

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Labor Supply and Personal Computer Adoption

Mark Doms

Federal Reserve Bank of San Francisco

and

Ethan Lewis

Federal Reserve Bank of Philadelphia

June 2006

Working Paper 2006-18

<http://www.frbsf.org/publications/economics/papers/2006/wp06-18bk.pdf>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Labor Supply and Personal Computer Adoption

June 3, 2006

Mark Doms, mark.doms@sf.frb.org
Federal Reserve Bank of San Francisco

Ethan Lewis, ethan.g.lewis@phil.frb.org
Federal Reserve Bank of Philadelphia

The authors would like to thank Meryl Motika and Shannon Mail for excellent research assistance. The authors would also like to thank Daron Acemoglu, Orley Ashenfelter, David Autor, Paul Beaudry, Marianne Bitler, Elizabeth Cascio, Mary Daly, Ellen Dulburger, Fred Furlong, Shane Greenstein, Erica Groshen, Chang-Tai Hsieh, Pete Klenow, Betsey Stevenson, Anita Todd, Robert Valletta, Phillip Vermuellen, Justin Wolfers, and seminar participants at the 2005 NBER summer institute, 2005 ASSA meetings, 2006 SOLE meetings, UBC, WRSA, and at the Federal Reserve Banks of Philadelphia and San Francisco for helpful comments. The views in this paper do not reflect the views of the Federal Reserve Bank of San Francisco, the Federal Reserve Bank of Philadelphia, or the Board of Governors of the Federal Reserve System.

Labor Supply and Personal Computer Adoption

---Abstract---

The positive correlations found between computer use and human capital are often interpreted as evidence that the adoption of computers have raised the relative demand for skilled labor, the widely touted skill-biased technological change hypothesis. However, several models argue the skill-intensity of technology is endogenously determined by the relative supply of skilled labor. We use instruments for the supply of human capital coupled with a rich dataset on computer usage by businesses to show that the supply of human capital is an important determinant of the adoption of personal computers. Our results suggest that great caution must be exercised in placing economic interpretations on the correlations often found between technology and human capital.

JEL Codes: J41, O33

Mark Doms
Federal Reserve Bank of San Francisco
101 Market Street, MS 1130

San Francisco, CA 94105
Mark.doms@sf.frb.org

Ethan Lewis
Federal Reserve Bank of Philadelphia
Research Department
10 Independence Mall,
Philadelphia, PA 19106
Ethan.G.Lewis@phil.frb.org

I. Introduction

The strong, positive correlations found between computer use and human capital are often interpreted as evidence that the adoption of computers have raised the relative demand for skilled labor, the widely touted skill-biased technological change hypothesis.¹ However, models such as Acemoglu (1998) and Beaudry and Green (2005) argue the skill-intensity of technology is endogenously determined by the relative supply of skilled labor.² By exploiting several instruments for the supply of human capital coupled with a rich dataset on computer usage by businesses, we find that the supply of human capital is indeed an important determinant of the adoption of personal computers (PCs). This result is robust to a wide range of measures of skill and to competing models of technology adoption. Overall, our results suggest that great caution must be exercised in placing economic interpretations on the correlations often found between technology and human capital.³

Our empirical approach exploits the wide cross-city variance in PC use and human capital during the 1990s, the period in which the personal computer came to dominate the information technology environment.^{4,5} In order for our approach to be successful, several estimation issues must be addressed. The first, and perhaps most important, issue is that cross-city differences in human capital are potentially endogenous. We must therefore develop instruments that capture differences in human capital supply and that do not affect technology use through other channels. Another issue is that our measure of human capital – the share of the workforce that is college educated – is admittedly crude and may be proxying for other, unobserved characteristics of the population. For example, cities with highly educated workforces may also possess large numbers of computer programmers or others with specific skills that make adopting computers more profitable. The third and final issue we confront is model misspecification. The technology diffusion literature is home to many models, especially

¹ One exception is Autor, Levy, and Murnane (2003), who argue that industries intensive in what they called “routine cognitive” tasks were most likely to adopt computers because these were the tasks computers replaced. There has been research on cross-country differences in technology adoption, which we address below.

² Acemoglu’s model reveals the bias inventors have toward developing technologies that raise the relative productivity of skilled labor when it is abundant. Beaudry and Green’s model emphasizes how a cost-minimizing firm’s choice to adopt a new more skill-intensive technology depends on the relative supply of skill.

³ Beaudry, Doms, and Lewis (2006) builds upon the results in the present paper with an endogenous model of technology adoption that simultaneously addresses how human capital affects technology adoption and how technology adoption affects labor demand. The primary goal of the present paper is to rigorously establish the human capital-technology adoption portion.

⁴ This spatial analysis approach is quite prevalent in research evaluating the labor market impact of immigration, but outside of cross-country studies it is rare in the literature examining recent technological change.

⁵ The large and surprisingly persistent differences in computer adoption across geographic markets in the U.S. has received little attention, although Nestoriak (2004) is a nice exception.

models that emphasize vintage effects and spillovers. For instance, our data show that the San Francisco Bay area is very computer-intensive outside of its high-tech industries; does that result arise because the Bay area is also home to Silicon Valley (which may generate spillovers) or because the Bay area has a highly educated workforce?

In terms of the first issue, endogeneity between technology and skills, we develop separate instruments for the level of human capital across cities and also for the change in human capital across cities, all based on pre-PC city characteristics. The instrument for cross-city level differences in the skill mix is based on pre-PC differences in the presence of colleges and universities, building upon instruments used by Moretti (2004). The logic behind these instruments is that residents of college-abundant localities will enjoy lower costs of attending college; such localities may also have amenities that college graduates value. The instrument we develop for the changes in college share is based on the locations of immigrant enclaves in 1980; immigrant enclaves can predict changes in college share because of the combined effects of the strong tendency of new immigrants to settle into existing enclaves (Bartel, 1989) and the boom in less-skilled immigration the U.S. has experienced in recent decades.⁶

To address the issue of whether our measure of the skill level of a city may proxy for other, more specific skills that are important in the PC adoption decision, we examine the robustness of our results to using more detailed occupational categories such as engineers and computer support workers. Further, we examine the age distribution of the workforce in each city to test whether younger, college-educated workers may be more PC-savvy, as suggested by Goldfarb (2005). Another variable that may capture the unobserved component of skill is the rankings of local computer science and electrical engineering departments; areas that have highly ranked departments may have workers with above average skills.

How intensively an area adopts PCs depends on factors other than the availability of certain types of labor. For information technology, spillovers (as in Goolsbee and Klenow, 2002) and vintage effects (as in Bresnahan and Greenstein, 1996) are additional factors in the adoption process. For spillovers, we examine how the diffusion of PCs is related to the presence of the IT producing sector. Additionally, we examine how technology of establishments is related to the company to which they belong. Vintage effects occur when the technology adopted at time t depends on the technology in place at time $t-1$. We address potential vintage effects in several ways, including looking at relatively long first differences, controlling for the

⁶ Card (2001) and Lewis (2005) also use instruments based on a similar concept.

computer capital stock in 1980 (a date before the introduction of the IBM PC), and controlling for the major computer systems in place in 1990.

Our basic finding is that establishments in areas with more college-educated labor adopt significantly more PCs per worker. The relationship is very robust to model specification and is also economically significant: half of the inter-city variation in computer use is explained by this single variable.⁷ A similar relationship emerges between changes in PC-intensity and changes in college share over the 1990s. The relationships are robust to the use of our instruments. We find little partial relationship between PC-use and an area's supply of young college graduates. Although the presence of certain specific types of highly skilled workers (engineers, programmers) and a software-producing sector is also associated with higher PC-intensity, controlling for these factors has little effect on the relationship between college share and PC use.

The evidence in this paper is consistent with previous research finding an association between computers and skilled labor across plants and industries, but makes a stronger case for the complementarity between the two by taking advantage of differences in factor mix which are more plausibly exogenous. This also arguably provides preliminary support for endogenous technical change models like Acemoglu (1998) and Beaudry and Green (2005), although factor substitution in a more standard framework can also account for this relationship. Distinguishing between these models to ascertain how PC adoption may be related to changes in relative wages is the subject of Beaudry, Doms, and Lewis (2006).

In addition, our findings are related to cross-country studies of technology adoption, which also tend to find a strong relationship between adoption and the education of the population (e.g., Caselli and Coleman, 2001, and Comin and Hobijn, 2004). Cross-country studies are potentially confounded by the large unobservable differences across countries. We believe this problem is diminished by looking across U.S. metropolitan areas and using instruments to assist in identification.⁸

⁷ This result is even more surprising given that the PC data and the labor data come from two independent sources.

⁸ Within the U.S. this paper is consistent with findings by Nestoriak (2004), who examines computer investment in a subset of U.S. counties, and Lewis (2005), who examines adoption of manufacturing automation technologies.

II. Personal Computers

II.1 Basic facts

In this paper we focus on one important facet of information technology, the personal computer (PC). On a heuristic level, the PC represents the epitome of changes in business computing over the past several decades, namely the migration from mainframe platforms to client/server platforms (see Bresnahan and Greenstein, 1996). The PC has also increased its role. Initially being mainly a stand-alone device used for office automation tasks such as spreadsheets and word processing, it has now become a communications device, a terminal, and in some cases, a server.

On an empirical level, nominal and real spending on PCs grew sharply during the 1990s, and spending on PCs outstripped spending on other types of computers. As shown in Table 1, real computer investment during the 1990s was propelled by strong gains in nominal spending (10 percent per year) and a substantial drop in price (18 percent). The segment within the computers and peripherals category that posted the largest increases was the PC, averaging a phenomenal 50 percent annual growth rate. Spending on other computers (including mainframes, midranges, and workstations) grew more slowly in both nominal and real terms.⁹ Additionally, during the whole of the 1990s, businesses spent nearly 90 percent more on PCs than they did on other computers, and PCs accounted for more than a third of all computer investment.¹⁰

II.2 Data on PCs used in this study

For this study, we use establishment-level data from 1990 to 2002 (even years only) from a private company, Harte-Hanks (HH). The HH data come from an establishment-level survey that collects information about specific technologies, such as the number of IBM mainframes (and what type), the number of Pentium IV PCs, what software applications are used, and so forth. Nearly all of the Harte-Hanks data are based on counts of particular technologies. For this study, there are approximately 80,000 usable observations per year.¹¹ Additionally, the HH data provide basic information on the establishment, such as the corporation to which it belongs,

⁹ The high rates of growth in real PC investment boosted the growth rate in the capital stock: according to Bureau of Labor Statistics, between 1990 and 2001, the real PC stock grew an average of 42 percent annually whereas the growth rate in the stock of other computers chalked up a lower, but still respectable, 25 percent average rate.

¹⁰ Additionally, the PC investment numbers do not include other IT investments that are complementary to PCs, such as printers, networking equipment, and prepackaged software.

¹¹ A number of studies have used the HH data to examine technology diffusion (e.g., Forman, Goldfarb, and Greenstein 2005) and the relationship between technology use and productivity (e.g., Brynjolfsson and Hitt 2003).

where the establishment is located, the size of the establishment (in terms of employment and revenue), and detailed industry codes of the establishment. In our analysis, we examine establishments with five or more employees and establishments in the private nonfarm sector.

Let $\gamma_{i,c,t}$ be the technology for establishment i in city c at time t (in this study we focus on PCs/employee). We estimate, using OLS, the following,

$$(1) \quad \gamma_{i,c,t} = \sum_{t=1990}^{2002} \left[\beta_{I,t} Ind_{i,t} * Size_{i,t} + \beta_{C,t} City_{i,t} + \beta_{Y,t} Year_{i,t} \right] + \varepsilon_{i,t}$$

where Ind , $Size$, and $City$ are vectors of dummy variables of industry (3 digit SIC) of the establishment, size of the establishment (eight employment size classes), where the establishment resides (over 200 consolidated metropolitan statistical areas (CMSAs), which we often refer to as “cities” or “metro areas”), and the year of the observation (even years including and between 1990 and 2002).^{12,13} We mostly focus on the $\beta_{C,t}$ coefficients, which vary over time and capture the mean differences in technology use across cities after controlling for over 950 industry/size interactions: we often refer to $\beta_{C,t}$ as the adjusted PC intensity or adjusted PCs per employee.¹⁴

One issue that arises from estimating (1) is that the HH data are not a random sample, and some industries are poorly covered. We drop observations where the HH coverage is particularly slim, such as retail, farming, and mining. We also excluded establishments in the IT producing sector for reasons that are explained below. Additionally, we estimated (1) with and without sample weights, where sample weights were derived from a match between the HH data and County Business Patterns. We also experimented with other modifications, such as examining establishments with more than 100 employees (the HH coverage of these establishments is much higher than for smaller establishments). Regardless of the various rules we used in estimating (1), the rank-order of cities was largely unaffected, though the point estimates varied some, especially for cities where sample sizes are small. However, the results presented throughout this paper appear to be very robust to the sample used in estimating (1).¹⁵

¹² The dependent variable in (1) is often truncated at 0 and at an upper bound dictated by employment. To address this distributional issue, two sided Tobit models were also estimated and the results generated are nearly identical.

¹³ We modified the official definition of several CMSAs to obtain a better concept of local labor markets.

¹⁴ This specification assumes that an establishment in a particular 3-digit industry is conducting the same activity across cities. For instance, a commercial bank in New York City is doing the same activity as a commercial bank in Houma, Louisiana. More refined measures of industry, if they were available, would increase the comparability of establishments across regions.

¹⁵ Comparing the HH data to data from the CPS and SIPP, and from statements from HH, it appears that establishments in the HH sample tend to be more technologically intensive than establishments not in the sample.

Figure 1 plots $\beta_{C,t}$ for the year 2000 against 1990 (where 1990 is actually the average of 1990 and 1992 and 2000 is the average of 2000 and 2002) in 230 metropolitan areas.¹⁶ The variation across cities in PC-intensity is large and highly persistent.¹⁷ For example, the San Francisco Bay area consistently ranks very high in nearly all measures of technology that we examined. In 1990 the mean establishment in San Francisco had 15.8 more PCs per 100 employees than the mean establishment in Hickory, North Carolina (a CMSA that frequently ranks low on technology use), even after controlling for industry differences across the two cities.¹⁸ In 2000, the difference in PC intensity between the Bay area and Hickory increased to 27.9. More generally, the correlation in $\beta_{C,t}$ between 1990 and 2000 is 0.62, almost identical to the Spearman rank correlation. In addition, a fitted line has a slope slightly larger than 1.0 (though not statistically distinguishable from 1.0), so on average cities' PC intensity diverged slightly between 1990 and 2000.¹⁹ Despite the persistence, there is quite a bit of variation around the line; for example, between 1990 and 2000 about one-third of cities (78) added more PCs per worker than did San Francisco.

Equation (1) only controls for industry and size. Further controls were added to reflect other factors that may influence the technology used at an establishment. In particular, we are interested in two potential influences on the PC-adoption decision that could affect the cross-city patterns in Figure 1. The first set of additional controls addresses whether vintage effects may be important and the second controls for firm ownership.

For the potential vintage effects, we are concerned about a story where, for instance, firms in the Boston area disproportionately relied on Digital Equipment Corporation (DEC) systems, and hence may have had a lower PC propensity than establishments in the San Francisco Bay area. The short answer to the question of whether the inter-city differences in PC propensity in 1990 are correlated with inter-city differences in other computer platforms appears to be “no.” The HH data provide information on the major system and the vendor of the major system used at the establishment, as well as whether mini, mainframe, and workstation

¹⁶ One outlier, Las Cruces, New Mexico, is not shown in the figure. We restrict the sample to establishments that are in CMSAs, as reliable labor market data are available only in metropolitan areas. The vast majority of U.S. employment is metropolitan. Non-metropolitan areas, we find, have lower PC intensity.

¹⁷ Large regional differences in technology adoption has been observed in other contexts, including hybrid corn (Griliches, 1957) and the use of beta-blockers (Skinner and Staiger, 2005).

¹⁸ Hickory, North Carolina is famous for its furniture manufacturing, and it now is also a major producer of fiber optic cable. We chose Hickory as an example because the estimate for the city is based on over 300 establishment-observations, greater than the sample sizes of other cities that ranked lower in PC use.

¹⁹ The regression of adjusted 2000 PC intensity on adjusted 1990 PC intensity, weighted by the square root of the 2000 sample size, has a slope of 1.12 and a heteroskedasticity-corrected standard error of 0.078.

computers are at the site. We modify (1) to control for the presence of mainframes, minis, and workstations, along with controls for the major vendors (IBM, HP, DEC, Wang, Unisys, or other) and controls for whether the major IBM system was a System 36 or AS/400.²⁰ Although these variables are good indicators of the PC intensity of an establishment, they explain very little of the cross-city differences in PC intensity. The differences between $\beta_{C,t}$ from (1) and from (1) modified to include major system controls tend to be less than 1 PC per 100 employees for most cities, and a regression between the two measures yields a coefficient extremely close to 1 and an R-squared of 0.94.

A similar result arises when we explore firm ownership. For example, a Wal-Mart store in Hickory, North Carolina, is likely to have technology based on its corporation and not just on the labor market conditions in Hickory (an area that has relatively low levels of education). More generally, areas that contain many branches of corporations that are relatively IT intensive may themselves become IT intensive.²¹ To derive information on the corporate structure across cities, we use the corporate identification information provided in the HH data. A corporate identifier is provided for establishments that belong to large corporations, and basic information about the corporation is provided (such as employment and revenue). Specifically, for each establishment, we measure the PC intensity of other establishments that belong to the firm outside of the city and include that term in (1). Not surprisingly, we find the coefficients on the corporate terms to be positive and significantly different from 0; that is, establishments that belong to PC intensive firms themselves are more PC intensive even after controlling for the city in which they are located and detailed industry and size class. However, for most cities, corporate spillovers alter the estimates of $\beta_{C,t}$ by less than 2 PCs/100 employees. In other words, in reference to Figure 1, the San Francisco Bay area does not have a high PC intensity simply because it has branches of corporations that are PC intensive, and Hickory, North Carolina, does not rank low on the PC intensity scale because it lacks branches of PC intensive firms.

In summary, after controlling for the industry and size of establishments, there is tremendous heterogeneity across cities in their propensity to use PCs, and this propensity is

²⁰ We chose these categories because they were the most abundant. In our data, there are literally hundreds of different types of major systems. An establishment with a particular type of DEC mainframe would have the dummy variables for DEC and MAINFRAME set equal to 1, but we do not distinguish between the different types of DEC mainframes.

²¹ One area where intra-corporate transfers of technology have been considered important is in the foreign direct investment literature.

somewhat persistent over time. Further, the inter-city heterogeneity does not appear to be the result of differences in the types of computer platforms across cities nor of differences in corporate ownership across cities.

III. Instruments for Human Capital

All else equal, it would not be surprising that a city with a highly skilled workforce would use some technologies more intensively than cities with lower skilled workforces. One reason might be that human capital and technology are complements, and, indeed, that theory lies behind several models, notably those developed by Beaudry and Green (2005) and Acemoglu (1998).²² The implications from both of these models are relatively straightforward: areas that have a relative abundance of skilled workers are also likely to have relatively higher levels of technology. In the case of the Acemoglu model, this is because high relative employment of skilled workers induces innovation directed toward raising the relative output of skilled workers; in the case of the Beaudry and Green model, it is because areas with more skilled labor have a comparative advantage in producing with a technologically intensive technique.

In a regression framework, an estimable equation could have the following form:

$$(2) \quad \beta_{c,t} = \alpha_0 + \alpha_1 HC_{c,t} + \alpha_2 X_{c,t} + \varepsilon_{c,t}$$

where $\beta_{c,t}$ is the adjusted city-level technology use in city c at time t estimated from equation (1) in the previous section, $HC_{c,t}$ is a measure of human capital in the city, $X_{c,t}$ is a vector of other characteristics of the city, and $\varepsilon_{c,t}$ are unobserved characteristics of the city that are related to technology adoption and measurement error. A problem with interpreting estimates of α_1 in (2) as the impact of human capital supply on technology use, as we would like to do, is that there may be determinants of technology adoption omitted from (2) (captured by $\varepsilon_{c,t}$) which are correlated with $HC_{c,t}$; if so, estimates of α_1 will be biased.²³ Our approach to this problem is to

²² Two papers recently have conducted empirical work examining technology use and local labor market conditions based on the Beaudry and Green and Acemoglu models. Lewis (2005) examines the use of a set of manufacturing technologies and Nestoriak (2004) examines computer investment.

²³ An example is if there are economies of scale in adoption of new technologies, then denser urban markets may be more likely to adopt first. If higher skill workers have a higher willingness-to-pay to live in dense urban markets, as seems to be true empirically, then one would observe (at least partially) spurious correlation between skill mix and technology use driven (in part) by market density. This is just an example, though: in fact, the relationship between city size and technology adoption is quite weak in our sample.

construct instruments which affect the supply of human capital but, we argue, have little or no direct effect on technology adoption.

Another strategy we pursue is to examine the determinants of changes in technology use over the 1990s, which reduces bias coming from any persistent (and time-invariant) city-level variables we do not observe.²⁴ A modified first difference of (2) is:

$$(3) \quad \Delta\beta_{c,t} = \tilde{\alpha}_0 + \tilde{\alpha}_1\Delta HC_{c,t} + \tilde{\alpha}_2X_{c,t-1} + \tilde{\alpha}_3\Delta X_{c,t} + \Delta\varepsilon_{c,t} + \Delta\eta_{c,t} .$$

As was the case for equation (2), we need instruments for $HC_{c,t-1}$ and $\Delta HC_{c,t}$ for equation (3).²⁵ Before discussing the instruments, we first need to define explicitly our primary measure of human capital, $HC_{c,t}$. Similar to other research examining the impact of skill-biased technological change (such as Katz and Murphy, 1992; Autor, Katz, and Krueger, 1998; and Card and DiNardo, 2002), we define our measure of human capital as the fraction of the area's workers who have a least a four-year college degree plus one-half of the fraction with one to three years of college education. We refer to this as the college-equivalent share, or college share. The college share of cities is highly persistent between 1990 and 2000; the regression coefficient is 0.99 and is not significantly less than 1.0. Also, the college share increased an average of five percentage points during the 1990s.

As a prelude to the results that follow, Figure 2 shows our measure of college share in 2000 with our measure of PCs per worker. The two measures come from different sources but, as discussed more fully below, are very strongly related. One issue that arises from Figure 2 is that the PC measure is adjusted for industry mix, while the human capital measure is not. One reason behind the inter-city variation in college share is the industry mix of cities; without making a statement about causality, cities with industries that favor skilled workforces may also be cities that have above-average skill levels. We constructed an industry-adjusted skill measure of college share based on the three-digit industry composition in each city, and Figure 3 shows a scatter plot between the adjusted and unadjusted measures. Although the industry composition is

²⁴ First difference estimates can be biased by time-varying city unobservables and because the impact of unobservables may vary over time. The required assumption for first difference estimates to suffer from less bias than cross-sectional estimates is that $Cov(\Delta HC_{c,t}, \Delta\varepsilon_{c,t})/Var(\Delta HC_{c,t}) < Cov(HC_{c,t}, \varepsilon_{c,t})/Var(HC_{c,t})$

²⁵ Note that this specification can be modified to include the level of human capital, which would allow us to evaluate the view that highly skilled workers are better able to learn about the new technology and employ the technology profitably, something suggested in the findings of several applied papers. In the case of computers, for instance, it is plausible that a college-educated workforce might be better prepared to switch over to the new PC technology during the 1990s, implying that areas with relatively high initial levels of human capital adopt more quickly than areas with lower levels of education. Another reason both levels and lags of the independent variables might appear in a first difference of (2) would be if the coefficients in (2) were allowed to be time-varying. However, the results that follow are insensitive to the inclusion of the lagged level of human capital, although sometimes the coefficient is positive and statistically significant.

related to the skill share, the correlation between the observed skill level and the industry adjusted skill level is .96. In the results that follow, we opt to use the unadjusted measure because our instruments are based on that measure.²⁶

Our instruments for $HC_{c,t}$ are based on the historical density of colleges in an area. The idea is that a local college shifts out the supply of educated labor by reducing the cost to the area's residents of obtaining higher education. Human capital theory predicts that otherwise similar individuals will complete more education when the cost of obtaining it falls (Card, 1999). Supporting this are many studies at the individual level that show that the distance a person lives from a college predicts their college attainment (e.g., Kane and Rouse, 1995; Card, 1995).²⁷ Additionally, cities with colleges may provide amenities that college graduates value. And, finally, cities with an abundance of colleges may have a higher college share in the workforce because of search costs; college graduates may stay close to where they attended college because finding a job nearby is less costly than finding a job elsewhere.

We have a collection of college instruments. One specific instrument we use, following Moretti (2004), is a dummy for whether or not the metropolitan area has a land-grant college. Land-grant colleges came into existence after Congress passed the Morrill Act in 1862. The legislation gave states land to fund the creation of university-level agricultural schools. As Nervis's (1962) history describes, after these land-grant colleges were founded, many changed from being strictly agricultural schools and developed into large universities (for example, University of Minnesota, University of California, and University of Maryland). These schools dramatically increased access to higher education: Moretti (2004) showed that areas with land-grant colleges even today tend to have a significantly higher college-educated share. Given the long lag from the founding of these schools until now and their original purpose of providing support for agriculture, it is reasonable to think that the location of these schools is unrelated to unobserved determinants of regional differences in technology and skill mix today. Owing to the fact that each of the 50 states received at least one land-grant college, the metropolitan areas that have land-grant colleges are quite regionally diverse. Our sample contains 35 such locations, which are listed in Table 2.

In addition to the land-grant colleges, our instrument set also includes lagged information on local college density. Among other things, there has been a dramatic growth in two-year

²⁶ Because the adjusted and unadjusted labor measures are so highly correlated, it is not surprising that our basic results and summary stand when using the adjusted measure.

²⁷ At an aggregate level, another channel through which college density may raise attainment is that people who go away to college may be more likely to search for a job in the labor market where they attended college.

colleges since World War II (documented in Kane and Rouse, 1999) that may have raised educational attainment in areas that received new schools. To capture the effect other colleges may have on local college share, we construct additional instruments using enrollment at two- and four-year colleges in 1971.²⁸ Specifically, we compute the number of college “seats,” CS_c , per working-age population in 1971:²⁹

$$CS_c = \frac{0.5Enr_{c,1971}^{2yr-coll} + Enr_{c,1971}^{4yr-coll}}{Pop_{c,1971}}$$

where $Enr_{c,1971}^{2yr-coll}$ and $Enr_{c,1971}^{4yr-coll}$ are total enrollment at two- and four-year colleges, respectively, and $Pop_{c,1971}$ is the total population aged 15-64, in area c in 1971.³⁰

Both the land-grant college dummy and the college seats instruments are very strong predictors of college share, as shown in the first five columns of Table 3; the coefficients possess the expected signs and are significant, and F-statistics are large even when conditioning on other controls used in the regressions (discussed further below). A visual representation for the results for land-grant colleges is shown in Figure 4, which plots kernel density estimates of college share in metropolitan areas with and without land-grant colleges. The distribution of college share is distinctively shifted to the right in areas with a land-grant college compared to areas without in both 1980 and 1990. Figure 5 similarly shows the strong upward sloping relationship between 1971 college seats per capita and college share in 1990 or 2000.

Are these instruments valid? One concern is that universities themselves are large employers of college graduates (though in most places they employ only a small proportion of the population). A small number of cities in our sample have student enrollment that makes up a large fraction of the population. However, dropping those cities from Figure 5 strengthens the first stage, rather than weakening it. Similarly, Moretti (2004) found that land-grant colleges continue to predict college share after dropping university-employed workers and college towns. Another concern about these instruments is that universities may have some direct effect on local technology adoption. Some IT firms are spin-offs from university labs, and IT firms may have local spillover effects on technology adoption. We discuss controls for such spillovers below.³¹

²⁸ Enrollment data are from the 1971 Higher Education General Information Survey (HEGIS) Institutional Characteristics file.

²⁹ Enrollment at land-grant colleges are excluded from this calculation.

³⁰ Unrestricted first-stage regressions support this restricted functional form for the instrument.

³¹ Another potential concern is that rather than being a supply shifter, the presence of universities could reflect demand for education. One a priori reason to doubt this is that most nonprofit universities in the U.S. cover a significant portion of their costs using nontuition sources (e.g., government grants, endowment earnings) which may

While the college capacity of an area can help predict the long-run average college completion rates of an area, we also need instruments for decadal change in college share, ΔHC_{ct} , for our first difference specification (3). Our strategy in this case makes use of the fact that the U.S. is in the midst of an immigration boom which has rapidly altered the skill mix of the workforce, particularly in markets where immigrants cluster. In the past 35 years, the share of U.S. workers who are foreign-born has risen from 5 percent (at its low point in 1970) to almost 15 percent. Along with the growth in volume of immigration has been a shift in the origins of immigrants to the developing world, which has resulted in a mix of immigrants that is less educated, in most places, than U.S.-born workers. Thus where immigrants concentrate, the college share may rise more slowly than in other cities.

Our immigration instrument relies on newly arrived immigrants having a strong tendency to cluster in markets where earlier waves of immigrants from the same part of the world settled (documented in Bartel, 1989). A place with a cluster of immigrants from some part of the world in 1980 tends to have continued migration from that part of the world in subsequent decades. The location of such immigrant “enclaves” has long been used as an instrument for changes in local skill mix in a long labor economics literature examining the impact of immigration on native-born wages and employment (Altonji and Card, 1991, is an early example), as well as more recent work examining the impact of local skill mix on on-the-job computer use and the adoption of advanced manufacturing technologies (Lewis, 2004, 2005).

To construct the instrument, let $N_{I,g,t}$ represent the total number of immigrants who arrived from country g (say, China) in the U.S. during some period $t-1$ to t (in practice, this will be 1990-2000). Further, let $\phi_{g,c,1980}$ be our enclave measure by representing the proportion of immigrants from g who lived in city c in 1980 (for example, if $\phi_{g,c,1980}=0.21$ for $c = \text{San Francisco}$ and $g = \text{China}$, then 21 percent of Chinese immigrants living in the U.S. in 1980 were settled in San Francisco). Imagining all new immigrants from each country continue to settle

not respond strongly to demand for education. One way this concern’s importance might be evaluated empirically is to check whether the presence of universities is correlated with high school completion, which might proxy for demand for education but should not be affected by the presence of universities. Unlike the rightward shift for college share, areas with land-grant colleges show little difference in the relative presence of high school graduates compared to areas without land-grant colleges. College seats are positively correlated with graduates/dropouts, but the relationship is not statistically significant in 1990 or 2000 and declines when outliers are removed. It therefore seems unlikely that the 1971 location of universities is strongly related to demand for education in 1990 or 2000, though this cannot be completely ruled out.

proportionately into the same mix of cities as they did in 1980, the enclaves predict that

$$\hat{N}_{I,c,t} = \sum_g \phi_{g,c,1980} N_{I,g,t} \text{ new immigrants will arrive in city } c \text{ during the decade.}$$

To bring skill mix into this, let $N_{I,g,t}^{CE}$ represent the number of college equivalents who arrived from g during the 1990s. The expected number of college equivalent immigrant arrivals to city c during the decade is, similarly, $\hat{N}_{I,c}^{CE} = \sum_g \phi_{g,c,1980} N_{I,g}^{CE}$. Thus $\hat{S}_{I,c}^{CE} = \hat{N}_{I,c}^{CE} / \hat{N}_{I,c}$ is the college share among “expected” immigrants. The instrument is then defined to be the amount by which a city’s college share would change if these expected inflows actually occurred and there were no other changes in the city’s population:³²

$$IM_{c,t} = \left(\hat{S}_{I,ct}^{CE} - S_{c,t-1}^{CE} \right) \frac{\hat{N}_{I,c,t}}{N_{c,t-1} + \hat{N}_{I,c,t}}$$

where $S_{c,t-1}^{CE}$ and $N_{c,t-1}$ are the initial college share in the city and number of workers in city c . In our work, all countries are partitioned into 17 country-of-origin groups, and all of the variables are derived from the appropriate decennial census files. These 17 groups, as well as the inflows $N_{I,g,t}^{CE}$ and $N_{I,g,t}$, are listed in Appendix Table 1.

Another way to describe this instrument is that if immigrants continued to settle in the same cities as they did in 1980, and there were no other changes in cities’ populations, a regression of the actual change in college share, ΔS_{ct}^{CE} , on this instrument would have a coefficient of 1.0.³³ In fact, the first-stage coefficient is 0.878, not statistically distinguishable from 1.0, as shown in column (6) of Table 3. This is consistent with research which shows that similarly skilled native-born workers do not move out of cities in response to immigrant inflows (e.g., Card, 2001). The first stage is also shown graphically in Figure 6. As can be inferred from the figure, most cities have too little immigration to have any noticeable impact on skill mix. In the ones where immigration is significant, however, it generally depresses the change in college

³² To see this, note that college share can be disaggregated into the parts due to past residents and to new immigrant arrivals, assuming no other changes in population:

$$S_{c,t}^{CE} = \frac{N_{c,t}^{CE}}{N_{c,t}} = \frac{N_{c,t-1}^{CE} + \hat{N}_{I,c,t}^{CE}}{N_{c,t-1} + \hat{N}_{I,c,t}} = \frac{S_{c,t-1}^{CE} N_{c,t-1} + \hat{S}_{I,c,t}^{CE} \hat{N}_{I,c,t}}{N_{c,t-1} + \hat{N}_{I,c,t}}$$

and given this, the change in college share gives the instrument:

$$\Delta S_{c,t}^{CE} = S_{c,t}^{CE} - S_{c,t-1}^{CE} = \frac{S_{c,t}^{CE} N_{c,t-1} + \hat{S}_{I,c,t}^{CE} \hat{N}_{I,c,t}}{N_{c,t-1} + \hat{N}_{I,c,t}} - S_{c,t-1}^{CE} = \left(\hat{S}_{I,c,t}^{CE} - S_{c,t-1}^{CE} \right) \frac{\hat{N}_{I,c,t}}{N_{c,t-1} + \hat{N}_{I,c,t}}$$

³³ We thank an anonymous referee on Lewis (2005) for suggesting this formulation of the instrument.

share as immigrants tend to be less skilled than native-born workers. One exception to this is Honolulu, which is the positive outlier in the figure. It has a disproportionately skilled immigrant mix, mostly from Asia. Honolulu lies very near the first-stage regression line.

Controls have little effect on this first-stage relationship. A nice specification check is that the college presence instruments, added in column (7), do not strongly predict changes in college share (and, if anything, move the point estimate on immigrant enclaves closer to 1). Symmetrically, the immigration instrument does not predict the level of college share (column 5). The validity of the immigrant enclave instrument is further demonstrated in Figure 7, which plots computer sales per worker in 1980 versus the instrument.³⁴ The figure shows that despite the large cross-city variation in computer sales per worker in 1980, the instrument has little power to predict it.³⁵ Thus the location of immigrant enclaves seems to be unrelated to pre-existing differences in technology adoption.

Finally, as with any instruments, these only provide credible identification to the extent that one accepts the assumption that they do not have a significant effect on the outcome (PC intensity) through a channel other than their effect on the independent variable (college share). There are reasons why this might not be the case (the most obvious of which we attempt to address) that could lead our results to be biased. However, the instruments rely only on city characteristics prior to the introduction of the IBM PC in 1981, and thus can at least be defined as predetermined if not truly exogenous.

Table 4 shows the results of bivariate regressions of PC-intensity on college share. The table shows both OLS and IV estimates from separate 1990, 2000, and first difference (1990-2000) regressions. Each regression uses weights equal to the square root of the number of observations in the Harte-Hank dataset, and standard errors are constructed to be robust to heteroskedasticity.³⁶ Columns (1)-(4) show the cross-section regressions. In each specification, the contemporaneous college share is estimated to have a large and highly significant effect on PC use. Not shown is that the results are nearly identical when lags of the college shares are used. Remarkably, the college share variable alone explains close to half of the variation in PC use across cities. Point estimates are around 0.5, i.e., a 10 percentage point increase in college

³⁴ Computer sales figures were obtained from the Computer and Business Equipment Manufacturers Association "All County Sales Report" for 1980, and were divided by employment estimated using the 5 percent public-use Census of Population from 1980.

³⁵ This computer sales variable is also significantly correlated with college share in 1980, and with PCs per worker in 1990. Correlations in both cases are roughly 0.4.

³⁶ The first difference-regressions are weighted by the efficient first-difference weight $(n_{1990}^{-1} + n_{2000}^{-1})^{-0.5}$ where n_t is the sample size in year t . The results in Table 4 and throughout the paper are not sensitive to weights used.

share (the difference between Dallas and San Francisco in 2000) is associated with 0.05 more PCs per worker.³⁷ IV estimates, which use the college instruments described above, are slightly smaller than OLS estimates, consistent with college share being partly endogenous. The estimates are larger in 2000 than in 1990, which is not surprising given that the variance of the dependent variable increased between these years.

Columns (5) and (6) show estimates from a regression of the 1990-2000 change in PC intensity on the 1990-2000 change in college share. OLS estimates, in column (5), are remarkably similar to the cross-section estimates. The IV estimate, which uses the immigrant enclave instrument and is shown in column (6), is considerably larger, though the standard error is large enough that the IV estimate is not statistically distinguishable from the OLS point estimate. Nevertheless, there are reasons why IV estimates may be larger than OLS in this case. One potential reason is that OLS estimates are attenuated by classical measurement error in college share, a problem that would be larger in first differences than in cross section. Another possibility is that the IV estimates identify a local average treatment effect that is larger than the average effect of college-share (Imbens and Angrist, 1994). Compared to U.S. natives, immigrants tend to have disproportionately low education and low English ability, so it is plausible that immigrant presence may have more of a depressing effect on PC use than the effect of native-born non-college workers. If so, immigration-driven variation in college share would have a larger effect on PC intensity than college share generally. (In fact, we have found that even OLS estimates are larger in the subsample of cities with a significant amount of immigration.)

The models in Table 4 are admittedly simple and are used to convey how our proposed instruments perform. The next two sections explore the robustness of our results when the meaning of skill is broadened and when other models of technology adoption are considered.

IV. Skill

Our measure of human capital, the college share, is crude and, as such, may be correlated with other characteristics of the population that are important in the PC adoption decision but were excluded from the regressions in Table 4. To ensure that we are capturing the effects of

³⁷ This is consistent, for example, with an individual-level gap in PC use between college equivalent and other workers of 50 percentage points. We know of no individual-level data on PC use at work, though the difference in computer use between college equivalent and other workers in the CPS computer-use supplement is smaller than this.

more specific skills, we examine the influence of detailed occupational categories such as engineers and computer support workers. We also examine the age distribution of the workforce in each city to test whether younger, college-educated workers may port computer skills acquired in college to the workplace. Finally, another variable that may capture the unobserved component of skill of the local workforce is the rankings of local computer science and electrical engineering departments; areas that have highly ranked departments may have workers with above-average computer skills.³⁸

The results are divided into the 2000 cross section (Table 5a) and the first difference over 1990-2000 (Table 5b).³⁹ For comparison, column (1) of each table repeats the OLS estimate of the bivariate relationship from Table 4. Both tables show OLS estimates because we do not have instruments for the various skill categories we examine, and, as shown in Table 4, the OLS results closely match the IV results.

The second column of each table disaggregates college graduate share into the share of college graduates 30 years and older and the share of younger college graduates. Goldfarb (2005) found that exposure to information technology during college years in the latter half of the 1990s led young college graduates to be early adopters of the Internet at home. In this case, however, we find that the share of older college graduates drives the college share-PC relationship. This result is consistent with individual-level data showing that on-the-job computer use rises with age (Card and DiNardo, 2002). It may reinforce the idea that computers complement general human capital, which rises with work experience.⁴⁰

Column (3) in Tables 5a and 5b includes controls for the presence and rank of a Ph.D. program in computer science and engineering in the area in 1981, as reported by Jones et al. (1982). The presence of a Ph.D. program is associated with roughly 0.02 more PCs per worker, but the quality of the program seems to have little impact.⁴¹ The point estimates on college share are only slightly diminished by these controls. So while computer science graduates may have an outsized effect on PC use, they do not seem to be driving the relationship between college share and PC-use.⁴² This fact should assuage some concerns that that the presence of colleges is

³⁸ We also examined wages of college graduates, high school graduates, and the ratio of the two. Wages may reflect differences in skill across cities. However, wages may also reflect differences in relative supply. The inclusion of wages does not affect our results, and a model that examines wages, changes in wages, and technology adoption is more fully explored in our companion paper with Paul Beaudry.

³⁹ Not shown are cross-sectional results for 1990; those results are remarkably similar to those for 2000.

⁴⁰ It might also mean that older workers have more specific computer skills.

⁴¹ We also examined many other specifications of the Ph.D. program rankings and found similar results.

⁴² Goolsbee and Klenow (2002) report that spillovers from experienced and intensive computer users are greater than from more average computer users, and that is one of the motivations for including this variable.

invalid as an instrument for college share because universities themselves have a direct effect on technology adoption through, e.g., spin-offs from university labs (e.g., Digital Equipment Corporation in Boston). While these estimates do not rule out that universities have a direct effect on an area's technology use, they suggest such spillovers are likely to be small in comparison to the effect universities have on the local supply of general human capital. The next section discusses spillovers further.⁴³

College share may also proxy for other job characteristics or worker skills associated with computer use. DiNardo and Pischke (1997), for example, argue that skilled workers are more likely to use computers simply because they are more likely to be in jobs (say, office jobs) that make use of computers; Autor, Levy, and Murnane (2003) also argue that computerization is a consequence of job attributes. The final columns of Tables 5a and 5b control for the occupation mix of an area; if college share is a proxy for occupation mix, these controls ought to soak up some of the effect of college share.⁴⁴ For the purpose of the regression, employment shares were constructed for nine broad occupational categories, listed in Tables 5a and 5b. The results in column (4) show that the inclusion of these occupational shares only slightly diminishes the coefficient on college share.⁴⁵ As for the occupational mix coefficients themselves, some align with expectations (for example, clerical share in the cross section and computer programmer share in first difference are positively associated with computer use), but most are not statistically distinguishable from the excluded category.⁴⁶ It is worth remembering, however, that the dependent variable already conditions on detailed industry. To be detected here, any effect of the supply of programmers, say, must operate through channels other than the compositional effects of industry mix (e.g., spillovers from the tech sector rather than the mere compositional effect of a larger tech sector).

In short, while we admittedly lack a method to identify the effect of occupation mix and education mix separately, the fact that the partial association between PC-use and occupation

⁴³ IV estimates using the university instruments are similarly unaffected by these Ph.D. program controls.

⁴⁴ On the other hand, occupation mix may be a consequence of skill mix, in which case regressions with occupation controls will understate the reduced-form effect of college share.

⁴⁵ Closely mimicking Autor, Levy, and Murnane (2003), we computed four measures of "skills" for each city based on the attributes of detailed occupational codes. Those measures are nonroutine cognitive/interactive tasks, non-routine manual tasks, routine manual tasks, and routine cognitive tasks. The results using these four skill measures as controls are similar to those reported in the last columns of Table 5a and Table 5b; the coefficient on the college share variable falls by nearly an identical amount and remains statistically significant. When both the nine occupational categories and the four skill measures are included simultaneously, the coefficient and the standard error on our college share variable remain nearly identical.

⁴⁶ This result corresponds to the results in Michaels (2006) that focus on the relationship between clerical workers and information technology.

mix is weak compared to the partial association with college share itself suggests college share may directly affect computerization, rather than only operating through occupation mix. The same jobs may be more likely to be carried out with a PC where there are more college-educated workers.^{47,48}

V. Model Misspecification

Many factors enter into the PC adoption decision of firms other than the availability of certain types of labor, as was emphasized in the previous section. For information technology, spillovers (as in Goolsbee and Klenow, 2002) and vintage effects (as in Bresnahan and Greenstein, 1996) are potentially important in the adoption process. In this section we explore how sensitive our results are to the inclusion of controls for spillovers and vintage effects.

One source of possible spillovers is from the IT-producing sector; employees in this sector may be able to transfer their knowledge about IT to firms in their area through social networks or through changing jobs.⁴⁹ Just how large spillovers from the IT-producing sector might be depends on the relative size of the sector. For each city in our sample, we construct the share of employment in the IT-manufacturing sector (the sector that makes computers, communications gear, and semiconductors) and the software sector. These controls are potentially interpretable as “spillovers” because the direct effect of industry is already controlled for in our measure of PC intensity, and the IT sectors are removed from the city-level average PC-intensity. Thus any observed effect of IT share is indirect, operating through the IT sector’s effect on other sectors.

Table 6 reports the results from including the IT sector variables. Columns (1)-(5) show results for the 2000 cross section and columns (6) – (11) show results for the 1990-2000 difference, with columns (1) and (6) repeating the OLS estimates from Table 4. Columns (2)-(3) and (7)-(9) show the effect of the size of a city’s IT sectors. The results in Table 6 show that the size of a city’s software sector in 1980 is somewhat positively associated with greater PC-intensity in 2000. The nearly identical results in columns (2) (based on all cities) and (3) (all

⁴⁷ DiNardo and Pischke’s (1997) argument against interpreting the computer-education correlation as evidence of a causal effect of computer use on skill demand came from the fact that other job tools that do not require special skills, like pencils, also have a strong association with education. While we lack data on pencils, in future versions of this paper we may subject data we have on phone lines/worker to the same analysis.

⁴⁸ One better way to address this issue would be to use the CPS computer use supplements. However, we conducted much of the analysis presented in this paper and in the companion paper based on CPS data. Unfortunately, the CPS data become quite thin at the city level beyond a handful of large cities.

⁴⁹ This argument is similar to the one in the previous section regarding more refined measures of human capital. We place the IT sector variables in this section primarily because spillovers may be important in and of themselves.

cities except San Francisco) suggest that the relationship is not driven by the outlier San Francisco Bay area, home to Silicon Valley. For the first-difference results, columns (7)-(9), we also include the level of the IT sector in addition to the change in the IT sector to pick up the possibility that the level of the IT sector may be related to the change in technology. Although all of the coefficients on the IT sectors are positive, most have large standard errors, with the occasional exception of a software variable.

It seems an abundance of software workers may have spillovers to other sectors. As mentioned above, a plausible mechanism for this spillover is the movement of former software workers to other sectors. An examination of worker flow data in the CPS shows that a large fraction of software workers move to a different sector within a year; it is a high turnover sector. IT-manufacturing workers, in contrast, we found to be much less likely to move, which is perhaps consistent with the lack of impact from that sector.⁵⁰

In any case, point estimates on college share are reduced somewhat by these controls, but not by a large amount. In results not shown, we also controlled for the size of the IT communications sector and the size of the finance sector (which is an intensive user of IT equipment, and hence, we imagined, might also have spillover effects). However, we never found any significant association between the size of these sectors and the PC-intensity of an area.

The other modeling issue that concerns us is vintage effects, namely, how the technology in place before PCs may affect the adoption of PCs. One argument could be that cities that had large computer investments before PCs may be slower to adopt PCs because of the fixed costs associated with switching to another technology. Alternatively, cities with high levels of pre-PC technology might be better prepared to incorporate PCs into their networks.⁵¹ We approach the problem of vintage effects in three ways. First, we examine relatively long first differences (from 1990-2000) instead of emphasizing short-run changes (say year-to-year) where vintage effects are no doubt more important. Second, as described in Section II, we estimate our city effects controlling for the major computing systems in place—we found our results were largely insensitive to this control.

Finally, we control for the computer spending from 1978 to 1980 (before the introduction of the IBM PC) using survey data collected by the Computer and Business Equipment

⁵⁰ These results are based on very small samples and hence we do not report the exact estimates.

⁵¹ See Bresnahan and Greenstein (1996) for a fuller description of the issues firms faced when switching from a mainframe mode of computing to the client-server approach (the approach that utilized PCs).

Manufacturers Association.⁵² Since this information is currently available only for a subset of our metropolitan areas, for comparison columns (5) and (10) show the relationship without controls in this subsample. Although investment in these early vintage computers is positively associated with PC-intensity later on (correlation with 2000 PC-intensity = 0.3), it has virtually no partial relationship with PC-intensity conditional on college share.

All of the controls in Tables 5-6 follow a similar pattern: they have at most a small effect on PC-intensity and only slightly diminish the relationship between college share and PC-intensity. But what happens when all of the controls are entered into the regression simultaneously? Does that diminish the college share effect even more? The answer is “no.” If one enters college share, occupation shares, rankings of computer science departments, and IT share into the same regression, the slope estimate (standard error) for college share is 0.608 (0.107) in the 2000 cross section and is 0.613 (0.168) in the 1990-2000 difference.⁵³ These estimates are larger than many of the estimates in Tables 5 and 6. The effect of college share thus is quite robust.

VI. Conclusion

There has been much research on the interaction of labor skills and technology, especially between education and computers. Properly identifying how the supply of skills affects technology adoption, and how technology adoption affects the demand for skills, has proven challenging. This paper has attempted to take a step toward identifying a portion of the labor-technology system by carefully approaching how the supply of skills in cities affects the decision of businesses to adopt personal computers. To undertake this step, several issues had to be addressed, namely deriving instruments for the college share across cities and instruments for the changes in college share. Further, we explored other definitions of “skill” and also explored how other factors may influence our results.

We find a tremendously strong relationship between the supply of skills and the adoption of PCs. This relationship is robust to a large number of other controls and is robust to several sets of instruments. These results are consistent with previous research finding an association

⁵² The data consist of nominal spending for 1978, 1979, and 1980. Although data before 1978 are not available, they are unlikely to affect our results noticeably because of the rapid price declines and high depreciation rates of IT equipment.

⁵³ In the subsample of cities in which we have 1980 computer investment data, the slope estimate (standard error) for college share controlling also for 1980 computer investment per worker is 0.601 (0.111) in the 2000 cross section regression and is 0.706 (0.173) in the 1990-2000 difference regression.

between computers and skilled labor across plants and industries but make a stronger case for the complementarity between the two by taking advantage of differences in factor mixes which are more plausibly exogenous. What these results do not shed much light on is the model that underlies the complementarity between skills and technology. Beaudry, Doms, and Lewis (2006) attempts to test between various endogenous technological adoption models, specifically models that build upon Beaudry and Green (2005), and other models that have been posited as underlying the skill-biased technological change hypothesis.

References

- Acemoglu, Daron (1998). "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Quarterly Journal of Economics* 113(4): November 1998, p. 1055-1089.
- Altonji, Joseph, and David Card (1991). "Effects of Immigration on Labor-Market Outcomes of Less-Skilled Natives." In John Abowd and Richard Freeman, eds., *Immigration, Trade, and the Labor Market*, National Bureau of Economic Research, 1991. p. 201-234.
- Autor, David H., Lawrence F. Katz and Alan B. Krueger (1998). "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113(4): November 1998, p. 1169-1213.
- Autor, David H., Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): November 2003, p. 1279-1334.
- Bartel, Ann (1989). "Where Do the New US Immigrants Live?" *Journal of Labor Economics* 7(4): October 1989, p 371-91.
- Beaudry, Paul, and David A. Green (2005). "Changes in U.S. Wages, 1976-2000: Ongoing Skill Bias or Major Technological Change?" *Journal of Labor Economics* 23(3): p. 491-526.
- Beaudry, Paul, Mark Doms, and Ethan Lewis (2006). "The IT Revolution at the City Level: Testing a Model of Endogenous Biased Technology Adoption," mimeo.
- Bresnahan, Timothy, and Shane Greenstein (1996), "Technical Progress and Co-Invention in Computing and in the Use of Computers." *Brookings Papers on Economic Activity: Microeconomics*, pp. 1-78.
- Brynjolfsson, Erik, and Hitt, Lorin (2003), "Computing Productivity: Firm-level Evidence," *Review of Economics and Statistics* 85(4): 2003, p. 793-808.
- Card, David (1995). "Using Geographic Variation in College Proximity to Estimate the Return to Schooling," in Louis N. Christofides, E. Kenneth Grant, and Robert Swidinsky, Eds., *Aspects of Labour Market Behaviour: Essays in Honour of John Vanderkamp*. Toronto: University of Toronto Press, p. 201-222.
- Card, David (1999). "The Causal Effect of Education on Earnings." *Handbook of Labor Economics*. Volume 3A, 1999, p 1801-1863. Amsterdam, New York, and Oxford: Elsevier Science, North-Holland.
- Card, David (2001). "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration." *Journal of Labor Economics* 19(1): January 2001, p. 22-64.

Card, David, and John E. DiNardo (2002). "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20 (4): October 2002, pp. 733-83.

Caselli, Francesco, and John Coleman (2001). "Cross-Country Technology Diffusion: The Case of Computers," *American Economic Review: Papers & Proceedings*, May 2001.

Comin, Diego, and Bart Hobijn (2004), "Cross-Country Technology Adoption: Making the Theories Face the Facts." *Journal of Monetary Economics* 51: 2004, p. 39-83.

DiNardo, John, and Jörn-Steffen Pischke (1997). "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *Quarterly Journal of Economics* 112, Feb. 1997, 291-303.

Forman, Chris, Avi Goldfarb, and Shane Greenstein (2005). "How Did Location Affect Adoption of the Commercial Internet: Global Village vs. Urban Leadership." *Journal of Urban Economics*, 58(3), 389-420.

Goldfarb, Avi (2005) "The (Teaching) Role of Universities in the Diffusion of the Internet." working paper, University of Toronto.

Goolsbee, Austan, and Peter Klenow (2002). "Evidence on Learning and Network Externalities in the Diffusion of Home Computers." *Journal of Law & Economics* 45: October 2002, 317-344.

Griliches, Zvi (1957), "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica* 25 (October): 501-522.

Imbens, G.W. and Angrist, J. (1994), "Identification and Estimation of Local Average Treatment Effects", *Econometrica*, 62, 467-475.

Jones, Lyle, Gardner Lindzey, and Porter Coggeshall, editors (1982). *An Assessment of Research-Doctorate Programs in the United States: Mathematical and Physical Sciences* National Academy Press, Washington D.C., 1982.

Kane, Thomas J., and Cecilia Elena Rouse (1995). "Labor Market Returns to Two- and Four-Year College." *American Economic Review* 85(3): June 1985, p. 600-614.

Kane, Thomas J., and Cecilia Elena Rouse (1999). "The Community College: Educating Students at the Margin Between College and Work." *Journal of Economic Perspectives* 13(1): Winter 1999, p. 63-84.

Katz, Lawrence F., and Kevin M. Murphy (1992). "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 112: 1992, p 291-303.

Lewis, Ethan (2004). "How Do Local Labor Markets in the U.S. Adapt to Immigration?" Mimeo, Federal Reserve Bank of Philadelphia, November 2004.

Lewis, Ethan (2005). "Immigration, Skill Mix, and the Choice of Technique." Federal Reserve Bank of Philadelphia Working Paper, May 2005.

Michaels, Guy (2006). "Technology, Complexity and Information: The Evolution of Demand for Office Workers," mimeo.

Moretti, Enrico (2004). "Estimating the Social Return to Higher Education: Evidence From Longitudinal and Repeated Cross-Sectional Data," *Journal of Econometrics*, 121(1-2): July-August 2004, p. 175-212.

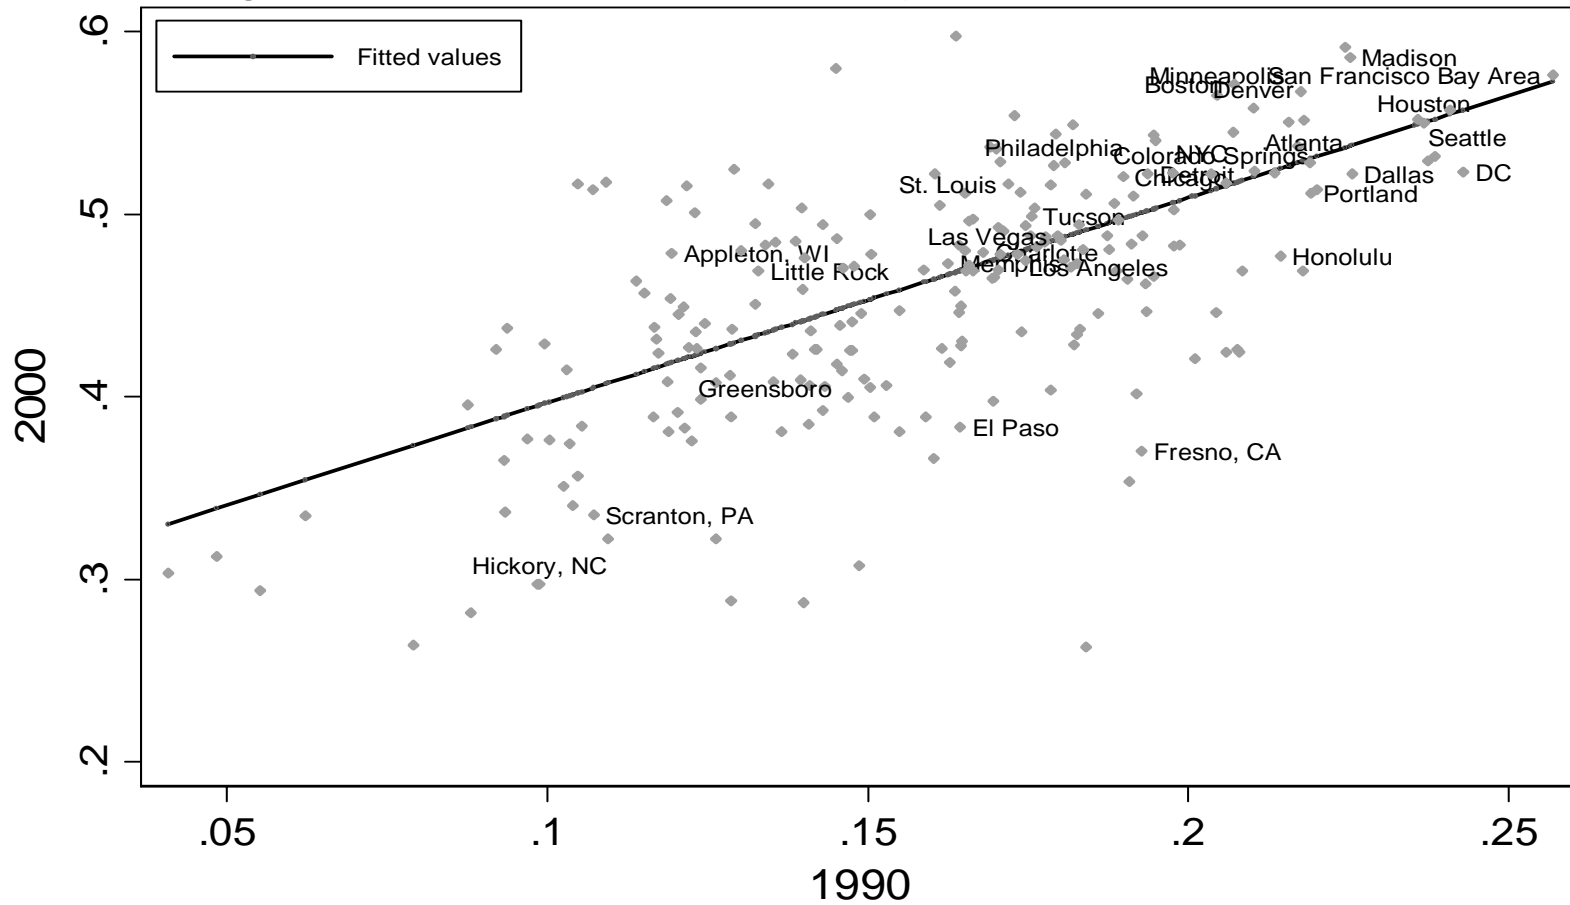
Nervis, Allan (1962). *The State Universities and Democracy*. Urbana, IL: University of Illinois Press, 1962.

Nestoriak, Nicole (2004). "Labor Market Skill, Firms, and Workers." Ph.D. dissertation, University of Maryland.

Skinner, Jonathan, and Douglas Staiger (2005). "Technology Adoption From Hybrid Corn to Beta Blockers." NBER Working Paper 11251, National Bureau of Economic Research, Inc.

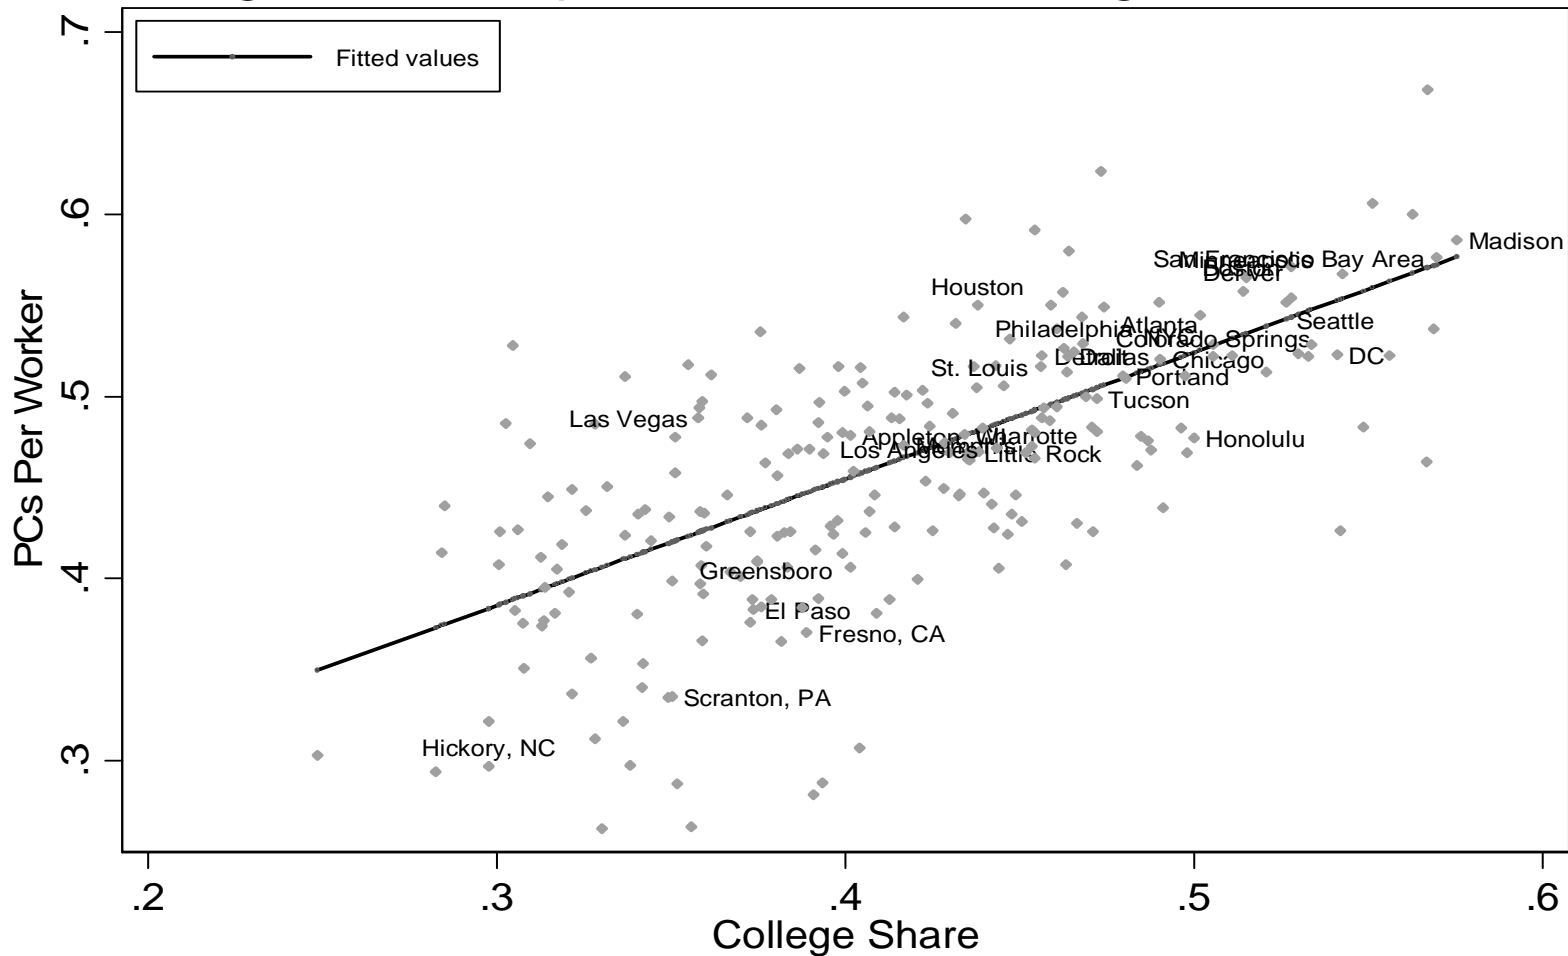
U.S. Department of Commerce, Bureau of the Census (1984). "Intercensal Estimates of the Population of Counties by Age, Sex, and Race [United States]: 1970-1980 [Computer file]." ICPSR version. Washington, DC: U.S. Dept. of Commerce, Bureau of the Census [producer], 1984. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2000.

Figure 1: PCs per Worker by CMSA, 1990 and 2000



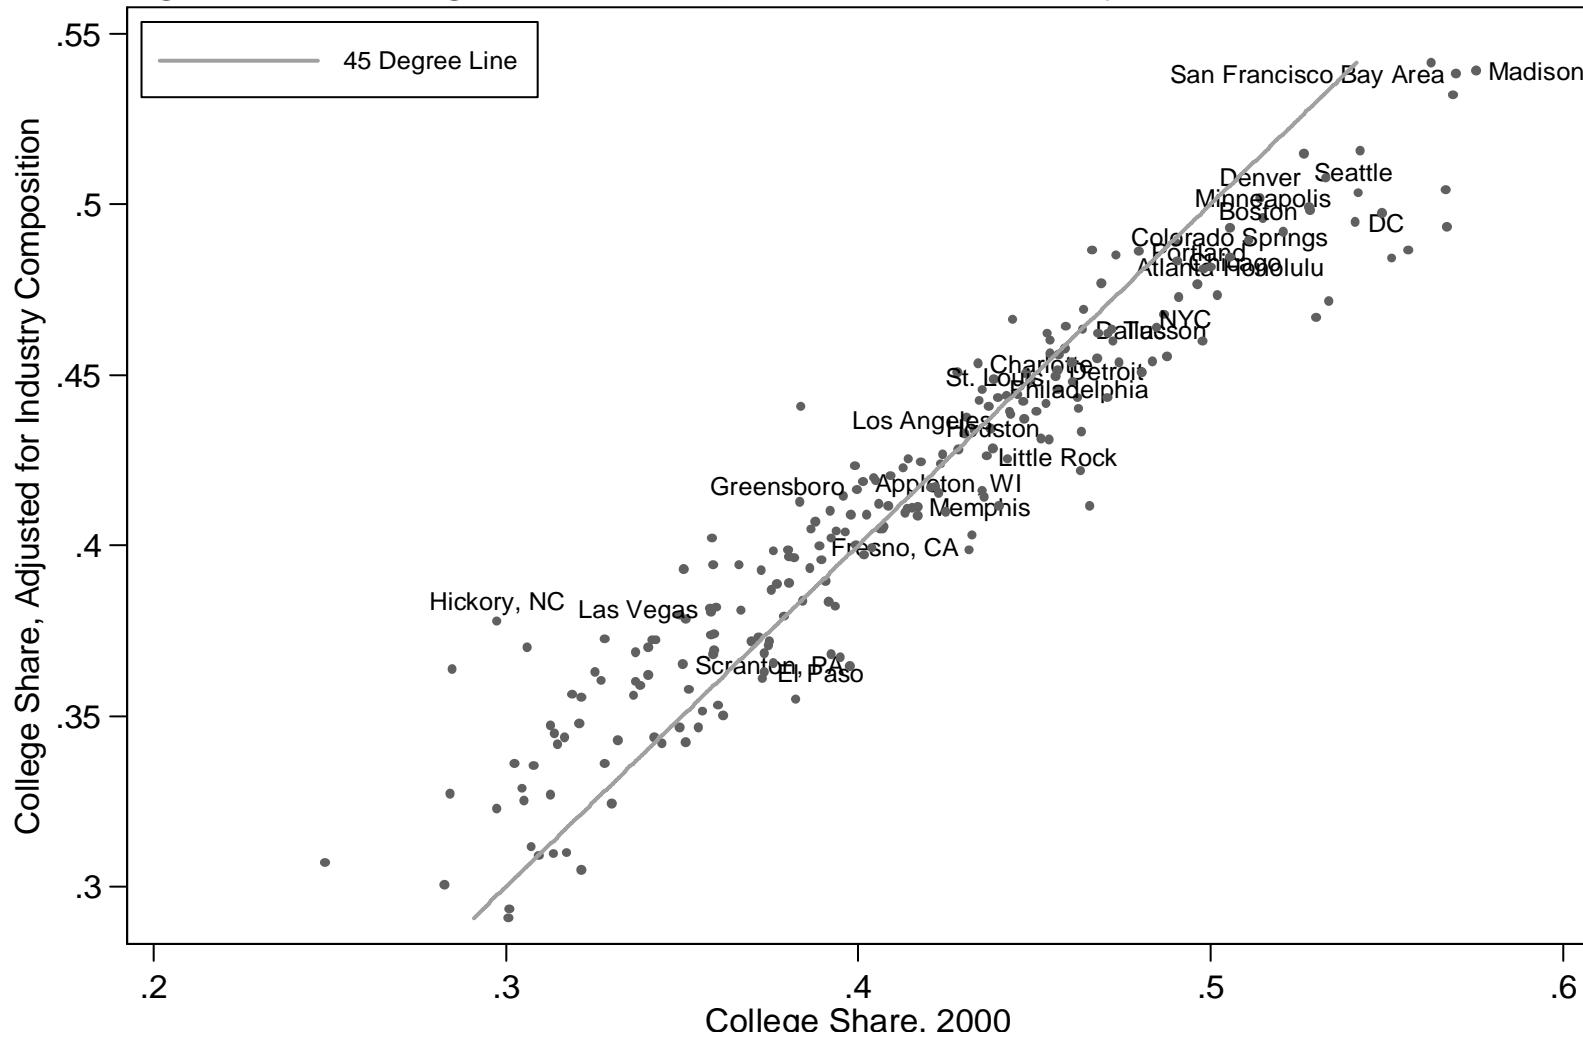
Notes: Mean number of PCs/employee in a CMSA after adjusting for 3-digit industry by plant size controls and year effects (see text). Data used come from Harte Hanks.

Figure 2: PCs per Worker and College Share, 2000




Notes: Mean number of PCs/employee in a CMSA after adjusting for 3-digit industry by plant size controls and year effects (see text). College share is derived from 5 percent public use files from the 2000 Census of Population. The sample includes employed people aged 16-65 with at least one year of potential work experience and not living in group quarters. College share is computed as 1/2 of the share of workers in the area with 1-3 years of college education plus the share with at least a four-year college degree.

Figure 3: College Share Adjusted for Industry Composition in 2000



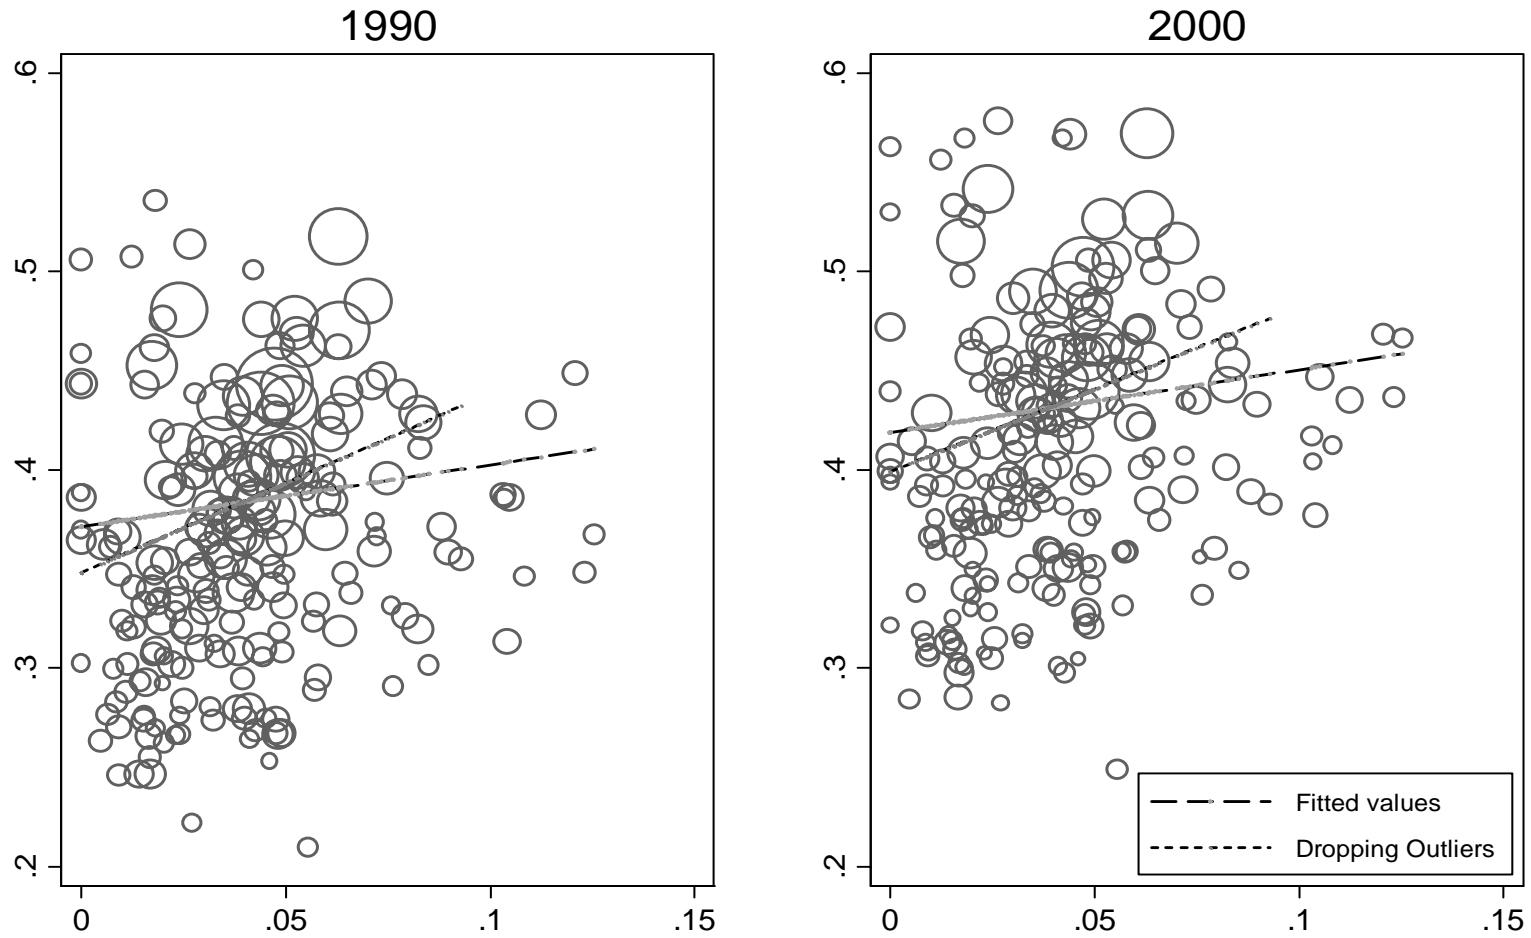
Notes: 2000 Census of Population 5% public-use files. Sample of 16-65 year olds with at least one year of potential work experience and not living in group quarters.

Figure 4. Kernel Density of College Share Across CMSAs



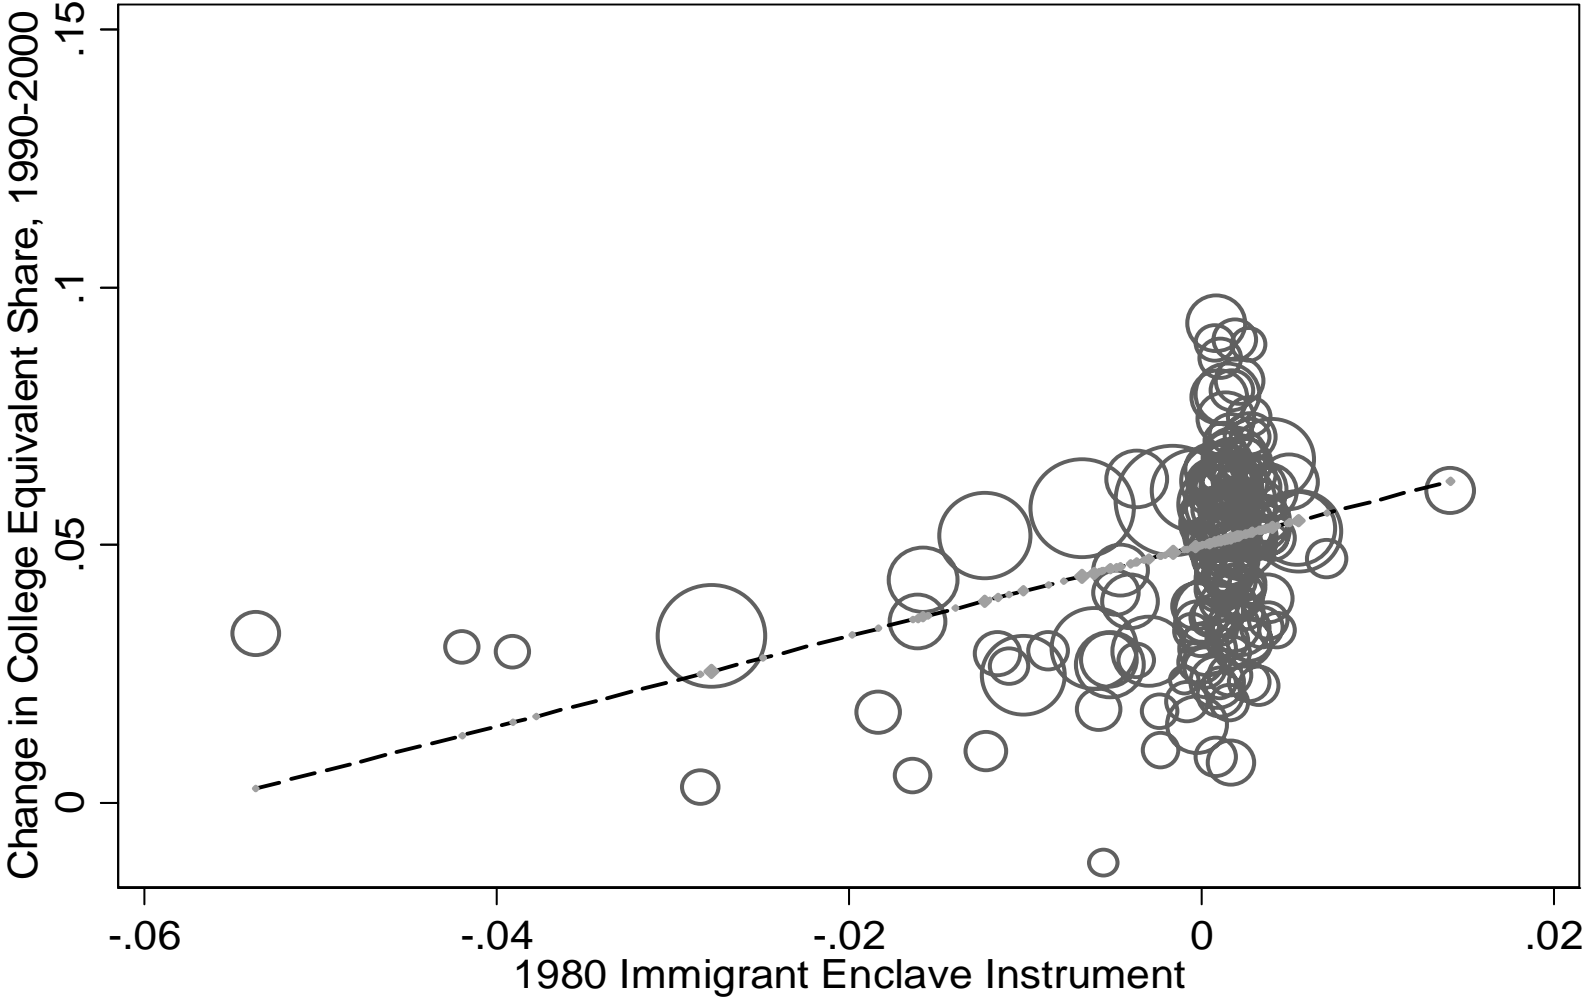
Notes: 2000 Census of Population 5% public use files. Sample of 16-65 year olds with at least one year of potential work experience and not living in group quarters. LGC=land grant college.

Figure 5: College Share vs. 1971 College Seats/Capita



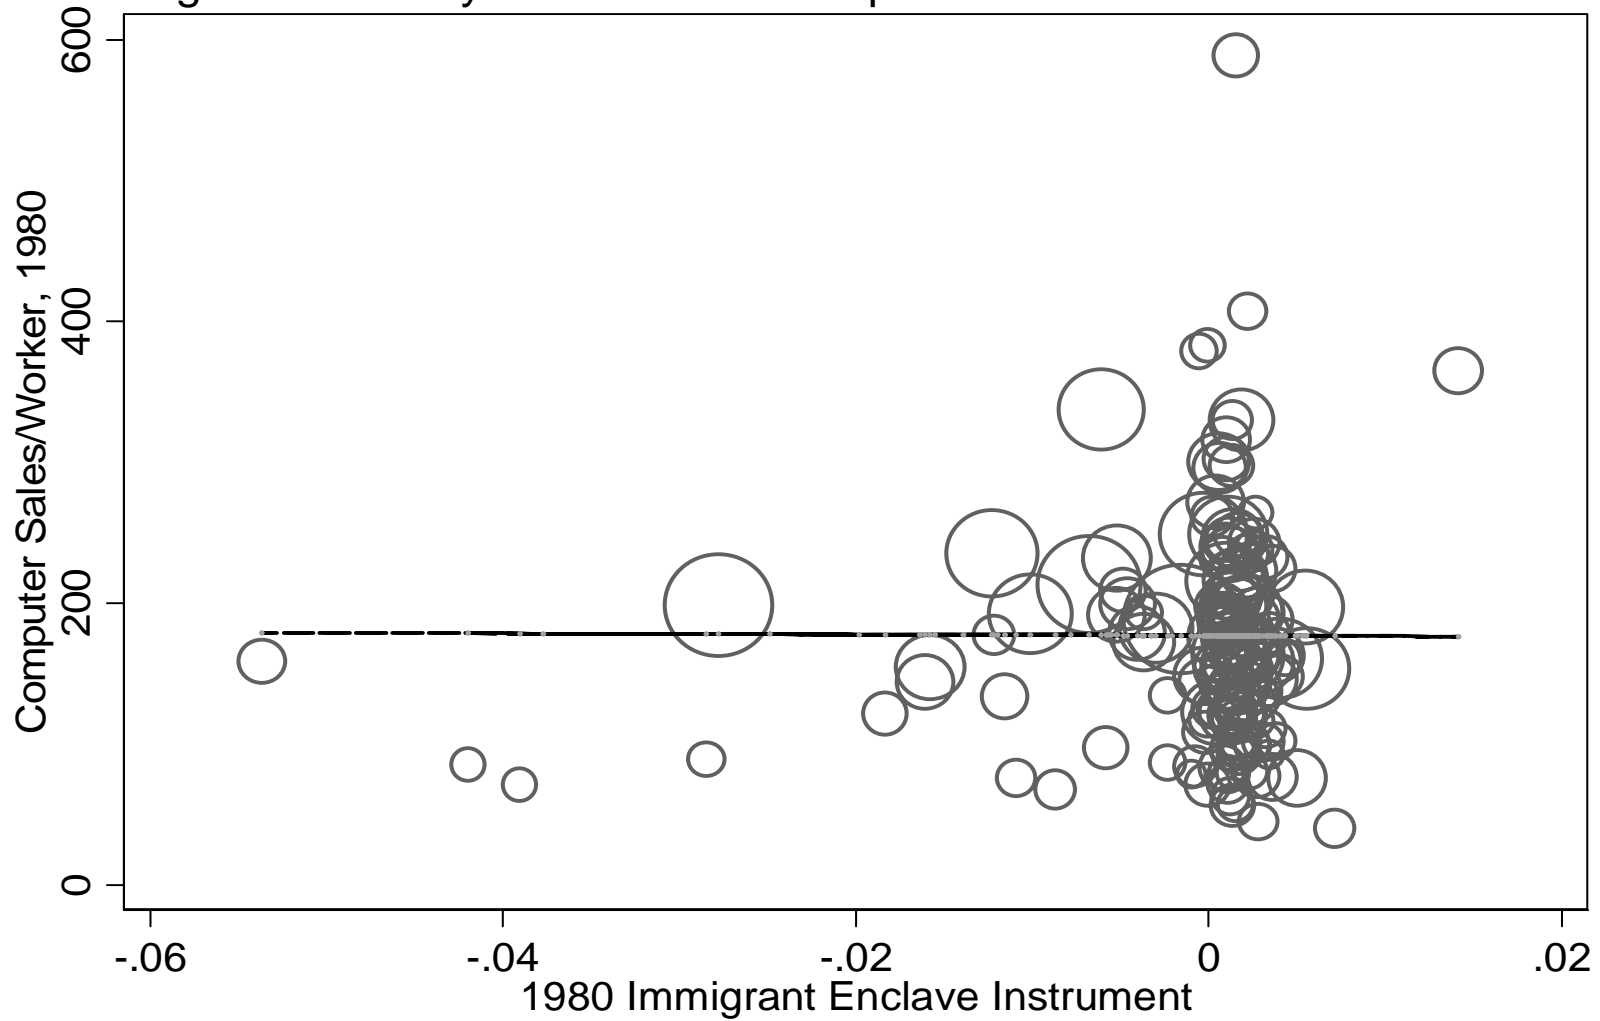
Notes: See Table 3.

Figure 6: Change in College Share vs. Instrument



Notes: See Table 3.

Figure 7. Validity Check: 1980 Computer Sales/Worker vs. Instrument



Data Sources: 1980 immigrant enclave instrument: see Table 3. 1980 Computer Sales/Worker: See Table 6.

Table 1: Information Technology Investment, 1990-2000 & 2000-2003

	Average annual percent change, 1990-2000:			Nominal spending (\$ billions)		
	Nominal investment	Prices	Real investment	1990	2000	1990 through 2000
IT	11.8	-7.0	20.2	131.5	401.7	2,371.7
Computers and peripherals	10.1	-17.7	33.9	38.6	101.4	725.6
Personal Computers	12.3	-25.1	50.0	10.6	34.0	267.3
Other computers	10.5	-20.2	38.4	9.4	25.5	141.6

	Average annual percent change, 2000-2003:			Nominal spending (\$ billions)		
	Nominal investment	Prices	Real investment	2000	2003	2000 through 2003
IT	-4.4	-6.5	0.8	401.7	350.8	1,454.3
Computers and peripherals	-2.1	-14.7	14.8	101.4	95.3	363.5
Personal Computers	0.3	-21.5	27.8	34.0	34.3	123.8
Other computers	-3.2	-21.8	23.7	25.5	23.1	86.9

Source: Kindly provided by BEA.

Table 2: CMSAs with Land-Grant Colleges

Albany-Schenectady-Troy, NY	Knoxville, TN	Providence, RI
Baton Rouge, LA	Lansing-East Lansing, MI	Raleigh-Durham, NC
Boston, MA	Lexington-Fayette, KY	Reno, NV
Columbus, OH	Lincoln, NE	Richmond-Petersburg, VA
Fayetteville-Springdale, AR	Los Angeles-Long Beach, CA	Sacramento, CA
Fort Collins-Loveland, CO	Macon-Warner Robins, GA	San Diego, CA
Gainesville, FL	Madison, WI	San Francisco Bay Area, CA
Greensboro--Winston-Salem--High Point, NC	Minneapolis-St. Paul, MN-WI	Springfield, MA
Greenville-Spartanburg, SC	Nashville, TN	Tallahassee, FL
Hartford, CT	New York, NY	Tucson, AZ
Houston, TX	Oklahoma City, OK	Washington, DC-MD-VA
Huntsville, AL	Philadelphia, PA-NJ	

Table 3: College Share and Instruments

	2000			1990		Change: 1990-2000	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
College seats per capita, 1971	0.316 (0.071)**		0.395 (0.069)**	0.390 (0.077)**	0.385 (0.075)**		0.017 (0.027)
Area has land-grant college=1, =0 otherwise		0.065 (0.011)**	0.070 (0.011)**	0.066 (0.010)**	0.066 (0.010)**		0.007 (0.003)*
1980 immigrant enclave instrument					-0.099 (0.498)	0.878 (0.157)**	0.943 (0.168)**
Constant	0.419 (0.007)**	0.416 (0.005)**	0.395 (0.007)**	0.348 (0.006)**	0.348 (0.006)**	0.050 (0.001)**	0.047 (0.002)**
Observations	230	230	230	230	230	230	230
R-squared	0.05	0.20	0.28	0.27	0.27	0.14	0.17
F-stat, instruments	19.86	36.16	32.79	31.91	21.40	31.19	10.71

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Notes: All regressions are run across 230 consolidated metropolitan areas defined according to their boundaries in 1999. Regressions in columns 1-3 are weighted by the square root of n_{2000} , the number of plants in the Harte-Hanks sample in 2000 and 2002. Regressions in columns 4-5 are weighted by the square root of n_{1990} , the number of plants in the Harte-Hanks sample in 1990 and 1992. Regressions in columns 6-7 are weighted by $(1/n_{1990} + 1/n_{2000})^{-0.5}$

Data sources. College share and 1980 immigrant enclaves: 5 percent public use files from the 1980, 1990, and 2000 census of population. The 1990 and 2000 sample include employed people aged 16-65 with at least one year of potential work experience and not living in group quarters. The 1980 sample consists of those at least age 16 not living in group quarters. College share represents 1/2 of the share of workers in the area with 1-3 years of college education plus the share with at least a four-year college degree. The immigrant enclave instrument is described in detail in the text. College seats per capita: 1971 Higher Education General Information Survey (HEGIS) for fall enrollment, and U.S. Dept. of Commerce, Bureau of the Census (1984) for population aged 15-64. It is the sum of four-year enrollment plus 1/2 of two-year enrollment divided by the population aged 15-64. Land-grant colleges: Nervis (1962).

Table 4: PC-Intensity and College Share

	2000		1990		Change 1990-2000	
	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)
College Share ^a	0.694 (0.044)**	0.595 (0.090)**	0.451 (0.029)**	0.432 (0.064)**	0.638 (0.150)**	1.213 (0.404)**
Constant	0.177 (0.019)**	0.220 (0.039)**	-0.001 (0.011)	0.006 (0.025)	0.275 (0.007)**	0.247 (0.020)**
Observations	230	230	230	230	230	230
R-squared	0.52	0.51	0.53	0.53	0.07	0.01

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Notes: All regressions are run across 230 consolidated metropolitan areas defined according to their boundaries in 1999. Regressions in columns (1)-(2) are weighted by the square root of n_{2000} , the number of plants in the Harte-Hanks sample in 2000 and 2002. Regressions in columns (3)-(4) are weighted by the square root of n_{1990} , the number of plants in the Harte-Hanks sample in 1990 and 1992. Regressions in columns (5)-(6) are weighted by $(1/n_{1990} + 1/n_{2000})^{-0.5}$. Data sources. PC-intensity: Harte-Hanks 1990-1992 and 2000-2002. Sample restricted to nonretail establishments with fewer than 1.2 PCs/employee. Dependent variable in columns (1)-(4) is the mean number of PCs/employee in a CMSA after adjusting for three-digit industry by plant size controls and year effects (see text). Dependent variable in columns (5)-(6) is the change in this adjusted PC-intensity variable between 1990 and 2000. College share: 5 percent public use files from the 1990 and 2000 census of population. The sample includes employed people aged 16-65 with at least one year of potential work experience and not living in group quarters. College share represents 1/2 of the share of workers in the area with 1-3 years of college education plus the share with at least a four-year college degree.

^a In columns (1)-(2), this is the college share in 2000; in columns (3)-(4) this is the change in college share in 1990; and in columns (5)-(6) this is the change in college share between 1990 and 2000.

Table 5a: PC-Intensity in 2000 and Additional Measures of Skill

	OLS (1)	OLS (2)	OLS (3)	OLS (4)
College share	0.694 (0.044)**		0.610 (0.053)**	0.621 (0.105)**
College share, age>30		0.795 (0.110)**		
College share, age<30		0.332 (0.273)		
Has ranked computer science Ph.D. program in 1981			0.018 (0.008)*	
In top-10 computer science departments, 1981			0.010 (0.010)	
In 11-20 computer science departments, 1981			0.009 (0.010)	
<u>Occupation category shares, 1990 (excluded group is teachers, professors, social science, artists, and performers)</u>				
Executives				0.410 (0.332)
Engineers				1.566 (0.521)**
Health care workers				1.221 (0.615)*
Technical				-0.790 (0.741)
Sales				0.741 (0.325)*
Clerical				0.912 (0.326)**
Low skill				0.502 (0.255)*
Computer programmers				0.826 (1.109)
Constant	0.177 (0.019)**	0.287 (0.015)**	0.206 (0.023)**	-0.338 (0.244)
Observations	230	230	230	230
R-squared	0.52	0.51	0.55	0.56

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Notes: All regressions are run across 230 consolidated metropolitan areas, defined according to their boundaries in 1999, and are weighted by the square root of n_{2000} , the number of plants in the Harte-Hanks sample in 2000 and 2002. Data sources. PC intensity: Harte-Hanks 2000-2002. Sample restricted to nonretail establishments with fewer than 1.2 PCs/employee.

Dependent variable is the mean number of PCs/employee in a CMSA after adjusting for 3-digit industry by plant size controls and year effects (see text). College shares and occupation shares: 5 percent public use files from the 1990 and 2000 census of population. The sample includes employed people aged 16-65 with at least one year of potential work experience and not living in group quarters. Ranking of Ph.D. Programs: Jones et al. (1982).

Table 5b: Change in PC-Intensity (1990-2000) and Additional Measures of Skill

	OLS (1)	OLS (2)	OLS (3)	OLS (4)
Δ College share, 1990-2000	0.638 (0.150)**		0.561 (0.148)**	0.597 (0.163)**
Δ College share, age>30, 1990-2000		0.933 (0.154)**		
Δ College share, age≤30, 1990-2000		0.622 (0.415)		
Has ranked computer science Ph.D. program in 1981			0.017 (0.009)+	
In top-10 computer science departments, 1981			0.006 (0.012)	
In 11-20 computer science departments, 1981			0.011 (0.009)	
<u>Change in occupation category shares, 1980-1990 (excluded group is teachers, professors, social science, artists, and performers)</u>				
Executives				-0.592 (0.638)
Engineers				-1.899 (0.880)*
Health care workers				-0.157 (1.149)
Technical				-2.243 (0.793)**
Sales				-0.672 (0.672)
Clerical				-0.480 (0.574)
Low skill				-0.511 (0.504)
Computer programmers				3.917 (1.312)**
Constant	0.275 (0.007)**	0.270 (0.007)**	0.271 (0.008)**	0.265 (0.012)**
Observations	230	230	230	230
R-squared	0.07	0.14	0.12	0.16

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Notes: All regressions are run across 230 consolidated metropolitan areas, defined according to their boundaries in 1999, and are weighted by $(1/n_{1990} + 1/n_{2000})^{-0.5}$, where n_{2000} and n_{1990} are, respectively, the number of plants in the Harte-Hanks sample in 2000-2002 and 1990-1992. Data sources. PC intensity: Harte-Hanks 1990-1992 and 2000-2002. Sample restricted to nonretail establishments with fewer than 1.2 PCs/employee. Dependent variable is the change in the mean number of PCs/employee in a CMSA after adjusting for 3 digit industry by plant size controls and year effects (see text). Dependent variable in columns (5)-(6) is the change in this adjusted PC-intensity variable between 1990 and 2000. College shares and occupation shares: 5 percent public use files from the 1980, 1990, and 2000 census of population; only employed respondents aged 16-65 with at least one year of potential work experience and not living in group quarters were included. Ranking of Ph.D. Programs: Jones et al. (1982).

Table 6: PC-Intensity and Measures of the IT Sector

	2000					Difference, 1990-2000					
	OLS (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)	OLS (6)	OLS (7)	OLS (8)	OLS (9)	OLS (10)	OLS (11)
College share ^a	0.694 (0.044)**	0.576 (0.057)**	0.576 (0.057)**	0.651 (0.044)**	0.631 (0.047)**	0.638 (0.150)**	0.597 (0.146)**	0.572 (0.152)**	0.554 (0.153)**	0.700 (0.155)**	0.687 (0.154)**
Software industry share, 1980		4.915 (2.383)*	4.915 (2.393)*				3.512 (2.012)+	2.187 (2.326)	1.535 (2.382)		
IT manufacturing share, 1980		0.050 (0.274)	0.050 (0.275)				0.188 (0.218)	0.139 (0.371)	0.138 (0.384)		
Δ Software industry share, 1990-2000								0.872 (0.608)	1.576 (0.793)*		
ΔIT manufacturing share, 1990-2000								0.109 (0.619)	0.150 (0.645)		
Computer sales/worker, 1980					0.000 (0.000)						0.000 (0.000)+
Constant	0.177 (0.019)**	0.211 (0.021)**	0.211 (0.022)**	0.197 (0.020)**	0.198 (0.020)**	0.275 (0.007)**	0.263 (0.009)**	0.264 (0.010)**	0.263 (0.010)**	0.274 (0.008)**	0.264 (0.010)**
Observations	230	230	229	206	206	230	230	230	229	206	206
R-squared	0.52	0.54	0.53	0.50	0.50	0.07	0.11	0.12	0.12	0.09	0.10
Sample:	All cities	All cities	No SF	1980 computers		All cities	All cities	All cities	No SF	1980 computers	

Robust standard errors in parentheses

+ significant at 10%; * significant at 5%; ** significant at 1%

Notes: All regressions are run across consolidated metropolitan areas defined according to their boundaries in 1999. Regressions in columns (1)-(5) are weighted by the square root of n_{2000} , the number of plants in the Harte-Hanks sample in 2000 and 2002. Regressions in columns (6)-(11) are weighted by $(1/n_{1990} + 1/n_{2000})^{-0.5}$ where n_{1990} is the number of plants in the Harte-Hanks sample in 1990 and 1992. **Data sources.** PC intensity: Harte-Hanks 1990-1992 and 2000-2002. Sample restricted to nonretail establishments with fewer than 1.2 PCs/employee. Dependent variable in columns (1)-(5) is the mean number of PCs/employee in a CMSA after adjusting for 3-digit industry by plant size controls and year effects (see text). Dependent variable in columns (6)-(11) is the change in this adjusted PC-intensity variable between 1990 and 2000. College-equivalent share, software share, and IT manufacturing share are derived from the 5 percent public use files from the 1980, 1990, and 2000 Censuses of Population. The sample includes employed people aged 16-65 with at least one year of potential work experience and not living in group quarters. College share represents 1/2 of the share of workers in the area with 1-3 years of college education plus the share with at least a four-year college degree. Computer sales/worker: Computer and Business Equipment Manufacturers Association "All County Sales Report" for computer sales, and 5 percent public-use Census of Population from 1980 for employment (sample described above).

^a In columns (1)-(5) this is the college share in 2000; in columns (6)-(11) this is the change in college share between 1990 and 2000.

**Appendix Table 1. Arrivals of Employed Immigrants by
Country of Origin and Skill, 1990-2000**

Country/Region of origin	Number of employed arrivals (1)	Number of college equivalents (2)	College equivalents as a percent of employed arrivals (3)
Mexico	1,641,103	130,022	7.9
Russia, Eastern Europe	457,832	236,315	51.6
South America	390,037	140,513	36.0
Central America	388,026	44,335	11.4
Caribbean	316,166	77,430	24.5
China,Hong Kong, Taiwan	306,347	190,583	62.2
Southeast Asia	291,393	88,369	30.3
India	284,458	228,986	80.5
Canada/UK/Australia	239,174	165,885	69.4
Philippines	224,921	142,321	63.3
Africa	205,787	99,961	48.6
Korea, Japan	171,887	112,747	65.6
Northwest Europe,Israel	145,819	103,638	71.1
Middle East	135,620	74,431	54.9
Central Asia	98,044	56,016	57.1
Cuba	91,501	27,431	30.0
Southwest Europe	56,570	29,880	52.8

Data source. 5 percent public use files from the 2000 Census of Population. The sample includes employed people aged 16-65 with at least one year of potential work experience not living in group quarters who report place of birth outside the U.S. and who report being noncitizens or naturalized citizens. College equivalents include workers with at least a four year college degree and 1/2 of those with 1-3 years of college education (computed using Census sample weights)