1990 college share (Panel A) and 1990 industrial diversity (Panel B) for 6 new work algorithms and 3 industrial diversity measures. Each cell is from a separate regression, though corresponding cells in each panel are from the same regression. *Herf.* and *Top 20 share* are standardized to have mean 0 and standard deviation 1, so the estimates in the last two columns of Panel B are comparable. I also report means and standard deviations of the new work variables in brackets.

In Panel A, the estimated effects of 1990 college share are robust to different measures of new work. In Panel B, a one-standard-deviation change in the 1990 Herfindahl index increases selection into new work by 0.02% to 0.16%, and about 0.1% for the baseline new work measure. A one-standard-deviation change in the 1990 top 20 industry share increases selection into new work by 0.09% - 0.22%, and about 0.2% for the baseline new work measure. These estimates, which explicitly take into account the relative sizes of city-industries, are more modest than the industry count measure, which does not. However, the estimated effects of industrial diversity are still positive and precisely estimated.⁴³

5.3 Region-specific characteristics and shocks

The main results are also robust to the inclusion of other city-specific characteristics.⁴⁴ In most specifications, I include the Blanchard and Katz et al. (1992) labor demand index as an additional regressor. This index is a weighted average of industry employment growth, where the weights are metro-specific industry employment shares.

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

The Blanchard-Katz index, $\hat{\eta}_j$, is the predicted growth in employment for metropolitan area j , based on 1990 employment shares ξ of industries k in metro j , and the change in log employment η_k in industry k (nationally) 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

⁴³There are two coefficient estimates that are positive but not significant at the 95% level; however, these occur in the case of the wide definition, which is an imprecise measure, and the baseline binary 2 definition, which takes a value of 1 for only 0.2% of workers.

⁴⁴As noted in Section 4, log population and land area, and metro some college and high school graduate share are included in all regressions. The coefficient estimates on these variables are generally not significantly different from zero when controlling for 1990 metro college share and industrial diversity.

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 $\hat{\eta}_j$ indicate that in 1990, the city was relatively specialized in industries that grew quickly during the 1990s. This variation is almost certainly correlated with metro education and industrial diversity, as well as the subsequent appearance of new work.

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 the effects of aggregate education and industrial diversity. Skilled workers who remain in Rust Belt cities as they decline may be more likely to select into new occupations as other skilled workers leave. To control for this, I include (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 the aforementioned patent data described in Section 2.5, I match patent counts of 2-digit industries between 1990-1999 to 1990 metro 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,

$$\hat{\pi}_j = \sum_k \xi_{jk} \pi_k \tag{5}$$

As before, ξ_k 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 clean out some additional endogenous variation across metropolitan areas.

The first four columns of Table 6 contain estimates from regressions with and without these variables. (Column (1) reproduces estimates from Table 4, column (4).) Both the Blanchard-Katz index and the predicted patenting index have positive estimated coefficients, consistent with fast-growing areas attracting new work. However, the standard errors are larger in magnitude than the point estimates, and their inclusion does not appreciably change the estimated coefficients on the main variables of interest.

There still may be a number of unobserved regional characteristics or shocks to regions related to the appearance of new work across regions that are correlated with aggregate education and industrial diversity.

One approach I take to control for remaining omitted variables is to include actual patent activity as an additional regressor. I use the (actual) patent data to control for additional unobserved variables related to knowledge creation. In Table 6, column (5), 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 (5) also illustrates orthogonal information in 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 the main variables of interest. Such an effect might lead to an overestimate of returns to metro education.⁴⁵ 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.⁴⁶ In column (6), 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 column (7) of Table 6, I include different measures of 1990 metropolitan industry-occupation structures. Additional controls include major (1-digit) industry and occupation shares, occupational diversity (as measured both by a Herfindahl index and occupation count), and lagged own-industry share. Column (8) further includes the 1990 metro share with a post-graduate degree and controls for climate: (log) heating degree-days, cooling degree-days, and precipitation, mean January and July temperatures, and average January lows and July highs.⁴⁷ Coefficient estimates for the main variables of interest, 1990 college share and

⁴⁵The presence of universities may also relate to an alternative hypothesis that some regions have a comparative advantage in using these new production techniques.

⁴⁶Moretti (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.

⁴⁷See Appendix A.3 for details.

industrial diversity, are similar to the baseline results even when all other metropolitan-level controls are included.

An alternative hypothesis consistent with the results presented so far is that some cities may have experienced idiosyncratic shocks over the period 1990-2000 that affected their ability to attract new work. For example, 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 existing 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 main results are due to migration, not increasing returns to the concentration of human capital.

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.

I classify manufacturing and information as traded industries and perform the baseline regression from Table 4, column (4) on only these industries.⁴⁸ These estimates, along with estimates on other industry sub-samples, are presented in Table 6, Panel A. I report the new work mean in each industry sub-sample at the top of each column. The estimated effects of the 1990 college share and industrial diversity appear to be stronger in traded industries relative to non-traded industries, and for information industries relative to services industries in general (which include both traded and non-traded industries). However, much of this difference in magnitude is driven by differences in the average new work share across industries.

I also decompose the sample by geography, to demonstrate that the main results are not region-specific. The first two columns for Panel B contain estimates for two sub-samples, the eastern and western U.S. In both cases, the estimated effects for college share and industrial diversity are similar. The point estimate for industrial diversity is slightly larger in the West, but it is less precisely estimated. This is possibly due to smaller sample size in the West (37 metropolitan areas).

⁴⁸Some other services are also traded. Alternative classifications, including wholesale trade, transportation, finance and insurance, professional services, management, higher education, and arts and entertainment, yield similar results.

There may be variation across space in what entails a “local” labor market. 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 the third column of Panel B are similar to the baseline results. In the last column, I restrict the sample to metropolitan areas that 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.

5.4 Sorting and new work

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 worker education by performing the baseline regression (Table 4, column (4)), on four separate samples: college graduates, workers with some college, high school graduates, and high school dropouts. Table 6, Panel C, 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.

Workers may sort into skilled metropolitan areas based on some unobserved component of ability. In a sense, this is tangential to the main theme of the paper, which is about which regions *attract* workers performing the newest activities, whether or not new workers are home-grown or move to regions based on some unobserved abilities. It may still be helpful, though, to identify whether new workers who move into new metropolitan areas are systematically different from others. In Table 6, Panel A, I divide the sample into

two sub-samples: those workers whose metropolitan area of residence changed between 1995 and 2000, and other workers who stayed in the same metro over the period.⁴⁹ Estimates on the main variables of interest do not appear appreciably different.

A second technique to correct for sorting is based on Evans, Oates and Schwab (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. In a similar spirit, I use 1970 and 1990 state-level characteristics to instrument for 1990 metropolitan characteristics. These state-level characteristics are matched to each worker's state of birth. These estimates are presented in Panel B of Table 6. The estimates using 1970 state characteristics may also be informative in mitigating concerns about (1990) unobserved metro characteristics that contribute directly to regional adaptiveness. If these unobserved characteristics are related to adaptiveness only contemporaneously, then the IV estimates may correct for the omitted-variables bias. Both sets of IV estimates are in line with the baseline results.

5.5 New work and historical data

In Section 4, I suggested that the initial distribution of educated labor and industrial diversity across regions reflects long-running historical processes. If this is in fact the case, then earlier measures of regional skill and industrial diversity should also predict worker selection into new work. In first column of Table 6, 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 future 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. These data are matched to 1970 and 1980 census microdata; as noted in Appendix A.2, the methodology relies on matching 3-digit DOCs across census years 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

⁴⁹The 2000 IPUMS contains no additional information on previous metropolitan area of residence, only an indicator of changed residence.

both the 1970 and the 1980 census; I use occupation outcomes in 1970 and 1980 matched to metropolitan data in 1950 and 1970. These results are displayed in the middle columns of Table 6. 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 from 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. Demonstrating that some of the same correlations exist in earlier periods is important, since it is likely that the 1990s were a unique period in terms of innovation and adaptation.

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 that may be related to regional innovativeness. Identification of the effect of lagged college share and industrial diversity comes from changes 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 the last column of Table 6, controlling for metropolitan fixed effects does not change the flavor of the main results.⁵⁰ Controls are the same as in Table 4. 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 aggregate education and industrial diversity.

⁵⁰Note that the large magnitude of these estimates relative to previous tables is due to differences in identification strategy across census years, and the resulting differences in variance in the dependent variable. See Appendix A.2 for details.

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 propose a measure of adaptiveness that more specifically characterizes how various forms of skill 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 precedes selection into new work may give us more insight into the adaptation process. In particular, it may be possible to use these 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 other data sources, in order to further understand high-frequency properties of innovation. Finally, the wage premia experienced by workers who select into new work may be important for understanding trends in wage inequality.

References

- Acemoglu, D. and J. D. Angrist (2000). “How Large Are the Social Returns to Education? Evidence from Compulsory Attendance Laws.” in B. S. Bernanke and K. Rogoff, eds., *NBER Macroeconomics Annual*.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration.” *The Quarterly Journal of Economics*. **118** (4), 1279–1333.
- , L. F. Katz, and A. B. Krueger (1998). “Computing Inequality: Have Computers Changed the Labor Market?” *The Quarterly Journal of Economics*. **113** (4), 1169–1213.
- Bacolod, M. and B. S. Blum (2005). “Two Sides of the Same Coin: U.S. ‘Residual’ Inequality and the Gender Gap.” Manuscript.
- , ———, and W. C. Strange (2007). “Skills in the City.” Manuscript.
- Bartelsman, E. J., R. A. Becker, and W. B. Gray (2000). *NBER-CES Manufacturing Industry Database*.
- Beaudry, P., M. Doms, and E. Lewis (2006). “Endogenous Skill Bias in Technology Adoption: City-Level Evidence from the IT Revolution.” NBER working paper 12521.

- 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.” *The Quarterly Journal of Economics*. **109** (2), 367–397.
- Bhidé, A. (2006). “Venturesome Consumption, Innovation and Globalization.” Manuscript.
- 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 (2007). “Thick-Market Effects and Churning in the Labor Market: Evidence from U.S. Cities.” Federal Reserve Bank of Philadelphia Working Paper 07-23.
- 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. and H. J. Holzer (2000). “Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s.” *Journal of Labor Economics*. **18** (1), 20–54.
- , J. Groen, G. Kézdi, and S. Turner (2004). “Trade in University Training: Cross-State Variation in the Production and Use of College-Educated Labor.” *Journal of Econometrics*. **121** (1-2), 143–173.
- Bresnahan, T. F. and M. Trajtenberg (1992). “General Purpose Technologies: ‘Engines of Growth?’.” *Journal of Econometrics*. **65** (1), 83–108.
- Carlino, G. A., S. Chatterjee, and R. M. Hunt (2007). “Urban Density and the Rate of Invention.” *Journal of Urban Economics*. **61**, 389–419.
- del Regato, J. A. (1995). “One Hundred Years of Radiation Oncology.” in J. Tobias and P. R. M. Thomas, eds., *Current Radiation Oncology*, Vol. 2, Oxford University Press.
- Duranton, G. and D. Puga (2000). “Diversity and Specialization in Cities: Why, Where and When Does it Matter?” *Urban Studies*. **37** (3), 533–555.
- and ——— (2001). “Nursery Cities: Urban Diversity, Process Innovation, and the Life-Cycle of Products.” *American Economic Review*. **91** (5), 1454–1477.
- and ——— (2004). “Micro-Foundations of Urban Agglomeration Economies.” in J. V. Henderson and J.-F. Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4, North-Holland, chapter 48, pp. 2063–2117.
- 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. E. Oates, and R. M. Schwab (1992). “Measuring Peer Group Effects: A Study of Teenage Behavior.” *The 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.
- and A. Saiz (2003). "The Rise of the Skilled City." NBER working paper 10191.
- , H. D. Kallal, J. A. Scheinkman, and A. Shleifer (1992). "Growth in Cities." *The Journal of Political Economy*. **100** (6), 1126–1152.
- Hall, B. H. (2005). "Exploring the Patent Explosion." *Journal of Technology Transfer*. **30**, 35–48.
- , A. B. Jaffe, and M. Trajtenberg (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 D. Pines, E. Sadka, and Y. Zilcha, eds., *Topics in Public Economics*, Cambridge University Press, pp. 33–54.
- Henderson, J. V. (2003). "Marshall's Scale Economies." *Journal of Urban Economics*. **53**, 1–28.
- , A. Kuncoro, and M. Turner (1995). "Industrial Development in Cities." *The Journal of Political Economy*. **103** (5), 1067–1090.
- Jacobs, J. (1969). *The Economy of Cities*, Random House.
- Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *The Quarterly Journal of Economics*. **108** (3), 577–598.
- Krugman, P. (1991a). "History Versus Expectations." *The Quarterly Journal of Economics*. **106** (2), 651–667.
- (1991b). "Increasing Returns and Economic Geography." *The Journal of Political Economy*. **99** (3), 483–499.
- Lucas, R. E. (1988). "On the Mechanics of Economic Development." *Journal of Monetary Economics*. **22**, 3–42.
- Marshall, A. (1890). *Principles of Economics*, Macmillan.
- Meyer, P. B. and A. M. Osborne (2005). "Proposed Category System for 1960-2000 Census Occupations." U.S. Bureau of Labor Statistics working paper 383.
- Mokyr, J. (2002). *The Gifts of Athena: Historical Origins of the Knowledge Economy*, Princeton University Press.
- Moretti, E. (2004). "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data." *Journal of Econometrics*. **121** (1-2), 175–212.
- Rauch, J. E. (1993). "Productivity Gains from 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.” Manuscript.
- Romer, P. M. (1990). “Endogenous Technological Change.” *The Journal of Political Economy*. **98** (5), S71–S102.
- Rosenthal, S. S. and W. C. Strange (2004). “Evidence on the Nature and Sources of Agglomeration Economies.” in J. V. Henderson and J.-F. Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4, North-Holland, chapter 49, pp. 2119–2171.
- Ruggles, S., 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]*, Minnesota Population Center.
- Saxenian, A. (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*, Harvard University Press.
- Schultz, T. W. (1975). “The Value of the Ability to Deal with Disequilibria.” *Journal of Economic Literature*. **13** (3), 827–846.
- Scopp, T. S. (2003). “The Relationship Between the 1990 Census and Census 2000 Industry and Occupation Classification Systems.” U.S. Census Bureau Technical Paper 65.
- U.S. Census Bureau (1950, 1960, 1970, 1980, 1990, 2000a). *Alphabetical Index of Occupations and Industries*, U.S. Department of Commerce.
- (1950, 1960, 1970, 1980, 1990, 2000b). *Classified Index of Occupations and Industries*, U.S. Department of Commerce.
- (1968). “Changes Between the 1950 and 1960 Occupation and Industry Classifications.” Technical Paper 18.
- (1972). “1970 Occupation and Industry Classification Systems in Terms of Their 1960 Occupation and Industry Elements.” Technical Paper 26.
- (1989). “The Relationship Between the 1970 and 1980 Industry and Occupation Classification Systems.” Technical Paper 59.
- (1994). *County and City Data Book*, U.S. Department of Commerce.
- U.S. Department of Housing and Urban Development (2005). *State of the Cities Data Systems*.
- U.S. Patent and Trademark Office (2000). *United States Patent Grants by State, County, and Metropolitan Area*.
- Venables, A. J. (1996). “Equilibrium Locations of Vertically Linked Industries.” *International Economic Review*. **37** (2), 341–359.
- Wu, T. (2006). “Weapons of Business Destruction: How a Tiny Little ‘Patent Troll’ Got BlackBerry in a Headlock.” *Slate*.
- Xiang, C. (2005). “New Goods and the Relative Demand for Skilled Labor.” *Review of Economics and Statistics*. **87** (2), 285–298.

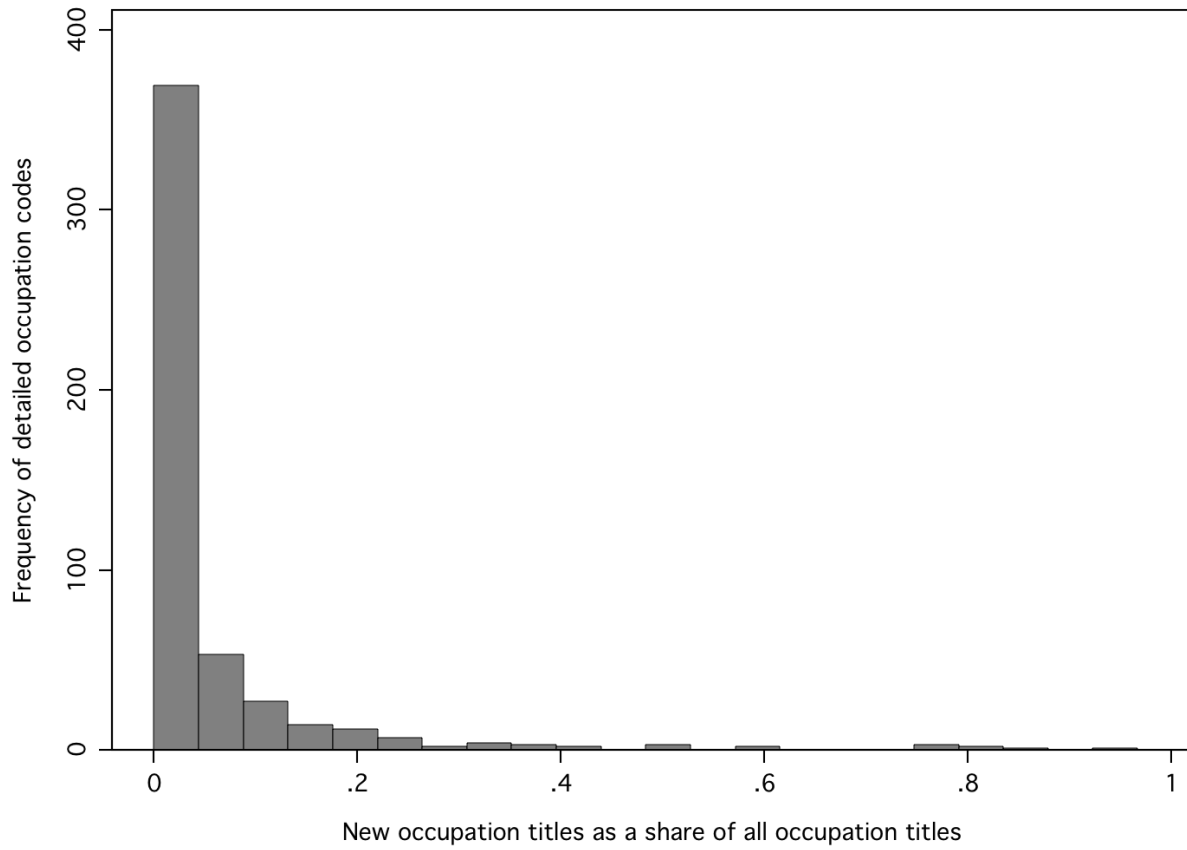


Figure 1: Frequency of DOCs by share of new titles in all titles.

Notes: This histogram describes the "newness" of 3-digit detailed occupation codes (DOCs). The horizontal axis is calculated by the share of new 5-digit titles in all titles within each DOC. About three-quarters of 2000 DOCs contain less than 5% new titles.

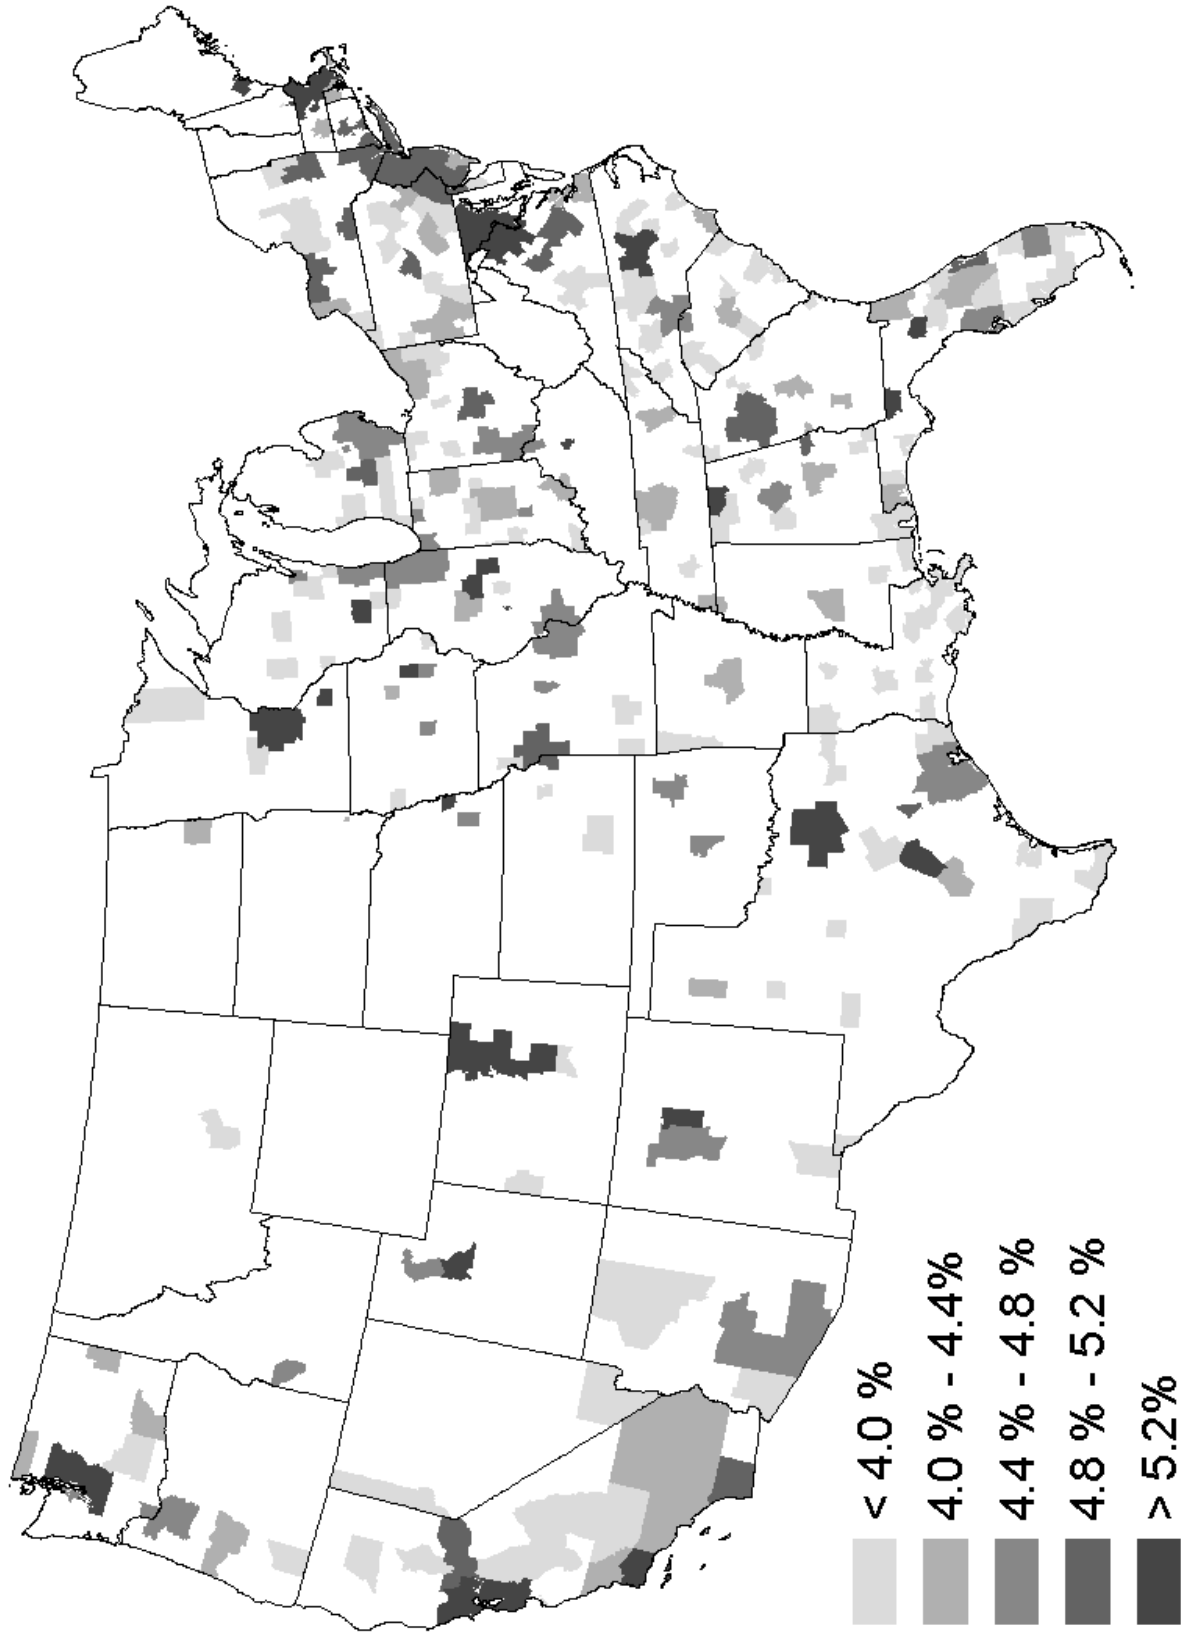


Figure 2: New work as share of local employment, U.S. consolidated metropolitan areas, 2000

Notes: Author's calculations using 2000 IPUMS, workers 20-65, with identified occupations. Areas are 1999-defined U.S. consolidated metropolitan areas. Areas not included in sample in white.

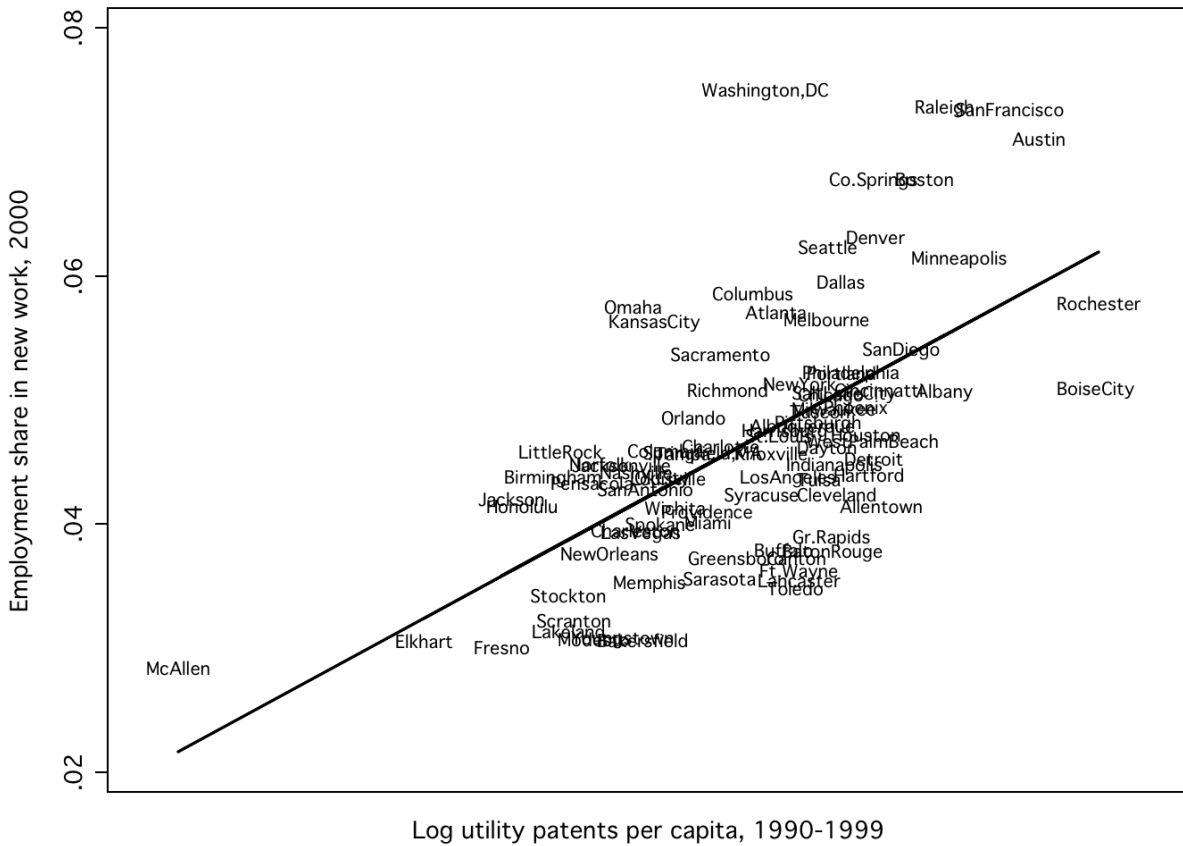


Figure 3: New work employment share in 2000 and utility patents per capita, 1990-99

Notes: Log utility patents (patents for invention) granted over the period 1990-1999, from the U.S. Patent and Trademark Office (2000). Location of each patent is determined by first-listed inventor. Areas are 1999-defined U.S. consolidated metropolitan areas.

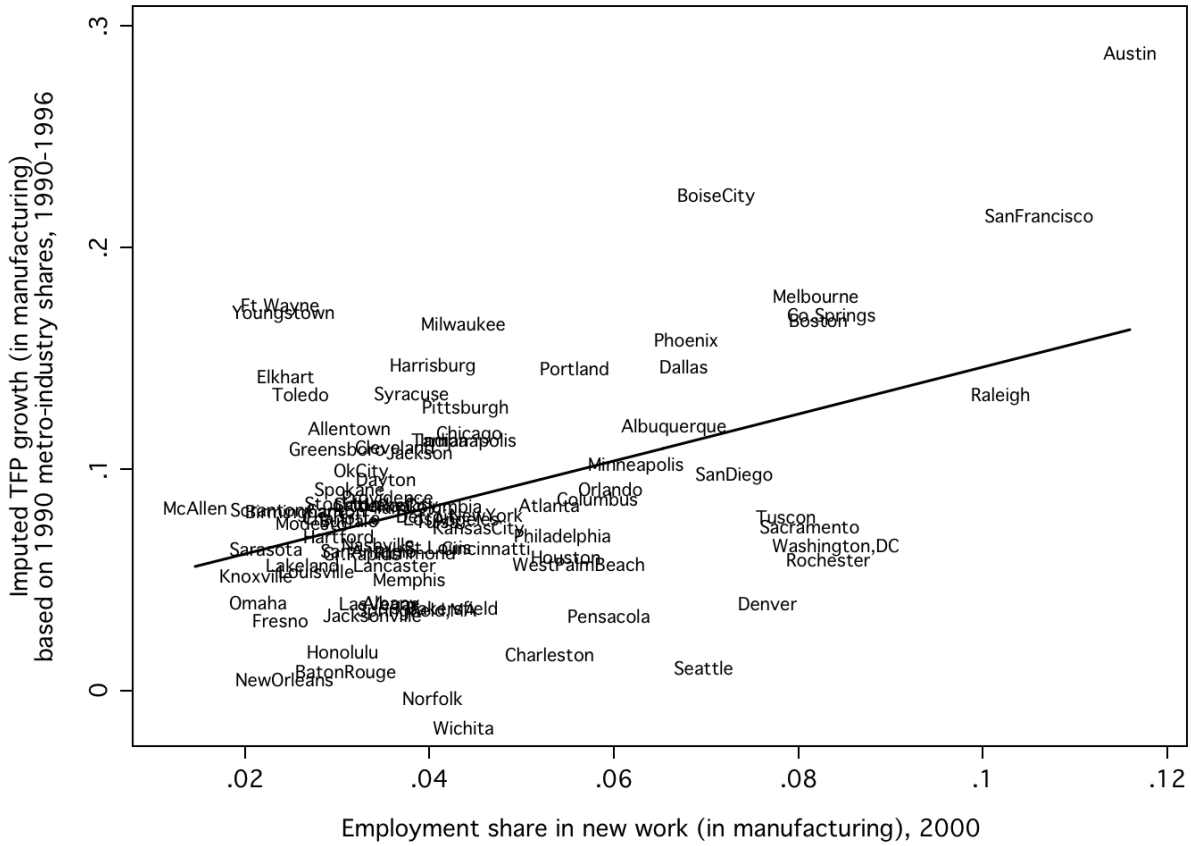



Figure 4: Imputed TFP growth in manufacturing, 1990-96, and new work employment share in 2000

Notes: Imputed TFP growth in manufacturing industries is calculated as a weighted average of estimated TFP growth, 1990-1996 (the latest year available), across 2-digit SIC industries, from the NBER-CES database. The weights are 1990 within-metro industry employment shares. New work share is in manufacturing industries only. Areas are 1999-defined U.S. consolidated metropolitan areas.

Case A. Unique stable equilibrium



Case B. Multiple stable equilibria




Figure 5: Innovation shocks and the location of new work

Notes: These graphs display the utility level of households in region 1 relative to region 2, u , and the share of production activities in region 1, v , which ranges from 0 to 1. The dashed and dotted lines, respectively, map utility–activity shares before and after an innovation shock. In Case A, with a unique stable equilibrium, $h_1/h_2 = 2$, $\sigma = 4$, $t = 6$, and μ goes from 0.67 (before) to 0.69 (after). In Case B, with multiple stable equilibria, $h_1/h_2 = 1$, $\sigma = 2$, $t = 4$, and μ goes from 0.67 (before) to 0.69 (after).

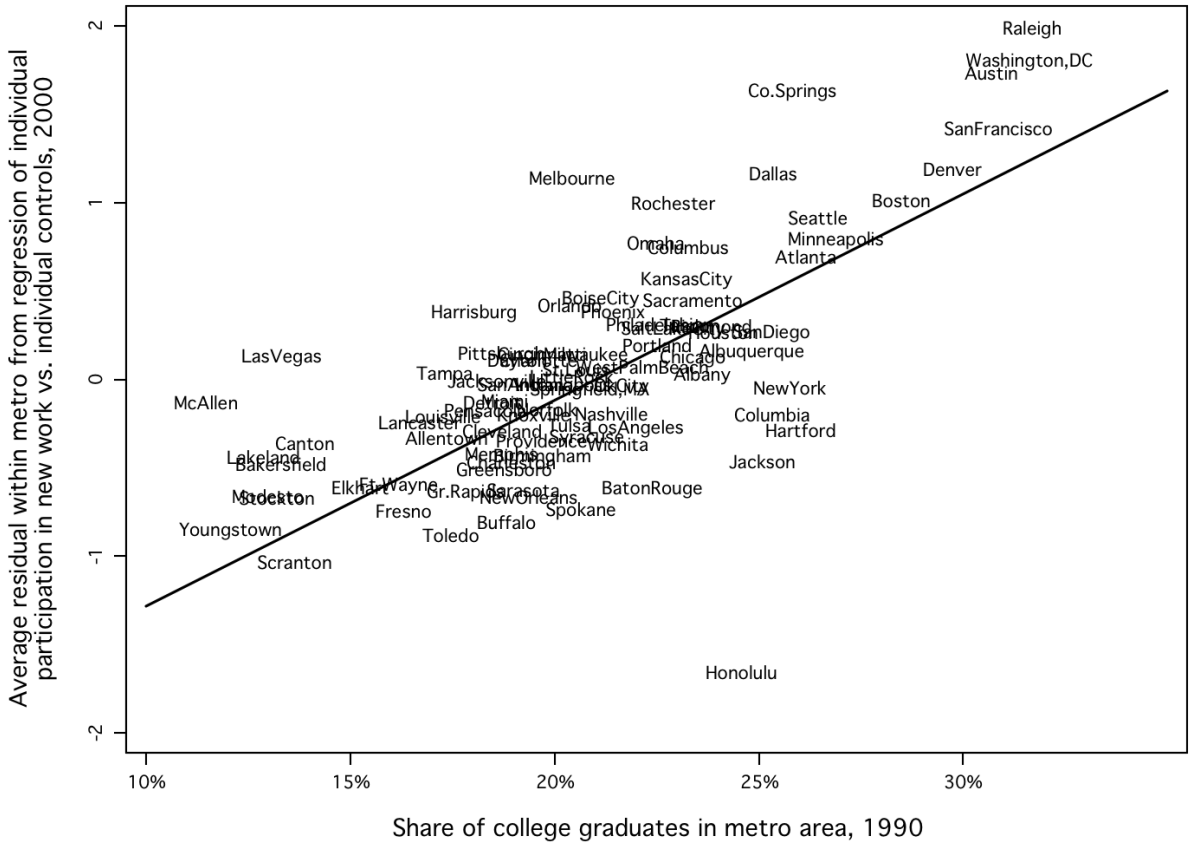


Figure 6: Selection into new work, 2000, and 1990 metro area college share

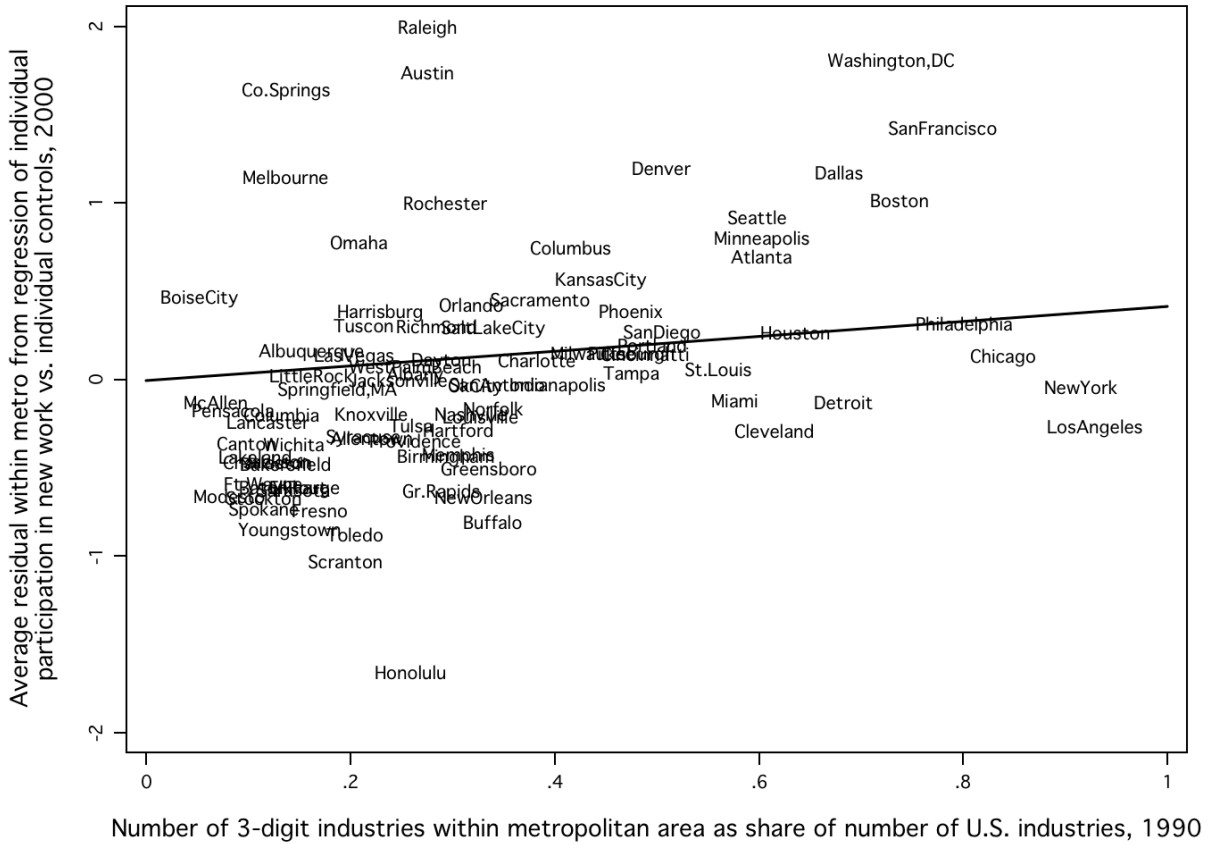


Figure 7: Selection into new work, 2000, and 1990 metro area industrial diversity

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

Census 2000 Detailed Occupation Code	New Title 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

Notes: 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 occupation titles. 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.

Table 2: Employment share in new work in 2000 for various categories

Universe	New work share	Top new DOC(s), in employment
United States, entire sample	4.8%	Computer software engineers
<i>By metropolitan status</i>		
Identified metropolitan areas	5.1%	"
Non-identified	3.7%	"
<i>By gender</i>		
Women	4.6%	"
Men	4.9%	"
<i>By age</i>		
Age < 26	4.3%	Computer support specialists
26-30	5.6%	Computer software engineers
31-40	5.6%	"
41-50	4.8%	"
51 < Age	4.3%	"
<i>By educational attainment</i>		
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	7.9%	Computer software engineers Computer programmers

Notes: Author's calculations using 2000 IPUMS, workers aged 20-65 in identified occupations. Share of employment (%) in new occupations is calculated using new title shares for each detailed occupation code (DOC) by employment in each DOC.

Table 3: Education, wages, TFP and new work

<i>Panel A. Sample mean educational attainment and wages</i>				
	New work	Existing work	<i>t</i>	
Average education, years	14.8	13.2	385.4 **	
College graduate share	0.470	0.214	338.2 **	
Log hourly wage	2.85	2.56	246.3 **	

<i>Panel B. Log hourly worker wages and selection into new work, 2000</i>				
	(1)	(2)	(3)	(4)
New work	0.757 (0.005) **	-	0.373 (0.004) **	0.212 (0.005) **
College graduate	-	0.745 (0.002) **	0.722 (0.002) **	0.623 (0.003) **
Some college	-	0.340 (0.002) **	0.329 (0.002) **	0.261 (0.002) **
High school graduate	-	0.165 (0.002) **	0.162 (0.002) **	0.125 (0.002) **
Potential experience	-	0.036 (0.000) **	0.036 (0.000) **	0.032 (0.000) **
Potential experience squared (x10000)	-	-0.006 (0.000) **	-0.006 (0.000) **	-0.005 (0.000) **
Census region dummies	-	YES	YES	YES
Other controls for race, age, nativity	-	YES	YES	YES
Metro X Industry fixed effects	-	-	-	YES
R-squared	0.02	0.22	0.23	0.30

<i>Panel C. Imputed metro TFP growth, 1990-1996, and new work in manufacturing, 2000</i>				
	(1)	(2)	(3)	(4)
New work share, manufacturing	1.053 (0.338) **	1.314 (0.457) **	0.709 (0.381) †	1.285 (0.447) **
Additional metro controls	-	X	-	X
Log utility patents per capita	-	-	0.015 (0.010)	0.002 (0.013)
Adj R-squared	0.17	0.20	0.20	0.32

Notes: ** - Significant at the 99% level of confidence, * - 95% level, † - 90% level. Robust standard errors in parentheses. In Panel A, sample mean characteristics are reported. New work column includes those workers in occupations with new title shares above the sample mean of $\nu=0.048$. Existing work column includes the balance of the sample. The third column reports *t*-statistic from a test on equality of means. In Panel B, the dependent variable is log hourly wage for workers, calculated from the 2000 IPUMS. Each column is a separate regression. Robust standard errors in parentheses. Each regression contains census region controls. $N=1,932,051$ for all regressions. In Panel C, the dependent variable is imputed TFP growth in manufacturing, 1990-1996, for 88 consolidated metro areas. $N=88$ for all regressions. TFP growth is calculated using TFP estimates for 3-digit SIC industries from the NBER-CES database matched to industry shares in metro areas in 1990. New work share in manufacturing includes employment in manufacturing industries only. Utility patents per capita are patents for inventions, accumulated over 1990-1999, assigned to the metro area of the first-listed inventor. Data on utility patents are from the U.S. Patent and Trademark Office (1999). Additional metro controls are as used in the worker-level regressions in Table 4, column (4).

Table 4: Main results on estimated effects on selection into new work

	[mean] [(s.d.)]	(1)	(2)	(3)	(4)
<i>1990 Metro area characteristics</i>					
College share (see note) (h.s. dropout share omitted)	[0.21] [(0.04)]	-	20.1 (1.9) **	12.5 (1.6) **	10.3 (1.7) **
No. of 3-digit industries as share of U.S. total	[0.34] [(0.21)]	-	3.0 (0.9) **	2.7 (1.0) **	3.5 (0.7) **
Additional metro controls		-	-	-	YES
<i>Worker characteristics</i>					
College graduate (high school dropout omitted)	[0.27]	4.7 (0.2) **	-	4.8 (0.2) **	4.7 (0.2) **
Some college	[0.31]	2.3 (0.2) **	-	2.5 (0.1) **	2.5 (0.1) **
High school graduate	[0.28]	0.6 (0.0) **	-	0.6 (0.0) **	0.6 (0.0) **
Male	[0.51]	1.0 **	-	1.4 **	1.4 **
Black	[0.12]	-0.4 **	-	-0.6 **	-0.6 **
Asian	[0.05]	1.3 **	-	1.1 **	1.2 **
Hispanic	[0.12]	-0.4 **	-	-0.5 **	-0.5 **
Self-employed	[0.10]	-1.1 **	-	-1.4 **	-1.4 **
Age 26-30 (20-25 omitted)		0.5 **	-	0.6 **	0.6 **
31-35		0.5 **	-	0.5 **	0.5 **
36-40		0.3 **	-	0.2 **	0.2 **
41-45		0.0	-	-0.1 †	-0.1 †
46-50		-0.3 **	-	-0.5 **	-0.5 **
51-55		-0.6 **	-	-0.7 **	-0.7 **
56-60		-0.8 **	-	-1.1 **	-1.1 **
61+		-1.0 **	-	-1.3 **	-1.3 **
Adj. R-squared		0.067	0.043	0.070	0.070

Notes: ** - Significant at the 99% level of confidence, * - 95% level, †- 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. 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 are census PUMS 2000, age 20-65, in identified occupations and metro areas. $N = 2,239,672$ in regression 1; 1,537,112 in regressions 2-4 (identified metro areas only). Number of (consolidated) metro areas = 88. Omitted categories are high school dropout and 1990 metro dropout share. Additional controls for marital status, nativity, and major industry dummies included in all regressions. Regressions 2-4 include metro some college and high school share, log population and log land area. These coefficient estimates are not significantly different from zero. Regression 4 includes controls for labor demand shock index and others as described in the text.

Table 5: Estimated coefficients using alternative measures of industrial diversity and new work

Panel A. Estimated coefficient on college share

<i>New work measure</i>	[mean] [(s.d.)]	<i>Industrial diversity measure</i>		
		Industry count	Herf.	Top 20 share
Baseline (manual review)	[5.29] [(12.5)]	10.3 (1.7) **	10.2 (1.5) **	11.0 (1.3) **
Census rules	[5.07] [(10.7)]	6.7 (1.3) **	6.8 (1.1) **	7.2 (1.1) **
Baseline & Census rules	[2.60] [(8.50)]	6.6 (1.0) **	6.7 (0.9) **	7.1 (0.8) **
Wide (string matching)	[12.6] [(16.1)]	10.1 (1.8) **	10.4 (1.5) **	10.2 (1.4) **
Baseline binary 1	[5.54] [(22.9)]	13.2 (2.6) **	13.8 (2.5) **	14.7 (2.0) **
Baseline binary 2 (> 0.9)	[0.20] [(4.45)]	1.8 (0.3) **	1.6 (0.2) **	1.7 (0.2) **

Panel B. Estimated coefficient on industrial diversity

<i>New work measure</i>	[mean] [(s.d.)]	<i>Industrial diversity measure</i>		
		Industry count	Herf.	Top 20 share
Baseline (manual review)	[5.29] [(12.5)]	3.5 (0.7) **	0.12 (0.05) *	0.17 (0.06) **
Census rules	[5.07] [(10.7)]	2.2 (0.6) **	0.07 (0.04) †	0.10 (0.04) *
Baseline & Census rules	[2.60] [(8.50)]	2.0 (0.4) **	0.07 (0.03) *	0.09 (0.03) *
Wide (string matching)	[12.6] [(16.1)]	2.4 (0.8) **	0.05 (0.05)	0.17 (0.05) **
Baseline binary 1	[5.54] [(22.9)]	5.9 (1.1) **	0.16 (0.08) *	0.22 (0.09) *
Baseline binary 2 (> 0.9)	[0.20] [(4.45)]	0.118 (0.122)	0.017 (0.007) *	0.019 (0.009) *

Notes: ** - Statistically significant at the 99% level of confidence; * - 95% level, † - 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. Each cell is a separate regression, using census weights. Dependent variable is selection into a new occupation, new occupation identification algorithm is as indicated by row headings. The *baseline* definition of new work is used throughout the paper and is based on a manual review of 5-digit occupation titles. The *census rules* definition of new work is based on 5-digit occupation titles in 2000 that (1) are not matched to a 3-digit 1990 DOC in census documents, and (2) are indicated new occupation titles for 2000 in census documents. Other definitions are explained in the text. Data and specification are the same as Table 4, column 4. Upper-left cell in each panel reproduces estimates from Table 4, column 4. Cells in corresponding positions from each panel display estimates from the same regression. *Panel A* reports estimated coefficients on 1990 metro college share. *Panel B* reports estimated coefficients on 1990 industrial diversity. Measure of industrial diversity is as indicated by column headings. *Industry count* is the number of identified 3-digit industries within each metro area in 1990, expressed as a share of the number of national 3-digit industries. *Herf.* is a Herfindahl-Hirschman index of employment across 3-digit industries within each metro area in 1990, adjusted for national employment across industries and inverted. *Top 20 share* is the share of employment in the top 20 3-digit industries within each metro area in 1990, inverted. The latter two industrial diversity measures have been standardized to have mean zero and standard deviation 1, so estimates in the last two columns of Panel B are comparable.

Table 6: Selection into new work controlling for additional metro characteristics

	Baseline	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>1990 Metro area characteristics</i>								
College share	10.3 (1.7) **	12.5 (1.6) **	10.3 (1.7) **	12.5 (1.6) **	11.5 (2.0) **	12.6 (1.6) **	7.2 (2.7) *	12.8 (3.5) **
No. of 3-digit industries as share of U.S. total	3.5 (0.7) **	2.8 (1.0) **	3.5 (0.8) **	2.7 (0.9) **	2.1 (1.0) *	2.9 (1.2) *	3.1 (0.8) **	2.2 (1.0) *
B-K predicted employment growth index (1990-2000)	3.5 (4.6)	0.7 (3.4)	-	-	-	-	-	4.0 (7.9)
Own metro-edu group growth (1990-2000)	0.19 (0.05) **	-	0.19 (0.05) **	-	-	-	-	0.18 (0.13)
Predicted patenting index (1990-2000)	0.05 (0.10)	-	-	-0.01 (0.08)	-	-	-	0.09 (0.13)
Log utility patents per capita (1990-1999)	-	-	-	-	0.15 (0.12)	-	-	0.13 (0.12)
Presence of a land-grant college	-	-	-	-	-	-0.04 (0.11)	-	-0.08 (0.08)
Industry & occupation structure	-	-	-	-	-	-	YES	YES
Temperature & precipitation	-	-	-	-	-	-	-	YES
Postgraduate share	-	-	-	-	-	-	-	YES
Adj. R-squared	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07

Notes: ** - Statistically significant at the 99% level of confidence; * - 95% level, † - 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. 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 is census PUMS 2000, age 20-65, in identified occupations and metro areas. $N = 1,537,112$ in each regression. Number of (consolidated) metro areas = 88. All regressions include additional controls for individual characteristics, metro some college and high school share, log 1990 metro population and land area, and census region, as in Table 4, columns 2-4. Baseline regression in column 1 is identical to Table 4, column 4, and reproduced here for reference. Regression (8) contains a total of 43 metro-level controls.

Table 7: Decompositions by industry, geography, and education

<i>Panel A. Decompositions by industry</i>							
	Traded industries		Services		Information		Non-traded industries
[New work mean]	[7.07]		[6.32]		[11.75]		[3.88]
College share, 1990	20.7 (2.7)	**	10.1 (2.2)	**	48.7 (10.8)	**	2.7 (0.8) **
Industrial diversity, 1990	4.5 (1.3)	**	3.6 (0.8)	**	9.9 (4.5)	*	2.1 (0.4) **
<i>Panel B. Decompositions by geography</i>							
	East		West		Metro areas		Consistently identified
[New work mean]	[5.14]		[5.10]		[5.13]		[4.90]
College share, 1990	10.7 (2.2)	**	11.6 (2.1)	**	9.9 (1.5)	**	10.2 (2.0) **
Industrial diversity, 1990	3.1 (0.8)	**	6.0 (3.0)	†	1.8 (0.8)	*	3.0 (0.7) **
<i>Panel C. Decompositions by educational attainment</i>							
	College graduates		Some college		High school graduates		High school dropouts
[New work mean]	[7.90]		[4.75]		[2.52]		[1.65]
College share, 1990	21.6 (2.7)	**	11.3 (1.5)	**	3.1 (0.7)	**	-1.8 (0.7) *
Industrial diversity, 1990	4.6 (1.3)	**	4.0 (1.3)	**	0.6 (0.4)		1.2 (0.4) **

Notes: ** - Statistically significant at the 99% level of confidence; * - 95% level, †- 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. Dependent variable is selection into a new occupation (range 0-100, sample mean 5.3), based on new title share of identified DOC. Each column is a separate regression, using census weights. Each regression contains controls identical to Table 4, column 4, on the sub-sample listed in the column heading. The dependent variable mean in each sub-sample is reported in brackets in each column.

Table 8: Sorting and new work

<i>Panel A. Migration and new work</i>				
	Movers (5 yr.)		Non-movers	
[New work mean]	[5.17]		[5.03]	
College share, 1990	10.3 (1.8)	**	10.2 (1.7)	**
Industrial diversity, 1990	3.2 (0.8)	**	4.2 (0.9)	**

<i>Panel B. State characteristics as IV</i>				
	SOB-1990		SOB-1970	
College share, 1990	10.6 (4.1)	*	8.2 (3.1)	**
Industrial diversity, 1990	3.5 (1.5)	*	2.6 (1.4)	†

Notes: ** - Statistically significant at the 99% level of confidence; * - 95% level, †- 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. Dependent variable is selection into a new occupation. Each column is a separate regression, using census weights. Each regression contains controls identical to Table 4, column 4. In Panel A, movers are those who lived in a different metropolitan area in 1995 relative to survey year reported metropolitan area. In Panel B, corresponding state-of-birth characteristics are used to instrument for metropolitan characteristics. Column 1 uses 1990 state characteristics; column 2 uses 1970 state characteristics.

Table 9: New work using historical data

<i>Microdata - New work year</i>	2000 PUMS		1970 PUMS		1980 PUMS	Pooled, Metro FE	
<i>Lagged metro characteristics year</i>	1970		1950		1970	-	
College share	20.5 (2.4)	**	0.37 (0.17)	**	0.90 (0.81)	21.2 (3.5)	**
Industrial diversity	1.6 (0.9)	†	0.04 (0.04)		0.32 (0.22)	1.6 (0.8)	*
<i>Worker educational attainment</i>							
College graduate	4.7 (0.2)	**	0.06 (0.02)	**	0.85 (0.06)	3.2 (0.3)	**
Some college	2.6 (0.1)	**	0.11 (0.01)	**	0.33 (0.03)	1.3 (0.3)	**
High school graduate	0.62 (0.05)	**	0.05 (0.01)	**	0.10 (0.03)	0.30 (0.31)	

Notes: ** - Statistically significant at the 99% level of confidence; * - 95% level, †- 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. Dependent variable is selection into a new occupation. Each column is a separate regression, using census weights. Each regression contains controls identical to Table 4, column 4. Regression in column (4) pools PUMS data from 2000, 1980, and 1970, and includes metropolitan and census year fixed effects.

Appendix A Further discussion of the 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.

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. The baseline definition of new work, used throughout the paper, comes from consultation with Technical Paper 65 (Scopp 2003) on 1990-2000 OCS revisions and a manual review of the remaining occupation titles.

I consider the baseline 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://jeffr.lin.googlepages.com/>).

A second (orthogonal) algorithm for identifying new work comes from an analysis of detailed internal census documents that track the transition between the 1990 and 2000 OCS. A database obtained from the Census Bureau contains, for each 2000 OCS title, an indicator for whether it was new to the 2000 OCS, corresponding 1990 OCS detailed occupation code(s), and further information about why changes, if any, occurred. I classify a title to be new under *census rules* if: (a) the database indicates that the title was added to the 2000 OCS, and (b) if the title does not have a corresponding 1990 OCS detailed occupation code. Results using this census rules definition appear in Table 6. This database is available upon request.

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 (del Regato 1995). These examples provide intuitive verification of the kinds of changes used in this paper.

A.2 Identifying new occupations 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" (Scopp 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 Table A, 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 B 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 Descriptives and other data

The main body of data is the Integrated Public Use Microdata Data Series (Ruggles and Sobek et al., 2004). These data contain 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 stratified 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 6 contains estimates using only a sample of completely and consistently identified metropolitan areas.

Metropolitan area data come 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 1990-1999 come from the U.S. Patent and Trademark Office (2000). Historical patent data come from the National Bureau of Economic Research (Hall et al. 2001). Data used to construct imputed TFP growth come from the NBER-CES database (Bartlesman, Becker, and Gray, 2000). I match 2-digit SIC manufacturing industries from the NBER-CES to 2-digit industry employment shares in metropolitan areas from the 1990 IPUMS. Data on metropolitan education and industry/occupation shares come from the State of the Cities Data Systems, maintained by the U.S. 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 come from

Moretti (2004). Data on climate come from the *County and City Data Book* (U.S. Census Bureau, 1994), which is then matched to the county of the principal city in each metropolitan area. Table C displays summary statistics for both metropolitan area worker characteristics, for most of the variables of interest.

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, in each region $i = 1, 2$:

$$p_i = [\sigma/(\sigma - 1)]\beta w_i \quad (\text{A})$$

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

$$x = x_i = (f/\beta)(\sigma - 1) \quad (\text{B})$$

The production function (equation (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 .

$$n_i = l_i/(f\sigma) \quad (\text{C})$$

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 be equal to total supply for each variety of traded good (from (1) and (B)):

$$\frac{1}{\beta}f(\sigma - 1) = \frac{p_1^{-\sigma}}{n_1p_1^{1-\sigma} + n_2(tp_2)^{1-\sigma}}\mu e_1 + \frac{t(tp_1)^{-\sigma}}{n_1(tp_1)^{1-\sigma} + n_2p_2^{1-\sigma}}\mu e_2 \quad (\text{D})$$

Each worker/consumer spends fraction $(1 - \mu)$ on housing; therefore, the aggregate value of housing services is $(1 - \mu)(e_1 + e_2)$. Aggregate income is labor income plus income from housing, or $w_1l_1 + w_2l_2 + (1 - \mu)(e_1 + e_2)$. Assume that housing stocks are equally owned by all workers, then total spending by residents of region i equals:

$$e_i = w_i l_i + \frac{l_i}{l_1 + l_2} \frac{1 - \mu}{\mu} (w_1 l_1 + w_2 l_2) \quad (\text{E})$$

Define $v \equiv n_1/N \equiv n_1/(n_1 + n_2)$ (i.e., the share of production activities located in region 1) = l_1/L (by (C)). Define $w = w_1/w_2$ (the wage in region 1 relative to the wage in region 2) = p_1/p_2 (by (A)). Substituting (A), (C), and (E) into (D), I obtain the first equilibrium condition:

$$1 = \frac{vw^{1-\sigma}}{vw^{-\sigma} + (1-v)t^{1-\sigma}} \left[\mu + (1-\mu)\left(v + \frac{1-v}{w}\right) \right] + \frac{(1-v)(tw)^{1-\sigma}}{v(tw)^{1-\sigma} + 1-v} \left[\frac{\mu}{w} + (1-\mu)\left(v + \frac{1-v}{w}\right) \right] \quad (\text{F})$$

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

$$\left(\frac{h_1}{l_1}\right)^{1-\mu} \left(\frac{\mu E_1}{l_1[n_1p_1^{1-\sigma} + n_2(tp_2)^{1-\sigma}]^{1/(1-\sigma)}} \right)^\mu = \left(\frac{h_2}{l_2}\right)^{1-\mu} \left(\frac{\mu E_2}{l_2[n_2p_2^{1-\sigma} + n_1(tp_1)^{1-\sigma}]^{1/(1-\sigma)}} \right)^\mu \quad (\text{G})$$

Substitute (A), (C), and (E) into (G) to yield relative utility $u \equiv u_1/u_2$:


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

Equations (F) and (H) 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 5, Case A, the unique equilibrium configuration, I set $h_1/h_2 = 2$, $\sigma = 4$, 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, Case B, the multiple equilibria configuration, I set $h_1/h_2 = 1$, $\sigma = 2$, 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 Figure A: New work employment share by 2-digit SIC industry, 2000

Notes: Author's calculations using 2000 IPUMS, workers 20-65, with identified industries and occupations.

Appendix Table A: New DOCs from earlier census years, 1960-1980

Panel A. 1960-1970 new DOCs, narrow definition

<i>DOC</i>	<i>Description</i>
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. 1960-1970 new DOCs, medium definition

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

Panel C. 1970-1980 new DOCs, narrow definition

<i>DOC</i>	<i>Description</i>
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 D. 1970-1980 new DOCs, medium definition (selected)

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

Notes: Based on comparison of 1960, 1970, and 1980 census occupation codes. For the period 1960-1970, there are 30 additional new DOCs under the wide definition. For the period 1970-1980, there are 18 additional new DOCs under the medium definition (Panel D) and 26 additional new DOCs under the wide definition. For an explanation of how these data were collected, see the discussion in the data appendix.

Appendix Table B: Employment in new work across census years

Panel A. Share of employment (%) in new occupations

	1970	1980	2000
(1) Narrow definition	0.06	0.02	5.29
(2) Medium definition	0.89	0.72	-
(3) Wide definition	2.84	3.19	12.6

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

	1970 narrow	1980 medium	2000
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	7.9

Notes: Employment shares, in percentage points, calculated from IPUMS 1970, 1980, and 2000, using all occupation-identified (and education-identified, in Panel B) workers, age 20-65.

Appendix Table C: Sample summary statistics

<i>Panel A. Metropolitan areas</i>			
	<i>1950</i>	<i>1970</i>	<i>1990</i>
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)
<i>Panel B. Individuals</i>			
	<i>1970</i>	<i>1980</i>	<i>2000</i>
<i>Workers</i>			
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

Notes: Data: IPUMS 1950, 1970, 1980, 1990, and 2000, and State of the Cities. Metropolitan samples are aggregated from respondents age 20-65 in identified metropolitan areas. Worker samples include all occupation-identified workers, age 20-65, in identified metropolitan areas. Standard deviations in parentheses.

Appendix Table D: Probit estimates for new work, 2000 PUMS

	<i>Industrial diversity measure</i>					
	Industry count		Herf.		Top 20 share	
<i>1990 Metro area characteristics</i>						
College share	0.098 (0.044)	*	0.090 (0.036)	*	0.093 (0.034)	**
Industrial diversity	0.032 (0.018)	†	0.0005 (0.0003)	†	0.002 (0.001)	*
<i>Worker characteristics</i>						
College graduate	0.139 (0.005)	**	0.139 (0.005)	**	0.140 (0.005)	**
Some college	0.090 (0.005)	**	0.090 (0.005)	**	0.090 (0.005)	**
High school graduate	0.029 (0.004)	**	0.029 (0.004)	**	0.029 (0.004)	**
Pseudo R-squared	0.0753		0.0753		0.0753	

Notes: ** - Statistically significant at the 99% level of confidence; * - 95% level, - 90% level. Robust standard errors, adjusted for clustering on metro area, in parentheses. Dependent variable is selection into a new occupation. Each column is a separate regression, using census weights. Each regression contains controls identical to Table 4, column 4. Marginal effects reported; for educational attainment, reported effect is for change of dummy variable from 0 to 1. $N = 1,537,112$ in all regressions.