The Internet and Local Wages: Convergence or Divergence?*

Chris Forman
College of Management
Georgia Institute of Technology
800 West Peachtree St. NW
Atlanta, GA 30308
chris.forman@mgt.gatech.edu

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St George St
Toronto, ON M5S 3E6
agoldfarb@rotman.utoronto.ca

Shane Greenstein
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston IL 60208
greenstein@kellogg.northwestern.edu

December 2008

Abstract

Did the diffusion of the internet lead to the convergence or divergence of local wages across the US? We offer empirical evidence on the relationship between business use of advanced internet technologies and regional variance in wage growth between 1995 and 2000. We show that business use of advanced internet technology is associated with wage growth. We find no evidence that the internet contributed to regional convergence of wages, however. Rather, business use of advanced internet technology is associated with wage growth in regions that were already well off in terms of income, education, population, and industry. We rule out any substantial role for the internet in wage growth in lower income, less urban, less educated areas with less IT-intensive industries. Overall, advanced internet technology explains one quarter of the difference in wage growth between the counties that were already well off and the others.

Keywords: wage growth, convergence, divergence, information technology, internet
JEL Classification: O33, R11

* We thank Mercedes Delgado and seminar participants at the University of Colorado, Washington University, Arizona State University, the NBER Productivity Lunch, and WISE 2008 for comments. We also thank Harte Hanks Market Intelligence for supplying data. All opinions and errors are ours alone.

1. Introduction

Widespread evidence indicates that investment in information technology (IT) in the 1990s produced gains in US productivity and economic growth at the national, industry, and firm levels. Equally substantial evidence raises questions about whether the benefits of IT investment have been experienced everywhere. In particular, new IT investments have had the greatest effects on productivity for industries that were already IT-intensive and for workers with more education and skills. Yet, those findings leave open some fundamental questions about variance in growth of incomes: Did those IT investments contribute to regional inequality in US wages? Specifically, did IT investments contribute to regional convergence or divergence in wages?

The question arises with special saliency in the 1990s because the new IT investments of that era – and particularly, the rise of the commercial Internet – facilitated long-distance communication. One view hailed the internet as a great enabler of economic growth, particularly for low-density regions. This perspective hypothesizes that increased communication between establishments would break the link between local investment, local productivity, and local wage growth, leading to convergence across regions. While this view has received support from publications such as Cairncross's (1997) *The Death of Distance* and Friedman's (2005) *The World is Flat* and has received some support at the international level from data on the globalization of services (OECD 2006; Arora et al. 2001), lack of regional data prevented systematic testing within the US.

A contrasting perspective cast the internet as a technology that exacerbates existing inequalities in wages between urban/rural and frontier/mainstream users of information technology (IT), leading to divergence across regions. This view argues that the Internet resembled prior generations of IT (Moss and Townsend 1997; Kolko 2002). Effective use of frontier IT relies on the presence of frontier skills in the labor market, and wage gains will be greatest for workers in skilled occupations that are more likely to be found in rich urban areas (Kolko 2002). This view implies that the internet might also lead to divergence. In line with this view, Wellman (2001) argues that the internet primarily benefits local communication because social networks are local. Glaeser and Ponzetto (2007) argue that low communication costs help rich, idea-producing areas more than poor, goods-producing areas.

While a lively debate has ensued, very little data on business use of the Internet has informed the discussion. We address this gap. We construct measures of regional investment in advanced internet use by businesses as of late 2000 from a comprehensive data set of business internet use. Our measure includes a set of frontier applications such as “e-commerce” or “e-business” that excludes basic applications such as e-mail or web browsing. In contrast to earlier studies, we study a margin of investment and a time period in which IT facilitated communication over long distances and so, potentially, effect economically isolated work. We focus on studying advanced internet, a margin of investment that may require the deep labor pools and other complementary resources found primarily in cities.

We connect our internet data to measures of local economic performance, particularly wages. Our econometric approach compares a location’s economic performance before advanced internet technology diffused (i.e., 1995) to its performance after diffusion (i.e., 2000). That is, we use a difference-in-difference econometric estimation approach to identify the relationship between variance in the extent of investment in advanced internet technologies and variance in regional economic outcomes.

Our initial specification assumes that aggregate investment decisions by local establishments are exogenous to wage growth. We find that advanced internet investment is associated with an increase in county-level wages. This positive correlation remains robust to numerous specifications and changes in controls.

We address the assumption that investment is exogenous. First, we add many controls for factors know to shape investment decision and the results do not change. Second, we directly address what we consider the most likely issue: omitted variables bias at the regional level. The timing of effects points to the Internet as a key driver. We find no positive relationship between areas that would later adopt advanced internet and wage growth between 1990 and 1994. Further, there is no relationship between early usage of advanced internet and wage growth from 1999 to 2005. The early wave of internet adoption appears to be associated with a one-time change in wage growth across US locations during the late 1990s.

Our most interesting finding suggests that the internet caused divergence. We find a stronger correlation between wage growth and advanced internet in counties that already were doing well on a variety of measures. In particular, we find advanced internet is especially correlated with wage growth in the 180 counties that, as of 1990, had a population over 100,000 and were in the top quartile in income, education, *and* fraction of firms in IT-intensive industries. Overall, while the internet explains just 1% of the wage growth in the average county in our sample, it explains a quarter of the difference in wage growth between the 180 counties that were already doing well and all other counties.

Once again, we consider omitted variable bias. We find it difficult to speculate about which unmeasured regional-specific mechanism led to the results we find. Counties that were not in this group did not experience wage growth associated with advanced internet, even counties that were leading adopters. At the same time we find it easy to provide an explanation that assigns causality to the Internet.

A scatterplot of the raw data forecasts our core results. Figure 1a shows the relationship between advanced internet use and local wage growth for all types of counties in the data. Careful observation will show that the slope of the regression line is upward sloping (it is also significantly positive), but advanced internet is clearly not a core explanation of wage growth in the full sample. In contrast, figure 1b compares the 180 counties that were already doing well with the other counties. For the 180 counties that were already doing well, advanced internet is strongly correlated with wage growth; for the other counties, there is no relationship between advanced internet and wage growth in the raw data.[1] Advanced internet allowed counties that were doing well to do even better. In contrast, widespread usage of advanced internet does not seem to be correlated with wage growth in smaller, poorer, less educated

[1] To construct figure 1, we truncated the picture and consequently removed some counties with very low and very high internet use. The results are qualitatively similar when we include these counties, though visually not as clean.

counties with fewer IT-intensive firms. Of course, our analysis goes far beyond this scatter-plot, but the intuition continues to hold after a wide battery of corrections and tests.

Our results have important public policy implications, suggesting that efforts to improve broadband access will have limited effects on local wage growth. For one, our results show that investments in basic internet such as internet access have little association with wage growth. Investments in internet access need to be coupled with additional investments in advanced internet for local wage growth to appear. Further, we show that investments in advanced internet are associated with wage growth only in those counties that are already doing well, strongly suggesting the potential for income gain did not exist in all locations.

1.1 Related Literature

Our study contributes to a large macroeconomic literature on regional convergence and divergence.[2] We complement recent work that has examined how cross-sectional variance in factors such as education and industry composition contribute to convergence and divergence (e.g., Higgins, Levy, and Young 2006). The most closely related paper, Glaeser and Ponzetto (2007), show that an increase in the share of skilled occupations is associated with greater local wage growth. They do not focus on IT. Indeed, more broadly, no research in the convergence/divergence debate has investigated the links between regional convergence/divergence and the internet-led investment boom of the 1990s. This gap in understanding requires attention because the IT investment boom was an important factor in the late 1990s economic experience. It also has been associated with growth and productivity gains at other levels of aggregation, such as the nation, industry, and firm.[3]

There is a literature on IT investment and regional growth, but it has not focused on the role of internet in fostering income convergence or divergence. For example, Beaudry, Doms, and Lewis (2006)

---

[2] See Magrini (2004) for a recent survey. For other research on the causes of convergence/divergence in regional growth see, e.g., Glaeser et al (1992), Barro and Sala-i-Martin (1991), and Higgins, Levy, and Young (2006). The literature linking technology to convergence across countries dates from Gershenkron (1962).

[3] IT-using industries and firms had exceptionally good macroeconomic performance in the 1990s, measured at the national (e.g., Jorgenson, Ho, and Stiroh 2005), industry (e.g., Stiroh 2002), firm (e.g., Brynjolfsson and Hitt 2003), and establishment levels (e.g., Bloom, Sadun, and Van Reenen 2007).

focus on how variations in PC use over 1980-2000 and local skills correlate with wage growth. Similarly, Kolko (1999, 2002) finds that IT use over 1977-1995 is associated with the fastest employment growth in agglomerated areas, though this is due to the presence of local skills. In comparison, our paper focuses on the internet as a distinct factor, and we examine a later time period, because we believed increasing communication capabilities of the internet would make convergence more likely. We also differ in our findings. We show that high skill levels are only part of the story: high income, high population, and a high number of firms in IT-intensive industries also shape whether the internet affected wage growth.

Another paper related to ours is Aral, Wu, and Morabito (2007), who use firm-level data from Italy to show that enterprise resource planning is most associated with productivity gains in regions with weak human capital and technological infrastructure. We find contrasting results to theirs at the local level, perhaps because we focus on the internet not ERP, perhaps because our econometric strategy leads to sharper estimates, or perhaps because our results simply reflect differences between Italy and the US.

While prior work has not examined whether the internet leads to convergence, a number of recent papers have demonstrated how communicating over the internet may lower the costs of engaging in economic activity in geographically isolated regions. For example, it has been shown that use of internet and related communication technologies lower the costs of retail shopping for isolated consumers (Forman, Goldfarb, and Ghose 2008) and stimulate greater job migration (Stevenson 2006). Further, it is widely accepted that lower communication costs enabled by IT and the internet have enabled the delivery of a set of tradable services at significant distances from the point of final demand (Arora and Gambardella 2005; OECD 2006).

Our paper shares similarities with the literature on skill-biased technical change, but there are also significant differences.[4] Research on skill-biased technical change generally tests the premise that changes in technology use alters the demand for skilled labor, thereby shifting the wage distribution in favor of skilled occupations (e.g., Autor, Katz, and Krueger 1998; Autor, Levy, and Murnane 2003). By

[4] See, e.g. Autor (2001) for a survey and Xiang (2005) for a recent example.

demonstrating that wage gains are greatest in locations with more highly skilled workforces, our results are consistent with this prior work. However, the variance in our data across regions explores a dimension often overlooked by prior studies. In our data variance arises from cross-sectional differences in average incomes, education, and other factors across locations, rather than variance in skills across occupations. Further, our results suggest that a highly skilled labor force alone is insufficient for a location to realize wage gains for internet investments: they must be present with other factors that shape local labor markets, the right combinations of high population, income, and industry composition.

While we make this novel connection, we do not fully connect it to the literature on biased technical change, partly due to data limitations. For example, we cannot examine whether wage gains are greatest for high or low skilled occupations within a county, nor can we examine how internet use changes the wage distribution within a location. Our findings will raise questions about such connections, and we leave that for further work.

2. The internet and the localization of growth

We measure the effects of advanced internet use on convergence by proceeding in two broad steps. We begin by measuring the average relationship between internet use and wage growth across all counties. Second, we identify convergence or divergence by examining whether advanced internet investment led to faster growth in high or low income areas.

Step 1: Advanced internet use and local wage growth. Basic economic reasoning combined with the findings from a range of prior studies suggests that use of the internet may be associated with accelerated local wage growth through two mechanisms. First, productivity advances at the establishments using advanced IT will raise demand for labor at those establishments. If enough local firms become more productive through internet use, local labor demand will rise and local wages should be higher. Second, internet use may lead to wage growth through narrower mechanisms. Increases in IT demand may put pressure on local markets for skilled IT workers such as programmers, database administers, and consultants. If the local supply of these occupations is inelastic, then increases in the demand for IT may translate into higher wages for these occupations. Over time, this mechanism also may put upward pressure on wages in other occupations which drew from a similar labor pool.[5]

To measure the impact of the internet on local wages, we use a difference-in-difference identification strategy, comparing a time period before advanced internet technologies diffused (1995) to a period where we observe use (2000). We take advantage of the fact that many features of regions that shaped labor markets and enterprises in 1995 did not change by 2000. Our endogenous variable will be Y_{it}, which represents the level of wages in a particular region (i) and year (t). Only a small set of research establishments employed advanced internet in 1995 and therefore we set this variable to zero in 1995. Our approach yields the panel regression:

$$(1) \qquad Log(Y_{it}) = \alpha_1 X_{it} + \alpha_2 Z_{it} + \beta Internet_{it} + \tau_t + \mu_i + \varepsilon_{it},$$

Here τ_t is a time dummy that captures average changes to wage levels over time, μ_i is a location-specific fixed effect that gets differenced out in the estimation, and $Internet_{it}$ measures the extent of advanced internet use by businesses in region i at time t.[6] We have assumed that ε_{it} is a normal i.i.d. variable.[7] We include two kinds of controls; X_{it} are controls for pre-existing factors that may affect wage growth such as income, population, and education. We set these to zero in 1995 so that they are not differenced out of the regression. This allows us to control for the degree to which *levels* of the variables affect *changes* in wages. Z_{it} are controls for changes in the factors that not directly related to income over time where we have data, as well as internet use by local households (the full list is in Table 1b).[8]

---

[5] We are unable to identify which of these explanations is most likely due to data constraints: we could find no reliable county-level data on 1995 wages for IT workers.

[6] As in Athey and Stern (2002) and Hubbard (2003), we treat the general diffusion of a new technology as an exogenous factor that leads to a change in economic outcomes. As in those papers, we examine the plausibility of the exogeneity assumption, i.e., whether unobservable factors correlated with adoption have a causal relationship with the endogenous variable.

[7] Since we estimate the standard errors using heteroskedasticity-robust methods, the two-period framework is especially appealing. Stock and Watson (2008) show that the standard fixed-effects heteroskedasticity-robust variance matrix estimator is inconsistent if T is fixed and greater than 2.

[8] Taking the first difference yields a standard growth equation of changes in wages on levels and changes in the covariates. Since we treat the standard errors appropriately, this means that the coefficients and standard errors are exactly equivalent to those estimated using a growth equation.

Our hypothesis is that increases in local business use of advanced internet will be associated with growth in local wages: a test of $\beta > 0$ against the null of $\beta = 0$.

We assume that the unobservable determinants of wages can be decomposed into an additively separable fixed component and a time-varying component that is constant across counties. We also assume that no systematic factors in ε_{it} are correlated with the unobserved difference in Y_{it}.

One potential concern in this model may be that unobservable changes to local firm or worker characteristics may be correlated both with wage growth and internet use. To fully eliminate these omitted variable bias concerns, we would need many control variables and a convincing instrument that is correlated with advanced internet use at the county level *but not* with wage growth. While we do have a long list of controls, in our view every factor that shapes Internet use by business also potentially shapes regional wage growth. Rather than make claims about instruments that lack credibility in this first step, we prefer to declare that we e have been unable to identify an appropriate instrument and circumscribe inferences with additional action.

We provide considerable suggestive evidence that, when combined, shows that advanced internet use by firms is strongly correlated with local wage growth. First, as noted above, we include many controls for the initial conditions of the county in order to address omitted variables bias at the regional level. Further, we include controls for changes in county characteristics such as population, racial composition, and age. We also include controls for changes in closely related margins of consumer and business IT investment such as basic internet investment, PCs per employee, and internet use at home, which also vary considerably across regions. If advanced internet is associated with wage growth but not these other margins of IT investment, then omitted variable bias must be specific to advanced internet investment.

The internet's unusual history also gives us an additional test for the role of regional-specific omitted variables: it enables us to employ a useful falsification test. Because the internet was originally developed to facilitate research collaboration, it did not become clear until late 1994 that commercial implementation of the internet could have a wide economic impact, and even at that point it was only apparent to internet insiders. Most incumbent vendors and users in communications and computing markets, as well as IT-intensive users in the economy as a whole, were caught by surprise in late 1994 and early 1995 as the commercial internet rapidly diffused. That (almost sudden) realization by so many firms contributed to a non-gradual response from both vendors and users. As a result, we should not see any affiliation between internet investment and local economic activity sooner than 1995, and probably it should arise later.[9] If our assumptions of the orthogonality between internet and changes in local unobservables are violated then our data will produce "false positive" associations between future use and growth in a period prior to 1996. If we find false positives, then it suggests that violations of our identification assumptions are artificially inducing $\beta > 0$. If not, then it boosts confidence in the exogeniety assumption embedded in (1).[10]

Step 2: Use of the internet and convergence/divergence. The internet's rapid diffusion pattern motivates examining divergence and convergence. The internet achieved the symptoms of near ubiquity in the US in a short period. By the end of the 1990s over half the households in the US had internet access and over 90% of medium and large establishments did as well. Since adoption was widespread, it is constructive to ask whether the changes in economic outcomes were too.

We test between the two starkest predictions of the impact of the internet on wage growth. One view predicts that the internet would improve growth prospects in many regions, and especially in low-

[9] We think it is safe to date the beginning of the investment boom in the internet to 1996. Dating the rise of the commercial internet is not an exact science, but a few well-known events provide useful benchmark. The National Science Foundation privatized the internet at about the same time that Tim Berners-Lee began to diffuse the basic building blocks for the World Wide Web. The first non-beta version of the Netscape browser became available in early 1995. The Netscape IPO occurred in August 1995 and it went spectacularly well for the founders and funders. Microsoft's announced its change in strategy on December 7, 1995. Certainly no serious vendor in IT markets was ignoring the commercial internet by December 1995. No large scale investor in IT applications was either by this point, but major investment tends to lag planning and change only slowly. As Forman (2005) discusses, "adjustment costs" slowed down deployment thereafter, but more than that was at work, as firms experimented and learned about new uses.

[10] We provide further details on our exploration of endogeneity and omitted variables bias in the results section below.

density, economically isolated regions where basic internet services were particularly valued for lowering communications costs. The contrasting view casts the internet as a technology that exacerbated existing inequalities in wages between urban/rural and frontier/mainstream users of IT.[11]

We begin by examining whether the benefits of internet use are greatest in areas with low or high income. Advanced internet use may contribute to divergence due to an overheating effect on local wage growth from local economic prosperity. One example can illustrate. During this time period establishments supplying parts for electronics goods and for automobiles had a high propensity to invest in advanced internet. However, an electronics parts supplier in San Jose, California, faces a very different local labor market than an automobile parts supplier in Akron, Ohio. While internet use in the face of tight local labor demand contributes to rising wages, similar investments in environments without tight local labor demand conditions do not. Cross sectional regional variance in income may also be correlated with variance in local skills, industry composition, or population size that could similarly contribute to divergence; below we explore the possibility that these other local characteristics could contribute to divergence.

Advanced internet use could similarly lead to convergence, or a "Robin Hood" forecast for the economy-wide impact from the diffusion of the internet. That is, as a communications technology with nearly instant universal availability, the technology might lead to widespread productivity advances across many facets of the economy. We do not dismiss this view out of hand as unrealistic for two reasons. First, it received considerable popular attention at the time, though no systematic test has ever confirmed or refuted any related prediction that goes beyond anecdote. Second, systematic statistical study of the diffusion of the basic Internet to business is consistent with a premise behind this view: other

[11] Cairncross (1997) was among the earliest work to forecast that the internet technology would lead to significant changes in the spatial distribution of economic activity. By reducing the costs of economic isolation, internet technology can shift economic transactions from locations in urban areas where average wages are relatively high (Glaeser and Mare 2001) to rural locations where wages are relatively lower. Forman, Goldfarb and Greenstein (2005) investigate the diffusion of the internet to business, but find evidence consistent with both views. They find some evidence consistent with basic internet technology diffusing first to rural and small urban areas, but also find evidence consistent with urban areas leading in the use of advanced internet. For a review of some of this literature see Forman and Goldfarb (2006) and Greenstein and Prince (2007).

works suggests that the basic Internet had high marginal benefit to businesses in isolated and low density areas, as well as low adaptation costs.

In our first approach we estimate the simplest version of this hypothesis:

(2) $$Log(Y_{it}) = \alpha_1 X_{it} + \alpha_2 Z_{it} + \beta Internet_{it} + \phi(Internet_{it} * HighIncome_i) + \tau_t + \mu_i + \varepsilon_{it},$$

Here, ϕ measures the difference for high and low income counties in the relationship between wages and advanced internet. Divergence caused by the internet will produce $\phi > 0$ and convergence $\phi < 0$ against a null of $\phi = 0$. Rejecting the null does not depend on β, but the estimate for β (combined with the estimate for ϕ) does shape the inference about the economic importance of the internet for wages.

In our next approach to estimating convergence/divergence we focus less on income per se'. Rather, we focus on the local factors that influence local labor market conditions, such as local skills, population (agglomeration), and industry composition.

We focus on skills because considerable evidence points towards complementarities between the use of advanced information technology and a skilled labor force, implying higher wages due to these complementarities.[12] We focus on population (agglomeration/density) because larger cities had thicker labor markets for complementary services, specialized skills, or specialized vendors. The presence of complementary resources also increased the net returns to IT-based process innovations, increasing the returns to productivity and growth from internet adoption for enterprises in such locations.[13] We focus on industry composition because the clustering of IT-intensive enterprises in the same industry, accentuated by similarities in investment behavior within industries (and, therefore, within regions), could lead to a simultaneous increase in demand for labor in agglomerated productive enterprises.[14]

---

[12] See e.g., Bresnahan, Brynjolfsson and Hitt 2002, Autor (2001), and Corali and Van Reenen (2001).

[13] A rich literature in urban economics has provided evidence on the presence of urban increasing returns and productivity benefits associated with location in an urban area (e.g. Rosenthal and Strange 2004). The reasons go back to Marshall's (1920) initial insights: thicker labor markets, input sharing, and knowledge spillovers. The same reasoning applies to urban areas with an agglomeration of IT-using industries and why they might have advantages over areas without such IT-using industries.

[14] Further, if IT-intensive enterprises earn greater productivity benefits from new IT use (Stiroh 2002) then industries with such enterprises will have the largest associated wage gains from internet use. When these productivity benefits spill over to other enterprises (e.g., Greenstone, Hornbeck, and Moretti 2008), then locations with the "right" industries will experience broad-based wage gains that are greater than other equivalent locations

We examine the extreme position that all of these factors matter and divide counties by skills, population, and the IT intensity of firms (in addition to income). We use this extreme position because it provides a way to simplify the underlying five-way interaction into a single interaction term. Call the high counties "high all factors." To investigate whether these other factors affect the divergence in incomes associated with internet adoption, we estimate:

(3) $Log(Y_{it})=\alpha_1 X_{it}+\alpha_2 Z_{it}+\beta Internet_{it}+\phi_1(Internet_{it}*HighIncome_i)+$

$\phi_2(Internet_{it}*HighEducation_i)+\ \phi_3(Internet_{it}*HighPopulation_i)+\ \phi_3(Internet_{it}*HighIT\text{-}$

$Intensity_i)+\ \phi_4(Internet_{it}*HighAllFactors_i)+\tau_t+\mu_i+\varepsilon_{it},$

Here, ϕ_1 measures the difference between counties with high and low incomes and ϕ_4 measures differences between counties with high and low all factors. If $\phi_1=0$ but $\phi_4>0$ then divergence in incomes is isolated to regions with high education, population, and IT intensity.

A finding of $\phi_1=0$ but $\phi_4>0$ also has implications for identification in the presence of the potential for omitted variables. If this result is a false positive caused by positive covariance between changes in ε_{it} and advance Internet, it suggests this covariance is isolated only to locations that are well off. While it is always possible that such unobservables may exist, we find it challenging to to identify what economic mechanism might produce such an unobservable in only a limited number of places and not others.

3. Data

To measure how internet investment influences growth in wages, we combine several data sources about medium and large establishments. Our IT data comes from the Harte Hanks Market Intelligence Computer Intelligence Technology database (hereafter CI database). The database contains establishment- and firm-level data on characteristics such as number of employees, personal computers per employee, and use of internet applications. Harte Hanks collects this information to resell as a tool for the marketing divisions of technology companies. This source has been used by other economic

---

that have invested in advanced internet.

researchers as a fruitful way to learn about enterprise IT use.[15] Interview teams survey establishments throughout the calendar year; our sample contains the most current information as of December 2000.

Harte Hanks tracks over 300,000 establishments in the United States. Because we focus on advanced internet applications, we exclude government, military, and nonprofit establishments. Our sample from the CI database contains commercial establishments with over 100 employees—in total 86,879 establishments.[16] While the sample only includes relatively large establishments, we do not view this as a problem because very few small establishments employed advanced internet in the late 1990s. The primary investors were large establishments making large scale enterprise-wide investments worth tens of millions, and, in some large multi-establishment organizations, hundreds of millions of dollars per year.[17]

We focus on those facets of internet technology that became available only after 1995 in a variety of different uses and applications. Our raw data include at least twenty different specific applications, from basic access to software for internet-based enterprise resource planning (ERP) business applications software. Advanced internet involves frontier technologies and significant adaptation costs. We identify use of advanced internet from the presence of substantial investments in e-commerce or e-business applications.[18]

We stress that the investments we consider include several aspects of an enterprise's operations, not just the most visible downstream interactions with customers. These often involve upstream communication with suppliers, and/or new methods for organizing production, procurement, and sales

[15] There is an increasingly long list of papers that have built on this data source and its predecessor from Computer Intelligence, including our own prior work.

[16] Establishments were surveyed at different times from June 1998 to December 2000. To control for increasing adoption rates over time, we reweight our adoption data by the ratio of average adoption rates in our sample between the month of the survey and the end of 2000.

[17] All our available evidence suggests that adoption monotonically increased in firm size, even controlling for many other determinants. Hence, our sample represents the vast majority of adopters, and this procedure leads to the best possible estimate of use in a location.

[18] In previous work we had labeled this variable "enhancement" because it enhanced existing IT processes and contrasted it with "participation", i.e., use of basic internet technologies, such as email or browsing (e.g. Forman, Goldfarb, and Greenstein 2002, 2005). In this paper the contrasts are not the central focus, so we simply call it advanced internet, and, when necessary, we will contrast it with basic internet use and PC use.

practices. We look for commitment to two or more of the following internet-based applications: ERP, customer service, education, extranet, publications, purchasing, or technical support.[19] Most often, these technologies involved inter-establishment communication and substantial changes to business processes. We have also experimented with a variety of alternative measures of business internet use and our results are qualitatively similar under these alternative definitions.

To obtain location-level measures of the extent of advanced internet use, we compute average rates of use for a location. Because the distribution of establishments over industries may be different in our sample than over the population, we compared the number of establishments in our database to the number of establishments in the Census. This data is only available at the level of the county. We calculated the total number of establishments with more than 50 employees in the Census Bureau's 1999 County Business Patterns data and the number of establishments in our database for each two-digit North American Industry Classification System (NAICS) code in each location. We then calculated the total number in each location. Therefore, to account for over- and under-sampling in the Harte Hanks data, we weight a NAICS-location by

$$\frac{\text{Total \# of census establishments in location-NAICS}}{\text{Total \# of census establishments in location}} \times \frac{\text{Total \# of establishments in our data in location}}{\text{Total \# of establishments in our data in location-NAICS}}$$

We sum the weighted establishment-level rates of use across establishments within a county to obtain county-level estimates of the extent of use of advanced internet.

Our prior research has shown that this measure has several attractive properties. For example, when aggregated to the industry level, this measure positively correlates with BEA measures of cross-industry differences in IT investment, but not perfectly, as it captures something distinct as well. Also, the cross industry differences appear plausible. Examples of industries that tend to be internet-intensive are electronics manufacturing, automobile manufacturing and distribution, and financial services (Forman,

[19] Additional details on the construction of this variable can be found in Forman, Goldfarb, and Greenstein (2002).

Goldfarb, and Greenstein, 2002). In addition, both industrial composition and features of local areas shape aggregate regional use in the direction economic intuition would forecast. Among the biggest cities, areas with high use are those where a high percentage of local employment is in internet-intensive industries (as well as IT-intensive), such as the San Francisco Bay Area, Seattle, Denver, and Houston (Forman, Goldfarb, and Greenstein 2005).

We obtain data on county average weekly wages, total employment, and total establishments from the Quarterly Census of Employment and Wages, a cooperative program of the Bureau of Labor Statistics and the State Employment Security Agencies. Matching these data to our internet data leaves a total of 2743 county observations. We dropped 372 of the total 3115 counties because we lack data on internet investment. We retain every urban and suburban county as well as most rural ones. The vast majority of the dropped counties come from lowest quartile of the population distribution.

In order to examine divergence, we construct a set of variables to interact with our measure of advanced internet investment. We focus on the role of income, education, population, and industry composition. The data on population, education, and income come from the 1990 US Census. For industry composition, we measure the fraction of firms in IT-using and producing industries in the county as of 1995 from the US Census County Business Patterns data. National aggregate data shows that such industries have unusually high returns from investment in IT in the 1990s. We define these industries using the classification reported in Jorgenson, Ho, and Stiroh (2005, p. 93). We call this “IT-intensity”.

We combine these data with county-level information from a variety of sources. This information allows us to control for the underlying propensity of the counties to grow and to innovate. First, the 1990 US Census provides county-level information on population, median income, percent graduated university, percent graduated high school, percent African American, percent below the poverty line, and percent over 65. We also use the 2000 US Census to control for changes in the non-income-related factors: population, percent graduated university, percent graduated high school, percent African American, net migration to the county, and percent over 65. The 2000 Current Population Survey (CPS)

Computer and Internet Use Supplement provide our data on the percentage of county households adopting the internet at home. We use four measures of county-level propensity to innovate: the number of students in Carnegie rank 1 research universities in 1990, the fraction of students enrolled in engineering programs, the percentage of the county's workforce in professional occupations in 1990, and the number of patents granted in the 1980s in that county, as found in the NBER patent database.[20]

Table 1a includes descriptive statistics on IT use and our measures of local wages, establishments and employment. Table 1b includes a description of control variables.

4. Empirical Results

We initially establish a link between investment in advanced internet and wages. We then show that no such link exists between advanced internet and employment. Next, we show that there is something different about advanced internet compared to basic internet applications and personal computer use. Next we show that investment in advanced internet contributes to divergence; in particular, advanced internet is only associated with wage growth in counties that have a combination of high income, education, population, and IT-intensive firms. We find no evidence that investment in advanced internet technology led to convergence.

4.1 Baseline Results

Table 2 shows the baseline results across counties. Column 1 shows the correlation between advanced internet investment and wages at the county level without any controls. As expected from the cross-tabs in Figures 1 and 2, there is a strong positive correlation with wage growth.

Column 2 includes county and year fixed effects, which alters the key estimates considerably, as we would expect. Column 3 provides what we view as our main specification. It includes controls for pre-sample demographics (such as county income and population in 1990), changes in non-income demographics (such as population from 1990 to 2000 and net migration from 1990 to 2000), measures of

[20] Downes and Greenstein (2007) showed that the first three factors help explain availability of internet infrastructure such as ISPs.

pre-sample innovativeness (such as patents granted in the 1980s) as well as changes in home internet adoption (effectively zero in 1995). Columns 4 though 7 provide a number of other specifications to show robustness. In every specification, we reject the null that advanced internet is not associated with wage growth.

In the main specification the coefficient on use of advanced internet technologies is 0.0252. That is, regions with an average level of internet use experienced 0.24% wage growth above what regions with no internet use experienced. A one standard deviation increase in the use of the internet is associated with 0.335% increase in wage growth. The data are quite skewed, so it is also interesting to look at the top decile, which is 0.216. That leads to a 0.32% increase in wage growth above the mean. Consistent with figure 1a, this suggests that the internet was not the primary force behind the 20% wage growth across all counties in our data from 1995 to 2000. Still, it is related to growth. As we show below, examining the average effect obscures variation across regions.

Even with such a small coefficient, omitted variables bias is an important concern in this analysis. As described in section 2, we take several steps to address this concern. First, we have included several controls for the initial conditions of the county. For example, if counties with a more educated population are more likely to experience a wage increase from 1995 to 2000 and are more likely to adopt advanced internet technologies then we may observe a positive correlation between internet technology and wage growth only because of this underlying correlation. For this reason, in Table 2 and subsequently, we control for several county-level characteristics including pre-sample population, racial composition, education, income, poverty, professional workers, enrollment in research universities, enrollment in engineering programs, age, and innovativeness as measured by patents granted in the county.

Second, we have included controls for changes in county characteristics: population, migration, racial composition, education, and age. We also control for internet use at home as measured by the 2000 CPS Computer and Internet Use Supplement. Therefore, we are examining the question of whether advanced internet use at work affects wages independent of individual-level propensities to use the internet on their own. Comparing columns 3 and 5 of Table 2 shows that this control has no substantive impact on the qualitative results.[21]

Third, we examine the timing of the relationship between advanced internet and wage growth to look for false positives. That is, advanced internet investment should only help firms in the latter half of the 1990s. Prior to 1995, the internet had not diffused. Therefore, in order to explore whether our measure of internet investment is simply capturing county-level propensity to grow, we show that our measure of internet investment is not correlated with wage growth prior to 1995.

Three versions of this falsification check are presented in table 2 columns 8 and 9 and in figure 2. Columns 8 and 9 replicate columns 3 and 4 but use logged wages in 1990 and 1994 as the dependent variables rather than logged wages in 1995 and 2000. Significance is lost and the coefficient magnitudes fall substantially. Thus counties with high levels of advanced internet in 2000 do not appear have grown faster prior to 1995. Figure 2 replicates the regression in column 3 using data for all years from 1990 to 2000. The controls are the same as in column 3 and the dependent variable is logged wages. We interact the measure of advanced internet (as of 2000) with year dummies from 1991 to 2000 in the same way and therefore we get a measure of the association between advanced internet and wage growth over the year period. Figure 2 clearly shows advanced internet investment is not correlated with wage growth until 1997 (when the internet began to diffuse widely). Between 1991 and 1996 the coefficient is statistically indistinguishable from zero in every year except 1992.

Table 3 examines endogenous variables other than wages. It shows that we find no consistent measurable relationship between advanced internet and employment or the number of establishments. These columns are representative of our more general finding through many more analyses that neither total employment nor establishments are correlated with internet use in any systematic way. These results suggest that the increase in wages is not related to a substantial negative effect on employment or the

[21] This data is only available for a subset of counties that can be identified in the CPS. For this reason, we also include a dummy variable that captures when this information is not available. We also show robustness to running the analysis on this subset of counties.

number of firms. Since we find no systematic pattern in the data for the relationship between internet use and employment or the number of establishments, for the rest of the study we focus on our wage results.

Finally, we ask whether advanced internet proxies for other kinds of information technology. Table 4 examines how county-level wages change with advanced internet investment, basic internet investment, and PCs per employee. These are all measured using the Harte Hanks data base and aggregated to the county level. Forman, Goldfarb, and Greenstein (2005) use the same measure of basic internet investment and found it to be widely adopted by 2000. The measure of PCs per employee resembles that found in Beaudry, Doms, and Lewis (2006).

The results suggest that advanced internet investment is distinct from other measures of IT. While PCs per employee are positively correlated with wage growth, this relationship is no longer significant once the controls are included. Further, including PCs per employee and basic internet investment as controls does not substantially change the marginal relationship between advanced internet investment and wages. This table suggests that advanced internet investment is not simply a surrogate measure of IT intensity but that the relationship between wage growth and advanced internet investment is related to advanced internet technology in particular. Indeed, since the correlation between advanced internet and PCs per employee at the county level is 0.20 and the correlation between basic and advanced internet is 0.18, the table also suggests that wages in some areas could especially diverge from others when both are high.

We have investigated the robustness of this finding and found no systematic relationship between basic internet technologies and growth in employment, establishments, or wages. This is surprising because levels of participation were high across establishments and locations by 2000. Revealed preference suggests the benefits were high, especially for a technology with so little use only five years earlier. We speculate that our intuition about revealed preference applies to an infra-marginal adopter. In other words, when just about everyone has adopted a technology, the data may be identifying an uninteresting margin in the benefits of participation. Said another way, with basic Internet technology there simply is too little variation in the independent variable.

Overall, our results suggest an association between advanced internet and wage growth in the late 1990s. Further, these results suggest that if our results are due to false positives arising from omitted variable bias, we can circumscribe the features of these omitted variables. They must be correlated with advanced internet but not other margins of IT investment, nor many other persistent regional features.

4.2 When was advanced internet technology related to local wage growth?

In this section, we provide evidence that adoption of advanced internet led to divergence in wages across counties. In particular, advanced internet adoption is primarily correlated with county-level wage increases in counties with high income, education, population, and a large percentage of IT-intensive firms.

Our regression results in Table 5 explore this pattern in several steps. Column 1 shows that advanced internet is only significantly associated with wage growth in counties in the top quartile of median income as of 1990. This means that counties in the top income quartile with high levels of advanced internet grew faster than counties in the top income quartile with low levels of advanced internet. In contrast, counties in other quartiles with high levels of advanced internet did not experience especially rapid wage growth. In short, advanced internet adoption contributes to divergence.

Columns 2 through 4 show how the impact of advanced internet is influenced by variation in local education, IT-intensity, and population. Like Column 1, Column 2 shows that advanced internet is associated with wage growth only for counties in the top quartile of higher education (percentage of the population that graduated university as of 1990). The similarity of results is not surprising since 60% of the counties overlap. Column 3 shows that counties with over 100,000 people display a strong association between advanced internet and wage growth.

Column 4 examines counties in the top quartile in IT-intensity. In this specification, advanced internet is not significantly correlated with wage growth for high IT-intensity counties. Still, we include

IT-intensity for three reasons. First and perhaps most importantly, IT-intensity has been emphasized in much of the previous literature linking information technology to average productivity (e.g. Jorgenson, Ho, and Stiroh 2005). Second, the coefficient is positive and when added to the coefficient on the main effect in the first row, it is significantly different from zero with 95% confidence. Third, we tried several specifications and the coefficient was sometimes significantly positive and never negative. Therefore, while we are concerned about the observational equivalence between IT-intensity and other observable regional attributes, we cannot reject a role for IT-intensity in the relationship between advanced internet and wage growth.

Column 5 shows that when we include all four measures of pre-internet county strength (income, education, population, and IT-intensity), none end up significant. This may not be surprising given that there is considerable overlap between the measures: Each measure contains roughly 680 counties, of which 180 are in the top group in all measures. Columns 6 and 7 show that it is in these 180 counties that advanced internet is correlated with wage growth. Column 7 shows that it is the combination of more than one factor that drives the relationship between advanced internet and wage growth. There is something different about the 180 counties with high income, education, population, and IT-intensity.

The core results of table 5 are robust to using continuous measures of income, education, population, and IT-intensity. Income loses significance and IT-intensity gains significance but the significance and importance of the interaction term remains. Furthermore, adding all two-way interactions to column 7 (i.e. high income and high education, high income and high population, etc.) does not change the core result: a large and significant coefficient for the 180 counties that were already doing well on all four measures.[22] These results are all included in the Appendix. Using the same method as figure 2, figure 3 provides a falsification check of the results in column 6. It shows that the relationship between advanced internet and wages begins in the late 1990s.

[22] Adding the complete set of three-way interactions leads everything to be insignificant. We believe there is too much overlap in the measures to get significant estimates.

These 180 counties also had higher wage growth than the other 2563 counties in the sample: 29.2% vs. 20.5%. For the 180 counties with high income, population, education, and IT-intensity, our results suggest that advanced internet use is related to 8.2% (2.4 percentage points) of the total wage growth. For the other counties, advanced internet explains just 1.1% (0.23 percentage points) of overall wage growth.[23] Combined, this means that advanced internet explains one quarter of the 8.7 percentage point difference in wage growth between these 180 counties and the other 2563 counties in the sample.

These results further circumscribe concerns about omitted variables. There is not a clear endogeneity story to explain the difference between regions with high all factors and other regions. For example, if otherwise unmeasured regional prosperity causes wages and investment to both rise, why is income growth only leading to internet investment in places that were already doing well? Income growth is unrelated to internet investment in other places even if they grew, and even if they were high adopters.

An analysis of outliers and "typical" cases among these 180 counties provides further details on the relationship between advanced internet and wage growth. Counties among the top 180 that have high internet and wage growth (both at least one standard deviation above the mean) include San Mateo and Santa Clara (both in San Francisco-Oakland-San Jose MSA); Boulder and Arapahoe, CO (Denver-Boulder-Greeley MSA); Fairfax, VA (Washington-Baltimore MSA); Travis (Austin-San Marcos MSA); and Washington, OR (Portland-Salem MSA). Those with high wage growth (one standard deviation above mean) but relatively low internet (below mean) include only Williamson, TX (Austin-San Marcos MSA) and Hudson, NJ (New York-Northern New Jersey-Long Island MSA). Those with high internet (one standard deviation above mean) but relatively low wage growth (below mean) include Madison, AL (Huntsville, AL MSA), Lake, OH (Cleveland-Akron MSA), Kalamazoo, MI (Kalamazoo-Battle creek MSA), and Middlesex, CT (New London-Norwich MSA). In short, counties with high internet and wage growth are often centers of IT production and use; counties with high internet but low wage growth are

[23] More precisely, these calculations use the coefficient estimates in table 5 column 6, the average Internet use for the 180 counties, and the average Internet use in all other counties.

often small cities where the labor markets are not very tight; and counties with high income but low wage growth are relatively rare.

We also stress these results can not arise due to the inordinate influence of canonical outliers. For example, we could remove Santa Clara or San Francisco from the data set and the results would not substantially change. In part, this should not be surprising; no single variable, not even advance Internet, could possibly explain the anomalous experience in Santa Clara in this time period (i.e., over 80% wage growth in five years). Mostly, however, the robustness of results to the exclusion of observations reflects the pattern in the data. There was a general experience found in a special set of urban counties outside Santa Clara. These 6% of US counties shared similar demographic and industrial traits prior to the Internet's diffusion and reacted to the diffusion of the Internet with similar economic experiences.

4.2 Additional implications of advanced internet

In this section we investigate whether the benefits of internet investment persist over time for *our* 180 top counties and also show whether these benefits can spill over to adjacent locations. These analyses will help us to circumscribe inferences, i.e., whether the effects of advanced internet are localized in time and space.

Table 6 explores whether early internet adopting counties continued to have higher wage growth once the diffusion of the internet slowed. It repeats the regressions in table 5 columns 6 and 7 but explores wage growth between 1999 and 2005 as the dependent variable. It shows that the difference between the 180 counties that were already doing well in 1990 and the other counties was coincident with the one-time diffusion of the internet. Advanced internet usage is related to rapid growth from 1995 to 2000 in places that were doing well in 1990. Then, these places maintained their new position in absolute terms. They did not grow faster, but the gains were not reversed either.

Table 7 examines whether the benefits of advanced internet in *HighAllFactors* counties can spill over to adjacent locations. We examine this question because local labor markets may extend beyond county boundaries. This is particularly likely in metropolitan areas, where workers frequently commute between counties. To investigate this possibility, we reran the regressions in columns 6 and 7 of table 5 adding a new variable that is equal to one when the county is located in an MSA with a *HighAllFactors* county but is not itself one of these counties. The coefficient on this new variable is positive but insignificant, suggesting that any spillover benefits to being located near a *HighAllFactors* county are positive, small, and not precisely estimated.

4.3 Open issues about biased technical change

We have introduced regional variation in wages into the discussion about the economic impact of the internet. Our findings raise questions about local variation in the productivity benefits of IT use both in and out of IT-intensive industries in particular regions. It also raises questions about local variation in the links between IT use and worker skills and education at a variety of levels.

Our findings stress the results for average wage growth, but motivate further work on the mechanisms at work. We do not fully connect our results to the literature on biased technical change, largely due to data limitations. For example, we cannot examine whether wage gains are greatest for high or low skilled occupations within a county, nor can we examine how internet use changes the wage distribution within a location. In particular, consistent county wage series for programmers were not kept earlier than 1999. Hence, we have been unable to link programmer and non-programmer wage experience to their changes before and after the deployment of the internet.

5. Discussion

In this study, we find evidence that an association between use of advanced internet technology and local wage growth. Further, we find that advanced internet is associated with divergence in wages: we only observe wage growth in locations in the top quartile of income. Probing this relationship further, we find that wage gains associated with advanced internet adoption were isolated to relatively populated locations in which IT production and use were concentrated, and where income and skills were high. This appears to have led to a one-time relative gain in wages for these locations. We also find little evidence that use of advanced internet was associated with growth in either employment or establishments.

Importantly, despite recent evidence that internet use may lower the costs of geographically isolated economic activity, there is no evidence in our data that advanced internet contributed to convergence in wages. In particular, our results do suggest the existence of a considerable divide in the benefits of advanced internet use across urban and rural areas; however, the results in this paper do not support the use of subsidies to build infrastructure to lower that gap. Rather, our results suggest that efforts to improve regional internet access would have little impact. Even if they are followed by investment in advanced internet within business, these investments only succeeded in raising wages in places that already had high levels of income, education, population, and IT-intensive firms.

Our work suggests that the returns to IT use may be higher when several factors appear together, such as an IT-friendly local industry, a skilled labor force, high local incomes, and a thick local labor market due to a high population. We think this changes the debate about the economic impact of IT and focuses attention on regional variation. It also points to the key role the Internet played in recent experience.

Considerable complementary evidence would be needed to overcome warranted caution about drawing too much from one exercise; however, our results raise many provocative questions in directions that no prior research has explored. We hope this inspires other research to continue to understand the underlying economic mechanisms.

References

Aral, Sinan, Lynn Wu, and Vincenzo Morabito. 2007. Regional Economic Context and the Value of Firm-Level IT Investments. Working paper, NYU.

Arora, Ashish, V.S. Arunachalam, Jai Asundi, and Ronald Fernandes. 2001. The Indian software services industry. *Research Policy* 30(8), 1267-1287.

Arora, Ashish and Alfonso Gambardella. 2005. *From Underdogs to Tigers: The Rise and Growth of the Software Industry in Brazil, China, Ireland, and Israel.* Oxford: Oxford University Press.

Athey, S., and S. Stern (2002) "The impact of information technology on emergency health care outcomes," *RAND Journal of Economics* 33(3), 399–432.

Autor, David, 2001, Wiring the Labor Market, *Journal of Economic Perspectives* 15(1): 25-40.

Autor, David, Lawrence F. Katz, and Alan B. Kreuger. 1998. Computing Inequality: Have Computers Changed the Labor Market? *Quarterly Journal of Economics* 113 (4), 1169-1214.

Autor, David, Frank Levy, and Richard J. Murnane. 2003. The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics* 118(4), 1279-1334.

Barro, R. J. and X. Sala-i-Martin. 1991 Convergence across States and Regions. *Brookings Papers on Economic Activity* 1: 107-158.

Beaudry, Paul, Mark Doms, and Ethan Lewis. 2006. Endogenous Skill Bias in Technology Adoption: City-Level Evidence from the IT Revolution. Federal Reserve Bank of San Francisco working paper #06-24.

Bloom, Nick, Raffaella Sadun, and John Van Reenen. 2007. Americans do IT Better: US Multinationals and the Productivity Miracle. CEP Discussion Paper No. 788.

Bresnahan, Timothy, Erik Brynjolfsson, and Lorin Hitt. 2002. Information Technology, Work Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *Quarterly Journal of Economics* 117: 339-376.

Brynjolfsson, Erik and Lorin Hitt. 2003. Computing Productivity: Firm-Level Evidence. *Review of Economics and Statistics* 85(4): 793-808.

Cairncross, F. 1997. *The Death of Distance*. Cambridge: Harvard University Press.

Corali, Eva, and John Van Reenen, 2001, "Skill Biased Organizational Change? Evidence from British and French Establishments." *Quarterly Journal of Economics,* 114(4), 1449-92.

Downes, Tom and Shane Greenstein. 2007. Understanding Why Universal Service Obligations May be Unnecessary: The Private Development of Local Internet Access Markets. *Journal of Urban Economics* 62(1): 2-26.

Forman, Chris, 2005. The Corporate Digital Divide: Determinants of Internet Adoption. *Management Science,* 51(4), pp 641-54.

Forman, Chris, Anindya Ghose, and Avi Goldfarb. 2008. Competition between Local and Electronic Markets: How the benefit of buying online depends on where you live. *Management Science,* Forthcoming.

Forman, Chris, and Avi Goldfarb. 2006. Diffusion of Information and Communication Technologies to Businesses. In (eds. Terrence Hendershott) Handbook of Information Systems, Volume 1: Economics and Information Systems. Amsterdam: Elsevier, p. 1-52.

Forman, Chris, Avi Goldfarb, and Shane Greenstein, 2002, Digital Dispersion: An Industrial and Geographic Census of Commercial Internet Use. NBER Working Paper #9287.

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

Friedman, Thomas. 2005. *The World is Flat: A Brief History of the Twenty-First Century*. New York: Farrar, Straus, and Giroux.

Gerschenkron, Alexander. 1962. *Economic Backwardness in Historical Perspective*, Belknap Press of Harvard University Press, Cambridge, MA.

Glaeser, E. L., H. D. Kallal, J. A. Scheinkman, and A. Shleifer. 1992. Growth in cities. *Journal of Political Economy* 100: 1126-1152.

Glaeser, Edward L., and David C. Mare. 2001. Cities and Skills. *Journal of Labor Economics* 19(2), 316-42.

Glaeser, Edward L., Giacomo A. M. Ponzetto. 2007. Did the Death of Distance Hurt Detroit and Help New York? NBER Working Paper 13710.

Greenstein, Shane and Jeff Prince, 2007. "The Diffusion of the Internet and the Geography of the Digital Divide," in (Eds. Robin Mansell, Danny Quah, and Roger Silverstone), *Oxford Handbook on ICTs*, Oxford University Press.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2008. Identifying Agglomeration Spillovers: Evidence from Million Dollar Plants. NBER Working Paper 13833.

Higgins, Matthew J., Daniel Levy, and Andrew T. Young. 2006. Growth and Convergence across the United States: Evidence from County-Level Data. *Review of Economics and Statistics* 88(4): 671-681.

Hubbard, Thomas N. 2003. Information, Decisions, and Productivity: On Board Computers and Capacity Utilization in Trucking. *American Economic Review* 93(4): 1328-1353.

Jorgenson, Dale, Mun S. Ho, and Kevin Stiroh. 2005. *Productivity Volume 3: Information Technology and the American Growth Resurgence*. MIT Press: Cambridge, MA.

Kolko, Jed. 1999. Can I Get Some Service Here? Information Technology, Service Industries, and the Future of Cities. Working Paper, Harvard University.

Kolko, Jed. 2002. Silicon mountains, silicon molehills: geographic concentration and convergence of internet industries in the US. *Information Economics and Policy* 14(2): 211-232.

Magrini, Stefano. 2004. Regional (Di)Convergence. In (eds. J. Vernon Henderson and Jacques-Francois Thisse) *Handbook of Regional and Urban Economics, Volume 4*. Amsterdam: Elsevier North-Holland, 2243-2292.

Marshall, Alfred. 1920. *Principles of Economics.* 8th edition. New York: Porcupine Press.

Moss, Mitchell L., and Anthony M. Townsend, 1997. Tracking the net: using domain names to measure the growth of the Internet in US cities. *Journal of Urban Technology* 4(3): 47-60.

OECD. 2006. *OECD Information Technology Outlook 2006.* Paris: OECD.

Rosenthal, Stuart S. and William C. Strange. 2004. Evidence on the Nature and Sources of Agglomeration Economies. In Handbook of Regional and Urban Economics, Volume 4, (eds.) J. Vernon Henderson and Jacques-Francois Thisse, Amsterdam: Elsevier, p. 2119-2172.

Stevenson, Betsey. 2006. The Impact of the Internet on Worker Flows. Working Paper, University of Pennsylvania.

Stiroh, Kevin J. 2002. Information Technology and the U.S. Productivity Revival: What do the Industry Data Say? *American Economic Review* 92(5): 1559-1576.

Stock, James H., and Mark W. Watson. 2008. Heteroskedasticity-Robust Standard Errors for Fixed Effects Panel Data Regression. *Econometrica* 76(1), 155-174.

Wellman, B. 2001. Computer networks as social networks. *Science* 29, 2031–4.

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

Table 1a: Descriptive statistics for measures of growth an productivity (for 2000)

Variable	Mean	Std. Dev.	Minimum	Maximum	Number of observations
Log(avg. weekly wage)	6.1531	0.2189	5.4931	7.3330	2743
Log(employment)	9.1898	1.4695	4.3175	15.0820	2743
Log(establishments)	6.6992	1.3143	2.7081	12.5900	2743
Advanced Internet	0.0890	0.1332	0	1	2743
Basic Internet	0.7869	0.4499	0	1	2743
PCs per employee	0.2253	0.1719	0	1.9372	2743

Table 1b: Description of control variables

Variable	Definition	Source	Mean
Internet use at home	Percentage of households with internet at home (2000)	Current Population Survey (CPS) Internet Use Supplement (Census)	0.444
No internet use at home data for county	Dummy indicating no data on home internet use	Current Population Survey (CPS) Internet Use Supplement (Census)	0.9213
Total Population	Total Population as of Decennial Census (1990)	Census	89173.03
% African American	% Population African American as of Decennial Census (1990)	Census	0.0908
% University Graduates	% Population University Graduates as of Decennial Census (1990)	Census	0.1379
% High School Graduates	% Population High School Graduates as of Decennial Census (1990)	Census	0.6996
% Below Poverty Line	% Population Below Poverty Line as of Decennial Census (1990)	Census	0.1622
Median Household Income	Median County Household Income as of Decennial Census (1990)	Census	24492.77
# enrolled in Carnegie rank 1 research university	Per capita number of Students enrolled in local PhD-granting institutions	Downes-Greenstein (2007)	0.0081
# in engineering program	Per capita number of Students enrolled in engineering programs at local Universities	Downes-Greenstein (2007)	0.0010
# patents granted in the country in the 1980s	Total number of patents from inventors located in county, 1980-1989	USPTO	155.73
% professional	% of County's Workforce Employed in Professional Occupations	Census	0.3258
Net Migration	Net migration to county	Census	123.54
% population over 65 years	% of County Population over 65 as of Decennial Census	Census	0.1452

Table 2: Wages increase with internet use

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	No fixed effects	County and year fixed effects	Main specification	MSAs only	No controls for home internet use	Only places where have home internet use data	Years 1991-2001	Falsification test: Years 1990-94	Falsification test: Years 1990-94 MSAs only
Advanced internet	0.5215	0.0370	0.0252	0.0672	0.0253	0.2393	0.0257	0.0105	0.0165
	(0.0481)**	(0.0132)**	(0.0128)*	(0.0364)+	(0.0128)*	(0.1116)*	(0.0141)+	(0.0098)	(0.0341)
Observations	5486	5486	5486	1686	5486	432	5486	5488	1686
(Within) R^2	0.05	0.85	0.86	0.92	0.86	0.95	0.94	0.83	0.90
Fixed effects	No	County	County	County	County	County	County	County	County
Controls	None	Year	All	All	All except home internet use	All	All	All	All

Dependent variable is logged wages. Unless otherwise stated, years are 1995 and 2000. Controls include year dummy, population in 1990, median income in 1990, percent African American in 1990, percent university graduate in 1990, percent high school graduate in 1990, percent below poverty line in 1990, per capita number of people attending Carnegie Type 1 schools in 1990, net migration into the county in 1995, number of patents granted to inventors located in the county in the 1980s, per capita number of students in engineering program in 1990, fraction professional in 1995, percent over 65 in 1990, change in total population between 1990 and 2000, change in percent university graduates between 1990 and 2000, change in percent high school graduates between 1990 and 2000, change in percent African American between 1990 and 2000, change in net migration into the county between 1995 and 2000, and change in percent over 65 between 1990 and 2000. Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 3: Employment and Establishments show no clear pattern of correlation with internet use

Dependent Variable→	*EMPLOYMENT*			*NUMBER OF ESTABLISHMENTS*		
	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
	No fixed effects	County and year fixed effects	Main specification with several further controls	No fixed effects	County and year fixed effects	Main specification with several further controls
Advanced internet	1.2483	-0.0023	-0.0181	1.1220	-0.0026	-0.0031
	(0.2573)**	(0.0201)	(0.0173)	(0.2210)**	(0.0147)	(0.0135)
Observations	5486	5486	5486	5486	5486	5486
(Within) R^2	0.01	0.27	0.44	0.01	0.42	0.58
Fixed effects	None	County	County	None	County	County
Other controls	None	Year	All	None	Year	All

Controls are the same as in table 2. Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 4: Is Advanced Internet different from other measures of IT use?

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*
	No Fixed Effects or controls	Compare all three measures of IT use	Compare Advanced internet and Basic internet	Compare Advanced internet and PCs per employee	Basic internet only	PCs per employee only
Advanced internet	0.0277	0.0232	0.0229	0.0244		
	(0.0413)	(0.0136)+	(0.0134)+	(0.0133)+		
Basic internet	0.5447	0.0127	0.0119			0.0153
	(0.0624)**	(0.0108)	(0.0103)			(0.0097)
PCs per employee	0.0702	-0.0014		0.0022	0.0057	
	(0.0185)**	(0.0078)		(0.0074)	(0.0071)	
Observations	5486	5486	5486	5486	5486	5486
(Within) R^2	0.23	0.86	0.86	0.86	0.86	0.86
Fixed effects	None	County	County	County	County	County
Other controls	None	All	All	All	All	All

Dependent variable is logged wages. Controls are the same as in table 2. Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%

Table 5: Effect primarily occurs in places that are already high income, education, IT-intensity, *AND* population

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*
Advanced internet	0.0150	0.0099	0.0213	0.0191	0.0007	0.0225	0.0029
	(0.0138)	(0.0127)	(0.0130)	(0.0160)	(0.0150)	(0.0130)+	(0.0152)
Advanced internet and	0.0891				0.0430		0.0378
High income county	(0.0367)*				(0.0499)		(0.0502)
Advanced internet and		0.1068			0.0824		0.0796
High education county		(0.0447)*			(0.0557)		(0.0557)
Advanced internet and			0.1903		0.0927		0.0298
High population county			(0.0680)**		(0.0756)		(0.0774)
Advanced internet and				0.0205	0.0198		0.0155
High IT-intensity county				(0.0232)	(0.0238)		(0.0241)
Advanced internet and High income, education,						0.1785	0.1232
IT-intensity, and population county						(0.0530)**	(0.0582)*
Observations	5486	5486	5486	5486	5486	5486	5486
(Within) R^2	0.86	0.86	0.86	0.86	0.87	0.86	0.87
Fixed effects	County	County	County	County	County	County	County
Other controls	All	All	All	All	All	All	All

Dependent variable is logged wages. Controls are the same as in table 2. Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%.

Table 6: Wage growth from 1999 to 2005 is not related to early use of advanced internet

	(1)	*(2)*
Advanced internet	-0.0081	-0.0017
	(0.0136)	(0.0156)
Advanced internet and		-0.0202
High income county		(0.0439)
Advanced internet and		-0.0624
High education county		(0.0695)
Advanced internet and		0.0757
High population county		(0.0791)
Advanced internet and		0.0127
High IT-intensity county		(0.0273)
Advanced internet and High income, education,	0.0003	-0.0040
IT-intensity, and population county	(0.0427)	(0.0471)
Observations	5486	5486
(Within) R^2	0.87	0.87
Fixed effects	County	County
Other controls	All	All

Dependent variable is logged wages. Years are 1999 and 2005. Controls are the same as in Table 2.
Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%.

Table 7: Benefits of early internet use do not spill over to adjacent locations

	(1)	*(2)*
Advanced internet	-0.0018	-0.0406
	(0.0290)	(0.0328)
Advanced internet and		0.0410
High income county		(0.0499)
Advanced internet and		0.0815
High education county		(0.0554)
Advanced internet and		0.0315
High population county		(0.0901)
Advanced internet and		0.0169
High IT-intensity county		(0.0239)
Advanced internet and High income, education, IT-	0.1967	0.1500
intensity, and population county	(0.0539)**	(0.0579)**
Advanced internet and in same MSA as High income,	0.0251	0.0439
Education, IT-intensity, and population county	(0.0294)	(0.0307)
Observations	5486	5486
(Within) R^2	0.86	0.87
Fixed effects	County	County
Other controls	All	All

Dependent variable is logged wages. Years are 1995 and 2000. Controls are the same as in Table 2.
Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%.

Figure 1a: Advanced Internet Use and Wage Growth by County

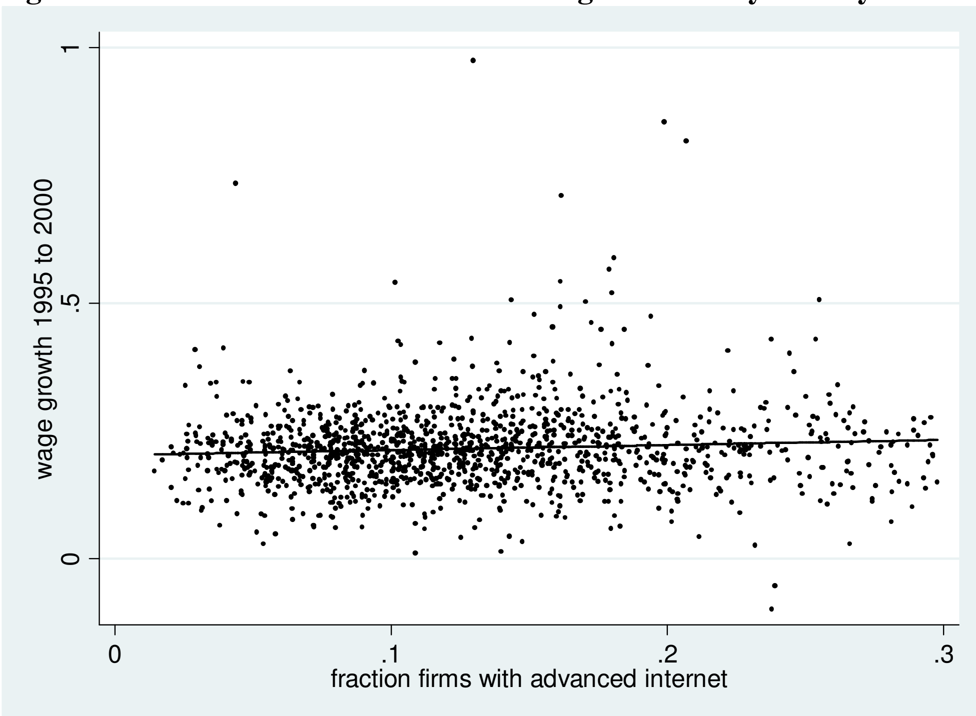

Figure 1b: Advanced Internet Use and Wage Growth by County Type

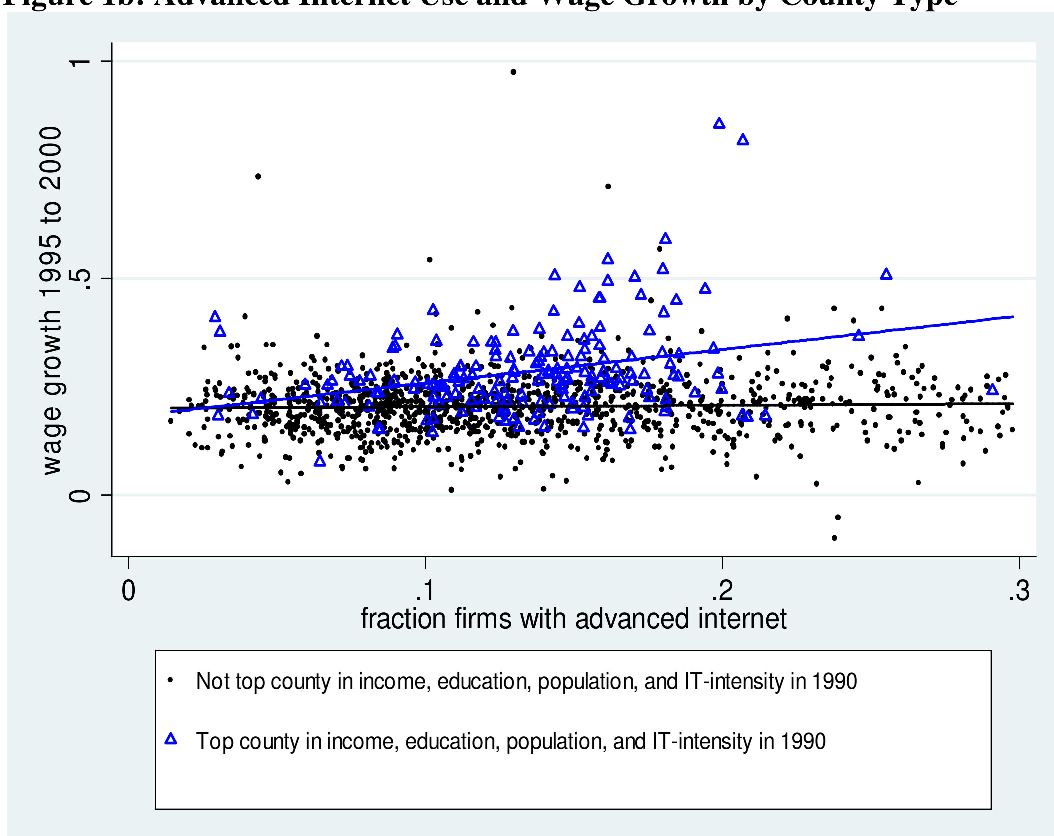

Figure 2: Marginal effect of advanced internet year-by-year

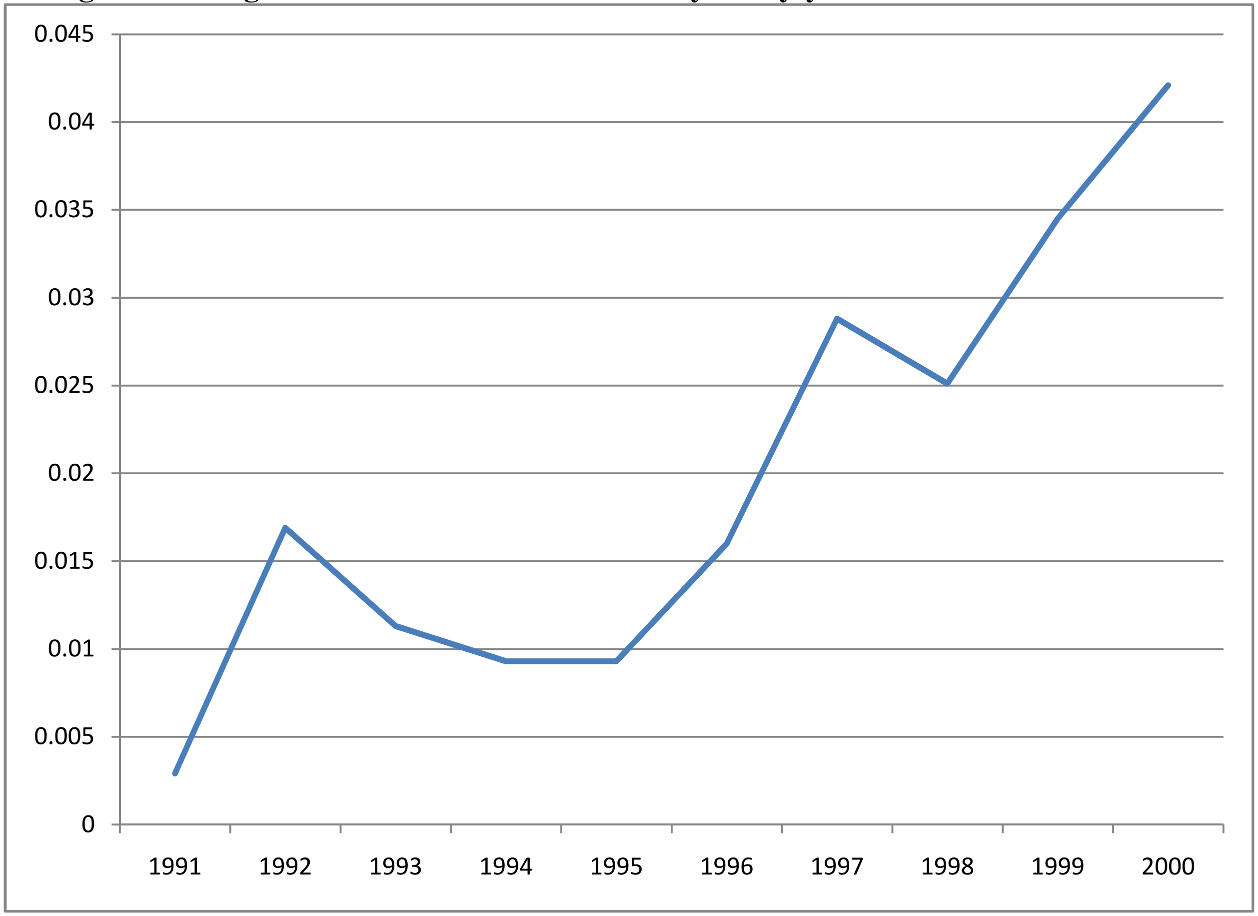

This is based on the regression model is table 2 column 3 except that each year from 1990 to 2000 is included in the regression and a separate effect of advanced internet (as of 2000) was estimated for each year using 1990 as the base. Controls are the same as in table 2.

Figure 3: Marginal effect of advanced internet year-by-year in top counties

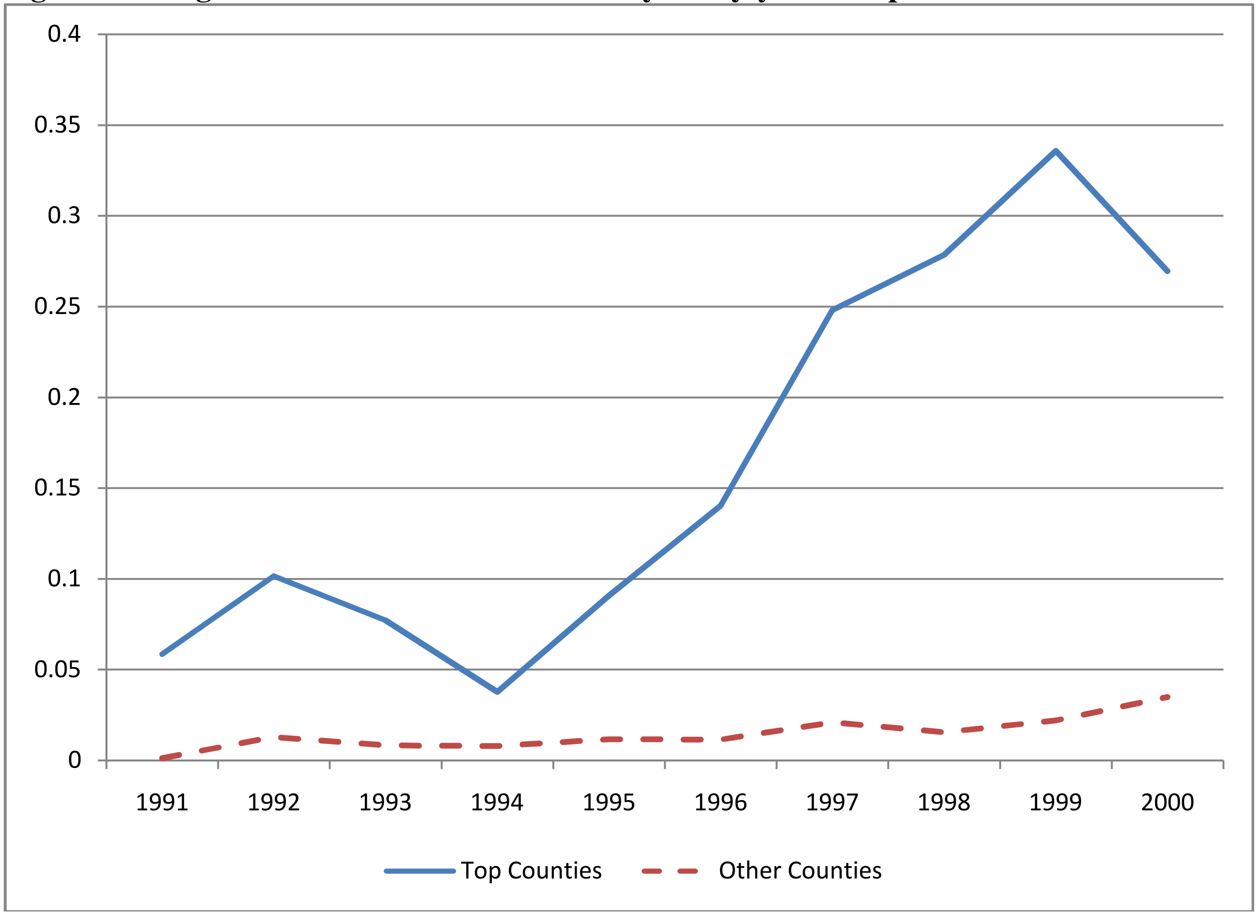

This is based on the regression model is table 5 column 6 except that each year from 1990 to 2000 is included in the regression and a separate effect of advanced internet (as of 2000) and the interaction was estimated for each year. Controls are the same as in table 2.

Online Appendix Table 1: Continuous measures for income, education, IT-intensive industry, and population

	(1)	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*
Advanced internet	-0.0317	-0.0421	0.0022	-0.0134	-0.0500	0.0213	-0.0226
	(0.0463)	(0.0361)	(0.0141)	(0.0301)	(0.0531)	(0.0129)+	(0.0532)
Advanced internet x	2.65e-06				-2.21e-06		-2.56e-06
county-level income	(2.07e-06)				(3.65e-06)		(3.63e-06)
Advanced internet x		0.5662			0.5497		0.5250
county-level education		(0.3025)+			(0.4691)		(0.4693)
Advanced internet x			1.07e-06		9.30e-07		4.27e-07
county-level population			(2.25e-07)**		(2.36e-07)**		(2.41e-07)+
Advanced internet x				0.1684	0.1650		0.1274
county-level IT-intensity				(0.1022)+	(0.1056)		(0.1063)
Advanced internet x income x						1.38e-10	1.03e-10
education x population x IT-intensity						(3.36e-11)**	(3.48e-11)**
Observations	5486	5486	5486	5486	5486	5486	5486
(Within) R^2	0.86	0.86	0.86	0.87	0.87	0.87	0.87
Fixed effects	County	County	County	County	County	County	County
Other controls	All	All	All	All	All	All	All

Dependent variable is logged wages. Controls are the same as in table 2. Heteroskedasticity-robust standard errors in parentheses.
+ significant at 10%; * significant at 5%; ** significant at 1%

Online Appendix Table 2: Further Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Includes two-way interactions	No weighting by time of survey		Dependent variable is wages, not logged		High population defined as top quartile of counties (> 63,657)	
Advanced internet	-0.0035	0.0119	-0.0020	11.23	-1.61	0.0227	0.0035
	(0.0154)	(0.0101)	(0.0096)	(8.35)	(10.79)	(0.0130)+	(0.0153)
Advanced internet and High income, education,	0.1706	0.1933	0.1278	216.92	154.30	0.1609	0.1211
IT-intensity, and population county	(0.0990)+	(0.0576)**	(0.0630)*	(46.19)**	(47.12)**	(0.0521)**	(0.0549)*
Advanced internet and	0.0779		0.0595		23.24		0.0460
High income county	(0.0501)		(0.0450)		(24.64)		(0.0505)
Advanced internet and	0.1145		0.0514		55.38		0.0837
High education county	(0.0672)+		(0.0438)		(32.58)+		(0.0554)
Advanced internet and High population	-0.0974		0.0333		89.66		-0.0363
county	(0.1026)		(0.0827)		(59.90)		(0.0439)
Advanced internet and High IT-intensity	0.0342		0.0269		5.13		0.0170
county	(0.0233)		(0.0186)		(13.24)		(0.0240)
Advanced internet and High IT-intensity and	0.0418						
population county	(0.0701)						
Advanced internet and High education and	-0.0936						
IT-intensity county	(0.0574)						
Advanced internet and High income and IT-	-0.0760						
intensity county	(0.0603)						
Advanced internet and High income and	0.0637						
population county	(0.0717)						
Advanced internet and High education and	0.1016						
population county	(0.0619)						
Advanced internet and High income and	-0.0659						
education county	(0.0725)						
Observations	5486	5486	5486	5486	5486	5486	5486
(Within) R^2	0.87	0.86	0.87	0.70	0.71	0.86	0.87
Fixed effects	County	County	County	County	County	County	County
Other controls	All	All	All	All	All	All	All

Dependent variable is logged wages unless otherwise stated. Time periods are 1995 and 2000. Controls are the same as in table 2. Heteroskedasticity-robust standard errors in parentheses.

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