

Twisting the Demand Curve: Digitalization and the Older Workforce¹

Erling Barth, Institute for Social Research, Oslo and NBER

James C. Davis, US Census Bureau

Richard B. Freeman, Harvard University and NBER

Kristina McElheran, University of Toronto

¹ This paper was presented at the October 2019 Cornell ILR Conference on “Models of Linked Employer-Employee Data” celebrating Abowd, Kramarz, and Margolis and their 1999 AKM decomposition model, and at the Stanford SIEPR “Working Longer and Retirement” conference, also held in October 2019. Thanks to Nicole Fortin and Paul Oyer for comments. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1571. (CBDRB-FY20-P1571-R8608). This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; we have received funding from the National Institute on Aging Grant AG018854; The Norwegian Research Council Grant #280307, the Alfred P. Sloan Foundation, Grant B-2017-9943-OWRR, the National Science Foundation Grant 1928616 FW-HTF-P (AIRL), and grants from the Ewing Marion Kauffman Foundation, the Canadian Social Sciences and Humanities Research Council (SSHRC).

Twisting the Demand Curve: Digitalization and the Older Workforce

Erling Barth, James C Davis, Richard B. Freeman, and Kristina McElheran

June 2020

Abstract

With a focus on older workers, this paper explores the relationship between rising business investment in software and earnings for workers of different ages. Using a large panel of linked firm and worker data from the U.S. Census Bureau and the AKM framework, we estimate log earnings regressions for workers by age group. We extend the analysis to include job-spell fixed effects to allow for a correlation between the worker-firm match and age, as well as time-varying firm effects to allow for a correlation between wage-enhancing productivity shocks and software investments. Within job-spell, we find a positive effect of software capital on earnings. This premium declines, however, almost linearly after the age of 50 to about zero after the age of 65. We find a consistent pattern for exit: software investment is associated with less exit, but decreasingly so from age 55 and up. We probe mechanisms by estimating earnings elasticities for non-IT equipment investment. These are, if anything, increasing for older workers – suggesting our results are related to differences in older worker productivity in higher-tech firms and less so to age-related differences in bargaining power. In addition, we find that software investments increase the earnings of high-wage workers more than those of low-wage workers, and that software investments increase the earnings in high-wage firms more than in low-wage firms. Thus, the current surge in software investment in the U.S. may tend to widen earnings inequality both within and across firms.

Keywords: Software investments; Earnings equations; Older workers; the AKM model; Age-biased technical change

Erling Barth (corresponding author)
Institute for Social Research
P.O. Box 3233 Elisenberg
0208 Oslo
Norway
and NBER
erling.barth@samfunnsforskning.no

James C. Davis
Boston Research Data Center
National Bureau of Economic Research
1050 Massachusetts Ave.
Cambridge, MA 02138
james.c.davis@census.gov

Richard B. Freeman
National Bureau of Economic Research
1050 Massachusetts Avenue
Cambridge, MA 02138
freeman@nber.org

Kristina McElheran
University of Toronto Scarborough
1095 Military Trail
Toronto, ON
Canada
M1C 1A4
k.mcelheran@utoronto.ca

1. Introduction

From 1980 to 2019, spending on business software in the United States increased from 5% to 33% of spending on equipment, reflecting a far-reaching shift in investment from machines run largely by humans to digital tools. Figure 1 shows that while investments in equipment capital doubled from 1998 to 2017, investment in software capital quadrupled during the same period. The digital nature of software allows it to spread more rapidly than physical capital, while the movement of information and communication technology (ICT) to the cloud has increased the speed with which firms can access the latest software.² The COVID-19 crisis is likely to further accelerate these trends as firms seek more ways for humans to interact digitally rather than in close physical contact.

Figure 1. Investments in Software and Equipment Source: St Louis Fed, Fred, software is BEA Account Code: B985RC equipment is BEA Y033RC.

What workers do in this increasingly digital economy depends more and more on the software with which they work. The vast majority of white-collar workers use software regularly on their jobs and an increasing number of blue-collar workers do as well. Improvements in software, some generated by advances in AI in recent years, not only automate routine business activities but also allow software programs to perform non-routine cognitive work beyond beating humans at chess, go, and other strategic games.³ U.S. Department of Labor data on software shows huge growth in the number of categories of software “necessary” for different jobs.⁴

How do older workers fare as production becomes more digital? Will they be overtaken by workers in younger age groups, or could they actually be in higher demand? To answer these questions, this paper explores in detail the relationship between business investments in software and the earnings of workers of different age groups. We utilize detailed firm and worker data from the United States, collected by the U.S. Census Bureau and made available for research in the Federal Statistical Research Data Centers. The link between firm and worker allows us to study the labor market implications of this burgeoning software investment both within and across firms. Our focus is on differential effects within firms of software investments on the earnings of workers of different ages—with associated differences in education and experience—who may be more or less complementary to new technology. We also explore exit patterns within firms, comparing the impacts of firm spending on software on the earnings and employment of workers across the age spectrum with the impacts of firm spending on traditional equipment. There are mainly two competing views of how software might affect workers of different ages.

² See Jin and McElheran (2017) and Byrne and Corrado (2017).

³ See e.g., Brynjolfsson and McAfee (2014), Autor (2015), Dillender and Forsythe (2019), and Webb (2020).

⁴ Hu and Freeman (2020).

The first view is that younger persons are more familiar with new software than older workers are, for instance by being offered computer programming and computer science courses in school and growing up with the latest smart phone or game playing programs, while older persons are mired in past technologies. Training to keep older workers technically proficient has earned a mixed record.⁵ Just as new vintages of machines embody the newest knowledge, recent vintages of workers will tend to embody frontier skills. In this case, software spending will twist the demand curve for labor in favor of younger workers, reducing the relative wages and/or employment of older workers.

The second view highlights that recent generations of software, with simplified user interfaces and smarter back-end technology, reduce the specialized knowledge required to run and interpret its output.⁶ In that case, the main beneficiaries of software spending may be older workers, who can bring tacit and other experience-based knowledge to complement the technical prowess of the software. For example, early users of statistical software packages benefited from knowledge of cutting-edge programming and statistics to develop the best analyses. Now, data availability, programming modules, and machine learning algorithms perform technical tasks in ways that complement difficult-to-transfer knowledge and experience among workers who must assess and communicate the meaning of results for the business. Which view – that software is more complementary to the skills of younger workers, or to the experience of older workers – better fits the recent period of increased investment in business software?

Most work on skill-biased technical change has focused on differences between workers in blue-collar versus white-collar occupations, or on workers who differ in years of schooling, finding that computer use and ICT spending benefits skilled workers more than less-skilled workers.⁷ A few studies use European data to examine the impact of new technologies on workers of differing ages, with mixed results.⁸ Gaining a clearer view of

⁵ Earlier research argues that workers fail to maintain relevant skills only if a technological advance is unexpected (Bartel and Sicherman 1993). Follow-on studies argue that investment in new technology accelerates skill obsolescence at a time when training yields lower returns due to pending retirement (Friedberg 2003; Ahituv and Zeira 2011). Low-skilled older workers (particularly those lacking basic computer literacy) may be less likely to receive ICT-related training (Behaghel and Greenan 2010). On the other hand, training has been useful for ameliorating some of the negative effects of technological and organizational change for older workers in Germany (Battisti, Dustmann and Shonberg 2020) and France (Behaghel, Caroli and Roger 2014).

⁶ Some evidence for this is found in the higher diffusion of touchscreens compared to a wide range of advanced business technologies across the U.S. in recent years (Zolas et al. 2020).

⁷ See e.g. Katz and Murphy 1992; Autor, Katz and Krueger 1998; Goldin and Katz 2009.

⁸ In Norway, Schøne (2009) finds age bias related to computer use in cross-sectional comparisons, but not in changes over time. Hægeland, Rønningen and Salvanes (2007) find that only the oldest male workers (above 60) retire early in response to “extraordinary” changes in process technologies. In Germany, Battisti et al. (2020) and Beckmann (2007) find evidence of age-biased technological change that may be ameliorated to some extent with

potential age- and skill-biased outcomes at a time of rapid digitalization is central to understanding widening earnings inequality worldwide, particularly as populations age across many countries worldwide.

Our paper addresses this gap in knowledge regarding the relative impact of software on workers by age, and creates a novel panel of employer-employee linked data from the U.S. for 2002 to 2014, by linking data for firm software and equipment investments from the Census Annual Capital Expenditure Survey (ACES) with data on individual workers in each firm from the annual Longitudinal Employment and Household Dynamics (LEHD) database. We use this detailed panel to analyze the impact of firm investments in software on the earnings of workers by age and on the retirement/exit behavior of older workers. By analyzing firm spending on software rather than on the computer usage prevalent in previous studies, we highlight the importance of modern applications that build on the waves of computer hardware and infrastructure diffusion in prior studies.⁹

Our analysis extends the Abowd, Kramarz, and Margolis (1999) AKM model that uses detailed fixed effects to sweep out the impact of invariant characteristics of firms and individuals, allowing us to better identify the relationship between software and worker earnings. We first augment the AKM model by including job-spell fixed effects, as in Barth, Davis and Freeman (2018), to control for the particular match between a worker and a firm. This may be especially important for identifying mechanisms related to worker age, as match quality may change over the career of a worker and distort the estimated age profiles. Next, we add time-varying firm effects, as per Lachowska, Mas, Raffaele and Woodbury (2020) and Engbom and Moser (2020), to address the possibility that time-varying productivity shocks may be correlated with software investments,¹⁰ which would bias estimates relying only on

training. In France, ICT investment has been found to be age biased (Aubert et al. 2006), and may disproportionately affect lower-skilled older workers (Behaghel and Greenan 2010, Behaghel et al. 2014). Recent work using Dutch micro data also shows negative impacts for older workers due to automation (Bessen, et al. 2019). In contrast, a U.S.-based study using CPS data on computer use in the 80s and 90s (Weinberg 2004), finds that computers complemented experienced workers who only completed high school; among college graduates, computers complemented younger workers.

⁹ Studies on PC use and labor market outcomes include Krueger (1993), Autor, Levy and Murnane (2003), Schöne (2009), Autor and Dorn (2013), and Hershbein and Kahn (2018), among others. Forman, Goldfarb and Greenstein (2012) document wage inequality related to internet use. Broadband is the focus of Akermann, Gaarder and Mogstad (2015) and Poliquin (2020). Complex software is implicated in increased demand for non-routine cognitive workers in U.K. small businesses (Gaggl and Wright 2017) and *reduced* wage polarization within Chilean firms (Almeida, Fernandes and Viollaz 2020). A few recent studies of how ICT affects the skill content of jobs include measures of modern software (Dillender and Forsythe 2019; Atalay et al. 2018; Webb 2020).

¹⁰ For instance, complementarities between ICT investments and unobserved changes in organizational design and management practices, observed in several studies (see e.g. Bartel, Ichniowski and Shaw 2007, Bresnahan, Brynjolfsson, and Hitt 2002, Brynjolfsson and McElheran 2019, Battisti et al. 2020), may have direct effects on productivity and correlate with software investments.

firm fixed effects. This approach is particularly useful in our setting, since it allows for identification of differences in the impacts of time-varying firm-specific variables, such as software investments, on different types of workers, such as age groups, within firms.

The AKM framework is also central in our exploration of the extent to which software investments tend to widen wage dispersion within firms. We leverage it to estimate individual fixed effects for the entire labor market, which we then interact with observed software investments. In this way, we may unpack whether software represents a typical skill-biased technical change, as discussed above, or adds to a “polarization” of jobs within firms, whereby workers in the top and the bottom of the occupational earnings distribution gain at the expense of the middle.¹¹ Similarly, we explore the extent to which software investments tend to increase earnings more in high-wage firms, as defined by the AKM decomposition, than in low-wage firms. In this way, we may provide direct evidence for one widely supposed mechanism behind the recent widening of wage dispersion across firms.¹²

The next section provides a description of the data, and some statistics that describe differences in characteristics between software-intensive firms and firms that use less software. Section 3 provides a theoretical framework for the analysis, and section 4 lays out our empirical strategy. Section 5 reports the results and section 6 concludes.

2. Data

The primary independent variable in our analysis is software capital per worker calculated using standard perpetual inventory methods. The amount of capital investment in software is reported in the Census Bureau's Annual Capital Expenditure Survey (ACES) – a nationally representative annual survey of 50,000 firms¹³ that collects basic data for all firm expenditures on new and used structures and equipment chargeable to asset accounts for which depreciation or amortization accounts are ordinarily maintained.

Beginning in 2002, Census asked respondents to report capital expenditures for computer software developed or acquired as either prepackaged or vendor-customized for internal use. Capitalized computer software expenditures consist of the cost of materials and services directly related to the development or acquisition of software; payroll and payroll-

¹¹ See e.g. Autor, et al (2003), Autor and Dorn (2013), Michaels, Natraj, and Van Reenan 2014, and Harrigan, Reshef and Toubal (2016).

¹² See e.g. Card, Heining and Kline 2013; Barth et al. 2016; Card et al. 2018; Song et al. 2019, and Cornwell, Shmutte, and Scur (2019) for recent evidence on rising earnings inequality within and between firms as well as potential causes.

¹³ ACES was expanded from 45 to 50 thousand employer firms in 2016. The survey also includes data for 15 to 30 thousand non-employer firms per year that are not used in this study.

related costs for employees directly associated with software development; and interest costs incurred while developing the software.¹⁴ The survey covers all companies with at least 500 paid employees with certainty and selects smaller companies randomly. The sample of firms that we use comes primarily from the certainty sample because we create measures of software capital stock from the expenditures using standard perpetual inventory methods.¹⁵ For our analysis, the expenditures must vary over time, ruling out most randomly selected smaller firms, which appear in the data set irregularly.

The ACES data asks firms about spending on software that firms capitalize as investment, which makes it comparable to other capital expenditures. This excludes software spending that is deductible as an expense of production or intermediate purchase and thus falls short of the National Income and Product Accounts (NIPA) estimate for private fixed investment based on sales of software reported by producing firms.¹⁶ It does, however, capture a key shift in the nature of ICT investment in recent decades from computer hardware or pre-packaged software towards firm-specific software.

This distinction matters because expensed software typically represents routine production costs similar to materials purchases (e.g., word processing and automated payroll applications). In contrast, large custom software investments and capitalized purchased software are argued to represent investments in innovation by the firm (e.g. Bessen and Righi 2020). A useful example is the custom logistics software system built and maintained by Walmart. These lumpy ICT investments – and the significant “co-invention” of the organizational design and production processes that typically accompanies them (Bresnahan and Greenstein 1996, McElheran 2015) – have been linked to significant productivity gains and shifts in skilled labor and wages at the firm level.¹⁷

¹⁴ Capitalized computer software is defined by the criteria in Statement of Position 98-1, Accounting for the Costs of Computer Software Developed or Obtained for Internal Use, issued by American Institute of Certified Public Accountants.

¹⁵ As initial conditions for the entry year, the industry average shares of software investments of total investments were combined with data on the initial gross value of assets. To accommodate for the shorter history and higher growth of software investments, the initial levels of software capital shares were set to .5 of the share of investments associated with software in each industry in the initial year.

¹⁶ St. Louis Federal Reserve Bank FRED series B985RC1A027NBEA. As most software spending is expensed, the capital data covers only about one-third of national investment in software. Appendix A describes a Bureau of Economic Analysis investigation of the issue. We have chosen to include only capitalized software investments in our analysis, and leave the treatment of expensed items, available in the Information & Communication Technology Survey (ICTS) for the years when it was conducted, to future exploration.

¹⁷ Brynjolfsson and Hitt 1996 & 2003; Bloom, Sadun, and Van Reenan (2012); Tambe and Hitt (2012), and Jin and McElheran (2017) document IT-driven productivity at the firm and establishment level. More skilled labor is observed in firms investing in ICT (Bartel et al. 2007; Bresnahan et al. 2002, Brynjolfsson and McElheran 2019). Innovation associated with technological change has also been linked to higher average firm wages and employment (e.g., Van Reenan 1996).

Recent work on automation and labor market outcomes emphasizes the importance of refined ICT measures to avoid confounding technologies that automate different types of tasks (e.g., Acemoglu and Restrepo 2019, Dillender and Forsythe 2019, Almeida et al. 2020). From 2002 on,¹⁸ ACES data also breaks out capitalized computer software into prepackaged, vendor-customized and internally-developed (including payroll) sub-categories. We ran all of the estimates presented in this paper using a software measure constructed as the sum of the first two sub-categories of prepackaged and vendor-customized. We found results consistent with those presented throughout the paper that use the ACES measure of total capitalized computer software. Given that it is unlikely that firms change their internal accounting practices from year to year, the ACES measure of capitalized software should sufficiently capture changes in software investment over time, particularly large-scale software investments that may impact outcomes for workers.

Our data links information on capitalized software investment, capital equipment investment, firm employment, and sales to workers in the annual Longitudinal Employer-Household Dynamics (LEHD) Employment History Files (EHF) data from 2002 to 2014 that follow individual workers in each firm (Abowd, Haltiwanger and Lane 2004; Abowd et al. 2009; Vilhuber 2018). The EHF is derived from firm reports to state unemployment insurance agencies for all employed workers. The LEHD EHF data contain the earnings of workers. We use annualized earnings for a worker's main job, defined as the longest job across quarters with the highest earnings.¹⁹ Worker age comes from the LEHD Individual Characteristics File (ICF), which obtains this from Social Security Number (SSN) applications to the Social Security Administration (SSA). The ICF imputes the education of workers using reported education from the 2000 decennial census. We construct shares of workers with different levels of education at the firm level from the imputed ICF values.²⁰ In 2012, 14% of workers in our sample have less than high school education, 26% have high school degrees, 31% have some college or trade school, and 29% have a college degree. Our LEHD data cover 22 states and the District of Columbia, representing half of the U.S. workforce in 2012, and appears representative of the country as a whole.²¹

¹⁸The ACES survey started in 1993.

¹⁹Jobs with earnings less than half the federal minimum wage in 2002 are excluded.

²⁰ This version of imputed education is reasonable for estimating the education distribution of workers at the firm but raises measurement error problems if used for individual-level earnings specifications.

²¹The 23 approved states for use in this study include AR, AZ, CA, CO, DC, DE, IA, IL, IN, KS, MD, ME, MT, ND, NE, NM, NV, OK, PA, TN, TX, VA and WA. Average earnings for the workers in covered states is 2.7% higher than for the 28 uncovered states. Employment of workers 45 and older in our 23 states is also half the total. The 0.506 employment share of these states nationally, 0.501 older worker share, and worker average

We group the 11.6 million workers (57% are male) in our samples by five-year age groups (20-24, 25-29, 30-34, etc.). Each age group represents around 12% of the distribution of workers, except for the three older groups, whose shares decline as workers exit the labor force. Workers aged 55-59 represent 9%, 60-64 year olds are 6%, and aged 65+ workers are 3% of the sample, respectively.

We work with two samples: (1) the *full* LEHD data for 23 states, which we use to run an AKM decomposition to obtain quintiles of the individual and firm fixed effects for all workers and firms in those states; and (2) our *analysis* sample, which contains all workers in the 23 LEHD states with jobs in firms in the ACES sample. Appendix B compares some attributes of the two samples.

If firms with differing software-capital-to-worker ratios hire workers from the same broad labor market, then the distribution of worker attributes among the firms should be informative about complementarities between software and worker attributes. In fact, a standard test for complementarities explicitly tests whether production inputs are adopted together across firms with similar production functions (e.g., Brynjolfsson and Milgrom 2013). Accordingly, we have examined the composition of the work forces of firms with differing amounts of software capital per worker.

Our examination reveals three high-level patterns in the data on which the ensuing econometric analyses focus:

1) Differing age compositions of the work force among firms by software intensity

Figure 2 shows the share of workers by age group for firms in the bottom quintile (Q1), median quintile (Q3), and top quintile (Q5) of software capital per worker.²²

Figure 2. Employment Share by Age Group. By Quintiles of Software Intensity.

Note: Average employment share across firms by quintile of software capital per worker. Q1 is the first quintile, Q3 the median quintile, and Q5 the top quintile. Employment weighted quintiles. Cross sectional data, 23 LEHD states in 2012.

earnings were calculated from the public use LEHD Quarterly Workforce Indicators (QWI) V4.6.0 downloaded 5/24/2020.

²² Note that there may be a lot of individual and firm demographics behind the age composition as well, see e.g. Autor and Dorn (2009) who argue that “occupations will ‘get old’ as their employment declines—that is, the mean age of an occupation’s workforce will rise, “because older workers are less likely to leave declining occupations, and younger workers will be hired into growing occupations.” The same argument may apply to firms, distorting effects related to skill complementarity in production unless estimated carefully.

Workers aged 20-29 are over-represented in the bottom quintile of software intensity, possibly suggesting that their labor skills make them more suitable employees for firms that use less software. Workers in the 35-49 age group are over-represented in the top quintile of software intensity firms, suggesting the opposite – that they have skills that complement software. The share of workers aged 50-64 declines steeply from the top quintile to the bottom quintile, and the oldest group, 65+, has the lowest representation in the most software-intensive firms.

The finding that workers in the 35-49 age group appear relatively more complementary with software than both younger and older workers is consistent with related work in France (Aubert, Caroli and Roger 2006), Germany (Beckmann 2007) and Norway (Schøne 2009). Since workers age over time, however, the cross-section pattern misses the dynamics of adjustment of workers' careers across firms as they age and firms' use of software over time, which we explore in more depth in section 4.

2) Software-intensive firms are bigger and have more educated workers and higher labor productivity than others

Table 1 shows average firm characteristics in 2012 by quintile of software intensity across firms. The first line displays a large dispersion across firms in logged software capital per worker. Looking at the distribution of employment, we see that the more software-intensive firms are larger, consistent with other work finding complementarities between scale and software investment.²³ Software-intensive firms also have more equipment per worker, are older, and are more likely to have multiple establishments. The employment shares by education groups show that firms with the greatest software intensity employ many more college-educated workers and many fewer workers who have not completed high school.²⁴

Table 1. Firm Characteristics by Quintile of Software Intensity

Note: The quintiles of software capital per worker are calculated over observations of firms (education shares over workers). Cross sectional data, 23 LEHD states in 2012.

Given these patterns, we would expect firms with greater software intensity to have higher labor productivity than others (e.g., Brynjolfsson and Hitt (2003), Bloom et al. 2012,

²³ See e.g., Tambe and Hitt (2012), McElheran (2015).

²⁴ This is consistent with the relationship between ICT adoption and skilled labor found in prior work where skill is measured in terms of formal education (e.g. Bresnahan et al. 2002).

Tambe and Hitt 2012). This is indeed what the last line in Table 1 shows. In particular, average revenue per worker in the top quintile of software intensity is more than 30 times the average revenue per worker in the bottom quintile of software intensity.

3) Software-intensive firms employ workers with high individual fixed effects in wages and are high-wage firms in terms of their firm fixed effects.

An alternative way to see where software-intensive firms and their work forces fit in the overall labor market is to exploit the panel aspects of our data to estimate the AKM decomposition of labor earnings into worker-fixed effects and firm-fixed effects, controlling for age and year effects. We do this for all workers in our full LEHD sample as detailed in the methodology section, below. Panel A of Figure 3 displays the results in terms of the average worker fixed effect.

Figure 3. Average Worker and Firm Fixed Effects by Quintile of Software Intensity

Note: Average worker and firm fixed effects across firms by quintile of software per worker. Q1 is the first quintile, Q3 the median quintile, and Q5 the top quintile. Employment weighted quintiles. Worker and firm fixed effects are first estimated from an AKM decomposition on the whole labor market of the LEHD states, controlling for worker age, age squared, and years of observation. Next, “Worker (Firm) FE res” is the residual, and “Worker (Firm) FE Xb” the predicted value of the worker (firm) fixed effects regressed on time invariant characteristics such as education, gender, and cohort (industry, average firm and establishment size, and firm age), see section 5 below.

There is a strong positive correlation between the individual fixed effects from the wage regression and firm software intensity, indicating that software-intensive firms disproportionately employ workers with attributes that help them achieve higher earnings. The figure also displays the average predicted value of the individual fixed effect, generated using gender, education, and age-cohort as time invariant covariates and the residual unobserved individual fixed effect. Note that the positive slope between the individual fixed effects and software intensity mainly arises from the unobserved part of workers’ time-invariant characteristics.

Panel B displays the result of the analogous calculation for firms from the same decomposition equation. The slope linking software intensity to the average firm fixed effect is steeper than the slope in Panel A, with a difference between the bottom and the top quintile of more than 0.5 log points. Both the unobserved part of the firm fixed effects as well as the observed part (predicted using average firm and establishment employment, firm age, and

industry) show a positive relation to software intensity. In this case, software intensity is positively correlated with both unobserved and observed productivity characteristics of the firm.

Several mechanisms may provide explanations for these observations, including wage effects of productivity-enhancing software capital, selection of software investments across firms, and selection of workers into firms. All of these factors may vary by the age group of workers. To answer our key research question, therefore, we need to separate out these different mechanisms. The fact that we have a matched panel of firms and workers, and that software intensity is time varying within firms, allows us to utilize an augmented AKM estimation methodology particularly suited to disentangle these factors by age. We first present our methodology in some detail.

3. Theoretical Framework

We examine three ways in which software capital may affect younger and older persons differently. First is the cohort of workers, which determines much of what they learned in formal education. Second is their age, with biological factors that primarily affect their physical capabilities. Third is the incentives and opportunities they have to invest in new knowledge for their job, where policies may be most effective. We contrast these mechanisms with expectations for investments in traditional non-ICT equipment to fix ideas for our empirical investigation.

Cohort effects evoke Johansen’s (1959) “putty-clay” framework for physical capital. When persons invest in human capital, they embody the technology and knowledge prevalent at that time, which often locks them into a particular occupation or area with, ideally, frontier skills for that period. As knowledge expands, the human capital of a cohort ages, with younger generations more closely connected to novel technologies than older generations. To the extent that software technologies change more rapidly than physical equipment, a rising share of software in capital will twist demand for labor toward younger persons due to cohort effects.²⁵

As workers age, they typically “learn by doing” and have more opportunities for training, including keeping up with technical changes to a certain degree.²⁶ However, aging

²⁵ Work on skill obsolescence dates back to Rosen (1975). Chari and Hopenhayn (1991) argue that new technologies require vintage-specific skills, so any increase in the rate of technological change increases returns for more-recent vintages and flattens the age-earnings profile.

²⁶ Recent work indicates that this is more difficult in contexts of rapid technological change and has been shifting returns to experience in the U.S. in recent years (e.g., Deming and Noray, forthcoming).

by itself changes the comparative advantage of workers to undertake different tasks. Biological processes give the young an advantage in tasks that require physical strength, dexterity, and a short reaction time, while the older may have an advantage in tasks that require careful deliberations, collaboration, and the ability to combine information with experience.²⁷ The physical wear-and-tear on the body makes it difficult for someone to keep driving a huge truck as they get older. However, the clearest examples of twists in demand with age occur in sports.

As workers age, their productivity and the tasks they do are thus likely to change, as is the fit of their skills with different types of firm investment. To the extent that age dictates that younger workers are more able to perform / OR? / are more likely to be assigned tasks involving physical labor than older workers, who are in turn more likely to be involved in organizational tasks, firm investments in labor-saving equipment are likely to twist demand against the young. How software capital investments shift demand for tasks undertaken by workers of different ages is less clear. Older workers may become more productive if they manage younger workers who interact with software.

The putty-clay analogy has limitations, moreover, when we consider training and experience. The key difference between human capital and vintages of physical (equipment) capital is that, through training and experience, workers acquire new skills and adapt to new technologies. This process is not frictionless, and training requires investments with payoffs that occur in the future. Older workers may have more to gain from training because their knowledge is of older vintage, but a greater distance to the technological frontier may require higher investments in training. Most importantly, the shorter time remaining until retirement skews the investment decision – by individuals, firms, and often public policies— to favor training the young (e.g., Friedberg 2003, Behaghel and Greenan 2010). Without intervention, incentives and opportunities to acquire frontier skills will tend to favor the young and improve their fit with novel technologies.

To explore these competing mechanisms with the data to hand, we need a framework that incorporates different productivities by age and that allows for different interactions between the two types of capital and worker age. Because software can affect the productivity of both equipment capital and workers, we develop a simple framework where software capital is potentially complementary with equipment capital and with efficiency

²⁷ For instance, Paccagnella (2016) provides evidence that individuals replace declining standardized proficiencies with other, more difficult-to measure, skills based on experience as they age.

units of labor, and where the efficiency of labor may vary between age groups. Finally, we allow for differential impacts of equipment capital and software capital on the efficiency of each age group.

Consider the simple production function (ignoring subscripts for firm j and time t):

$$Y = \Omega K^{\alpha_K} S^{\alpha_S} \tilde{L}^{\alpha_L}$$

where Y is revenue, Ω is total factor productivity, K is capital (without software, but including computer hardware capital), S is a measure of software capital, and \tilde{L} is efficiency units of labor. Assuming that workers of different age groups are perfect substitutes, but with potentially different productivities, efficiency units of labor is given by $\tilde{L} = \sum_g (e^{\omega_g} L_g)$, where ω_g is the productivity of each age group, g .

To investigate the extent to which software capital affects the relative productivity of each type of labor, g , we assume that: $\omega_g = \pi_g + \delta_g^s \ln s + \delta_g^k \ln k$, where π_g is an age-specific productivity term, $s=S/L$ and $k=K/L$ are capital intensities, and δ_g^k , $\kappa=s, k$, represents the relative impact of software and equipment capital intensities on age group g . The marginal productivity of labor of age g is given by:

$$\frac{\delta Y}{\delta L_g} = \alpha_L \frac{Y}{\tilde{L}} \frac{\delta \tilde{L}}{\delta L_g} = \alpha_L \Omega k^{\alpha_K} s^{\alpha_S} L^{(\alpha_K + \alpha_S + \alpha_L - 1)} \left(\frac{L}{\tilde{L}} \right)^{1 - \alpha_L} e^{\pi_g + \delta_g^s \ln s + \delta_g^k \ln k}$$

Software capital per worker affects the productivity of labor of type g in two ways. First it affects the average productivity of efficiency units of labor. In addition, it affects the relative productivity of each age group of workers. $\alpha_S > 0$ suggests complementarity between the two types of capital, while the effect of software on the efficiency units of labor differs between groups, and is given by $\alpha_S + \delta_g^s$ which may take any sign. A similar argument applies to equipment capital. We discuss issues of normalization in the empirical implementation below.

Wage Determination

We allow for firm-specific wage determination, and consider a simple monopsony model (see Manning, 2002; Card, Cardoso, Heining and Kline 2018). Let w_g be the wage of age group g in the firm. The labor supply facing each firm is given by $L_g(w_g)$, where $\frac{\delta L_g}{\delta w_g} > 0$. Let $\varepsilon > 0$ be the elasticity of labor supply facing the firm. Profits are given by:

$$\Pi(\mathbf{w}) = Y(K, S, \tilde{L}(\mathbf{w})) - \sum w_g L_g(w_g)$$

Where \mathbf{w} is the vector of wages. Profit maximization w.r.t w_g gives:

$$(1) \quad w_g = \frac{\varepsilon}{1 + \varepsilon} \frac{\delta Y}{\delta L_g} = m_g \alpha_L \frac{Y}{L} \frac{L}{\tilde{L}} e^{\pi_g + \delta_g^s \ln s + \delta_g^k \ln k}$$

Where $m = \frac{\varepsilon}{1+\varepsilon}$ is the “monopsony discount” of wages on marginal productivity. The assumption that age groups are perfect substitutes implies that the marginal productivity of each group is independent of its relative size, and we note that the term $\alpha_L \frac{Y_L}{L}$ is firm specific and common for all age groups within the firm.

4. Empirical Implementation

The age-specific wage in firm j at time t , w_{gjt} , is given by equation (1). Let the wage of individual i (of age group g), in firm j at time t , be given by $w_{ijt} = e^{c_i + X_{it}b + \xi_{ij}} w_{gjt}$. $c_i = Z_i b^z + \tilde{c}_i$ is an individual fixed component, including observables, Z_i , such as gender and education, and an unobserved fixed effect, \tilde{c}_i . X is a vector of time-varying individual characteristics, ξ_{ij} a match specific component between the individual and firm.

We represent the monopsony discount of group g in firm j at time t by the decomposition $m = e^{\mu_g + b_{jt}}$, where μ_g is age specific (determined by factors such as the relative size of each age group in the labor market) and b_{jt} is firm specific (determined by factors such as the firm’s position in the wage distribution across firms in the labor market). Adding an error term, u , we obtain:

$$(2) \quad w_{ijt} = \alpha_L \frac{Y_{jt} L_{jt}}{L_{jt} L_{jt}} e^{c_i + X_{it}b + \xi_{ij} + \mu_g + b_{jt} + \pi_g + \delta_g^s \ln s_{jt} + \delta_g^k \ln k_{jt} + u_{ijt}}.$$

Identification Within Job-Spell

We start out by estimating a log wage regression on the observed time-varying covariates, including dummies for age groups, the interaction terms between age groups and the two measures of capital (S and K), log of employment, year dummies, and a job-spell fixed effect:

$$(3) \quad \ln w_{ijt} = X_{it}b + \mu_g + \pi_g + (\alpha^s + \delta_g^s) \ln s_{jt} + (\alpha^k + \delta_g^k) \ln k_{jt} + (\alpha_s + \alpha_k + \alpha_L - 1) \ln L_{jt} + \ln \alpha_L + \psi_{ij} + \gamma_t + e_{ijt}$$

The job-spell fixed effect, $\psi_{ij} = c_i + \varphi_j + \xi_{ij}$, incorporates the AKM specification, with a worker fixed effect, c_i , a firm fixed effect, φ_j , and an additional match fixed effect, ξ_{ij} , related to the specific match between the firm and the worker.²⁸ We discuss the underlying parameters absorbed by the job fixed effects below.

²⁸ ψ_{ij} may easily be decomposed into its three parts by defining the match effect as orthogonal to the individual and firm effects (but not necessarily to the other covariates of the model) and appropriate normalizations.

For each age group, g , we obtain the estimate for $\tilde{\delta}_g^\kappa = (\alpha^\kappa + \delta_g^\kappa)$, which incorporates the effect of capital on the labor productivity of the firm, as well as the group-specific parameter. Even if we are not able to identify the capital-specific parameters separately from the group-specific parameters, we may identify the difference between the group-specific parameters by taking the difference between the estimates for different age groups. Since the firm specific term is constant across age groups, it drops out of the difference: $\tilde{\delta}_g^\kappa - \tilde{\delta}_h^\kappa = \delta_g^\kappa - \delta_h^\kappa$.

Adding Time-Varying Firm Effects

We may write log wages given by (2) as:

$$(4) \quad \ln w_{ijt} = X_{it}b + \pi_g + \mu_g + \varphi_{jt} + \psi_{ij} + \gamma_t + \delta_g^s \ln s_{jt} + \delta_g^k \ln k_{jt} + u_{ijt}$$

Where we have added a time-varying firm effect that absorbs all time-varying firm-specific factors:

$$(5) \quad \varphi_{jt} = b_{jt} + \omega_{jt} + \alpha^s \ln s_{jt} + \alpha^k \ln k_{jt} + (\alpha_s + \alpha_K + \alpha_L - 1) \ln L_{jt} + (1 - \alpha_L)(\ln L_{jt} - \ln \tilde{L}_{jt})$$

which includes both the unobserved productivity and monopsony discount, and the observable time-varying firm-level covariates.

In the previous job-spell fixed effect equation (3), firm-specific shocks to b_{jt} and ω_{jt} (deviations from their job-spell-specific means that are not captured by the common year effect γ_t) are included in the error term, e . To the extent that such shocks are correlated with investments in software or equipment, they will bias the within job-spell estimates. A similar concern arises with regard to the unobserved time-varying efficiency units of labor, \tilde{L}_{jt} .

Adding time-varying firm effects removes such biases.

In equation (4), both $X_{it}b$ and $\pi_g + \mu_g$ are observable time-varying individual variables (as also g may vary over time within individual in the panel). The component $\delta_g^s \ln s_{jt} + \delta_g^k \ln k_{jt}$ varies between groups within the firm over time as well as at any given point in time, and may be estimated using interaction terms between age group and capital intensities. Estimating equation (4) with time-varying firm effects effectively controls for both observable and unobservable wage shocks that affect wages of every worker in the firm in any year. In particular, the time-varying firm effect includes two unobserved time-varying firm-specific terms: b_{jt} and ω_{jt} . These represent temporal variation in the monopsony discount of

the firm (i.e. market power or bargaining power in the labor market) and potential temporal variation in the unobserved productivity of the firm, respectively.

The interpretation and identification of the time-varying firm effects and the job-spell fixed effects require normalization. Excluding a reference year ($t=0$), i.e. the first year a firm appears in the panel, from the φ_{jt} vector in the estimation, the job-spell fixed effects will absorb the worker fixed effect, the firm fixed effect *for the reference year*, and the job-spell fixed effect: $\psi_{ij} = c_i + \varphi_{j0} + \xi_{ij}$. The estimated **time-varying firm effects** will then be defined in terms of the difference from the reference year: $\tilde{\varphi}_{jt} = \varphi_{jt} - \varphi_{j0}$.

The last terms, $\delta_g^k \ln \kappa$ for $\kappa=S, K$, represent the impact of software and equipment capital of age group g , in addition to the overall effect absorbed by φ_{jt} . Again, we are not able to separately identify the overall impact and the group-specific impact of capital without some normalization. We thus normalize the overall efficiency of the firm to that of a base age group, $L_{g=0}$, and exclude the interaction with the base group in the estimation of equation (4). The parameters $\delta_{g \neq 0}^k$ of equation (4) should then be interpreted as a relative effect for age group g compared to the base age group.²⁹

5. Results

Software Earnings Premiums by Age

The elasticity of earnings with respect to software capital is strongest among workers between 30 and 59 years of age. Figure 4 shows the estimated elasticity of individual earnings with respect to software capital per worker, estimated for each age group in models that also control flexibly for equipment capital by age group.³⁰ The figure clearly shows a hump-shaped relationship by age. High-wage industries are more software intensive, and within industries, high-wage firms are more software intensive. This is visible in the figure by the downward shifts as additional fixed effects are added to the model.

Figure 4. Software Elasticity of Earnings by Age Group. FE: Region + Industry + Firm

Note: Dependent variable: Log earnings. The figure shows the coefficients for log software capital per worker interacted with ten age groups. The models control flexibly also for equipment capital by age group, and include controls for age group \times gender, year dummies, three dummies for level of education, log firm employment, log

²⁹ The parameters δ_0^k ($k=S, K$) are incorporated into α^k of the time-varying firm effect, and $\pi_0 + \mu_0$ is added into the constant term of the equation.

³⁰ The models also include indicators for age groups interacted with gender, and controls for year, education, firm and establishment size (employment), and firm age and firm age squared.

establishment employment, firm age, firm age squared. Region, Industry, and Firm fixed effects are added successively.

The idea behind the AKM model is to control for systematic selection of *both* firms and workers. For instance, high-wage workers may be more likely to work in high-wage firms, and high-wage workers may be more likely to work in software-intensive firms, conditional on the wage premium of the firm. Notably, this selection may also be systematically different by age group, as workers sort into different firms over their careers. Adding dummies for both individuals and firms effectively controls for such selection.

Our next step is thus to see how the estimated elasticities respond to the inclusion of such controls. We go one step further and add job-spell fixed effects: indicators for the specific job-spell of one particular worker in a particular firm, which comprise the individual fixed effect, the firm fixed effect, and any match-specific fixed effect.

Table 2. Earnings Elasticities by Age. Dependent variable: log Earnings. Job-spell FE

Note: The table shows estimated elasticities of earnings (standard errors) with respect to software and equipment capital, both estimated from the same model with job-spell fixed effects. All capital variables are measured per worker in logs. The model also includes controls for age group×gender, year dummies, log firm employment, log establishment employment, and firm age squared. Standard errors clustered by firm-year.

The hump shape is retained when we estimate the elasticity of earnings with respect to software capital per worker within job-spell. The first column of Table 2 reports the estimated elasticity by age group. The model includes the same covariates as in Figure 4 (adjusted to accommodate for collinearity with the fixed effects, e.g. by dropping indicators for education and the linear term for firm age). Only the oldest workers (65+) have lower earnings in years with higher software intensity. The largest earnings premiums associated with software capital per worker accrue to workers in the 30-49 age range, with the highest elasticity for workers between 35 and 39 years of age. Notably, controlling for individual and match specific fixed effects in addition to the firm fixed effect in Table 2 has a positive impact on the estimated coefficients among the young (particularly below 30 years of age), and a negative impact on the coefficients for older workers (from 40 years onwards), compared to the previous results from Figure 4.

Bargaining Power or Differential Productivity Impact?

The different effects estimated by age group clearly suggests that software is complementary to workers in the 30-49 age group, enhancing their productivity; but older workers appear to gain less and even lose out when the use of software goes up. However, an alternative

hypothesis could be that the productivity gains accrue to all, but only workers in the middle of the age distribution are able to capture part of these productivity gains in terms of higher wages. Both young and older workers may have lower bargaining power than workers in the middle may: the young may be easily replaceable by other youth of similar expected qualifications, and the old may lack credible alternatives outside of the firm, perhaps due to firm-specific human capital, deferred payment schemes within the firm, or different types of age discrimination in hiring.

Comparing the age profiles of equipment capital to that of software capital provides a useful contrast.³¹ Both are expected to improve revenue per worker, and a similar profile would suggest that differential bargaining power could be at work. Column 2 in Table 2 reports the profile for equipment capital per worker. We find an almost linearly increasing elasticity of earnings with respect to equipment capital by age group. It starts out with a negative elasticity for the young, and only workers above 45 years of age seem to gain from higher equipment capital, ending at almost 2 percent for the oldest. These different patterns suggest that the particular age profile of software elasticities, in particular the strong decline by age after 45, is not due simply to differential bargaining power between age groups.

A further check is provided by adding log sales per worker into the equation, also interacted with age groups. These interaction terms would capture differential bargaining power by age group. Table 3 reports the results from such a regression model, which is still the same model with job-spell fixed effects, but with an added interaction term for log sales per worker. First of all, conditional on the two capital measures and log employment (both at the firm and establishment level), the coefficients for log sales per worker shows the expected pattern, with lower coefficients for the young and the old. However, the elasticities are rather small. The coefficients and patterns of earnings premiums by age for the two capital measures hardly change from those in Table 2, estimated without log sales per worker. It appears that differential bargaining power by age group does not contribute much to our estimated pattern. The adjusted R squared remains at the same level, reinforcing this observation.

Table 3. Earnings Elasticities by Age. Dependent variable: log Earnings

Note: The table shows estimated earnings elasticities by age group, estimated from the *same model with job-spell fixed effects*. All variables measured per worker in logs. The model also includes controls for age group×gender, year dummies, log firm employment, log establishment employment, and firm age squared. Standard errors clustered by firm-year.

³¹ We have experimented with measures of equipment capital excluding computer investments, and our main results appear not to be sensitive to this change.

Patterns of Exit Among the Old

Having established that older workers gain less than workers in the middle of the age distribution from investments in software, but seem to gain more from investments in equipment, it may be instructive to investigate how exit patterns by age vary with investment in the two types of capital. To do this, we estimate a simple linear model of the probability of exit in the subsequent period, conditional on the two forms of capital per worker interacted with age.

In 2012, 21% of workers 45 and over exit their job, with one-third taking a new job at another one of our ACES firms within our 23 covered states. We treat the other 2/3rd as leaving the labor force.³² As exit is defined as exit from our ACES sample from the LEHD states, it is not a clean measure of retirement or exits out of the labor force, and as job-to-job moves are defined as mobility to another firm within our ACES sample, it is not a comprehensive measure of job-to-job moves.³³ Still, both measures are well defined and do count exits from the firms in the LEHD states, and even if we cannot completely rule it out, the bias arising from leakages out of the sample is not likely to be related to the interaction between software investments and age. The inclusion of job-spell fixed effects should alleviate some of these concerns, as well.

Figure 5. Exit and job-to-job mobility patterns for older workers (45+). Estimated coefficients for software and equipment by age.

Note: The figure shows the estimated coefficients by age for software capital per worker and equipment capital per worker in regressions of exit and job-to-job mobility. Dependent variables: *Exit* is measured as the last year an individual is observed in our ACES sample (set to missing the last year of our panel). *Job-to-Job Move* is measured as the last year an individual is observed in a firm, conditional on being employed in our ACES sample next year (set to missing the last year of our panel). The analysis includes individuals older than 44 years of age only. The linear probability models include job-spell fixed effects, age groups interacted with gender, log firm and establishment employment, firm age squared, and year dummies.

For software, we find a negative but upward-sloping impact on exits for all age groups, approaching zero for the 65+ group. This pattern is highly consistent with the earnings gains that workers of different ages obtain from software investments. Software investment tends to reduce exit, but the magnitude of the effect is diminishing with increases in worker age. Job-

³²The LEHD ICF national indicator of individuals' employment in any state can be used to separate exits into those leaving the labor force versus those taking a new job. See Hahn, Hyatt, Janicki and Tibbets (2017) and Hyatt, McEntarfer, Ueda and Zhang (2018) for details on using the LEHD data for job-to-job transitions.

³³ In the next version of this paper, exit and job-to-job mobility will be defined relative to employment in any firm in any state. Because of the temporary closure of the RDC in connection with the Covid-19 virus, we have not been able to undertake this exercise for the current version of this paper.

to-job moves are positively impacted, and also with an upward slope. One interpretation consistent with our findings on earnings is that workers become more attractive to other firms as software capital goes up,³⁴ and that older workers take the options more often because they are not as generously rewarded from software capital in their current firms.³⁵ The pattern for equipment is opposite, in line with the opposing patterns in earnings.

High-Wage Workers and High-Wage Firms: Polarization or SBTC

To explore the extent to which investments in software capital benefit workers located at different places in the earnings distribution, we first use an AKM decomposition on the entire labor market,³⁶ and place each worker in quintiles of the distribution across individuals as well as in quintiles of the distribution across firms³⁷ in our ACES sample. To ease the discussion, we follow the terminology of AKM and refer to workers in the top (bottom) quintile among individual fixed effects as high-wage (low-wage) workers, and to workers in firms in the top (bottom) quintile of the firm fixed effects as workers in high-wage (low-wage) firms.

Next, we use our ACES sample, estimate the log earnings equation including software and equipment capital per worker, interacted with age groups as before, and add interactions between the capital measures and quintiles of worker and firm fixed effects. The results are displayed in Figure 6.³⁸

The elasticity of earnings with respect to software is strongly increasing in the workers' rank in the distributions of both the individual and firm fixed effects. The left panel shows the pattern across quintiles of individual fixed effects. The lines are drawn for workers in the middle of the firm distribution: high-wage workers in median firms display an elasticity of 3 percent, whereas low-wage workers in median firms display an elasticity of 0.7 percent.

³⁴ This would be consistent with findings in Agrawal and Tambe (2016), which leverages linked employer-employee data to provide evidence that IT investment associated with private equity acquisitions creates transferable IT-related human capital.

³⁵ The differential results for the oldest workers may be related to prior findings focused on training and firm-specific human capital. Bartel et al. (2007) and Battisti et al. (2020) report complementarities between ICT and on-the-job training. Workers in the top age group in our data may not have the runway to invest in sufficient training or learning-by-doing with new software prior to retirement (Friedberg 2003), and prefer to stay where firm-specific human capital is most valuable (e.g., Violante 2002).

³⁶ This means all jobs in our sample of LEHD states are included, and not only those that are participants in the ACES sample. The AKM decomposition is done after estimating a step one regression with job-spell fixed effects, conditional on year fixed effects and age squared, and then a step two, regressing the job-spell fixed effects on a dummy for each age measured at year 2000.

³⁷ The quintiles of both workers and firms fixed effects are calculated over all observations of workers in the ACES sample of our LEHD states (approximately 161 million observations)

³⁸ Regression results available on request.

The right panel shows the pattern across quintiles of firm fixed effects. The pattern is very similar: workers in high-wage firms display a much larger elasticity than do workers in low-wage firms, suggesting that software investments are more productive in high-productivity firms.

Since the lines in each panel are drawn for median values (3rd quintile) in the other panel, and the model is additive linear, the effects may be added together using the difference from the median value from the other panel. For instance, high-wage workers in high-wage firms display an elasticity of approximately 5 percent ($0.03 + (0.032 - 0.012) = 0.05$), while low-wage workers in low-wage firms have an elasticity of approximately zero ($0.007 + (0.005 - 0.012) = 0$).

Figure 6. Earnings Elasticities by Quintiles of Individual and Firm Fixed Effects.

Note: The figure shows the elasticity of earnings with respect to the two types of capital, estimated by quintiles of individual (left panel) and firm (right panel) fixed effects. The firm and individual fixed effects displayed along the x-axis are estimated on all workers in the economy (LEHD states), but the quintiles are calculated on our ACES sample. The model also allows for separate age effects, and includes the same covariates as in Figure 5 above. Estimated with job-spell fixed effects. The effects are calculated for a worker between 35 and 39 years of age, and placed in quintile 3 (median) of the other panel (i.e. the elasticity for high-wage workers is calculated for a worker in a median firm).

The estimated individual effects may be decomposed into an observed part and an unobserved part. In a second-step regression, we estimate the fixed effect as a function of individual fixed characteristics, such as the level of education and gender, and retain the residual. The patterns for the quintiles of residual fixed effects are very similar to the ones shown for the total fixed effects. Thus, the pattern displayed in Figure 3 mainly arises from different elasticities for workers in different quintiles in the distribution of the unobserved fixed effects.

The estimated firm fixed effects may similarly be decomposed into an observed part, such as industry, and an unobserved part. The pattern for the quintiles of the residual fixed effects are similar, but somewhat attenuated compared to the pattern displayed for the total fixed effects. Again, the coefficients and standard errors are reported in the appendix. The pattern displayed in Figure 6, thus, for the most part, arises from differences between quintiles in the distribution of the unobserved firm fixed effects.

The pattern for equipment capital is quite the opposite: only the first three quintiles of the distributions across both individuals and firms display positive coefficients. Overall, this suggests that increasing software investment is likely to increase earnings inequality: high-wage workers in high-wage firms gain the most, whereas low-wage workers in low-wage firms gain

the least. Recall that this pattern shows up conditional on the age profile of the elasticity, displayed earlier, so that we need to take the age-earnings profile into account when we consider the full effects on inequality.

Time-Varying Firm Effects

A potential worry using the AKM model with firm fixed effects is possible correlations between shocks in the error term and changes in software capital. One conceivable scenario could be that the firm chooses increased software investments when it experiences a positive shock in sales. This would be an endogenous response, possibly leading to a positive bias in the estimated impact of software investments. If such shocks are correlated with the age profile of the firm and its workers, for instance that they more often occur in expanding firms, with a younger workforce, and rarely in older firms, with an ageing work force, then the estimated age profiles could be biased as well.

Another, related, concern is that much prior research points to the complementarities between ICT investment and changes in organizational design and management practices (e.g. Bartel et al. 2007, Bresnahan et al. 2002, Brynjolfsson and McElheran 2019, Battisti et al. 2020). Similar to most studies in this area, we do not observe these shifts, much less important intangible increases in complementary organizational capital (e.g., Saunders and Brynjolfsson 2016). This omission could lead us to overestimate the productivity gains – and associated wage increases – associated with ICT investment. Time-varying firm controls should absorb such effects.

As outlined in the methodology section, we control for such firm-specific shocks, and all other time- and firm-specific confounding factors, by adding a time varying, firm effect to the model.³⁹ Adding time-varying firm effects to the model enables us to identify the differential response of each age group to a reference age group, whereas the response of the reference age group will be absorbed in the time-varying firm-fixed effect. Table 4 presents the results.

We have chosen workers from 25 to 49 years of age as the reference group,⁴⁰ and the coefficients obtained in step 1 thus measure the difference between the elasticity of the particular age group and the elasticity of the 25-49 years old. We find that the youngest workers

³⁹ See Engbom and Moser (2020) and Lachowska et al (2020) for a discussion of the identification of time-varying firm effects.

⁴⁰ We experimented with excluding the two first groups in the reference group: 25-29 and 30-34, from the reference group in step 1, but obtained larger coefficients for those groups than for the reference groups, rather than smaller, and chose to include them in the reference group. Results available on request.

(20-24 years old) have a return to software capital which is .0056 lower than that of the reference group. We compare this difference to the difference between the coefficients for the youngest group and the 25-49 years old in the fixed job-spell model, previously reported in Table 2. Using the average estimate for 25-49 years old in the job-spell fixed effect model (0.0185), the difference is almost four times larger: $-0.0018 - 0.0185 = -0.0203$ compared to -0.0056 here. For older workers, on the other hand, the differences from the 25-49 group are practically the same as in the job-spell model.

The results for equipment capital are much less pronounced than in the job-spell model. The return for the young is very similar to that of the reference group, and the increase in the return for 50+ is much more modest than what we found in Table 2.

Table 4. Earnings elasticities from different job-spell FE models

Note: *Step 1 model*: Unit of observations: Individuals. Dependent variable: ln Earnings. The model also includes controls for age group X gender, job-spell fixed effects and time-varying firm effects. *Step 2 model*: Unit of observation: Firms. Dependent variable: Firm X year fixed effects from step 1. The model also includes year dummies, log firm employment, log establishment employment, firm age squared and firm fixed effects.

We need to set the level of the elasticities for the reference group. This is done in step 2 as outlined in the methodology section, by estimating the relationship between the firm-year effects estimated in step 1 and the log of the two capital measures, using firm-year as the unit of observation. Note that while the relative impact of capital intensity on different age groups within the firm, estimated in step 1, uses time-varying firm effects, the second step is conducted with firm fixed effects only. This means that the effect for the reference group, which determines the level of the elasticities, is only identified with standard AKM fixed effects assumptions.⁴¹ The last two columns of Table 4 show the results. The reported coefficients represent the elasticities for the reference age group (25-49 of age). We find elasticities of 0.0293 and -0.0001 for software and equipment respectively, to be compared with the averages⁴² of 0.0185 and -0.0031 in the job-spell model from Table 2. The step 2 models include firm fixed effects and thus effectively include controls for both individual and match effects, from step 1, and firm fixed effects from step 2.

For equipment capital, including time-varying firm effects changed the results also for older workers. The average coefficient for the reference group in the fixed job-spell model from

⁴¹ It is possible to address this concern using IV- or other methods in the step 2 regression, estimated at the firm-year level.

⁴² Calculated from the averages of the coefficients for software and equipment for age groups 25-29, 30-34, 35-39, 40-44, and 45-49 in Table 2.

Table 2 is -.0031. Compared to the job-spell model, the age profile of the elasticity of equipment capital flattens out as firm year effects are included.

Figure 7. Elasticities estimated with and without time-varying firm effects.

Note: Dependent variable: ln Earnings. The Job FE models show coefficients from Table 2, estimated with job-spell fixed effects. The Job+Firm_yr models show coefficients from step 1 and step 2 of Table 4, where time-varying firm effects are added in a step 1.

Figure 7 compares the implied elasticities from step 1 and step 2 (Job+Firm_yr) to the estimated elasticities from the job-spell model (Job FE). The difference in overall level arises from the difference in the estimated step 2 coefficients to the average effects estimated for the groups between 25 and 49 years of age. The pattern for the young and for the old arises from adding the coefficients in step 1. For software, we clearly see that while the pattern for the young is different between the models, with much smaller decline in the firm-year model, the pattern for the old is very much the same, with a steady decline from 50 years of age. For equipment, when estimated with firm-year effects, the earnings elasticity is small, and the age pattern much more moderate.

The changing pattern between specifications for the very young (20-24), suggests that there is a difference in the bias arising from the fixed firm effect assumption between age groups. Such differences could arise from heterogeneity in terms of the correlation between productivity shocks and software investments across firms with different age profiles. Most likely, there is an overrepresentation of young workers in certain segments of firms. We have already noted that they are overrepresented in low-software intensity firms (Figure 2), many are students, and they are also more likely to work in low-skill firms, younger firms, and in low-wage firms. It may be that in these segments of the labor market, the correlation between productivity shocks and software investments is more negative, inducing a negative bias for the young in the job-spell specification.

6. Conclusion and Discussion

Combining data on software investments by firms with LEHD unemployment insurance records on the age and earnings of workers, and using an AKM decomposition framework to deal with econometric problems due to the potential impact of unobservable productivity on firms investment in software and of unobservable individual characteristics of workers in

software intensive firms, we find that investments in software capital twist the demand curve for labor toward workers 25-49 compared to older workers and younger workers.

Specifically, we find that:

1) Firm investments in software capital raise the earnings of workers in the age group of 25-49 by an elasticity of about 3 percent, but the earnings premium declines almost linearly from 50 years of age through age 65 to essentially nothing for those 65 and older.

2) Higher software capital is also associated with lower exits among workers, but again with less impact on older workers. The share of employment of older workers does not vary much among firms with differing investments in software while high investors in software hire relatively fewer 20-24 year-old workers.

As best we can tell, the smaller gains for older workers from working with more software capital is because their skills are less complementary to software technologies than younger workers, as opposed to having lower bargaining power.

3) The biggest beneficiaries of software investments among workers inside a firm are those with high unobserved (fixed effects) earnings characteristics, whose wages are raised more by software investments than other workers' wages.

4) The biggest beneficiaries of software investments among workers across firms are those in firms with high unobserved or observed time invariant (fixed effects) earnings characteristics.

Our analysis augments the AKM framework in two directions that may be useful in other studies of worker-firm panel data. First, we added a match-specific term to the earnings equation, by estimating the effects of software on earnings within job-spell. This effectively removes a potential correlation between the match effects and age, which is important if workers' match improves over time in the labor market. Second, we added time-varying firm fixed effects to the model. This allows us to estimate the relative impact of software on older workers, relative to younger workers, with full control for time-varying productivity shocks or changes to the bargaining power of workers.

Overall, by raising earnings most for high-wage workers in the middle of the age distribution within firms, and most for workers in high-wage firms, software investments tend to increase earnings inequality. It is a form of skill-biased technical change that operates both within and across firms.

REFERENCES

- Abowd, J. M., J. Haltiwanger and J. Lane (2004). Integrated longitudinal employer-employee data for the United States. *American Economic Review* 94(2), 224–29.
- Abowd, J. M., F. Kramarz and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67, 251–333.
- Abowd, J. M., B. Stephens, L. Vilhuber, F. Andersson, K. L. McKinney, M. Roemer and S. Woodcock (2009). The LEHD infrastructure files and the creation of the quarterly workforce indicators. In T. Dunne, J. B. Jensen and M. J. Roberts (Eds.), *Producer Dynamics: New Evidence from Micro Data*, 149–230. Chicago, IL: University of Chicago Press.
- Acemoglu, D., and D. Autor (2011). Skills, tasks and technologies: implications for employment and earnings. In O. Ashenfelter and D. Card (Eds.), *The Handbook of Labor Economics Volume IV*, 1043–1171. Amsterdam: Elsevier.
- Acemoglu, D. and P. Restrepo (2019) Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*. 33(2): 3-30.
- Agrawal, A. and P. Tambe (2016). Private equity and workers’ career paths: the role of technological change. *Review of Financial Studies* 29, 2455–89.
- Ahituv, A. and J. Zeira (2011). Technical progress and early retirement. *Economic Journal* 121, 171–93.
- Akerman, A., I. Gaarder and M. Mogstad (2015). The skill complementarity of broadband internet. *Quarterly Journal of Economics* 130, 1781–824.
- Almeida, R. K., A. M. Fernandes and M. Viollaz (2020). Software adoption, employment composition, and the skill content of occupations in Chilean firms. *Journal of Development Studies* 56, 169–85.
- Atalay, E., P. Phongthientham, S. Sotelo and D. Tannenbaum (2018). New technologies and the labor market. *Journal of Monetary Economics* 97, 48–67.
- Aubert, P., E. Caroli and M. Roger (2006). New technologies, organisation and age: Firm-level evidence. *Economic Journal* 116, F73–F93.
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives* 29, 3–30.
- Autor, D. H. and D. Dorn (2009). This job is “getting old”: Measuring changes in job opportunities using occupational age structure *American Economic Review* 99(2), 45-51.
- Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the U.S. labor market. *American Economic Review* 103(5), 1553–97.

- Autor, D. H., L. F. Katz and M. S. Kearney (2008). Trends in U.S. wage inequality: revising the revisionists. *Review of Economics and Statistics* 90, 300–23.
- Autor, D. H., L. F. Katz and A. B. Krueger (1998). Computing inequality: have computers changed the labor market? *Quarterly Journal of Economics* 113, 1169–213.
- Autor, D. H., F. Levy and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118, 1279–333.
- Bartel, A. P., C. Ichniowski and K. L. Shaw (2007). How does information technology really affect productivity? Plant-level comparisons of product innovation, process improvement and worker skills. *Quarterly Journal of Economics* 122, 1721–1758.
- Bartel, A. P. and N. Sicherman (1993). Technological change and retirement decisions of older workers, *Journal of Labor Economics* 11, 162–83.
- Barth E., J.C. Davis and R. Freeman (2018). Augmenting the human capital earnings equation with measures of where people work, *Journal of Labor Economics*, Vol 36(S1):S71-S97.
- Barth, E., A. Bryson, J. C. Davis and R. Freeman (2016). It's where you work: increases in the dispersion of earnings across establishments and individuals in the United States *Journal of Labor Economics* 34, S67–S97.
- Battisti, M., C. Dustmann and U. Shonberg (2020). Technological and organizational change and the careers of workers. Working Paper, University of Glasgow.
- Beckmann, M. (2007). Age-biased technological and organizational change: firm-level evidence and management implications. Discussion Paper 05/07, Wirtschaftswissenschaftliches Zentrum, University of Basel.
- Behaghel, L., E. Caroli and M. Roger (2014). Age-biased technical and organizational change, training and employment prospects of older workers. *Economica* 81, 368–89.
- Behaghel, L., and N. Greenan (2010). Training and age-biased technical change. *Annals of Economics and Statistics* 99/100, 317–42.
- Bessen, J.E., M. Goos, A. Salomons, and W. van den Berge (2019) Automatic Reaction: What Happens to Workers at Firms that Automate? Boston Univ. School of Law, Law and Economics Research Paper No. 19-2 (2019). Available at https://scholarship.law.bu.edu/cgi/viewcontent.cgi?article=1585&context=faculty_scholarship.
- Bessen, J. E. and C. Righi (2020) Information Technology and Firm Employment. Boston Univ. School of Law, Law and Economics Research Paper No. 19-6 (2019). Available at SSRN: <https://ssrn.com/abstract=3371016> or <http://dx.doi.org/10.2139/ssrn.3371016>.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012). Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review* 102(1): 167-201.

- Bresnahan, Tim and Shane Greenstein (1996), Technical Progress and Co-Invention in Computing and in the Use of Computers. *Brookings Papers on Economics Activity: Microeconomics*, 1-78.
- Brynjolfsson, Erik and Lorin Hitt (1996). Paradox Lost? Firm-level Evidence on the Returns to Information Systems Spending. *Management Science* 42(4): 541-558.
- Brynjolfsson, Erik and Lorin Hitt (2003). Computing Productivity: Firm-level Evidence. *Review of Economics and Statistics* 85(4): 793-808.
- Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: firm-level evidence. *Quarterly Journal of Economics* 117, 339–76.
- Byrne, D. and C. Corrado (2017). ICT services and their prices: what do they tell us about productivity and technology? Working Paper, Finance and Economics Discussion Series 2017-015, Federal Reserve Board, Washington, DC.
- Brynjolfsson, E. and A. McAfee (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W. W. Norton & Company.
- Brynjolfsson, E. and K. McElheran (2019). Data in action: data-driven decision making and predictive analytics in U.S. manufacturing. Working Paper 3422397, University of Toronto.
- Brynjolfsson, E. and P. Milgrom P (2013) Complementarity in organizations. *The Handbook of Organizational Economics*. 11-55.
- Card, D., A.R. Cardoso, J. Heining and P. Kline (2018). Firms and labor market inequality: Evidence and some theory *Journal of Labor Economics* 36(S1): S13-S70.
- Card, D., J. Heining and P. Kline (2013). Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics* 128: 967–1015.
- Chari, V. V. and H. Hopenhayn, H. (1991) Vintage human capital, growth, and the diffusion of new technology, *Journal of political Economy* 99(6), 1142–1165.
- Cornwell, C., I.M. Shmutte, and D. Scur (2019). "Building a Productive Workforce: The Role of Structured Management Practices," CEP Discussion Papers dp1644, Centre for Economic Performance, LSE.
- Deming, D.J. and K.L. Noray. (forthcoming) Earnings dynamics, changing job skills, and STEM careers. *Quarterly Journal of Economics*.
- Dillender, M. and E. Forsythe (2019). Computerization of white collar jobs. Working Paper 19-310, W. E. Upjohn Institute for Employment Research.
- Engbom, N. and C. Moser (2020). Firm pay dynamics. Working Paper 3531250, SSRN.

- Forman, C., A. Goldfarb and S. Greenstein (2012). The internet and local wages: a puzzle. *American Economic Review* 102(1), 556–75.
- Friedberg, L. (2003). The impact of technological change on older workers: evidence from data on computer use. *Industrial and Labor Relations Review* 56, 511–29.
- Gaggl, Paul, and Greg C. Wright (2017). A Short-Run View of What Computers Do: Evidence from a UK Tax Incentive. *American Economic Journal: Applied Economics* 9 (3): 262-94.
- Goldin, C. D. and L. F. Katz (2009). *The Race between Education and Technology*. Cambridge, MA: Harvard University Press.
- Grimm, B. T., B. R. Moulton and D. B. Wasshausen (2005). Information-processing equipment and software in the national accounts. In C. Corrado, J. Haltiwanger and D. Sichel (Eds.), *Measuring Capital in the New Economy*, 363–402. Chicago, IL: University of Chicago Press.
- Hægeland, T., D. Rønningen and K. G. Salvanes (2007). Adapt or withdraw? Evidence on technological changes and early retirement using matched worker-firm data. Discussion Paper 509, Research Department, Statistics Norway.
- Hahn, J., H. Hyatt, H. P. Janicki and S. Tibbets (2017). Job-to-job flows and earnings growth. *American Economic Review* 107(5), 358–63.
- Harrigan, J, Reshef A and Toubal F (2016) The march of the techies: Technology, trade, and job polarization in france, 1994-2007. *National Bureau of Economic Research Working Paper*.
- Hershbein, B. and L. B. Kahn (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review* 108(7), 1737–72.
- Hu, X. and R. B. Freeman (2020). Software usage by occupation. Unpublished.
- Hyatt, H., E. McEntarfer, K. Ueda and A. Zhang (2018). Interstate migration and employer-to-employer transitions in the U.S.: new evidence from administrative records data. *Demography* 55, 2161–80.
- Jin, W. and K. McElheran (2017). Economies before scale: learning, survival and performance of young plants in the age of cloud computing. Working Paper 3112901, SSRN.
- Johansen, L. (1959) Substitution versus fixed production coefficients in the theory of economic growth: a synthesis. *Econometrica* 27, 157–76.
- Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963–1987: supply and demand factors. *Quarterly Journal of Economics* 107, 35–78.

- Krueger, A. B. (1993). How computers have changed the wage structure: evidence from microdata, 1984–1989. *Quarterly Journal of Economics* 108, 33–60.
- Lachowska M., A. Mas, D. S. Raffaele and S. A. Woodbury (2020). Do firm effects drift? Working Paper 26653, National Bureau of Economic Research.
- Manning, Alan (2003) *Monopsony in Motion*, Princeton University Press
- McElheran, K. (2015). Do market leaders lead in business process innovation? The case(s) of e-business adoption. *Management Science* 61, 1197–216.
- Michaels, G., A. Natraj, and J. Van Reenen (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics* 96, 60–77.
- Moylan, C. E. (2001). Estimation of software in the U.S. national income and product accounts: new developments. Bureau of Economic Analysis, U.S. Department of Commerce.
- Paccagnella, M. (2016) “Age, Ageing and Skills: Results from the Survey of Adult Skills”, *OECD Education Working Papers*, No. 132, OECD Publishing, Paris.
<http://dx.doi.org/10.1787/5jm0q1n38lvc-en>
- Poliquin, C. W. (2020). The wage and inequality impacts of broadband internet. Working Paper, University of California, Los Angeles.
- Rosen, S. (1975) *Measuring the obsolescence of knowledge, Education, income, and human behavior*, NBER, pp. 199–232.
- Reed, K. (2015). Capitalization of software development costs, Memorandum 201549024, Office of Chief Counsel, Internal Revenue Service.
- Saunders, A. and E. Brynjolfsson (2016). Valuing IT-related intangible assets. *MIS Quarterly* 40, 83–110.
- Schøne, P. (2009). New technologies, new work practices, and the age structure of the workers: correlates or causality? *Journal of Population Economics* 22, 803–26.
- Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2019). Firming Up Inequality. *The Quarterly Journal of Economics* 134(1): 1-50.
- Syverson, C. (2011) What determines productivity? *Journal of Economic Literature* 49, 326–65.
- Tambe, P. and L. M. Hitt (2012). The productivity of information technology investments: new evidence from IT labor data. *Information Systems Research* 23, 599–617.

- Van Reenen, John (1996). The Creation and Capture of Economic Rents: Wages and Innovation in a Panel of UK Companies. *Quarterly Journal of Economics* 111(1): 195-226.
- Vilhuber, L. (2018). LEHD infrastructure S2014 files in the FSRDC. Working Paper CES-WP-18-27, Census Center for Economic Studies.
- Violante, G. L. (2002). Technological acceleration, skill transferability, and the rise in residual inequality. *Quarterly Journal of Economics* 117, 297–338.
- Webb, M. (2020). The impact of artificial intelligence on the labor market. Working Paper, Stanford University.
- Weinberg, B. A. (2004). Experience and technology adoption, Discussion Paper 1051, Institute for the Study of Labor (IZA), Bonn.
- Zolas, N., Z. Kroff, E. Brynjolfsson, K. McElheran, D. Beede, C. Buffington, N. Goldschlag, L. Foster, and E. Dinlersoz (2020). Advanced technology adoption and use by U.S. firms: Evidence from the Annual Business Survey. Mimeo, Center for Economic Studies, U.S. Census Bureau.

Appendix A: Relation between ACES capital investment in software and total software spending

There is a striking difference between aggregate ACES national totals for capitalized software investment and all software investment reported in the National Income and Product Accounts (NIPA) for private fixed investment.⁴³ The NIPA data show about three times the spending as the ACES data. Part of the aggregate difference is that NIPA includes investment by government and farms that are excluded by ACES, which surveys non-farm private firms. From the BEA 2012 benchmark input-output tables, this is approximately 6% of software expenditures. In addition, another 21% is accounted for by software exports. But most of the difference between the two series occurs because firms report most software expenditures as direct purchase of intermediate inputs rather than as capitalized investment. The ACES Information & Communication Technology Survey (ICTS), available for some years and now discontinued, asked respondents for expensed software expenditures. ICTS noncapitalized software reporting accounts for 25-35% of the NIPA figures depending on the year. Thus, in summary for 2013, ACES capitalized software expenditures represent 32% of the NIPA figures, ICTS noncapitalized software 28%, and from the BEA input-output tables, exports approximately another 21%, and agriculture and government approximately 6%. There remains a difference of 14%, potentially due to the differences in measurement approaches as discussed below. Another potential difference could be related to the treatment of software as an intermediate material input. An example of this would be when a computer laptop vendor bundles the operating system as part of the hardware product sold. The operating system is purchased from a software supplier and embedded into the hardware product for resale. These material inputs would not be reported by firms in the ACES nor ICTS survey.

Grimm, Moulton and Wasshausen (2005) and Moylan (2001) describe the way the ACES and NIPA differ in the sources of their estimate. ACES is a “demand-side” survey of firms’ capital expenditures.⁴⁴ NIPA uses a commodity-flow “supply-side” approach that trace commodities from their domestic production and imports to their final purchase. A software producing firm would report all of its sales as software without distinguishing between purchases that a customer might capitalize and those that it would expense. Because IRS regulations allow firms to treat smaller expenditures that could be capitalized as expenses, many choose to expense them. According to Grimm et al (2005) “Internal Revenue Service (IRS) regulations allow for low-value items (under \$17,500 for 1998) that fit the criteria for capital investment to be expensed. It is possible that much software falls into this category and that ACES respondents follow IRS guidelines when determining what is a capital investment” ... (suggesting) “that capital expenditures for software are significantly understated in the ACES estimates.” Moylan notes that “For software expenditures to be capitalized, firms must view these expenditures as significant and they must have a useful life of more than one year. Annual site licenses are expensed, but multi-year licenses should be capitalized. Firms decide for themselves what is maintenance and what is a major improvement that requires capitalization.” In 2015 the IRS issued a Chief Counsel Advice (CCA) related to the tax treatment of software development costs (Reed, 2015). This summarized that software purchase and customization are capitalized, and internal software development costs are deductible as current expenses. In this CCA guidance however, IRS also cautioned that some taxpayers were improperly expensing all software costs.

Another difference between ACES and NIPA data is how they treat internally developed software. Since 2002, ACES breaks out a category of internally-developed capitalized computer software, but there is substantial non-response for this item. This is

⁴³St. Louis Federal Reserve Bank FRED series B985RC1A027NBEA.

⁴⁴BEA uses ACES data to allocate by industry its total investment and investment by type of asset estimates.

consistent with Moylan's (2001) summary that "*almost no own-account software is capitalized*, while some prepackaged and custom software are capitalized. Firms not in the business of producing software for commercial sale view own-account expenditures as an expense." and "Although in theory, prepackaged software purchases with a useful life of at least one year should be capitalized, most are treated as an expense. For example, a Fortune 500 firm said that its policy was to expense all single software purchases of \$250 or less, as well as all site licenses or combined purchases that are less than \$10,000." Grimm et al (2005) and Moylan (2001) conclude that the software investments firms elect not to capitalize on their books can explain a substantial proportion of the differences between the NIPA and ACES estimates.

Appendix B:

There are some differences between the full LEHD data for 23 states used to run an AKM decomposition for all workers and firms in those states and our analysis sample that includes all workers in the 23 LEHD states with jobs in firms in the ACES sample. In the analysis sample we have 11.6 million workers in 2012 employed in approximately 16,000 ACES firms that represent a 20% share of all LEHD workers in our 23 covered states. Our sample of ACES firms are older – 27 years in business on average – and 70% are multi-location firms in 2012; average headcount is 600 employees per firm.

Linked to our samples, in 2012 the average worker's employer had \$112,000 in sales per worker, \$29,400/worker in capital equipment (excluding software), and \$3,800 in software capital stock per worker.

Table 1. Firm Characteristics by Quintile of Software Intensity

	All	Q1	Q2	Q3	Q4	Q5
Log Software/Worker	1.35	0.01	0.32	1.01	1.89	3.5
Log Equipment/Worker	3.38	0.54	2.66	3.87	4.47	5.36
Log Firm employment	6.40	5.90	6.44	6.33	6.55	6.77
Firm Age	27.22	20.73	26.54	28.95	30.00	29.89
Multi-unit Firm	0.70	0.50	0.68	0.69	0.77	0.84
Log Sales/Worker	4.72	2.57	4.34	5.08	5.50	6.10
Less than high school	0.14	0.20	0.16	0.15	0.12	0.09
High School	0.26	0.28	0.27	0.27	0.25	0.22
Some College	0.31	0.30	0.31	0.32	0.32	0.30
College education	0.29	0.22	0.26	0.27	0.31	0.39

Note: The quintiles of software capital per worker are calculated over observations of firms (education shares over workers). Cross sectional data, LEHD states in 2012.

Table 2. Earnings Elasticities by Age. Dependent variable: log Earnings, Job-spell FE

	Elasticity of Earnings with respect to	
	Software Intensity	Equipment Intensity
20-24	-0.0018 (0.0027)	-0.0143*** (0.0017)
25-29	0.0118*** (0.0022)	-0.0078*** (0.0015)
30-34	0.0199*** (0.0018)	-0.0063*** (0.0014)
35-39	0.0221*** (0.0016)	-0.0040** (0.0014)
40-44	0.0210*** (0.0015)	-0.0007 (0.0012)
45-49	0.0179*** (0.0014)	0.0031* (0.0012)
50-55	0.0143*** (0.0014)	0.0075*** (0.0013)
55-59	0.0088*** (0.0014)	0.0123*** (0.0015)
60-65	0.0034* (0.0016)	0.0153*** (0.0017)
65 +	-0.0067*** (0.0019)	0.0186*** (0.0020)
Job-spell effects	Yes	
Adj. R squared	0.862	
N	1.61e+08	

Note: Estimated elasticities of earnings (standard errors) with respect to software and equipment capital per worker for each age group, both estimated from *the same model with job-spell fixed effects*. The model also includes controls for age group X gender, year dummies, log firm employment, log establishment employment, and firm age squared. Standard errors clustered by firm-year.

Table 3. Earnings Elasticities by Age. Dependent variable: log Earnings

	Elasticity of earnings with respect to		
	Sales per worker	Software intensity	Equipment intensity
20-24	-0.0012*** (0.0003)	-0.0009 (0.0027)	-0.0139*** (0.0017)
25-29	0.0009*** (0.0002)	0.0121*** (0.0022)	-0.0084*** (0.0015)
30-34	0.0017*** (0.0002)	0.0199*** (0.0018)	-0.0072*** (0.0015)
35-39	0.0018*** (0.0002)	0.0221*** (0.0016)	-0.0049*** (0.0014)
40-44	0.0019*** (0.0002)	0.0210*** (0.0015)	-0.0016 (0.0012)
45-49	0.0017*** (0.0002)	0.0179*** (0.0014)	0.0024* (0.0012)
50-55	0.0016*** (0.0002)	0.0142*** (0.0014)	0.0069*** (0.0013)
55-59	0.0013*** (0.0003)	0.0088*** (0.0014)	0.0119*** (0.0015)
60-65	0.0012*** (0.0003)	0.0033* (0.0016)	0.0151*** (0.0017)
65 +	0.0003 (0.0004)	-0.0066*** (0.0019)	0.0188*** (0.0020)
Job-spell effects	Yes		
R squared	0.862		
N	1.61e+08		



Note: The table shows estimated earnings elasticities by age group, estimated from the *same model with job-spell fixed effects*. All variables measured per worker in logs. The model also includes controls for age group×gender, year dummies, log firm employment, log establishment employment, and firm age squared. Standard errors clustered by firm-year.

Table 4. Earnings elasticities from different job-spell FE models.

	Step 1		Step 2	
Unit of obs.	Individuals \times year		Firm \times year	
Dep. Var.	Ln Earnings		Firm-year fixed effects from step 1	
Fixed eff.	Job + Firm-year		Firm	
Elasticity with respect to	Software intensity	Equipment intensity	Software int.	Equipment int.
20-24	-0.0056*** (0.0002)	0.0019*** (0.0001)		
25-49			0.0293*** (0.0009)	-0.0001 (0.0007)
50-55	-0.0056*** (0.0001)	0.0026*** (0.0001)		
55-59	-0.0106*** (0.0001)	0.0039*** (0.0001)		
60-65	-0.0154*** (0.0002)	0.0031*** (0.0002)		
65 +	-0.0239*** (0.0002)	0.0032*** (0.0002)		
R squared	0.857		0.560	
N	1.32e+08		2.10e+05	


Note: *Step 1 model*: Unit of observations: Individuals. Dependent variable: Ln Earnings. The model also includes controls for age group \times gender, job-spell fixed effects and year specific firm effects. *Step 2 model*: Unit of observation: Firms. Dependent variable: Firm \times year fixed effects from step 1. The model also includes year dummies, log firm employment, log establishment employment, firm age squared, and firm fixed effects.

Figure 1. Investments in Software and Equipment




Source: St Louis Fed, Fred, software is BEA Account Code: B985RC, equipment is BEA Y033RC

Figure 2. Employment Share by Age Group. By Quintiles of Software Intensity.




Note: Average employment share across firms by quintile of software capital per workers. Q1 is the first quintile, Q3 the median quintile, and Q5 the top quintile. Employment weighted quintiles. Cross sectional data, LEHD states in 2012.

Figure 3. Average Worker and Firm Fixed Effects by Quintile of Software per Worker



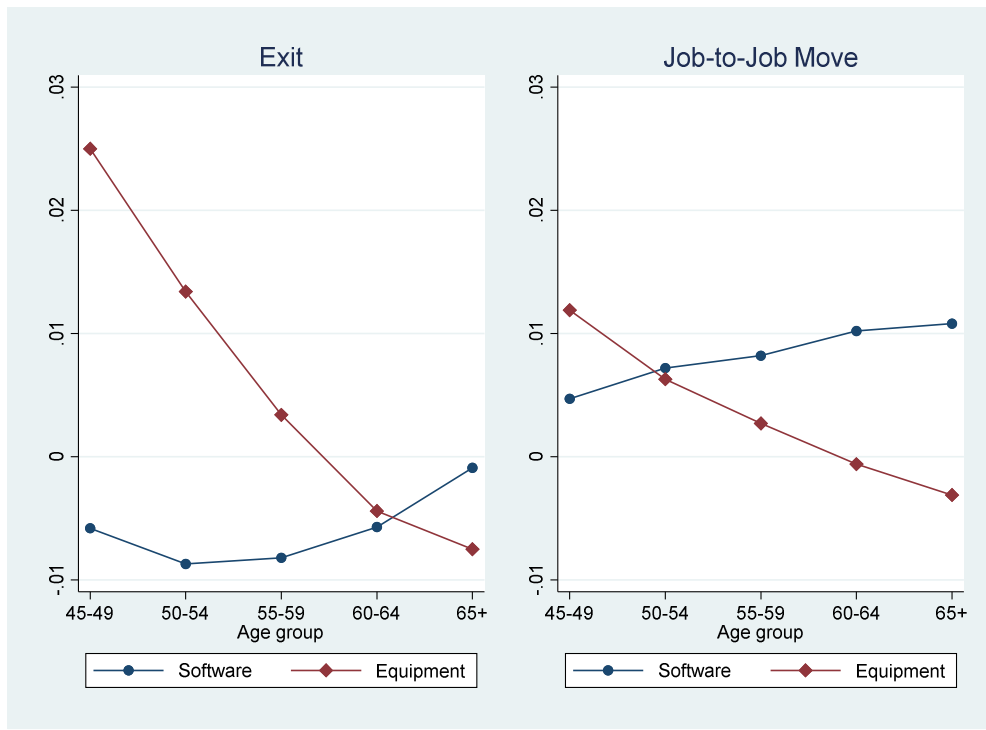
Note: Average worker and firm fixed effects across firms by quintile of software per worker. Q1 is the first quintile, Q3 the median quintile, and Q5 the top quintile. Employment weighted quintiles. Worker and firm fixed effects are first estimated from an AKM decomposition on the whole labor market of the LEHD states, controlling for worker age, age squared, and years of observation. Next, “Worker (Firm) FE res” is the residual, and “Worker (Firm) FE Xb” the predicted value, of the worker (firm) fixed effects regressed on time invariant covariates such as education, gender, and cohort (industry, average firm and establishment size, and firm age), see section 5 below.

Figure 4. Software Elasticity of Earnings by Age Group. FE: Region + Industry + Firm



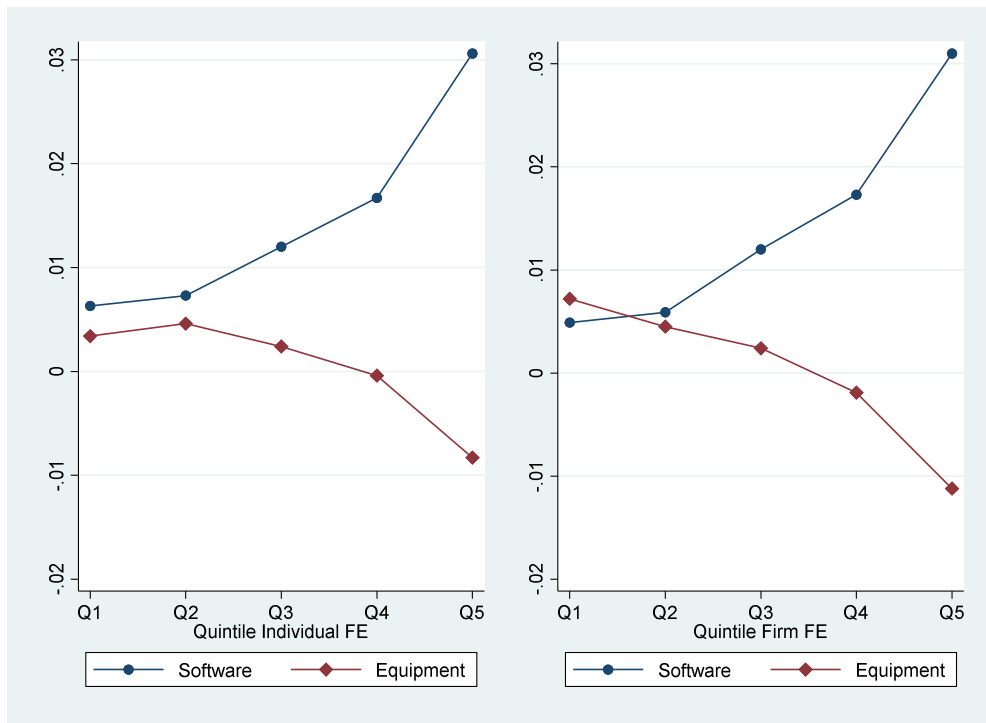
Note: Dependent variable: Log Earnings. The figure shows the coefficients for log software capital per worker interacted with ten age groups. The models control flexibly also for equipment capital by age group, and include controls for age group×gender, year dummies, three dummies for level of education, log firm employment, log establishment employment, firm age, and firm age squared. Region, industry, and firm fixed effects are added successively.

Figure 5. Exit and job-to-job mobility patterns for older workers (45+). Estimated coefficients for software and equipment by age.




Note: The figure shows the estimated coefficients by age for software capital per worker and equipment capital per worker in regressions of exit and job-to-job mobility. Dependent variables: *Exit* is measured as the last year an individual is observed in our ACES sample (set to missing the last year of our panel). *Job-to-Job Move* is measured as the last year an individual is observed in a firm, conditional on being employed in our ACES sample next year (set to missing the last year of our panel). The analysis includes individuals older than 44 years of age only. The linear probability models include job-spell fixed effects, age groups interacted with gender, log firm and establishment employment, firm age squared, and year dummies.

Figure 6. Earnings Elasticities by Quintiles of Individual and Firm Fixed Effects.



Note: The figure shows the elasticity of earnings with respect to the two types of capital, estimated by quintiles of individual (left panel) and firm (right panel) fixed effects. The firm and individual fixed effects displayed along the x-axis are estimated on all workers in the economy (LEHD states), but the quintiles are calculated on our ACES sample. The model also allows for separate age effects, and includes the same covariates as in Figure 5 above. Estimated with job-spell fixed effects. The effects are calculated for a worker between 35 and 39 years of age, and placed in quintile 3 (median) of the other panel (i.e. the elasticity for high-wage workers is calculated for a worker in a median firm).

Figure 7. Elasticities estimated with and without time-varying firm effects.



Note: Dependent variable: ln Earnings. The Job FE models show coefficients from Table 2, estimated with job-spell fixed effects. The Job+Firm_yr models show coefficients from step 1 and step 2 of Table 4, where time-varying firm effects are added in a step 1.