

NBER WORKING PAPER SERIES

SECULAR STAGNATION? THE EFFECT OF AGING ON ECONOMIC GROWTH
IN THE AGE OF AUTOMATION

Daron Acemoglu
Pascual Restrepo

Working Paper 23077
<http://www.nber.org/papers/w23077>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2017

We are grateful to David Autor, Ricardo Caballero, Emmanuel Farhi, Alp Simsek, and participants in the 2017 AEA Annual Conference for comments and suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Daron Acemoglu and Pascual Restrepo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation
Daron Acemoglu and Pascual Restrepo
NBER Working Paper No. 23077
January 2017
JEL No. E30,J11,J24,O33,O47,O57

ABSTRACT

Several recent theories emphasize the negative effects of an aging population on economic growth, either because of the lower labor force participation and productivity of older workers or because aging will create an excess of savings over desired investment, leading to secular stagnation. We show that there is no such negative relationship in the data. If anything, countries experiencing more rapid aging have grown more in recent decades. We suggest that this counterintuitive finding might reflect the more rapid adoption of automation technologies in countries undergoing more pronounced demographic changes, and provide evidence and theoretical underpinnings for this argument.

Daron Acemoglu
Department of Economics, E52-446
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and CIFAR
and also NBER
daron@mit.edu

Pascual Restrepo
Department of Economics
Boston University
270 Bay State Rd
Boston, MA 02215
and Cowles Foundation, Yale
pascual.restrepo@yale.edu

The rapid aging of the population of both developed economies and much of the rest of the world, depicted in Figure 1, is seen as one of the most dangerous economic ills of the next several decades. An increasingly popular thesis, building on Alvin Hansen’s famous 1938 presidential address to the AEA, views developed economies as being afflicted by “secular stagnation,” partly because an aging population creates an excess of savings relative to investments (Alvin Hansen, 1939, Lawrence Summers, 2013, and the essays in Coen Teulings and Richard Baldwin, 2014). A different but related challenge is emphasized by Robert Gordon (2016), who identifies demographic change as the first “headwind” slowing down economic growth in the developed world, for an older population will reduce labor force participation and productivity (workers’ earnings, and presumably productivity, peak in their 40s, e.g., Kevin M. Murphy and Finis Welch, 1990).

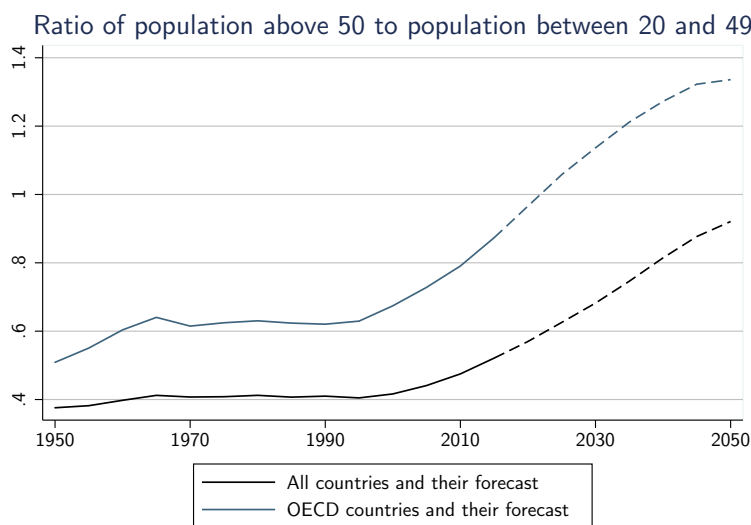


Figure 1: Aging from 1950 to 2015 and projections until 2050 (from the UN data). Aging is measured by the ratio of the population above 50 years old to the population between 20 and 49.

Though both perspectives imply that countries undergoing faster aging should be suffering more from these economic problems,¹ we show that since the early 1990s or 2000s, the periods commonly viewed as the beginning of the adverse effects of aging in much of the advanced world, there is no negative association between aging and lower GDP per capita.² Figure 2 provides a glimpse of the

¹Three qualifications to this conclusion should be noted. First, there are non-demographic factors, such as increased levels of inequality and slower technological progress, that have also been suggested as potential causes of demand-side secular stagnation. Second, this type of secular stagnation could be partially offset by monetary policy. Third, with international capital flows, aging in one country might affect GDP per capita in others as well.

²Thomas Lindh and Bo Malmberg (1999) and James Feyrer (2007) investigate the relationship between demographics and aggregate productivity or growth, focusing on pre-1990 data. Both papers find some evidence supporting the notion that the fraction of the population above 50 contributes negatively to GDP per capita. Their findings motivate our baseline choice of demographic variable as the ratio of the population above 50 to those between the ages of 20 and 49.

relevant pattern by depicting the raw correlation between the change in GDP per capita between 1990 and 2015 and the change in the ratio of the population above 50 to the population between the ages of 20 and 49. In the next section, we show that even when we control for initial GDP per capita, initial demographic composition and differential trends by region, there is no evidence of a negative relationship between aging and GDP per capita; on the contrary, the relationship is significantly positive in many specifications.

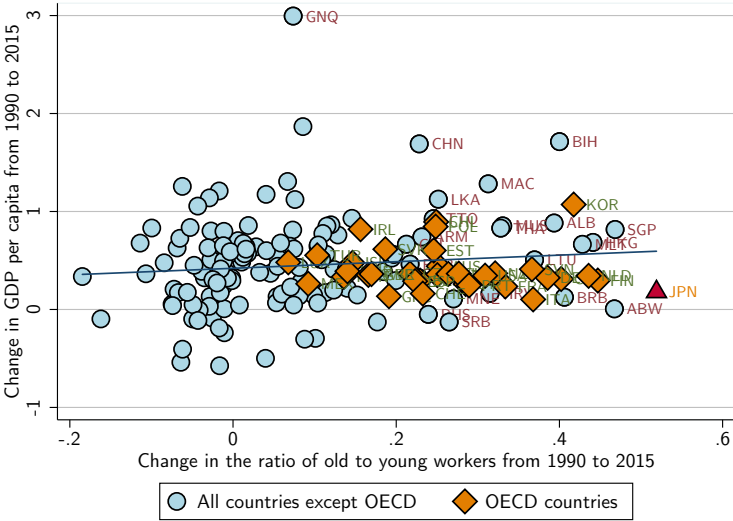


Figure 2: Correlation between aging and growth in GDP per capita (in constant dollars). Aging is defined as the change in the ratio of the population above 50 years old to the population between 20 and 49. GDP per capita data from the Penn World Tables.

The lack of a strong negative association between changes in age structure and changes in GDP per capita is surprising. So what explains it?

The post-1990 era coincides with the arrival of a range of labor-replacing technologies, most recently robotics and artificial intelligence, which provide a wide variety of options for firms to automate the production process. In Section II, we show that countries undergoing more rapid demographic change are more likely to adopt robots (see also Acemoglu and Restrepo, 2017).³ In Section III, we show that when capital is sufficiently abundant, a shortage of younger and middle-aged workers can trigger so much more adoption of new automation technologies that the negative effects of labor scarcity could be completely neutralized or even reversed.

³The recent working paper by Nicole Maestas, Kathleen J. Mullen and David Powell (2016) shows a negative association between aging and economic growth across US states. To the extent that US states have more similar technologies and more coordinated adoption decisions than countries, the countervailing effects of technology adoption we emphasize would be absent or much muted in this sample.

I. Aging and GDP Per Capita: The Cross-Country Evidence

In this section, we start by showing that the relationship depicted in Figure 2 is robust. We use data on GDP per capita from the Penn World Tables and population by age from the United Nations. Our main results are shown in Table 1, which presents regressions of the change in (log) GDP per capita from 1990 to 2015 on our baseline measure of population aging, the change in the ratio of the population above 50 to those between the ages of 20 and 49. Our baseline sample includes 169 countries for which we have data. Panel A reports OLS regressions in changes (long differences) with robust standard errors. Column 1 shows the raw correlation, already depicted in Figure 2. We see a positive but insignificant relationship. The rest of the table investigates the robustness of this relationship.

Table 1: Estimates of the impact of aging on GDP per capita from 1990 to 2015 and from 2000 to 2015.

	SAMPLE OF ALL COUNTRIES					OECD COUNTRIES		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A. aging measured by the ratio of the population above 50 to population between 20 and 49.</i>								
Change in ratio of old to young workers (from 1990 to 2015)	0.335 (0.210)	1.036*** (0.257)	1.162*** (0.276)	0.773** (0.322)	1.703*** (0.411)	-0.262 (0.352)	0.042 (0.346)	1.186*** (0.458)
Initial GDP per worker in 1990		-0.153*** (0.039)	-0.138*** (0.042)	-0.156*** (0.046)	-0.190*** (0.045)		-0.205*** (0.072)	-0.260*** (0.092)
First-stage F statistic					19.36			7.38
Overidentification test p-value					0.51			0.44
Observations	169	169	169	169	169	35	35	35
<i>Panel B. aging measured by the average age of the population above 20.</i>								
Change in average age (from 1990 to 2015)	0.042*** (0.015)	0.090*** (0.017)	0.092*** (0.018)	0.065*** (0.024)	0.103*** (0.023)	0.017 (0.046)	0.001 (0.034)	0.059* (0.035)
Initial GDP per worker in 1990		-0.167*** (0.038)	-0.157*** (0.045)	-0.167*** (0.048)	-0.188*** (0.046)		-0.193** (0.079)	-0.187** (0.089)
First-stage F statistic					26.77			8.01
Overidentification test p-value					0.63			0.38
Observations	169	169	169	169	169	35	35	35
<i>Panel C. Long differences from 2000 to 2015</i>								
Change in ratio of old to young workers (from 2000 to 2015)	0.051 (0.210)	0.950*** (0.238)	1.028*** (0.282)	0.215 (0.259)	0.809* (0.453)	-0.151 (0.294)	0.009 (0.285)	0.373 (0.463)
Initial GDP per worker in 2000		-0.136*** (0.022)	-0.127*** (0.024)	-0.077** (0.031)	-0.095** (0.037)		-0.259*** (0.057)	-0.256*** (0.054)
First-stage F statistic					19.95			8.06
Overidentification test p-value					0.62			0.10
Observations	169	169	169	169	169	35	35	35
<i>Differential trends by:</i>								
Population and initial age structure			✓	✓	✓		✓	✓
Region				✓	✓			

Notes: The table presents long-difference estimates of the impact of aging on GDP per capita in constant dollars from the Penn World Tables for all countries (columns 1 to 5) and OECD countries (columns 6 to 8). Panels A and C define aging as the change in the ratio of the population above 50 to the population between 20 and 49. Panel B defines aging as the change in the average age of the population above 20. Panels A and B present results for the long differences between 1990 and 2015; while panel C presents results for the long differences between 2000 and 2015. Columns 5 and 8 present IV estimates in which we instrument aging using the birthrate in 1960, 1965, . . . , 1980. The bottom rows indicate additional controls included in the models but not reported. The population and age structure controls include the log of the population and the initial value of our aging measure. We report standard errors robust to heteroskedasticity in parentheses.

Column 2 includes initial log GDP per capita on the right-hand side, while column 3 also adds the initial demographic composition (the ratio of the population above 50 to those between 20

and 49 and log population in 1990). Column 4 in addition includes a set of dummies for World Bank “regions” (which are Latin America, East Asia, South Asia, Africa, North Africa and Middle East, Eastern Europe and Central Asia, and Developed Countries), thus allowing for differential regional trends. With these controls, the relationship between aging and GDP per capita becomes less positive but remains statistically significant at 5 percent. For example, in column 4, the coefficient estimate is 0.773 (standard error = 0.322).⁴ Column 5 estimates the same relationship using birthrates for the 1960, 1965, 1970, 1975 and 1980 cohorts as instruments for our demographic change variable, thus purging it from variation due to migration or changing mortality, which could be endogenous to changes in GDP per capita. In this case, the coefficient of interest becomes even more positive and significant, 1.703 (standard error = 0.411).

Columns 6-8 report the same regressions for 35 OECD countries. In this case the OLS estimates are imprecise, though the IV estimates are once again positive and similar to those in the whole sample—1.186 (standard error = 0.458).

Panel B shows similar patterns with a different measure of aging—change in the average age of the population above 20. Panel C shows that the broad picture is also similar when we focus on the post-2000 sample (2000-2015), where concerns about secular stagnation have become more prominent.

Table 2 extends the sample to 1965 and reports regressions with two differences of 25 years for each country stacked together. Columns 1 and 2 mimic columns 1 and 3 from Table 1 and present very similar results. In addition, columns 3-6 include country dummies, which are equivalent to country-specific linear trends in levels, and report OLS and IV estimates of this more demanding specification. The estimates again point to a positive and statistically significant relationship between population aging and economic growth in the full sample (and a less positive and insignificant one in the OECD).

II. Aging and Robots

Why is there not a strong negative relationship between aging and GDP per capita as predicted by a range of theories, including recent ones on secular stagnation? One possible answer is that technology adjusts so as to undo this potential negative effect. We argue that this answer is

⁴Figure 2 shows that Equatorial Guinea is an outlier, but leaving it out has essentially no impact on the regressions reported here. For example, in the equivalent specification to column 4 without Equatorial Guinea, the coefficient estimate is 0.615 (standard error = 0.290). Equatorial Guinea is an oil producer, and we also confirmed that this has no effect on our results by including the share of oil revenues in GDP in 1990 (from the World Bank) as an additional control.

Table 2: Estimates of the impact of aging on GDP per capita from 1965 to 1990 and 1990 to 2015.

	SAMPLE OF ALL COUNTRIES				OECD COUNTRIES	
	(1)	(2)	(3)	(4)	(5)	(6)
Change in the ratio of old to young workers	0.305 (0.196)	0.346 (0.274)	0.775 (0.491)	2.168*** (0.510)	0.072 (0.322)	0.361 (0.292)
First-stage F statistic				8.64		4.69
Overidentification test p-value				0.44		0.41
Observations	282	282	282	282	63	63
Countries	169	169	169	169	35	35
<i>Differential trends by:</i>						
Initial GDP per capita		✓	✓	✓	✓	✓
Population and initial age structure		✓	✓	✓	✓	✓
Region		✓	✓	✓		
Country			✓	✓	✓	✓

Notes: The table presents stacked-difference estimates of the impact of aging on GDP per capita in constant dollars from the Penn World Tables for all countries (columns 1 to 4) and OECD countries (columns 5 to 6). We define aging as the change in the ratio of the population above 50 to the population between 20 and 49. We report estimates for the periods from 1965 to 1990 and 1990 to 2015. Columns 4 and 6 present IV estimates in which we instrument aging using the birthrate in 1960, 1965, . . . , 1980. The bottom rows indicate additional controls included in the models but not reported. The population and age structure controls include the log of the population and the initial value of our aging measure. We report standard errors robust to heteroskedasticity and serial correlation within countries in parentheses.

plausible in two steps. First, in this section, we draw on Acemoglu and Restrepo (2017) to show that countries experiencing more rapid aging are the ones that have been at the forefront of the adoption of one important type of automation technology: industrial robots.

The relationship between aging and adoption of robotics technology is established in Acemoglu and Restrepo (2017) using data from the International Federation of Robotics (IFR), which provides information on industrial robots across a range of industries for 49 countries. We use the same data in the next figure to show the basic cross-country pattern, which reveals a strong correlation between the same measure of demographic change used in our analysis so far—the change in the ratio of the population above 50 to those between 20 and 49, and the change in the number of robots (per million of labor hours) between the early 1990s and 2015.⁵

Acemoglu and Restrepo (2017) document that this cross-country pattern is robust; it holds if we exclude Korea (a clear outlier) and it holds within the OECD countries. Crucially, as would be expected from a simple model of directed technological change, we also show that it is more pronounced in industries that employ younger workers and those in which there are more opportunities for automation.

⁵This figure excludes Japan, since the IRF notes that Japanese data are not comparable over time because of a change in classification.

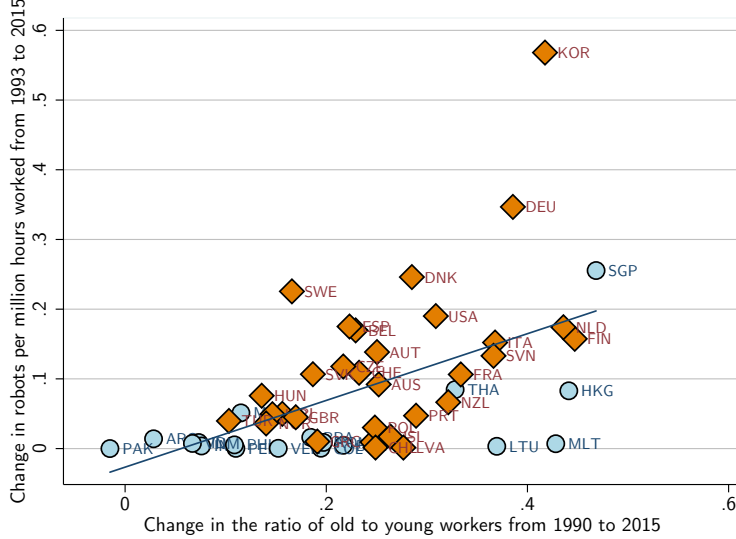


Figure 3: Correlation between change in the ratio of old to young workers between 1990 and 2015 and change in robots per million hours worked between 1993 and 2014 (from the International Robotics Federation).

III. Can Labor Scarcity Lead to Higher GDP Per Capita?

In this section, we undertake the second, theoretical step in our argument. Drawing on Acemoglu (2010) and Acemoglu and Restrepo (2016), we demonstrate that the scarcity of younger and middle-age labor can trigger sufficient adoption of robots (and other automation technologies) so as to actually increase aggregate output, despite the reduced labor input.

For illustration purposes, we use a static model. Suppose that the aggregate production technology is given by the following Cobb-Douglas aggregate over the services of a range of tasks,

$$\ln Y = \int_0^1 \ln y(i) di. \quad (1)$$

Each task i can be produced with capital or labor combined with their specialized intermediates, $q(i)$. In particular, as in Acemoglu and Restrepo (2016), we assume that tasks $i \leq \theta$ are automated and can be produced using capital or labor, with the production function

$$y(i) = q(i)^\eta (k(i) + l(i))^{1-\eta}. \quad (2)$$

Tasks $i > \theta$ can only be produced using labor, and their production function takes the form

$$y(i) = q(i)^\eta l(i)^{1-\eta}, \quad (3)$$

where $\eta \in (0, 1)$. Intermediates, the $q(i)$'s, can be produced at the marginal cost of one unit of the final good, and are supplied by a monopolist which charges a constant markup of $\chi \in (0, 1)$.

The monopolist also chooses $\theta \in [0, 1]$ at cost $C(\theta)Y$, which is interpreted as a (domestic) technology choice to adapt robotics or other automation techniques to the conditions of the country in question. We assume that C is twice differentiable, strictly increasing (reflecting the fact that automating more tasks is costly for the monopolist), strictly convex (with a positive second derivative everywhere), and satisfies the Inada condition $\lim_{\theta \rightarrow 1} C'(\theta) = \infty$.

We assume that capital and labor are inelastically supplied, with supplies given, respectively, as K and L , and

$$K > L. \tag{4}$$

This implies that capital is abundant and cheap relative to labor, which is plausible given the very low interest rates around the world at the moment. This assumption ensures that automating tasks will be profitable and increase aggregate output. In mapping the model to data, we think of L as the supply of younger and middle-aged workers, so that population aging will correspond to a reduction in L —a phenomenon to which we also refer to as an increase in labor scarcity.

Following the same steps as in Acemoglu and Restrepo (2016), equilibrium aggregate output can be expressed as $Y = K^\theta L^{1-\theta}$. Then the profit maximization problem of the monopolist is

$$\max_{\theta \in [0,1]} K^\theta L^{1-\theta} \Gamma(\theta),$$

where $\Gamma(\theta) = \eta\chi - C(\theta)$ is strictly decreasing, has a negative second derivative everywhere, and satisfies $\lim_{\theta \rightarrow 1} \Gamma'(\theta) = 0$. The presence of the term $\eta\chi$ reflects the profits of the monopolist from the markup on the intermediates.

The profit maximization of the monopolist, combined with the Inada condition on C , implies

$$\ln K - \ln L + \Gamma'(\theta) = 0.$$

Differentiating this relationship yields

$$\frac{d\theta}{d \ln L} = \frac{1}{\Gamma''(\theta)} < 0,$$

since Γ has a negative second derivative. This establishes that labor scarcity—i.e., a lower L —encourages further automation as in Acemoglu (2010) and Acemoglu and Restrepo (2016).

What is the effect of labor scarcity on aggregate output? To answer this question, let us totally differentiate the expression for $\ln Y$, taking into account the indirect effect of $\ln L$ working through

additional automation:

$$\begin{aligned}
\frac{d \ln Y}{d \ln L} &= 1 - \theta + \frac{\partial \ln Y}{\partial \theta} \frac{d \theta}{d \ln L} \\
&= 1 - \theta + \frac{\ln K - \ln L}{\Gamma''(\theta)} \\
&= 1 - \theta - \frac{\theta}{\varepsilon_{\Gamma}},
\end{aligned} \tag{5}$$

where $\varepsilon_{\Gamma} = \Gamma''(\theta)\theta/\Gamma'(\theta) > 0$ is the elasticity of the derivative of the Γ function.

In view of condition (4), the second term is negative. Thus a lower L creates a direct effect which is to decrease GDP because of the reduction in the labor input, but also a positive effect through additional automation. If the second term, which is negative, is sufficiently large, then the scarcity of labor caused by an aging population can increase GDP. This will be the case if the gap between K and L is sufficiently large (from the second line of (5)), making capital much cheaper than labor, or if the elasticity ε_{Γ} is small (from the third line). Hence, the aging of the labor force, which reduces the available supply of workers to perform productive tasks in the economy, need not reduce GDP per capita, and may in fact increase it, once we take the response of technology into account.

IV. Conclusion

This paper establishes that, contrary to a range of theories including recent ones on demographics-based secular stagnation, there is no negative relationship between population aging and slower growth of GDP per capita. This is a major puzzle for several theories that have become very popular over the last several years.

One possible explanation for this pattern is the endogenous response of technology—in particular, the adoption of technologies performing tasks previously undertaken by labor. We document that countries undergoing more rapid population aging have adopted more robots, although we do recognize that this evidence is neither causal nor does it establish that the adoption of robots is the mechanism that neutralizes the potential negative effects of population aging on economic growth. We also demonstrate that models of directed technological change can account for the lack of such a negative relationship, and could generate a positive relationship, between population aging and economic growth.

There is a clear need for future work that systematically investigates the relationship between demographic change and GDP growth as well as the channels via which this relationship works.

References

Acemoglu, Daron and Pascual Restrepo (2016) “The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment” NBER Working Paper No. 22252.

Acemoglu, Daron and Pascual Restrepo (2017) “Demographics and Robots: Theory and Evidence” work in progress.

Feyrer, James (2007) “Demographics and Productivity” *Review of Economics and Statistics*, 89(1), 100-109.

Gordon, Robert (2016) *The Rise and Fall of American Growth*, Princeton University Press, Princeton New Jersey.

Hansen, Alvin (1938) “Economic Progress and the Declining Population Growth” *American Economic Review*, 29(1), 1-15.

Lindh, Thomas and Bo Malmberg (1999) “Age Structure Effects and Growth in the OECD, 1950-1990” *Journal of Population Economics*, 12(3), 431-449.

Maestas, Nicole, Kathleen J. Mullen and David Powell (2016) “The Effect of Population Aging on Economic Growth, the Labor Force and Productivity” NBER Working Paper No. 22452.

Murphy, Kevin M. and Finis Welch (1990) “Empirical Age-Earnings Profiles” *Journal of Labor Economics*, H(2), 202-229.

Summers, Lawrence (2013) “Why Stagnation Might Prove to Be the New Normal” *The Financial Times*.

Teulings, Coen and Richard Baldwin (2014) *Secular Stagnation: Facts, Causes and Cures*, CEPR Press.