Richard Startz is a professor of economics at the University of Washington.

Commentary

Richard Startz

When I was an undergraduate, I was told that maybe a little inflation was a good thing because nominal wages are downward rigid. What's more, back then this accounted in part for the Phillips curve being curved—being relatively flat at low inflation rates and steeper at high inflation rates. At low inflation rates, you had to have a lot of unemployment to get rid of a little inflation because nominal wage cuts are so painful. Downward wage rigidity is regarded as one of those truths that are so self-evident that there is no need to look at the data. Professor McLaughlin has actually looked at the facts and said, "If it's so self-evident, why doesn't it show up in the data?" Or, at the very least, said, "It's not nearly so simple."

Let me outline six points for consideration.

- Where does wage rigidity fit into macro?
- Do nominal wage changes move one-for-one with inflation?
- Are wage changes skewed?
- Does the skewness change with inflation?
- Is the spike at zero big? Does its size change with inflation?
- What about the spread of the distribution? Does it change with inflation?

Let me begin by saying that Professor McLaughlin's topic is really important, that it is really important to the Federal Reserve, and that it is really important at exactly this point in history. If you take a policy window of the last two or three years and the next two or three years (assuming a similar economy) a critical question for monetary policy is whether we should aim for a steady zero inflation rate or for a steady 2 percent inflation rate. The argument for a zero inflation rate is that zero is a magic number for political reasons and reasons of transparency. The argument for 2 percent is that nominal wages are downward rigid, and that 2 percent allows for more flexible real wages. So McLaughlin's paper bears precisely on the central question of medium-term monetary policy.

Second, using micro data is exactly the right way to answer this sort of question. It also is a lot of work. You have to worry about measurement error. You have to worry about exact definitions of survey data. And so forth. McLaughlin's paper is very well done and deserves a great deal of appreciation for both the quantity and the quality of the work.

By way of final preface, where does wage rigidity fit into macro? Specifically, in a recession, why don't wages drop to clear the market? Let me give the old-fashioned answer. There are at least four places where wage rigidity fits:

First, from an old-fashioned Keynesian viewpoint, the labor market is driven by the demand side as in:

$$L = L^D\left(\frac{w}{p}\right). \tag{1}$$

If the real wage doesn't drop, too much unemployment results. In this situation, if nominal wages are downward rigid, then real wages surely will be.

Second, consider a model in which the real money supply matters: perhaps quantity theory, perhaps Keynesian with a Pigou effect, or perhaps ISLM.

$$Y = v\left(\frac{M}{p}\right) \quad \text{or} \quad Y = C\left(Y, \frac{M}{p}\right) + I + G \tag{2}$$

or ISLM

$$P = \mu W \tag{3}$$

Suppose firms set prices as a markup over wages—maybe a competitive markup, maybe something different. Here, nominal

Figure 1

Illustration of Phillips Curves with Nominal vs. Real Wage Rigidity

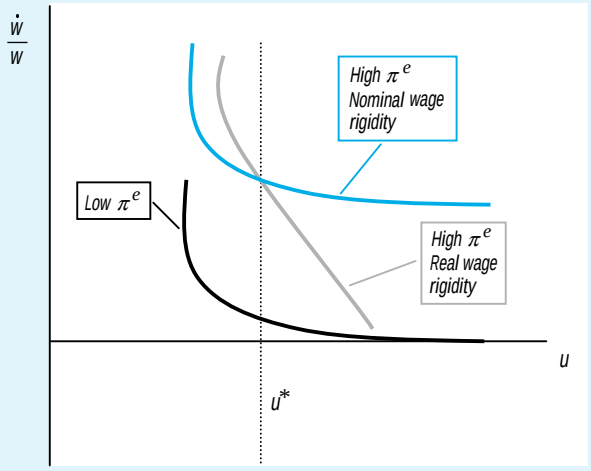

wage rigidity causes nominal price rigidity and prevents the consumer price index from dropping to pull the economy out of recession.

The third place where wage rigidity shows up is in the slope of the Phillips curve at low inflation rates. (See Figure 1.) If it is hard to force down wages, then the sacrifice ratio is really bad at low or negative inflation rates. One way to think about real vs. nominal rigidity is to ask whether the sacrifice ratio worsens specifically near zero or just at inflation rates close to expected inflation.

Fourth, and last, rigidity just messes up allocational efficiency.

The paper is excellent. Nonetheless, nobody gives a discussant credit for saying nice things; so I want to discuss the one part with which I disagree and then make a few suggestions for other ways of working with the data.

The only part of the paper with which I take issue is the conclusion that "nominal wage changes move one-for-one with anticipated inflation, and are even closely linked to unanticipated inflation." This is a question about averages, that is, macro data. The author has averaged his micro wage change data. This is neither better nor worse than using the usual macro data except that it is limited to 21 annual data points. We all agree that, in the long run, the real wage is neutral with respect to inflation. But the short to medium run matters, and 21 data points are not enough to answer the question. I'm not saying that the answer is wrong, just that you cannot get a definitive answer this way.

I decided to do a very small amount of sensitivity analysis. The numbers in Table 1, 0.84 and 0.88, are the two numbers I think the author wanted to emphasize. Regressing wage on price inflation results in a coefficient that is marginally statistically different from one. The coefficient also says that 7 percent inflation lowers the real wage about 1 percent. I guess that's a small effect—but I'd like to know if a sustained 7 percent inflation would lower the real wage 1 percent every year. We then split the effect into anticipated and unanticipated and see that the anticipated effect is somewhat smaller.

Eschewing the daunting task of trying to replicate the micro data, I chose the first likely looking variable from the DRI (Data Resources, Inc.) database. The left-most panel of Table 1 shows McLaughlin's numbers. The middle panel replicates his regressions using my data to demonstrate that the data differences are not important. In the right-hand panel, I augmented the author's specification with very simple dynamics in the form of a Koyck lag. With the augmented specification, 7 percent inflation lowers the real wage by 2.3 percent, which is a lot. Even using anticipated inflation, the coefficient *is* statistically below one and economically is really far from one. This does not mean the right panel is better than the left; it just means that 21 annual data points are not the right way to answer this question.

Are wage changes skewed? Absolutely! Some people have taken this as evidence that the lower tail is truncated. Of course, while truncation may imply skewness, skewness need not imply truncation. The author did something clever; he looked for skewness far from zero and found it. So while there is some truncation, there are other factors producing skewness. So we

Table 1

Regression of Wage Inflation on Price Inflation

Variable	Author Inflation Process		Macro Data Inflation Process		Macro Data Inflation Process	
	ARIMA(0,1,1)	AR(3)	ARIMA(0,1,1)	AR(3)	ARIMA(0,1,1)	AR(3)
Inflation	.840 (.103)		.865 (.083)		.666 (.123)	
Anticipated Inflation	.880 (.101)	.928 (.113)	.876 (.082)	.930 (.098)	.722 (.139)	.719 (.186)
Unanticipated Inflation	.584 (.176)	.592 (.805)	.678 (.161)	.710 (.152)	.554 (.191)	.625 (.172)
Lagged-Wage Inflation					.284 (.139)	.183 (.174)	.238 (.186)
AR(1)					-.060 (.260)	.169 (.331)	.026 (.333)

Notes on macro data:
price = GDP price deflator (DRI GDNFPC)
wage = Hourly compensation nonfinancial corp. (DRI LCPB)

should be careful about interpreting skewness *per se* as evidence of wage rigidity.

Is this skewness correlated with inflation? The author's Table 4 says no. If skewness results from wage rigidity and inflation reduces wage rigidity, then inflation should reduce skewness. Strikingly, there is no evidence for inflation reducing skewness. I wonder if there is a way to estimate the power of these tests, in the economic rather than statistical sense. One suggestion is to write a simple model of nominal wage setting calibrated twice—once for parameters where all would agree that nominal rigidity is important and one with the opposite assumption. Then one would generate simulated data from each, compute the statistics in Table 4, and determine the extent to which the two simulation results differ.

Is there a big spike at zero? There is a spike, but what metric tells us whether the spike is large? In addition to the size of the spike, the author discussed the amount of censoring. It would be useful to have a table or some guided comparison of the zero spike and estimated censoring against inflation to see if it tells the same story.

Let me turn to the question of the variance of wage changes rather than the skewness. Take an example from life. I wear two hats. I spend 80 percent of my time as an economist and another 80 percent as department chair. I assign wage changes in my department, so for the University of Washington I know what the process is. The University of Washington faculty is about a third the size of the PSID sample. Straight-time wages are downward rigid. By this I mean there is no place on the form I complete to lower wages. It literally cannot happen. And I hazard the same thing is true for half the attendees of the October meeting.

When I give raises I have a dollar-budget constraint as well as a non-negativity constraint. So downward rigidity cannot have any effect on the average wage change. The histogram in Figure 2 gives a fictional, but I think accurate, picture of what happens at different inflation rates. At a 10 percent average increase, there are maybe two raises at 6 percent to one raise at 18 percent. At a 4 percent average increase, there might be two raises at 2 percent to one raise at 8 percent. At high inflation, there's a 12 percent change in relative real

Figure 2

Illustrative (but Fictitious) Histogram of Wage Increases in High vs. Low Inflation Regimes

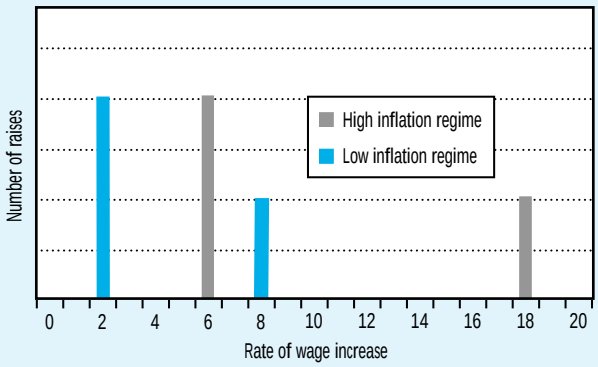

wages. At 4 percent inflation there is only a 6 percent change in relative real wages. The former correctly reflects productivity. Thus, inflation does give us more efficient wage setting, but lower inflation in our case squeezes both tails. I'm not sure the squeezing of both tails would show up in any of the measures used in this paper. Further research might include comparing data from institutions with known nominal wage rigidity with data from institutions known to have flexible wages.

During the discussion at the October meeting, Bill Poole offered a telling objection to my illustrative histogram. Some people are denied tenure; effectively, their wage change is –100 percent! Much of the wage rigidity literature carefully looks only at "nonmovers." If we think carefully about the role of wages in labor markets, we know that labor force adjustments occur on both the intensive (hours) and extensive (hire/fire) margins. McLaughlin's paper focuses on nonmovers, as it should; but perhaps there is more work to be done on the linkage between wage rigidity and the intensive versus extensive margin of labor force adjustment.

In summary, this is a stimulating paper that leaves the reader begging for more. It also is a great example of narrowly focused scientific work that is enormously informative on a fundamental policy question; in this case, "How much inflation should we have?"