11-2016

Information Technology and Learning On-the-Job

James Bessen

INFORMATION TECHNOLOGY AND LEARNING ON-THE-JOB

Boston University School of Law
Law & Economics Working Paper No. 16-47

James Bessen
Boston University School of Law

Information Technology and Learning On-the-job

By James Bessen, BU School of Law

10/2016

Abstract: Economists disagree how much technology raises demand for workers with pre-existing skills. But technology might affect wages another way: through skills learned on the job. Using instrumental variables on 9 panels of workers from 1989 to 2013, this paper estimates that workers who use information technology (IT) have wage growth that is about 2% greater than non-IT workers, all else equal, implying substantial learning. This effect persists over time, implying sustained productivity growth from IT. Also, it benefits workers both with and without college degrees. Because many more college-educated workers use IT, college wages grow faster, contributing to economic inequality.

Keywords: information technology, skills, human capital, wage inequality

JEL codes: J24, J31, O33

Thanks to Tim Simcoe and participants at the NBER productivity seminar for comments.

Does information technology (IT) raise the wages of those workers who have skills needed to use this technology? In the skill-biased technical change account, computers complement certain groups of workers, such as college graduates, raising demand and wages for these workers and increasing wage inequality. Card and Dinardo (2002) and Mishel, Schmitt, and Shierholz (2013) argue against this hypothesis as a major cause of inequality.

Aside from their critiques, the skill-biased technical change account has a significant shortcoming: it assumes, implicitly or explicitly, that some groups of workers already have the needed skills; often, "skill" is equated with a college education.[1] However, it might be that the skills needed to work with new technology are not pre-existing, but must, instead, be developed by working with the technology on the job. Indeed, a variety of evidence suggests that computer adoption involves organizational change and substantial investments in new skills, often learned on the job.[2]

This might seem like a minor distinction, but it is important for at least two reasons. First, it implies different policy remedies. If IT adoption increases the college wage premium, signaling an undersupply, then perhaps policy needs to boost college graduation rates (Goldin and Katz, 2004). On the other hand, if critical skills are learned on the job for all users, then it might be better to target policies to increase opportunities for workers who lack college degrees (Bessen 2015).

---

[1] Acemoglu (2002) reviews the literature, however, the review does not mention the possibility that the new skills might be learned on the job.

[2] Bresnahan and Greenstein (1996) find substantial investments in knowledge by firms adopting computers. Some of the learning involves not the technology itself, but new organizational procedures (Bresnahan 1999). Brynjolfsson, Hitt, and Yang (2002) find large investments in organizational capital with computer adoption. Juhn et al. (1993) find that much of the growth in income inequality is not explained by education or other observed worker characteristics. More generally, Abowd et al. (2002) find that education and observed characteristics account for only a small part of human capital. Bessen (2015) reviews historical evidence.

Second, mounting evidence suggests that the main impact of IT on wages comes with experience on the job rather than from increasing the returns to pre-existing skills. In the "canonical" model of skill-biased technical change (Acemoglu 2002), IT augments skilled workers, raising their marginal productivity. Because pre-existing skills are not firm-specific, workers using IT should have higher wages than comparable non-using workers. Yet the best estimates suggest that this wage premium is quite small.

Most of the empirical support for skill-biased technical change comes from studies using aggregate data. While these studies cannot measure the wage premia associated with IT use, a number of papers have used micro-data. The micro studies find that the difference in wages, after controlling for unobserved worker heterogeneity, is small—typically only a few percent—too small to have much impact on wage inequality. Initially, Krueger (1993) regressed log wages against a dummy variable for computer use and control variables, finding a 10-15% wage premium; he concluded that the computer wage premium accounted for as much as half of the growth in the college-high school wage premium. However, his estimates could suffer from sample selection bias—computers might tend to be assigned to workers who have unobserved ability that is also correlated with higher wages.[3] DiNardo and Pischke (1997) found that pencil use is also correlated with higher wages. Since pencils likely do not require specialized skills, the positive coefficient on pencil use suggests substantial selection effects.

Several studies have attempted to correct for selection bias by using longitudinal data from select short panels of worker data, estimating fixed or random effects models. They find evidence of substantial selection bias and only a small wage premium. Entorf and

---

[3] Krueger attempted to control for ability by using data on achievement test scores and school performance.

Kramarz (1997), using a panel of French workers from 1985-7, find that computer users with no previous computer experience earn 6% more on average. Entorf, Gollac, and Kramarz (1999) use a French panel from 1991-3 and find that a worker using a computer without previous experience earns 1% more. Pabilonia and Zoghi (2005) use a Canadian survey from 1999-2002. Using a fixed effects analysis, they estimate a wage premium of 1% for workers without computer experience; using an instrumental variable estimation, their wage premium estimate is not significantly different from zero. Dostie, Jayaraman, and Trepanier (2010) use the same dataset but they use a model with random effects for firms and workers; they find a 4% wage premium exclusive of computer experience.[4]

Other research is consistent with the finding of only a small wage premium for IT use. Borghans and ter Weel (2004) do not control for unobserved heterogeneity, but do control for computer skills, finding no significant relationship between these skills and the computer wage premium. Using a difference-in-differences design for the rollout of broadband in Norway, Akerman, Garder, and Mogstad (2015) find that local wages for skilled workers rise with the availability of broadband, but the effect is small.[5]

It seems hard to reconcile these findings of low IT wage premia with the skill-biased technical change hypothesis. How can a one-time IT wage premium of a few percent explain a rise, say, in the college-high school wage premium of 30% or more since 1980? The skill-biased technical change hypothesis holds that technology augments labor, but, presumably it

[4] One concern with all of these estimates is the short panel duration. With short panels, identification in a fixed effects estimation comes from the small number of workers who switch computer use during the panel, possibly leading to measurement error and imprecise estimates. For this reason, Pabilonia and Zoghi (2005) use instrumental variables and Dostie, Jayaraman, and Trepanier (2010) use a model with random effects. Their estimates are not markedly larger, however.

[5] They estimate that a 10% increase in the availability of broadband in a municipality increases skilled wages there by 0.2%, suggesting that the total effect of broadband use generates a wage premium on the order of 2%. Forman, Goldfarb, and Greenstein (2012) also find that advanced Internet applications are associated with faster local wage growth, however, their estimates do not translate easily to wage premia.

is the workers who actually use information technology who should be the main beneficiaries of higher productivity—computer tools primarily boosts the productivity of the workers who use them; when the Internet improves communication, it boosts the productivity of those workers who use it to communicate. While IT use in one occupation can affect workers in other occupations and industries (see Bessen 2016), the main impact on worker productivity should be measurable among the skilled workers who use these tools. And the longitudinal studies suggest that this impact is small when measured for pre-existing skills.

The picture is very different, however, if we consider skills learned through experience on the job. All of the panel data surveys above asked how many years each worker had been working with computer technology. Although this was not the main focus of these studies, the estimated impact of computer experience in each study is substantially larger than the wage premium associated just with computer use. The estimates of the wage premia associated with computer experience increase 1-2% *per year*.[6]

These findings hint at a significant growth in the returns to IT use over time, however, they suffer from two shortcomings. First, the panels are short and none are very recent, making it difficult to infer much about how this effect might have worked over the last few decades. For example, the most recent survey is from 1999-2002 and respondent's mean experience with computers was only 6 years. The panel regressions also include a quadratic term in computer experience, however, the short panel lengths and the low values

[6] The regressions included computer experience plus computer experience squared, so the annual increases vary. Entorf and Kramarz (1997) report a 1-2% per annum increase in productivity with experience but that lasts a couple years. Entorf, Gollac, and Kramarz (1999) report a total 2% gain over 1-3 years. Pabilonia and Zoghi (2005) obtain estimates of 1.2% and 2.0% productivity growth with computer experience in the first years. Dostie, Jayaraman, and Trepanier (2010) estimate returns to experience grow 1.0-1.6% per year during the first years.

of computer experience make extrapolation unreliable. Furthermore, these studies do not control for the selection of computer users and applications. Early computer users ("early adopters") are very likely different from those workers who begin using computers later on, using different applications. Inferences about the time pattern of wage growth based on self-reported computer experience in cross-sectional data might therefore be misleading.

The second problem is that wages might increase with job tenure for reasons other than growing computer skills. If these factors correlate with IT use, then estimates of the returns to IT experience will be biased even in fixed effects and random effects models.[7]

Measuring the growth in IT wage premia

This paper estimates the growth of wage premia for IT users by employing instrumental variables and data that cover a much longer time span: it uses 9 two-year longitudinal panels that cover computer and Internet use from 1989 through 2013. The data come from the Current Population Survey (CPS), which asked questions about computer and Internet use in a number of supplemental surveys. Although the CPS was designed primarily for cross-sectional research, it is possible to track individual workers and their wages over two years (Drew, Flood, and Warren 2014).

Ideally, one could measure growth in an IT wage premium by regressing the change in wages against IT use. However, several other delayed payment mechanisms might also generate a rising wage-tenure profile including: self selection of prospective employees (Salop and Salop 1976), insurance for risk adverse employees when productivity is uncertain (Harris and Holmstrom 1982, Freeman 1977), and incentives to reduce shirking when productivity is costly to monitor (Lazear 1979). These mechanisms all imply that wages will

[7] Pabilonia and Zoghi (2005) use instrumental variables but do not instrument computer experience.

rise faster than worker productivity with tenure. To the extent that these mechanisms might be correlated with IT use—and there are good reasons to think this is the case—then the estimated effect of IT on wage growth will be biased.

To obtain unbiased estimates, I need to distinguish between wage growth that arises from worker productivity growth and wage growth that exceeds productivity growth.[8] This can be done with an instrumental variable that is correlated with IT use, but is uncorrelated with delayed payment mechanisms that raise wages faster than productivity. Below I select an instrumental variable based on institutional considerations and I validate that this variable is uncorrelated with an indicator of delayed payment. Although this variable is not randomly assigned, the institutional analysis supports its validity as an instrument.

The instrumental variable estimates for wage premium growth are substantial, a 1.3%-2.6% increase per year in real hourly wages for computer users and Internet users. These estimates are in line with previous research, however, this paper finds that this wage growth does not appear to diminish substantially over time. The implication is that large costs are associated with learning to use IT. It also suggests that the critical skills are not those associated simply with using computers or the Internet per se, but rather skills associated with the ongoing stream of software, system improvements, and associated organizational changes.

Substantial growth in the IT wage premium also has significant consequences for wage inequality because access to IT is highly unequal. A firm's decision to assign a computer system to a given worker is a classic technology adoption problem (Hall and Khan 2003). Theoretically, firms assign IT to workers when the productivity payoff exceeds the

[8] I also distinguish between growth in productivity associated with IT and growth in the returns to other measured and unmeasured variables that might be correlated with IT.

adoption cost. All else equal, more highly skilled workers will have larger payoffs to adoption so that skills should be correlated with IT adoption. Estimates of probit equations for computer and Internet use confirm this hypothesis. However, this means that IT-related learning will tend to exacerbate pay differentials. For example, because relatively more college-educated workers use IT than workers with just a high school education, the wages of college-educated workers grow faster as a group. In fact, on the job learning of IT-specific skills can account for the majority of the growth in the college-to-high school wage ratio since 1980.

Note, moreover, that this wage gap grows because more college-educated workers are given the opportunity to use IT, not because a college diploma is "required" to use IT. The wages of non-college workers grow at the same rate or faster if they, too, are assigned computers and Internet. In this account, rising inequality does not signal an undersupply of college graduates; rather it suggests that the costs of learning new skills or other adoption costs may be holding back wage growth for many workers.

Model

Skills and productivity

To explore IT adoption and wage growth, it is helpful to begin with a model of productivity. Consider the marginal productivity of worker i at time t. Let X represent measured worker characteristics, let U represent unmeasured characteristics that are correlated with wages, and let $C = 1$ if the worker uses IT and $C = 0$ if not. We observe different cohorts of workers, each at time, t, each cohort observed for two periods, $T = 0, 1$. Their productivity might change with experience. Then the worker's marginal productivity can be written

(1) $$\ln V_{it} = \alpha_t X_i + (\beta_T + \delta_t) C_i + \gamma_t U_i + \epsilon_{it},$$

where δ_t captures IT-related changes in productivity over cohorts, β_T captures IT-related productivity changes with experience/learning, and ϵ_{it} represents random errors that are unrelated to worker characteristics, such as random measurement error in *V*.

The two time dimensions—*t* and *T*—capture different ways that skills can change with IT use:

- If IT makes workers more efficient but without requiring any new skills, then firms will adjust employment so that marginal productivity equals the market wage. Then, equilibrium marginal productivity will remain unchanged, $\beta_T = \delta_t = 0$.
- If workers who use IT are more productive regardless of their experience with the technology—that is, IT complements pre-existing skills—then δ_t will be positive.
- If IT users learn on the job, then β_T will increase with worker's time on the job. This can be seen by reference to standard human capital models. On-the-job learning can be realized through investments training, mentoring, or learning-by-doing. In the classic human capital model (Becker 1964), these investments take the form of training costs or the opportunity costs of lost production. Figure 1a illustrates the productivity of a worker who begins learning new skills with the use of new technology in her fourth year on the job. That year productivity drops because of training costs or production lost while learning, but it is more than offset by gains the following years. The difference between productivity in years 5 and later compared to the early years represents the return on the human capital investment. In this case, $\beta_4 < 0 < \beta_5$.

Many information technologies improve sequentially. New versions of software or entirely new products are run on new and improved hardware. If such ongoing improvements each require some new skills, then the pattern would be repeated although perhaps with somewhat smaller magnitudes. This is illustrated in Figure 1b where β continues to increase after the initial year of IT use.

With both a single innovation and sequential innovation, β increases over tenure on average, despite some down years, since the net return to human capital is positive. The data in this paper capture two consecutive years labeled 0 and 1, but the data lack information on the number of prior years the worker has used IT. In effect, for IT users, the data capture an average over the years with IT use. And so, on average, if IT users learn on the job, then they will have

(2) $$\beta_0 < \beta_1,$$

especially when IT innovation is sequential. This will not be the case for non-IT users and for pre-existing skills.

Wages

In a companion paper we will address the degree to which worker skills are firm specific or general. Abstracting away from these issues, I assume that in each period workers and firms split the value of the worker's output,

(3) $$w_{it} = \theta_T V_{it}, \qquad \pi_{it} = (1 - \theta_T) V_{it} \text{ where } 0 < \theta < 1$$

and θ might change with *T* for reasons explained below. Then

(3a) $$\ln w_{it} = \alpha_t X_i + (\beta_T + \delta_t) C_i + \gamma_t U_i + \ln \theta_T + \epsilon_{it} \text{ and}$$

(3b) $$\ln \pi_{it} = \ln(1-\theta_T) + \alpha_t X_i + (\beta_T + \delta_t) C_i + \gamma_t U_i + \epsilon_{it}.$$

Adopting information technology

First, consider the adoption decision, that is, whether the firm chooses to assign IT to worker *i*. Suppose that there is an adoption cost, *A*. This cost might reflect the difficulty of developing an application for the given occupation or the firm's cost of training.

The firm will assign IT to worker *i* if the total profit over periods 0 and 1 exceeds *A*. In the simplest case, assume that the returns to skill are constant over time so that adoption will occur if (suppressing the unchanging time subscripts and the error term)

(4) $$\pi_{i0} + \pi_{i1} = (1-\theta) e^{\alpha X_i + \gamma U_i} \left(e^{\beta_0} + e^{\beta_1}\right) > A.$$

This equation shows that workers with high *U*—unmeasured ability or other unmeasured characteristics that are correlated with wages as in (3)—should also be more likely to be assigned to work with information technology, all else equal.

It is helpful to express that correlation as an explicit linear relationship, $U_i = C_i + \mu_i$, where *U* is normalized so the coefficient of *C* is 1. Then, redefining the error term, (3a) can be written

(5) $$\ln w_{it} = \alpha_t X_i + (\beta_T + \delta_t + \gamma_t) C_i + \ln \theta_T + \epsilon_{it}.$$

Expressed this way, the coefficient of *C* in an OLS regression will clearly be a biased estimate of the returns to computer use. Probit estimates of IT use below provide evidence that unmeasured characteristics that correlate with wages are also correlated with IT use.

Learning effects

However, the goal of this paper is to estimate $\Delta\beta \equiv \beta_1 - \beta_0$, rather than the total return to IT use. One could estimate a first-differenced version of (5),

(6) $$\Delta \ln w_i = \Delta\alpha \cdot X_i + (\Delta\beta + \Delta\delta + \Delta\gamma) \cdot C_i + \Delta \ln \theta + \omega_i, \quad T = 1.$$

This approach has two problems. First, the coefficient of C clearly includes more than $\Delta\beta$. The term $\Delta\delta + \Delta\gamma$ represents the change in returns to pre-existing computer skills and to unmeasured skills. Below I develop separate estimates of these. The second problem is that $\Delta \ln \theta$ might be correlated with IT use. Since we have no independent measure of this variable, it cannot be included. Estimating (6) without it gives rise to omitted variable bias. Below I develop instrumental variable estimates to overcome this limitation.

Data and variables

The main data source for this paper is the Current Population Survey (CPS) of the Bureau of Labor Statistics. In select years, a supplemental survey asked whether the respondent "directly used a computer at work;" in other years, the survey asked "do you access the Internet from work?"[9] The universe for these questions was employed persons of 15 years of age or older. I use these survey items as the measure of information technology use. Because these measures are self-reported, they may undercount cases where information technology is embedded in hardware. For example, radiologists might not call digital X-ray machines computers; similarly, cashiers might not call computerized check out terminals

[9] Computer use was surveyed in October of 1984, 1989, 1993, 1997, and 2003 and in September of 2001; Internet use was surveyed in December 1998, August 2000, and July 2011, 2013, and 2015. The 1984 and 2015 data were not used in this study because wage data is not available for 1984 and the longitudinal panel for 2015 was incomplete at this date.

computers. Nevertheless, reported levels of IT use are quite high, suggesting that most IT use is reported.

Also, these measures provide only the most basic indicator of technology use. Workers who use IT at work can be applying a wide range of software and systems that have varying impacts on worker productivity. IT use may also involve substantial organizational changes. Moreover, these systems likely changed dramatically between 1989 and 2013. The coefficients estimated on the IT variables are thus at best crude averages across different types of systems and organizations. Forman, Goldfarb, and Greenstein (2012) argue that most of the wage benefit from the Internet comes from advanced applications, not general use. If so, my estimates understate the impact of advanced technologies and the associated measurement error might attenuate coefficient estimates.

For the basic data set, I use the sample of all workers who responded to the IT questions aged 15 through 65, excluding imputed responses. This sample includes 318,547 computer users and 216,300 Internet users. Figure 2 shows the share of IT users over time.

Estimating equation (6), however, requires longitudinal data that can be differenced. Although the CPS was not designed as a longitudinal survey, I am able to construct short longitudinal panels using the method of Drew, Flood, and Warren (2014) and implemented in IPUMS (Ruggles et al. 2015). The CPS surveys households for four months and then, after a hiatus of eight months, they are surveyed for another four months. Wage data are collected in the fourth and final months (the so-called outgoing rotation groups). For each worker in my panel, I include two or three months: the two outgoing rotation months and the month that the information technology question was asked (might also be an outgoing rotation month). For this analysis, I necessarily exclude workers who do not have two wage observations or have wage data allocated. This will exclude workers who relocated during the year, who were unemployed, or who did not respond for other reasons. The wage panels consist of 109,154 observations for computer use and 60,940 observations for Internet use.

Table 1 presents sample summary statistics. The panel samples are substantially smaller and have somewhat higher IT use, raising the possibility that the panel sample might not be representative of the entire population. Below I explore sample selection issues.

Occupational characteristics, including the instrumental variable, come from the Dictionary of Occupational Titles (1977). The US Department of Labor has sought to define aspects of some 14,000 distinct jobs; England and Kilbourne (2013) have mapped these to Census detailed occupation codes, averaging them to this higher level of aggregation. One characteristic is STRENGTH, which rates the physical demands of the job on a scale of 1, for sedentary occupations, to 5, for very heavy work.

Identifying the Learning Effect

Trends in productivity coefficients

In order to use (6) to estimate $\Delta\beta$, it is necessary to estimate $\tau \equiv \Delta\delta + \Delta\gamma$, which represents the change in returns to pre-existing computer skills and to unmeasured skills. Because my data covers different cohorts over decades, the average trend growth in these terms can be estimated by using a version of the levels equation (5) with a trend term. Assuming that these trends grow at a roughly constant rate, we can write

$$\delta_t = \delta_0 + \overline{\Delta\delta} \cdot t, \quad \gamma_t = \gamma_0 + \overline{\Delta\gamma} \cdot t$$

so that equation (5) can be re-written

(7) $$\ln w_{it} = \alpha_t X_i + (\delta_0 + \gamma_0)C_i + \tau \cdot C_i \cdot t + \rho \cdot I(T = 1) + \epsilon_{it}$$

where I capture the tenure-related changes with a simple dummy variable here to focus on the long-term trends.

Table 2 shows regressions of this equation for computer use (left) and Internet use (right). The samples are the pooled (T=0,1) observations of the wage samples. Time is converted from years and months to fractional years. Columns 2 and 4 add trend interactions for education and age (coefficients not shown). The first row shows the coefficient of *C*, which, as discussed, represents an upwardly biased estimate of the returns to IT use. The second row estimates τ, the average annual increase in $\Delta\delta + \Delta\gamma$, the returns to pre-existing and unmeasured skills. The estimates are about 0.2% per year. As we shall see, this figure is an order of magnitude smaller than the estimates of $\Delta\beta$ obtained below.

Rising wages, job tenure, and IT use

A second hurdle to estimate learning effects using equation (6) is the possibility that $\Delta \ln \theta$ is correlated with IT use, creating an omitted variable bias. A significant literature provides reasons why firms might increase $\ln \theta_T$ with tenure and these reasons might be correlated (perhaps negatively) with IT use. In particular, the major theoretical explanations for why wages might rise faster than worker marginal productivity depend on some sort of incomplete information about worker productivity that is revealed over time. Yet worker productivity might be easier to measure in those occupations that tend to use IT, making the error term negatively correlated with IT use. There is, indeed, evidence of such as association.

It is helpful first to look at several major theoretical explanations for why wage growth might diverge from productivity growth. In Lazear's model of shirking (1979), employers have difficulty monitoring worker productivity, although productivity is revealed over time. By paying workers less than their marginal productivity during their early years on the job and more than their marginal productivity later, workers will not shirk during the early years because they risk forfeiting their large future payout. In Salop and Salop (1976), rising wages act as a screening device to weed out less productive prospective employees when their productivity is private information. Again, productivity is revealed after some time on the job. In Carmichael (1983), wages diverge from marginal productivity when employers have private information on productivity in a model with firm specific training. In Freeman (1977) and Harris and Holmstrom (1982), productivity information is symmetric but unpredictable; workers and firms learn a worker's productivity over time and rising wages act as a kind of insurance policy for risk-averse workers.

In all of these cases, a delayed payment mechanism is used because information about worker productivity is revealed over time. Yet one might expect occupations to differ in the degree to which productivity can be easily observed and/or predicted. Hutchens (1987) first proposed that occupational characteristics might be related to occupational differences in monitoring. He proposed specifically that occupations that involve repetitive tasks are easier to monitor and therefore should be negatively correlated with features of delayed payment contracts such as longer job tenure, mandatory retirement, pensions, and higher wages for more experienced workers. He found evidence to support those hypotheses using data for 1971.

Interestingly, some other research suggests that repetitive work is associated with computer use. Autor, Levy, and Murnane (2003) find an association between computer use and the degree to which an occupation involves routine tasks. This means that computer use is likely negatively correlated with delayed payment. If so, then the error term in (5) is likely to be negatively correlated with IT use, biasing estimates of the learning effect.

Validating an instrumental variable

Instrumental variable estimation can correct for this bias, in particular, one using an instrumental variable that is correlated with IT use, but uncorrelated with delayed payment. Hutchen's insight on the link between delayed payment and occupational characteristics such as job tenure provides a means for identifying occupational characteristics that are uncorrelated with delayed payment and which, therefore, might serve as instruments.

I begin by re-creating Hutchens's result using current data for occupations from the CPS and measures of job characteristics taken from the 1977 edition of the Dictionary of Occupational Titles. While legal changes have largely eliminated mandatory retirement rules and have changed pensions to include employee-financed plans, job tenure still provides a useful indicator of delayed payment. If wages rise above marginal productivity, more experienced workers will have above-market wages, this should make them less likely to quit, and they should therefore exhibit longer job tenure on average.

Table 3 shows several occupational characteristics grouped by their anticipated correlation with productivity monitoring. Two characteristics that would seem to make productivity measurement easier are the importance of repetitive tasks and routine tasks; two characteristics that might make monitoring more difficult are abstract and creative activities and jobs that involve making "generalization, evaluations, and decisions based on sensory or judgmental criteria." Worker productivity is likely easier to measure for repetitive and routine tasks and harder to measure when output is abstract, creative, or depends on individual judgment.

The correlations shown in Table 3 verify these interpretations of occupational characteristics, showing correlations with job tenure.[10] Repetitive occupations are, indeed, negatively and significantly correlated with job tenure, as Hutchens found. Highly routine jobs also tend to have shorter job tenure. On the other hand, harder to monitor occupations have longer tenure, suggesting that delayed payment is used as an incentive or insurance mechanism in occupations with these characteristics.

The last row of Table 3 shows the variable STRENGTH, which measures the degree of exertion required in the occupation, ranging from sedentary to heavy work. I propose this as an instrumental variable because sedentary occupations are obviously easier to adapt to use of desktop computers. For this reason, it turns out that STRENGTH is negatively correlated with computer use (correlation coefficient of -.63).

But is this variable uncorrelated with delayed payment? There is little reason to expect that worker productivity can be more readily determined for sedentary occupations than for occupations that require standing or physical exertion. Perhaps workers who don't sit are more mobile and therefore harder to observe. However, a test rejects the significance of job mobility.[11] Moreover, IT use is not likely to influence this measure because the 1977 edition of the Dictionary of Occupational Titles largely predates widespread computer use.[12]

The bottom row of the table verifies that the STRENGTH variable is, in fact, uncorrelated with job tenure, implying that it is not correlated with the use of delayed

---

[10] The job tenure data come from a supplement to the Current Population Survey from February 1996, 1998, and 2000, and January 2002, 2004, 2006, and 2008 including civilians in the labor force. I excluded later years because the recession of 2008 might distort mean job tenure. I calculate a weighted average of job tenure for each detailed occupation in each year.

[11] I ran the regressions in Table 4, columns 3 and 6, excluding workers in outdoor occupations; the estimates changed little.

[12] The current paper use the 1977 edition; a future revision will use the 1965 revision.

payment mechanisms. Regardless of the particular mechanism responsible for an increase in $\ln\theta$ with tenure, such an increase should be associated with greater average job tenure; wages that rise faster than productivity discourage quitting. Because STRENGTH is *not* correlated with job tenure, it appears that this variable is independent of $\ln\theta$ and is therefore a valid instrumental variable.

Estimates of Learning Effects

Basic estimates

Table 4 shows regressions of equation (6), for the computer use on the left side and for Internet use on the right. Columns 1 and 4 show a simple uninstrumented regression. Because the dependent variable, the one-year change in log wages, likely contains significant measurement error, I use Least Absolute Deviation estimates to temper the effect of outliers. The estimates of the learning effect are 0.6% and 0.4% for computer users and Internet users respectively (standard errors 0.1% and 0.2%). These estimates are statistically significant, but are notably smaller than those from the literature using fixed effects or random effects models. This could be because differencing increases measurement error in the dependent variable, attenuating coefficient estimates (Griliches and Hausman 1986).

As noted above, the wage sample includes only about one third of the workers who report IT use. From the summary statistics (Table 1), the wage sample appears to be slightly better educated and older. It may be that workers with low wage growth are more likely to move or to drop out of the labor force and hence drop out of the wage sample. To test for sample selection bias, I used a Heckman model (not shown). For the selection equation, I used 5 education dummy variables, 6 age dummy variables, gender, race, and part time status. A Wald test for both computer use and Internet use strongly rejected the null hypothesis that the regression was independent of selection. The estimates of the learning effect were 1.1% for computer use and 0.6% for Internet use, both slightly larger than the comparable estimates in columns 1 and 4. Sample selection bias appears to be small and to bias the estimates downward.

Columns 2 and 5 show basic instrumental variables regressions.[13] For the remaining regressions, I trimmed the sample of 5% tails in the dependent variable to reduce the influence of extreme outliers and mismeasurement. The lower part of the table shows the first stage regression coefficient of STRENGTH. This variable is strongly correlated with IT use. Wald tests reject the null hypothesis that the error term is uncorrelated with IT use. In general, the estimates for Internet use are less precise than those for computer use. As the analysis above suggested, IT use is negatively correlated with unmeasured selection variables, ρ being the correlation coefficient. Thanks to the negative correlation, the IV estimates of the learning effect are larger than the LAD estimates and more in line with those obtained in the fixed effects and random effects models. For computer use the estimate is 1.9% wage growth per year and for Internet use, 2.4% per year.

Columns 3 and 6 add additional controls for education, age, gender, race, whether the worker is in a union or covered by a union contract, and part-time status. With these the learning effect estimates are 1.5% per year for computer use and 2.8% per year for Internet use (standard errors of 0.4% and 0.8% respectively). Taking into account the 0.2% contribution from trend growth in the returns to pre-existing and unmeasured skills, τ, the best estimates for the learning effect are 1.3% annual real wage growth for computer use and 2.6% for Internet use.

---

[13] I use Stata's etregress routine. This has the advantage of handling a binary treatment variable, IT use; results using standard IV regression were similar.

Interactions

These estimates are in line with the estimates reported by studies using fixed or random effects models. However, those studies employed short duration panels and had only a limited ability to infer longer time trends in wages. Table 5 explores long-term trends and other characteristics of learning effects by interacting the IT use variable with other, categorical variables.

The top section of the table shows interactions with the panel year, computer use on the left, Internet use on the right. The table only shows coefficients for the interaction terms; the other control variables are those in columns 3 and 6 of Table 4. While the time interaction coefficients appear variable and the individual year estimates imprecise, there is no clear trend in IT use, although the coefficients for Internet use are mostly higher than those for computer use. Given that many workers who used a computer in 2003 also used one in 2001 and 1997, etc., this suggests that workers continue to learn new productivity-enhancing skills perhaps as those computers are running new and more advanced software applications and using new, improved hardware.

This finding might seem contrary to Entorf and Kramarz (1997) and Entorf, Gollac, and Kramarz (1999) who see wage growth only lasting for the first few years of self-reported computer experience.[14] However, their finding might be misleading because the payoff to IT applications might change over time, especially during the early years of adoption. If so, the apparent slowing of wage growth with computer experience in cross-section might not reflect longitudinal time trends.

---

[14] Pabilonia and Zoghi (2005) and Dostie, Jayaraman, and Trepanier (2010) do not find such a dramatic decline wage growth with computer experience.

While my data do not capture individual worker's computer experience, we can make inferences about differences between newbie IT users and more experienced ones based on whether most workers in the occupation use IT or not. That is, a higher portion of IT users will be first-time users in occupations where few workers use IT; in occupations where most people use IT, most IT users will be experienced. The second panel in Table 5 shows the wage growth coefficient interacted with a dummy variable that is 1 if over half the workers in the occupation use IT. Both groups of occupations exhibit similar learning effects. This is consistent with the view that new IT technologies are requiring new skills on an ongoing basis.

The next panel interacts IT use with worker age. Older workers do not exhibit significantly slower wage growth. Conditional on their being assigned IT equipment, older workers apparently learn IT applications as quickly as younger workers.

The fourth panel explores whether more highly educated workers show greater wage growth. It is sometimes argued that college educated workers are better able to learn on the job and might therefore exhibit higher wage growth. These regressions interact IT use with a dummy variable that is 1 if the worker has some post-secondary education. The estimates show that more educated workers do not have larger learning effects; their wage growth with IT tenure is slightly lower, in fact, although not significantly so. Conditional on using IT, high school educated workers appear to learn just as well on the job.

The final panel explores whether learning might be particularly strong in occupations that are directly involved with computer, including computer programming, engineering, and mathematical occupations.[15] Results suggest that computer occupations do exhibit stronger

---

[15] Specifically, I create a dummy variable that is 1 if the 1990 occupation is between 43 and 68 inclusive, excluding 66, 229, 213-216, 233, and 308-9.

learning effects, particularly for Internet applications, although the difference is not statistically significant.

Who uses information technology?

It is sometimes asserted that younger or better-educated workers are better at learning new technologies. The findings in Table 5 suggest that these groups of workers are not better able to acquire IT-related skills than older and less-educated workers, conditional on those workers being assigned to use IT. However, there are major differences in IT use across different groups of workers. This is not surprising because the analysis of technology adoption implies that the payoff to IT use may be greatest where the opportunity is greatest, that is, among more highly paid workers.

Table 6 explores IT use across different groups. Columns 1 and 4 show mean IT use across all sample years for computers and Internet respectively. Columns 2 and 5 show the marginal effects, evaluated at sample means, from a Probit estimation on computer and Internet use. More educated workers are more likely to use IT as are white, native, female, and fulltime workers. Teenage workers and workers over 60 are less likely to use IT.

With the exception of gender and retirement age workers (where the time horizon to capture returns on human capital investment might be limited), characteristics associated with higher wages also tend to have higher IT use. Some of these differences might be driven by unobserved worker characteristics, *U*, in equation (4), that might also be correlated with observed characteristics. Columns 3 and 6 include log real wages as an independent variable. However, because IT use is associated with wage growth, log real wages will be correlated with the error term. Consequently, I use an instrumental variables probit in these columns. I instrument each worker's wage with the quartile of the mean occupational wage for that worker's occupation in 1980 (and year dummies). The findings imply that unmeasured characteristics that affect wages are significantly correlated with IT use as well.

Discussion

My estimates are in line with estimates of the general rise of wages with job tenure. After controlling for general changes in wage levels over time, Topel (1991) estimates that wages increase 2.5% for each additional year of tenure; Altonji and Williams (2005) estimate a rise of 1.1%. These are estimates for all workers, not just IT users, and these estimates include wage increases associated with delayed payment mechanisms, where my estimates use instrumental variables that are independent of delayed payment. Nevertheless, the similarity suggests my estimates are not unreasonable.

A variety of comparisons imply that the estimated learning effects correspond to a significant investment in human capital. Assuming that the increase in wages with IT experience can be attributed to returns on human capital, the absolute magnitude of the increase should be at least as large as the human capital investment. That is, the returns to human capital should be at least as large as the investment. Since the increase in pay represents the worker's share of the returns and since the worker's share is less than the total returns, the absolute increase in pay provides a lower-bound estimate for the investment in human capital associated with IT investment. For the 2013 cohort of the wage sample, the mean wage of Internet users is $24.11. Assuming 2080 hours per year and a 1.0% increase in wages associated with learning,[16] this implies an annual human capital investment of about

[16] Taking the 1.2% estimate for 2013 from Table 5 less 0.2% for the trend estimated in Table 2.

$500 per IT user; using the alternative estimate of 2.6%,[17] the human capital investment is $1300 per year.

These figures roughly equal the annual amount of gross private investment in computers and peripheral hardware, which is about $800 per IT user.[18] Investment in software is substantially larger, about $2970 per Internet user per year. These comparisons suggest that human capital investment comprises a significant, but not predominant, component of total investment in information technology.

The investment in IT skills learned on the job appears quite large compared to formal firm investments in training. A survey of employers that train conducted by the Association of Talent Development (2014), a human resources trade organization, found that these firms spent $1,208 per year per employee. However, most employees do not receive formal employer-funded training during any given year. The Council of Economic Advisers reports that in 2008, only 11.2% of employees received training paid for by the employer and 8.4% received formal on-the-job training (CEA 2015). By comparison, 63% of workers in 2013 used the Internet. So although the amounts spent per worker are comparable, many more workers are involved with informal learning on the job regarding IT than are involved with formal job training. Given 99 million Internet users in 2013, total investment, total human capital investment comes to about $50 - $130 billion per year. This is comparable to total Federal spending on higher education, $75.6 billion (Pew 2015).

---

17 Using the 2.8% estimate from Table 4 less the 0.2% trend estimate from Table 2.

18 From BEA NIPA accounts, Table 5.3.5 Private Fixed Investment by Type, annual private nonresidential investment in computer and peripheral hardware is $79.1 billion; from the CPS, there were 99.2 million Internet users at work in 2013. $79.1b/99.2m = $797. Annual gross private nonresidential investment in software is $294.6b, or $2970 per user. Internet users understate the total number of IT users.

The wage increases associated with IT use are also significant. In particular, because the use of IT is distributed unequally across different groups of workers, these increases affect wage inequality. For example, one common measure of inequality is the college-high school ratio, the ratio of the mean wage of workers with four years of post-secondary education to the mean wage of workers with only 12 years of schooling. Using the CPS outgoing rotation groups for 1980 and 2013, excluding self-employed workers, this ratio grew by 0.73% per year. Part of this rise could be explained by the rise of wages by IT users. Because more college graduates than high school graduates are IT users—39% more, see Table 6—the rise of wages with IT use will contribute to a growing gap in mean wages between college and high school workers. All else equal, an annual 1.3% increase in wages for IT users will increase the college-high school wage gap by 1.3% x 39% = 0.51% per year.[19] In other words, the rise of wages associated with IT use could account for most of the rise in the college-high school wage gap. This back-of-the-envelope calculation does not take into account how wages change when workers switch jobs. Nor does it take into account a variety of other factors, including relative supply and demand. Nevertheless, this crude calculation suggests that on-the-job learning provides a substantial tailwind toward increasing wage inequality.

Conclusion

Workers who use information technology tend to experience rising wages, relative to non-users and after controlling for a range of observed variables and instrumenting IT use.

---

[19] Suppose $\ln w_t^j = \ln w_0^j + .01 \cdot s^j \cdot t$, where s^j is the share of IT users among workers of education *j* = high school, college. Then the rate of change of $\ln w_t^{college} - \ln w_t^{high\ school}$ equals $.013 \cdot (s^{college} - s^{high\ school})$. I use 1.3% by taking the 1.5% estimated from column 3, Table 4 and subtracting 0.2% trend growth from Table 2.

These increases are substantial—1.3 to 2.6% per year—and they have been persistent since 1989, even in occupations that already have high levels of IT use. These findings suggest that IT requires valuable skills that are learned on the job

Much discussion of technology and skills has focused on pre-existing skills such as those skills acquired through formal schooling. But previous research has only found a small wage premium associated with computer use exclusive of experience, suggesting that computers do little to enhance the productivity of workers without computer experience. In contrast, this paper finds a much larger impact of IT use over time on the job. Moreover, workers without a college diploma benefit from substantial wage gains as much as college-educated workers when they use IT. These findings suggest that IT does complement skills, but the key skills are learned on the job rather than prior to employment.

This difference has important implications for policy regarding training and skills development and also for several other issues. For example, it suggests a different view of the effect of IT on wage inequality. The skill-biased technical change hypothesis sees IT as increasing the relative demand for college graduates. But regardless of whether IT raises the demand for workers with college diplomas, a much higher proportion of college graduates are assigned to work with IT, so their wages rise faster thanks to this greater opportunity.

Also, the sustained pattern of wage growth through 2013 suggests that the impact of IT on productivity might not have "lost steam" after 2004 as some observers have suggested (Gordon 2016, Fernald 2014). IT spending surged during the late 1990s and subsequently declined, leading some observers to infer declining productivity benefits from IT. But the wages of workers using IT have continued to rise at a roughly steady rate, suggesting these workers are experiencing productivity improvements at the roughly same rate as in the past.

Further research is needed to understand the timing of the link between IT investment, IT-related human capital investment, and the returns to these investments.

Also, depending on the degree to which IT skills are firm-specific, large human capital investments in IT-related skills have important implications for worker mobility and employment dynamics, possibly causing a skills mismatch and slow job growth (Restrepo 2015). Labor force participation might also be affected if many older workers' skills become obsolete.

The key question here is how new technology interacts with worker skills to unlock productivity benefits and wage increases. The answer is important for a host of issues today and also to understand the impact of future technologies.

References

Abowd, John M., Paul A. Lengermann, and Kevin L. McKinney. 2002. "The Measurement of Human Capital in the U.S. Economy." Longitudinal Employer-Household Dynamics Technical Papers 2002—09, Center for Economic Studies, U.S. Census Bureau, 2002.

Acemoglu, Daron, "Technology and the Labor Market" Journal of Economic Literature, 40, (2002), 7-72.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. "The Skill Complementarity of Broadband Internet." The Quarterly Journal of Economics 130.4 (2015): 1781-1824.

Altonji, Joseph G., and Nicolas Williams. "Do wages rise with job seniority? A reassessment." Industrial & Labor Relations Review 58.3 (2005): 370-397.

Association of Talent Development, State of the Industry, 2014. https://www.td.org/Publications/Magazines/TD/TD-Archive/2014/11/2014-State-of-the-Industry-Report-Spending-on-Employee-Training-Remains-a-Priority.

Autor, David H., Frank Levy and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration." Quarterly Journal of Economics, 116(4), 2003.

Becker, Gary S. Human capital: A theoretical and empirical analysis, with special reference to education. University of Chicago Press, 1964.

Bessen, James (2015), Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth, Yale University Press.

Bessen, James E., How Computer Automation Affects Occupations: Technology, Jobs, and Skills (May 16, 2016). Boston Univ. School of Law, Law and Economics Research Paper No. 15-49. Available at SSRN:http://ssrn.com/abstract=2690435

Borghans, Lex, and Bas Ter Weel. "Are computer skills the new basic skills? The returns to computer, writing and math skills in Britain." Labour Economics 11.1 (2004): 85-98.

Bresnahan, Timothy, Shane Greenstein. "Technical progress and co-invention in computing and in the uses of computers." Brookings Papers on Economic Activity. Microeconomics (1996): 1-83.

Bresnahan, Timothy. "Computerisation and Wage Dispersion: An Analytical Reinterpretation." Economic Journal 109.456 (1999): F390-415.

Brynjolfsson, Erik, Lorin M. Hitt, and Shinkyu Yang. "Intangible assets: Computers and organizational capital." Brookings papers on economic activity 2002.1 (2002): 137-198.

Card, David and John DiNardo. "Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles" Journal of Labor Economics, 20, (2002), 733-783.

DiNardo, John, and Jörn-Steffen Pischke. "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?." Quarterly Journal of Economics 112.1 (1997): 291-303.

Dostie, Benoit, Rajshri Jayaraman, and Mathieu Trepanier. "What (if any) are the returns to computer use?." Applied Economics 42.30 (2010): 3903-3912.

Drew, Julia A. Rivera, Sarah Flood, and John Robert Warren. "Making full use of the longitudinal design of the Current Population Survey: Methods for linking records across 16 months." Journal of economic and social measurement 39.3 (2014): 121.

Entorf, Horst, and Francis Kramarz. "Does unmeasured ability explain the higher wages of new technology workers?." European Economic Review 41.8 (1997): 1489-1509.

Entorf, Horst, Michel Gollac, and Francis Kramarz. "New technologies, wages, and worker selection." Journal of Labor Economics 17.3 (1999): 464-491.

Fernald, John. Productivity and Potential Output before, during, and after the Great Recession. No. w20248. National Bureau of Economic Research, 2014.

Forman, Chris, Avi Goldfarb, and Shane Greenstein. "The Internet and local wages: A puzzle." The American Economic Review 102.1 (2012): 556-575.

Freeman, Smith. "Wage trends as performance displays productive potential: A model and application to academic early retirement." The Bell Journal of Economics (1977): 419-443.

Goldin, Claudia Dale, and Lawrence F. Katz. The race between education and technology. Harvard University Press, 2009.

Gordon, Robert J. The rise and fall of American growth: The US standard of living since the civil war. Princeton University Press, 2016.

Griliches, Zvi, and Jerry A. Hausman. "Errors in variables in panel data." Journal of econometrics 31.1 (1986): 93-118.

Hall, Bronwyn H., and Beethika Khan. Adoption of new technology. No. w9730. National Bureau of Economic Research, 2003.

Harris, Milton, and Bengt Holmstrom. "A theory of wage dynamics." The Review of Economic Studies 49.3 (1982): 315-333.

Hutchens, Robert M. "A test of Lazear's theory of delayed payment contracts." Journal of Labor Economics (1987): S153-S170.

Juhn, Chinhui, Kevin M. Murphy and Brooks Pierce, "Wage Inequality and the Rise in Returns to Skill" Journal of Political Economy, 101 (1993) 410-442.

Krueger, Alan B. How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989. The Quarterly Journal of Economics, Vol. 108, No. 1 (Feb., 1993), pp. 33-60.

Lazear, Edward P. "Why is there mandatory retirement?." The Journal of Political Economy (1979): 1261-1284.

Lucas, Robert E. "On the mechanics of economic development." Journal of monetary economics 22.1 (1988): 3-42.

Mishel, Lawrence, John Schmitt, and Heidi Shierholz. "Assessing the job polarization explanation of growing wage inequality." Economic Policy Institute. Working Paper (2013).

Pabilonia, Sabrina Wulff and Cindy Zoghi. "Returning to the returns to computer use." The American economic review 95.2 (2005): 314-317.

Pew Charitable Trusts. “Federal and State Funding of Higher Education: A changing landscape, 2015. http://www.pewtrusts.org/~/media/assets/2015/06/federal_state_funding_higher_education_final.pdf.

Restrepo, Pascual. “Skill Mismatch and Structural Unemployment.” Working paper, 2015.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. Integrated Public Use Microdata Series: Version 6.0 [Machine-readable database]. Minneapolis: University of Minnesota, 2015.

Salop, Joanne, and Steven Salop. "Self-selection and turnover in the labor market." The Quarterly Journal of Economics (1976): 619-627.

Topel, Robert. "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority." Journal of Political Economy 99.1 (1991).

Whitehouse, Economic Report of the President, 2015.

Tables and Figures

Table 1. Summary Statistics

	Computer use		Internet use	
	Total	Wage sample	Total	Wage sample
IT user	49%	56%	37%	43%
Education				
< high school	13%	9%	11%	7%
High school	33%	33%	29%	27%
Some post-sec.	28%	29%	30%	31%
College	18%	19%	20%	23%
Post-college	9%	10%	10%	12%
Age groups				
15-19	5%	2%	5%	2%
20-29	22%	17%	21%	17%
30-39	27%	29%	24%	27%
40-49	25%	29%	24%	28%
50-59	16%	18%	19%	21%
60-65	5%	4%	7%	5%
Part time	10%	9%	11%	9%
Union	18%	20%	16%	18%
Male	54%	52%	53%	50%
White	85%	87%	83%	86%
Log wage ($1982)		2.10		2.15
Δ log wage		2.7%		3.3%
N	318,547	109,154	216,300	60,940

Mean values weighted by population weights.

Table 2. Level regressions on log real wage with time trends.

	Computer use		Internet use	
IT user	0.186 (0.005)**	0.185 (0.006)**	0.187 (0.006)**	0.190 (0.007)**
IT user x *t*	0.0022 (0.0003)**	0.0022 (0.0004)**	0.0020 (0.0004)**	0.0017 (0.0005)**
T=1	0.020 (0.002)**	0.021 (0.002)**	0.024 (0.003)**	0.023 (0.003)**
Education				
< high school	-0.711 (0.005)**	-0.636 (0.012)**	-0.794 (0.007)**	-0.805 (0.013)**
High school	-0.533 (0.004)**	-0.475 (0.010)**	-0.583 (0.005)**	-0.563 (0.011)**
Some post-sec.	-0.408 (0.004)**	-0.351 (0.010)**	-0.441 (0.005)**	-0.388 (0.011)**
College	-0.152 (0.004)**	-0.112 (0.011)**	-0.185 (0.005)**	-0.152 (0.011)**
[> College omitted]				
Age groups				
[15-19 omitted]				
20-29	0.209 (0.006)**	0.307 (0.014)**	0.146 (0.007)**	0.150 (0.014)**
30-39	0.420 (0.006)**	0.530 (0.013)**	0.374 (0.007)**	0.348 (0.013)**
40-49	0.477 (0.006)**	0.595 (0.013)**	0.444 (0.007)**	0.395 (0.013)**
50-59	0.490 (0.006)**	0.632 (0.014)**	0.463 (0.008)**	0.380 (0.014)**
60-65	0.447 (0.008)**	0.605 (0.018)**	0.410 (0.010)**	0.333 (0.020)**
Part time	-0.154 (0.005)**	-0.158 (0.005)**	-0.186 (0.006)**	-0.186 (0.006)**
Union	0.169 (0.002)**	0.169 (0.002)**	0.135 (0.004)**	0.136 (0.004)**
Male	0.258 (0.002)**	0.257 (0.002)**	0.211 (0.003)**	0.211 (0.003)**
White	0.084 (0.003)**	0.084 (0.003)**	0.067 (0.004)**	0.067 (0.004)**
Year dummies	✓	✓	✓	✓
Education x *t*		✓		✓
Age x *t*		✓		✓
N	218,160	218,160	121,771	121,771
R-squared	0.416	0.417	0.414	0.415

Note: robust standard errors in parentheses. "**" = significant at the 1% level; "*" = significant at the 5% level. Regressions also include year dummies and population weights. Time, *t*, is measured in years, including fractional months. The second and fourth regressions include interaction terms for education variables with time and age variables with time, coefficients not shown.

Table 3. Correlations with occupation characteristics

	Job tenure (years)
Characteristics making it easier to monitor	
Repetitive work	-0.133**
	0.000
Routine activities	-0.214**
	0.000
Characteristics making it harder to monitor	
Abstract and creative activities	0.087**
	0.000
Decisions based on sensory or judgmental criteria	0.089**
	0.000
Prospective instrumental variable	
Strength (exertion vs sedentary)	-0.003
	0.896
N	2,208

Note: Table shows correlation coefficients with probability values of the null hypothesis (no correlation) underneath. "**" = significant at the 1% level; "*" = significant at the 5% level. Sample is occupation averages for each detailed occupation for each panel year, even years from 1996-2008.

Table 4. Basic Regressions on One-year Change in Log Real Wages

	Uses computer at work			Uses Internet at work		
	1	2	3	4	5	6
	LAD	IV	IV	LAD	IV	IV
IT use	0.006 (.001)**	0.019 (0.004)**	0.015 (0.004)**	0.004 (.002)*	0.024 (0.006)**	0.028 (0.007)**
<u>Education</u>						
High school			0.005 (0.004)			0.006 (0.004)
Some post-sec.			0.008 (0.004)*			0.005 (0.005)
College			0.013 (0.004)**			0.010 (0.006)
Post-college			0.015 (0.004)**			0.007 (0.006)
Part time			-0.006 (0.004)			-0.008 (0.005)
Union			0.004 (0.002)			-0.002 (0.003)
<u>Age group</u>						
20-29			-0.024 (0.008)**			-0.025 (0.010)*
30-39			-0.050 (0.008)**			-0.049 (0.010)**
40-49			-0.055 (0.008)**			-0.059 (0.009)**
50-59			-0.064 (0.008)**			-0.066 (0.009)**
60-65			-0.070 (0.008)**			-0.070 (0.011)**
Male			-0.002 (0.002)			0.003 (0.002)
White			0.003 (0.003)			-0.003 (0.004)
N	110,350	103,368	103,319	61,346	56,700	56,691
ρ		-0.044 (0.011)**	-0.036 (0.011)**		-0.054 (0.017)**	-0.059 (0.018)**
Wald test (P value)		0.000	0.000		0.002	0.003
<u>First stage regression</u>						
STRENGTH		-1.002 (0.076)**			-0.725 (0.065)**	

Standard errors clustered by occupation in parentheses. "**" = significant at the 1% level; "*" = significant at the 5% level. Year dummies not shown. Regressions use population weights. IT use is instrumented using the degree to which the worker's occupation involves STRENGTH and dummies for panel years (not shown). IV estimates use sample trimmed of 5% tails.

Table 5. Interactions

	Computer use	Internet use
1. Year		
1989	0.021 (0.005)**	
1993	0.004 (0.008)	
1997	0.021 (0.008)**	
1998		0.030 (0.009)**
2000		0.034 (0.011)**
2001	0.018 (0.008)*	
2003	0.018 (0.008)*	
2011		0.032 (0.011)**
2013		0.012 (0.012)
2. Share of occupation using IT		
<= 50%	0.018 (0.005)**	0.030 (0.009)**
> 50%	0.021 (0.006)**	0.027 (0.010)**
3. Age groups		
15-19	0.028 (0.014)*	-0.009 (0.025)
20-29	0.012 (0.020)	0.030 (0.034)
30-39	0.018 (0.019)	0.028 (0.034)
40-49	0.014 (0.019)	0.027 (0.034)
50-59	0.013 (0.020)	0.029 (0.034)
60-65	0.018 (0.021)	0.021 (0.035)
4. Schooling		
High School or less	0.020 (0.004)**	0.031 (0.008)**
Post-secondary	0.016 (0.006)**	0.030 (0.010)**
5. STEM occupation		
No	0.015 (0.004)**	0.023 (0.008)**
Yes	0.016 (0.006)**	0.038 (0.010)**

Robust standard errors in parentheses. "**" = significant at the 1% level; "*" = significant at the 5% level. Only the interaction terms are shown for regressions corresponding to columns 3 and 6 of the previous table.

Table 6. Worker characteristics and IT use

	Computer use			Internet use		
	1 Mean use	2 Probit marginal effects	3 IV Probit marginal effects	4 Mean use	5 Probit marginal effects	6 IV Probit marginal effects
Full sample	49%			37%		
<u>Education</u>						
< high school	13%			7%		
High school	36%	0.20 (0.02)	0.34 (0.11)	20%	0.09 (0.01)	0.09 (0.05)
Some post-sec.	55%	0.39 (0.02)	0.60 (0.15)	37%	0.24 (0.01)	0.35 (0.08)
College	75%	0.59 (0.02)	0.70 (0.24)	59%	0.45 (0.02)	0.40 (0.14)
Post-college	77%	0.63 (0.02)	0.57 (0.28)	68%	0.56 (0.02)	0.40 (0.17)
<u>Age</u>						
15-19	19%			9%		
20-29	45%	0.06 (0.01)	-0.14 (0.06)	33%	0.11 (0.01)	0.10 (0.04)
30-39	53%	0.12 (0.02)	-0.34 (0.12)	41%	0.18 (0.01)	-0.08 (0.08)
40-49	54%	0.11 (0.02)	-0.45 (0.13)	40%	0.17 (0.01)	-0.23 (0.09)
50-59	51%	0.08 (0.02)	-0.53 (0.14)	41%	0.13 (0.02)	-0.30 (0.10)
60-65	40%	0.00 (0.02)	-0.59 (0.12)	31%	0.05 (0.05)	-0.29 (0.09)
Part time	36%	-0.16 (0.01)	-0.18 (0.06)	23%	-0.15 (0.01)	-0.20 (0.04)
Male	43%	-0.14 (0.03)	-0.68 (0.07)	35%	-0.03 (0.02)	-0.39 (0.05)
White	50%	0.08 (0.01)	0.09 (0.04)	37%	0.04 (0.01)	0.02 (0.02)
US born	56%	0.14 (0.01)	0.22 (0.04)	37%	0.07 (0.01)	0.08 (0.03)
Ln real wage (instrumented)			1.41 (0.24)			1.43 (0.14)
N		318,497	246,045		216,289	163,370

Probit regressions show the marginal effects calculated at sample means. In columns 3 and 6, the log real wage is instrumented using the quartile of the worker's mean occupational wage in 1980. Probits use population weights.

Figure 1. Worker's marginal productivity with on-the-job learning.
A. Single learning episode

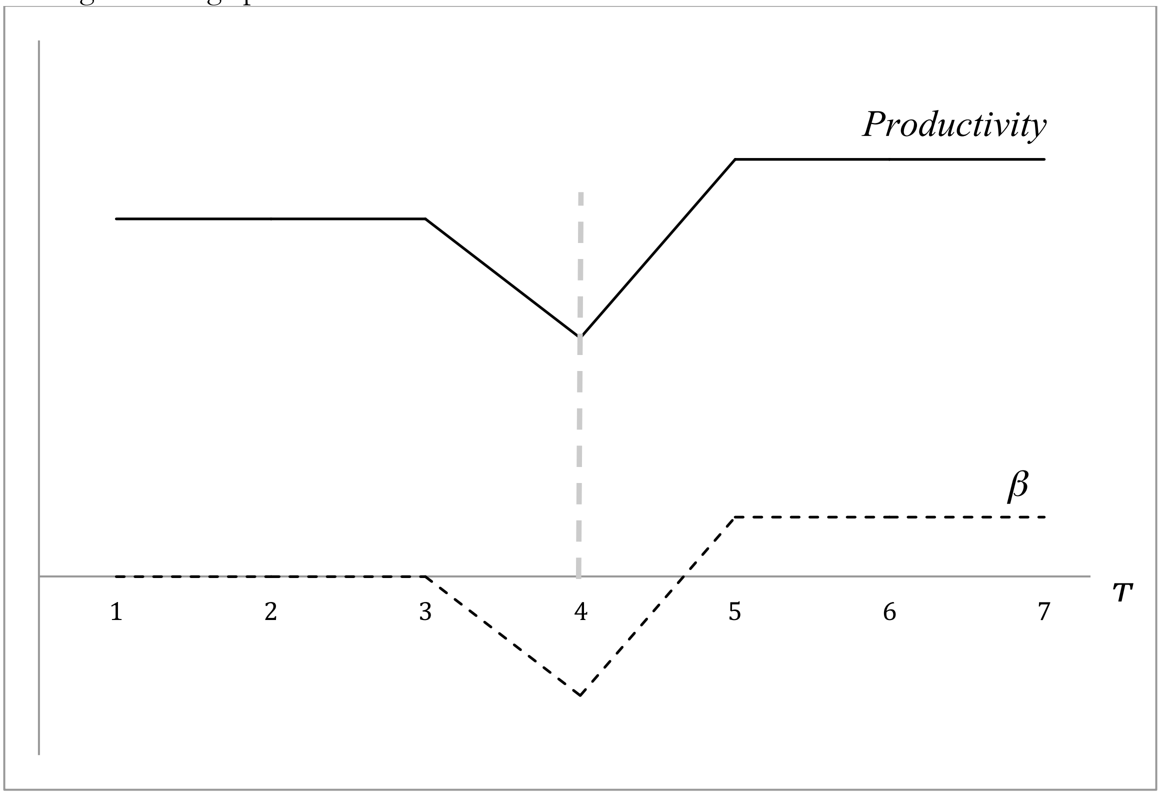

B. Stream of learning episodes

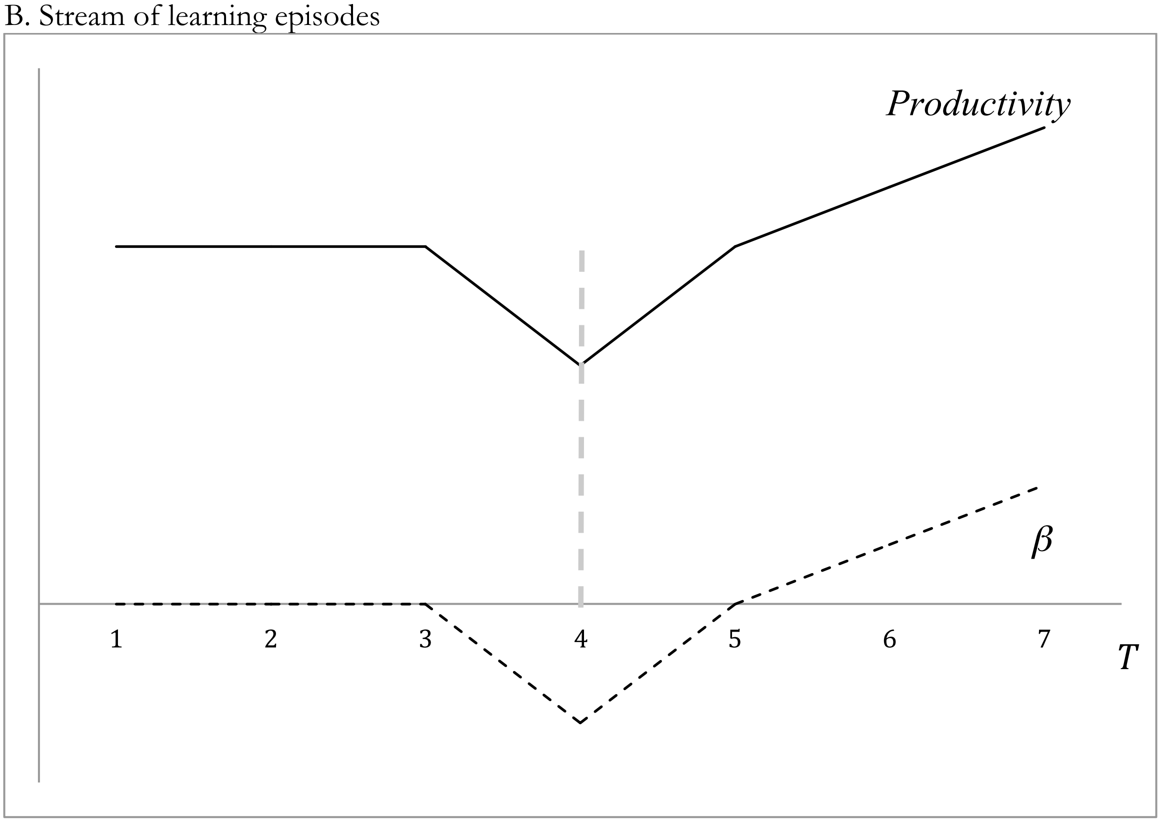

Figure 2. Share of workers using information technology by type

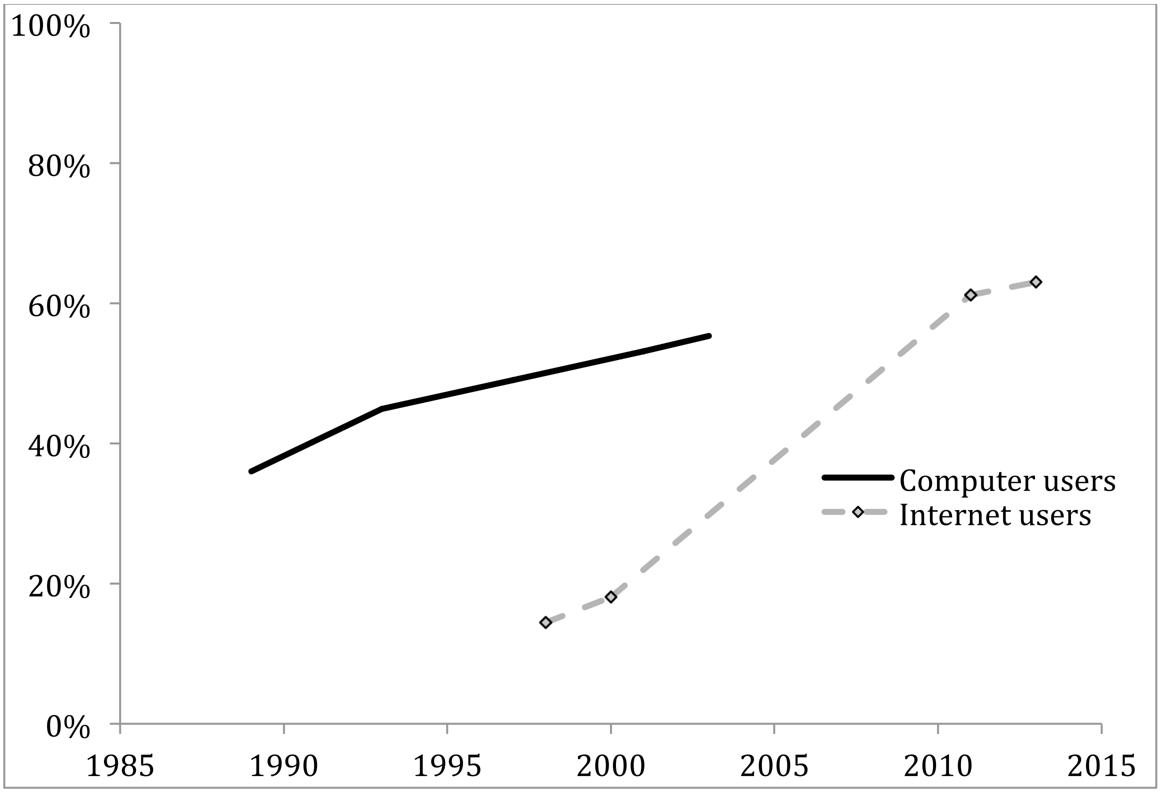

Population weighted means for each of 9 panels.