ECONOMIC INNOVATION GROUP

The American Worker Project

What isn't the matter with American workers?

Paul Krugman

June 2024

Editor's note: This is one of a pair of essays responding to the Economic Innovation Group's report, "The American Worker: Toward a New Consensus", by Adam Ozimek, John Lettieri, and Benjamin Glasner. The other is by Michael R. Strain of the American Enterprise Institute.

Life got immensely better for American workers for the first few decades after World War II. Real wages soared, hours fell, working conditions improved.

Progress then stumbled for a while in the face of oil shocks and stagflation, although real incomes held up better during the 1970s than legend has it. But what happened to workers after inflation was tamed and it became Morning in America?

We know that despite all the hype around technology, productivity growth has been much slower since 1980 than it was in the postwar generation. We also know that income inequality rose, so that some of the gains from rising productivity were siphoned off from the working class, accruing to a relatively small group at the top instead.

But did workers actually lose ground, as many suggest? Ozimek, Lettieri, and Glasner say no. Despite everything, they argue, the median worker has experienced substantial real wage gains, while workplaces have become safer and more pleasant, and job security has at least modestly improved.

So, are they right? (Yes.) If they are, why do so many believe otherwise? (I have some ideas.) And what does all of this imply? ("Better than you may have thought" doesn't mean "good enough".)

Have workers actually made progress?

Ozimek et al show that median wages have clearly outpaced inflation since 1980 unless you use a measure of consumer prices that has well-known flaws; this is true even without considering the surely under-measured gains from new technologies. It's important to note that the main dispute over these unmeasured gains is whether they have been bigger in recent decades than in the past; the reality of such gains isn't in dispute. In general, the authors' discussion of the measurement issues is solid and persuasive.

Still, whenever there is dispute about what is really happening to the economy, it's always helpful to kick the tires by seeing whether other sources of evidence are consistent with the standard data. For example, right now there are many people claiming either that official measures greatly understate inflation or that the recent decline in the inflation rate is a myth, so it's helpful to cross-check the Bureau of Labor Statistics numbers against various private-sector sources, from McDonald's news releases to surveys of purchasing managers, to make sure that they're more or less consistent with the official data. (They are.)

So do other measures show a rise in living standards consistent with what Ozimek et al say about rising worker incomes? Yes. For example, the size of the median new single-family house has gone from less than 1,700 square feet in the 1970s to approximately 2,300. The great majority of Americans already owned cars by 1980, but today's cars are much better (I'm old enough to remember). The number of U.S. air travelers tripled between 1980 and 2019, even though the population rose only 55 percent, strongly suggesting a democratization of what used to be elite travel. And so on down the line. I believe that if you could somehow take a time machine back to observe the daily life of ordinary American workers in 1980 you'd be somewhat shocked by how cramped and relatively shabby it seems.

So yes, real standards of living are up. It's not a transformative change like that of the postwar boom, but it's significant.

The authors also document the huge improvement we've seen in workplace safety. I'd just add that the air and water have become much cleaner, which is also part of the standard of living. And the data saying that workers remain overwhelmingly satisfied with their jobs is new to me, but also striking.

So are American workers better off than they were in 1980? Almost surely. Any calculation indicating otherwise fails the credibility test, in that it's at odds with a lot of other evidence.

In that case, however, why do so many people believe that things were better in the past?

Why we think too well of the past

Imagining that the past was better, especially compared with the present, than it actually was has been common throughout history. The ancients believed in bygone golden ages. When the reformer William Cobbett wrote his classic "Rural rides" in the 1820s, he provided invaluably acute observations of what life was really like in the country—but also managed to convince himself that England, already embarked on the Industrial Revolution, had grown poorer over the past generation.

To the extent that many of us suffer from a similar misconception about what has happened to American workers, I'd like to speculatively suggest three stories, by no means mutually exclusive. And of course there may be other stories I'm missing.

First is what I think of as the *Bridgerton* effect.

What is your image (if you have one) of life in Regency England? My guess is that it largely involves the life of aristocrats or, at best, that of landed gentry. These groups weren't representative! Most people's lives were poor, nasty, brutish and short—but those people aren't the subject of Netflix dramas or even Jane Austen novels.

Obviously the 1960s and 1970s are still within living memory, and portraits of American lives in that era covered a wider spectrum than the literature of earlier eras—it wasn't all Ozzie and Harriet. Yet even though there were TV shows and movies about working-class Americans, we rarely hear about the lives of migrant farm workers and urban day laborers, classes that almost surely were larger then than they are now, or that of the millions of workers with poorly paid, unpleasant and/or dangerous jobs. That is, our cultural memory puts too much weight on workers who more or less got to live the American Dream, which makes our image of the past condition of workers rosier than it should be.

This bias toward remembering the good stuff may be reinforced by a second factor: the natural tendency to judge an economy by its most prominent enterprises. During the great postwar boom, America's largest private employer (other than the AT&T system) was General Motors—a corporation with both great market power and a powerful union, which in effect shared some of its monopoly rents with employees. Some people who worked there called it Generous Motors; it certainly offered good pay, good benefits, and considerable job security.

Today our biggest private-sector employers are service-sector giants, especially Walmart and Amazon, which are notorious for low wages, poor working conditions and high turnover. (Walmart has gotten better, but is still hardly an encouraging model.)

But even the biggest employers account for only a small fraction of the work force, and need not be representative. In 1970, when the U.A.W. staged a major strike against G.M., it involved 400,000 workers—a big number, but only 0.6 percent of nonfarm employment. Whole books were written—John Kenneth Galbraith's *The New Industrial State* comes to mind—largely based on the premise that GM represented the essence of the U.S. economy. But it never did.

Our current biggest conventional private employers have, for whatever reason, not established a comparable grip on public perceptions. But in some ways Uber, which isn't exactly an employer but certainly has a large number of people working via its platform, has established a similar iconic status; quite a lot of discussion I encounter about the economy sees modern work as defined by gig work "like Uber". There has been a huge statistical debate over the extent of gig work, which I am in no position to referee; but my sense is that the portion of employment involving Uber-like arrangements consists in large part of … Uber.

The hype over gig work is, arguably, part of my third point: There has been a big change in what we might call the ideology of business that glamorizes what business leaders do but may simultaneously create negative perceptions about the state of workers.

Here's what I mean: In the 1950s and 1960s the popular image of a top business executive was, as far as I can tell, The Organization Man—a colorless bureaucrat shuffling paper. This was true even though those were decades of rapid technological progress and substantial structural change.

Since the 1980s, however, business leaders have increasingly been portrayed—or portrayed themselves—as agents of radical change, of disruption, of creative destruction and so on. They do this even though, at least as measured by total factor productivity, technological change has slowed, not accelerated.

As Ozimek et al document, it's not at all clear that today's economy really is more dynamic and unstable than the economy pre-1980; certainly job security has been stable or even improving. So where is this ideology of constant change coming from? Part of it may be the explosion of executive compensation: CEOs who are being paid 300 times as much as their workers may, perhaps unconsciously, feel the need to see themselves as heroes contending with a world in constant turmoil rather than mere bureaucrats, even if that's basically what they are.

But to the extent that commentators buy into the vision of modern business as a wild world of constant struggle in which only the paranoid survive, they may also buy into a vision of highly unstable work—a world that is very stressful and insecure for workers—which turns out not to be true.

Finally, while median real wages have risen, they have lagged well behind wages at the 90th percentile, let alone the incomes of the 1 percent. Does this adversely affect workers' welfare? I would argue yes, in a couple of ways.

First, there are surely ways in which the affluence of a minority can have concrete negative impacts on the lives of those less affluent, albeit in ways that may not be captured by measures of real wages.

Earlier I used the growing size of houses as evidence of rising living standards, but the Census data do show significant retrogression on one dimension. Almost two-thirds of houses built in the 1960s and before had access to public transportation, but only around a third of those built in the 21st century. There are several reasons for this decline, but isn't it likely that as multiple car ownership spread among the affluent, public transit atrophied, hurting those less affluent?

Even the growth in house sizes is a bit ambiguous. To what extent is the median worker effectively forced to buy a bigger house because high-wage workers are adding square feet? You might say trust the market—but higher-income households may set norms, sometimes enforced by zoning and building codes, that hurt lower-income households.

Beyond these material issues, status matters; it's no use telling people that it shouldn't. Real incomes are up, but there has been a gradual decline in the percentage of Americans who consider themselves middle class. So you can debunk claims about falling wages while still acknowledging that rising inequality is a problem.

These are fairly disparate stories about why we might have the impression that workers have lost ground. It's clearly important to set the record straight and point out that real wages are actually up. But it's equally clear that we'd be feeling better if real wages had risen faster. What, if anything, could we do to accelerate the pace of wage gains?

What to do

What could we do to improve real wage growth going forward? Obviously this isn't a policy manifesto. Instead, let me offer a ranking, from hardest to easiest, of things we might try to accomplish.

One major reason real wage growth has been slower since 1980 than it was before is that productivity has risen at a slower pace. So why not try to increase productivity growth?

The answer is that for the most part we don't know how to do that. It's extremely hard to come up with policies that we would expect to have more than a marginal impact on productivity growth. Robert Solow famously noted that attempts to explain slow British growth usually end in a "blaze of amateur sociology". The same is true for changes in the rate of productivity growth over time. We should do what we can, but not expect much.

While major increases in productivity growth seem hard, raising median real incomes by reducing inequality seems much more achievable. Redistribution via taxes and transfers is obviously feasible. Even market inequality seems a lot more amenable to policy than many economists used to believe; we've seen a dramatic narrowing of wage gaps in recent years, probably simply by finally running the economy near full employment.

Finally, as Ozimek et al point out, while we don't know much about how to raise productivity, we know a lot about how to make housing—a large part of the consumption basket—cheaper. To a large extent it's just a matter of removing the regulatory barriers that inhibit construction.

So what's the moral? The common view that American workers have actually suffered a decline in living standards is wrong. But gains have been much slower since 1980 than before, so we shouldn't just lecture people about accepting how the economy has served workers. And while nobody knows how to have another generation of surging living standards comparable to the postwar decades, that still leaves plenty of room to do better.

Explore the Economic Innovation Group's American Worker Project here.