

Wage outcomes and macroeconomic conditions: what's the connection? Jared Bernstein, Center on Budget and Policy Priorities¹

This version: 8/16/16

Probably because I've been working for years on the relationship between full employment and various economic outcomes, the organizers of this conference asked me to present my thoughts on wages and macroeconomic conditions. What follows is a collection of the findings I and others have generated in this area of inquiry, designed to both show a set of relevant empirical relationships and engender a conversation with this uniquely knowledgeable group.

My main finding is that the pervasive absence of full employment, or, if you prefer, the excess of labor market slack that's prevailed in recent decades, has been an important source of both wage stagnation and inequality. In an open economy with very little in the way of unions and collective bargaining, full employment is one of the only ways many middle- and low-wage workers muster bargaining power. Only in tight job markets must employers bid up compensation to get and keep the workers they need.

To be clear, slack is but one explanation, and I do not discount the many other wage determinants—employers' technology-induced skill biases, expanded trade with low-wage countries (a development I discuss below), eroding labor standards (minimum wage, overtime), and others. But it is curious to me that given their importance, macro-conditions get relatively little attention in wage trend debates relative to these other determinants. Reweighting the debate such that slack plays more prominent in policy discussions among economists concerned about wage stagnation and inequality is another goal of this work.

The paper/presentation proceeds with a brief discussion of theories of wage determination, followed by a historical review of the extent of slack in US labor markets. I then use the variation in slack over time and across states to generate statistical evidence linking macro-conditions and wage growth, along with the growth of hours worked per week and weeks worked per year. Next, I discuss recent work on the morphing tradeoff between slack and price growth. The last section discusses policies to help promote full employment, taken from work done on behalf of the Center on Budget and Policy Priorities Full Employment Project.

Many theories of wage formation assume full employment. That's a problem.

Most of what follows is empirical evidence of the influence of labor market conditions on wage growth. But I wanted to motivate that information with a short conversation about theories of wage formation. My point is not at all to provide an exhaustive survey but instead to divide theories into two groups: those that assume full employment and those that do not. Evidence I present in the subsequent section suggests that theories of wage determination that assume full employment may be of limited use in understanding US wage trends.

¹ I thank Sandy Black, Jesse Rothstein, Ben Spielberg, Chad Stone, and David Wilcox for helpful comments. Mistakes are my own. Thanks also to Keith Bentele for preparing ASEC data.

Perhaps the dominant theory in this space (at least in intro textbooks) maintains that profit-maximizing firms hire workers up to the point where their additional contribution to the firm's revenue fails to cover their cost – that is, up to the point where the marginal revenue product of the last worker hired is equal to their wage. To hire beyond that point would generate an unnecessary cost to the firm; to hire below that point would leave money on the table (as the firm's technology and market share could profitably absorb more production).

Marginal product is a microeconomic concept in that it refers to the wage formation process at the individual or firm level. In what is perhaps the dominant macroeconomic theory—the neoclassical growth model, which also assumes full employment—aggregate productivity plays a central role in wage levels and growth. In this theory, average compensation grows at the rate of (total factor) productivity, which itself is a function of the interaction of capital (e.g., equipment, structures, hardware, and software), technology, and worker skills. Things that raise the level of productivity, which could be smarter workers or innovations that generate more output per hour (i.e., productivity), will raise the average level of compensation.

While this theory has some empirical support—there are clearly time periods when average compensation grew at the rate of productivity—it too has shortcomings. First, it is mathematically the case that when compensation grows at the rate of labor productivity, compensation as a share of national income will remain constant. However, it is widely recognized that in recent years this has not been the case; the compensation share of national income declined fairly sharply. More recently, the compensation share has been rising again, suggesting a cyclical component which again contradicts the full employment assumption.

In fact, the inability of both of these wage theories to explain cyclical changes is, as I'll show, empirically problematic, as both nominal and real wages are often cyclical; empirically, their trajectories can be to a degree predicted by slack variables. Another limitation of the macro theory from the perspective of this paper is that it is limited to the average, whereas a central focus of what follows seeks to understand connections between macro-conditions and wages *throughout* the pay scale. Importantly, lower wages tend to be more sensitive to the business cycle than higher wages. This reality, which I argue below relates to differences in bargaining clout across the wage scale, is also problematic for the theories discussed so far, as they assume away such dynamics.

Based on the pervasiveness of labor market slack in recent decades, models that do not assume full employment—i.e., the allow for slack—and/or provide a role for bargaining power are more likely to describe actual wage outcomes across the pay scale. In a later section, I estimate “wage curve” models that explicitly link changes in labor market slack to nominal and real wage outcomes (the former is the Phillips Wage Curve and is another common model active use today; I devote considerable attention to this model below).

“Search models” of wage formation are also useful, as they provide an explicit role for bargaining power. When potential workers and employers bargain over the wage offer, the outcome of the deal is partially a function of their “threat points,” or outside options that give them either more or less room to maximize their position in the bargaining process. In this model, an employer hiring in a slack economy/job market faces little in the way of

unmet product demand and has many choices of applicants, leaving them ample time to “shop around” for the best worker at the lowest price (wage).

There’s also the “efficiency wage theory,” which maintains that, under certain conditions, employers will adjust pay above that of a worker’s outside options, given their skill level and experience. The reason for the above-market wage is to make it costly for workers to lose their job, thereby providing incentive to increase the worker’s effort or, by increasing their allegiance to the firm, reduce turnover and thus avoid losing sunk costs associated with hiring and training. However, in a slack labor market, workers will already fear losing their job (due to the increased difficulty of obtaining another one), and, as a result, firms would not have to pay higher wages to generate the same benefits of the efficiency wage model.

Another model that is consistent with the persistence of slack is hysteresis. The theory starts from a simple growth model discussed above where the economy’s potential growth rate is a function of factors such as productivity, labor supply and capital deepening. But it adds an important wrinkle: in this model, persistent slack can, in and of itself, reduce these supply-side quantities, and thus permanently lower productivity, labor supply, capital stock, and, through the productivity channel, wages. Essentially, hysteresis argues that cyclical shortfalls that persist can become structural, through channels like long-term unemployment that deteriorates workers’ skills and labor market connections or weak capital investment that permanently lowers productivity growth and thus wage growth. Empirical studies suggest these dynamics are operating in various economies today.²

Finally, “institutionalist” theories of wage formation posit that entrenched societal institutions, laws, and norms play a key role in how earnings are determined and distributed. Moreover, these institutionally determined outcomes have less to do with marginal product than any of the theories above would dictate. Unions, political power, the ideology of policy makers from Congress to the Federal Reserve, the setting and enforcement of labor standards (minimum wages, overtime rules, workplace safety), immigration practices—all of these are determinant forces in wage formation.

Institutionalist theory interacts with my main theme—the role of labor market slack in wage formation—in an interesting way. The theory itself says little about slack in wage formation, focusing instead on role of institutions like unions and labor standards regardless of the cycle. But there’s a linkage worth considering: because the institutions supporting workers have eroded, tight labor markets are an essential way to boost worker bargaining power in the absence of institutional support. Full employment, in this sense, plays a similar role to unions or other labor policies that create pressures for firms to more equitably share profits.

In the next section, I show considerable evidence of the persistent absence of full employment labor markets in recent years. Thus, we might expect that predictions of this last bucket of theories better explain wage trends, particularly among lower-wage workers with diminished bargaining clout.


² See <http://www.nber.org/papers/w20185> and <https://www.imf.org/external/pubs/ft/wp/2015/wp15230.pdf>.

The pervasive absence of full employment

Econometricians have long argued that, to the extent that we can identify a “natural rate” of unemployment—the lowest rate consistent with stable prices—we can only do so within a large confidence interval, where “large” means: too wide to reliably guide policy. In a later section, I’ll show one measure of how that interval has exploded of late. But for now, suspend such disbelief and imagine that the Congressional Budget Office’s historical natural rate series is ballpark accurate. Then we can ask a simple question: how often has the job market been at full employment over the last 60-plus years?

Given the pervasive assumption of full employment in many of our theories and analyses, we might expect full employment conditions to prevail most of the time. But the next two figures show this hasn’t been the case since around 1980. The line chart plots the CBO’s natural rate against the unemployment rate, with blue for periods when the unemployment rate was below the natural rate and red when it was above. The bar chart shows that from the 1950s through the late-1970s, the unemployment rate was at or below the CBO’s natural rate 72 percent of the time. Since then, job markets have been tight by this metric only 29 percent of the time.

Figure 1



Conditional, for now, on our acceptance that the CBO is in the ballpark with their natural rate estimates, this is a serious and wasteful economic problem. That’s a tremendous amount of output “left on the table.” More to the point from the perspective of this presentation, the absence of full employment has negative implications for real wage stagnation and earnings inequality.


The next set of figures show that, during the period when we were mostly at full employment, median wages tracked productivity and Census-based family incomes doubled

for low, middle, and high-wage families. Of course, these macro-dynamics are one of many factors in play here, and I want to be careful not to suggest that the whole story of wage stagnation is a function of the absence of full employment. But neither do I think these patterns are coincidental.


Figures 2 & 3

The “Wedge” Between Growth and Productivity

Cumulative Percent Change Since 1948



Source: Economic Policy Institute analysis of Bureau of Labor Statistics and Bureau of Economic Analysis National Income and Product Accounts data.



As some may share my skepticism regarding our ability to nail down a point estimate of a natural rate—even a moving one—it might be useful to consider another measure of slack.


Economist Andy Levin has derived a measure that I’ve found to be useful in predicting wage movements (as shown below), so it is worth considering in this context. Levin’s “total employment gap” has three components: the gap between the unemployment rate and CBO’s natural rate estimate (the unemployment gap); the gap between the actual and (CBO’s) potential labor force, which is intended to capture workers who have stopped looking for jobs but might start again (the participation gap)³; and the gap between the trend number of part-time workers who would rather have a full-time job and the actual number of workers in that category (normalized into full-time equivalents and called the underemployment gap).

As the next figure shows, for the most part, the Levin measure follows the more common $u-u^*$ gap (unemployment minus the natural rate). That makes sense, as CBO’s natural rate is embedded in the Levin measure. But the recent difference, circled in the figure, is notable and important, as it represents both the participation gap and the still-elevated share of the

³ Labor economists have hotly debated the magnitude of the labor force participation gap. Shift-share analysis I’ve included in appendix A suggest that of the approximately 3 percentage point decline in the LFPR between 2007q4 and 2016q2, about 2 points are due to aging leaving 1 point to other factors, including weak demand.

employed who are involuntary part-timers, as described above. Meanwhile, the CBO natural rate registers the job market as having hit full employment ($u-u^*=0$).


Figure 4



As shown in the next figure—the first in my presentation to directly link less slack to faster wage growth—the underemployment dynamic captured by Levin’s measure puts downward pressure on wage growth. The figure shows a principal-components-derived composite wage series “mash-up,” which includes quarterly ECI wage and compensation data, compensation from the BLS productivity series, median earnings for full-time workers, and the production, non-supervisory wage from the payroll survey. The forecast is based on regressing year-over-year quarterly changes on the Levin slack variable with four lags, and a lag in the dependent variable.

Figure 5

Simple Slack Model of Nominal Wage Growth with Full Employment Simulations



The circle at the end of the figure shows two forecasts. The lower forecast allows the Levin measure to fall at its recent rate, hitting zero – full employment by this metric – in the second half of 2017. At that point nominal wage growth settles at around 3.5%, which notably happens to be the wage target set by Fed chair Janet Yellen, i.e., the wage-growth rate she views as consistent with stable inflation. The higher forecast path allows the Levin variable to fall to 1 percentage point below full employment, and this ends up generating nominal wage growth of about 4%.

In sum, using a conventional and widely accepted measure of full employment, slack has largely been the norm in recent decades. Moreover, a simple model using a more comprehensive measure of slack suggests that achieving full employment would quickly raise nominal yearly wage growth to its Fed target of 3.5%. Yet there is also more convincing evidence of the impact of slack on real wage trends, and not just at the average, but across the wage scale.

Yes, there's an hourly wage curve and its slope varies by wage decile. There's also an hours and weeks curve.

This section presents results from two data series on the relationship between real earnings and labor market slack. First, I use a panel data set with state observations on real hourly wages by decile since 1979 (I use the CPI-U-RS to deflate wages). Second, I turn to national time series data from the ASEC CPS files to model the impact of slack measures on hours and weeks worked.

All of the regressions in this section break the data down to quintiles, thirds, or deciles, thus enabling me to identify different impacts of slack on dependent variables across the wage or

hours-worked scale. The main finding from these models is a persistent gradient showing that diminished slack disproportionately benefits less well-off workers or households. In fact, achieving or even moving towards full employment does little for higher wage earners or wealthy households. This is not unexpected, as their unemployment rates are known to be lower than average with little variance, such that their wages and hours are not particularly responsive to the business cycle.

The first set of results is derived from a panel of 50 states (and DC) since 1979, so I have more than 1,800 observations. I run fixed effects models with robust standard errors, clustered by state, using real wage deciles by state from the Economic Policy Institute and BLS information on unemployment, employment rates, and underemployment (aka “U-6” in the monthly employment report; available only since 1994).


The regressions are very simple: they regress the log real hourly wage on the slack measure (logged and lagged one year), along with fixed effects. Their most notable aspect is the gradient of the elasticities, moving from relatively high and always statistically significant coefficients for low-wage workers to relatively small and insignificant coefficients for high-wage workers.

The next figure plots the coefficients on unemployment and employment rates (“EPOPs”) times one standard deviation in the independent variable. The aforementioned gradient is clear and the magnitude of the coefficients is economically meaningful.⁴ For example, as the US job market moved to full employment over the 1990s, the jobless rate fell from 7.5% in 1992 to 4% in 2000. Over the period, 20th percentile real wages grew 10% and median real wages grew 4%, implying that about 70% of each increase is explained by the unemployment decline.

Figure 6


⁴ For unemployment, coefficients from the 6th decile up are insignificant; for EPOPs, the top two deciles’ coefficients are insignificant (at the <0.05 level).

Percent change in real hourly wage associated with a one-standard-deviation increase in a given slack variable (1979-2015 panel data)



The next figure shows the same gradient for state underemployment rates. While these elasticities are somewhat larger in absolute value than those in the last chart, that is mainly because the sample is restricted, since we only have underemployment data beginning in 1994. Unemployment coefficients, for example, look much the same as these when I restrict the sample to 1994-2015.

Figure 7



Full employment, by definition, also boosts the supply side of the labor market, increasing weeks worked per year and hours worked per week. But does the same gradient exist for those variables? Using annual, national data on single-mother households from the March CPS, I show evidence of the same gradient – i.e., more evidence of the equalizing impact of full employment.


The table below shows coefficients and t-stats from simple, time-series regressions of hours worked per week and weeks worked per year, respectively, on the unemployment rate and the lagged dependent variable (all in logs) for single mom households. I focus on the bottom third and top third, by income. The coefficients for the bottom-third group are 4 to 5 times those for the top group (though “ceiling effects” for the latter group are always in play in this sort of comparison). The next figure shows how well one simple variable—unemployment—explains the fairly complex phenomenon of annual hours worked.⁵

Table 1

⁵ Note that this is a “dynamic” forecast, meaning the lagged dependent variable (annual hours, the product of hours/week and weeks/year) is predicted for each year.


Single Mom Households			
		coefficient	t-statistic
Hours	Bottom 3rd	-0.180	-5.438
	Top 3rd	-0.032	-2.264
Weeks	Bottom 3rd	-0.193	-4.002
	Top 3rd	-0.045	-3.425

Figure 8



The forecast at the end of the series shows what happens to annual hours worked by low-income (i.e., bottom third of the income scale) single-mom households when I let the jobless rate decline to 4.5, 4, and 3.5 percent in 2016, 2017, and 2018, respectively. Finally, in the next figure, I hold these households' hourly wage constant at its observed 2015 level (about \$10/hour) and, using 2015 dollars, multiply that wage by predicted annual hours. This adds \$3,700 to earnings between 2015 and 2018, about \$1,200 per year.

Figure 9



In one sense, that's a low-ball estimate; I'm holding the hourly wage constant despite the fact that earlier regressions predict that the wage would grow. On the other hand, it is likely the Fed would push back on unemployment as low as I've assumed.⁶


Since, along with the tightening job market, welfare reform was in the mix regarding single-mom's hours and weeks worked, I run the same regressions are forecasts for low and high-income African-American households. The results, shown below, are much the same.

Table 2

⁶ Another interesting wrinkle to consider here is the extent to which higher earnings for these single mom families would mean the gain or loss of certain benefits conditioned on earnings.

African American Households			
		coefficient	t-statistic
Hours	Bottom 5th	-0.243	-4.738
	Top 5th	-0.016	-1.184
Weeks	Bottom 5th	-0.261	-3.861
	Top 5th	-0.018	-1.420


Figure 10



A final figure in this section shows another simple use of these types of findings, in this case, one focused on low-wage black workers. Suppose the FOMC saw this presentation and decided that running a “hot” job market would help minorities gain more employment, wages, and bargaining clout (the next section deals with the inflationary implications of their

actions). So they let the jobless rate fall a point below their current natural rate estimate of about 5 percent. That would be expected to lower the corresponding black rate by about 1.7 percentage points, down to around 7 percent, which would be the lowest rate since the last time the job market was at truly full employment in the late 1990s. Mapping that onto my model's predictions for their 20th percentile real wage generates the pattern in the figure, an increase of almost 3 percent in their real wage.

Figure 11



These are all admittedly simple regressions and simulations; variables that also impact wage outcomes, like education, demographics, and tenure, are omitted. But the results certainly underscore my contention about the critical role of slack in both wage stagnation and inequality. The absence of full employment shown in the earlier section is implicated in these deep problems. Moreover, as this particularly fractious election season progresses, these wage issues are no longer solely a matter of living standards; they are now playing a potent and potentially explosive role in our politics.

How does one incorporate the evidence that the Phillips Curve may be out-of-order into this work?

The evidence presented thus far broadly shows the benefits of favorable macro-conditions, specifically tight labor markets, to not just wage growth in general, but to wage and hours growth for the least advantaged in particular. Since wage stagnation and the patterns of inequality have generally been toughest on those with lower wages and incomes, this evidence suggests that economists and policy makers concerned about these issues should aggressively push for full employment. Put differently, the fact that periods of full


employment have been the exception in recent decades is a major economic problem with deep political repercussions.

Among some economists, this setup invokes concerns about the tradeoff between inflation and slack, often represented as the price Phillips Curve (PC). Without wanting to stray too far into a hotly contested area of contemporary macro, this section summarizes why one might be skeptical that the alleged tradeoff should be invoked to slow or halt progress towards getting to and staying at full employment.

In fact, as I'll show in a moment, by some traditional measures the PC for price growth (vs. wage growth) is very hard to accurately identify (though other analysts have found specifications and samples that do identify the curve). To the extent that it is difficult to identify the PC, it is also difficult to identify the so-called "natural rate" of unemployment.

The first figure in this section shows that as unemployment has fallen by more than half, inflation has been quiescent, even as nominal wage growth has slightly accelerated. By the end of the figure, the actual unemployment rate is coincident with the Fed's natural rate estimate (which has come down a bit), yet core PCE inflation is flat (the overall PCE has showed much slower inflation since 2015 due to falling energy prices), as are [inflationary expectations](#) (which remain, as the Fed puts it, "well-anchored"). The facts that inflation failed to fall sharply when slack was high or rise much when slack was low suggests a flat PC.


Figure 12



A more formal look at these issues comes from the Economic Report of the President, written by the Council of Economic Advisers. The ERP tends to be an important and influential book, at least in DC circles. So when this year's ERP highlighted the breakdown of the PC and the NAIRU (the non-accelerating inflation rate of unemployment), many who follow this sort of thing took notice.

The CEA shared their data with me and I made the next two figures to give you a look at rough correlations since 1990. The first shows a relatively flat negative regression line fitted between core CPI inflation minus inflationary expectations and a fixed-weight unemployment rate.⁷ The second figure replaces price inflation with wage inflation, using the same "mash-up" series used elsewhere in this presentation (a composite of five wage series extracted using the first principal component). The slope is considerably steeper, which is really just a variation of the finding in Figure 5 above.


Figures 13/14



The next three figures are right out of the CEA's 2016 ERP. The first figure shows the inflation variance explained by the unemployment rate over subsequent 20-year time periods. Peak correlation occurred in the early 1990s, but the relationship has weakened since then, and the end of the figure shows that unemployment explains almost none of the variation in inflation in recent years.


Figure 15

⁷ CEA uses q4/q4 changes. Inflationary expectations are the lagged values through 2007 and the FOMC's expectation after that. See 2016 ERP, box 2-9 for more details.




This is again consistent with a very flat PC, as shown in the next figure, which plots the coefficient on unemployment from the same set of 20-year rolling regressions. Over the full period, the elasticity was significant: a one point increase in unemployment led inflation to fall -0.4 percentage points. But the elasticity is far from constant, and by the end of the period, it too is just about zero.

Figure 16



Since the NAIRU, or natural rate, is the lowest unemployment rate consistent with stable prices, if we set the left side of the formula in the figure to zero (and assign zero to the random error term), the NAIRU is backed out of this equation as α/β . That's plotted in the third figure, underscoring the point I've made throughout this presentation: given the confidence interval CEA computes towards the end of the figure, estimating a natural rate that could reliably serve as a policy guidepost right now may be beyond economists' capabilities.

Figure 17



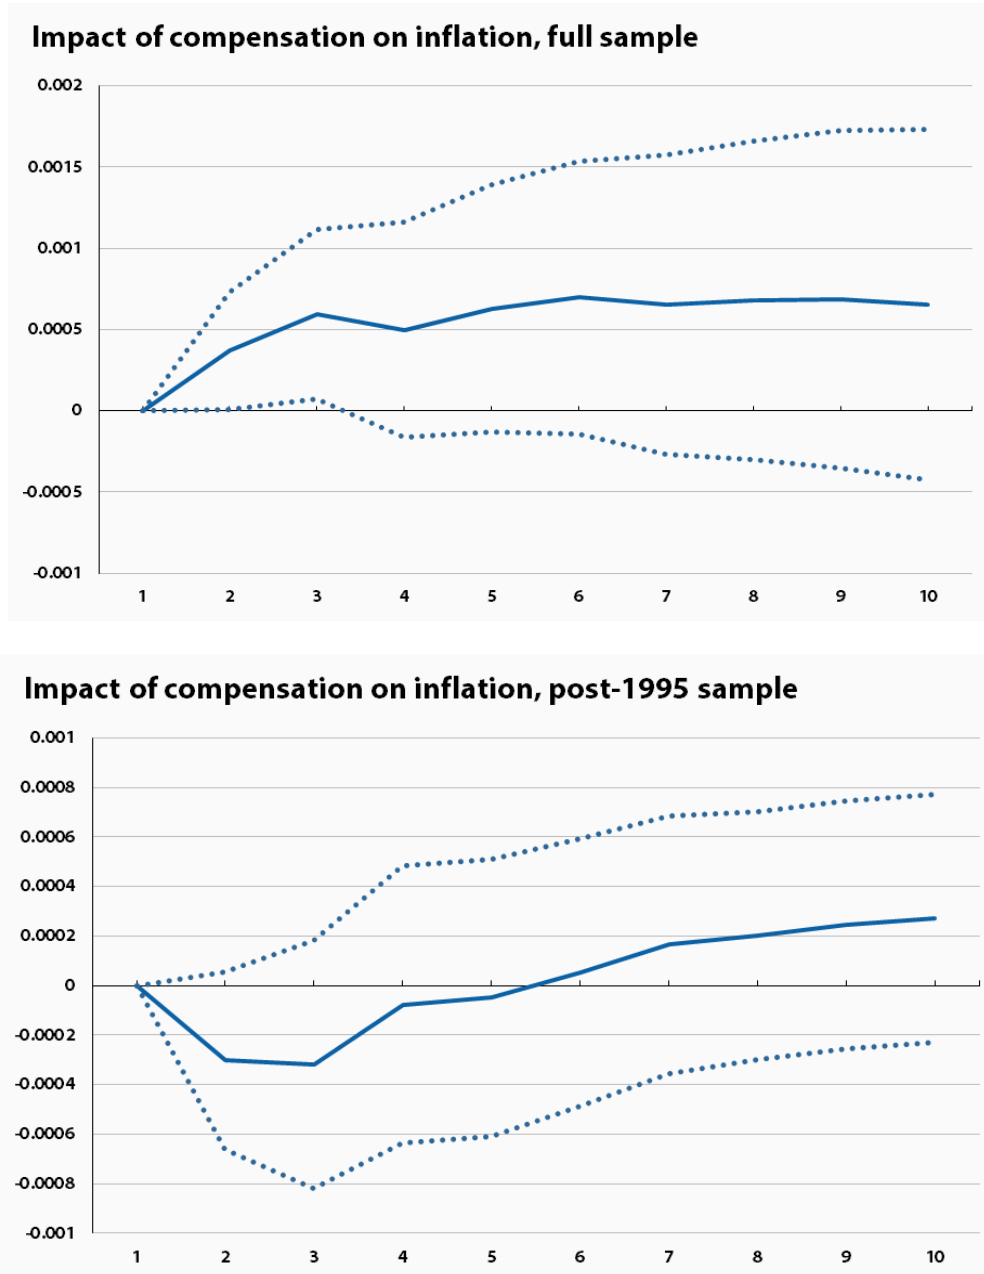
Many of us learned to be wary of such guideposts back in the 1990s, when economists thought the lowest you could go on unemployment was around 6 percent. But Fed chair Alan Greenspan recognized other moving parts in the economy, most importantly faster productivity growth, that meant unemployment could fall below the supposed natural rate without juicing inflation. He was right: unemployment was 4 percent in 2000, and inflation was well-behaved.

But do these insights create a new problem? Slack may be less correlated with faster inflation, but it's still correlated with wages. So don't we still have to worry about "wage-push" inflation?


In fact, evidence of wage growth bleeding into price growth is increasingly hard to find in the data. Economists at the Cleveland Fed found that adding wage growth to inflation forecasting equations was of little help in improving their [forecasts](#). Other Fed economists used time-varying models to arrive at a [similar finding](#): even controlling for productivity growth, evidence for wage-driven inflation is elusive.

The next few figures show the [results](#) of vector autoregressions that include prices and wages, along with the trade-weighted dollar index as a control. They show no pass-through, even when I control for productivity growth by using unit labor costs as my compensation variable.

Figures 18-20



Impact of ULCs on inflation, full sample



Some analysts argue that the PC is alive but hiding; you’ve just got to know where to look for it. [Blanchard et al.](#) look across countries and find a stable but historically low correlation between slack and price growth. Larry Ball [identifies](#) a PC coefficient using median core inflation and smoothed capacity measures. Goldman Sachs economists present compelling evidence that a PC can be well-identified using a panel data of metro-areas over time, as cities with low unemployment have seen inflation accelerate (they also show, however, that this is largely a function of housing price movements).

This is not the place to adjudicate this issue, but the CEA results cannot be ignored, especially as it is working off a very standard definition of the PC. Fed Chair Janet Yellen herself, who presumably used to hang around these parts, recently [argued](#)—wisely, I’d contend—that because of dynamics of the sort shown in Figure 12, “significant uncertainty attaches to Phillips curve predictions, and the validity of forecasts from this model must be continuously evaluated in response to incoming data.”

Putting it all together, my argument is that the results of the last two sections—1) the general absence of full employment, and 2) the gradient of the coefficient on the hourly wage deciles, a similar gradient for hours and weeks worked, and the flat PC by some measures—points toward the benefits of pursuing a full employment agenda that pushes back on real wage stagnation and wage inequality.

Getting to full employment

Much of the above is, if not well known, then at least none too surprising. Assuming you’re with me, the findings beg the question: how do we get to and stay at full employment? While a deep policy dive is beyond the scope of this presentation, I will make a few points on which attendees may want to follow up.

First, CBPP’s Full Employment Project has commissioned around 20 papers so far to address that question. I suspect many of you would enjoy [perusing the website](#). I’d also

welcome any pitches for new papers on policies to help us get to and stay at full employment.

Second, I will note a few findings from recent papers that may pique your interest.

--*The trade deficit as a barrier to full employment:* Since the trade deficit is by definition a drag on GDP growth, all else equal, it potentially contributes to output gaps and economic slack (see Dean Baker's [paper](#)). But all else is far from equal, and while we've run historically large trade deficits consistently since the mid-1970s, we've also hit full employment (or close to it) in some of those periods. In 2007, for example, unemployment was 4.6 percent while the trade deficit was a historically large -5 percent of GDP. So it is a mistake to assume that trade deficits preclude full employment.

But there are three important factors to consider. First, one of the main factors driving down unemployment in the 2000s was the housing bubble, which, as Ben Bernanke argued [in 2005](#), related to important international trade and financial dynamics. Certain trade competitors suppressed consumption, boosted savings, and managed their currencies to export savings and import labor demand. Their current account surpluses mirrored our current account deficits. As noted, we handily offset the demand loss, but through a housing bubble inflated by robust capital flows (leavened by inadequate regulatory oversight). Mercantilism and currency manipulation thus played a highly destructive macroeconomic role.

Second, when interest rates are already very low across the globe, as they are now, and when the Fed is at the zero lower bound, these dynamics are much harder to offset. Capital, starved for even a mildly positive return, quickly books a ticket for wherever yields are safest and highest. In normal times, with strong demand and high "animal spirits" among investors, such incoming capital gets put to use—hopefully good, productive use—by financing domestic investment, jobs and growth that offset the economic drag created by the larger trade deficit. But what if, as has been the case in our economy even before the recent apparent slowdown in job creation, demand is none too strong and investors have little use for extra capital?

If the Fed wasn't stuck near zero, it could lower rates until the price of investing was low enough to meet the weakened demand. But when they're stuck at zero, we're stuck with an imbalance that has to come out somewhere, and if the excess capital can't be absorbed through lower interest rates, it gets absorbed through bigger trade deficits and weaker growth.

Moreover, when low global rates block other economies from making similar, needed adjustments, we're into a zero sum game where everyone's trying to steal demand and jobs from everyone else, trying to drive down their exchange rates and export their way out of the slump. But at the end of the day, one country's exports are another country's imports. The solution over the past couple of business cycles has been for the United States to be the world's "growth engine." In recent years, that's meant demand-zapping trade deficits and large capital flows inflating one bubble after the next. [New work](#) by Eggertsson et al. underscores these dynamics in the current interest rate context.

Third, as Autor et al [have shown](#), which countries we have a trade deficit with matters. Deficits with low-wage countries trigger “factor-price equalization impacts,” which Josh Bivens [shows](#) spills over to the earnings of non-college educated workers in the non-traded sector as well.

Pushing back on currency manipulation through enforcement rules in trade agreements – or at least other policies outside of trade agreements, like countervailing duties or reciprocity in capital purchases (if the Chinese can buy dollars, we can buy yuan) – is necessary and complementary to achieving full employment and avoiding the economic “shampoo cycle:” bubble, bust, repeat.

--Monetary and fiscal policy: These are the main tools at policy makers’ disposal, and while monetary policy has arguably aggressively pushed for full employment in recent years, its impact is diminished by the lower bound problem noted above. Meanwhile, political gridlock has blocked fiscal policy that would help hasten the arrival of full employment. A paper by [DeLong et al.](#) for our project argued that with low borrowing costs and “hysteresis effects” (wherein persistent cyclical weakness does permanent, structural damage to growth inputs), fiscal policy can come close to paying for itself (relative to a counterfactual of hysteresis-induced lower growth).

In addition, while prominent monetary policy authorities, notably Chair Yellen, have been admirably careful and data driven regarding the tightness of the labor market (looking well beyond just the unemployment rate), it is also the case that traditional Fed guideposts have been uniquely unhelpful in recent years. In this regard, their current campaign toward “normalization” of the federal funds rate is viewed as misguided by many advocates of truly full employment.

Numerous papers for our project have underscored this last point. One important [contribution](#) from Josh Bivens makes a strong case for nominal wage targets in assessing labor market conditions. Larry Ball makes a compelling [case](#) for monetary policy that pushes the unemployment rate well below the Fed natural rate in the interest of “reverse hysteresis” (i.e., repairing some of the structural damage to labor supply and capital investment). Danny Blanchflower and Andy Levin [call for](#) using benchmarks, like the Levin measure featured above, that do a better job than the unemployment rate in picking up “hidden unemployment.”

--The challenges of deep poverty, isolated neighborhoods, and the criminal justice system: It is important not to conclude that full employment can reach everyone who is underemployed or disconnected from the labor market. Place-based research identifies neighborhoods that remain “job deserts” even at full employment. Those stuck in deep poverty often are beset by a variety of problems and deep skill deficits that create steep barriers to work. Similar barriers confront the millions of people with criminal records trying to get into, or back into, the labor market.

Given the magnitudes of these groups, we are extremely unlikely to get to full employment without direct job creation policies and criminal justice reform. Donna Pavetti [writes](#) about the former for our project, citing the notable success of a highly subsidized employment program in place during the Great Recession. Maurice Emsellem and Jason Ziedenberg

[explore](#) the policies at the intersection of full employment and criminal justice, including rules to “ban the box” from the beginning of the job application process and strategies to reduce the number of people caught up in the criminal justice system in the first place.

Other papers explore the underappreciated role of [apprenticeships](#), [worksharing](#) (a form of unemployment insurance that keeps people on the job but with reduced hours), and labor [standards](#) (like minimum wages, overtime, going after wage theft, and establishing a level playing field for union organizing) targeting job quality, as opposed to job quantity.

Conclusion

The evidence suggests three conclusions: 1) full employment has been the exception as opposed to the norm in recent years; 2) tighter job markets help offset wage stagnation and do so in a progressive manner, meaning they also push back on wage inequality; and 3) the inflationary costs of pursuing full employment may be considerably lower than at some periods in the past.

Moreover, the evidence suggests that macro-conditions matter not just for hourly wages, but for supply-side variables too, notably annual hours worked, and again with a strong, progressive gradient. These findings imply that the benefits of full employment flow disproportionately to the least advantaged, including low-wage workers and minorities.


While a more aggressive pursuit of full employment invokes inflationary concerns for some, I’ve argued that the price Phillips Curve has flattened in recent years, implying a diminished correlation between slack and changes in price growth. There is credible research that disagrees, but even some of this work (e.g., see Ball citations throughout) suggests that a tradeoff between less labor market slack and slightly faster inflation would be useful at both the micro (wage trends) and macro (“reverse hysteresis”) levels.

Finally, though there’s been real progress, it is still the case that many economists concerned with wage determinants pay short shrift to the impact of macro-conditions and the absence of full employment. The tendency is to focus much more on the benefits of education and the related impact of “skill-biased technological change.” Of course, these are very important issues in their own right, but it’s not at all obvious to me that they’re more important than the fact that the job market has been slack for most of the last 35 years!

This insight motivates our Full Employment Project at CBPP, which is dedicated to identifying and elevating policies that can help the economy get to and stay at full employment. I hope the evidence I’ve presented above inspires others here to join us in pursuit of that cause.

Appendix A: Shift-share analysis

Following a recent Federal Reserve analysis,⁸ the figure below uses three different techniques to decompose the decline in LFPR into aging and other (the residual after accounting for aging) factors for four age groups and both genders. The first bar holds each group's LFPR constant in the base year, the second holds population shares constant, and the third shift-share analysis is a hybrid of the other two approaches that holds each group's LFPR constant for only one-month periods at a time.



⁸ <http://www.federalreserve.gov/pubs/feds/2014/201464/201464pap.pdf>