CAMBRIDGE
UNIVERSITY PRESS

SPECIAL ISSUE ARTICLE

Forever young: where older workers keep on working

Steven G. Allen* and Ting Wang

Department of Economics, Poole College of Management, NC State University, Raleigh, NC, USA
*Corresponding author. Email: sgallen@ncsu.edu

Abstract
This paper examines inter-industry patterns of the employment of older workers over the last 20 years to understand where employment opportunities have grown the most. The underlying premise is that firms strategically align their age mix depending on production function and labor cost parameters. The industries that had the largest increases in the percentage of older workers were those that had the broadest pension coverage and those that made the greatest use of high-tech capital. There also is evidence in 2001–07 that the percentage of older workers increased more in the industries most exposed to increased Chinese imports.

Older workers are increasingly likely to delay retirement. The employment–population ratio for men 55 and over increased from 35.9% in 1993 to 45.1% in 2019. Over the same period, the ratio for women rose from 22.0% to 34.0%.[1] These upward trends stand in contrast to the downward trajectory of employment odds for younger and middle-aged persons.[2] Longer careers, combined with the large birth cohorts after World War II, have resulted in an aging labor force.

Workforce aging has several important implications. On the surface, longer careers would mean more output and more labor income, thereby leading to higher GDP. Yet questions have been raised about the productivity of older workers and whether it impacts aggregate productivity, especially as technology mastery becomes a more critical skill.[3] Delayed retirements alter human resource decisions by organizations. The uncertainty associated with their future retirement date complicates planning for the future. What happens if they stay 'forever' and what happens if they leave all at once? Are older workers blocking promotion opportunities for others? Government budgets are being impacted as well. Longer careers translate into more earnings leading to greater tax revenue; they also have significant implications for Social Security and Medicare.

There have been numerous studies examining why more older individuals are delaying retirement and working longer. Social Security, private pensions, and retiree health insurance create incentives for individuals to retire, often at specific ages. Personal health and accumulated savings also are critical factors, along with the availability of satisfying and rewarding work. The consensus to date is that people are working longer because of the shift from defined benefit (DB) to defined contribution (DC) pensions, cutbacks in retiree health insurance, Social Security reforms, and increased education.[4]

---

[1]Economic Report of the President 2020.
[2]Abraham and Kearney (2020).
[3]See, for instance, studies by Maestas *et al.* (2016) of the USA and Aiyar *et al.* (2016) of the EU.
[4]Coile (2019).

While we have clear indications of how labor supply factors have led to more employment of older workers, little is known about what has happened on the employer side of the equation. Within an establishment, increased willingness to work longer is unlikely to be matched equally by employer demand. In which sectors of the economy have older workers been able to successfully continue their careers or start new ones? What common characteristics do these sectors have? These are the central questions addressed in this research. The motivation behind this paper is to present a set of stylized facts about employer characteristics (including wages, pension coverage, and technology choices) and employment of workers in different age groups over a 20-year period. As will be discussed in more detail below, there are no truly exogenous variables in the regression analysis, making it difficult to identify specific shifts in labor demand (or supply).

From an employer perspective an increased supply of older workers is a mixed blessing. Older workers create value when they have critical experience, knowledge and skills related to customers, internal operations, market behavior, and suppliers. There also is a tie-in between equipment and software installed 20 or more years ago and the skills that older workers provide to operate and maintain such essential capital. Longer careers mean reduced turnover, leading to savings in search and hiring costs and stronger incentives for training investments. At the same time firms face higher wages and benefit expenses for older workers, along with potential cognitive and physical declines in productivity associated with aging. Ameriks *et al.* (2020) find that older workers prefer more flexible work schedules. This workplace amenity can be provided at low cost in some workplaces, but not so easily in others. In some cases, there is the question of whether older workers can adapt to rapidly changing technology.

This paper contributes to our understanding of how older workers fit into today's organizations by focusing on three key questions: (1) What factors determine the age structure of an organization and how do older workers fit in? (2) What are the stylized facts about the inter-industry age structure for 2001–19? (3) What industry characteristics are associated with changing employment shares of older workers? This paper focuses on economy-wide data on the age structure and work-force characteristics for 60 industries pulled from the American Community Survey (ACS) for 2001–19, matched with data on different types of capital from the Bureau of Economic Analysis and firm and worker characteristics (pension coverage, firm size, unionization) from the Current Population Survey (CPS). The analysis will examine how indicators of demand and supply for older workers within these industries are related to changing employment patterns. There will not be any attempt to estimate structural parameters; instead, the focus will be on developing measures of observable variables that will provide some insights into the varying degrees of growth in the percentage of older workers across different sectors of the economy. Further analysis of data sets that contain sources of exogenous labor demand or supply variables will be needed to determine channels of causation.

What do we know from previous studies? In the 1970s and 1980s a series of papers were written about labor demand for workers from different age groups using translog production (or cost) functions and share equations to estimate elasticities of complementarity (or substitution). In his summary of these studies Hamermesh (1993) concluded that the weight of the evidence indicated that older and younger workers were substitutes. More recent studies by Bianchi *et al.* (2021) and Carta *et al.* (2021) for Italy and Mohnen (2021) for the USA focus on how delayed retirements affect younger workers. The studies by Bianchi *et al.* (2021) and Mohnen (2021) find older and younger workers tend to be substitutes whereas Carta *et al.* (2021) find they are complements.

Regarding the adaptability of older workers to changes in technology, there are two recent studies of note. Hudomiet and Willis (2021) examined the labor market impact of computers on older workers from 1984 to 2017. They found that older workers initially lagged their younger counterparts in terms of computer knowledge but eventually caught up in most arenas. The gaps were largest in the 1980s and 1990s but were still present among workers 65 and over after 2000. Barth *et al.* (2020) focus on software capital intensity and find that middle-aged workers gain more than young and older workers from software investments in terms of wages and employment.

This paper is also related to recent studies by Acemoglu and Restrepo (2020, 2022) in that it uses industry data to examine the interaction of technology and labor market outcomes. Acemoglu and Restrepo (2020) focus on robots and find negative effects on wages and employment. They do not examine the possibility of differential impacts on workers of different age groups. In most of the literature on technological change and labor markets, the presumption is that technological change happens exogenously. Acemoglu and Restrepo (2022) argue that in some instances demographic challenges induce investments in automation. Using data on manufacturing industries for different countries, they show that workforce aging is associated with increased use of robots.

This paper begins in Section 1 with evidence on how the employment of older workers (defined here as those aged 55 and over) has changed since 2000 with a particular focus on their distribution across industries. The results show that there is considerable variation in the degree to which different industries have increased their utilization of older workers. To understand what is driving these changes, Section 2 develops some simple frameworks to explain (1) the determinants of an industry's optimal age structure and (2) the factors driving changes in the age structure over time. The key underlying questions concern why some industries make more intensive use of older workers than others and, given rising labor force participation rates among the elderly, why have some industries hired or retained more older workers than others? The empirical model, a reduced-form employment share equation, is explained in Section 3, with particular attention paid to the choice of independent variables and to potential sources of misspecification. The data sources and variable definitions are summarized in Section 4. The empirical results are reported in Section 5, while Section 6 concludes.

The key findings are that the industries that had the largest increases in the percentage of older workers were the industries that had the broadest pension coverage and the industries that made the greatest use of high-tech capital. There also is evidence in 2001–07 that the percentage of older workers increased more in the industries most exposed to increased Chinese imports (albeit this really means that younger workers suffered larger job losses in those industries than older workers). We do not find evidence of (1) any shift in demand toward older workers in certain industries (as indicated by relative wages and employment moving in the same direction) and (2) any decline in employment shares of younger workers in the industries that showed the largest increases in the shares of older workers.

1. Descriptive background

Older workers, defined as those who are aged 55 or above, represent a rapidly growing segment of the workforce. In 2000 older workers accounted for 12% of all employees; by 2019 their share had increased to 22%. Over the same period, the share of younger workers (16–29) held roughly constant in the 26–27% range whereas the share of middle-aged workers (30–54) fell from 60% to 52%. The drop in the share of middle-aged workers largely reflects the aging of the baby boomer population.

Before turning to the distribution of older workers by industry, it is instructive to take a quick look at how worker and job characteristics have changed over the last 20 years for this age group. Appendix Table A1 reports data on older workers from the American Community Survey public use files for 2001 and 2019. As the older cohorts exit from the labor market and are replaced by younger cohorts, older employees in 2019 have more years of schooling than their counterparts in 2001. The occupational mix adjusted as well, with more managerial and professional jobs and fewer positions in sales, support, and production occupations. The overwhelming majority of older employees work 35 or more hours per week. The shape of the age distribution of older employees changed between 2001 and 2019, with a shrinking percentage in the 55–59 bracket and a growing percentage in the 60 and over bracket. This reflects both cohort aging and longer careers.

Average hourly earnings of older workers rose by 66.6% ($20.89–34.80) between 2001 and 2019. As a benchmark the annual personal consumption expenditure deflator rose by 38.0% over this period. It should be kept in mind that the composition of the older workforce changed considerably over this period, with higher levels of education and a larger percentage of professional workers as noted above. Over the same period, average hourly earnings increased by 55.4% for 16- to 29-year-olds

($12.14–18.88) and by 58.7% for 30- to 54-year-olds ($20.00–31.76). On the surface this might infer that demand for older workers increased over this period in that they gained in terms of both employment share and relative earnings. It also is possible that educational attainment among older workers increased more for older workers during this period than it did for middle-aged or younger workers.

Labor mobility, or lack thereof, is a critical driver of any changes to be observed in the age distribution of older workers across industries. Allen (2023) demonstrated two notable trends over the last 20 years. First, job tenure for older workers increased for those 65 and over. Among women the percentage who have been with their employer for 20 or more years increased from 23% to 26% between 2000 and 2020. The same share for 65 and older men increased from 28% to 34% between 2000 and 2008. It then dropped to 30% in 2010 as a consequence of the Great Recession and essentially stayed at that level through 2020. There are also slight increases in long-term job holding for men (30.6–31.4 years) and women (24.3–24.6) in the 55–64 age bracket.

Second, the odds that an older worker is in a new job have decreased substantially. Among men, 11.2% of those aged 55–64 and 14.4% of those 65 and over were in new jobs in 2000. By 2020, these percentages dropped to 9.8% and 9.7%. The percentage of women with a new employer declined as well, from 11.3% to 9.5% for the 55–64 age group and from 11.5% to 8.2% for the 65 and over group. The main driver for any changes in the age distribution of workers in an industry will be workers staying in their jobs longer as opposed to a rising share of older workers pursuing new work opportunities.

The percentage of workers aged 55 and above in an industry is an intuitive measure at a given point in time of how intensively older workers are employed in that industry. A key question in this paper is how and why this percentage varies across industries and within industries over time. Figure 1 displays a scatterplot of this percentage for 60 industries in 2001 and 2019; the raw data are reported in Table A2. It shows that there is a very wide range in the share of older workers across industries in both years. It also displays a strong positive correlation between the percentage of older workers in a given industry in 2001 and 2019; nearly 40% of the variation in the 2019 values can be explained by the 2001 values in a simple ordinary least squares (OLS) regression.

As a benchmark, the average industry had 13% older workers in 2001 and 23% in 2019, a 10-point increase. Table A3 shows that it increased at a slightly higher rate in the 17 industries with the highest average share of older workers (11.0 points) as opposed to the 14 with the lowest (8.2 points). Among the industries that have the highest share of older workers, the increase in their share ranges from as low as 2 percentage points (railroads) to as much as 15 points (paper, utilities). Looking at the industries that have the lowest share of older workers, the change ranges from a low 3 points in motion pictures and restaurants to a high of 13 points in rental and leasing services. *The change in the age mix of employees is far from a uniform process across industries.*

A simple shift-share analysis shows that changes in the distribution of all workers across industries has had no impact on the overall change in the percentage of workers who are 55 or older. Using 2001 employment weights, the mean percentage of older workers in ACS is 12.6%. Continuing to use the 2001 values of the percentage of older workers but shifting to 2019 employment weights, the mean percentage of older workers is 12.5%. This indicates that we need to focus on within industry changes, not shifts in the mix of industries.

With the percentage of older workers growing in every industry, the percentage of workers younger than 55 must fall correspondingly. One issue of concern is whether the rising share of older workers has made it more difficult for younger workers to get hired. Such displacement could take place in organizations if older workers delay retirement and the establishment's headcount remains constant. This is a modern twist on the lump of labor argument and does not consider (1) overall economic growth and (2) births and deaths of establishments. Alternatively, the rising share of older workers could merely reflect the aging of baby boomers. Appendix Figures A1 and A2 display the raw data for the changes in the percentages of younger (16–29), middle-aged (30–54), and older workers by industry. There is no correlation in Figure A1 between the growth in the percentage of older workers and the change in the percentage of younger workers across industries. This casts initial doubt
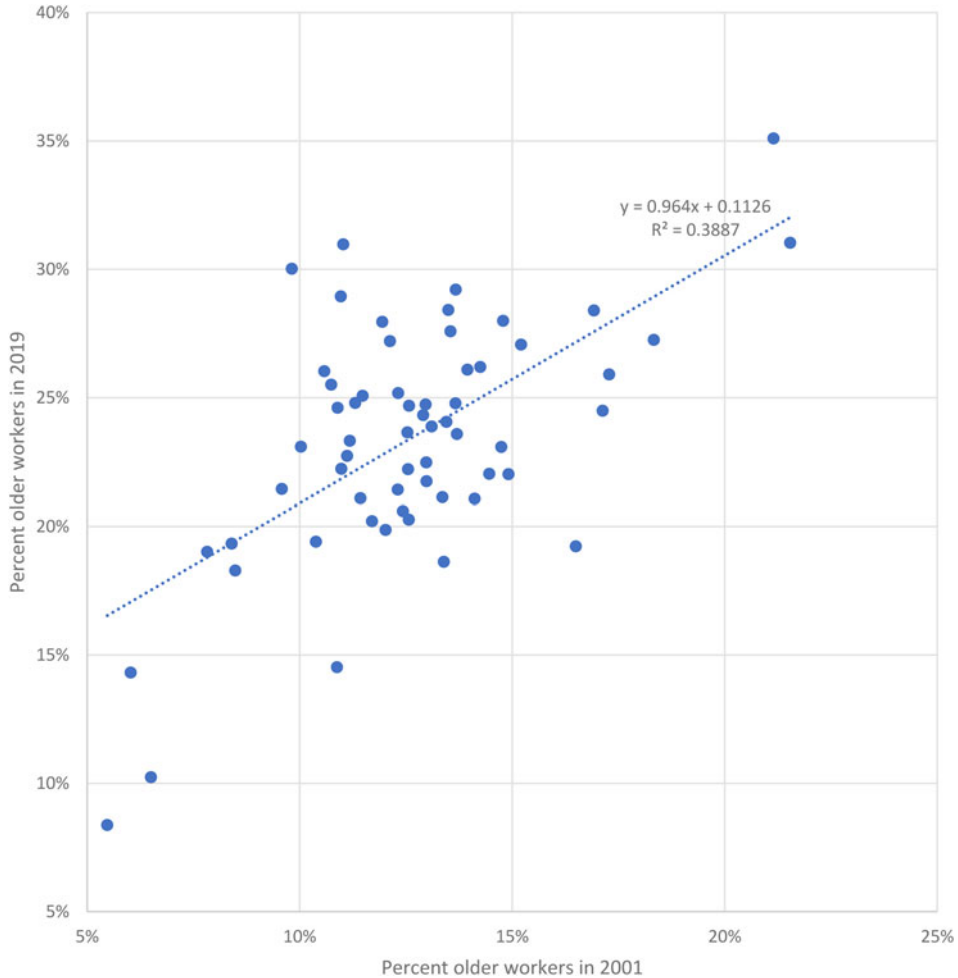
Figure 1. Percentage of older workers by industry, 2001 to 2019 American Community Survey.

concerning the issue of potential displacement of younger cohorts by older generations. By necessity, this leaves us with an inverse relationship between the growth in the percentage of older workers and the change in the percentage of middle-aged workers in an industry, as displayed in Figure A2. This most likely reflects the aging patterns within each industry; most middle-aged workers in 2001 fall into the 55-plus category by 2019.

2. What determines an industry's age mix?

This study starts with the basic premise of labor demand: firms select optimal combinations of different types of labor and capital based on their relative productivity and cost, while also considering which combinations work best and the adjustment costs associated with changing the mix of inputs. The age-mix of any company or establishment hinges upon a wide range of considerations, including decisions about organizational boundaries, job design, hiring, training, and compensation.[5] The age structure of an industry reflects at an aggregated level the hierarchies and organizational structures

---

[5]Clark and Ghent (2010) and Grund and Westergård-Nielsen (2008) are early studies that examined age structures of firms.

of the establishments in that industry. Keep in mind that the mix of firms within each industry changes over time. Facebook and Tesla are not part of the data sets in the early 2000s whereas Enron and Lehman Brothers are not there in the 2010s.

Here we will consider younger, middle-aged, and older workers as separate inputs into a production function. This function also includes multiple types of capital, including specific types of equipment, structures, and intellectual property as inputs. The age-mix of the workforce reflects these considerations:

The productivity and cost of different types of labor: In the simplest production function framework, firms adjust the mix of age groups just as they would any set of inputs so that the ratio of marginal product to the cost of each input is equalized. Suppose initially that workers of all ages are doing the same jobs. Theory and evidence summarized in Allen (2023) both indicate an upward sloping age–earnings profile, albeit with some flattening around age 62. Non-wage labor costs associated with DB pensions, health insurance, and paid time off also increase with age. If there are not corresponding increases in productivity, a firm will use relatively more younger workers. The evidence on age and productivity, summarized in Allen (2023), does not indicate that there is a corresponding upward sloping age–productivity profile.

Complements or substitutes: In most organizations there is a hierarchy of jobs, and within that hierarchy there is a sorting by age group into the various positions. If the hierarchy has multiple levels and the knowledge and skills needed vary across these levels, then older and younger workers would serve complementary roles. For instance, younger workers might have a relative advantage in production tasks whereas older workers might have an advantage in managing customer or supplier relationships. Even within a job category, Bersin and Chamorro-Premuzic (2019) argue that a mix of older and younger workers leads to increased cognitive diversity and improved performance. In situations where jobs are homogeneous, younger and older workers would be substitutes.

Hiring costs and human capital: There are fixed costs associated with bringing a new worker into an organization and training that person. These costs vary across and within organizations. In cases where it is easy to find new workers and there are negligible training costs, then we would expect firms to be indifferent about their age mix as long as wage rates are being held constant. Firms faced with sizable up-front investments in new hires have a retention incentive that should result in more older workers on the payroll. Firms that make significant investments in firm-specific human capital will be especially more interested in hiring young workers and retaining older workers.

Customer demand: Although they must be mindful of age discrimination issues, some employers consider customer preferences when making staffing decisions about customer-facing jobs, especially in youth-oriented industries such as entertainment and hospitality. There is a flip side. Customer preferences help explain why shoe repair, retail florists, religious organizations, and funeral homes are among the industries with the highest percentages of older workers.[6]

Technology: Whether based on perceptions or reality, the adaptability of workers to changes in technology influences employer decisions about age mix. Younger persons have more recent training, as well as a longer time horizon to benefit from investments in new knowledge and skills, than older workers. In contrast, older workers have more perspective from their experience and will accurately realize in some cases that the latest new way of doing things need not be the most appropriate.

In light of these considerations, what would we expect age distributions to actually look like? The age distribution would be flat in situations where the firm stays the same size, hires at the entry level, emphasizes training and promotion from within, and sets new hires equal to retirements. In firms that hire at multiple levels of the age distribution, there would still be a flat distribution if hires equal separations within each age group. The age structure should tilt heavily toward young workers if human capital investments are modest or if there is an inverted-pyramid hierarchy with few opportunities to advance. Such a tilt also would be observed in growing firms if most new hires come

---

[6]The age distribution of employees by industry for each year since 2011 is available at https://www.bls.gov/cps/demographics.htm#age.

from the lower tail of the age distribution. A tilt toward older workers would take place in situations where there has been little hiring for a decade or more, where wage differentials by age are less than the productivity differentials, or where customer preferences dictate a tilt toward older workers.

2.1 Changes in the age structure

From a strategic perspective, the discussion so far has focused on how employer decisions about capital, employment, and compensation determine the age structure of an organization. This structure changes over time as the firm executes on that strategy through hiring and separation decisions. In any given period, there will be persons in different age brackets who leave by quitting, being terminated or retiring. At the same time there will be others who arrive as either new hires or rehires. The age structure will tilt toward older workers if many young workers leave or more older workers get hired, whereas it will tilt toward younger workers if many older workers exit, or more younger workers get hired. Demographics accentuate these forces. Many organizations hired lots of baby boomers in the 1960s and 1970s, creating a bulge in the left tail of the age-distribution at that time, followed by a bulge in the middle 20 years later and a bulge on the right in the 2010s.

The underlying economic forces behind changes in the age structure can best be categorized into a set of supply and demand considerations. It is entirely possible that these changes can be accounted for by decisions by individual workers to delay retirement (along, of course, with decisions by their employers to hire or retain them). The Social Security reforms of 1983 created incentives to delay retirement across all industries. In addition, many firms switched from DB to DC pensions in the 1980s and 1990s and very few new firms that have opened in the last 40 years have set-up DB plans. DB plans typically have strong incentives for retirement before age 65, whereas the benefit stream from a DC plans is age neutral. As a result, we would expect to see more delayed retirement in industries where the shift from DB to DC plans has been the greatest. Employee health and longevity also could be a factor leading to changes in the age distribution at different rates across industries. Delayed retirement is more likely to happen among workers who are healthier and expect to live longer.

On the demand side, there are several forces that could have differential impacts on employment by age group. Relative wages play a central role in labor demand theory. In a textbook world, an increase (decrease) in the wages of older workers relative to that of younger or middle-aged workers leads to a decrease (increase) in demand for older workers. The wages we observe in the data reflect a complex selection process. A rising percentage of workers within a cohort leaves the labor market as that cohort ages. It is quite unlikely that labor market exits are pulled evenly throughout the wage distribution. Those with the highest earnings potential may tend to retire early to purchase more leisure. Alternatively, they may delay retirement more than others because their work is more financially rewarding. Similar arguments apply to those with the lowest earnings potential.

Wages reflect employer decisions as well. Some firms pay higher than average wages as part of a conscious strategy to attract and retain the best workers. Firms that become concerned about retaining key senior employees may choose to increase their wages, whereas those that wish to encourage exits may flatten the wage profile. We also should be aware that the observed wage rate of workers in each age category in an industry reflects their educational and occupational composition.

The question of how well older workers interact with changing technology is one of the oldest questions in the social sciences. New technologies such as robotics reduce the physical demands of work and thereby should allow older workers to have longer careers. Often new technologies, especially those related to data and information, require new human capital investments. It could be argued that firms are reluctant to invest in training older workers because they are closer to retirement. However, turnover rates among young workers are considerably higher than those of older workers, so it is unclear *ex ante* which group has the longer expected job duration. Bartel and Sicherman (1993) note that technological change tends to be persistent in many industries, which can result in a sorting of workers in all age categories in terms of their ability to adjust to change. If the change is rapid enough then the shorter time horizon for investment in older workers matters less.

A final potential source of disruption on the labor demand side is import competition and outsourcing. Since 2000 we have seen realignment of supply chains as well as increased imports, especially from China. In some cases, entire establishments closed and there was a reduction in industry employment across all ranges of the age distribution. In others, there were vastly fewer production positions, but managerial and professional jobs survived to some extent, resulting in a larger percentage of older workers.

3. Empirical framework

The central question to be examined in this research is what factors have determined why the share of older workers has risen significantly in some industries between 2001 and 2019 and relatively little in others. A natural choice for a dependent variable is the change in the percentage of older workers in an industry. The appeal of this measure is that it provides information about which sectors are adding or retaining more older workers relative to younger workers. A challenge is that it is difficult to interpret directly as an index of labor demand. An industry could have a rising percentage of older workers simply by having markedly fewer younger or middle-aged workers. This can happen if there are employment reductions and seniority is the key element in retention decisions. It also can happen if an industry stops growing and stops hiring.

Another way to examine the determinants of the changing age mix of employment by industry is to examine reduced form models of changes in employment by age group. To be specific, suppose that the change in employment for each age group a (where y = young, m = middle-aged, and o = old) in industry j can be expressed as

$$dL_{aj} = \theta_a X_j + \gamma_a K_j + \delta_a w_{aj} + \varepsilon_{aj}, \tag{1}$$

where dL indicates change in log employment for age group a in industry j, X is a vector of control variables reflecting industry characteristics, K is the capital–labor ratio, and w is the log wage rate. All right-hand variables are measured at the start of the sample period. Taking first differences of (1) between the demand for older workers and the demand for young and middle-aged workers respectively, we have the estimating equations:

$$dL_{oj} - dL_{yj} = (\theta_o - \theta_y)X_j + (\gamma_o - \gamma_y)K_j + (\delta_o w_{oj} - \delta_y w_{yj}) + (\varepsilon_{oj} - \varepsilon_{aj}) \tag{2a}$$

$$dL_{oj} - dL_{mj} = (\theta_o - \theta_m)X_j + (\gamma_o - \gamma_m)K_j + (\delta_o w_{oj} - \delta_m w_{mj}) + (\varepsilon_{oj} - \varepsilon_{mj}) \tag{2b}$$

An underlying assumption in this framework is that the wage coefficient varies for each age group. If instead $\delta_o = \delta_m = \delta_y$, then we have a simpler framework:

$$dL_{oj} - dL_{yj} = (\theta_o - \theta_y)X_j + (\gamma_o - \gamma_y)K_j + \delta_o(w_{oj} - w_{yj}) + (\varepsilon_{oj} - \varepsilon_{aj}) \tag{3a}$$

$$dL_{oj} - dL_{mj} = (\theta_o - \theta_m)X_j + (\gamma_o - \gamma_m)K_j + \delta_o(w_{oj} - w_{mj}) + (\varepsilon_{oj} - \varepsilon_{mj}) \tag{3b}$$

A key aspect of this approach is that the change in employment is posited as a function of initial values of the right-hand variables. These models are designed to answer the question of what is likely to happen to the employment of older workers in the future based on what we know at a given point in time. These are by no means to be interpreted as structural models of labor demand. However, they should be informative about certain matters. For instance, did the employment of older workers relative to other groups increase, not change, or decrease in industries with high capital–labor ratios or those making intensive use of information technology? Similarly, how does the relative employment of different age groups vary by employer characteristics such as firm size, occupational mix, or pension coverage?

The model includes the overall capital–labor ratio along with interaction terms that allow the capital–labor coefficient to vary for types of capital associated with information technology, instrumentation, and intellectual property. One challenge in this framework is the issue of whether initial values of K can truly be considered exogenous. On the surface one could argue that, since many of the workers aged 55 and over in 2001 were hired in the 1970s and 1980s, their employers could not have foreseen when they were hired the changes in information technology that impacted the workplace in the 2000s and 2010s. However, there are industries where employers have long planning horizons and changes in employment, both overall and by age group, are intricately related to their investment decisions. Because the length of planning horizons and the pace and predictability of technical change varies across industries, it is very difficult to structurally model how investment decisions impact employment of workers in different age groups. The goal here is to observe whether there is a correlation between technology and the employment share of older workers. Because of the strong likelihood that future values of K are related to both (1) initial values of K and (2) changes in employment patterns in some sectors, it will not be possible to make a structural interpretation of the K coefficients.

Two things can be established by examining the impact of including the wage rate variable. First, the raw data show that older workers became a growing share of the labor force over the last 20 years. If the inter-industry correlations show a corresponding increase (decrease) in relative wages over that period, it would provide a signal as to whether a demand (supply) increase was an important factor. Second, it would be worthwhile to determine whether the results for other variables are robust to inclusion of w. The log ratio of average weekly earnings of older workers compared to younger workers (aged 16–29) measures w. To control for worker heterogeneity across industries, controls for worker education, occupation, and gender are included in the model.

Other variables are included in the model to capture market, worker, and workplace characteristics that would likely influence changes in the age structure. Previous studies have shown that the shift from DB to DC pensions has contributed to the rising employment–population ratio for older workers. In the context of this study, this calls not just for the inclusion of a measure of pension coverage, but also for consideration of how many workers are covered by which type of plan. There should be relatively few older workers in industries where DB plans dominate because of the incentives for early retirement often found in such plans. In contrast, DC plans rarely have provisions that spike pension wealth at specific ages or seniority dates, so we should expect a higher percentage of older workers in industries where most workers are covered by DC plans.

Union density is included because firms have fewer degrees of freedom to manage the age mix of employees when covered by a collective bargaining agreement. We also include controls for firm size and the percentage of workers with 20 or more years of tenure. Large firms are likely to provide more career opportunities and paths, so we might expect them to have a higher percentage of older workers. Some firms commit to long-lasting careers for their workers and there would be a higher percentage of older workers in such firms if these relationships are stable over time.

Lastly, we include the rate of output growth and its variance over the 10 years preceding the sample period in the model. Between 2001 and 2019, the log change in real value added ranged between 2.047 in data processing, internet publishing, and other information services and −0.8 in apparel and leather and allied products. During growth periods firms typically add proportionally more young and middle-aged employees than older employees, whereas during periods of industry decline most separations occur among younger workers. The variation in output growth has an independent effect on the age structure. Imagine two industries with the same average growth over a given period, one on a steady path and the other with repeated ups and downs. The latter industry will end up with a sizable contingent of older workers sheltered from the ups and downs, accompanied by fewer younger and middle-aged workers.

4. Data sources

Because the concern here is labor market trends rather than year-to-year fluctuations, the focus will be on three intervals: 2001–2007, 2007–2019, and 2001–2019. The choice of intervals controls for the business cycle; the initial and final year in each period is a business cycle peak.

The main data source for this study is the ACS for 2001 through 2019. ACS data are available for 2000 but the sample size is considerably smaller that year, resulting in noisy measures for smaller industries that become noisier still when first differenced over 19 years. Employees are defined as wage and salary workers; the self-employed and military are not included. The ACS was used to calculate total employment in each industry, along with its age breakdown into three groups (16–29, 30–54, and 55 and over). Average hourly earnings are estimated for each age group using continuous measures of wage and salary income, usual weekly hours, and weeks worked. Measures of percentage college graduates, percentage female, and percentage in managerial or professional occupations also were pulled from the ACS. These percentages are calculated across employees from all age groups combined.

Capital data come from the Bureau of Economic Analysis (BEA) of the US Census, which publishes Current-Cost Net Capital Stock of Private Non-residential Fixed Assets data for 62 industries and 96 categories of capital.[7] In the empirical analysis total capital is defined as the sum of equipment, structures, and intellectual property products (IPP). To determine whether the employment of older workers is impacted by advanced technology, the capital coefficient is allowed to vary with different types of capital associated with information technology, instrumentation, and intellectual property. The measure used in the first set of results reported here includes 14 types of computing, instruments, and office equipment along with five types of IPP. The robustness of the results to alternate definitions is examined below.

Real value added by industry as calculated by BEA was used to construct two control variables: output growth over the 10 years preceding the sample period and the variance of output growth over the same interval.[8]

Industry measures of percentage covered by union contracts, percentage working in large firms, percentage covered by pensions, and percentage of employees with 20 or more years of service were pulled from the Current Population Survey. Lagged values from the 1990s were used to minimize issues related to reverse causation with the dependent variable. The union variable comes from the outgoing rotation groups in the merged 1992–93 CPS public use files. The firm size measure comes from the Annual Social and Economic Supplement (ASEC) for the same years. The years of service data come from the 1996 and 1998 Job Tenure and Occupational Mobility Supplements. The pension coverage data come from the 1999 CPS ASEC so that estimates of the percentage in DB and DC plans could be obtained from the public use files of Internal Revenue Service (IRS) Form 5500, using data on plan participants and plan characteristics. Public use files for Form 5500 are not available for download for years before 1999.

Industry definitions were established by matching North American Industry Classification System (NAICS) codes in the ACS with those used in BEA data sets on output and capital. Governments were not included as separate industries because the capital data pertain solely to the private sector. The empirical analysis is based on 60 industry categories which are listed in the data appendix.[9] The limiting factor for industry definitions was the 62 industry categories in the BEA capital data set. Two pairs of industries had to be combined. Securities, commodity contracts and investments (NAICS code 523) and funds, trusts, and other financial vehicles (NAICS code 525) have the same detailed industry code in the ACS. Hospitals (NAICS code 622) and nursing and residential care facilities (NAICS code 623) are combined in the pre-1997 data used for lagged GDP values.

---

[7]Downloads available at https://apps.bea.gov/national/FA2004/Details/Index.htm.

[8]https://apps.bea.gov/iTable/index_industry_gdpIndy.cfm.

[9]A concordance between the 1990 Census industry codes in the CPS and the NAICS codes in the ACS was developed to allow the CPS data to be merged with the ACS data. In two cases a 1990 Census industry was split into two NAICS industries: (1) oil and gas extraction and mining and (2) banking and credit intermediation and securities, commodities, funds, trusts, and other financial instruments. The same lagged values were used for each pair of industries.

5. Empirical results

Table 1 reports regressions of the first difference in the percentage of older workers for three different time periods: 2001–2007, 2007–2019, and 2001–2019. Two models are reported for each period: one excluding and one including the ratio of log wages of older workers to younger workers. The key results are as follows:

(1) Pension coverage is the strongest predictor of which industries had the largest growth in the share of older workers. The growth in the percentage of older workers in an industry is largest in those industries where most workers are covered by a pension plan. Compare two industries – one where no workers are covered by a pension and one where all are covered by a pension. Growth in the percentage of older workers was 6.1 percentage points higher between 2001 and 2007 and 12.8 points higher between 2007 and 2019 in an industry where all workers were covered by a pension as compared to an industry where none were covered.

(2) The share of older workers grew most between 2007 and 2019 in the industries with the highest ratios of high-tech capital to labor. The capital–labor ratio itself was weakly and inversely related to the growth in the share of older workers in an industry. To examine the practical magnitude of this relationship, consider the following industry comparison. The mean high-tech capital–labor ratio across all industries in 2007 was 19.5; the standard deviation was 28.5. Compare two industries: one with the mean high-tech capital–labor ratio and one where the ratio was twice as high. Considering that a doubling of the high-tech capital–labor ratio also leads to an increase in the overall capital–labor ratio, the growth in the

Table 1. Regression results: first difference in percent older workers

Dependent variable	Change in percentage of workers 55 and above					
Years	2001–2007	2001–2007	2007–2019	2007–2019	2001–2019	2001–2019
Percent college graduates	0.019	0.0203	−0.108***	−0.122***	−0.108***	−0.110***
	(0.019)	(0.0187)	(0.0390)	(0.0409)	(0.0379)	(0.0376)
Percent female	0.021*	0.0187	−0.0367	−0.0309	−0.00854	−0.00521
	(0.012)	(0.0124)	(0.0231)	(0.0236)	(0.0281)	(0.0285)
Percent professional	−0.079	−0.0688	−0.0791	−0.0812	−0.135	−0.149
	(0.053)	(0.0547)	(0.102)	(0.106)	(0.120)	(0.124)
GDP growth	−0.014*	−0.0181*	−0.0161**	−0.0143**	−0.0126	−0.00760
	(0.008)	(0.00939)	(0.00640)	(0.00682)	(0.0225)	(0.0233)
GDP variance	0.120	0.143	0.0386	0.0149	0.00495	−0.0267
	(0.109)	(0.111)	(0.108)	(0.110)	(0.318)	(0.278)
Ln(K/L)	−0.0012	−0.000807	−0.00534	−0.00576	−0.00547	−0.00603
	(0.0018)	(0.00172)	(0.00470)	(0.00485)	(0.00494)	(0.00492)
Ln(HiTecK/L)	0.0027	0.00340	0.0126***	0.0117**	0.0120**	0.0110*
	(0.0021)	(0.00223)	(0.00421)	(0.00450)	(0.00561)	(0.00603)
Pension coverage	0.058***	0.0591***	0.128***	0.125***	0.200***	0.199***
	(0.017)	(0.0171)	(0.0413)	(0.0412)	(0.0497)	(0.0505)
Percent in large firms	−0.015	−0.0156	−0.0139	−0.0124	−0.0350	−0.0337
	(0.014)	(0.0134)	(0.0374)	(0.0380)	(0.0427)	(0.0445)
Percent long-term workers	0.005	−0.0116	−0.247**	−0.240**	−0.243**	−0.220**
	(0.040)	(0.0401)	(0.101)	(0.0936)	(0.102)	(0.100)
Union coverage	0.011	0.0132	−0.0378	−0.0372	−0.0344	−0.0378
	(0.018)	(0.0166)	(0.0336)	(0.0360)	(0.0407)	(0.0425)
Ln wage ratio 55+ to 16–29		−0.0138		0.0307		0.0189
		(0.00941)		(0.0493)		(0.0333)
Constant	0.010	0.0144	0.0931***	0.0819***	0.0968***	0.0905***
	(0.012)	(0.0124)	(0.0222)	(0.0301)	(0.0274)	(0.0312)
N	60	60	60	60	60	60
R^2	0.563	0.583	0.435	0.443	0.434	0.441

Standard errors in parentheses.
*p < 0.10, **p < 0.05, ***p < 0.01.

percentage of older workers is 0.85 percentage points higher in the industry with the higher high-tech capital ratio. Admittedly this is not a huge difference, especially over a 12-year period. However, the key takeaway is that the employment share of older workers was growing fastest in high-tech industries.

(3) The percentage of older workers in an industry decreased in those industries with the most rapid output growth in the 10 years preceding the sample period. This no doubt reflects the fact that expanding firms hired mostly young workers. The magnitude of the coefficient is small; a 10% increase in output growth is associated with a 0.1 percentage point decline in the percentage of older workers.

(4) The growth in the share of older workers in an industry is unrelated to the relative wage differential between the old and the young. The coefficient of the wage variable is both small and measured with little precision. Further, the coefficients of the other variables are not sensitive to the addition of the wage variable. There is no evidence of increased demand for older workers within industries.

(5) There were two counterintuitive results. The growth in the share of older workers was lowest in industries with the largest share of college graduates and in industries that had relatively large numbers of workers with 20 or more years of tenure in the 1990s. This could reflect higher initial values of the percentage of older workers in these industries that limited opportunities for large first differences.

(6) There is not a stable relationship between the change in the share of older workers and the independent variables examined here. The coefficients for 2001–2007 and 2007–2019 are often quite different, sometimes with opposite signs.

These results are based on first difference analysis of employment levels by age group in 60 industries. Employment levels vary considerably across industries. The average industry in the sample had 1.9 million workers in 2001, with a range from 33,000 (pipeline transportation) to 14 million (retail trade). If some very large industries ended up being statistical outliers, the results could change considerably if the regression were weighted by industry size. Also, hours per person vary across industries, with some making intensive use of part-time workers and others relying heavily on overtime. To address these considerations, the models in Table 1 were re-examined by (1) weighting each observation by the sum of employment at the beginning and end of the sample period and (2) recalculating the dependent variable so that it is measured in labor hours instead of employees. The results, reported in Tables A4 and A5 in the Appendix, show that neither weighting the regression nor changing the dependent variable made any meaningful difference in the findings.

The pension coverage variable does not distinguish between DB and DC plans. Estimates of the percentage covered by each type of plan were derived from IRS Form 5500 data for 1999 by counting the number of active participants in DB and DC plans for each of the 60 industry groups. The models in Table 1 were then re-estimated using separate variables for the percentage of workers covered by DB and DC plans; the results are reported in Table A6 in the Appendix. In 2001–2007 and 2007–2019 the coefficient for DC plans is roughly the same as the coefficients for all plans in Table 1, whereas the coefficient for DB plans is not measured precisely. For the 2001–2019 period, the coefficient for DB coverage is 50% larger than the coefficient for DC coverage but the hypothesis that the two coefficients are equal cannot be rejected. Although the impact of DB and DC plans on retirement and retention might be expected to be different, this does not show up in this exercise.

One might question how robust the results for high-tech capital in Table 1 are to potential changes in the classification of which types of capital are considered high tech. The measure used in the regression models includes 14 types of equipment and five types of intellectual property products (IPP). The BEA provides data on 20 additional categories of IPP. To examine the sensitivity of the results to the inclusion of IPP, three additional high-tech capital–labor ratio variables were examined. One of the variables included high-tech equipment but no IPP. One added software-related IPP, namely

prepackaged software, custom software, and own account software, while omitting software publishing and computer systems design. The third added all 25 IPP categories to the high-tech equipment, including IPP in aerospace, motor vehicles, and pharmaceuticals. As shown in Table A7, the coefficients of the first two alternate variables in 2007–2019 are like that of the one in Table 1. In contrast, the coefficient for the most inclusive measure is consistently smaller than its standard error. Overall, it is clear that as long as high-tech capital is defined as the usage of information technology then the results for this variable are quite robust. On the other hand, the presence of other types of IPP in an industry does not seem to have any relationship to its increased usage of older workers.

This empirical framework has focused on aggregate industry data. It is entirely possible that the changes in the percentage of older workers within an industry vary by skill; managerial and professional workers may have delayed retirement but those in other jobs have not. To explore this issue, the industry cells in ACS were split by education level, with one set of data points corresponding to those with 12 years or less of completed schooling and another set consisting of those with more than 12 years of schooling. Industries with a high percentage of older workers in one schooling group also tend to have a high percentage of older workers in the other group; the correlation coefficients for the 60-industry sample are 0.57 in 2001 and 0.70 in 2019. Also, industries with the biggest increases in the percentage of older workers who have higher education levels between 2001 and 2019 tend to be the industries with the largest increases in that percentage for workers with less education; the correlation coefficient is 0.58. The message to draw from this exercise is that the aging patterns within an industry tend to be similar across education levels.

Although aging patterns within industries tend to be similar across different education levels, there is still the question of whether the coefficients for key variables such as pension coverage and high-tech capital vary by schooling levels. The model in Table 1 was re-estimated over two different samples: one where the variables from the ACS all pertained to those with 12 years of schooling or less and another where the ACS variables all pertain to those with more than 12 years of schooling. This framework allows for complete interactions between all the independent variables and the two schooling categories. The results for the pension and high-tech capital–labor ratio variables are reported in Table A8.

Pension coverage continues to be strongly associated with a rising percentage of older workers in an industry for both groups over the entire sample period and for those with 12 years or less of schooling in 2007–2019. The association between pension coverage and aging is weaker for both schooling groups in 2001–2007 and for the group with higher schooling in 2007–2019. One possible explanation for this pattern is that pensions for those with no post-secondary schooling shifted from DB to DC plans over the sample period whereas pensions for those with post-secondary schooling were DC throughout, resulting in more delayed retirement among those with no post-secondary schooling.

The high-tech capital–labor ratio results from Table 1 continue to hold for workers with post-secondary schooling, but not for their counterparts with less schooling. The coefficients in Table A8 of this variable for 2007–2019 for the two groups are quite close to each other; they are also quite close to the coefficient in Table 1. However, the standard errors for the group with no post-secondary schooling are considerably larger. The argument that older workers are complements to high-tech capital appears to hold much more strongly for those with post-secondary schooling, perhaps because more schooling is needed to be an effective complement to such capital.

The results reported so far reflect changes in the share of older workers compared to the combined share of younger and middle-aged workers. The last logical step is to compare the growth in employment of older workers to that of each of the other two groups by estimating equations (3a) and (3b). These estimates are reported in Tables 2 and 3, using the same independent variables as before.

Overall, the main results from the analysis of inter-industry differences in the percentage of older workers hold when we separately examine young-old and middle-aged-old differences. Employment growth of older workers was more rapid than employment growth of younger workers between 2001 and 2007 in industries with higher levels of pension coverage. This was the case for the employment growth differential between older and middle-aged workers for the entire sample period. Industries with the highest ratios of high-tech capital to labor saw the largest growth in the

Table 2. Regression results: relative log employment growth of older and younger workers

Dependent variable Years	Log employment change of workers 55 and over – log employment change of workers 16–29					
	2001–2007	2001–2007	2007–2019	2007–2019	2001–2019	2001–2019
Percent college graduates	−0.0495	−0.0285	−0.471	−0.623**	−0.611	−0.615
	(0.289)	(0.275)	(0.296)	(0.293)	(0.375)	(0.377)
Percent female	−0.0321	−0.0643	−0.368*	−0.305*	−0.348	−0.341
	(0.203)	(0.193)	(0.189)	(0.174)	(0.271)	(0.277)
Percent professional	0.340	0.469	−0.284	−0.307	−0.0648	−0.0921
	(0.588)	(0.628)	(0.729)	(0.738)	(1.082)	(1.094)
GDP growth	0.000295	−0.0484	−0.132***	−0.113**	0.00389	0.0142
	(0.0953)	(0.115)	(0.0396)	(0.0450)	(0.120)	(0.141)
GDP variance	−1.422	−1.115	0.941*	0.680	−2.625	−2.690
	(1.348)	(1.326)	(0.519)	(0.509)	(1.613)	(1.645)
Ln(K/L)	−0.121***	−0.116***	−0.0148	−0.0195	−0.124**	−0.125**
	(0.0348)	(0.0351)	(0.0298)	(0.0304)	(0.0473)	(0.0501)
Ln(HiTecK/L)	0.0772**	0.0870**	0.0921**	0.0823**	0.150**	0.148**
	(0.0356)	(0.0372)	(0.0407)	(0.0394)	(0.0700)	(0.0720)
Pension coverage	0.583**	0.595**	−0.0389	−0.0659	0.666	0.663
	(0.285)	(0.278)	(0.314)	(0.325)	(0.525)	(0.535)
Percent in large firms	0.0456	0.0332	0.125	0.142	0.140	0.143
	(0.270)	(0.262)	(0.252)	(0.245)	(0.441)	(0.450)
Percent long-term workers	−0.353	−0.569	−1.037	−0.955	−1.464	−1.419
	(0.655)	(0.666)	(0.631)	(0.684)	(0.981)	(1.068)
Union coverage	−0.187	−0.154	−0.616***	−0.609***	−0.875***	−0.882***
	(0.255)	(0.252)	(0.221)	(0.217)	(0.297)	(0.306)
Ln wage ratio 55+ to 16–29		−0.183		0.339		0.0386
		(0.148)		(0.256)		(0.220)
Constant	0.533***	0.595***	0.750***	0.626***	1.236***	1.223***
	(0.179)	(0.177)	(0.185)	(0.169)	(0.220)	(0.217)
N	60	60	60	60	60	60
R^2	0.458	0.480	0.460	0.479	0.423	0.423

Standard errors in parentheses.
*p < 0.10, **p < 0.05, ***p < 0.01.

employment of older workers relative to employment of the other age groups (although the relationship for middle-aged workers is somewhat weak in 2001–2007). Output growth narrowed the spread in employment growth between older workers and their middle-aged and younger counterparts. The anomalous results for industries with high percentages of college graduates and industries with a tradition of having long employment relationships continue to appear.

Three new patterns do appear in Tables 2 and 3. First, the growth of employment for older workers was much slower than that for younger workers in industries with high percentages of unionized workers. These industries tend to still have DB pensions in place for production workers, potentially limiting the growth of employment of older workers. Union density has a mean of 19% and a standard deviation of 16%. Compare two industries, one with 20% union members and one with 40% union members. The industry with more union members would have 0.12 higher log employment growth among younger workers than older workers between 2007 and 2019. Second, the growth in employment of older workers compared to middle-aged and younger workers was lower in industries with high percentages of women. This could reflect higher percentages of women in the 2 younger age groups. Third, the change in employment of older workers relative to young or middle-aged workers was lowest in industries with the highest overall capital–labor ratios. This relationship was particularly strong for the old versus young results in Table 2 for 2001–2007. Perhaps older workers are substitutes for traditional capital equipment and complements for high-tech capital.

In summary, two major themes arise from this analysis. First, pension coverage and pension characteristics are strongly related to changes in the age structure of industries over time. One possible interpretation of the results is that as the first generation of workers covered largely by DC plans

Table 3. Regression results: relative log employment growth of older and middle-aged workers

Dependent variable Years	Log employment change of workers 55 and over – log employment change of workers 30–54					
	2001–2007	2001–2007	2007–2019	2007–2019	2001–2019	2001–2019
Percent college graduates	0.268	0.267	−0.615***	−0.785***	−0.437*	−0.466**
	(0.192)	(0.195)	(0.223)	(0.231)	(0.230)	(0.229)
Percent female	0.0325	0.0341	−0.296**	−0.225*	−0.244*	−0.200
	(0.116)	(0.120)	(0.114)	(0.114)	(0.144)	(0.142)
Percent professional	−0.546	−0.552	−0.169	−0.195	−0.470	−0.648
	(0.446)	(0.443)	(0.591)	(0.615)	(0.709)	(0.701)
GDP growth	−0.113*	−0.111*	−0.0797**	−0.0578	−0.0877	−0.0206
	(0.0581)	(0.0596)	(0.0387)	(0.0408)	(0.125)	(0.109)
GDP variance	1.225	1.210	0.779	0.486	0.869	0.447
	(0.893)	(0.869)	(0.730)	(0.701)	(2.271)	(1.699)
Ln(K/L)	0.00259	0.00232	−0.0600**	−0.0652**	−0.0521	−0.0597*
	(0.0170)	(0.0176)	(0.0299)	(0.0323)	(0.0315)	(0.0313)
Ln(HiTecK/L)	0.0176	0.0171	0.0628**	0.0518*	0.0562*	0.0427
	(0.0181)	(0.0189)	(0.0279)	(0.0296)	(0.0332)	(0.0336)
Pension coverage	0.375**	0.374**	0.691**	0.661**	1.166***	1.149***
	(0.164)	(0.165)	(0.268)	(0.254)	(0.295)	(0.298)
Percent in large firms	−0.0530	−0.0524	0.0492	0.0680	−0.0172	−0.000118
	(0.131)	(0.133)	(0.222)	(0.216)	(0.249)	(0.263)
Percent long-term workers	−0.0917	−0.0810	−1.355**	−1.263**	−1.421**	−1.124*
	(0.350)	(0.380)	(0.659)	(0.579)	(0.630)	(0.646)
Union coverage	0.00741	0.00577	−0.247	−0.240	−0.294	−0.340
	(0.134)	(0.136)	(0.199)	(0.214)	(0.223)	(0.244)
Ln wage ratio 55+ to 30–54		0.00907		0.379		0.252
		(0.0744)		(0.283)		(0.157)
Constant	0.0952	0.0922	0.732***	0.594***	0.778***	0.694***
	(0.104)	(0.110)	(0.146)	(0.189)	(0.164)	(0.170)
N	60	60	60	60	60	60
R^2	0.548	0.548	0.443	0.477	0.478	0.516

Standard errors in parentheses.
*p < 0.10, **p < 0.05, ***p < 0.01.

hit the later stages of their working career, they decided to stay on the job longer. At the same time the union results in Table 2 suggest that DB plans work in the opposite direction, generating earlier exits for older workers and more opportunities for younger workers. Second, there does not seem to be a technological backlash toward the employment of older workers, especially those with post-secondary schooling. To the contrary, there was actually more growth in the percentage of older workers in those industries that use the most high-tech capital.

5.1 Are younger workers being squeezed out?

As more older workers delay retirement, some have speculated that this had resulted in fewer opportunities for younger workers. The simple scatterplots in Figures A1 and A2 indicated that there did not seem to be a relationship between the change in the share of older workers in an industry and the share of younger workers. Now that we know variables such as pension coverage and the capital–labor ratio have an impact on changes in the deployment of older workers in an industry, it would be logical to revisit that issue in a regression framework.

The challenge is that the shares of older, middle-aged, and younger workers in an industry are not independent variables. Further, there is no genuine source of exogenous variation in the ability of firms to hire or retain workers in different age brackets across different industries. What can be done here is to see what the data tell us about any possible tradeoff between jobs for older and younger workers within a given industry. This is done in three ways in Table 4. First, to capture the raw data pattern we report a simple regression of the change in the share of younger workers on the change in the share of older workers. Second, the percentage of older workers will be used as an independent

Table 4. Regression results: relationship between shares of older and younger workers

Dependent variable	Change in percentage of workers 16–29					
Time period	2001–07	2001–07	2007–19	2007–19	2001–19	2001–19
Estimation method	OLS	2SLS	OLS	2SLS	OLS	2SLS
	Coefficient (S.E.) of percent 55 and above					
OLS simple regression	−0.196		−0.022		−0.172	
	(0.225)		(0.114)		(0.150)	
Multiple regression without wage variable	−0.492	−0.018	0.091 (0.140)	−0.091	−0.251	−0.368
	(0.387)	(0.475)		(0.310)	(0.152)	(0.412)
Multiple regression with wage variable	−0.233	−0.364	−0.010	−0.103	−0.240	−0.475
	(0.380)	(0.406)	(0.026)	(0.315)	(0.147)	(0.393)

All equations were estimated over the 60-industry sample. Instrumental variables for 2001–07 included percentage female, GDP growth and pension coverage. Instrumental variables for 2007–19 included percentage college graduates and percentage of employees with 20 or more years of tenure in 1980s.

variable in an OLS model of changes in the percentage of younger workers over time. This will tell us whether covariates have any impact on the relationship. Third, even though there are no *a priori* reasons to exclude some of the independent variables in the model for older workers from the model for younger workers, there are correlation patterns suggesting that they could be excluded because of their explanatory power for one age group and lack of explanatory power for the other. The control variables are the same as those used in the previous models; results with and without the relative wage of older to younger workers are reported.

The key finding across all the results reported in Table 4 is that there is no evidence of any declines in the share of younger workers being related to increases in the share of older workers. The coefficients for 2007–2019 are quite small, ranging between 0.1 and −0.1. The coefficients for 2001–2007 and 2001–2019 are larger, running from 0 to −0.5. The null hypothesis cannot be rejected in any of the models. The conclusion that can best be drawn from this exercise is that, despite the best efforts being made to find a relationship between the shares of older and younger workers across industries, there is no such relationship to be found in this data set.

5.2 Imports and older workers

China joined the World Trade Association in 2001 and the resulting reduction in trade barriers led to a surge of imports into the USA. In 2001, the USA imported $8.5 billion worth of goods from China, increasing to $26.8 billion by 2007. Chinese imports peaked at $44.9 billion in 2018 and subsequently dropped to $37.6 billion in 2019 after both China and the USA took retaliatory trade measures.

Studies such as Autor *et al.* (2013) have established that the surge of imports from China resulted in a decline in manufacturing jobs in the USA. The question examined here is how older workers were impacted relative to younger workers. The 2001–2007 sample period used in our data set coincides with the period of the most rapid increase in Chinese imports, so this is likely to be the period where the impact of expanded trade with China has the largest impact. The 2007–2019 sample contains both the Great Recession and the ramping up of tariffs and other trade barriers after the 2016 US Presidential Election, so the results for that period may be less clear-cut.

One challenge that arises when examining this question concerns how to handle the 37 industries in the data set outside of primary goods and manufacturing. Here two approaches are considered. The first is to restrict the sample to the 23 industries for which trade data are tabulated. This requires a reduction in the number of independent variables to allow enough degrees of freedom to obtain precise estimates of the import penetration variables.

The second is to assign values of zero to the import penetration variables for non-traded goods as well as services. This is a reasonable approach for personal services such as haircuts and spa treatments. On the surface this approach would seem less reasonable for services such as travel;

Table 5. Regression results: impact of Chinese and other imports

Dependent variable	Change in percentage of workers 55 and above					
Sample	23-industry sample			60-industry sample		
Years	2001–07	2007–19	2001–19	2001–07	2007–19	2001–19
Ratio of Chinese import growth to initial output	0.052** (0.016)	0.004 (0.068)	−0.001 (0.067)	0.040* (0.019)	0.168 (0.140)	0.114 (0.069)
Ratio of other import growth to initial output	0.037 (0.024)	−0.103** (0.030)	−0.075 (0.083)	−0.015 (0.016)	0.004 (0.028)	−0.062 (0.048)

Independent variables in 23-industry sample include GDP growth, log capital–labor ratio, log high-tech capital–labor ratio, pension coverage, and log wage ratio of older to younger workers. Independent variables in 60-industry sample are the same as in Tables 1–3.

American tourists in Shanghai are in effect importing Chinese services. However, there is no evidence of any surge in Chinese service 'imports' that is in any way comparable to what has happened in manufacturing. Further, the service industries with zero imports are not total outliers in the distribution. Manufacturing industries with little to no change in Chinese import penetration during this period include food products, petroleum and coal products, chemicals, and primary metals.

Separate measures for Chinese and all other imports were created from USA Trade Online to estimate the impact of increased import competition on the age structure of employees by industry. One variable is the change in Chinese imports in a given industry over the sample period divided by gross output in that industry in the initial year of the sample period.[10] The second variable is the change in imports from all other countries divided by initial gross output. Both variables are included to allow the impact of Chinese imports on employment to vary from that of the impact of the imports from other countries.

The empirical models have so far included as many as 12 independent variables, a strategy that is not likely to be successful when adding two more variables to a data set with 23 observations. The set of independent variables in Table 5 for the 23-industry sample is restricted to those that have proven to be of critical importance to the model: GDP growth, overall capital–labor ratio, high-tech capital–labor ratio, pension coverage, and the relative wage of older to younger workers. The full set of independent variables is used for the 60-industry sample.

The results for both samples show that the percentage of older workers increased more in those industries with the largest surges in Chinese imports in 2001–2007. Compare two industries one with a 20% change in import penetration and one with no change in import penetration. The growth in the percentage of older workers is 0.8–1.0 percentage points higher between 2001 and 2007 in the industry with the higher degree of new import competition. In the 23-industry sample there is no correlation between Chinese import penetration and the share of older workers between 2007 and 2019. The coefficient for Chinese import penetration for the 60-industry sample is actually more than twice as large in 2007–2019 but the standard error increases to an even greater degree, making it unwise to draw any conclusions for that sample in that time period.

Focusing further on the 2001–2007 period, separate equations estimating the change in the percentage of younger workers were estimated with the same independent variables for the two samples. In the 23-industry sample, the China imports coefficient (S.E.) for young workers was −0.038 (0.053), but it was −0.085 (0.031) in the 60-industry sample. This indicates the surge in Chinese imports led to employment reductions among younger workers.

As for import penetration from countries other than China, the results are mixed. In the 23-industry sample, the percentage of older workers grew more slowly in 2007–2019 in industries with the fastest growth in imports from other countries. However, there was no relationship in any

---

[10]The import measure is the value of goods imported for consumption including cost, insurance and freight and excluding duties.

period between import growth from these countries and the change in the share of older workers in the 60-industry sample.

The takeaway from this analysis is that Chinese import growth had a sizable impact on the age structure of industries between 2001 and 2007, mainly by reducing the employment of younger workers relative to older workers. This impact appears to have dissipated in 2007–2019, perhaps because the rate of import growth slowed down.

6. Conclusion

More older people are working than ever before. This study has demonstrated that there is considerable inter-industry variation in the percentage of older workers at a given point in time, as well as in the rate at which those percentages change over time. The empirical analysis has attempted to determine the characteristics of those industries where they have found the greatest opportunities.

Pension coverage is the single strongest predictor of whether an industry will have a growing percentage of older workers during the first two decades of this century. Now that DC plans dominate in the private sector, workers in industries with high pension coverage can make their organizational exit and retirement decisions without having to consider spikes in pension wealth at certain ages that take place under DB plans. Better health and higher longevity appear to be translating into longer careers for these workers.

Older workers now comprise a larger share of the employees in industries that make the most intensive use of high-tech capital. Even if they may not be proficient in Python or R, they are making contributions in managerial and professional roles. There also is evidence that older workers had less employment growth than other age groups in industries that make intensive use of traditional types of capital, namely equipment and structures.

Employment patterns by age react to changes in output and trade patterns. Older workers became a smaller percentage of the workforce in industries with the fastest growth of output, reflecting firm's decisions to hire younger or middle-aged workers with a longer potential training horizon. In the industries that were most impacted by Chinese imports in the 2001–2007 period, the employment of all age groups declined but older workers had smaller proportional job losses than the other two age groups.

In closing, we also should note some things we did not find. There is no evidence the percentage of older workers is rising more in industries where the wage gap between older and younger workers was greatest. In other words, there does not seem to be a demand shift toward older workers. Also, there is no evidence that an increase in the percentage of older workers in an industry leads to a decrease in the percentage of younger workers. Further work needs to be done to examine more detailed data on hiring patterns by age and how it relates to the growth of the employment of older workers by industry. It also will be helpful to study more carefully changes in employment levels by age group, rather than just focusing on employment shares.

The focus here has been to reveal basic patterns in the data. To establish causality more rigorously, Acemoglu and Restrepo (2020) and Autor *et al.* (2013) used values from other countries as instruments for technology and trade variables. Matching industry codes over multiple data sets over this period (pre- and post-NAICS industry codes) for the USA was no small challenge for this study. One option would be to examine the EU KLEMS data set. It provides capital variables for dozens of other countries but with a more highly aggregated set of industry definitions.

This study intentionally ended with annual data for 2019. A logical next step will be to examine what happened to the employment of older workers by industry month-by-month in the subsequent years. The aggregate data show that there was a sharp downward spike in labor force participation starting in April 2020. It remains to be seen whether many of those who left the labor force during the pandemic will return. Some important clues may very well be found in their inter-industry employment patterns.

References

Abraham KG and Kearney MS (2020) Explaining the decline in the US employment-to-population ratio: a review of the evidence. *Journal of Economic Literature* **58**(3), 585–643.

Acemoglu D and Restrepo P (2020) Robots and jobs: evidence from US labor markets. *Journal of Political Economy* **128**(6), 2188–2244.

Acemoglu D and Restrepo P (2022) Demographics and automation. *Review of Economic Studies* **89**(1), 1–44.

Aiyar S, Ebeke C and Shao X (2016) The impact of workforce aging on euro area productivity, IMF Country Report 16/220. Available at https://www.imf.org/external/pubs/ft/scr/2016/cr16220.pdf.

Allen S (2023) Demand for older workers: what do we know? What do we need to learn?. *Journal of the Economics of Ageing* **24**(2), 100414.

Ameriks J, Briggs J, Caplin A, Lee M, Shapiro MD and Tonetti C (2020) Older Americans would work longer if jobs were flexible. *American Economic Journal: Macroeconomics* **12**(1), 174–209.

Autor DH, Dorn D and Hanson GH (2013) The China syndrome: local labor market effects of import competition in the United States. *American Economic Review* **103**(6), 2121–2168.

Bartel A and Sicherman N (1993) Technological change and retirement decisions of older workers. *Journal of Labor Economics* **11**(1), 162–183.

Barth E, Davis J, Freeman R and McElheran K (2020) Twisting the demand curve: digitalization and the older workforce, NBER Working Paper No. 28094.

Bersin J and Chamorro-Premuzic T (2019) The case for hiring older workers, Harvard Business Review. Available at https://hbr.org/2019/09/the-case-for-hiring-older-workers.

Bianchi N, Bovini G, Li J, Paradisi M and Powell M (2021) Career spillovers in internal labor markets, NBER Working Paper 28605.

Carta F, D'Amuri F and von Wachter TM (2021) Workforce aging, pension reforms, and firm outcomes, NBER Working Paper No. 28407.

Clark R and Ghent L (2010) Strategic HR management with an aging workforce: using demographic models to determine optimal employment policies. *Population Research and Policy Review* **29**(1), 65–80.

Coile C (2019) Working longer in the United States: trends and explanations. In Coile CC, Milligan K and Wise DA (eds), *Social Security Programs and Retirement Around the World: Working Longer*. Chicago: University of Chicago Press, pp. 299–324.

Grund C and Westergård-Nielsen N (2008) Age structure of the workforce and firm performance. *International Journal of Manpower* **29**(5), 410–422.

Hamermesh D (1993) *Labor Demand*. Princeton, NJ: Princeton University Press.

Hudomiet P and Willis R (2021) Computerization, obsolescence, and the length of working life, NBER Working Paper No. 28701.

Maestas N, Mullen KJ and Powell D (2016) The effect of population aging on economic growth, the labor force and productivity, NBER Working Paper No. 22452.

Mohnen P (2021) The impact of the retirement slowdown on the U.S. youth labor market. Available at https://sites.google.com/view/paulmohnen/research?authuser=0.