

Faster productivity growth would be great. But don't count on it to raise middle-class incomes.

By Jared Bernstein

The stagnation of middle-class incomes is one of the most important and influential trends in American economics. Politicians, from President Obama to Hillary Clinton and Jeb Bush have consistently argued that loosening the middle-class squeeze is a major goal of their policies. Prominent economist Larry Summers recently [wrote](#) that in crafting “global as well as domestic” economic policy, “the middle class matters the most,” warning against approaches that offer “little to those in the middle.”

In recent weeks, largely motivated by an analysis by President Obama's economic team, a solution to stagnant middle-income growth has been elevated: faster productivity growth. The President's Council of Economic Advisors (CEA) tells us the [following](#):

“What if productivity growth from 1973 to 2013 had continued at its pace from the previous 25 years? In this scenario, incomes would have been 58 percent higher in 2013. If these gains were distributed proportionately in 2013, then the median household would have had an additional \$30,000 in income.”

That's a huge increase—more than 50 percent—over the current median household income of about \$52,000. Given an increase of this magnitude, it's no surprise that the idea that higher productivity growth is an important answer to stagnant middle-class incomes has been touted by [various commentators](#) in this debate. Based on the CEA's [calculation](#), Tyler Cowan goes as far as to claim that faster productivity growth would not just boost middle class incomes, but would increase their economic mobility, implying they'd gain relative to other groups.

There are, however, two problems with all this cheerleading. First, there is a large and persistent gap between productivity growth and middle-class incomes: we cannot realistically assume that faster productivity growth would reach the middle as opposed to doing an end-run around them on its way to the top. Second, while it would be great to have faster productivity growth, I'm afraid we must be humble about our ability to make that happen.

In what follows, I document the inequality-induced split between median family income and productivity growth, with a focus on the point that, while productivity did slow in the mid-1970s, it still grew much faster than middle incomes. In this regard, income inequality has increasingly played a wedge-like role between middle-class income and productivity growth.

This observation on the role of inequality in the income/productivity split is key to my argument. Productivity is really another way of looking at overall, or macroeconomic growth—it is *aggregate output* per hour. A central, even definitional, characteristic of increased inequality is that less growth reaches the middle class relative to the wealthy. Thus, to assign the benefits of faster productivity growth to the middle class is analogous

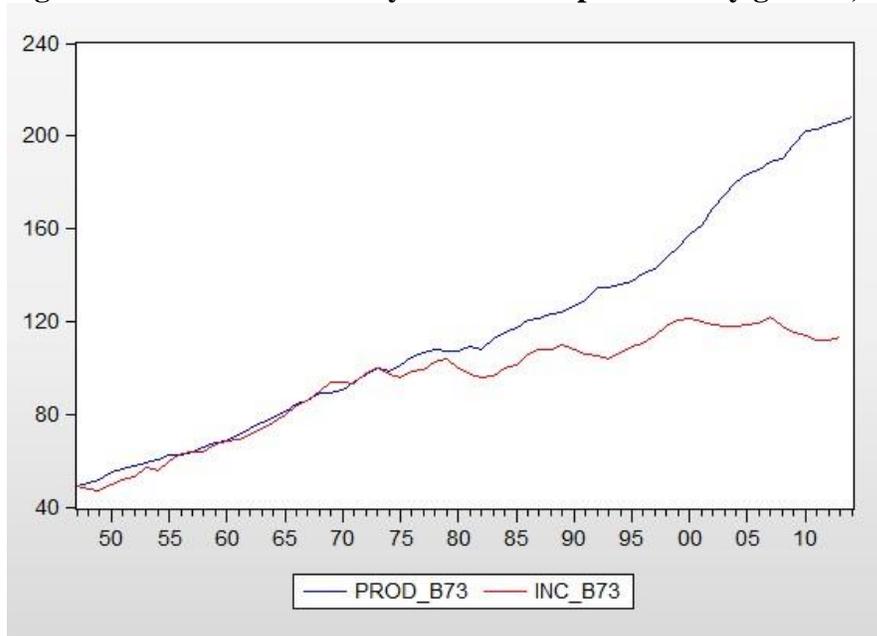
to assuming away growing inequality, an obvious “diagnostic” mistake with the potential to distract policy makers who want to help middle-class families from the importance of policies designed to push back against rising inequality.

To be clear, faster productivity growth would be a highly welcomed development, and would surely help boost median incomes to some degree. But as long as inequality continues to play a wedge-like role between aggregate growth and the incomes of low- and middle-income households, analysts need to be much more cautious in their assumptions regarding the benefits of faster productivity to the middle class.

The Income/Productivity Split

Figure 1 plots real median family income against productivity growth from 1947 to 2013. Both series are indexed to 100 in 1973 to put them on the same scale; thus, these indexes show percentage changes relative to the base year. Both series just about doubled between 1947 and 1973. Since then, productivity has doubled again, while median family income is up only 13 percent.

Figure 1: Real median family income and productivity growth, 1973=100



To be sure, the pace of output-per-hour has slowed. Productivity doubled in about 25 years in the earlier, “growing together” period, and about 37 years in the post-1973 “growing apart” period. But as the figure shows, its growth slowed much less than that of median income, leading to the large and growing gap. As inequality has increased, more of whatever level of productivity growth there is has gone to those at the top of the income scale.

What happened? Why did productivity slow and why did it cease to reach the median income?

The first question, while obviously critical, is beyond my scope as there is no consensus among economists to explain the deceleration in productivity growth. Analysts have cited higher energy prices, measurement issues (underestimating efficiency improvements), labor quality, and weaker technological gains, but our understanding of the mid-70s slowdown is no further along than that of the speedup in the mid-90s¹ or of the recent slowdown (note in Figure 1 how the slope on the productivity variable rises less quickly post-2010).

On the other hand, we know more about the causes of the slowdown of middle-class incomes. Especially for working-age families, stagnant earnings, especially for men, have been identified as a key factor. This in turn relates to weaker labor markets—full employment was the exception over the “growing apart” period and the norm over the “growing together” period—persistent trade deficits causing the loss of manufacturing jobs, declining union power, a fall in the real minimum wage (for low-income families), and, at least until around 2000, the growth in the wage premium for more highly educated workers. Remarkably, as EPI’s Josh Bivens has emphasized, this income slowdown occurred even as families significantly increased their time spent in the paid labor market (see Table 2.17 [here](#)).

So, is it wrong to claim faster productivity growth would lift the middle class?

Not necessarily. In their original analysis, CEA was perfectly explicit about their approach, which was simply to say that if productivity growth a) hadn’t slowed, and b) continued to reach middle-income households the way it used to, those households would be more than 50% better off. That’s a legitimate thought experiment, but one that can perhaps be improved upon.

Another approach is to use statistical model that uses a number of variables, including productivity growth, to explain the growth in real median income. Then one can alter the path of key predictors, like productivity, to engage in thought experiments similar to that of the CEA. To be clear, both family income and productivity growth are complex processes and the type of time-series analysis I present here is a blunt tool to explain them. But the approach still yields some useful insights.

The model predicts the annual change in real family income using productivity growth, unemployment, inflation, inequality, and women’s labor force participation (to capture the impact of increased female labor supply on family income). See the short data appendix for details on the statistical procedure (as well as a link to the underlying data in case you want to fool around with them yourself).

By far, the trickiest variable to model in this equation is the one we care most about here: productivity growth. The relationship between productivity and *average compensation* is well known to economists: average compensation divided by productivity equals, by definition, labor’s share of national income (see appendix). For many years, the labor

¹ To be fair, careful analysis suggests investment in information technology played a significant role in the mid-1990s speedup, though as I show later, that speedup appears to have been curiously short-lived.

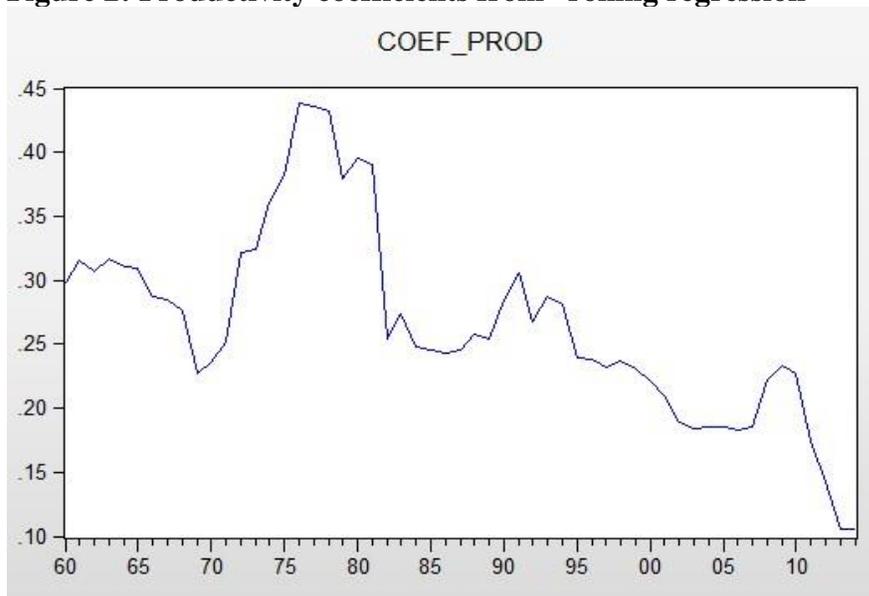
share of income has been fairly constant, implying a tight relationship between average (as opposed to median) pay and productivity growth (the fact that this hasn't been the case in recent years—even average comp has fallen behind productivity growth—is yet another indictment of the notion that faster productivity would reach the middle class).

But a characteristic of inequality is the divergence of the average from the median, and there is no simple equation that links output-per-hour to the median wage. Moreover, income involves a lot more than wages—it includes hours of labor supplied by the family as well as other, non-labor sources of income (more on that below). What we do know, as per Figure 1, is that productivity used to grow along with median income; then, when inequality began to grow, the two variables diverged.

At any rate, the simple specification does a good job explaining income growth and the coefficients on each variable all have the expected signs. Productivity growth, however, is insignificant (see appendix), which is unsurprising in light of its diminished correlation with income over time, as well as the likely faulty assumption that productivity growth contemporaneously feeds into income growth (though various lag structures don't help).

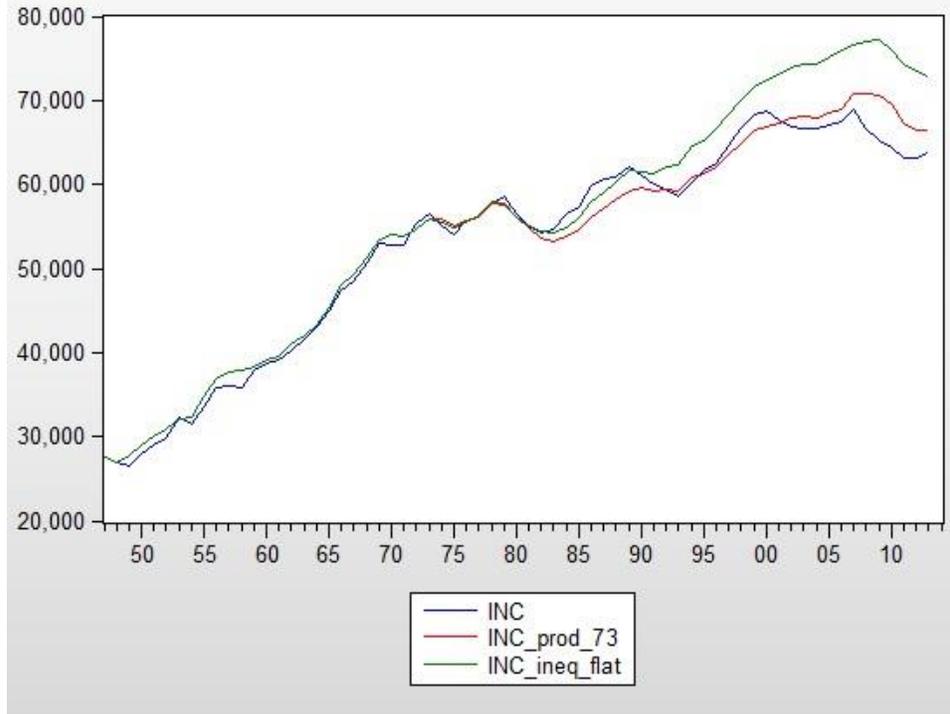
A good way to see the evolution of the relationship between income and productivity is to run a “rolling regression,” which starts with a short sample and adds years one-by-one, pulling off the coefficients on productivity as it goes. If the correlation between middle-class income and productivity growth were relatively constant, adding years to the equation wouldn't change much, but as shown in Figure 2, the coefficient on productivity growth is much higher—though never statistically significant—before inequality began to grow, hitting about 0.45 at its peak in the mid-1970s (meaning a 10% increase in productivity growth might be expected to generate slightly less than a 5% increase in real median income). In other words, to the extent that productivity growth can lift median incomes, its ability to do so has clearly diminished over time.

Figure 2: Productivity coefficients from “rolling regression”



As noted, I can use the model to simulate what would have happened had productivity not slowed or inequality not sped up. Figure 3 shows the results of these simulations. Allowing productivity to keep growing at its pre-1973 pace of 2.8% per year boosts middle income growth a little, but not nearly as much as holding inequality at its 1973 level.

Figure 3: Simulated effects of productivity and inequality on median income.



Based on these simulations, the productivity slowdown cost the median family about \$2,600 by 2013, while the rise inequality cost them \$9,000.

As intimated above, it is not obvious that my approach is the right one to model productivity and family income. One might suppose that a longer-term changes in productivity growth would correlate more closely with middle-class income growth, although that approach didn't help in the model. On the other hand, long differences in both income and productivity are positively correlated ($r=0.61$ for 10-year percent changes) suggesting perhaps a slow-moving, longer-term relationship between the two variables. However, even this correlation fades if we begin the sample in the late 1970s, as inequality takes hold.

Do Other Income Series Tell a Different Story?

The first panel in the table below, using the same data in Figure 1, again shows that productivity more than doubled both pre- and post-1973, though it did so much more quickly in the earlier period. The number that stands out in the upper panel is 13%, revealing a sharp slowing in median family income growth post-1973.

Along with the regression exercise, it is on this basis that I argue that policy makers should be less sanguine than the commentators noted above that faster productivity growth, as compared to actions to reduce inequality, will raise middle-class incomes.

**Real Median Family Income Growth and Productivity Growth
1947-2013**

	<u>Income</u>	<u>Productivity</u>
1947-1973	104%	105%
1973-2013	13%	107%

**Addendum: CBO Comprehensive Median Household Income
1979-2011**

Pretax/transfer	17%	89%
Post-tax/transfer	47%	

Sources: Census, BLS, CBO

However, some critics argue that the Census data inadequately capture median income growth they leaving out various income sources. The bottom panel of the table above thus shows the growth in real median household income from 1979-2011 using another income series from the Congressional Budget Office.

“Market” incomes are up 17% in this series, compared to 89% for productivity growth, so switching income series does not meaningfully change the productivity/income split.

CBO also provides a more comprehensive income measure that accounts for the impact of (federal) government taxes and transfers. Post-tax and transfer income growth is a lot faster, as shown in row 2 of the bottom panel, though there’s still a 40% gap between productivity and median income growth.

Moreover, we would likely expect productivity growth to lift middle-class incomes through higher earnings, not income redistributed through government sources. That is, I don’t think those arguing that productivity growth will lift median incomes mean to argue: “we must raise productivity growth so that we can offset higher inequality by redistributing more income from those at the top to those in the middle.” Instead, they more likely believe that faster productivity growth should raise the wages/market incomes of the middle-class. If so, these data too point to the importance of attacking the sources of inequality.

Conclusion

There are numerous problems with the assumption that faster productivity growth will raise middle-class incomes:

--rising inequality has sharply lowered the correlation between the two variables, so we cannot assume that faster productivity growth will necessarily reach the middle class;

--while we don't know how to reliably generate faster productivity growth, we have a better idea of policy interventions to boost the bargaining power and earnings of low- and middle-class workers;

--there's no clear, conceptual model linking productivity and median income growth in a high-inequality environment, which raises the evidentiary bar for those making broad claims about how productivity will boost middle-class incomes; what model do they have in mind and is it empirically justified?

None of these points should deter us from the pursuit of faster productivity growth, but that unfortunately remains somewhat of a black box for economists. On the other hand, raising the minimum wage, pursuing full employment through fiscal and monetary policy, boosting collective bargaining, and other such interventions have all been shown to raise the pay and bargaining clout of middle- and low-wage workers. Whatever the pace of productivity growth, measures like these are a lot more likely to lift the incomes of the middle class.

Data appendix

Regression output:

Dependent Variable: DLOG(INC)
 Method: Least Squares
 Date: 04/15/15 Time: 15:16
 Sample (adjusted): 1949 2013
 Included observations: 65 after adjustments

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	-0.053782	0.027057	-1.987722	0.0515
DLNFB	0.105410	0.161112	0.654267	0.5155
LOG(UNRT)	-0.025590	0.008718	-2.935420	0.0047
DLOG(RS)	-0.543885	0.109139	-4.983400	0.0000
D(LFPR_W)	0.028527	0.004939	5.775820	0.0000
D(GINI_SM)	-1.442918	0.430509	-3.351654	0.0014
R-squared	0.643220	Mean dependent var		0.013279
Adjusted R-squared	0.612985	S.D. dependent var		0.026348
S.E. of regression	0.016391	Akaike info criterion		-5.296364
Sum squared resid	0.015852	Schwarz criterion		-5.095651
Log likelihood	178.1318	Hannan-Quinn criter.		-5.217170
F-statistic	21.27363	Durbin-Watson stat		1.771123
Prob(F-statistic)	0.000000			

Notes: The dependent variable is the log change in real median family income (Census Bureau, CPI-RS adjusted). DLNFB is the log change in non-farm business productivity (BLS); UNRT is the unemployment rate (BLS); RS is the CPI-RS deflator, entered as a log change (BLS); LFPR_W is women's labor force participation rate, entered as a change (BLS); GINI_SM is the Gini coefficient for family income, entered as a change (Census). In 1993, a top-coding change led to a large jump in this Gini measure, so I do a log-linear interpolation phasing the 1992-93 increase over 1992-95.

The text notes that average compensation divided by productivity is equal to labor's share of aggregate income. In the equation below, w is the average wage (or compensation), L is hours worked, and Y is total income. Thus, the first fraction equals labor's share of national income. Divide both numerator and denominator by L and we end up with the result, since income/hours is the definition of productivity.

$$\frac{wL}{Y} = \frac{wL/L}{Y/L} = \frac{w}{Prod}$$

There is, however, no similar identity for medians.

Here is the data set used in this study: http://jaredbernsteinblog.com/wp-content/uploads/2015/04/faminc_ts.xlsx