

Print Edition

Thursday 20 August 2026

EDITION NO. 28 · 1 PIECE · 10 PAGES

IN THIS EDITION

1. Economists should be worried about birth rates · Noahpinion
-

1

Economists should be worried about birth rates

BY NOAH SMITH · NOAHPINION · 19 AUGUST 2026



Photo by Age Cymru on Unsplash

The Economist has an interesting article this week, in which it goes after Daron Acemoglu. Acemoglu is probably the top economist in the world at this point, having just earned a Nobel prize, and sitting right at the top of the list of most cited economists. So it's kind of funny that when it talks about people who doubt Acemoglu's research, the one person it mentions is...me!1

Give an economist a few drinks, however, and some of them will venture their true opinions about this giant. "Much of his theoretical work is useful, but he uses his models to inform populist policies

that have been tried before and failed,” blasts one well-known economist. Some commentators do not require Dutch courage. “I’ve been yelling about Acemoglu for literally a decade,” Noah Smith, an economics blogger, has said, in response to a flurry of online criticism of Mr Acemoglu’s work.

It is true that I have been pretty critical of much of Acemoglu’s work over the years. In 2012 I argued strongly against a paper he wrote in which he claimed that America is more entrepreneurial than Sweden because America’s weaker safety net forces people to work harder (in actuality, Sweden is more entrepreneurial by most measures). In 2022 I pointed out that a famous Acemoglu paper claiming that robots destroy jobs was actually an outlier, and listed a bunch of other papers that reach the opposite conclusion. I liked his book *Why Nations Fail*, but I also admit that its empirical foundations are fairly shaky. I didn’t like Acemoglu’s more recent book *Power and Progress*, and I was highly critical of a recent paper Acemoglu wrote about AI and productivity.² I also think Acemoglu’s argument that skilled immigration has crippled the American education system — now being heavily cited by Breitbart and such — is extremely implausible.

But this is hardly newsworthy. I go after Acemoglu’s work because I know he can take it; he’s a titan of the economics field, and I am but a lowly blogger. Nothing I say is going to affect his reputation or his prestige, even if someone at *The Economist* mentions my critiques in an article. Indeed, after the article came out, top figures in the profession rushed to condemn it and to defend Acemoglu. *The Economist* does list a bunch of — usually justified — criticisms of Acemoglu’s work, but this does not mean Acemoglu has been discredited as a researcher or exposed as overrated; indeed, it would be difficult to name a top economist whose body of work does not contain a variety of questionable theoretical assumptions, motivated reasoning, and/or shaky empirical results. This is a problem with the field itself, not the man; it’s part of a broader crisis of unreliability throughout much of academia.

But for that same reason, despite the *Economist* article’s lack of newsworthiness, I think it’s good that the magazine decided to come at the king. Economics is far too hierarchical and closed of a profession. Younger and less accomplished researchers routinely defer to the authority of famous and senior figures, and critics from outside the field are typically brushed off. Exactly how it got to be this way is a topic that deserves a longer post, but I believe the hierarchical, closed culture of econ has resulted in a research literature that has been too skewed toward the priorities and intuition of top authority figures. I thus think it’s generally a healthy thing to tweak the tails of those Olympian figures, as long as criticisms are grounded in substance. As they say, “science is the belief in the ignorance of experts.”

So in that spirit, let’s critique another Acemoglu paper.

First, the background. Recently, a lot of people (including myself, but also prominent economists) have begun worrying about low fertility

rates. Across all countries — poor as well as rich — fertility just keeps going down and down, with no floor in sight.

This presents two distinct dangers: *population aging* and *population shrinkage*. Aging, rather mechanically, creates a burden for young people, because you have more retirees who have to be supported by each worker, either through taxes or through family support. It also *might* reduce productivity, for example if older managers are less innovative, but that's more speculative. Population shrinkage, meanwhile, is a threat to total GDP, which you might care about if you want your country to be more powerful. There's also the possibility that a smaller population might reduce growth — by reducing the opportunities for specialization, or by reducing the available pool of researchers.

But some people argue that a shrinking population is no problem — or might even be a good thing. When you make human workers scarce, it creates an incentive to invest in labor-saving technology, which boosts productivity. This mechanism has been proposed by some growth theorists, and some economic historians, like Robert Allen, even think this is what caused the Industrial Revolution! This effect might be strong enough to cancel out the aggregate effect of population aging, so that society stays just as rich — or even gets richer — due to low fertility rates.

This is actually a plausible mechanism. I have to say, I'm skeptical that it works in the general case. Human beings aren't just labor *supply*; they also create labor *demand*. The incentive for businesses to buy new machine tools, robots, AI, etc. is that *someone is going to buy the stuff* they produce with those tools. If you have fewer people, you have fewer consumers. This is probably the reason why the effect of immigration on wages is typically close to zero. Babies are not that different from immigrants. A larger market size can also create an incentive for more rapid innovation — in fact, Acemoglu and Linn (2004) find that faster population growth increases pharmaceutical innovation.³

But anyway, the idea that population scarcity stimulates innovation is the thesis of a new paper by Daron Acemoglu, David Autor, Keelan Beirne, and Andrew Scott. This is from the abstract:

Contrary to the widespread expectation that [low birth rates] hamper economic growth, we find lower birth rates are associated with higher growth in GDP per working-age adult across countries and **higher wage growth** across US commuting zones, with no negative impact on aggregate GDP or earnings. These patterns are not explained by educational upgrading, rising female labor force participation, the declining importance of agriculture, or neoclassical-Solow mechanisms. We argue that they reflect the endogenous, labor-saving response of technology to the scarcity of younger workers. Consistent with this interpretation, countries and regions with lower birth rates exhibit more labor-saving patents and growing high-tech activity. There is also **higher TFP growth** across countries and industries. [emphasis mine]

It's a plausible result. But before we go on to the details of the paper, let's talk about what this result would *imply*, if it does turn out to be true.

One of the paper's basic theses is that **technological automation raises wages**. That flies in the face of the empirical work that Acemoglu has done on robots.⁴ It also more broadly seems to contradict what Acemoglu has been saying about AI, both in his research papers and in his public statements. In his 2024 paper, "The Simple Macroeconomics of AI", Acemoglu writes that "predicted TFP gains over the next 10 years [from AI] are... predicted to be less than 0.53%." In his 2021 paper, "Harms of AI", Acemoglu writes:

I argue that if AI continues to be deployed along its current trajectory and remains unregulated, it may produce various social, economic and political harms. These include: damaging competition, consumer privacy and consumer choice; excessively automating work, fueling inequality, inefficiently **pushing down wages**, and **failing to improve worker productivity**. [emphasis mine]

This stands in direct contradiction to Acemoglu's new paper with Autor, Beirne, and Scott. If modern automation technologies push down wages without raising productivity, it cannot compensate for population aging in the way that Acemoglu, Autor, Beirne, and Scott claim that it must — and therefore, population aging is dangerous in a way that it wasn't before.

Now, it's fine for a researcher to find two seemingly contradictory things — or even put forward two contradictory theories — in two different papers. The real world is messy and complicated, and there are plenty of "puzzles" in the literature. But if Acemoglu goes around simultaneously telling us:

1. not to worry about *population aging*, because automation will compensate for it with higher productivity and higher wages, and
2. to worry a lot about automation, because it pushes down wages without raising productivity much...

...then we have a problem.⁵

But anyway, on to the actual details of Acemoglu et al. (2026). First of all, I'm not convinced by their result that lower birth rates raise living standards at the country level. The reason is that once they control for a bunch of stuff — education, urbanization, geographic region, etc. — the result loses statistical significance. Here's their Table 1, with the estimates I'm talking about circled in red:

Table 1: Effect of Birth Rates on Growth in GDP per Working-Age Adult (Countries)

	(1))	(2)	(3)	(4)	(5)	(6)
	Base	Educ	Urban	Region	All	FE
Panel A: Long Differences						
birth rate (t-20)	-0.23 (0.11)	-0.17 (0.11)	-0.21 (0.11)	-0.17 (0.12)	-0.04 (0.12)	
N	107	107	107	107	107	
R-sq	0.494	0.518	0.506	0.575	0.622	
Panel B: Stacked Differences						
birth rate (t-20)	-0.22 (0.05)	-0.18 (0.05)	-0.21 (0.05)	-0.15 (0.06)	-0.11 (0.06)	-0.09 (0.08)
N	229	229	229	229	229	214
R-sq	0.406	0.415	0.408	0.443	0.460	0.838
Panel C: Long Differences - weighted						
birth rate (t-20)	-0.36 (0.15)	-0.17 (0.11)	-0.36 (0.15)	-0.30 (0.19)	-0.04 (0.12)	
N	107	107	107	107	107	
R-sq	0.751	0.518	0.752	0.768	0.622	
Panel D: Stacked Differences - weighted						
birth rate (t-20)	-0.34 (0.05)	-0.35 (0.05)	-0.34 (0.05)	-0.32 (0.05)	-0.32 (0.06)	-0.20 (0.16)
N	229	229	229	229	229	214
R-sq	0.642	0.642	0.642	0.656	0.658	0.841

Notes: This table reports estimates of β^h in equation (10), estimated in our cross-country dataset. The dependent variable is $\Delta_h y_{c,t+h}$, the change in log GDP per worker over horizon h . In each panel we restrict the sample to be consistent across specifications. Panels A and C set $h = 50$ and restrict the sample to be a single cross section, where $t = 1970$. Panels B and D set $h = 25$ and restrict the sample to two non-overlapping cross sections, where $t \in \{1970, 1995\}$. Panels A and B weight countries equally, while Panels C and D weight countries by working-age population in year t . The vector $\mathbf{x}_{c,t}$ includes initial log GDP per working-age adult and labor supply in t , measured by log population aged 20–45 and log population aged 45–70. Column 2 additionally controls for average years of schooling among the population aged 45–55 in year $t - 20$. Column 3 controls for urbanization, measured by the share of people residing in urban areas in t . Column 4 controls for continent fixed effects, while Column 5 adds all controls simultaneously. Column 6 (Panel B only) includes country fixed effects, not identified in the long-difference specifications. Standard errors, clustered by country, are reported in parentheses.

Source: Acemoglu et al. (2026)

In their abstract, the authors declare that “these patterns are not explained by educational upgrading, rising female labor force participation, the declining importance of agriculture, or neoclassical-Solow mechanisms.” But at least some of these patterns can apparently be explained by other things — initial education levels, urbanization, and so on — which makes the claim in the abstract a lot less impressive than it sounds.⁶

Also, note that the more controls the authors add, the weaker the estimated effect becomes. That’s generally a red flag in empirical papers.


No matter how smart and careful you are, there are always things you don't control for, especially in a cross-country regression, since countries differ in so many ways. So if the things you do control for tend to weaken your headline result by a lot, you should be worried that the rest of your result can be explained by the controls you left out.

For example, what about institutions, which Acemoglu has spent much of his career telling us are the main determinants of development? If the legacy of colonialism can be canceled out by passing out free condoms, why did Acemoglu win a Nobel prize? The word “institutions” does not even appear once in this new paper!⁷

A more realistic possibility is that lower birth rates in 1940 were correlated with other things — maybe a lot of other things — and that it's these other things, rather than birth rates, that give rise to the correlation in the paper.

Of course, what we ultimately care about isn't GDP per working-age adult⁸ — it's GDP per capita. The authors have data on GDP per capita, but they barely mention it in the analysis;⁹ I'm not sure why. They do look at total GDP, and here they find no correlation, but a big standard error:

Figure 2: 1950 Birth Rates and Income Growth from 1970 to 2020



Notes: This figure plots Frisch-Waugh residuals from the regressions $\Delta_{50}y_{c,2020} = \alpha + \beta b_{c,1950} + \mathbf{x}'_{c,1970}\gamma + \varepsilon_c$ where $\Delta_{50}y_{c,2020}$ is income growth between 1970 and 2020, $b_{c,1950}$ is the birth rate in 1950 and $\mathbf{x}_{c,1970}$ is a vector of controls including initial income in 1970 alongside the log number of younger workers (20–45) and the log number of older workers (45–70). In Panel A y is log GDP per working-age (20–70) adult across countries, while in Panel B y is the composition-adjusted wage rate across commuting zones. In Panel C y is log aggregate GDP across countries and in Panel D y is the aggregate earnings across commuting zones. To minimize the impact of the COVID-19 pandemic, for countries we proxy for 1970–2020 income growth using 1970–2019 income growth, re-scaled to a 50-year period. Cross-country regressions are unweighted, while commuting-zone regressions are weighted by working-age population in 1970.

Source: Acemoglu et al. (2026)

We shouldn't interpret absence of evidence as evidence of absence. Cross-country regressions have small samples and tons of heterogeneity, so their standard errors tend to be huge — they just have trouble explain-

ing much about the world. But the authors sort of blur this line, saying things like “no negative impact on aggregate GDP or earnings”, when really they should say “we can’t find a negative impact.” There could be a positive relationship between birth rates and growth hiding in that cloud of data points. (Sadly, this is standard practice in economics, but it makes a difference in how results get sold.)

Anyway, the result for wages in U.S. commuting zones holds up somewhat better. It’s a bit more statistically significant, and it doesn’t get attenuated as much by adding controls. But here, I’m worried about different issues: sorting and clustering. Basically, the result shows that coastal cities — places like New York City, San Francisco, Boston, etc. — had both lower birth rates in 1940 and faster wage growth from 1970 to 2020. The authors’ interpretation is that businesses in NYC, SF, and Boston had fewer workers, and so were forced to embrace labor-saving automation, while businesses in places like Birmingham, Alabama or Gary, Indiana refused to automate because they were endowed with plentiful cheap labor thanks to the legacy of an extra-large Baby Boom.

I am suspicious of this result, because *American cities are not independent of each other*. People can pretty easily move from Gary, Indiana to Boston, and vice versa! Ideas and capital can move even more easily — it’s pretty trivial for a company that patents a robot in Boston to actually put that robot to work in Gary. In fact, things like this happen all the time in America. Acemoglu and Restrepo (2017) note that the commuting zones with the greatest exposure to robots include Detroit, Lansing, Saginaw, Defiance, OH, Lorain, OH, Muncie, IN, Racine, WI, and Wilmington, DE.

Just to take one example of how different regions in America are dependent on each other, suppose that over the period from 1970 to 2020, Americans with lots of talent tended to move from Gary to Boston in order to take advantage of the increasing number of knowledge-industry jobs there. And suppose that at the same time, less talented Americans moved from Boston to Gary, in order to take advantage of the cheap land there. And suppose that high-value industries, like biotech and robotics and software, simultaneously decided to put their research labs in Boston rather than Gary because that’s where the best workers were headed.

In this case, you’d see Boston grow more than Gary, for reasons that had nothing to do with birth-rate-driven labor supply. Other research — for example, by Enrico Moretti — has shown that the sorting of talent and knowledge industries was strongly correlated with income divergences between American regions after 1970. Now, Acemoglu et al. do control for initial local education levels, but if sorting really began after 1970 (as Moretti has shown), that doesn’t really help. They also use “composition-adjusted” wages that try to model what local wages would be if education levels were held constant. But even if that adjustment is valid,¹⁰ it’s likely that a lot of skills aren’t captured by simply measuring who has a college degree.¹¹

So in Acemoglu et al.’s telling, a substantial part of the reason the San Francisco Bay Area became Silicon Valley is not because the tech industry naturally tends to cluster in certain locations and draw in talented engineers and entrepreneurs, but because people in the Bay Area in 1940 didn’t

have many kids, which forced Bay Area businesses to embrace automation instead. I simply don't buy that story. I might eventually be *persuaded* to buy that story, but I don't think this paper has sufficient evidence to persuade me yet.

Anyway, I am not claiming that Acemoglu et al. (2026) is a crappy paper. It's clearly an important addition to the literature, and it deals with a very hard area of research where definitive results are very scarce. The basic idea is worth taking seriously. Maybe in a world where AI is growing by leaps and bounds, a scarcity of human workers should be the last thing on our minds!

But we shouldn't regard this paper as definitive. And it contradicts a bunch of other research:

- Kremer (1993) finds evidence that regions with faster population growth had faster technological progress.
- The endogenous growth literature argues that because ideas are non-rival, larger populations are able to sustain faster technological progress — this is a basic assumption in Romer (1990), Jones (2022), and many other models.
- Maestas, Mullen, and Powell (2023) find that aging was negatively correlated with productivity growth among U.S. states from 1980 to 2010; 12 Ozimek, DeAntonio, and Zandi (2018) find something similar.
- Aksoy et al. (2019) do a panel time-series regression across OECD countries from 1970 to 2014, and find that low fertility and low population growth are correlated with slower economic growth.

And even if Acemoglu et al. (2026) is right about the past, that doesn't mean we can extrapolate it to the *future*. Jesús Fernández-Villaverde points out that you can't necessarily do a linear extrapolation of effects at above-replacement fertility to effects at below-replacement fertility:

The fertility collapse now underway has no historical precedent. The seven decades of data the [Acemoglu et al. (2026)] paper analyzes, however carefully, contain nothing remotely comparable to the ultra-low fertility rates we observe today...

[T]he birth and fertility rates of 1950 provide little information about the consequences of the ultra-low birth and fertility rates we are experiencing today. Observations from a very different range of values are unlikely to capture the nonlinear cumulative effects of ultra-low birth and fertility rates...[E]ven in 1980, birth and fertility rates were not as low as they are today...A TFR of 1.56 means each generation is 76% the size of the last: the population declines about 0.9% per year, so after 100 years you retain roughly 40% of the initial population. A TFR of 0.75 means each generation is 37% of the last: about a 3.35% decline per year, leaving roughly 3.5% after a century. Hence, we are dealing with a factor of about 11. A TFR of 1.56 and a TFR of 0.75 both count as "low fertility" in a regression, but they are qualitatively different regimes. The 1980 variation the paper identifies lives in the gentle-decline world; nothing in their data resembles the second...

[W]e should not fool ourselves: we are sailing into demographic terra incognita, and in these waters, the pretense of knowledge is the most dangerous temptation.

In fact, Hayashi (2025) documents just such a nonlinear relationship between aging and growth, where countries with above-replacement fertility seem to benefit from lower birth rates, but countries with below-replacement fertility seem to do worse when their birth rates fall even further.¹³

External validity and nonlinearity are big problems in much of growth economics, development economics, economic history, and so on (and often in other fields too). But it means that there's a danger, when the world's top economist writes a paper making big bold claims, that we end up deferring to that individual's intellectual authority instead of recognizing the true uncertainty of the unprecedented changes our world is facing. I'm still very worried about the shrinkage of the human race, and I think you should probably still be worried too.

1

For the record, *The Economist* did not actually interview me for this article. Their quote of me was from a tweet from June, which they did not link to in the article.

2

Because I'm being critical of these works, I'm violating a standard rule of etiquette, which is to always list co-authors equally. Acemoglu's co-authors have mostly chosen to eschew the role of outspoken public intellectual, and most are not as untouchably famous as he is, so I don't want to put them "on blast" on a popular blog. But for the sake of completeness, Acemoglu's 2012 paper on "cuddly capitalism" is with James Robinson and Thierry Verdier, *Why Nations Fail* is written with Robinson, *Power and Progress* is written with Simon Johnson, and Acemoglu's 2017 paper on robots and jobs is with Pascual Restrepo. My apologies for not including these co-authors in the main text.

3

See also Acemoglu (2009), which theorizes that labor scarcity discourages the creation of technology that complements human labor:

The main result of the paper shows that labor scarcity will encourage technological advances if technology is strongly labor saving. In contrast, labor scarcity will discourage technological advances if technology is strongly labor complementary.

4

Acemoglu and Restrepo (2017) find that robots depress wages, whereas the entire mechanism in Acemoglu et al. (2026) relies on automation raising wages. Of course, Acemoglu might claim that robots are special, and not representative of automation technologies in general. But that would make the 2017 result a lot less important and generalizable than it was sold as being, back in 2017!

5

Also note that Acemoglu et al. (2026) can be read as an *anti-immigration* result. If young population growth stifles automation, doesn't that mean we should restrict immigration in order to raise wages? That doesn't mean the result is *wrong*, and it will certainly make MAGA types happy, but it does mean that it stands in conflict with the existing empirical literature on immigration and wages.

6

The "stacked differences" model does retain statistical significance after the inclusion of controls, but just barely. I wouldn't put much faith in a borderline result like that.

7

The word "institutional" does appear once, but not as an explanatory variable for past patterns of growth:

At the same time, the demographic changes currently underway are accompanied by increased life expectancy (Scott, 2021, 2024), which may spur institutional changes, policies, and additional human capital investments that complement longer lifespans. These forces may counteract any negative effects of aging[.]

8

Actually, this is a bit beside the point, but I wish economists would stop using GDP per working-age adult as a measure of *anything*. Working-age adults are not some magical resource that creates GDP just by eating and breathing; they must *actually work*. Data on output per hour of work is readily available, and if it's not available, output per employed person is usually a good proxy. I don't see any good reason to use output per working-age adult. And you should NEVER mislabel output per working-age adult "output per worker", as Acemoglu et al. (2026) do throughout their paper. Bad!!

9

Except in one section later in the paper, about WW2.

10

Post-1970 clustering effects could change the relationship between education and wages, in which case the adjustment wouldn't hold.

11

In fact, the authors note the greater accumulation of college degrees in the coastal metros with low 1940 birth rates after 1970, but they attribute this — without justification, as far as I can tell — to deliberate local choices and policies rather than geographic sorting.

12

Acemoglu et al. (2026) have an appendix arguing that the Maestas et al. (2023) result is not robust.

13

Acemoglu et al. (2026) have an appendix arguing that the Hayashi (2025) result is not robust.