Has the Phillip Curve Died?

New Evidence from Cross-Geographical Data

By Daniel J. Wilson

Summary

- Low wage and price inflation during the current recovery has caused many observers to ask whether the wage and/or price Phillips Curves have flattened in recent years.
- I offer new evidence on this question using cross-geographical Phillips Curve (PC) estimations, based on MSA-level panel data from 1991 to 2015.
    - I investigate both the price and wage PCs.
- The evidence points to a flattening of the cross-geographical price and wage PCs over time, especially during the current recovery.
    - The evidence for flattening is especially strong for the wage PC.
    - The cross-MSA evidence suggests that flattening of the price PC during the current recovery has lowered price inflation pressures by about one-tenth of a percentage point, while the flattening of the wage PC during the current recovery has lowered wage inflation pressures by about two-tenths.
- I discuss the implications of this cross-geographical PC flattening for the aggregate PC.
    - While these slopes can differ due to cross-geographical labor mobility and tradability, these factors together have trended over recent years in such a way as to suggest the aggregate slope has flattened at least as much as the cross-geographical slope.

I. Has the Phillips Curve flattened in recent years?

A. The price Phillips Curve

- As in Leduc & Wilson (2015), I estimate the cross-geographical price Phillips Curve using the following specification:

$$\pi_{it} = \alpha_t + \alpha_i + \rho\pi_{i,t-1} + \beta\left(u_{it} - \bar{u}_{it}\right) + \varepsilon_{i,t} \ , \qquad (1)$$

  where π_{it} denotes the inflation rate in MSA *i* in year *t*, α_t represents time fixed effects, α_i represents MSA fixed effects, u_{it} denotes MSA *i*'s actual rate of unemployment and $\bar{u}_{it}$ denotes MSA *i*'s natural rate of unemployment. The natural rate is allowed to be MSA- and time-varying.

  - I measure MSA-level inflation using the annual percentage change in the CPI price index.
  - I measure the natural rate by a 10-year trailing average of the area's unemployment rate.
  - The year fixed effects will capture aggregate movements in inflation expectations, national or global supply shocks, and other macroeconomic factors (including monetary policy).

- I first update the results of Leduc & Wilson, extending their sample through 2015.[1] The estimated slope – the coefficient on the unemployment gap – for the full 1991-2015 sample period is shown in the first row of column (1) of **Table 1**.

- The estimated full-sample PC slope is -0.23, which is statistically significant at below the 1% level. This slope implies that a one percentage point (p.p.) decline in a MSA's unemployment gap is associated with an increase in its inflation rate of 0.23 p.p.

[1] State-level price inflation data is only available from 2008 onward, so I do not use them here for investigating how the price PC slope has changed over time. Nonetheless, I have also estimated the cross-*state* price PC for the period 2009 – 2015, finding a precisely estimated slope of -0.06, consistent with the hypothesis that the cross-geographical PC is close to flat over this period.

Table 1: Cross-MSA Price and Wage Phillips Curve Slopes
Estimates of $\hat{\beta}$ from $\pi_{i,t} = \alpha_i + \alpha_t + \gamma\pi_{i,t-1} + \beta(u_{i,t} - \bar{u}_{i,t}) + \epsilon_{i,t}$

	Price Inflation	Wage Inflation
1991-2015	-0.234***	-0.347***
	(0.030)	(0.056)
1991-2008	-0.280***	-0.447***
	(0.038)	(0.072)
2009-15	-0.155***	-0.140
	(0.060)	(0.098)

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

- This full sample slope, however, masks consider variation across subperiods, as shown in the second and third rows of **Table 1**. Over 1991 – 2008, the estimated price PC slope in the cross-MSA data is -0.28. Yet, for the current recovery period (2009 – 2015), the slope is significantly flatter, at -0.16. The difference in slopes is statistically significant at the 10% level.
    - To put this flattening in context, note that the national unemployment gap fell by five p.p. between 2009Q4 and 2016Q4, or about 0.7 p.p. per year. A PC slope of -0.28 implies that a drop in a city's unemployment gap of that magnitude should increase its annual inflation rate by 0.20 p.p. By contrast, a slope of -0.16 implies an increase in annual inflation of 0.11 p.p. Hence, this flattening amounts to about one-tenth lower inflation pressure during the current recovery.
    - This relative flattening of the cross-MSA Phillips Curve is robust to alternative measures of slack.[2]
    - The flattening is also broad-based across subcategories of inflation.[3]

[2] **Appendix Table A1** shows results where slack is measured by the current unemployment rate, the lagged unemployment rate, or the lagged unemployment gap. In each case, the price PC slope is flatter over 2009-2015 than over the full sample.
[3] **Appendix Table A2** shows analogous results where inflation is split into subcategories. Five of the six subcategories have a flatter PC slope during the current recovery compared with the full 1991-2015 sample period. The other category, nondurable goods, shows essentially no change in its PC slope.

- To investigate whether the cross-geographical price PC slope has changed over time more generally, I estimate equation (1) using 7-year rolling samples for 1991 – 1997, 1992 – 1998, and so on, through 2009 – 2015.
- The PC slope coefficients ($\hat{\beta}$) and their 90% confidence intervals are shown in **Figure 1**. The evidence on flattening is mixed. The slope was generally stable over time, hovering around -0.2 to -0.3, until recent years. A clear increase in the slope did not occur until very recently, with sample periods ending in 2014 or 2015.
- **Figure 2** shows analogous results using the unemployment rate, instead of the gap, as the measure of slack. Note that this graph starts a few years earlier because the unemployment rate data is available earlier than the gap data.
    - Using the unemployment rate, the evidence of a flattening of the slope in recent years is much stronger.

Figure 1: Slope of Cross-MSA Price Phillips Curve Over Time
Slack Measured by Unemployment Gap

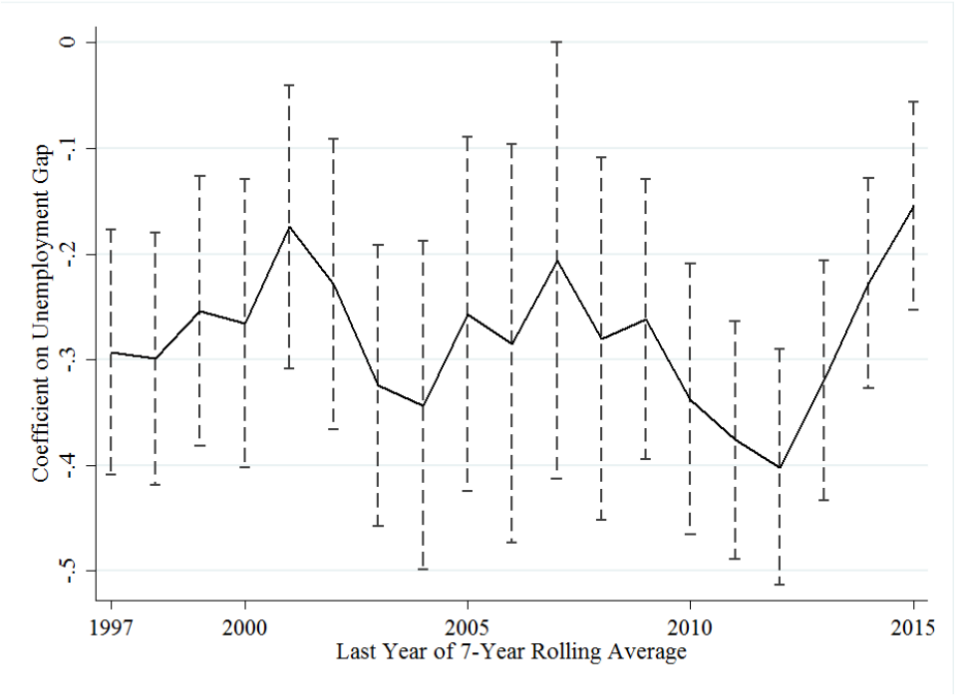

Graphs show the coefficients and 90% confidence interval on the unemployment gap from a regression of CPI inflation on lagged CPI inflation, year fixed effects and geographical fixed effects for the preceding seven years.

Figure 2: Slope of Cross-MSA Price Phillips Curve Over Time
Slack Measured by Unemployment Rate

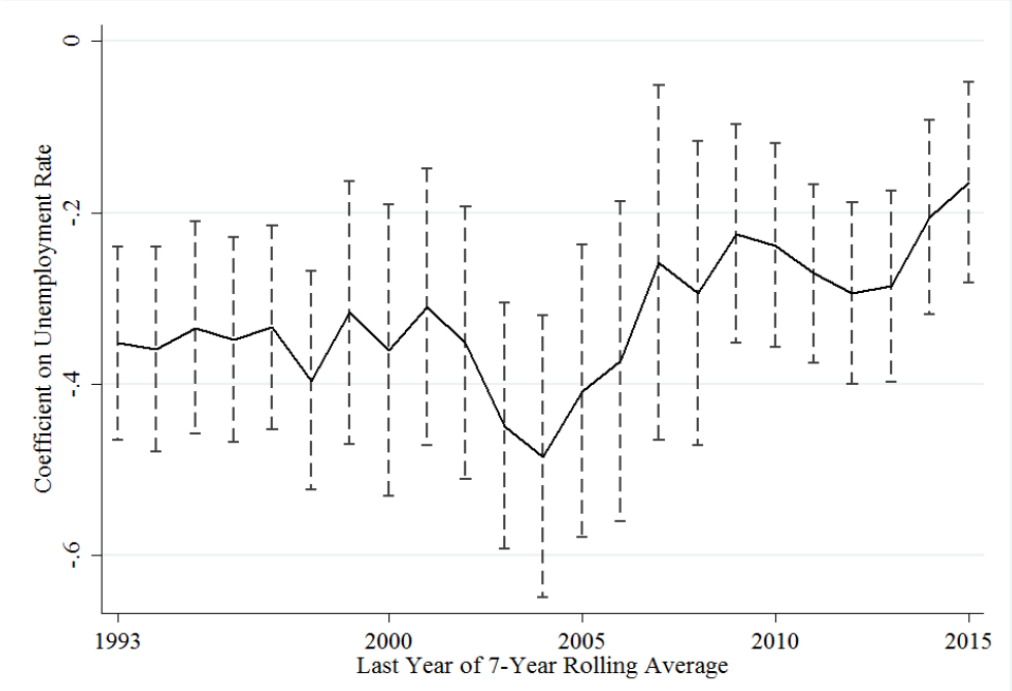

Graphs show the coefficients and 90% confidence interval on the unemployment rate from a regression of CPI inflation on lagged CPI inflation, year fixed effects and geographical fixed effects for the preceding seven years

B. The wage Phillips Curve

- Next I turn to the cross-geographical *wage* Phillips Curve, estimating a standard wage PC specification:

$$\pi_{it}^{w} = \alpha_t + \alpha_i + \rho\pi_{i,t-1} + \beta\left(u_{it} - \bar{u}_{it}\right) + \varepsilon_{it} \ , \qquad (2)$$

  where π_{it}^{w} denotes the wage inflation rate in MSA *i* in year *t*, and other variables are as defined earlier.
    - My primary measure of annual "wages" for a given area is total wage and salary income divided by total employment (wage income per job). The data come from the BEA Regional Economic Accounts.
- As I did for the price PC, I estimate the wage PC (equation (2)) using MSA panel data over the full sample period 1991-2015 and over the subperiods 1991-2008 and 2009-2015. The results are shown in column (2) of **Table 1**.
- For the full sample period, the estimated cross-MSA wage PC slope is -0.35, which is statistically significant at below the 1% level. This slope implies that a one percentage point decrease in a MSA's unemployment gap is associated with an increase in its wage inflation rate of 0.35 percentage point.
- However, as was the case for the price PC results, these full sample results mask consider variation across subperiods. The estimated cross-MSA PC slope over 1991-2008 is -0.45. Over 2009-2015, the slope is just -0.14 and statistically insignificant.
    - To put this in context, a 0.7 p.p. per year drop in the unemployment gap – such as that observed at the national level during the current recovery – implies an increase in annual wage inflation of 0.32 p.p. if the wage PC slope is -0.45 but just 0.10 p.p. if the slope is -0.14. Hence, this flattening amounts to about two-tenths lower wage inflation pressure during the current recovery.
- This evidence of flattening is robust to using alternative measures of slack.
    - **Appendix Table A3** shows results using the current unemployment rate, the lagged unemployment rate, or the lagged unemployment gap. In each case, the slope of the wage PC is statistically significantly negative over the full sample period but it is closer to zero and statistically insignificant over the 2009 - 2015 period.

- As I did for price inflation, I also investigate whether the cross-geographical wage PC slope has changed over time more generally using 7-year rolling-sample regressions of equation (2).
- The wage PC slope coefficients ($\hat{\beta}$) and their 90% confidence intervals, based on MSA data, are shown in **Figure 3**. The analogous figure using the unemployment rate instead of the gap is shown in **Figure 4**.[4] Both figures point to a flattening of the cross-MSA wage PC slope starting around the 2002 - 2008 sample and continuing through the latest (2009 - 2015) sample period.
    - This flattening is quite robust. It is insensitive to omitting time fixed effects, replacing lagged price inflation as a regressor with lagged wage inflation (in the MSA regressions), measuring slack with the unemployment rate instead of the gap (as shown), and down-weighting outliers (via quantile/median regression).
- Evidence of flattening in the wage PC also holds in cross-*state* regressions.
    - Analogous BEA wage inflation data is available back to 1991 for states. **Appendix Figure A1** (based on the unemployment gap) and **A2** (based on the unemployment rate) show the estimated wage PC slope using the state data. One sees a clear flattening of the slope in recent sample periods.
    - At the state level, I also have constructed wages using average hourly earnings (AHE) directly from the CPS micro data. Rolling-sample wage PC slope results based on AHE are shown in **Appendix Figures A3-A4**. Again, one sees strong evidence of flattening in recent periods.

[4] Note that this graph starts a few years earlier because the unemployment rate data is available earlier than the gap data.

Figure 3: Slope of Cross-MSA Wage Phillips Curve Over Time
Slack Measured by Unemployment Gap

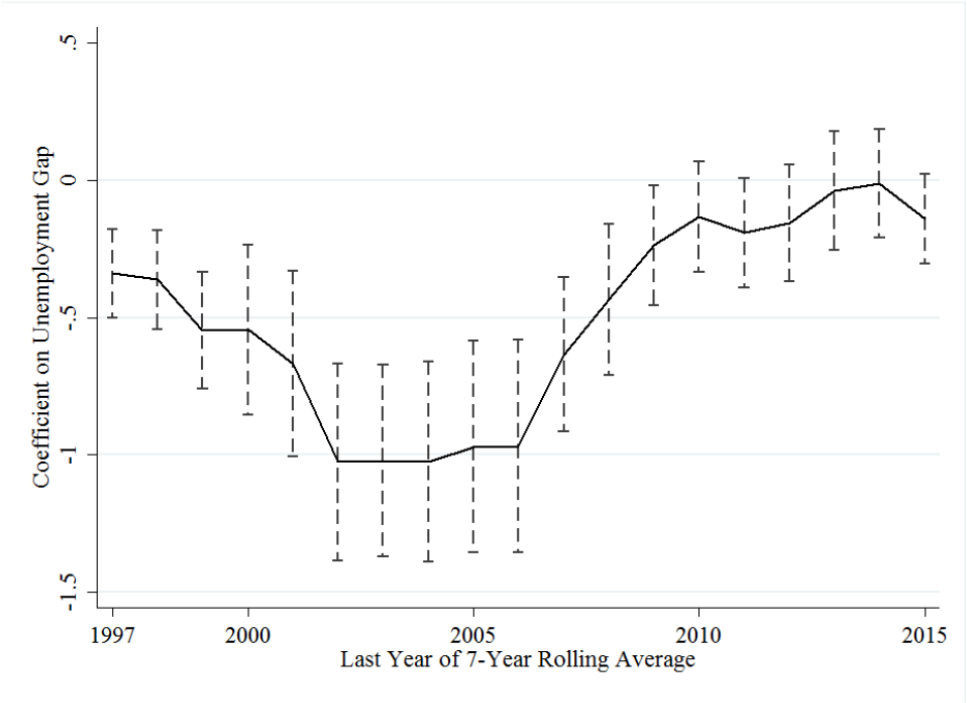

Graphs show the coefficients and 90% confidence interval on the unemployment gap from a regression of wage inflation on lagged CPI inflation, year fixed effects and geographical fixed effects for the preceding seven years.

Figure 4: Slope of Cross-MSA Wage Phillips Curve Over Time
Slack Measured by Unemployment Rate

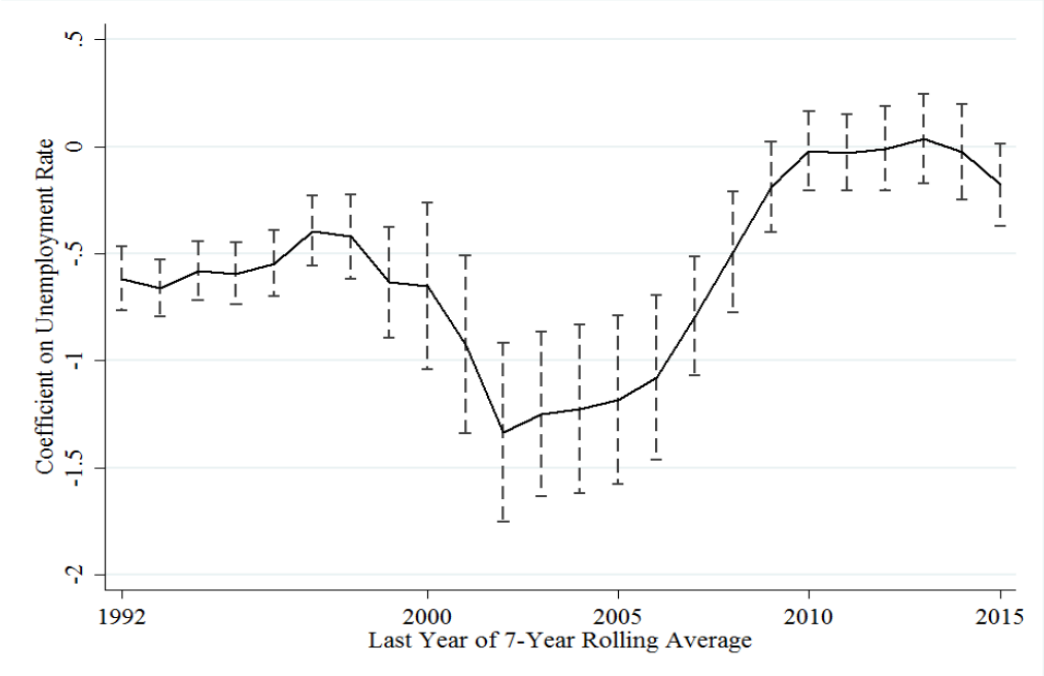

Graphs show the coefficients and 90% confidence interval on the unemployment rate from a regression of wage inflation on lagged CPI inflation, year fixed effects and geographical fixed effects for the preceding seven years

II. Implications for the aggregate price and wage Phillips Curves

A. Does Cross-Geographic PC flattening imply aggregate PC flattening?

- At any given point in time, the cross-geographical Phillips Curve slope need not be equal to the aggregate Phillips Curve slope. There are two key reasons why they might differ:[5,6]
  - **Labor mobility**: Blanchard and Katz (1992) document that interstate migration in the U.S. is a key mechanism by which state economies return to equilibrium after regional booms and busts. In effect, in-migration of job seekers in response to a local boom represents a labor supply shock that helps offset the wage pressures stemming from the initial increase in labor demand. This will tend to mute the effect of labor market tightness on wage inflation relative the effect at the national level, where in-migration of labor from abroad is very limited.
  - **Tradability**: A local "aggregate demand" shock will bid up prices of local non-tradables such as housing, but have little effect on the prices of locally produced tradable goods, the prices of which are set by national and international demand and supply. Thus, the greater the share of local final consumption that is produced non-locally (i.e., the higher is tradability), the lower is the effect of local demand on local inflation. Of course, the same logic applies at the national level in an international context, but international tradability is far lower than intra-national tradability.
- *A priori*, it is possible that increases over recent years in cross-geographical labor mobility or output tradability could have caused a flattening in the cross-geographical price PC *without a corresponding flattening in the aggregate price PC*. However, available evidence on these two factors suggests the opposite has occurred.

[5] Another issue is monetary policy. In cross-geographical comparisons, either with a single cross-section or pooled cross-sections with time fixed effects, monetary policy is held fixed. By contrast, in national time series PC regressions, monetary policy must be explicitly controlled for or it becomes an omitted variable. If it is omitted, estimates of the national PC slope will be downward biased estimates of the true causal effect of demand shocks on inflation because monetary policy generally "leans against the wind" – offsetting the inflationary pressures from aggregate demand shocks by raising interest rates in response to reductions in resource slack. Thus, national time-series estimates of the PC slope that do not explicitly account for monetary policy could be lower than estimates of the cross-geographical PC slope.

[6] For similar reasons, these same factors are key sources of difference between a cross-geographic fiscal multiplier and an aggregate multiplier, as shown in Nakamura and Steinsson (2014) and Leduc and Wilson (2013). Blanchard and Katz (1999) discuss how internal labor mobility can cause cross-geographic PC estimates to be biased estimates of the aggregate PC: "This source of bias is likely to be especially important for the United States, where labor mobility is a major source of adjustment to state labor-market shocks."

  - First, cross-state labor mobility in the U.S. has fallen over time (Kaplan and Schulhofer-Wohl 2017), which should have *steepened* the cross-geographical PC slope over time.
  - Second, there is no evidence suggesting that the tradable-good share of local economies' final consumption has changed over recent years. Also, note that the tradability factor only affects the price PC slope, so changes in tradability could not explain the observed flattening of the cross-geographical wage PC.
- Thus, changes in these two factors together cannot explain the flattening of the cross-geographical PC slope, implying that the aggregate PC slope – holding monetary policy fixed – has flattened at least as much.

B. Why might the PC have flattened?

- The cross-geographical PC estimates can be useful in discriminating between alternative theories of *why* the aggregate PC slope may have flattened during the current recovery.
  - One theory is that inflation expectations have become unanchored. This explanation is hard to square with the cross-geographical evidence presented in this briefing because inflation expectations are generally thought to be tied to the monetary policy inflation objective, which is common across all states and cities.
  - Another theory is that the economy has been hit by a string of negative inflation shocks, involving several different idiosyncratic factors (see other briefing in current binder). However, the factors to which analysts have pointed, such as factors involving the dollar, oil, pharmaceutical products, cellular services, etc., tend to be global and national in nature and hence they are unlikely to explain the cross-geographical PC flattening.
  - Similarly, disinflation due to technological and structural changes in the retail sector cannot explain this cross-geographical flattening.
  - In addition, these idiosyncratic factors relate to *price* inflation and offer no explanation of the observed flattening of the cross-geographical *wage* PC.
- Theories that *could* potentially explain both a cross-geographic PC flattening and an aggregate PC flattening generally involve changes in structural parameters characterizing the wage and price setting processes (e.g., in a New Keynesian model).
  - For example, a decline in worker bargaining power (combined with unchanged pass-through of unit labor costs to final prices) – consistent with reduced unionization, the

decline in the labor share, and increased industry concentration ratios – is a potential explanation. However, these trends generally pre-date the start of the flattening of the cross-geographical PC.

  - Another possibility is *downward nominal wage rigidities*. Daly and Hobijn (2014) show that downward nominal wage rigidities are more likely to bind in recessions, and particularly in the Great Recession. They show theoretically that such rigidities can mute the downward pressure on wage growth during these periods when labor market slack is especially high, thus flattening the wage Phillips Curve. Moreover, pent-up demand for real wage cuts by firms can linger into the recovery period and delay the normalization of the Phillips Curve even as slack returns to normal.

- The explanation for the flattening, of course, is crucial for predicting whether the PC slope will soon return to prior levels or remain flat in coming years.
  - Explanations involving secular declines in worker bargaining power suggest the flat slope will persist, while explanations such as downward nominal wage rigidities suggest the PC slope should steepen going forward.
- Monitoring how the cross-geographical PC evolves in the years ahead should help reveal the underlying causes of the recent flattening.

References

Blanchard, Olivier, and Lawrence F. Katz, "Wage Dynamics: Reconciling Theory and Evidence," *American Economic Review*, May 1999, 89(2), pp.69-74.

Daly, Mary, and Bart Hobijn (2014), "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money, Credit, and Banking*, 46(S2), pp.51-93.

Kaplan, Greg, and Sam Schulhofer-Wohl (2017). "Understanding the Long-Run Decline in Interstate Migration," *International Economic Review*, 58(1), pp.57-94.

Leduc, Sylvain, and Daniel J. Wilson (2015), "Regional Phillips Curves," FRBSF pre-FOMC Briefing, June 2015.

Leduc, Sylvain and Daniel Wilson, "Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment." In *NBER Macroeconomics Annual* 2012, Acemoglu, et al., editors (2013) pp. 89-142.

Nakamura, Emi and Jon Steinsson, "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 2014, 104 (3), 753-92.

Appendix. Supplementary Results

Table A1: Slope of Cross-MSA Price Phillips Curve
Alternative Measures of Slack
Estimates of $\hat{\beta}$ from $\pi_{i,t} = \alpha_t + \gamma\pi_{i,t-1} + \beta(u_{i,t} - \bar{u}_{i,t}) + \epsilon_{i,t}$

Measure of $(u_{i,t} - \bar{u}_{i,t})$:	$urate_t$	$urate_{t-1}$	$ugap_{t-1}$
Pooled 2009-15	-0.165**	-0.056	-0.106
	(0.071)	(0.074)	(0.068)
Pooled 1991-2015	-0.251***	-0.226***	-0.262***
	(0.030)	(0.031)	(0.031)

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Pooled regressions include year and geographical fixed effects. Unweighted OLS.

Table A2: Estimates of $\hat{\beta}$ by Inflation Category, MSA Level
$\pi_{i,t} = \alpha_t + \gamma\pi_{i,t-1} + \beta(u_{i,t} - \bar{u}_{i,t}) + \epsilon_{i,t}$

	1991-2008	2009-2015	1991-2015
Total	-0.280***	-0.155***	-0.234***
	(0.038)	(0.060)	(0.030)
Durable	-0.013	0.074	-0.022
	(0.076)	(0.120)	(0.061)
Non-Durable	-0.144***	-0.172*	-0.144***
	(0.051)	(0.090)	(0.042)
Food (Home)	-0.097	-0.004	-0.084
	(0.076)	(0.095)	(0.058)
Commodities	-0.135***	-0.073	-0.121***
	(0.044)	(0.077)	(0.037)
Rents	-0.646***	-0.483***	-0.579***
	(0.071)	(0.114)	(0.058)
Other Services	-0.025	0.198*	0.039
	(0.068)	(0.113)	(0.056)

Standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Regressions include year and geographical fixed effects. Unweighted OLS.

Table A3: Slope of Cross-MSA Wage Phillips Curve
Alternative Measures of Slack
Estimates of $\hat{\beta}$ from $\pi^w_{i,t} = \alpha_t + \gamma\pi_{i,t-1} + \beta(u_{i,t} - \bar{u}_{i,t}) + \epsilon_{i,t}$

Measure of $(u_{i,t} - \bar{u}_{i,t})$:	$urate_t$	$urate_{t-1}$	$ugap_{t-1}$
Pooled 2009-15	-0.178	0.068	0.023
	(0.116)	(0.120)	(0.111)
Pooled 1991-2015	-0.360***	-0.200***	-0.253***
	(0.057)	(0.060)	(0.062)

Standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Pooled regressions include year and geographical fixed effects. Unweighted OLS.

Figure A1: Slope of Cross-State Wage Phillips Curve Over Time
Slack Measured by Unemployment Gap

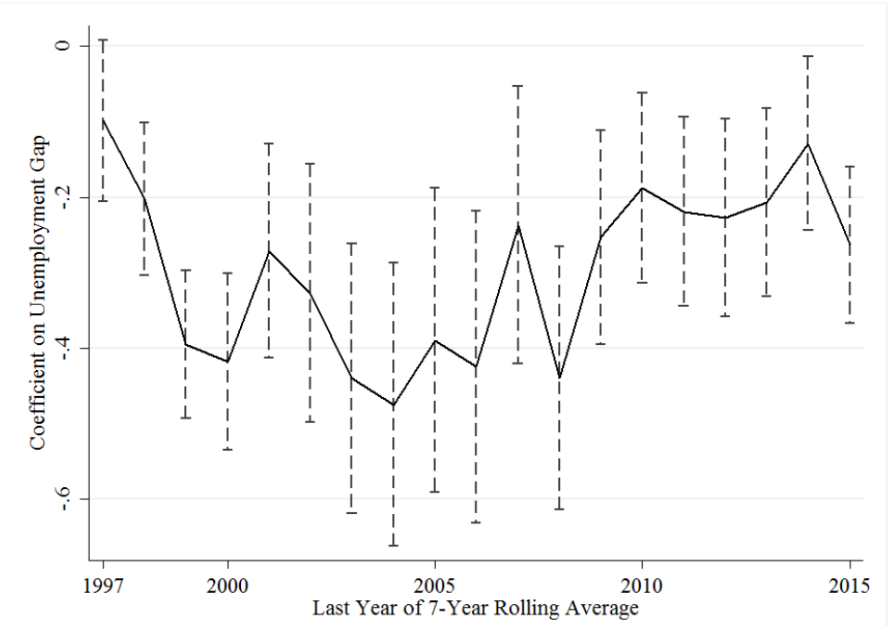

Graphs show the coefficients and 90% confidence interval on the unemployment gap from a regression of wage inflation on lagged wage inflation, year fixed effects and geographical fixed effects for the preceding seven years.

Figure A2: Slope of Cross-State Wage Phillips Curve Over Time
Slack Measured by Unemployment Rate

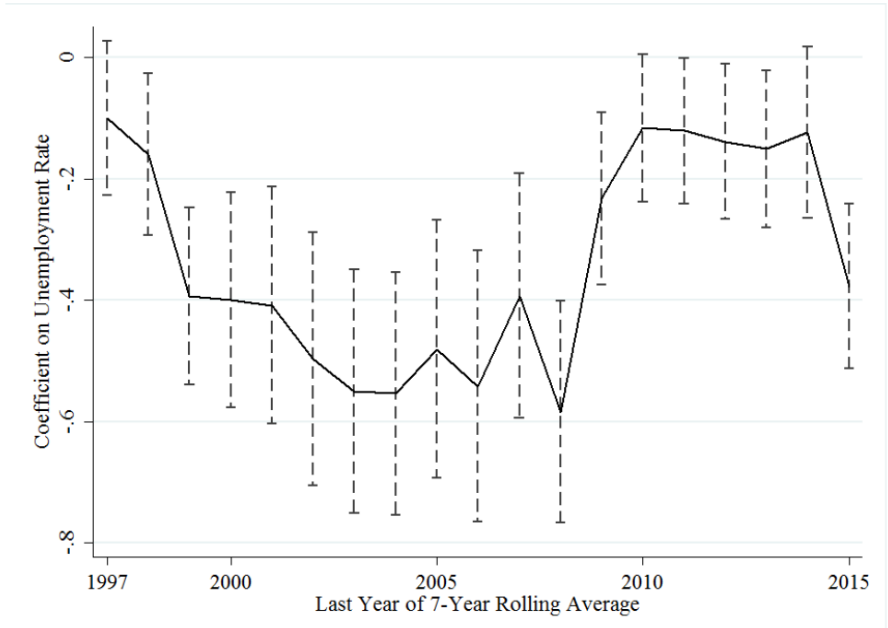

Graphs show the coefficients and 90% confidence interval on the unemployment rate from a regression of wage inflation on lagged wage inflation, year fixed effects and geographical fixed effects for the preceding seven years

Figure A3: Slope of Cross-State Wage Phillips Curve Over Time
Wage Measured by Average Hourly Earnings
Slack Measured by Unemployment Gap

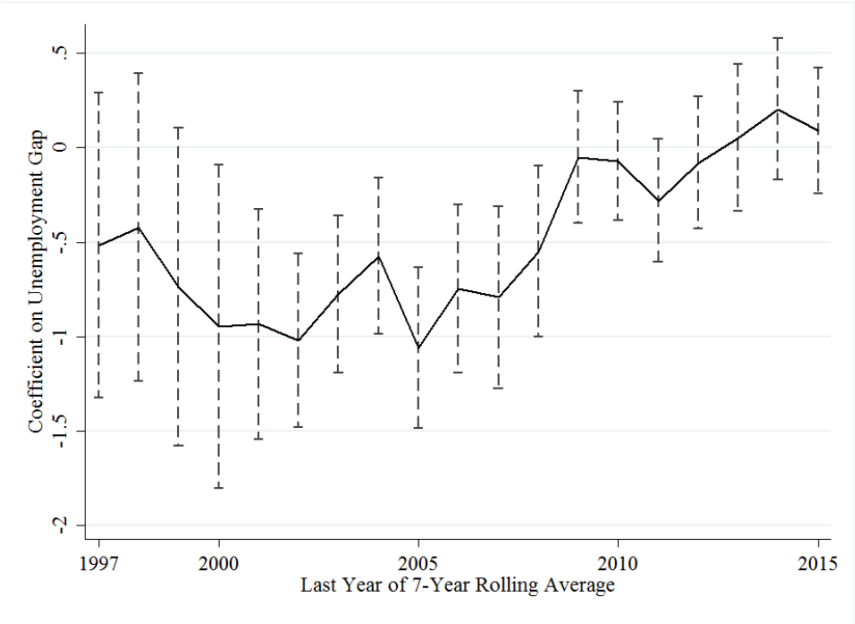

Graphs show the coefficients and 90% confidence interval on the unemployment gap from a regression of average hourly earnings inflation on lagged average hourly earnings inflation, year fixed effects and geographical fixed effects for the preceding seven years

Figure A4: Slope of Cross-State Wage Phillips Curve Over Time
Wage Measured by Average Hourly Earnings
Slack Measured by Unemployment Rate

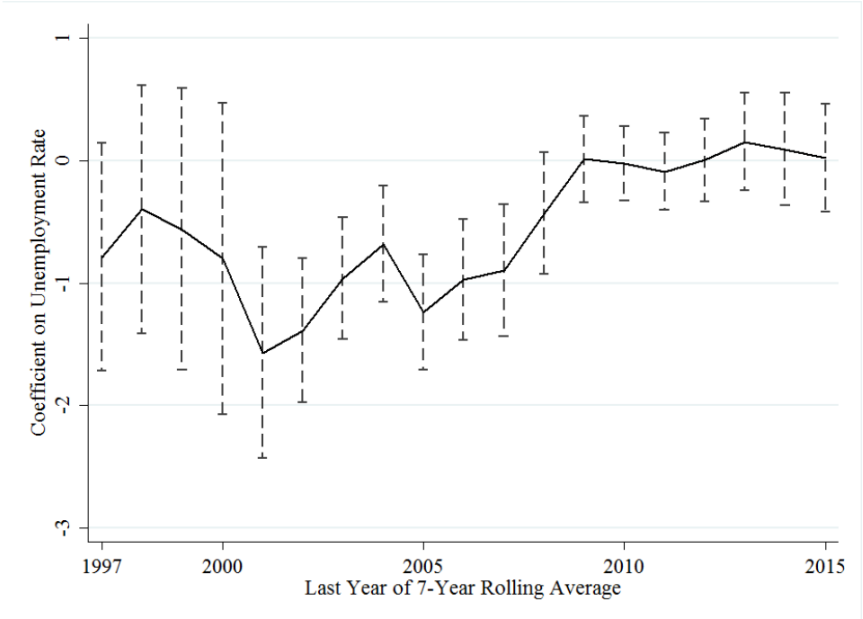

Graphs show the coefficients and 90% confidence interval on the unemployment rate from a regression of average hourly earnings inflation on lagged average hourly earnings inflation, year fixed effects and geographical fixed effects for the preceding seven years