

FROM NY TO LA: A LOOK AT THE WAGE PHILLIPS CURVE USING CROSS-GEOGRAPHICAL DATA

SYLVAIN LEDUC AND DANIEL WILSON

ABSTRACT. This paper estimates the cross-geographical wage Phillips Curve (PC) and relates this object to the aggregate wage PC through the lens of a New Keynesian model of regions within a monetary union. We argue that a well-identified cross-geographical PC, combined with a theoretical mapping from this object to the aggregate PC, provides an appealing alternative to estimating the latter from time-series variation. We employ this approach to study the recent debates over whether the wage PC slope has flattened in recent years and whether the wage PC is nonlinear. We find substantial evidence of a flattening of the wage PC during the recovery from the Great Recession, using both state and city panel data. We find no evidence of any economically meaningful nonlinearity. As our theoretical model shows, a flattening cross-geographical wage PC need not imply a flattening aggregate PC if intra-national labor mobility has risen and/or if monetary policy has become less passive. However, evidence points to the opposite, suggesting that the aggregate PC slope flattened at least as much.

Preliminary. Do Not Quote. January 2018

We thank Chitra Marti for superb research assistance. The paper has benefitted from numerous helpful comments and discussions with colleagues at the Bank of Canada and the Federal Reserve Bank of San Francisco, especially those of Mary Daly, Oscar Jorda, Adam Shapiro, and John Williams. The views expressed in this paper are our own, and do not reflect those of the Bank of Canada, the Federal Reserve Bank of San Francisco, or the Federal Reserve System.

I. INTRODUCTION

This paper proposes a general approach to analyzing the wage Phillips Curve that combines parameter identification from cross-geographical panel data with a micro-founded, structural multi-region monetary union model. We argue that this combination leverages the empirical identification advantages of panel data variation for local labor markets while providing insight into the national Phillips Curve, which is of critical importance to monetary policymakers.

Our approach allows us to address one of the biggest macroeconomic policy questions in recent years: why did price and wage inflation remain puzzlingly low during the recovery from the Great Recession despite the steady strengthening of the labor market and the notion of a stable wage Phillips curve (PC)? This puzzle has prompted many economists and policymakers to rethink their understanding of the wage PC. In particular, two hypotheses regarding the relationship between labor market slack and wage growth have gained prominence recently as potential explanations for the observed weakening of the slack-wage growth correlation. The first hypothesis is that the slope of the wage PC has declined in recent years. The second hypothesis is that the wage PC is non-linear, specifically convex, such that it is flatter during states of the world in which slack is high. If so, current forecasts of a tight labor market going forward would also predict accelerating wage growth. Of course, these two hypotheses need not be mutually exclusive.

However, identification challenges and limited time series variation greatly constrain the options for credibly testing these hypotheses with aggregate data alone. We argue that the panel data estimation of the cross-geographical Phillips Curve provides a cleaner identification strategy for estimating the causal relationship between labor market slack and wage growth.

Empirically, we find a strong negative slope for cross-geographical wage PC on average over our 1991 to 2016 sample period. However, we also find that this slope has varied over time. In particular, it flattened significantly over the 2009 to 2014 timeframe. This period corresponds to the recovery period from the Great Recession, with the national unemployment rate falling swiftly from roughly 10% to below 6%. These results are robust to different specifications, different measures of wage growth and labor market slack, and alternative geographical categories.

On the other hand, we find no evidence of a convex wage PC, unlike a number of other recent cross-geographical studies. The convexity result found in these prior studies, which even those studies acknowledge is economically modest, appears to stem from using more restrictive functional forms and measuring slack with the unemployment rate instead of the unemployment gap. In particular, we adopt a flexible, non-parametric, approach. We

provide evidence that the impact of labor market slack on wage growth rises linearly as the degree of slack in the labor market diminishes.

We then extend the closed economy framework of Galí (2011) with nominal wage rigidities to one of a federation with two economic regions. Importantly, the model allows for labor mobility across regions. While we show that the two curves can differ partly as a result, we also demonstrate an equivalence between the regional and national wage PCs that arises under relatively standard conditions. For instance, the slopes of the two PCs are identical when asset markets are complete and households in different regions have the same labor supply elasticity, even when allowing for labor mobility. Absent these conditions, labor mobility will generally imply that the cross-geographical PC slope will be an upper bound (in absolute value) on the aggregate PC slope.

Our cross-geographical approach provides an alternative view on the wage PC relative to the time-series literature. A number of important studies of the wage PC have explored its stability over time. In particular, Galí (2011) finds that the wage PC slope was strongly negative in the 1960s and in the post-1985 “Great Moderation” period (his sample went through 2009Q3), while being roughly flat in between. Galí attributed this flattening in the wage PC slope to variation in inflation regimes. An attractive feature of the cross-geographical approach is that it essentially hold inflation regimes constant (as monetary policy does not vary across cities and states).

Daly and Hobijn (2014) explore the potential of downward nominal wage rigidities (DNWR) to cause a flattening of the wage PC during recessions and their early recoveries. They show theoretically that DNWR can mute the downward pressure on wage growth during these periods when labor market slack is especially high, thus flattening the PC. Moreover, pent-up demand for real wage cuts by firms can linger into the recovery period and delay the normalization of the PC even as the level of the unemployment gap returns to normal. They show empirically that the aggregate U.S. wage PC was flatter in the aftermath of the past three recessions than in other years. This result is potentially consistent with our evidence that the cross-geographical PC slope has flattened during the current recovery. It is plausible that many local labor markets over 2009-2014 had high enough slack to experience binding DNWR, leading to a flattening of the PC.

On the other hand, the DNWR theory may seem inconsistent with the lack of evidence across states and cities for nonlinearity of the PC. However, as noted above, Daly and Hobijn’s theory and evidence indicate that DNWR during recessions cause wage growth to be unexpectedly slow in the initial years of the recovery, because of pent-up demand for real wage cuts, even as the level of the unemployment gap returns to normal. Thus, DNWR is really more about the relationship between wage inflation and the unemployment gap at the

peak of the most recent recession (and the amount of cumulative price inflation between then and now), and not so much about the relationship between wage inflation and the current unemployment gap.

A number of other recent studies have explored the cross-geographical evidence on the wage or price Phillips curves. Though each focuses on different aspects of the PC, Fitzgerald and Nicolini (2014), Kiley (2015), Babb & Detmeister (2017), and Murphy (2017) all estimate versions of a price PC using cross-MSA data. Fitzgerald and Nicolini focus on the temporal stability of the PC slope (assuming linearity) and conclude that it was stable over their 1977 – 2010 sample. Both Babb and Detmeister (2017) and Murphy (2017) test for nonlinearity in the cross-MSA price PC. Both studies find evidence of statistically significant convexity in the PC, but both also conclude that the degree of convexity is not economically significant. Neither of the latter papers look at the stability of the wage PC slope over time.

Kumar & Orrenius (2016) look at the wage PC using cross-state data. Similar to Babb and Detmeister’s (2017) and Murphy’s (2017) analyses of the price PC, they test for nonlinearity in the wage PC and find statistical evidence of convexity. But as in the price PC studies, the convexity appears to be economically insignificant in that they find that a wage PC forecasting model with convexity performs approximately the same as one without convexity.¹

Below we first describe, in Section II, the data and empirical methodology we use for estimating cross-geographical wage PC slope using cross-MSA and cross-state data. We then investigate the temporal stability (Section III) and the nonlinearity (Section IV) of the wage PC. In Section V, we present a model of regions within a monetary and fiscal union. We use the model to derive both the regional and aggregate wage PC. We then calibrate the model in Section VI and use the calibrated model to analyze the implications of our results on the cross-geographical wage PC slope for the aggregate wage PC slope. Section VII concludes.

II. ESTIMATING THE CROSS-GEOGRAPHICAL WAGE PHILLIPS CURVE SLOPE

II.1. Methodology. We estimate the cross-geographical wage PC slope using a dynamic panel model with area and time fixed effects:

$$\pi_{it}^{\omega} = \alpha_t + \alpha_i + \rho\pi_{it-1}^{\omega} + \beta(u_{it} - \bar{u}_{it}) + \epsilon_{it}, \quad (1)$$

where π_{it}^{ω} denotes the wage inflation rate in area i in year t , α_t represents time fixed effects, α_i is an area fixed effects, u_{it} denotes area i ’s actual rate of unemployment and $\bar{u}_{it}$ denotes

¹Kumar and Orrenius estimate a very restrictive nonlinear PC specification, in which lagged inflation is assumed to have a coefficient of one, slack is measured by the unemployment rate rather than the unemployment gap, and the slope is characterized by a linear spline with a single known kink at the sample average unemployment rate.

area i 's natural rate of unemployment. The natural rate is allowed to be area- and time-varying. Specifically, we assume that it is given by a 10-year trailing average of the area's unemployment rate. Note that α_t will capture aggregate movements in inflation expectations as well as other macroeconomic factors, including monetary policy. To account for potential serial correlation and contemporaneous spatial correlation in the error term, we two-way cluster by state and by year.

Abstracting from state and time fixed effects, equation (1) is standard (see, e.g., Blanchard and Katz [1999]) and also can be derived from a New Keynesian multi-region monetary union model under standard assumptions, as shown in Section V. In our cross-geographical panel, expectations that are not already captured by area-specific lagged wage inflation are assumed to be national in scope and hence captured by time fixed effects.

Our baseline identification assumption is that local labor market slack ($u_{it} - \bar{u}_{it}$) is orthogonal to the error term, ϵ_{it} , in equation (1). We assume that labor supply shocks, such as shocks to the discount factor or preferences governing the labor/leisure choice, are national or global in nature and hence are absorbed by the common year fixed effects. In addition, the baseline model assumes that persistent differences between areas in preferences or other labor supply factors will be captured by the area fixed effects. However, we also show that the results are robust to including area-specific time trends (in case such factors are trended) or including measured labor productivity growth, a key potential omitted variable.

In subsection II.4 below, we also test this baseline identification assumption using instrumental variables (IV). We instrument for local slack using a demand shock based on an interaction between the value of the dollar and the export orientation of each state. The instrument and the exclusion restrictions involved are discussed in subsection D.

II.2. Data. We measure yearly wages for a given area using a few alternative data sources. First, we calculate the average annual wage for each state and each Metropolitan Statistical Area (MSA) by taking total annual wage and salary income divided by total employment. Data on both the numerator and denominator are available from the BEA Regional Economic Accounts for states and for MSAs. For states, we also calculate Average Hourly Earnings (AHE) from the Current Population Survey (CPS) microdata. Specifically, we use the monthly outgoing rotation groups (CPS-MORG) microdata to calculate individual level hourly earnings.² We then calculate the average and median AHE across all individuals within each state using CPS sampling weights.³ The results presented below are based on the average AHE, but results are very similar using the median AHE. As of the current

²We follow Schmitt's (2003) methodology to account for top-coding, sampling weights, and outliers.

³We do not calculate AHE by MSA because a temporally consistent MSA geocoding is not available in the publicly-available CPS microdata.

writing, we have data for all of these series from 1989 to 2016. Because wage growth requires prior-year's wage and we include lagged wage growth in the regressions, these allows for a regression sample period of 1991 to 2016. Note that in the national time series, annual growth rates in average wage, AHE, and ECI are all very highly correlated.

Unemployment rates come from the Bureau of Labor Statistics are available back at least to 1981, allowing one to calculate unemployment gaps (using 10-year trailing averages of unemployment rates) from 1991 onward.

II.3. Baseline Full-Sample Results. Before investigating the constancy (over time) and linearity (over labor market tightness) of the Phillips Curve, we begin by estimating the average PC slope over our full sample period. The estimates of this slope, $\hat{\beta}$, along with the wage growth persistence parameter, $\hat{\rho}$, from estimating equation (1) for the full 1991 – 2016 sample period are shown in Table 1. We also report the implied long-run PC slope, $\hat{\beta}/(1 - \hat{\rho})$. Results are shown for three different sets of data. Those in Panel A are based on average hourly earnings (AHE) growth at the state level; Panel B is based on growth in the average wage (wages and salaries per worker) at the state level; and Panel C is based on growth in the average wage at the MSA level. In each panel, the first column shows our preferred specification, which includes both region (either state or MSA) fixed effects and year fixed effects. The second column excludes the year fixed effects and the third column excludes both sets of fixed effects.

In all cases, we find an economically and statistically significant cross-geographical Phillips Curve slope. In our preferred specification, we obtain slope estimates ranging from -0.24 using state average wages to -0.44 using state AHE. The estimate using MSA average wages is in between at -0.33. The long-run slope estimates range from -0.31 to -0.39. The range of long-run slope estimates is more narrow because growth in state AHE is negatively autocorrelated, leading the long-run PC slope to be somewhat flatter than the short-run slope, while state and MSA growth in average wages is positively autocorrelated, leading the long-run slope to be steeper than the short-run slope.

Comparing across columns, it seems that the slope estimate is rather insensitive to the presence or absence of region and year fixed effects. The insensitivity to region fixed effects is likely due to the fact that our measure of the local natural rate of unemployment – a 10-year trailing average of the local unemployment rate – is slow-moving and hence may capture much of the heterogeneity in permanent state characteristics.

Columns (4) and (5) address the potential concern that unobserved factors such as shocks to local productivity growth or labor supply could lead to biased estimates of the PC slope. For instance, a positive shock to local productivity growth could cause both a decline in the local unemployment gap and an increase in local wage growth, leading to a negative bias on

the estimated PC slope. Column (4) adds region (state or MSA) specific linear time trends, which will absorb any trends in productivity growth or labor supply that could be correlated with trends in slack and/or wage growth. We find that allowing for such trends has very little effect on the estimated PC slope. Column (5) adds measured labor productivity growth – the year-over-year growth in state/MSA GDP per worker – as a control variable. The estimated PC slope is again virtually unchanged.

Table 2 shows analogous results where the unemployment gap is replaced with the unemployment rate. The PC slopes estimated from the preferred, two-way fixed effects model are slightly larger, but generally very similar, to those in Table 1. And as in Table 1, conditioning on either region-specific time trends or measured productivity growth has very little effect on the estimated PC slope, suggesting that our baseline estimates are unbiased.

II.4. Instrumental Variables Results. To further check whether our baseline estimates of the PC slope are unbiased, we consider an instrumental variables (IV) approach. We construct an instrument for local slack based on an interaction between the value of the dollar and the export orientation of each state. The motivation is that movements in the value of the dollar can have big effects on the economies (and hence labor market slack) of export-oriented states. For example, Figure 2 shows the value of the dollar over time plotted against Washington state’s unemployment rate relative to the US unemployment rate. In general, when the value of the \$ rises, making US exports less competitive, export-oriented states like WA are especially hurt.

The instrument we construct is the interaction between the real trade-weighted broad dollar index (from the Federal Reserve Board’s G.5 series) and a state’s average export share of GDP:

$$Z_{it} = \$_t * mean_i(X_{i\tau}/GDP_{i\tau}) \quad (2)$$

where $\$_t$ is the real trade-weighted value of the dollar and $X_{i\tau}/GDP_{i\tau}$ is the nominal value of exports originating from state i divided by nominal GDP in state i in year τ . The variable $mean_i(X_{i\tau}/GDP_{i\tau})$ is the mean of that ratio over time in state i . Data on exports by state start in 1996, so these means are calculated over 1996 to 2016. Also, because exports data are not available by MSA, we obtain IV results only at the state level.

The results are shown in Table 3. The first column shows results based on measuring wage growth using average hourly earnings while the second columns uses the average wage. Panel A is based on measuring slack using the unemployment gap while Panel B using the unemployment rate. The first-stage F statistic is around 14-16, well above standard critical values for weak instrument bias. The coefficient on the unemployment gap in Column (1) is nearly identical to the analogous OLS result in Table 1 (Panel A, Column (1)). However,

it is imprecisely estimated. The IV slope estimate in Column (2) is statistically significant and somewhat more negative than the analogous OLS case in Panel B of Table 1. The IV results using the unemployment rate to measure slack (Panel B) yield larger negative slopes. However, the first-stage fit of the instrument for the unemployment rate is weaker which appears to contribute to large standard errors.

All in all, the IV results are somewhat imprecise but generally confirm the hypothesis that the slope of the cross-geographical Phillips Curve is strongly negative.

III. THE PHILLIPS CURVE OVER TIME

III.1. Rolling-Sample Results. To investigate whether the cross-geographical wage PC slope has changed over time, for each of our data sets we estimate a series of panel regressions, including region and year fixed effects, using 7-year rolling samples for 1991 – 1997, 1992 – 1998, ..., through 2009 – 2015. Each regression uses our baseline OLS specification in equation (1). The PC slope coefficients ($\hat{\beta}$) and their 90% confidence intervals are shown in Figure 3. The panels on the left-hand side use the unemployment gap to measure slack while the panels on the right-hand side use the unemployment rate. The top two panels measure wages using state AHE; the middle two panels use state average wage; and the bottom two panels use MSA average wage. There is a strikingly similar pattern across all six cases: we see a flattening of the PC starting around the Great Recession (i.e., when the 7-year panel starts to include data from 2007 onward).

This flattening tends to persist through around 2014, after which the PC slope appears to steepen (become more negative) again. These rolling sample regressions suggest there may have been a structural break in the PC slope around 2009 and lasting through around 2014, which corresponds to the recovery period from the Great Recession.

III.2. Panel Structural Break Tests. To formally test for a structural break in the PC slope, we follow the methodology proposed in Baltagi, Feng, and Kao (2016) for estimating an unknown common (to all cross-sectional units) structural break and for testing its statistical significance. We apply this methodology to our several cases, varying the wage PC specification by whether we use state AHE, state average wage, MSA average wage, or Canadian province AHE data, and whether we measure slack with the unemployment gap or rate.

This methodology is essentially a panel extension to the well-known Andrews (1993) test for a single unknown break in a time series model. That test involves repeatedly estimating the model allowing for a break at date for every possible break date t_B in the sample. The estimated t_B yielding the minimum sum of squared residuals (SSR) (if using a Wald test) is

the most likely break date and Andrews (1993) providing a limiting distribution for the test statistic defined by this minimum SSR.

Similarly, the Baltagi, et al. (2016) panel extension involves repeatedly estimating the panel regression with a common break for every possible break date. In our case, the estimating equation takes the following form:

$$\pi_{it}^w = \alpha_t + \alpha_i + \rho\pi_{i,t-1}^w + \beta_0(u_{it} - \bar{u}_{it}) + \beta_1\mathbf{1}[t > t_B] \cdot (u_{it} - \bar{u}_{it}) + \epsilon_{it}, \quad (3)$$

where $\mathbf{1}[t > t_B]$ is an indicator equal to one if year is greater than the prospective break date, t_B . We estimate this regression separately for every potential break year between 1994 and 2012. There are too few years in our 1991 – 2016 sample to reasonably test for break dates in the first or last three years. We identify which break date yields the minimum SSR, $\hat{t}_B$. Baltagi, et al. show that, conveniently, this minimum SSR test statistic has “the same asymptotic distribution as if the true change points were known.” In other words, the test of the null hypothesis that there is no structural break relative to the hypothesis of a single break at $\hat{t}_B$ is simply the t-test of β_1 .

The results are shown in Table 4. The results for the preferred specification, using the unemployment gap, are shown in Panel A. The results generally align with the graphical rolling-sample results. The structural break test identifies a significant flattening in 2008 when using state AHE data or MSA average wage data. In both cases, the slope is estimated to be more positive by about 0.3. Using state average wage data, the structural break test picks up the significant steepening of the PC slope around 1998 that is visually apparent in Figure 2. This negative break statistically dominates any more modest positive/flattening break occurring around 2008. Across these cases, the common structural break estimation generally identifies a break around 2007 or 2008. The results based on the unemployment rate are shown in Panel B and they are similar to those based on the unemployment gap.

III.3. The Wage PC Slope During the Great Recession Recovery. The results above point to a possible flattening of the wage PC slope during the recovery from the Great Recession. Thus, here we attempt to further quantify the magnitude of the flattening by estimating the PC slope separately for years up to 2008 and for years from 2009 to 2014. These subsample results, using the unemployment gap to measure slack, are shown in Table 5. The regressions underlying Panel A are unweighted; those underlying Panel B weight states or MSAs by population. In the table, we report the short-run PC slope, $\hat{\beta}$ from equation (1), as well as the long-run slope, $\hat{\beta}/(1 - \hat{\rho})$.

For the period up through 2008, both the short- and long-run slopes are consistently negative and highly significant in all cases. For the Great Recession recovery period (2009-2014), the slope is significantly flatter. In fact, the slope is close to zero and statistically

insignificant during the latter period. The results are similar when weighting by population. They are also similar when using the unemployment rate instead of the gap to measure slack, as shown in Table 6.

IV. IS THE PHILLIPS CURVE NONLINEAR?

The evidence above suggests that the cross-geographical wage Phillips Curve weakened substantially during the recent recovery period, consistent with some time-series evidence pointing to a recent flattening in the aggregate wage and price Phillips Curves. Some observers have argued that the PC might be nonlinear, specifically convex, such that it is much steeper at very low levels of slack, which have not yet been seen in the current recovery but could be coming. Indeed, a number of recent cross-geographical studies, looking at both wage inflation and price inflation have found evidence consistent with a convex PC. Thus, we investigate here the nonlinear shape of the PC, using both wage and price inflation and both MSA and state data. While prior studies have estimated very restrictive functional forms (such a quadratic function or a piecewise linear function with 1 or 2 kinks), we employ a much more flexible semi-parametric specification allowing the inflation-unemployment relationship to vary by each decile of the unemployment gap distribution:

$$\pi_{it}^{\omega} = \alpha_t + \alpha_i + \rho\pi_{it-1}^{\omega} + \sum_{p=1}^{10} \beta_p D_{it}^p + \varepsilon_{it} \quad (4)$$

where D_{it}^p is an indicator function that equals 1 if the unemployment gap for that area*year is in the p^{th} decile (of the unemployment gap's distribution over the regression sample period) and 0 otherwise; so $D_{it}^1 = 1$ means the unemployment gap in area i in year t was in the lowest/tightest ten percent of area*year labor markets in the sample and $D_{it}^{10} = 1$ means the gap in area i in year t was in the highest/loosest ten percent of area*year labor market observations.⁴ The estimated $\hat{\beta}_p$ from equation (3) traces out the full nonlinear Phillips Curve.

The results are shown in panels (a)-(f) of Figure 3. All cases use the full 1991 – 2016 sample. The panels on the left-hand side define deciles using the unemployment gap, while those on the right-hand side use the unemployment rate. First, note that the estimated nonlinear Phillips Curves are qualitatively consistent with our earlier findings in that they are all strongly negatively sloped.

⁴For reference, the distribution of the unemployment gap across all MSA*year observations from 1991 to 2016 is shown in Appendix Figure A4 (along with the distribution of MSA unemployment rates and the gap and rate distributions across state*year observations in Figures A5-A6). The vertical lines indicate the decile cut-offs. So, for instance, $\beta_1 = 1$ for MSA*year observations in which the unemployment gap is -1.64 or below. In terms of MSA unemployment rates, the first decile is 3.57 or below.

Second, we find no evidence of economically significant nonlinearities in general and no evidence of convexity in particular. The evidence for convexity in the PC is slightly stronger if one measures labor market slack using the unemployment rate instead of the unemployment gap. As the right-hand side panels of Figure 3 shows, there is evidence of a slightly steeper slope around the lowest unemployment rate decile. However, as is visually apparent, even using the unemployment rate the estimated convexity is economically minor.⁵ A simple back-of-the-envelope calculation illustrates the point. The estimates in panel A (the cross-state AHE case) imply that a labor market moving from the 2nd lowest to the 1st lowest unemployment rate decile – which corresponds to a reduction of 1.54 percentage points based on the midpoints of these deciles – would tend to see wage growth rise by about 0.9 percentage point (from 7.5% to 8.4%). Let us compare that to the implied increase in inflation from a 1.54 p.p. reduction in the unemployment rate according to the estimated linear slope from Table 2, panel A. That slope estimate is -0.47 (see Appendix Table A2), which implies that a reduction in the unemployment rate of 1.54 p.p. increases inflation by 0.39 percentage point ($1.54 \times 0.47 = 0.72$). Thus, this nonlinearity at the lowest end of the unemployment rate distribution adds just three tenths (8.4p.p. versus 8.1p.p.) to the wage growth rate relative to that implied by a linear PC.

For all intents and purposes, the cross-geographical wage Phillips Curve is approximately linear.⁶ How then have some prior studies found evidence in support of a convex cross-geographical PC? There are two parts to the answer. First, those studies measure slack using the unemployment rate instead of the unemployment gap, which we saw above leads to a slightly more convex PC toward the tightest end of the unemployment rate distribution.⁷

⁵Moreover, very few economic forecasters expect the national unemployment rate to reach a level as low as that represented by the lowest decile in these figures. Specifically, note that the minimum unemployment rate projection of any FOMC participant in the June SEP was 3.8% (for 2019), which falls in decile 2 of both the MSA and state distributions. The current (July) national unemployment rate is 4.35%, which falls in decile 3 of both distributions. Hence, in terms of the key policy question of whether current forecasts of a “hot” labor market could imply nonlinear inflation pressures, the most optimistic SEP labor market projection should be thought of as moving from the 3rd to the 2nd decile of these unemployment rate distributions. While Figure 3 suggests some potential steepening of the PC between the 2nd and 1st deciles, there’s no evidence of any significant steepening between the 3rd and 2nd deciles.

⁶The results are similar using an alternative measure of the unemployment gap in which the natural rate for a given area and year is defined as the CBO’s national NAIRU in that year plus the average difference between that area’s unemployment rate and the national unemployment rate over 1991 – 2015.

⁷The studies typically include time and area fixed effects, so they are implicitly assuming that all variation in area-specific natural rates of unemployment are captured by time and area fixed effects. If, on the other hand, the natural rate varies within area over time and across areas within a year, then their Phillips Curve estimation will be misspecified.

Second, those studies find some evidence of statistically significant nonlinearity (with rather restrictive functional forms), but even the authors acknowledge that the nonlinearity is not economically significant, consistent with the back-of-the-envelope calculation above.

V. MODEL OF LOCAL WAGE PHILLIPS CURVE

Our empirical results indicate a significant decline in the cross-geographical wage PC with the onset of the Great Recession and ensuing recovery, although our evidence also points to a linear relationship between local labor market slack and nominal wage growth. Our approach has the benefit of capturing more variations in the data allowing us to more clearly identify the relationship of interest. However, one issue is the extent to which changes in the cross-geographical wage PC translate into similar changes in the national one. In this section, we examine this question through the lens of a theoretical model. In particular, we derive conditions under which the regional wage PC is isomorphic to the national one.

The model consists of a monetary union composed of two regions of equal size, H and F . Each region specializes in one type of tradable good, produced under perfect competition. Firms producing the goods use labor as the only input to production and can hire local labor as well as labor from the other region. Assets markets are complete at the regional and national levels. We relax this assumption later.

We define $C_{H,t}$ as the Home agent's consumption of the Home good at time t ; similarly, $C_{F,t}$ is the Home agent's consumption of the imported good f . The full consumption basket, C_t , in each country, aggregates Home and Foreign goods according to the following standard CES function:

$$C_t \equiv \left[a_H^{1/\phi} C_{H,t}^{\frac{\phi-1}{\phi}} + a_F^{1/\phi} C_{F,t}^{\frac{\phi-1}{\phi}} \right]^{\frac{\phi}{\phi-1}}, \quad \phi > 0, \quad (5)$$

where a_H and a_F are the weights on the consumption of Home and Foreign traded goods, respectively and ϕ is the constant (trade) elasticity of substitution between $C_{H,t}$ and $C_{F,t}$.

Each region is also composed of a large representative household encompassing a continuum of members. Labor is indivisible, so that all variations in labor input arises from the number of members working, i.e., from the extensive margin. Each member is specialized in a type of differentiated labor services. Labor services are indexed by i , with Home workers supplying labor services for $i \in [0, n]$, with Foreign workers supplying services from $i \in [n, 1]$. We assume that households faces nominal wage rigidity à la Calvo, so that each period there is a probability θ_ω that wages can be readjusted. Note that the probability is independent of either the time since the last adjustment or the type of labor. Labor services are supplied at wage rates $W_t(i)$ and $W_t^*(i)$ in the Home and Foreign regions, respectively, to intermediate goods producers who aggregate them into a composite labor input according to the following

technology

$$N_{H,t}^d = \left[\int_0^n N_{H,t}(i)^{\frac{\epsilon_\omega - 1}{\epsilon_\omega}} di + \int_n^1 N_{H,t}^*(i)^{\frac{\epsilon_\omega - 1}{\epsilon_\omega}} di \right]^{\frac{\epsilon_\omega}{\epsilon_\omega - 1}}$$

where $N_{H,t}(i)$ denotes labor services supplied by Home workers, $N_{H,t}^*(i)$ is the supply of labor services by Foreign workers in the Home market, and $N_{H,t}^d$ is the total labor demand. The elasticity of substitution between differentiated labor inputs is represented by ϵ_ω . In turn, the cost of the composite labor input is given by

$$W_t = \left[\int_0^1 W_t(i)^{1-\epsilon_\omega} di \right]^{\frac{1}{1-\epsilon_\omega}}.$$

The intermediate goods producers demand for each type of labor is obtained by minimizing expenditures taking wages as given, which yields the standard expression

$$N_{H,t}(i) = \left(\frac{W_t(i)}{W_t} \right)^{-\epsilon_\omega} N_{H,t}^d \quad (6)$$

and

$$N_{H,t}^*(i) = \left(\frac{W_t(i)}{W_t} \right)^{-\epsilon_\omega} N_{H,t}^d. \quad (7)$$

The household utility is a function of its members' consumption, C_t , and disutility from work. An employed household member has disutility from work given by χj^φ , whereas it is zero for an unemployed member. The disutility of work is partly governed by the preference shifter χ and by φ , which governs the marginal disutility of work. As in Galí (2011), the household per period utility corresponds to the integral of its members' utilities

$$\begin{aligned} U(C_t, \{N_t(i)\}) &= \log C_t - \chi \int_0^n \int_0^{N_{H,t}(i)} j^\varphi dj di - \chi \int_n^1 \int_0^{N_{F,t}(i)} j^\varphi dj di \\ &= \log C_t - \chi \int_0^n \frac{N_{H,t}(i)^{1+\varphi}}{1+\varphi} di - \chi \int_n^1 \frac{N_{F,t}(i)^{1+\varphi}}{1+\varphi} di, \end{aligned}$$

where $N_{F,t}(i)$ denote the fraction of Home household members specialized in type i labor who are employed in Foreign region in period t . Importantly, we assume full risk sharing between household members.

The household from working at Home, $\int_0^n W_t(i) N_{H,t}(i) di$, and in the Foreign region, $\int_n^1 W_t^*(i) N_{F,t}(i) di$, from dividend income from ownership of the Home firms, Π_t , and from holdings of a riskless one-period bond, B_{t-1} . The household uses its income to purchase a consumption bundle at price P_t and to buy the riskless bond at price Q_t . The household budget constraint is thus given by

$$P_t C_t + Q_t B_t = B_{t-1} + \int_0^n W_t(i) N_{H,t}(i) di + \int_n^1 W_t^*(i) N_{F,t}(i) di + \Pi_t. \quad (8)$$

A household member will decide to supply labor at a given wage rate only if the real wage is greater or equal to the disutility of work, expressed in terms of consumption, as given by the marginal utility of consumption

$$\frac{W_t(i)}{P_t} \geq C_t \chi j^\varphi.$$

Thus the marginal supplier of labor of type i , denoted by $L(i)$, is implicitly given by

$$\frac{W_t(i)}{P_t} = C_t \chi L_{H,t}(i)^\varphi.$$

Integrating over labor types, we can rewrite this expression as

$$\frac{W_t}{P_t} = C_t \chi L_{H,t}^\varphi, \quad (9)$$

where $\int_0^n L_{H,t}(i) di = L_{H,t}$ can be interpreted as the labor force participation of Home workers in the Home market and $\int_0^n W_t(i) di = W_t$ is the average wage rate in the economy. In turn, we define wage inflation in the Home region as $\pi_t^\omega = \frac{W_t}{W_{t-1}}$, where wages evolve according to the following law of motion

$$W_t = W_t^{o1-\theta_\omega} W_{t-1}^{\theta_\omega},$$

where W_t^o is the optimal wage rate chosen when wages can be reset, which occurs with probability $1 - \theta_\omega$.

V.1. Optimal wages and the local wage Phillips curve. When workers reoptimize their wages in a one-region, “closed,” economy as in Erceg, Henderson, and Levin (2000) and Galí (2011), they do so by choosing a wage rate that maximizes household expected utility subject to constraints. In contrast, in a two-region, “open,” economy, workers set the optimal wage rate to maximize household expected utility in the two regions, since that wage rate will impact workers in these different places. Therefore, when updating the wage rate, workers chooses the optimal wage, W_t^o (W_t^{*o}), to maximize the household expected utility in the two regions

$$\max_{W_t^o, W_t^{*o}} E_t \sum_{k=0}^{\infty} (\beta \theta_\omega)^k \left(\begin{aligned} &\log C_{t+k|t} - \chi \int_0^n \frac{N_{H,t+k|t}(i)^{1+\varphi}}{1+\varphi} di - \chi \int_n^1 \frac{N_{F,t+k|t}(i)^{1+\varphi}}{1+\varphi} di \\ &+ \log C_{t+k|t}^* - \chi \int_0^n \frac{N_{H,t+k|t}^*(i)^{1+\varphi}}{1+\varphi} di - \chi \int_n^1 \frac{N_{F,t+k|t}^*(i)^{1+\varphi}}{1+\varphi} di \end{aligned} \right)$$

subject to the budget constraint (8), the demand for labor given by (6) and (7), as well as their respective counterparts in the Foreign region. The household's discount factor is given by β , where $0 < \beta < 1$. The subscript $t+k|t$ indicates a variable at time $t+k$ given the optimal wage rate that was set in period t , but still prevailing at $t+k$. The first-order condition to this problem dictates that workers set the wage rate such that

$$\sum_{k=0}^{\infty} (\beta\theta_{\omega})^k E_t \left[\begin{aligned} & N_{H,t+k|t} \left\{ \frac{\epsilon_{\omega}}{(1-\epsilon_{\omega})} MRS_{t+k|t} - \frac{W_t^o}{P_{t+k}} \right\} \\ & + N_{H,t+k|t}^* \left\{ \frac{\epsilon_{\omega}}{(1-\epsilon_{\omega})} MRS_{t+k|t}^* RER_{t+k} - \frac{W_t^o}{P_{t+k}} \right\} \end{aligned} \right] = 0$$

where $MRS_{t+k|t}$ ($MRS_{t+k|t}^*$) is the marginal rate of substitution between consumption and labor supply in the Home (Foreign) region and RER_{t+k} is the real exchange rate defined as the ratio of relative prices across regions, $\frac{P_{t+k}^*}{P_{t+k}}$.

abor supply of domestic workers at time $t+k$ given the optimal wage rate set in period t

Log-linearizing the expression around the non-stochastic steady state, we get an expression for the optimal wage

$$\widehat{W}_t^o = (1 - \beta\theta_{\omega}) \sum_{k=0}^{\infty} (\beta\theta_{\omega})^k E_t \left[\left(s \widehat{MRS}_{t+k|t} + (1-s) \left(\widehat{MRS}_{t+k|t}^* + \widehat{RER}_{t+k|t} \right) \right) + \widehat{P}_{t+k} \right]$$

From this problem, we can derive a log-linearized expression for the wage Phillips curve

$$\widehat{\pi}_t^{\omega} = \beta E_t \widehat{\pi}_{t+1}^{\omega} + \frac{(1 - \theta_{\omega})(1 - \beta\theta_{\omega})}{\theta_{\omega}(1 + s\varphi\epsilon_{\omega} + (1-s)\varphi^*\epsilon_{\omega})} (s\widehat{\mu}_t^{\omega} + (1-s)\widehat{\mu}_t^{\omega*}) \quad (10)$$

where $\widehat{\pi}_t^{\omega}$ is wage inflation and $\widehat{\mu}_t^{\omega}$ denotes the wage markup, and where hatted variables represent deviations from steady state (Appendix A provides details of the derivation). The average wage markup represents the wedge between the average real wage and the average marginal rate of substitution between consumption and labor supply, given, in linearize form, by

$$\widehat{\mu}_t^{\omega} = \widehat{MRS}_t - (\widehat{W}_t - \widehat{P}_t) = \widehat{C}_t - \varphi \widehat{N}_{H,t} - (\widehat{W}_t - \widehat{P}_t).$$

Following Galì (2011), we can write the wage Phillips curve in terms of unemployment gaps. Linearizing equation (9) and defining the unemployment rate of Home workers in the Home market as $U_{H,t} = \frac{L_{H,t} - N_{H,t}}{L_{H,t}}$, which for relatively small unemployment rates can be approximated by $\widehat{U}_{H,t} = \widehat{L}_{H,t} - \widehat{N}_{H,t}$, we obtain

$$\widehat{W}_t - \widehat{P}_t - \widehat{C}_t - \varphi \widehat{N}_{H,t} = \varphi \widehat{U}_{H,t} = \widehat{\mu}_t^{\omega}.$$

Under flexible prices, we have a constant desired wage markup, $\widehat{\mu}^{\omega}$. In this case, the natural unemployment rate is also constant

$$\widehat{U}_{H,t} = \frac{\widehat{\mu}^{\omega}}{\varphi} = 0.$$

Taken all these together, we can write the wage PC as

$$\widehat{\Pi}_t^{\omega} = \beta E_t \widehat{\Pi}_{t+1}^{\omega} + \frac{(1 - \theta_{\omega})(1 - \beta\theta_{\omega})}{\theta_{\omega}(1 + s\varphi\epsilon_{\omega} + (1-s)\varphi^*\epsilon_{\omega})} \left(s\varphi \widehat{U}_{H,t} + (1-s)\varphi^* \widehat{U}_{H,t}^* \right). \quad (11)$$

The slope of the wage Phillips curve in (11) is flatter the higher the degree of wage rigidity, θ_ω , the lower is the Frisch elasticity of labor supply $\frac{1}{\varphi}$ ($\frac{1}{\varphi^*}$), i.e., the higher is φ (φ^*), and the higher is the elasticity of substitution between types of labor ϵ_ω . Note also that the slope is impacted by labor mobility, represented by the share of domestic and foreign workers employed in the Home region, s .

V.1.1. *Equality between the regional and national wage PCs.* We show that the regional wage PC curve collapses to the aggregate one if households have the same Frisch labor supply elasticity and if there is full risk sharing across regions. Given the utility function, full risk sharing entails that the real exchange rate equals relative consumption

$$RER_t = \frac{C_t}{C_t^*}$$

As a result, the wage markup of Foreign workers in the Home market becomes

$$\begin{aligned}\hat{\mu}_t^{\omega*} &= \widehat{W}_t - \widehat{P}_t^* - \widehat{C}_t^* - \varphi^* \widehat{N}_{H,t}^* \\ &= \widehat{W}_t - \widehat{P}_t - \widehat{C}_t - \varphi \widehat{N}_{H,t}^*.\end{aligned}$$

Since labor is demand determined and demand is the same for either domestic or foreign sources of labor (since the wage rate is the same for all workers, irrespective of regions), we have that $\widehat{N}_{H,t}^* = \widehat{N}_{H,t}$, implying that $\hat{\mu}_t^{\omega*} = \hat{\mu}_t^\omega$. Thus, if households in different regions have identical Frisch labor elasticities (i.e., $\varphi = \varphi^*$), then $\widehat{U}_{H,t} = \widehat{U}_{H,t}^*$. In this case the regional wage PC is isomorphic to the national, closed economy, one, as derived by Gali (2010):

$$\widehat{\Pi}_t^\omega = \beta E_t \widehat{\Pi}_{t+1}^\omega + \frac{(1 - \theta_\omega)(1 - \beta\theta_\omega)}{\theta_\omega(1 + \varphi\epsilon_\omega)} \varphi \widehat{U}_{H,t}. \quad (12)$$

Under these two assumptions, our empirical results on the flattening of the regional wage PC would thus also be indicative of a similar flattening at the national level. From the slope of the wage PC in equation (12), such flattening could be due to a higher the degree of wage rigidity, θ_ω , a lower Frisch elasticity of labor supply, or to a higher is the elasticity of substitution between labor types, ϵ_ω . However, it would not be affected by labor mobility across regions, in contrast to the more general case in expression (11).⁸

Finally, note that this PC also emerges in the absence of labor mobility across regions, which arises when setting $s = 0$ in the regional PC given by equation (11) above.

⁸Note that empirically the total unemployment rate in a region may differ from $U_{H,t}$ in the model, since the Home household may send workers to the Foreign region who may become unemployed there. The Current Population Survey (CPS), for instance, would classify these unemployed workers as unemployed in the Home region, since the survey is residence based.

VI. IMPLICATIONS AND CONCLUDING REMARKS

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: (1) cross-geographical labor mobility and (2) “leaning against the wind” behavior of monetary policy. The first factor reduces the cross-geographical PC slope relative to the national slope, while the second does the opposite.

A priori, it is possible that increases over recent years in cross-geographical labor mobility could have caused a flattening in the cross-geographical price PC without a commensurate flattening in the aggregate wage PC. However, available evidence on these two factors suggests the opposite has occurred. 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. Thus, change in this factor 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.

Of course, monetary policy has not necessarily been fixed over recent years. If anything, it was more passive than normal during the Great Recession and the first few years of the recovery due to the zero lower bound. Passive monetary policy implies a flatter national PC slope.

In sum, while at any given point in time, the aggregate PC slope may be greater or less than the cross-geographical slope, the factors causing this gap collectively suggest that the aggregate PC slope has flattened even more than the observed flattening of the cross-geographical slope.

Given this implication, the cross-geographical PC estimates can be useful in discriminating between alternative theories of why the aggregate PC slope seems to 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. Put differently, if households and firms believe the Fed has become more tolerant of low inflation – that is, it has either reduced its implicit target (while keeping its stated target unchanged) or reduced its weight on inflation in the monetary policy rule – then expected inflation should fall equally in all regions.

Another possibility is that the unemployment rate or conventional measures of the unemployment gap are poor measures of labor market slack and that Phillips Curve estimates based on better measures show no flattening of the PC slope. (See, for example, studies looking at short- vs long-term unemployment or part- vs full-time unemployment.) This

theory is also hard to square with the cross-geographical evidence since this measurement error in the unemployment gap at the state or MSA levels seems likely to be independent of local area inflation or wage growth, though one would want to test this empirically before fully ruling it out.

One remaining theory that could be consistent with the cross-geographical evidence is a decline in worker bargaining power (combined with unchanged pass-through of unit labor costs to final prices). This also is consistent with the secular decline in the labor share, reductions in unionization, and increases in industry concentration ratios, although some of these trends started well before the current recovery. More generally, any theory relating to changes in structural parameters characterizing the wage and price setting processes (e.g., in a New Keynesian model) are candidates for explanations that could simultaneously explain the aggregate and cross-geographical flattening of the Phillips Curve. For instance, as mentioned above, 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. According to this theory, the slope of the Phillips Curve will steepen going forward as wages return to equilibrium levels.

VII. REFERENCES

Andrews, D. W. K. (1993) “Tests for parameter instability and structural change with unknown change point.” *Econometrica* 61: 821–856.

Babb, Nathan R., and Alan K. Detmeister (2017). “Nonlinearities in the Phillips Curve for the United States: Evidence Using Metropolitan Data,” Finance and Economics Discussion Series 2017-070. Washington: Board of Governors of the Federal Reserve System.

Baltagi, Badi H., Qu Feng, and Chihwa Kao (2016). “Estimation of heterogeneous panels with structural breaks,” *Journal of Econometrics*, 191, pp. 176-195.

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.

Demyanyk, Yuliya, Elena Loutskina, and Daniel Murphy (2017). “Fiscal Stimulus and Consumer Debt.” Federal Reserve Bank of Cleveland. Unpublished Manuscript.

Gali, Jordi (2011) “The Return of The Wage Phillips Curve,” Journal of the European Economic Association, 9(3), pp.436-461.

Gordon, Robert J. (2013) “The Phillips Curve Is Alive and Well: Inflation And The Nairu During The Slow Recovery,” NBER Working Paper #19390.

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

Kiley, Michael. (2015) “An evaluation of the inflationary pressure associated with short- and long-term unemployment.” Economics Letters, 137, pp. 5-9.

Kumar, Anil and Pia Orrenius (2016) “A Closer Look at the Phillips Curve Using State Level Data” Journal of Macroeconomics. 47, pp 84-102.

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

Murphy, Anthony (2017) “Is the U.S. Phillips Curve Convex? Some Metro Level Evidence.” Federal Reserve Bank of Dallas. Unpublished manuscript.

Nakamura, Emi and Jón Steinsson (2014). “Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions.” American Economic Review 104(3): 753-792.

Nalewaik, Jeremy (2016). “Non-Linear Phillips Curves with Inflation Regime-Switching.” Finance and Economics Discussion Series 2016-078. Board of Governors of the Federal Reserve System (U.S.).

Roberts, John M. “Monetary Policy and Inflation Dynamics” (2006). International Journal of Central Banking, 2(3), pp. 193-230.

Stroebel, Johannes and Joseph Vavra, 2014. “House Prices, Local Demand, and Retail Prices,” NBER Working Paper 20710.

Williams, John C. (2006). “The Phillips Curve in an Era of Well-Anchored Inflation Expectations,” unpublished working paper, Federal Reserve Bank of San Francisco, September.

Wilson, Daniel J. “Answering Macro Questions with Geographical Data: Pros and Cons,” FRBSF pre-FOMC Briefing, June 2014.

Bank of Canada

Federal Reserve Bank of San Francisco