

NBER WORKING PAPER SERIES

COMPUTERIZATION, OBSOLESCENCE, AND THE LENGTH
OF WORKING LIFE

Péter Hudomiet
Robert J. Willis

Working Paper 28701
<http://www.nber.org/papers/w28701>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2021

Research for this paper was supported by the Alfred P. Sloan Foundation Grant G-2016-7259. An earlier version of this paper was presented at a “Working Longer” Conference at the Stanford Institute for Economic Policy Analysis. We thank Nicole Maestas and other participants for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w28701.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Péter Hudomiet and Robert J. Willis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Computerization, Obsolescence, and the Length of Working Life

Péter Hudomiet and Robert J. Willis

NBER Working Paper No. 28701

April 2021

JEL No. J14,J24,J26,J31

ABSTRACT

This paper analyzes how computerization affected the labor market outcomes of older workers between 1984 and 2017. Using the computerization supplements of the Current Population Survey (CPS) we show that different occupations were computerized at different times, older workers tended to start using computers with a delay compared to younger workers, but computer use within occupations converged to the same levels across age groups eventually. That is, there was a temporary knowledge gap between younger and older workers in most occupations. We estimate how this knowledge gap affected older workers' labor market outcomes using data from the CPS and the Health and Retirement Study. Our models control for occupation and time fixed effects and in some models; we also control for full occupation-time interactions and use middle aged (age 40-49) workers as the control group. We find strong and robust negative effects of the knowledge gap on wages, and a large, temporary increase in transitions from work to non-participation, consistent with a model of creative destruction in which the computerization of jobs made older workers' skills obsolete in birth cohorts that experienced computerization relatively late in their careers. We find larger effects on females and on middle-skilled workers.

Péter Hudomiet
RAND Corporation
1776 Main Street
Santa Monica, CA 90407
hudomiet@rand.org

Robert J. Willis
Professor of Economics, Emeritus
University of Michigan
1001 Laurel St.
Apt. 323
San Carlos, CA 94070
and NBER
rjwillis@isr.umich.edu

A data appendix is available at
<http://www.nber.org/data-appendix/w28701>

A supplementary data on knowledge gap, mapping of O*NET to occupations is available at <https://sites.google.com/site/phudomiet/research>

1. Introduction

Technological progress in the past few centuries has had an enormous impact on the labor market, shifting workers between occupations and sectors and changing the types of tasks workers are required to do. Even though it has been frequently argued that machines may eliminate human jobs altogether “rendering the population redundant”¹, the labor force has been able to adapt to all these new technologies and employment-to-population ratios did not change much in the past. In fact, technological progress overall benefitted the society by increasing workers’ productivity and wages and requiring workers to do less repetitive and strenuous work. The transitions, however, may not always have been easy for workers with specialized skills that suddenly became obsolete or not competitive with new machines or other workers. The transitions were likely most difficult for older workers, who had accumulated more specialized skills compared to their younger peers and who expected to spend a shorter time on the labor market to benefit from adopting the new technologies. When an older worker experiences the obsolescence of her skills, she may decide to retire earlier than she originally planned.

In this paper we focus on a particular innovation: computerization, arguably the most important technological change of our era, and we estimate how it affected the labor market outcomes of older workers during the last 35 years from a longitudinal perspective that emphasizes the potential mismatch between the computer-related skills possessed by workers and the skills needed to carry out the tasks demanded by employers of workers in those occupations. Our main research question is: “How has the advance of computer technology and the introduction of personal computers (PCs) and the Internet since the 1980s affected the

¹This expert opinion about the “the substitution of machinery for human labour” was rendered at the beginning of the 19th century (*The Economist*, June 25-July 1, 2016).

decisions of older workers about retirement and the length of working life in those occupations that use computers most intensively relative to workers in less computer intensive jobs and why?” We define the state of computer technology broadly to include the current capabilities of existing types of hardware (e.g., mainframes, PCs, mobile phones); software including word processors, spreadsheets, statistical programs, graphics programs; storage including floppy disks, hard drives, cloud; internet bandwidth; connectivity and communication through the internet, etc.

Over time the astounding increases in CPU speed, storage capacity and communication speeds has enabled a proliferation of ways of producing new and existing products and services which, in turn, has had major effects on the distribution of employment and wages across occupations. Autor, et al. (2003) argue that these changes stem from the fact that worker productivity depends on the set of tasks required by their job and that the productivity of computers and software depends on the tasks that computer routines can be programmed to carry out. The impact of computerization on the distribution of employment and wages depends on both supply-side effects on worker productivity and demand-side factors that influence own- and cross-elasticities of demand for labor of each type.

Computerization affected birth cohorts at different part of their life cycles. We define “Retiree Cohorts” as those birth cohorts who were first affected by computerization on the labor market, after they finished their formal education. By 2017 people who had reached age 65 and thus were eligible for both Social Security and Medicare benefits were born in 1952. These people reached high school graduation at age 18 in 1970 at a time when mainframe computers were just beginning to be used for scientific and accounting work in large organizations. This implies that nearly all Americans born before 1952 who were observed to retire before 2017—that is, nearly all Americans who have ever retired—finished their high school degree before

computerization became a significant influence in the job market. In addition, as we show in this paper, the introduction of the IBM PC in 1982 and shortly after by the Internet was associated with an acceleration of computer use at work and home. This cohort perspective implies that the human capital developed by these Retiree Cohorts during their education and early careers may have been poorly suited for jobs that required the use of computers. Facing competition from younger cohorts, we show that members of the Retiree Cohorts who did not learn to use computers in occupations in which computer use was prevalent suffered wage decreases which reduced their incomes and the marginal value of work. In effect, computerization was a form of creative destruction that imposed a loss of the value of human capital on members of the Retiree Cohorts, likely lowering the capital value of the lifetime earnings of the Retiree Cohorts through obsolescence of their human capital, while creating new opportunities for younger cohorts.

To our knowledge, the literature on the growing inequality of wages and the polarization of the job market has not addressed the implications of computerization for the labor market and retirement behavior of older workers.² However, the analysis in this paper is related to a wide range of theoretical and empirical literature in labor economics on inequality and polarization of occupations, education and wage.

In the empirical part of this paper, we first track the fraction of workers who used computers and the Internet at work and at home between 1984 and 2017 in detailed occupation and age groups, using the computerization supplements of the Current Population Survey (CPS). We measure the size of the computerization shock by the gap between the fraction of younger

² Prior literature has investigated the effects of computerization on wage-inequalities and employment patterns in different age groups (Burstein et al., 2019, Belbase and Chen, 2019). A few articles studied the effect of PC use at work on labor supply, but the results were mixed (Biagi et al., 2013; Friedberg, 2003; Schleife, 2006). Another line of research investigated the effect of IT investments by firms on employment and wages, and generally found strong associations (Aubert et al., 2006; Gaggle and Wright, 2017; Freeman, 2018).

and older workers in a given year and occupation who used computers at work. We show that this “*knowledge gap*” tended to be larger in highly skilled occupations that were computerized earlier, but importantly it was temporary, and it largely disappeared by 2017, as computer use among older workers and younger workers converged. We made this knowledge gap measure available online and it can be merged to other datasets.³

We hypothesize that the knowledge gap signaled a pressure on older workers at the time to either adopt the new technology or to leave these jobs or the labor market altogether. Then we test how the knowledge gap affected the labor market outcomes of older workers at the time when it affected their own occupation and birth cohort using millions of observations from the CPS, and additional data from the Health and Retirement Study (HRS). Our empirical models use time-series variation in computerization for identification. A growing recent literature has used occupation measures, based on the O*NET or the Dictionary of Occupational Titles (DOT), to study how workers’ job characteristics affected their labor market outcomes. These measures, however, only provide cross-sectional variations based on the workers’ occupations. We, instead, directly control for skill and time fixed effects in our models. These specifications hold workers’ skills constant and compare birth cohorts that were affected by computerization at different times in their careers, while controlling for general economic conditions with year fixed effects.

We find strong evidence that computerization (i.e. the knowledge gap) temporarily shortened the working lives of older workers and it decreased their wages. The main results of the paper are robust to many methodological assumptions, such as the definition of “skills,” the type of control groups, or the use of control variables. We also look at other outcomes, such as transitions into less skilled occupations, and into part-time jobs; or being depressed. Finally, we

³ The datafiles and instructions are available at <https://sites.google.com/site/phudomiet/research>.

also carry out detailed subgroup analysis, to understand if the effect of computerization varied by demographics (such as gender) or job types. Our paper starts in Section 2 with a brief overview of the history of computerization and the relevant economic theory about how it may affect retirement. Section 3 introduces the CPS and the HRS datasets that we use in this project. Section 4 presents our main empirical results, and Section 5 concludes.

2. Background

2.1. A brief history of computerization

In this paper, we use CPS and HRS data to study how the use of computers by individual workers affected workers who were born between 1915 and 1968. In Table 1, we see that members of the earliest cohorts in our data (born around 1920) were still at the beginning of their careers when digital computers were invented in the 1950s and were already more than 15-20 years into their working careers by the time mainframe computers were commercially available in the 1960s. While mainframes were very useful in scientific research and engineering work and allowed large businesses to automate many payroll and accounting functions, the spread of business and scientific applications accelerated with the introduction of minicomputers during the 1970s. The introduction of the IBM personal computer in 1982 produced a rapid spread in the application of computers and software to existing tasks such as word processing, payrolls and accounting beyond large organizations to small businesses and households. The scale and scope of computerization was vastly enhanced by the parallel development of the Internet beginning in 1974.

The strong complementarity between innovations in hardware, software and networking led to an unprecedented spread of new services, products, smart phones, motion pictures, retail

sales and so on that reduced the demands for the traditional skills required in many occupations while increasing the demands for other kinds of skills.

Members of the 1920 cohort reached retirement age about the time that the personal computer was introduced. The first cohort of the HRS, born in 1931-41, were already between 40 and 50 years of age when the PC was introduced. Many male members of this cohort needed to learn to type before they could attempt to exploit the potential of the computer. More generally, the computer revolution was well underway before 1982 when older cohorts were first exposed to the PC and the myriad applications it spawned. Indeed, we show that virtually all workers who have retired since the invention of the digital computer learned to use computers on the job if they learned at all. Conversely, beginning with cohorts that graduated from high school after 1982, many high school students were aware of business and scientific applications of computers and began having access to PCs both at school and home creating expectations that influenced choices about high school curricula, college majors and occupations. As time has gone on, the most recent cohorts are first exposed to computerized devices as toddlers. In addition, the use of computers and the Internet for games, communication, shopping, hobbies, etc. has fostered ubiquitous use of computers and the Internet outside the workplace.

The diffusion of computer and Internet use at work and at home over the working life cycle by successive cohorts after the introduction of the PC is illustrated in the four panels of Figure 1. The figure shows estimates based on data from the CPS computerization supplements that will be explained in detail in the results section of the paper. Despite very large initial cohort differences in the fraction of individuals who use computers or the Internet, we see dramatic convergence to virtually constant fractions of use of computers and the Internet by successive cohorts of about 60-70 percent at work and 85-90 percent at home. Clearly, use at home grew

much more rapidly than in the workplace as the price of PCs fell over time. Since, at least in its early phase, use of the Internet required use of computers it is not surprising that the growth curves in Panels C and D of Figure 1 are almost identical. We can also see that use of computers at home and of the Internet at both work and home was strongly related to the introduction of the PC. For example, in Panel D computer use at home that was about 30 percent for workers with 10 years of experience in 1980-89 jumped to 70 percent during the next 10 years and was also at about 70 percent in those years for the next younger cohort in 1990-99. Figure B1 in the appendix shows that computer use at work was strongly related to schooling: workers with higher levels of education started using computers earlier, and computer use converged to higher levels compared to workers with lower levels of education.

2.2. Technological change and the labor market

Variations in the rewards to work during the past century in the United States have been explained as the outcome of the “The Race between Education and Technology” by Goldin and Katz (2008). Skill-biased technological change has continually raised the relative demands for more skilled workers which, in turn, has increased the derived demand for human capital investments via increased levels of formal education and on-the-job training. Economists have long argued that a major reason that technological change is skill biased lies in complementarity between new forms of physical capital (such as computers) and the skills needed for workers to learn to utilize capital effectively. (See Berman, Bound, and Griliches, 1994, for early contributions.)

More recent literature, surveyed by Acemoglu and Autor (2011), suggests the need for a more nuanced view of the relationship between the cognitive abilities and skills of workers and the tasks they perform in order to understand how new technologies affect their productivity and

employment. This literature documents a growing polarization of the occupational skill distribution created by a relative decline in the number of people in middle-skilled—and middle class—occupations relative both to people in highest skilled occupations and those in the lowest skilled occupations (Autor et al., 2006). Autor, et al. (2003) famously argued that computerization lead to a “hollowing out” of middle skill occupations because the primary tasks carried out many such jobs were “routine” in the sense that computer code could be written to perform these tasks more cheaply and rapidly than humans. In contrast, advances in computerization tended to complement the productivity of workers in many high skill “knowledge jobs” while it had little relevance to most tasks required of workers in low skill manual and service jobs. As a result, the distributions of employment and of wages became increasingly polarized.⁴

Computerization has created ongoing dramatic changes in the organization of work; the division of labor within and between firms; the creation of new products and new ways of producing and delivering old goods and services. As Leiner, et al. (1997) presciently stated: “*The Internet is at once a world-wide broadcasting capability, a mechanism for information dissemination, and a medium for collaboration and interaction between individuals and their computers without regard for geographic location.*” These features of computerization provided enormous advantages to knowledge workers that largely account for the rapid growth of wages and employment of highly educated workers in occupations such as finance, management, at STEM jobs in which

⁴ If demand is sufficiently elastic and computerization significantly lowers the cost of production, employment will tend to grow and often new products or new methods of production of old products will emerge. The elimination of routine tasks creates the opportunity for different, more complex human skills to be used in production. For example, employing artists who drew thousands of slightly different pictures, it took three years to produce *Snow White* in 1937. Computerization allowed artists using programs to generate apparent motion in highly complex ways nearly instantaneously. The new technology changed the ways to produce animated movies, led to new products such as computer games and increased the premium on artistic imagination relative to draftsmen’s skills and opened up opportunities for high monetary rewards for successful artists.

collaboration and communication are important components of productivity. The decoupling of geographic location and highspeed interactive communication among workers within a firm; between employees of different firms; and between sellers and buyers has profoundly transformed the organization of production and consumption at scales ranging from the relationship between a boss and a secretary or salesclerk and customer to the outsourcing of specialized functions previously done in house and offshoring of manufacturing. These changes have drastically reduced the demand in many specific occupations such as travel agents, department store clerks, or workers in a wide variety of manufacturing jobs that are largely unrelated to the computer skills of the displaced workers. Over time, some of these workers used their existing computer skills or acquired new skills to take jobs created by computerization while others took jobs that did not require the use of computers.

Complementarity between knowledge work and computerization had important downstream effects on the workers who provide services for knowledge workers. For example, as we illustrate later, the use of computers by managers and secretaries followed identical growth curves that accelerated with the introduction of PCs, word processors and spreadsheets in the early 1980s. Bosses learned to type their own letters and reports while secretaries learned to finish up the final electronic products, send them to their destination and organize electronic archives in place of filing cabinets. As office hardware and software continued to become easier to use purely secretarial tasks became more efficient so that one secretary could service more managers. However, integrated office software also continually increased the range of things that bosses might like to do and, therefore, the range of tasks former secretaries who became office assistants and managers whose responsibilities and required software skills that ranged widely depending on whether the boss was a doctor, lawyer, professor, CEO, salesman and on

and on. Similar changes have occurred in a myriad of other occupations ranging from architects to zoologists and they will surely continue in both predictable and disruptive ways into the indefinite future.

2.3. Computerization and obsolescence

Computerization must have come as a surprise to cohorts of workers who made educational and occupational choices without, most likely, much foreknowledge of the vast implications for the skills they would need to have in the workplace of the future. As Autor et al. (2003) argue, those with skills to perform what turned out to be “routine tasks” that could be performed by a computer program (e.g., bookkeepers) suffered a loss of productivity in their chosen occupation and there was often a decrease in demand for workers in such occupations which led to hollowing out of both the wage and occupational distributions. From a cohort perspective, this was a classic case of “creative destruction” in which the bundle of skills possessed by the “pre-computer” cohorts whose jobs involve routine tasks suffered a capital loss while workers who possessed skills in “cognitive jobs” whose productivity was enhanced by computers saw their earnings growth accelerate.

It is clear that advances in computer technology—hardware, software, generality and specificity of application, etc.—provide opportunities for workers in many different occupations to improve their productivity, and for firms to either increase or decrease the number of workers they require depending on the elasticities of substitution in production and final outputs⁵. It is

⁵ For example, Hamermesh (2013) finds that papers in top economics journals by older economists nearly quadrupled over two decades while the fraction of papers in pure theory fell significantly. Complementarity between advances in statistical and econometric theory and computer programs; data collection and dissemination; training in using these methods and sharing analysis and writing within teams doubtlessly played an important role in lengthening the academic careers of economists and changing the comparative advantage of the best and the brightest between pure theory and empirical analysis.

also clear that obsolescence is an especially serious threat to experienced workers in computer-intensive jobs. These workers are in competition with younger workers who have learned cutting edge skills in their formal education or during their first years in the labor market when investments in on-the-job training and learning-by-doing are at their highest levels. This was especially true for workers who finished their education and began their careers before computerization began. These workers either needed to learn the new technology from scratch or find a different job that could make use of the skills they already possessed. As the worker ages, costs to the worker and firm of investments to offset obsolescence may well increase because of the increasing time separating the worker's computer-related skill and the skills possessed by younger workers. In addition, the benefits of continued investment in a worker's human capital diminish as retirement approaches and the time to capture these benefits diminishes. Both factors suggest that, other things equal, workers in relatively computer-intensive occupations will tend to retire earlier. A possible exception to this prediction occurs if the advance of computer technology is sufficiently complementary to the worker's skills to raise her productivity. An example is an older economics professor collaborating with a young economist with cutting edge skills. In this case, the professor may capture more of the gains from investments in keeping her knowledge of her field current by delaying retirement.

Recently, Ahituv and Zeira (2010) and Deming and Noray (2018) have focused attention on the effect of technological change on the careers of STEM workers. They extend Rosen's (1975) pioneering model of equalizing differences due to obsolescence in several ways and use new data to show how the tasks conducted by STEM graduates change over the working life cycle as a consequence of technological change. Like Rosen (1975) and most other subsequent related literature, they build a formal theoretical model in which there is a constant rate of

technological change and derive steady state patterns of life cycle changes in a competitive labor market in which workers sort according to ability and other characteristics, and wages and occupational choices provide equal present values of life time earnings net of the direct and opportunity costs of investment. While Ahituv and Zeira, (2010) and Deming and Noray (2018) provide rich theoretical and empirical frameworks for studying many of the same issues that concern us, they assume that workers have perfect foresight about the benefits and costs of investing in skills and knowledge to compensate for obsolescence due to technological change. In contrast, we focus on implications for the wages, labor supply and retirement in the Retiree Cohorts who suffered a mismatch between the skills and knowledge embodied their human capital relative to that embodied in the human capital of younger cohorts that was caused by their lack of foreknowledge of computerization when they were making educational and occupational choices earlier their careers.

3. Data

3.1. The basic monthly Current Population Survey

The CPS is a large monthly survey that represents the non-institutionalized adult US population. It collects information about demographics, labor market activities, and education, among other things. The main advantage of the CPS is its large size, roughly 60,000 households each month.

The CPS uses a rotating panel survey design,⁶ which allows researchers to estimate monthly or yearly transition probabilities between labor market states, such as employment → out of the labor force transition probabilities. These so-called semi-panels (or short panels) have

⁶ See https://cps.ipums.org/cps/sample_designs.shtml for details about the sampling design.

been used by many researchers in the past to analyze patterns in worker turnover (see e.g. Blanchard and Diamond, 1990; Drew et al., 2014; Shimer, 2012). Following the methodology discussed in Shimer (2012) we built such a panel dataset by linking all CPS monthly surveys from January 1984 to January 2018.

Each person is in the sample a maximum of eight times. If the first interview was in month 1, then the person is interviewed in months 1, 2, 3, 4, 13, 14, 15, 16. These eight months are sometimes called rotation groups indexed from 1 to 8. We estimate yearly transition probabilities by tracking changes in labor market status between months 1→13, 2→14, 3→15, and 4→16. More than two thirds of the sample can be successfully linked by this methodology. For example, of the 10,617,199 individuals, who were age 25-68 between 1985 and 2017, worked for pay, and were in rotation groups 1-4, we could link and detect a valid labor force status for 7,325,859 (69%) individuals. The rest of the sample either moved (the CPS does not follow those who move) or were lost to follow-up for other reasons. Because job-losers are more likely to move than non-movers, the estimated job-loss probabilities may be downward biased in the linked data. We expect this bias to be relatively small in the 50-69-year old sample, however. Additionally, our econometric model identifies the effect of computerization from longitudinal variation, and we are not aware of any reason why changes in the moving-related bias, if any, would correlate with changes in computerization.

The CPS asks about individuals' detailed occupations using the 3-digit census occupation classifications. Most recently it covered more than 500 different occupations. There was one major change in the classifications in 2003, and there were two minor changes in 1991 and 2011. Occupational crosswalks between these alternative definitions are available, such as the one by

IPUMS⁷ and the one developed and used by us in prior research⁸ (Carr et al., 2020; Hudomiet, 2015; Sonnega et al., 2018). The IPUMS crosswalk is more detailed, but it is not balanced as many occupations are only available in certain years. The categories developed in our earlier research has fewer categories (191), but it is balanced and more consistent over time. We use these 191 categories in this project.

3.2. CPS Annual Social and Economic Supplement

The basic monthly CPS does not collect information about wages and income. The CPS-ASEC, collected every March asks individuals about wages and income received in the previous calendar year. We use the version of CPS-ASEC downloaded from the IPUMS website.

We restrict the sample to those who worked in the prior calendar year, and we use the occupations based on their longest jobs in that year. For the wage models we use individuals' annual wage and salary income, and we sometimes restrict the sample to full-year workers, defined as working 50-52 weeks a year.

3.3. CPS Computerization Supplement

The CPS also includes other supplementary modules that are less frequently asked. One includes questions about computer and Internet use either at home, at work, or at other places. These questions have been asked for a total of 14 times between 1984 and 2017. In this project we use four variables that were more-or-less consistently measured over time:

1. The person uses a computer at work
2. The person uses the Internet at work

⁷ See https://usa.ipums.org/usa/volii/occ_ind.shtml for details.

⁸ Available at <https://sites.google.com/site/phudomiet/research>.

3. The person uses a computer at home
4. The person uses the Internet at home

The primary focus of this paper is whether workers use computers or the Internet at work.

Computer use at work is available in six supplements: 1984, 1989, 1993, 1997, 2001, and 2003. Computer use at home is available in the same waves, as well as in 2015, and 2017.

Internet use at work is also available in more recent years: first in 1997, then in 1998, 2000, 2001, 2003, 2011, 2013, 2015, and 2017. Finally, Internet use at home is available in the same waves as Internet use at work, and, in addition, in 2010 and in 2012.

There were slight changes in the wording of the questions. After investigating we decided to drop the 1998 and 2000 waves from the analysis, because the CPS used a different skip patterns in those years that disrupted the time-series.⁹ The rest of the measures and waves were kept in our analysis sample.

The CPS computerization supplements asked additional information about computer and Internet use at home, at work and at other places, but those measures have poor comparability over time, and they were not used in this project.

3.4. The Health and Retirement Study

The HRS is a popular dataset for studying the retirement process as it has a large sample of the relevant 50 and older age range, and it has very detailed panel information on individuals. It collects detailed information about individuals' demographic characteristics, SES, wages, income, labor force outcomes, health, and other variables. Moreover, the HRS also elicits the

⁹ In 1998 and 2000 the CPS first asked if individuals used the Internet outside their homes. Those who answered yes were asked whether they used it at work. In other years, the CPS asked about Internet use at work directly.

retirement expectations of workers by asking “What do you think the chances are that you will be working full-time after you reach age 62?” A similar question asks about age 65. As Hurd (2009) argues, these probabilistic expectations data tend to predict actual future realizations well in most cases. Importantly, expectations data are available in even younger birth cohorts, such as the Early Baby Boomers (born between 1948 and 1953) and Mid Baby Boomers (born between 1954 and 1959).

The HRS provides information on workers’ occupations at the three-digit-census level. The classification changed in 2006 from the 1980 to the 2000 census classifications, but the same cross-walk described in 3.1. is used to make the data consistent. The publicly available version of HRS only has a rudimentary one-digit occupation categorization. The detailed occupation of workers is considered restricted data, but researchers can obtain access through a procedure described on the MiCDA website.¹⁰

The rest of the measures we used will be discussed in the results section before we present our findings.

3.5. O*NET occupational characteristics

The O*NET (the Occupational Information Network¹¹) was developed under the sponsorship of the U.S. Department of Labor to replace the DOT which was the public standard description of occupations prior to 2000. The O*NET collects information about hundreds of occupational characteristics in detailed three-digit occupations that can be linked to the CPS and the HRS. The occupational characteristics include measures of skills and abilities required to perform well on the job, as well as detailed description about work activities. The main

¹⁰ <http://micda.psc.isr.umich.edu/enclave>.

¹¹ <https://www.onetcenter.org/about.html>.

advantage of the O*NET is its large size and great precision. However, the data cannot be used to identify changes in occupational characteristics, as most of the data was collected in the mid-2000s.

The O*NET has a direct measure of computerization, which asks if workers need to “Interact with computers” at work. We use this measure in our computerization model as a high-quality predictor variable of computer use at work in the CPS. Our coding procedure follows other papers in the literature such as Autor et al. (2003), Firpo et al. (2011) and Carr et al. (2020). For each occupation, the O*NET provides information on the “importance” and “level” of required work activity. We assign a Cobb-Douglas weight of two thirds to “importance” and one third to “level” to create single measures of computer usage in occupations. These measures are then averaged for the somewhat less detailed census occupation level. Finally, in order to collapse the information to our 191-category occupations, a weighted average of the measure is taken, where the weights are the relative frequencies of the occupations in the 2000 census within the 191-category occupations.

3.6. Descriptive Statistics

Table B1 in the appendix shows descriptive statistics of the four datasets used in this project. The first three columns represent 25-69-year old workers in the basic monthly CPS, the computerization supplements, and in the CPS-ASEC respectively. The last column represents 55-59-year old workers in the HRS. Each sample is restricted to workers with a valid occupation, and the basic monthly CPS is further restricted to individuals who were successfully matched to themselves in the 12-month follow-up survey.

The basic monthly CPS has more than seven million person-year observations. The average age of the workers is 44.8, a little less than half of the sample is female, less than 10% are high school dropouts, while the rest of the sample is evenly split between high school graduates, individuals with some college education but without a degree, and college graduates. About 19% of the sample belongs to a racial or ethnic minority group, 70% are married, the same fraction work in a private firm, and there is a reasonable distribution of workers in different aggregate occupations.

The computerization supplements and the CPS-ASEC are also large datasets, but they are noticeably smaller compared to the basic monthly CPS with about 625 thousand and 2.6 million observations respectively. The distributions of the basic variables are very similar to the basic monthly CPS.

The HRS sample is considerably smaller with about 12 thousand individuals. This sample is older by design. The HRS oversamples racial minorities which explains the relatively smaller fraction of non-Hispanic whites, and a lower mean education level compared to the CPS. The distributions of the other variables, including the aggregate occupations, however, are very similar to the CPS samples.

4. Results

4.1. Trends in computerization by age and skill

We start by documenting the penetration of computers into different skill and age groups from the mid-1980s until today using the CPS computerization supplements. We are primarily interested in documenting differences in computer use between older and younger workers within the same skill groups (education, occupations), which we interpret as a knowledge gap

between the age groups. We are also interested in the time-series patterns in knowledge gap, and how they vary by skill.

The eight panels of Figure 2 show trends in the fraction of computer and Internet users at work by age and education. Very few high school dropouts use computers or the Internet at work, even though their fraction steadily increased over time. Moreover, we do not see strong age differences among high school dropouts, the fraction of computer and Internet users remained low in all age groups. At the other end of the skill spectrum, the fraction of college graduates who used computer or the Internet at work was non-negligible even in the '80s (20-50% depending on age), and it continued to increase over time. Most recently about 80% of college graduates use the Internet at work. Similar to earlier literature, we also see significant differences in computer and Internet use by age (Borghans and ter Weel, 2002, Friedberg, 2003). The gap appeared to be the largest in the '80s and early '90s, then it gradually shrank and completely disappeared by the mid-2010s. The patterns are similar, but less strong among workers with some college education. The fraction of computer and Internet users is somewhat below that of college graduates, but the gap between younger and older workers also peaked at earlier times among workers with some college education. The fraction of high school graduates who use computers or the Internet at work is in between high school dropouts and workers with some college education. The knowledge gap between younger and older workers was not as great in the '80s since very few high school graduates used computers at that time, but the gap increased, peaked in the '90s, and then started shrinking again.

The eight panels of Figure B2 in the appendix show the fraction of computer and Internet users at home. The patterns among the two higher-educated groups are similar to those in Figure 2: there was a large gap in computer and Internet use by age in the earlier periods, but the gap

gradually shrank and completely disappeared by the mid-2010s. The patterns, however, are different among high school dropouts and high school graduates. Even though many of these workers do not use computers and the Internet at work, the fraction who used the Internet at home rapidly increased in the 2000s, especially among younger workers. More recently there is still a noticeable gap in computer use at home between younger and older workers.

Figures B3-B4 in the appendix show similar results in four selected occupations (instead of education). Overall, the patterns are similar. In the highest skill professional occupations, computer use was moderately high even in the '80s, and the age-gap was high. As computer and Internet use continued to increase, the age-gap continued to shrink, and it completely disappeared in the mid-2010s. In the least skilled food and cleaning service occupations, very few people use computers or the Internet at work, and the gap seems to be small. At the same time, computer use at home still shows a quite noticeable gap between younger and older food and cleaning service workers. Office workers are similar to professionals, and production workers are similar to high school graduates.

To summarize:

1. More skilled workers were more likely to use computers and the Internet both at home and at work as compared to less skilled workers throughout the entire time period.
2. Computer and Internet use at work is still negligible among the least skilled workers, though many of them use computers/Internet at home at this juncture.
3. The age-gap in computer use at work was the largest in higher skilled occupations at earlier times: in the '80s and early '90s.

4. The age-gap among less skilled workers peaked a bit later, but it never was as large as among higher skilled workers.
5. The knowledge gap between age groups was temporary, most outcome variables (computer and internet use at home and at work) converged to the same levels across age groups within the same occupation.

Overall, we found interesting heterogeneity in computer use and knowledge gap by time and skill.

4.2. Estimating knowledge gap in computerization

4.2.1. Model setup

In this section we provide an overview of the model we used to estimate a knowledge gap measure that enters our retirement models in the next sections. Appendix A discusses further technical details.

In the data we observe computer and internet use at work and at home by detailed skill groups. We fit a set of probit models on these data and recover a measure of “knowledge gap”, which, under some assumptions, can be interpreted as the estimated probability that an older worker (age 50-69) does not use computers at his/her job even though the computers would be useful. A knowledge gap, then would potentially indicate a mismatch of the older worker with his/her job. The usefulness of computers is identified from computer use among middle-aged individuals (age 40-49) in the same skill group.

Let’s introduce the following notation:

- Event G^C : Computers are useful at the person’s job
- Event K^C : The worker knows how to use computers

- Event W^C : The worker uses computers at work
- A: Age
- X: Skill of worker

The knowledge gap is the estimated joint probability that a worker does not know how to use computers, even though they are useful at his/her job. Formally it is the probability:

$$\Pr(G^C, \bar{K}^C) = \Pr(G^C)(1 - \Pr(K^C | G^C)) \quad (1)$$

Our basic model only uses one of the four outcome variables: whether the person uses computers at work, W^C . The knowledge gap is the product of two terms, the probability that computers are useful at the person's job, $\Pr(G^C)$, and the conditional probability of knowing how to use computers if they are useful, $\Pr(K^C | G^C)$. The single outcome variable does not identify these two equations and the knowledge gap. We use the following three identifying assumptions for estimation:

1. $\Pr(W^C) = \Pr(G^C, K^C)$: Workers use computers at work if and only if they know how to use them and they are useful at their jobs.
2. $\Pr(K^C | G^C, A < 50) = 1$: Workers younger than age 50 always use computers when it is useful in their jobs. That is, workers below 50 always learn how to use computers if it is useful at their jobs.
3. $\Pr(G^C | X, 40 \leq A < 50) = \Pr(G^C | X, A \geq 50)$: Conditional on skill, the usefulness of computers is the same among middle age (age 40-49) and older (age 50-69) workers.

Assumption 3 establishes that middle-aged workers form a valid control group for older workers, and it is motivated by the observation in the previous section that computer and internet use at work converged to the same levels across age groups in most skill groups. Assumption 2 normalizes the knowledge gap at zero in younger groups, which is likely a strong assumption in reality. If assumption 2 does not hold true, our derived knowledge gap measure can only be interpreted as the difference in computer use between middle-aged and older workers within the same skill group. This paper does not intend to recover any deep structural parameters, and for our main purposes, this alternative interpretation is also fine.

After functional form assumptions, these three identifying assumptions allow us to estimate the knowledge gap. The usefulness of computers at work, $\Pr(G^C)$, is assumed to be a flexible probit model including detailed occupations, education, race, age groups, a cubic function of time, a cubic function of time interacted with the O*NET computerization measure¹², and a linear time trend interacted with education, race, age groups, and a cubic function of the O*NET computerization measures. The probability of knowing how to use computers if they are useful, $\Pr(K^C | G^C)$, is also a flexible probit with education, race, age groups, O*NET computerization, a linear time trend interacted by all the measures from above, and a quadratic term in the O*NET computerization measure.

This model can be estimated from observations on W^C , and then we can use (1) to estimate the gap. To increase precision, we developed an extended version of the model by

¹² The detailed occupation dummies, thus, appear as a baseline predictor, but they are not interacted with the time trends. When we tried to interact the occupations dummies with time, the probit model usually did not converge. In our preferred specification we used the O*NET computerization measure interacted with time in order to reduce the complexity of the estimation model.

incorporating Internet use at work, as well as computer and Internet use at home. The extended model further assumes that Internet users all use computers as well. Computer and Internet use are specified using probit models with flexible functions of skill, age, and interview years.

Appendix A discusses details.

4.2.2. Results

Figures 6-13 in Appendix B shows model fit: predicted and observed patterns in computer and Internet use at work and at home by age and years, in selected education, and occupation groups. The predicted probabilities are smoother and less noisy than the observed probabilities, but the fit of the model is very good for all four outcome variables.

Figure 3 shows the estimated knowledge gap by age and years by education. Figure B5 shows the gap by major occupation groups. The gap clearly increases with age, because the oldest workers are the least likely to use computers when their younger peers do. We also found that the knowledge gap was the largest among the higher skilled workers in the late '80s and early '90s, while the gap was smaller and peaked a bit later among lower skilled workers. These results aligned with our findings in Section 4.1 when we looked at the raw data. We also looked at the knowledge gap in even finer occupation categories, and we found additional nuances. First, some highly skilled occupations, such as computer programmers, never experienced a knowledge gap, since they all used computers, even in the 1980s. Second, older workers in other moderately skilled occupations, such as managers of properties and real estate, various sales jobs, and supervisors of mechanics or foremen, experienced sizable knowledge gap even in the late 2000s and early 2010s.

4.3. The effect of knowledge gap on yearly transition probabilities

4.3.1. Methods

We use two econometric models to estimate the effect of the knowledge gap on yearly labor market transition probabilities observed in the monthly CPS. The first one uses age, skill and time fixed effects:

$$y_{iast} = \beta_0 + \beta_1 Gap_{ast} + \alpha_a + \gamma_s + \delta_t + \varepsilon_{ist}, \quad (2)$$

where i indexes individuals, s indexes skill groups, and t indexes time. y_{iast} is an outcome variable, such as an employment \rightarrow out of the labor force transition probability, the Gap_{ast} is the knowledge gap measure, which varies by skill, age, and time, and the model has single-year age, skill, and time fixed effects. The skill dummies are full interactions between detailed occupations, education, race and ethnicity, and 5-year age groups.

The model is estimated using the sample of 50-69-year-old workers. Because we control for skill and time fixed effects, the effect of computerization is identified from within skill group patterns in the knowledge gap. That is, we test if we see above average transitions into retirement, for example, in years when the knowledge gap was higher than average in a particular occupation. For simplicity, we will sometimes refer to this model as a DD (difference-in-differences) model, because of the similarity in the control structure to a DD model. A standard DD model would use a binary knowledge gap variable. Our empirical model, however, uses a continuous rather than binary treatment variable (the gap). The main, non-testable, identification assumption is that the gap does not correlate with the error term in (2). This would be violated if there were differential trends in the outcomes across skill groups, and these trends correlated with the trends in the knowledge gap. This would be the case, for example, if some outside factor induced workers in some occupations to start using computers and to change their

retirement behaviors at the same time. We are not aware of any economic arguments that would justify this concern, but nevertheless it is a possibility.

To check the robustness of our findings, we replicate the DD model within various population subgroups, such as by education or occupations. The identification assumptions are more likely to hold within these subgroups, because the “treated” and “control” workers are more similar within such groups.

Our second model relaxes some of the identification assumptions in the DD model:

$$y_{iast} = \beta_0 + \beta_1 Gap_{ast} + \alpha_a + \gamma_s + \delta_t + \gamma_s \times \delta_t + \varepsilon_{ist}, \quad (3)$$

The extra term is the full interaction between the skill and the time dummies ($\gamma_s \times \delta_t$).

These extra interaction terms leave no variation in the knowledge gap measure in the 50-69-year-old sample, and therefore it cannot be estimated in that sample. Instead, we estimate the model on the 40-69 or the 25-69-year-old samples, and these younger groups serve as the controls¹³.

The model identifies the effect of the gap by tracking the difference between younger and older workers’ outcomes within the same skill groups and same year. That is, the model tests if older workers are more likely to transition out of work compared to younger workers when the knowledge gap is high, while we control for a full interaction between skill groups and time. For simplicity we will sometimes refer to this model as a DDD (triple difference-in-differences) model, because of the similarity of the control structure to a triple diff-in-diffs model in which the 40-49-year old groups (or the 25-49-year-olds) form the third “diff”. Again, the main

¹³ The skill dummies include 5-year age dummies from age 50-54, to 65-69. Workers younger than age 50 were randomly assigned to one of these four age groups when we created the skill dummies, γ_s . The age fixed effects, α_a , however, are based on workers’ own age.

difference between a regular DDD and our model is that our treatment variable is continuous rather than binary.

Table B2 in the appendix illustrates the DD and DDD models by estimating models of computer use at work as a function of the knowledge gap and the various control variables. The coefficient on knowledge gap is expected to be negative when we control for time and skill fixed effects. We use two data sources: the CPS computerization supplement, and the HRS. The latter asks workers how often they are required to use computers at work, and we coded their answers to 1 (vs. 0) if they reported “most of the time” or “all the time” (vs. “some of the time” or “never”). As expected, all models predict a strong negative effect of the gap on computer use, though the magnitude of the coefficients varies by the sample and estimation method. We want to note that in column [3], the estimated coefficient is precisely -1, which was expected, because the knowledge gap was basically an estimate of the difference between 40-49-year-old and older workers’ computer use, and the DDD model uses the same variation for identification. When we control for age fixed effects, the coefficient on the knowledge gap shrinks, but remains strongly negative and statistically significant. We include age fixed effects in all models below.

4.3.2. Main results

Table 2 shows DD and DDD estimates of the yearly employment → out of the labor force, and employment → unemployment transition probabilities. We found strong, statistically significant, and robust negative effect on leaving the labor force. Older workers were more likely to leave the labor force when their jobs were computerized, and they did not know how to use computers. For a knowledge gap of 0.1 (or 10%, which was not uncommon in higher skilled occupations in the late ‘80s and early ‘90s), the coefficient would imply an increase in the

transition probability of 1-1.4 percentage point, while the mean of the outcome variable is 8 percentage points. This is a quite large effect.

We found a small and non-robust negative effect on the employment → unemployment transition probability. The negative effect would mean that workers who experience a knowledge gap were less likely to search for a new job, and they, instead, left the labor force altogether. Our preferred estimate, however, is zero, because the more reliable DDD models show quite precise zero effects.

In Table 3 we restricted the sample to those who worked in both years and estimated transition models between job types. We found robust evidence that workers with a knowledge gap were more likely to switch from a full-time job to a part-time job, which may mean that their skills were less valuable for their employers and they were forced to downgrade. We also looked at transitions into less skilled occupations, or “occupational downgrading”. It is possible that workers who experienced a gap had to switch jobs and move into less skilled occupations. First, we estimated occupational wages, by taking the occupational means of log annual wage and salary income of full-year workers in the CPS-ASEC. A transition into a lower paid occupation was defined as a transition into an occupation with at least 10% lower wages on average. We also looked at transitions into less computerized occupations, which were defined as moving into an occupation which had at least 0.1 lower value of the O*NET computerization measure compared to the original job. The 0.1 change is about a half standard deviation in the measure. These transition probabilities have surprisingly high means, above 10 ppts, which is likely due to noise: before 1994, the CPS coded occupations independent of their previous values, which led to many spurious occupation changes in the linked data. Overall, we found mixed results. The DD models predicted no changes in occupational skills, while the DDD models predicted that the

gap decreased the propensity to downgrade. The negative DDD estimates are somewhat counter-intuitive, but it may be explained if older workers who experienced a knowledge gap were less likely to search for any jobs, including less skilled jobs.

4.3.3. Sensitivity analysis

We estimated a number of alternative specifications of the main empirical models to test the robustness of our findings to certain assumptions. The preferred specification defined skill groups as the full interaction between 191 occupations, four education, four race, two gender, and four age groups, totaling $191 \times 4 \times 4 \times 2 \times 4 = 24,448$ groups. The knowledge gap measure was estimated for each of these groups; we entered this many fixed effects terms into our DD models; and even more to the DDD models (because of the year interactions). Table B3 in the appendix shows estimates of the employment \rightarrow non-participation transition probabilities using far simpler skill group definitions. In the simplest case we only used four education and four age groups, totaling only 16 groups. We also considered two other models with 10 occupations and four age groups (=40 groups). In each case we re-estimated the “knowledge gap” using these more aggregate skill groups, and then we re-estimated the DD and DDD models. The table shows that the main results are qualitatively the same, and quantitatively very similar. For example, when we only use four education groups the DD model yields a coefficient of 0.086 vs. 0.103 of the preferred specification. Both coefficients are significant at any conventional levels. The DDD models are also similar: 0.146 vs. 0.122. Overall, the main results are robust to the definition of the skill groups.

We also tested if the results are robust to controlling for the sizes of the occupational labor force. Autor and others found that middle-skilled occupations have shrunk in recent decades partly due to computerization. It is possible that our results are related to this change,

and perhaps they are driven by it. Table B4 in the appendix shows models of the yearly employment → non-participation and employment → unemployment transition probabilities when we control for the size of the occupational work force. More precisely we estimated the ratio of the number of workers in a given occupation in a given year (using each monthly CPS surveys in that year) to the number of workers in that occupations in 2002 (before 2002) and 2003 (after 2003). This ratio was, by design, equal to one in 2002 and 2003, and the variable showed trends in the number of workers in the 191 occupations.¹⁴ Table B4 shows that the DD models are virtually identical to the inclusion of these occupational share variables. The DDD estimates are not shown, but they were found to be perfectly identical to our preferred specifications, because the occupational shares do not vary by age (i.e. they are the same for the 40-49 and the 50-69-year old groups), and the DDD models identify from those differences.

4.3.4. Subgroup analysis

Table 4 shows subgroup analysis of the main results in Tables 2-3. Each model is re-estimated on the given subsample; that is, they are not just interactions between the gap measure and the group identifiers. The subgroup estimates scatter around the estimate for the total sample, which strengthens the validity of the identification assumption in the total sample (because the subgroup models' validity is more likely guaranteed). We found larger effects on females, workers with some college education, workers aged 60-64, and workers in office jobs. Females may have been more affected by computerization because they were less likely to learn to use computers than men. Office workers may have been more strongly affected than workers in other occupations, because computerization wiped out some of those middle-skilled jobs (Autor et al,

¹⁴ We verified that, similarly to the literature, the constructed measure showed a decline in the share of workers in middle skilled occupations, such as office and administration, and it showed increasing shares in high skilled production and managerial, as well as low skilled personal care jobs.

2003). Moreover, workers in office jobs may have had less freedom in choosing whether they want to work with computers, compared to workers in more skilled professional jobs, for example. We also found that the effects were strongest among 60-64-year old workers. The oldest 65-69-year-old workers may have experienced a larger knowledge gap than 60-64-year old workers, but the fact that they continued working to an older age may reveal that they are not average workers.

4.4. The effect of knowledge gap on wages

We use the same DD and DDD type models to estimate the effect of the knowledge gap on log annual wage and salary income of older workers in the CPS-ASEC. Table 5 shows the results. The first three columns do not restrict the sample by the number of weeks workers worked, and therefore they mix wage and labor supply effects. The second three columns restrict the sample to full-year workers.

The first three columns (no weeks restriction) show very large and quite robust effects on wages. A 0.1 increase in knowledge gap, for example, is estimated to decrease annual wage and salary income by 6-7 percent. When we restrict the sample to full-year workers, the DD estimate does not change, but the more reliable DDD estimates shrink to about 0.25. This implies that a 0.1 increase in the knowledge gap decreases wages by about 2.5 percent. This is a relatively small effect, but nevertheless, is highly statistically significant.

In Table 6 we report subgroup estimates, in which we re-estimated the models within each of the population subgroups. Similar to the transition probabilities, we found larger effects on females, workers with some college education, and office workers. We found that 65-69-year-old workers suffered the largest negative wage effects. The model also predicted a counter-

intuitive positive wage effect on the self-employed as well as on operators and laborers, but the standard errors on these coefficients were quite large.

4.5. The effect of knowledge gap on workers' beliefs, mental health, and retirement

We finally present results based on the HRS that contain additional information about individuals' mental health, retirement expectations and job characteristics. All estimates are based on the DD methodology, since the HRS only includes workers above age 50.

Table 7 is based on 55-59-year-old workers between 1992 and 2014, and for some outcomes we follow these individuals to older ages in the panel. The first four columns of the table investigate outcome variables at the baseline age window of 55-59. P62 and P65 mean workers' own subjective reports of the probability that they would work after age 62 or 65 measured at baseline. As expected, the knowledge gap had a strong negative effect on these probabilities, implying that these workers expected to retire earlier. Then we looked at wages, and similar to the CPS-ASEC model, we found a strong negative effect on wages. Finally, in columns 3-4 we looked at job satisfaction and mental health, but we found no statistically significant effects.

Next, we make use of the panel structure of the HRS more directly. We took workers at age 55-59 from the 1992-2004 waves, and then followed them until they turned age 65. Columns 5-7 of Table 7 show how the gap at baseline (at age 55-59) predicted various outcomes at age 65. We found that workers were significantly less likely to work in a full-time job at age 65. The effect on any work, including part-time jobs, was also negative, but not statistically significant. We also found that workers with a knowledge gap worked significantly fewer years between the

baseline year and age 65. Overall, we robustly found that workers with a knowledge gap retired earlier.

5. Discussion and conclusion

This paper explored the implications of the cognitive demands posed by computerization for the length of working life of older workers. While technological change may reduce the physical demands of work, allowing workers to retire later; rapid advances in technology may increase the cognitive demands of work, making it harder for some older workers to continue working until the age they planned to. The development of computers and the ubiquitous spread of computerization throughout the economy largely occurred in the 1980s and 1990s, after members of older birth cohorts had completed their formal education and, in many cases, after they had chosen an occupation during their early labor market careers. Economic theory implies that rapid technological change will lead to obsolescence of a worker's skills unless the worker (and the employer) makes investments in human capital to keep up with and adapt to the new technology. For older workers who are close to the end of their working lives, however, it may not be optimal to make such investments, and they may be forced into early retirement.

We first used data from Current Population Survey supplements to track the fraction of workers who used computers and the Internet at work and at home between 1984 and 2017 in different skill and age groups. Skill groups were defined by the detailed 3-digit occupations of the workers as well as their gender, education, and race. We showed that workers in different skill groups started using computers at different times; older workers tended to start using computers with a delay compared to younger workers; but the difference largely disappeared by the early 2010s. We hypothesized that this knowledge gap between younger and older workers placed pressure on older workers to adopt the new technology.

In the second part of the paper, we tested how the estimated knowledge gap in computerization between younger and older workers affected the labor market outcomes of older workers. To do so, we used millions of observations from the basic monthly CPS, the CPS Annual Social and Economic Supplement, and the Health and Retirement Study. Importantly, our empirical models used time-series variation in computerization for identification. A growing recent literature has used occupation measures, such as based on the O*NET or the DOT, to study how workers' job characteristics affected their labor market outcomes. These measures only provide cross-sectional variations based on the workers' occupations. Our empirical specifications, however, directly controlled for skills (i.e. 3-digit occupations, education, etc.) and time fixed effects. These models hold workers' skills constant and compare birth cohorts that were affected by computerization at different times in their careers, while controlling for general economic conditions with year fixed effects. We also developed an even more general empirical framework, in which we directly controlled for skill-year interactions, and we used younger (40-49 or 25-49-year-old) individuals in the same skill group and year as the "control group". . Essentially these models tracked the differences in the outcome variables between younger and older workers (such as the wage gap between them) as a function of the knowledge gap and controlling for a full interaction between the skill and time fixed effects.

We found robust evidence that the knowledge gap shortened the working lives of older workers, it pushed many full-time workers into part-time jobs, and it lowered their wages. For example, a knowledge gap of 0.1 (or 10%, which was not uncommon in higher skilled occupations in the late '80s and early '90s), increased yearly employment → not-in-labor-force transition probabilities by about 1-1.4 percentage points among 50-69-year-old workers. This is

quite a large effect, because the sample mean of this transition probability is only 8 percentage points.

We ran detailed subgroup analysis in which the same empirical models were estimated on narrower population subgroups, such as on workers with the same level of education, or workers in the same 1-digit occupations. The primary reason to estimate these models were to probe heterogeneity in the estimated effects. The secondary reason was to test the robustness of our findings because the identification assumptions of the statistical models were more likely satisfied in narrower, more homogeneous population subgroups. We found that the subgroup effects scattered around the total effect, and that the effects were larger on females, on middle-skilled workers (such as workers with some college education and office workers), and workers in the 55-64 age range. Middle-skilled and office workers may have been more strongly affected by computerization than workers in other occupations, because computerization wiped out more jobs in these occupations as discussed by Autor et al (2003) and many follow-up works. It is also possible that these workers had less freedom to choose whether they want to work with computers, compared to workers in more skilled professional jobs, which led to a larger effect in these groups.

We estimated models on other outcomes, but we found small, and inconsistent effects on occupational downgrading, such as transitions into lower paid occupations, and on the mental health of workers measured in the HRS. We ran a number of alternative specifications of the empirical models, such as defining skill groups in narrower or broader ways (to derive the knowledge gap measure and to control for skill fixed effects), or controlling for different sets of control variables, and we found the main results of the paper to be robust to these assumptions.

Computer technology improved the productivity of many workers over the past decades. Our results, however, indicate that many older workers retired earlier than “normal” when computerization first penetrated their occupations. This finding is consistent with a model of obsolescence in which older workers are in competition with younger workers who have learned cutting edge skills in their formal education or during the earlier years in their careers.

Even though our empirical models were guided by economic theory, a limitation of our study is that we did not estimate an equilibrium model of skill-formation, occupation choice, and retirement. Estimating such a model is beyond the scope of this paper, but it would be an interesting extension of our work. For example, it would be interesting to study why observationally equivalent older workers decided to learn or not learn how to use computers.

Nevertheless, a major empirical finding of the paper is that computerization did, on average, have a significant effect on inducing earlier retirement. The negative effects of computerization due to knowledge gap for cohorts that have already retired are unlikely to have a major effect on future cohorts because the gaps have largely disappeared during the past decade. Obsolescence due to continuing rapid progress (AI, robots, etc.) will continue to challenge older workers. However continued reductions in the physical demands of work resulting from computerization may offset some of the pressure for early retirement on workers in the post-Retiree cohorts who will enter their 60s with computer knowledge and skills required by their occupations and less competition from relatively smaller younger cohorts.

We believe that by studying the effect of computerization, arguably the most important technological advancement of our times, we gain a better understanding of the broader question of how technological change, which is a continuous and ongoing process, helps or discourages

the elderly to keep up with technology by investing in their human capital, and ability to work longer. An open question that we did not address in this paper is the extent to which the effect of computerization on the increase of inequality in the wage distribution by education and occupation since the late 1970s will be reduced for future cohorts. These cohorts have been introduced to computer use as toddlers. They may grow up to regard computers as just another tool to use in combination with their uniquely human capabilities as they pursue careers in the economy of the future. Or, alternatively, will the prediction of the redundancy of labor finally come to pass?

References

- Acemoglu, D. and Autor, D. H. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics Volume 4*, Orley Ashenfelter and David E. Card (eds.), Amsterdam: Elsevier: 1043–1171.
- Aubert, P., Caroli, E, Roger, M. (2006). New Technologies, Organisation and Age: Firm-Level Evidence. *The Economic Journal*, 116(509): F73–F93.
- Autor, D. H., Levy, F., Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics* 116(4): 1279–1333.
- Autor, D. H., Katz, L. F., Kearney, M. S. (2006). The Polarization of the US Labor Market. *American Economic Review Papers and Proceedings*, 96(2): 189–194.
- Autor, D. H. (2014). Skills, Education and the Rise of Earnings Inequality among the “Other 99 Percent.” *Science*, 344(6186): 843–851.
- Ahituv, A. and Zeira, J. (2010). Technical Progress and Early Retirement. *The Economic Journal*, 121 (March), 171–193.
- Belbase, A. and Chen, A. (2019). How Did Computerization Since the 1980s Affect Older Workers? *Issues in Brief*, ib2019-19, Center for Retirement Research.
- Berman, E., J. Bound, and Z. Griliches (1994). Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers. *Quarterly Journal of Economics*, 109(2): 367–397.
- Biagi, F., Cavapozzi, D., and Miniaci, R. (2013). Employment Transitions and Computer Use of Older Workers. *Applied Economics*, 45(6): 687–696.
- Blanchard, O. J. and P. Diamond (1990). The Cyclical Behavior of the Gross Flows of U.S. Workers. *Brookings Papers on Economic Activity*, 21(2): 85–156.

- Borghans, L. and ter Weel, B. (2002). Do Older Workers Have More Trouble Using a Computer than Younger Workers? In de Grip, A., van Loo, J. and Mayhew, K. (Ed.) *The Economics of Skills Obsolescence, Research in Labor Economics*, Emerald Group Publishing Limited, Bingley, (21): 139–173. [https://doi.org/10.1016/S0147-9121\(02\)21009-2](https://doi.org/10.1016/S0147-9121(02)21009-2)
- Burstein, A., Morales, E., and Vogel, J. (2019). Changes in Between-Group Inequality: Computers, Occupations, and International Trade. *American Economic Journal: Macroeconomics*, 11(2): 348–400.
- Carr, D. C., Castora-Binkley, M., Kail, B. K., Willis, R. J., and Carstensen L. (2020). Alternative Retirement Paths and Cognitive Performance: Exploring the Role of Preretirement Job Complexity. *The Gerontologist*, 60(3): 460–471.
- Deming, D. J. and Noray, K. L. (2018). STEM Careers and the Changing Skill Requirements of Work. *NBER Working Paper* No. 25065.
- Drew, J. A., Flood, S., & Warren, J. R. (2014). Making Full Use of the Longitudinal Design of the Current Population Survey: Methods for Linking Records Across 16 Months. *Journal of economic and social measurement*, 39(3): 121–144. doi:10.3233/JEM-140388.
- Firpo, S., Fortin, N. M., and Lemieux, T. (2011). Occupational Tasks and Changes in the Wage Structure. Working paper.
- Freeman, R. (2018). Has AI/Robotics Contributed to the Shift in Labor Market to Older Workers? Working paper presented at the 2018 Working Longer and Retirement Conference, SIEPR.

- Friedberg, L. (2003). The Impact of Technological Change on Older Workers: Evidence from Data on Computer Use. *Industrial and Labor Relations Review*, 56(3): 511–529.
doi:10.2307/3590922
- Gaggl, P. and Wright, G. C. (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. and Katz, L. F. (2008). *The Race Between Education and Technology*. Cambridge, MA: Belknap Press of Harvard University Press.
- Hamermesh, Daniel S. (2013). Six Decades of Top Economics Publishing: Who and How? *Journal of Economic Literature*, 51(1): 162–72.
- Hudomiet, P. (2015). The role of occupation specific adaptation costs in explaining the educational gap in unemployment. Working paper.
- Hurd, M. D. (2009). Subjective Probabilities in Household Surveys. *Annual Review of Economics*, 1. 543–562.
- Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Kleinrock, L., Lynch, D. C., et al. (1997) “A Brief History of the Internet.” Internet Society, downloaded from <https://www.internetsociety.org/internet/history-internet/brief-history-internet/>.
- Rosen, S. (1975). Measuring the Obsolescence of Knowledge, Education, Income, and Human Behavior. NBER, 199–232.
- Schleife, K. (2006). Computer Use and Employment Status of Older Workers — An Analysis Based on Individual Data. *Labour*, 20(2): 325–348.
- Shimer, R. (2012). Reassessing the Ins and Outs of Unemployment. *Review of Economic Dynamics*, 15(2): 127–148.

Sonnega, A., Helppie-McFall, B., Hudomiet, P., Willis, R. J., & Fisher, G. G. (2018). A

Comparison of Subjective and Objective Job Demands and Fit with Personal Resources as Predictors of Retirement Timing in a National U.S. Sample. *Work, aging and retirement*, 4(1), 37–51. doi:10.1093/workar/wax016.

Weinberg, B. (2002). New Technologies, Skills Obsolescence, and Skill Complementarity. In de

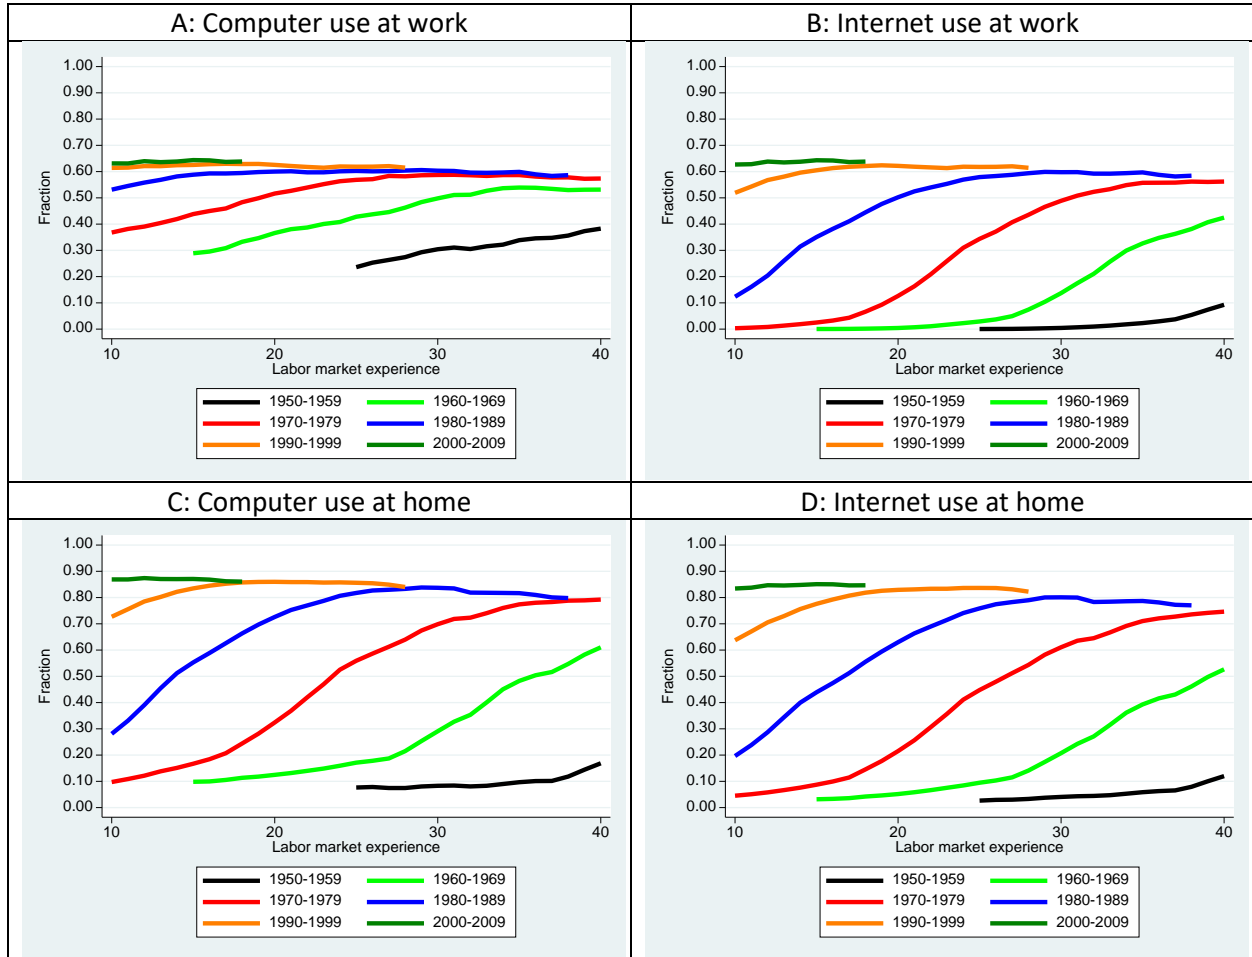
Grip, A., van Loo, J. and Mayhew, K. (Ed.) *The Economics of Skills Obsolescence*,

Research in Labor Economics, Emerald Group Publishing Limited, Bingley, (21): 101–

118. [https://doi.org/10.1016/S0147-9121\(02\)21007-9](https://doi.org/10.1016/S0147-9121(02)21007-9)

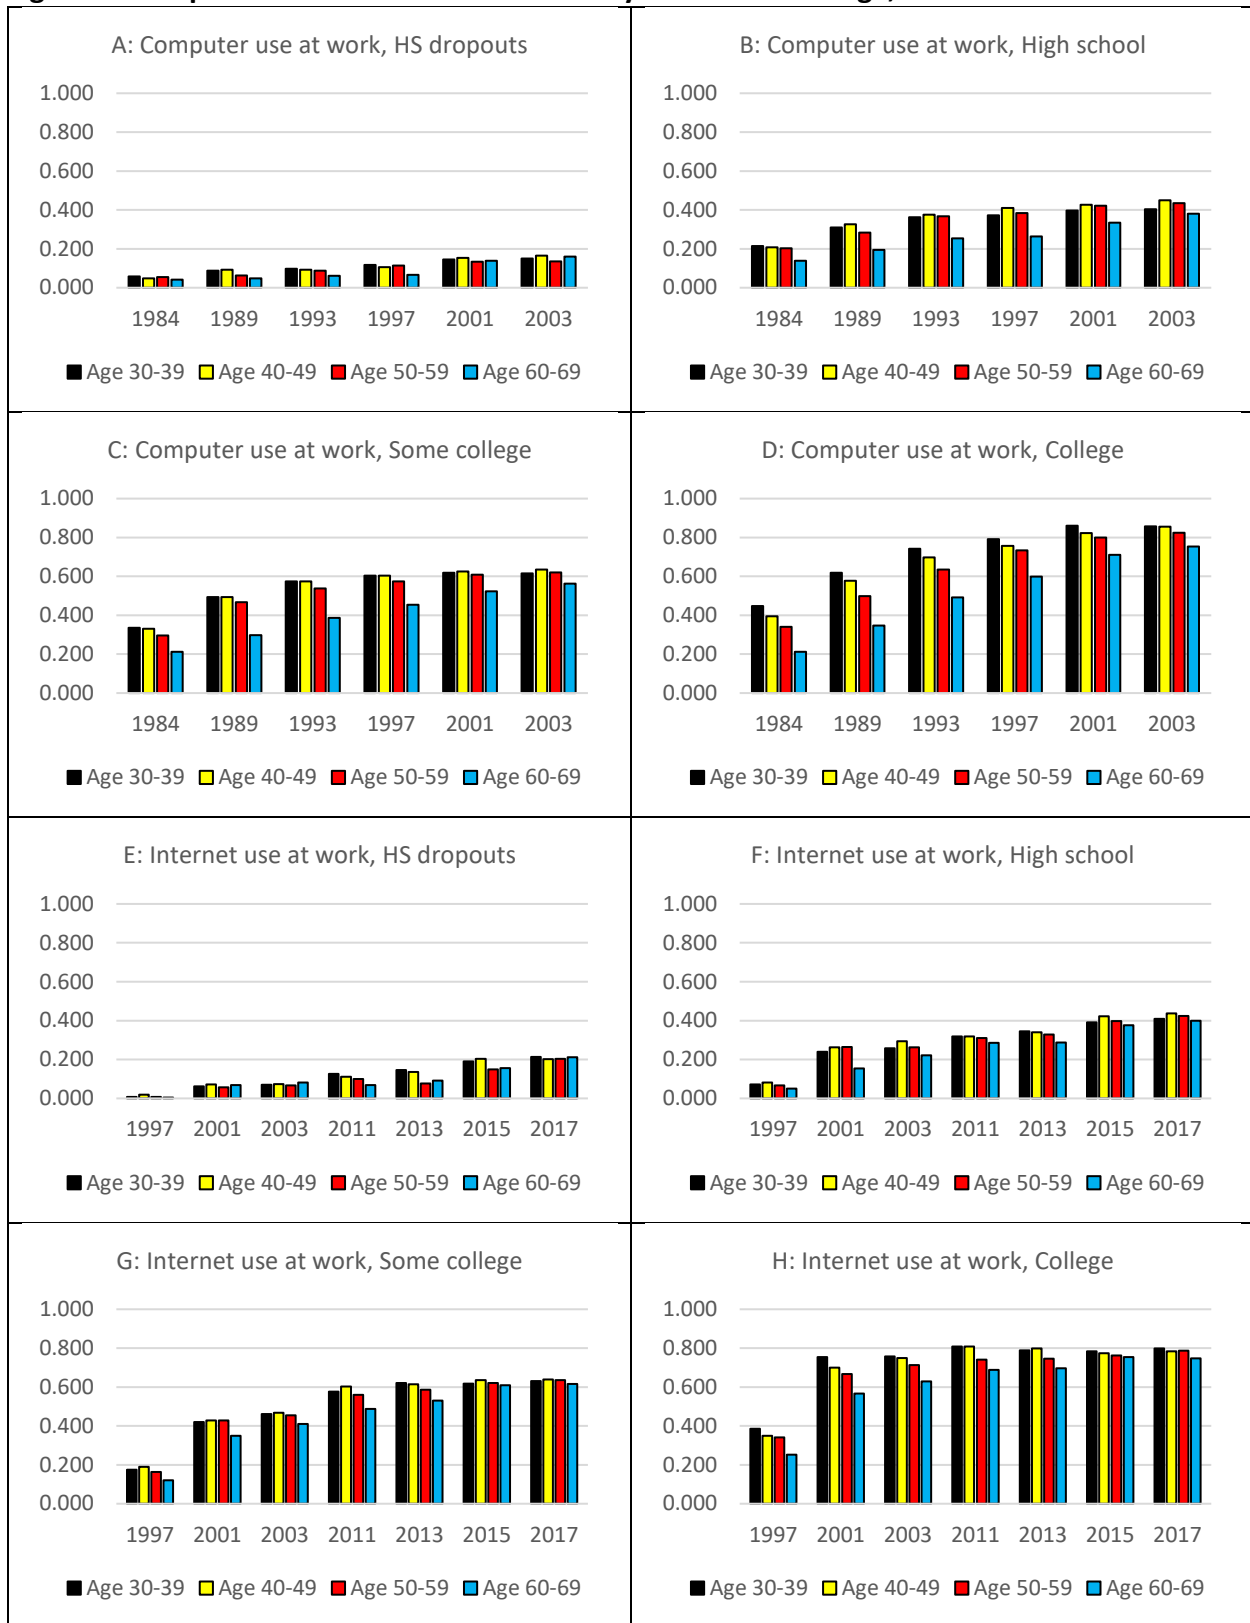
Figures and Tables

Figure 1. Estimated computer and Internet use by the year of labor market entry and labor market experience, CPS 1984-2017



*Model estimates based on data from the CPS Computerization supplements, 1984-2017. The x-axis represents Mincerian labor market experience defined as age minus years of education minus 6. The six lines correspond to cohorts that entered the labor market in consecutive decades.

Figure 2. Computer and Internet use at work by education and age, CPS 1984-2017



*Data: CPS Computerization supplements, 1984-2017

Figure 3. Estimated knowledge gap in computer use by education and age, CPS 1984-2017




Table 1. Computer Innovations by Life Cycle Stage of Workers

Birth year	Years when workers...				
	... graduated HS	... graduated college	... gained 10 years of experience**	... gained 20 years of experience	... turned 65
1920	1938	1942	1952	1962	1985
1930	1948	1952	1962	1972	1995
1940	1958	1962	1972	1982	2005
1950	1968	1972	1982	1992	2015
1960	1978	1982	1992	2002	2025
1970	1988	1992	2002	2012	2035
1980	1998	2002	2012	2022	2045

*Innovations in computer technology

	Digital computer invented		World wide web launched
	Mainframe introduced		Growth of internet
	Mini-computer introduced		Smart phones, apps
	Personal computer introduced		

** The table shows potential labor market experience of a college graduate worker.

Table 2. Linear probability models of yearly transition probabilities between labor market states

	Employment → Nonparticipation			Employment → Unemployment		
	Age 50-69 DD model [1]	Age 40-69 DDD model [2]	Age 25-69 DDD model [3]	Age 50-69 DD model [4]	Age 40-69 DDD model [5]	Age 25-69 DDD model [6]
Knowledge gap	0.103 [0.010]***	0.122 [0.009]***	0.140 [0.008]***	-0.012 [0.003]***	-0.004 [0.003]	0.001 [0.003]
Skill, year, month, age FE	Yes	Yes	Yes	Yes	Yes	Yes
Skill X year FE	-	Yes	Yes	-	Yes	Yes
R-squared	0.075	0.182	0.143	0.025	0.132	0.104
N	2,635,562	4,782,190	7,380,573	2,635,562	4,782,190	7,380,573
Mean outcome	0.080	0.059	0.054	0.018	0.019	0.021

* Basic monthly CPS, 1984-2018. Skill fixed effects include full interactions between gender, education, race & ethnicity, occupations (191 categories), and five-year age categories from age 50-54 to 65-69. In the DDD models individuals younger than age 50 are randomly assigned to one of the four age categories. Age fixed effects include single age dummies, without rewriting the ages of individuals below age 50. Skill X year fixed effects include full interactions between the skill fixed effects and years.

Table 3. Linear probability models of yearly transition probabilities between job types among those who remain employed

	Full-time → Part-time		Employment → Lower paid occupation		Employment → Less computerized occ.	
	Age 50-69 DD model [1]	Age 40-69 DDD model [2]	Age 50-69 DD model [3]	Age 40-69 DDD model [4]	Age 50-69 DD model [5]	Age 40-69 DDD model [6]
Knowledge gap	0.057 [0.009]***	0.061 [0.008]***	0.001 [0.010]	-0.046 [0.010]***	-0.003 [0.009]	-0.029 [0.009]***
Skill, year, month, age FE	Yes	Yes	Yes	Yes	Yes	Yes
Skill X year FE	-	Yes	-	Yes	-	Yes
R-squared	0.060	0.164	0.150	0.249	0.126	0.232
N	2,026,269	3832080	2,375,686	4,407,937	2,375,686	4,407,937
Mean outcome	0.053	0.046	0.157	0.159	0.112	0.113

* Basic monthly CPS, 1984-2018. See Table 2 for definitions of the control variables.

Table 4. Linear probability models of yearly transition probabilities in various population subgroups, coefficients on knowledge gap

<i>Subgroups</i>	Emp → OLF		Emp → Unemp		Full → part-time	
	coef	s.e.	coef	s.e.	coef	s.e.
Total sample	0.122	[0.009]***	-0.004	[0.003]	0.061	[0.008]***
Males	0.073	[0.010]***	-0.001	[0.004]	0.041	[0.009]***
Females	0.178	[0.016]***	-0.010	[0.005]**	0.091	[0.016]***
High school	0.132	[0.019]***	0.027	[0.007]***	0.082	[0.019]***
Some college	0.188	[0.020]***	-0.005	[0.007]	0.120	[0.018]***
College	0.106	[0.011]***	-0.019	[0.004]***	0.036	[0.010]***
Age 50-54	0.053	[0.017]***	-0.014	[0.013]	0.022	[0.017]
Age 55-59	0.142	[0.015]***	-0.015	[0.009]*	0.031	[0.013]**
Age 60-64	0.169	[0.014]***	-0.005	[0.005]	0.042	[0.011]***
Age 65-69	0.099	[0.014]***	0.000	[0.004]	0.089	[0.014]***
Private sector	0.145	[0.011]***	-0.008	[0.004]*	0.089	[0.010]***
Public sector	0.148	[0.021]***	-0.002	[0.005]	0.009	[0.016]
Self-employed	0.028	[0.018]	-0.001	[0.005]	0.019	[0.021]
Management	0.170	[0.015]***	-0.034	[0.006]***	0.071	[0.012]***
Professionals	0.129	[0.017]***	-0.005	[0.005]	0.016	[0.015]
Sales	0.130	[0.021]***	-0.021	[0.009]**	0.130	[0.022]***
Office/Administration	0.237	[0.023]***	-0.011	[0.008]	0.154	[0.023]***
Precision production	0.097	[0.043]**	0.045	[0.019]**	-0.138	[0.031]***
Operators/laborers	0.025	[0.064]	0.102	[0.030]***	-0.135	[0.052]***

* Basic monthly CPS, 1984-2018. Each coefficient corresponds to a different estimation model. The specification is the DDD model including skill, age, year, month, and skill X year fixed effects, and including 40-69-year old workers. Each model was estimated on the population subgroups represented by the rows of the table, except for the age subgroups that always include the 40-49-year old control group. See Table 2 for definitions of the control variables.

Table 5. Models of log annual earnings

	All workers			Full-year workers		
	Age 50-69 DD model [1]	Age 40-69 DDD model [3]	Age 25-69 DDD model [4]	Age 50-69 DD model [1]	Age 40-69 DDD model [3]	Age 25-69 DDD model [4]
Knowledge gap	-0.646 [0.043]***	-0.744 [0.044]***	-0.730 [0.038]***	-0.611 [0.036]***	-0.259 [0.037]***	-0.243 [0.032]***
Skill, year, age FE	Yes	Yes	Yes	Yes	Yes	Yes
Skill X year FE	-	Yes	Yes	-	Yes	Yes
Constant	9.907 [0.010]***	10.107 [0.004]***	9.770 [0.004]***	10.258 [0.009]***	10.322 [0.003]***	10.052 [0.003]***
R-squared	0.358	0.492	0.451	0.409	0.537	0.497
N	674,794	1,361,999	2,429,273	539,748	1,110,225	1,938,382

* CPS-ASEC, 1984-2018. Full-year workers are defined as working for at least 50 weeks a year. Occupations correspond to individual's longest jobs during the previous calendar year. See Table 2 for definitions of the control variables.

Table 6. Models of log annual earnings in various population subgroups

<i>Subgroups</i>	All workers		Full-year workers	
	coef	s.e.	coef	s.e.
Total sample	-0.744	[0.044]***	-0.235	[0.038]***
Males	-0.256	[0.055]***	0.032	[0.047]
Females	-1.278	[0.076]***	-0.586	[0.064]***
High school	-0.897	[0.091]***	-0.268	[0.078]***
Some college	-1.169	[0.096]***	-0.662	[0.084]***
College	-0.553	[0.062]***	-0.058	[0.051]
Age 50-54	-0.255	[0.152]*	-0.275	[0.132]**
Age 55-59	-0.605	[0.105]***	-0.136	[0.089]
Age 60-64	-0.446	[0.074]***	0.013	[0.063]
Age 65-69	-0.990	[0.072]***	-0.421	[0.064]***
Private sector	-0.781	[0.052]***	-0.325	[0.044]***
Public sector	-1.028	[0.118]***	-0.257	[0.075]***
Self-employed	0.430	[0.203]**	0.475	[0.205]**
Management	-0.610	[0.078]***	-0.170	[0.065]***
Professionals	-0.666	[0.089]***	-0.194	[0.069]***
Sales	-0.286	[0.129]**	0.310	[0.119]***
Office/Administration	-1.404	[0.113]***	-0.636	[0.096]***
Precision production	0.046	[0.176]	-0.175	[0.148]
Operators/laborers	1.082	[0.297]***	0.671	[0.288]**

* CPS-ASEC, 1984-2018. The specification is the DDD model including skill, age, year, and skill X year fixed effects, and including 40-69-year old workers. Each model was estimated on the population subgroups represented by the rows of the table, except for the age subgroups that always include the 40-49-year old control group. Full-year workers are defined as working for at least 50 weeks a year. Occupations correspond to individual's longest jobs during the previous calendar year. See Table 2 for definitions of the control variables.

Table 7. Models of additional outcomes from the HRS

	Outcomes measured at baseline age 55-59				Outcomes measured at age 65		
	P65 [1]	log hourly wage [2]	Enjoys work [3]	CESD depression [4]	Works [5]	Works FT [6]	Years worked [7]
Knowledge gap	-0.661 [0.214]***	-0.955 [0.412]**	0.050 [0.325]	-0.141 [1.648]	-0.677 [0.571]	-1.157 [0.492]**	-7.844 [3.367]**
R-squared	0.251	0.515	0.272	0.300	0.293	0.288	0.305
N	11,761	11,041	9,366	9,597	6,162	6,162	6,162
Mean outcome	0.311	2.673	0.877	1.192	0.457	0.242	7.299

* HRS, 1992-2014, age 55-59. The models include skill, year, month, and age fixed effects. The skill fixed effects include full interactions between gender, education, race & ethnicity, occupations (191 categories), and five-year age categories from age 50-54 to 65-69. Baseline outcomes are measured at the first HRS wave when the person is observed in the 55-59 range. Age 65 outcomes are measured at the first time the individual is observed after the age of 65. P65 measures individuals' subjective expectations about the probability of working after age 65. Works FT refers to working full-time. Years worked is measured as the number of years worked between age 65 and the baseline wave.

