Aging and the Productivity Puzzle

Adam Ozimek[1], Dante DeAntonio[2], and Mark Zandi[3]

[1]Senior Economist, Moody's Analytics
[2]Economist, Moody's Analytics
[3]Chief Economist, Moody's Analytics

September 4, 2018

Abstract

Research on the effect of an aging population on economic growth has tended to focus on labor force participation and the dependency ratio, and the limited work that has examined the effect of aging on productivity remains inconclusive. This analysis contributes to the aging and productivity literature and finds a strong negative effect of aging on productivity using both state-industry aggregate data and matched employer-employee microdata. The aggregate data show a clear relationship between an older workforce and lower productivity at the state-industry level, in both cross-section and panel models. The results are confirmed using employer-employee linked data from private sector human resource company ADP that show having older coworkers reduces an individual's wages. Robustness tests include three-digit ZIP code geographic fixed effects and firm fixed effects for a sample of large nationwide firms. The results suggest that between a quarter and full percentage point of the slowing in annual productivity growth in recent years may be due to aging, an effect that is likely to persist for the next decade.

I. Introduction

The workforce has aged significantly over the past decade. Two main factors are driving this. Most important is demographics. The baby boom generation, which includes anyone born from 1946 to 1964, is larger than both the generation that came before it and the one after it. As a result, as baby boomers age, so too does the working-age population. Most notably, starting about a decade ago baby boomers began becoming senior citizens, and as a result senior citizens now make up an increasingly large share of the adult population.

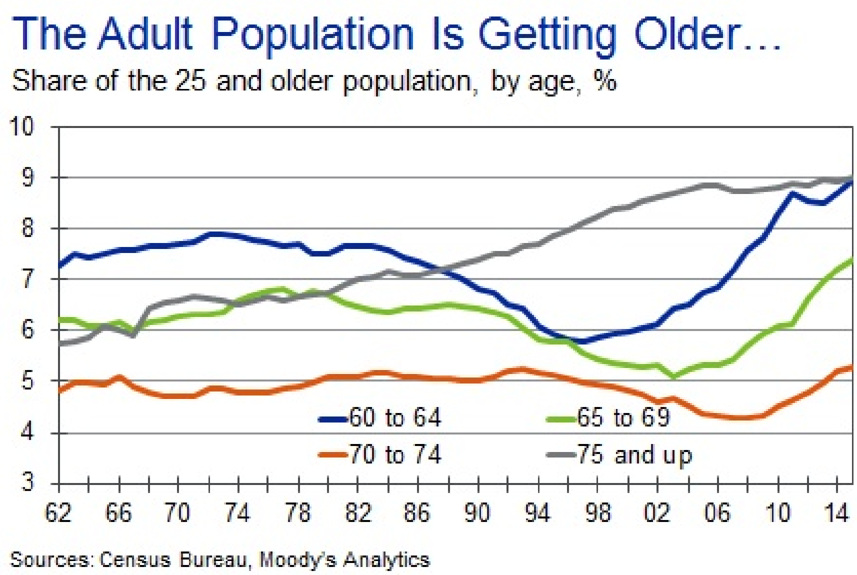

Also behind the aging workforce are increasing lifespans and the lack of retirement savings, which are keeping older people in the workforce at higher rates than in the past. This is true for younger seniors as well as older ones.

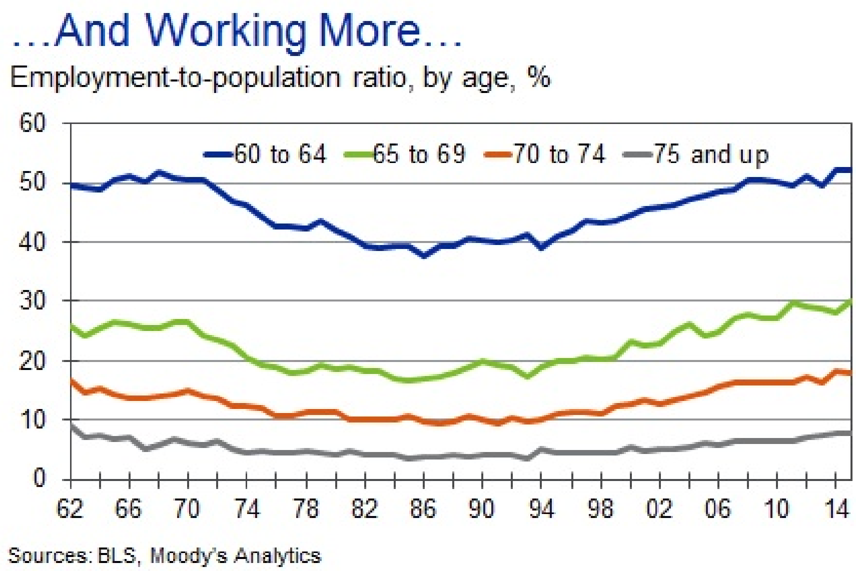

The combined effect of an older population and people working more later in life is that the senior share of the workforce has increased significantly, doubling from 3% to 6% in the last 15 years to the highest level on record.

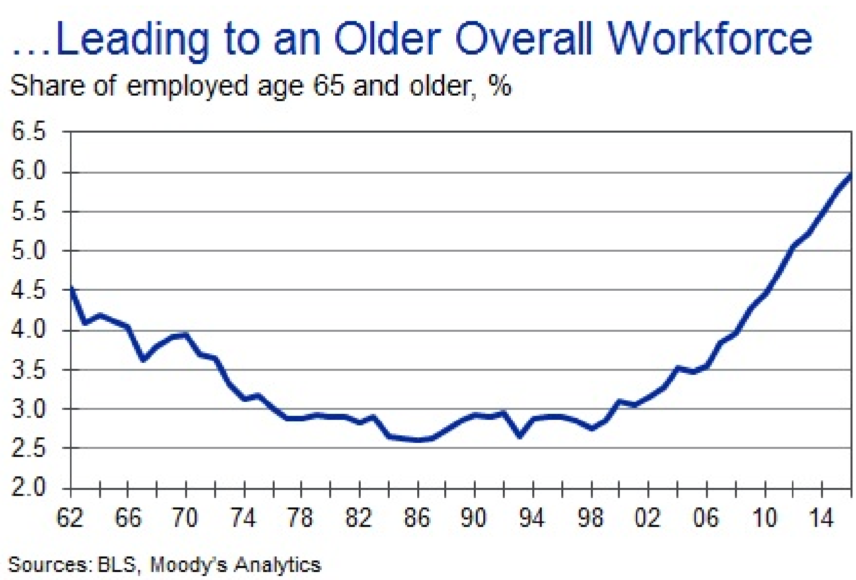

The graying of the workforce has occurred alongside declining productivity, raising questions about whether aging is partly responsible for some of the recent slowdown in productivity. However, aside from timing, the theoretical basis for this connection is not obvious. Past work has been vague or agnostic about the mechanism through which an aging workforce might affect productivity.[1] A direct effect of aging on a worker's own productivity is not supported by lifecycle wage data, where the lack of a sharp drop off in pay for older workers suggests marginal productivity does not drop off either. As a result, it is more plausible that any observed relationship between aging and productivity is due to negative productivity spillovers rather than composition effects of an aging workforce.

One plausible mechanism for productivity spillovers is that an older workforce may slow the adoption of productivity-enhancing technologies that require learning on the part of workers. The benefits from such learning will likely be lowest for older workers, due to shorter remaining careers that provide less time to recoup the upfront costs. The costs may also be higher if older workers have a harder time learning new skills. If firms reap the benefits and pay the costs of the training, it may therefore be profit-maximizing to forgo adopting new technologies. Alternatively, if workers bear some of the costs or benefits, they may block technology adoption via x-inefficiency.

The effect of an aging workforce on productivity is theoretically ambiguous, and empirical results have been mixed as well. Maestas et al. (2016) look across U.S. states and find that an increase in the share of the population age 60 and older slows productivity growth significantly. Feyrer (2007) finds that the age distribution of the workforce affects growth in cross-country panel models, but his results suggest positive effects of the 40 to 50 age group rather than negative effects of older workers. Acemoglu and Restrepo (2017) look at more recent cross-country data and find that countries with more aging have higher productivity.

[1] Maestas et al. (2016) provide a model of economic output where human capital productivity of the labor force depends on the age distribution of the workforce, however they provide little intuition or detail about why this might be. Feyrer (2007) discusses the possibility of human capital spillovers due to experience, however "is agnostic as to the mechanisms through which demographic change and productivity are related."

They explain that a scarcity of younger and middle age workers drives adoption of new automation technologies that increase productivity. Older workers may also generate firm-specific human capital that is not easily replaced by the firm. Using data on 34,000 deaths at German firms, Jäger (2016) shows that, consistent with economist Gary Becker's notion of firm-specific human capital, the deaths of high-skilled or managerial workers have negative effects on coworker wages.

The theoretical ambiguity and mixed empirical results suggest more research is needed. This analysis contributes to the literature by first showing a clear relationship between aggregate aging and productivity in cross-sections of state-industry data for the United States. To do this, annual Quarterly Workforce Indicators data on the age of the workforce for state-industry pairs is matched to Bureau of Economic Analysis data on productivity. After removing industry and state fixed effects, which control for U.S.-level changes in industry productivity and state-level changes such as a growing dependency ratio, there is a clear relationship between an older workforce and lower productivity. The results are confirmed in state-industry panel regressions with data from 2000 to 2015, resulting in an elasticity of productivity to the share of workers 65 years and older of between 2.3 and 9.2. The aggregate data are suggestive, but remain open to reverse causality and a variety of omitted variables.

The results are then confirmed in a more rigorous analysis using matched employer-employee data from ADP, a private sector human resource company that processes payrolls for millions of workers. These data measure within-firm spillover effects directly, ruling out composition effects alone. Separate regressions for each age group confirm that having older workers causes negative spillovers to workers of all ages.

In addition, the ADP data allow for local geographic controls that rule out omitted variables at the three-digit ZIP code level. As a further robustness test, a sample of large national firms with employees located throughout the country is examined, which allows the inclusion of fixed effects to rule out firm-level omitted variables. The consistent results across models, datasets and robustness tests suggest that older workers are causing lower productivity. Simulations using these coefficients suggest that the aging of the workforce can explain between 0.3 and 0.7 percentage points of the slowdown in per annum productivity growth over the last 15 years. Forecasts of population aging suggest that these effects will continue for more than a decade.

II. State-Industry Models

Given the theoretical indeterminacy of the relationship between aging and productivity, it is worth examining whether aging and low productivity appear correlated in aggregate at the state-industry level. The age distribution of workers by industry and state is available from the Census Bureau's Quarterly Workforce Indicators (QWI), which provide detailed data on matched worker and firm characteristics.[2] These data reveal that the aging of the workforce is a broad-based phenomenon and has not affected just some industries and states.

From 2000 to 2015, every industry and state has seen an increase in the share of workers age 65 and older. Out of 755 state-industry pairs for which data are available, all but two saw an increase in the senior share of the workforce from 2000 to 2015. The more complete sample available from 2005 to 2015 shows that 872 out of 880 state-industry pairs saw an increase.[3]

[2]The QWI comes from the Census Longitudinal Employer-Household Dynamics data, which are microdata that link employers and employees and cover 95% of private sector employment. The data are available for most states starting in the early 2000s, with the first quarter of data availability varying by state. Industries are categorized by NAICS code.

[3]Of the eight observations with a decrease in the senior share, five are in mining and three are in North

That almost every industry and state are aging suggests that neither industry-specific nor state-specific factors are driving the entirety of the phenomenon. If only northeastern states or only services industries were aging, it would make cross-sectional or panel analysis more difficult.

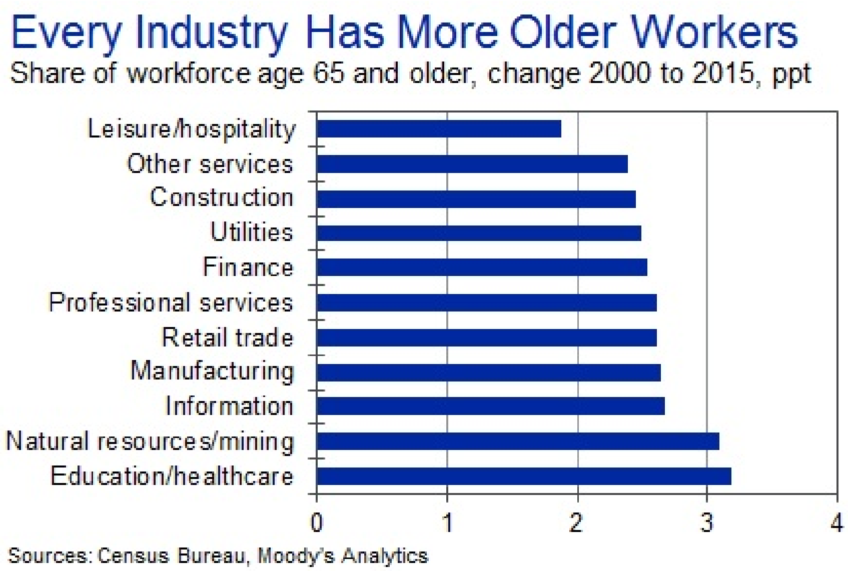

Annual state-industry QWI data are matched to Bureau of Economic Analysis data on state-level productivity by industry. Overall, the matched data cover 50 states and Washington DC and 18 industry aggregations. States began participating in the QWI at different times, creating a staggered start period with 83% of the 918 state-industry observations[4] available with QWI data in 2000, and 98% or higher for 2005 through 2015.[5]

Dakota.

[4]918 possible observations are equal to 18 industries x 50 states and DC.

[5]Complete QWI data for Wyoming are available only through 2014, leaving 898 observations in 2015. Productivity and QWI data for natural resources and mining in Washington DC end in 2005, and QWI data for Massachusetts are available only beginning in 2010.

1. The Cross-Sectional Relationship

The relationship between aging and productivity is clearly visible in simple cross-section comparisons that ignore changes over time and more complicated panel econometrics. First, both productivity and the share of workers age 65 for every state-industry pair in the year 2015 were regressed on state and industry fixed effects. For outcome variable Y in industry i, state j, time t, the following residual is estimated using the constant, α, and industry and state fixed effects, α_i and α_j:

$$\epsilon_{ijt} = Y_{ijt} - \alpha - \alpha_i - \alpha_j \tag{1}$$

This removes the variation in productivity that is due to effects common to all industries or to all states. The resulting residuals have a clear relationship when plotted against each other, with a higher senior share of workers associated with lower productivity at a correlation of -0.29. The results are similar for 2010 and 2005, with correlation coefficients of -0.28 in these years.

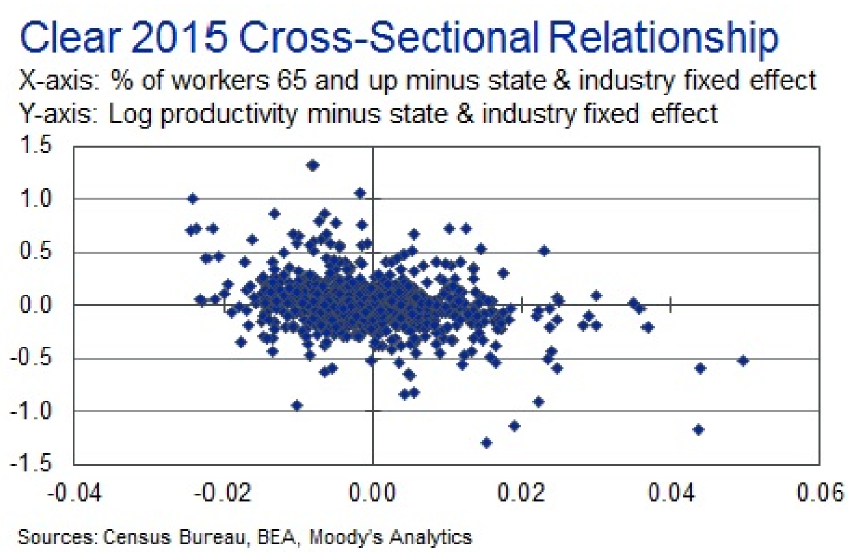

Cross-sectional regressions for 2000, 2005, 2010 and 2015 confirm that the relationship between aging and the workforce is statistically significant. Despite the sample changing each year due to the staggered start period for each state, the coefficient is consistent across cross-sections, ranging from -6.56 in 2000 to -9.21 in 2010.

A ten-year first difference model from 2005 to 2015 shows that changes in productivity are statistically significantly related to changes in the senior share. This result is robust to the inclusion of state-industry pair job growth. While the coefficient is smaller than the cross-section, it is of a similar magnitude. These results suggest a one percentage point increase in the senior share of the workforce reduces productivity by 3.84%. Finally, a first difference from 2000 to 2015 with industry-state specific fixed effects shows a statistically significant -2.35 coefficient.

Table 1: State-Industry Models

VARIABLES	Model 1 ln(prod.)	Model 2 ln(prod.)	Model 3 ln(prod.)	Model 4 ln(prod.)	Model 5 ln(prod.)	Model 6 ln(prod.)	Model 7 ln(prod.)
Pct. workforce 65+	-6.56	-8.53	-9.21	-7.82	-4.78	-3.82	-2.35
p-value	0.00	0.00	0.00	0.00	0.00	0.01	0.00
log(employment)						0.079	0.066
p-value						0.179	0.015
Constant	4.55	5.19	5.33	5.18	-0.055	-0.090	3.40
p-value	0.00	0.00	0.00	0.00	0.406	0.188	0.00
First difference?					X	X	X
Emp. controls?						X	X
Years	2000	2005	2010	2015	2005-2015	2005-2015	2000-2015
Observations	755	896	915	897	877	877	877
R-squared	0.916	0.911	0.902	0.905	0.366	0.369	0.594

Sources: Census Bureau, BEA, Moody's Analytics
Notes: Standard errors clustered at the state level. All models include industry and state fixed effects.

Because the senior share of the workforce is still generally a minority share, the large effects on productivity suggest an impact that goes beyond merely the composition effect of low productivity of older workers on average productivity. Even if the productivity of age 65-and-older workers was zero, then going from 3% to 6% of the workforce would decrease productivity by only 3.1% and not the 4% implied by Model 5 in Table 1.[6] Overall, it is difficult to get large productivity changes from composition effects alone given the small share of the workforce that seniors represent.

While composition effects are unlikely to explain the entirety of the effect, they could explain a portion. In addition, while the QWI state-industry analysis is suggestive, omitted variable bias remains a possible concern. Finally, while a range of plausible models suggest an association between state-industry productivity and workforce aging, it is possible to construct panel models where the effect of aging on productivity is statistically insignificant. For example, five-year changes with job growth controls result in insignificant effects. As a result, while the QWI analysis is suggestive, more robust microdata analysis is useful to more plausibly establish causality.

III. Worker-Level Analysis

Individual level matched employer-employee data can be used to ensure that compositional effects are not driving the results by focusing on the effects of having older coworkers on an individual worker's wage, controlling for that worker's own age. The detailed micro-level

[6] Algebraically, letting the average non-senior wage = W, then the change in the average from going from 3% to 6% with a zero productivity of senior workers is: (6% x 0 + 94% x W) - (3% x 0 + 97% x W) / (3% x 0 + 97% x W) = (94% - 97%) / 97% = -3.1%

data also allow for a variety of robustness tests. The data used in the analysis come from ADP payroll records. This includes quarterly information on all of the workers at a firm that uses ADP's payroll processing services, including age, gender and the ZIP code of a worker's residence. The data also contain firm-level industry classifications and allow for the calculation of the share of workers at a firm in a specific age group and total firm job growth over time.

Mincer regressions are used to test the effect of firm workforce characteristics on individual worker pay while controlling for individual worker level characteristics. The basic cross-sectional regression model is defined as:

$$ln(Y_{i,j,t}) = \alpha + \beta_1 X_{i,t} + \beta_2 \pi_{j,t} + \beta_3 \theta_{j,t} + \mu_{j,t} + \epsilon_{i,j,t} \quad (2)$$

where Y is the hourly wage for individual i at firm j in period t. The vector $X_{i,t}$ includes worker-specific controls, and the vector $\pi_{j,t}$ includes firm-specific controls. The variable $\theta_{j,t}$ measures the share of workers at firm j who are age 65 and older.

The basic model indicates that working at a firm with a larger share of older workers is associated with a lower wage. Model 1 in Table 2 shows the results of the model using data from the fourth quarter of 2016, but the results are robust to the time period of consideration. The results show that an increased share of firm employment age 65 and older has a statistically significant negative effect on employee wage levels. A one percentage point increase in the share of workers age 65 and older reduces wages by around 1%.

In addition, to control for persistent common factors at the firm and individual levels, the model can also compare changes in wages over a period of time to firm-level changes in

Table 2: Worker-Level Analysis - Basic Model

VARIABLES	Model 1 ln(wage)	Model 2 dln(wage)
65+ share of firm employment	-1.008***	
	(-0.0365)	
Change in 65+ share of firm employment		-0.0857***
		(-0.0132)
Observations	914,627	914,627
R-squared	0.374	0.09

Sources: ADP, Moody's Analytics
Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All models include controls by 3-digit ZIP code, firm industry classification, as well as worker characteristics. Standard errors clustered at the 3-digit ZIP code level in parentheses. Sample includes all firms with 20-500 employees.

the share of workers age 65 and older at the firm. The results are given in Model 2 in Table 2 and are consistent with the results from the wage-level analysis. The results shown are for a three-year time period from the fourth quarter of 2013 to the fourth quarter of 2016, but are robust across one-year, two-year, and three-year changes. The results again show a statistically significant negative effect on wage growth as a result of an increase in the share of workers at a firm age 65 and older. A one percentage point increase in the share of workers age 65 and older reduces wage growth by 0.0857 percentage point.

1. Albatross and Wise Man Theories

The results of these models are consistent with two opposing theories: 1) The albatross theory holds that older workers have negative productivity spillover effects, and so having more of them brings down productivity; or 2) The wise man theory holds that older workers have positive spillover effects, and firms with additional older workers have more retiring workers, and more retirements are what brings down productivity. To determine which theory is appropriate it is useful to break the stock of older workers in period t into flows from period $t-1$ to period t:

$$\theta_{j,t} = \theta_{j,t-1} + H_{j,t} + A_{j,t} - E_{j,t} \tag{3}$$

where $\theta_{j,t}$ is the number of older workers in period t at firm j, $H_{j,t}$ is the number of older workers hired in period t, $A_{j,t}$ is the number of workers who aged into the older working group and were employed at firm j in period t and $t-1$, and $E_{j,t}$ is the number of older workers who were employed at firm j in period $t-1$ but exited in period t. This last term would include retirements.

Instead of including $\theta_{j,t}$ in the wage growth regression, the individual components can be included, all converted to shares by dividing by the average firm employment in periods t and $t-1$. If older workers have negative spillovers, then $\theta_{j,t-1}$ and $A_{j,t}$ will be negative, and $E_{j,t}$ will be positive. A negative sign on $H_{j,t}$ would also be consistent with negative spillovers. However, given that firms are unlikely to hire new workers who have negative productivity spillovers, rather than simply fail to fire them as in $\theta_{j,t-1}$ and $A_{j,t}$, this suggests an indeterminate sign for $H_{j,t}$. If older workers have positive spillovers, then $\theta_{j,t-1}$ and $A_{j,t}$ will be positive and $E_{j,t}$ will be negative.

The results are consistent with negative spillovers from older workers (see Table 3, Model 3). The coefficients on $\theta_{j,t-1}$ and $A_{j,t}$ are negative and statistically significant, and $E_{j,t}$ is positive. Wage growth is reduced as more of the workforce at a given firm ages into seniority and wage growth picks up when older workers exit the firm. The coefficient on $H_{j,t}$ is positive, suggesting that firms on average hire older workers who do not have negative productivity spillovers, suggesting new hires are systematically different from the average older worker, although the effect is not statistically significant in this model.

Table 3: Worker-Level Analysis - Flow Model

VARIABLES	Model 3 dln(wage)
65+ share in 2013m12 ($\theta_{j,t-1}$)	-0.0674***
	(-0.0111)
65+ hires from 2013m12 to 2016m12 ($H_{j,t}$)	0.0242
	(-0.019)
workers aged into 65+ share ($A_{j,t}$)	-0.0856***
	(-0.0166)
65+ exits from firm over 3-yr period ($E_{j,t}$)	0.0287***
	(-0.0111)
Constant	0.277***
	(-0.00414)
Observations	914,627
R-squared	0.09

Sources: ADP, Moody's Analytics
Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All models include controls by 3-digit ZIP code, firm industry classification, as well as worker characteristics. Standard errors clustered at the 3-digit ZIP code level in parentheses.

IV. Blocked Technology as a Mechanism

To test the mechanism that older workers tend to resist productivity-improving technologies, it is useful to investigate how the effect of an aging workforce varies by education level. Given the evidence for skill-biased technological change, it is plausible that firms that face more options for productivity-improving technologies would be those that have a more skilled workforce (Goldin and Katz, 2007). As a result, an aging workforce that blocks new technologies would be expected to have a greater impact in more skill-dependent industries.

What's more, high skilled workers should also face a greater productivity loss from blocked technologies if those technologies tend to be skill-biased, on average.

These theories can be tested using 2016 5-year data from the American Community Survey. We estimate a similar Mincer equation as above for all full-time, year-round worker employed in private industry. The models include controls for demographic characteristics such as race, ethnicity, education, and nativity as well as fixed effects for 260 metros, 251 industries, 474 occupations, and single-year age. The effect of older workers on own-wage is captured by measuring the share of workers in an industry and metro area that are 65 and up. For example, we measure the share of workers 65 and up in the truck transportation industry in Colorado Springs, CO. The advantage of ACS data is that the models are able to control for skill levels of individuals as well as industries using information on education level.

The baseline model supports the results from the ADP data. That is, each percentage point increase in the senior share of workers in ones' industry and metro area is associated with a 0.19 percentage point reduction in pay (see Table 4, Model 1). When this effect is interacted with education, it shows that the effects are greater for more educated workers. The effect on workers with less than a high school degree is negative but not statistically significant, and the coefficient rises to a statistically significant 0.59 for workers with a graduate degree or more (see Table 4, Model 2).

The results also show that the effect is stronger in industries where the overall education level of the workforce is higher. To show this, the share of workers with a college degree or higher is estimated for all industries at the U.S. level. Industries are then grouped into quintiles of education level. In the lowest quintile, 13% or less of the workforce has a college degree. In the highest quintile, 51% or more have a college degree. Interacting dummies for the five quintiles of average education level shows that workers in higher skilled industries are hurt the most by an older workforce.

Overall, these results confirm the effect of an older workforce on wages. They also suggest that the effects are strongest for more skilled workers and those in industries where the overall skill level is higher. The greater impact for skilled workers and industries is consistent with blocked technological progress playing an important role in the link between an older workforce and lower productivity.

V. Robustness Tests

While first differences and firm controls help reduce omitted variable bias, it remains a possibility that firms with increasing shares of older workers experience some other negative shock over the same period that reduces wage growth but is imperfectly measured by total firm-level job growth. An additional robustness test can be performed by focusing on variation in geographic labor markets within large national firms. While the ADP data do not contain information on establishments, large firms can be divided into local labor markets by using the three-digit ZIP codes of the workers' home addresses. There are approximately 900 three-digit ZIP codes in the U.S., or about one-third the number of counties.[7] The assumption is that workers at the same firm but in different three-digit ZIP codes work at different establishments within the same firm.

This analysis repeats Model 1 above, but measures the age 65 plus share of the workforce

[7]According to the 2010 Census, there were 3,143 county or county equivalents in the United States.

Table 4: Worker-Level Analysis Using ACS Data

VARIABLES	Model 1 ln(wage)	Model 2 ln(wage)	Model 3 ln(wage)
Share 65 and up	-0.186***		
	(0.00)		
Share 65 and up x			
< High school		-0.023	
		(0.81)	
High school		-0.107***	
		(0.00)	
Some college		-0.146***	
		(0.00)	
College		-0.254***	
		(0.00)	
Graduate school and up		-0.590***	
		(0.00)	
Share 65 and up x			
Industry education quintile 1			-0.188**
			(0.02)
Industry education quintile 2			-0.148***
			(0.00)
Industry education quintile 3			-0.141***
			(0.00)
Industry education quintile 4			-0.241***
			(0.00)
Industry education quintile 5			-0.302***
			(0.00)
Observations	65,696,792	65,696,792	65,696,792
R-squared	0.507	0.507	0.507

Sources: ACS, Moody's Analytics
Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All models include fixed effects for metro area, industry, occupation, and individuals controls for eduction, race, age, ethnicity, and nativity. P-values in parentheses.

at firm j in a particular three-digit ZIP code. Importantly, this allows for the inclusion of firm-specific fixed effects and focuses only on within-firm wage and workforce variation. This ensures that firm-specific shocks or omitted variables are not driving the results. The sample is limited to firms with workers who live in more than 100 three-digit ZIP codes. The large-firm only model is robust to the addition of firm-specific fixed effects (see Table 5, Model 5), although it does slightly reduce the impact that older workers have on wages.[8]

Table 5: Worker-Level Analysis - Large Firm Test

VARIABLES	Model 4 ln(wage)	Model 5 ln(wage)
65+ share of firm employment	-1.489***	-0.626***
	(-0.0999)	(-0.0936)
Firm-level fixed effects?	No	Yes
Observations	1,412,164	1,412,164
R-squared	0.601	0.712

Sources: ADP, Moody's Analytics
Notes: *** p < 0.01, ** p < 0.05, * p < 0.1. All models include controls by 3-digit ZIP code, as well as worker characteristics. Standard errors clustered at the 3-digit ZIP code level in parentheses. Sample includes firms with workers in 100 or more 3-digit ZIP codes.

Finally, the models are run separately for workers in 10-year age groups (see Table 6). These results reveal that having more senior coworkers has a negative effect on wages and wage growth across all age groups. The effect on wage levels is smallest for young workers, while the effect on wage growth is largest for young workers. This is consistent with younger workers having the lowest base pay and fastest wage growth, and vice versa for older workers. Importantly, these results suggest that having more senior workers at a firm reduces wages

[8]Using the sample of larger firms but not including firm-specific fixed effects (see Table 5, Model 4) yields results that are similar in magnitude to Model 1, which used a sample of smaller firms. Therefore, firm size does not appear to have a significant impact on the results.

(or wage growth), even for older workers themselves.

Table 6: Worker-Level Analysis - Age Cohorts

VARIABLES	Ages <30	Age 30-39	Age 40-49	Age 50-59	Age 60-69
Dep. variable: ln(wage)					
65+ share of firm emp.	-0.419***	-0.764***	-1.021***	-1.077***	-1.018***
	(-0.0426)	(-0.0271)	(-0.0263)	(-0.0245)	(-0.0319)
Observations	53,625	197,210	238,468	262,815	142,059
R-squared	0.5	0.391	0.37	0.353	0.319
Dep. variable: dln(wage)					
Δ in 65+ share of firm emp.	-0.213***	-0.170***	-0.0840***	-0.0698***	-0.00657
	(-0.0332)	(-0.0192)	(-0.0141)	(-0.013)	(-0.0156)
Observations	53,625	197,210	238,468	262,815	142,059
R-squared	0.072	0.043	0.02	0.018	0.021

Sources: ADP, Moody's Analytics
Notes: *** p $<$ 0.01, ** p $<$ 0.05, * p $<$ 0.1. All models include controls by 3-digit ZIP code, firm industry classification, as well as worker characteristics. Standard errors clustered at the 3-digit ZIP code level in parentheses. Sample includes all firms with 20-500 employees.

VI. Wages and Productivity

It is possible that the measured effect of aging on wages does not reflect changes in productivity, and instead could reflect changes in rent-sharing within the firm. One plausible theory consistent with a rent-sharing effect is that older workers may prevent younger workers from receiving promotions, as they occupy top management spots and concomitant higher pay.

However, the results show that having older coworkers reduces the wages of older workers as well (see Table 6), who are least likely to be crowded out of management by their own cohort. Alternatively, older workers' higher health costs may be spread to all workers and therefore reduce wages. If this were the case, the negative effect of having coworkers who are 70 and older, rather than 65 and older, would be diminished as they enroll in Medicare, but instead preliminary evidence shows that the effects grow larger with cohort age.

We cannot rule out the theory that aging changes how rents are shared within a firm without affecting productivity. However, the QWI results suggest aging does have a direct effect on productivity and not only wages.

VII. Productivity Growth Forecast

These results suggest that the slowing in U.S. productivity growth over the past 15 years and the aging of the workforce are more than coincidental, and the effects of aging on productivity growth are significant. Based on the state-industry and worker-level models, the elasticity of productivity growth with respect to the share of the workforce over 65 years old, ranges from approximately 1% to 3%. Given this, the aging of the workforce has reduced productivity by between 3% and 9%, equal to between 0.25% and 0.7% per annum.[9]

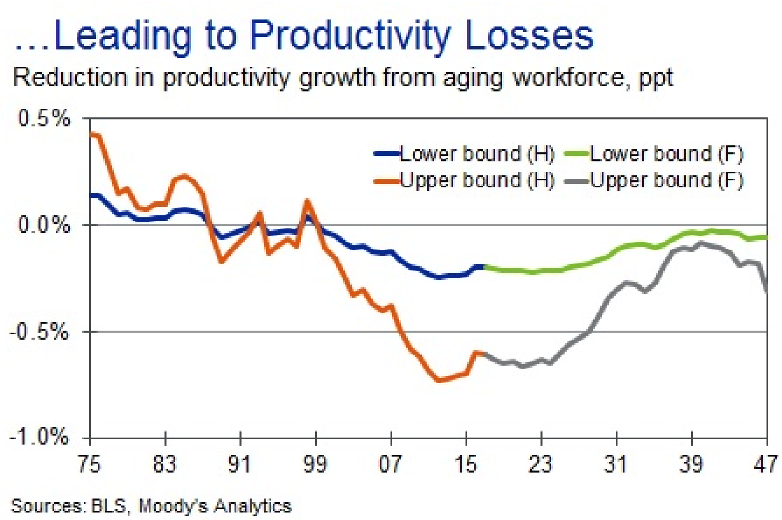

[9]The future share of the workforce that is 65 and older is estimated by holding constant the current employment-to-population ratio for that age group and applying it to Moody's Analytics forecast of population growth for that cohort.

For context, nonfarm business productivity growth during the current nearly 10-year long economic expansion has been close to 1% per annum. This is a full percentage point below the 2% per annum growth in productivity growth experienced in the post-World War II period up until this expansion. Our results suggest, that between one-fourth and almost three-fourths of the productivity slowdown in this expansion is due to the aging workforce.

Even more important, our results suggest that productivity growth will continue to be significantly constrained in the coming more than a decade, as the share of the workforce that is 65 and older will continue to increase at a rate similar to that in the past decade.

VIII. Conclusion

The aging of the population has enormous implications for performance of the U.S. economy. Most obvious is the impact it has had on the growth in the labor force and thus the economy's growth potential. Less obvious, but based on our work, equally as important is the impact aging is having on productivity and potential growth.

The link between aging and productivity is clear from our state-industry panel model which uses data that cover close to the full universe of private firms and workers in the U.S. The effects are too large to be plausibly driven by worker composition effects. While these aggregate results can best be interpreted as suggestive of an association, our use of private sector administrative payroll data on millions of workers allows more precise measures of the effects of firm-level workforce aging on individual wages.

Our work can offer no definitive conclusions as to the mechanisms causing aging to weigh on productivity, but a plausible theory for which we have shown suggestive evidence is that older workers may resist productivity-improving technologies. Further efforts to distinguish possible mechanisms for why aging appears to affect productivity remains an open task for future research.

Understanding why older workers are reducing productivity growth is essential if ways to mitigate the effects are to be found. The youngest baby boomers are only 53 years old, and the aging of the workforce is far from over. This bodes poorly for the outlook for productivity growth. Policymakers should not only focus on steps to incent older workers to remain in the workforce longer, but also ways to mitigate their impacts on productivity.

References

Acemoglu, D. and Restrepo, P. (2017). Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation. *National Bureau of Economic Research*, No. w23077.

Feyrer, J. (2007). Demoographics and Productivity. *The Review of Economics and Statistics*, 89(1):100–109.

Goldin, C. and Katz, L. F. (2007). The Race Between Education and Technology: The Evolution of US Educational Wage Differentials, 1890 to 2005. *National Bureau of Economic Research*, No. w12984.

Jäger, S. (2016). How Substituable are Workers? Evidence from Worker Deaths. *Working Paper*.

Maestas, N., Mullen, K. J., and Powell, D. (2016). The Effect of Population Aging on Economic Growth, the Labor Force and Productivity. *National Bureau of Economic Research*, No. w22452.