Shelter and The Phillips Curve

James T. Smith III

js2749a@student.american.edu, American University

Abstract

This article examines inflation by Consolidated Metropolitan Statistical Area and by component of the Consumer Price Index in order to understand the Phillips curve. The shelter component of the CPI is the most sensitive to unemployment. Cities that are more land constrained have steeper Phillips curves. This suggests that housing is central to the price Phillips curve.

Keywords:
Business Fluctuations; Cycles (E30), Price Level, Inflation, Deflation (E31)

1. Introduction

In questioning before the Congress on July 10th, 2019, Federal Reserve Chairman Jerome Powell said of the Phillips curve that "...it's gotten weaker..." but that "It's there; you can see it in state–level data". Powell may well have been thinking of a small literature that finds a negative relationship between inflation and unemployment in state and city data (Kumar and Orrenius, 2016; Leduc and Wilson, 2017; Fitzgerald et al., 2013; Kiley, 2015). This article adds to that literature by noting the important role of rent.

It is the shelter component of the Consumer Price Index that is driving the relationship. This is shown by examining the estimated relationship in a panel of 16 American Consolidated Metropolitan Statistical areas (CMSAs) between 1990 and 2017. The correlation between shelter inflation and unemployment also exists in American macroeconomic data going back to 1978.

Macroeconomic theory since Friedman has emphasized price stickiness and expectations as the explanation for any inflation/unemployment correlation. Macroeconomics should instead consider rents. The prices of goods that are in short supply should be more sensitive to fluctuations in income, especially if they are as necessary as housing.

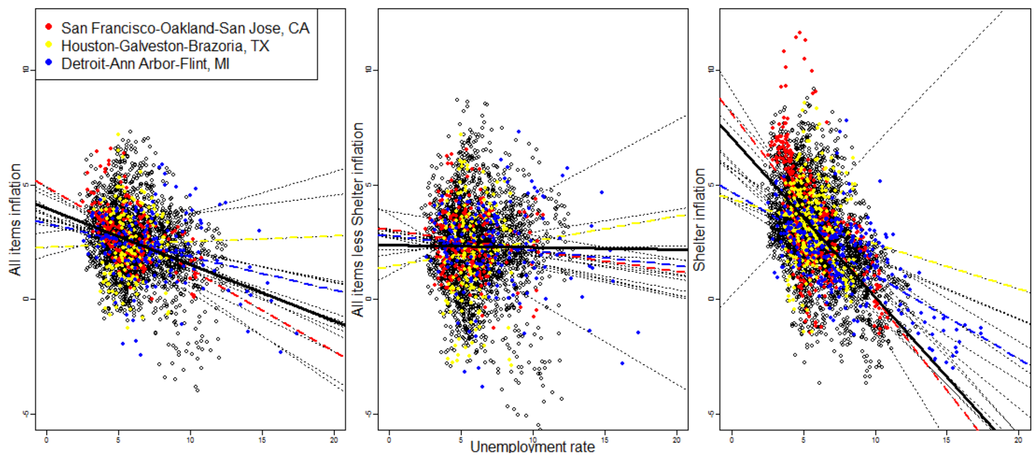

Figure 1: Scatter plot of CPI components versus Unemployment rate with CMSA specific regression lines. The panel regression fit is shown as the thick black line. Color coding is limited to three interesting cases.

The central finding of this study comes from an examination of the CPI inflation rate and local area unemployment rate in a panel of American Consolidated Metropolitan Statistical Areas (CMSAs) over the 1990–2017 period. A CMSA is a set of counties around a central city (see the data section for more detail). County–level unemployment data is aggregated to the CMSA–level and compared to the published CPI series. The leftmost panel of figure 1 shows that the all items CPI inflation rate is decreasing in the unemployment rate. It also appears that San Francisco has a steeper Phillips curve than Houston. The center panel shows that the relationship disappears when shelter is removed from the index. The shelter component of CPI is presented in the rightmost panel with a steeply negative relationship to unemployment. Shelter is crucial for understanding the Phillips curve.

An implication of housing's central role is that core inflation should be measured using all items in the inflation index less food, shelter, and energy (SA0L12E). Clark (2001) sets out the following objectives for a measure of core inflation: it ought to measure trend price changes, it ought not reflect changing relative prices, and it should be useful for forecasting. The standard core inflation removes food and fuel in order to remove volatile components. Shelter should also be removed because it is correlated with the business cycle. SA0L12E is about as volatile and about as useful for forecasting as CPI less food and fuel (CPILFESL). The major difference between the two series is that SA0L12E has run one half of a percentage point lower in recent years. This difference could reflect the fact that shelter has taken up a growing share of consumer expenditures over the past three decades. Economists and policymakers need to understand this secular change and its short run relationship with unemployment.

2. Literature Review

There are four direct precedents for this study: two published articles (Kiley, 2015; Kumar and Orrenius, 2016), a brief published by the Federal Reserve Bank of San Francisco (Leduc and Wilson, 2017), and a preliminary working paper published by the Federal Reserve of Minneapolis (Fitzgerald et al., 2013). Kiley uses a panel of cities to determine if there are distinct effects of short–term and long–term unemployment on price inflation. Kumar and Orrenius use a state panel to examine non–linearity in the real wage Phillips curve. Leduc and Wilson are concerned chiefly with the flattening of the wage Phillips curve. Fitzgerald et al. examine a panel of cities out of their belief that aggregate data cannot identify the Phillips curve's slope. All of these authors have important insights, but this letter will only discuss the slope estimates.

Short–term and long–term unemployment affect the shape of the wage growth distribution. Kiley finds approximately the same negative effect of short–term and long–term unemployment on a city's price inflation. Kumar and Orrenius find a difference between short–term and long–term in a state's median real wage but not on its average real wage. Short–term unemployment has the expected negative effect. Long–term unemployment has no significant effect on average wages but it does affect median wages. The coefficients for short–term and long–term unemployment in their real median wage equation are approximately equal (about -0.33). Kumar and Orrenius hypothesize that the significant effect of long–term unemployment on median real wages is due to the greater exposure of low wage workers to long–term unemployment. These results suggest that median wages and prices depend on total unemployment, but average wages are pulled up by high wage workers who are less exposed to long–term unemployment.

These findings justify not breaking unemployment into short–term and long–term components. Equal sized coefficient estimates on long and short term unemployment imply that the coefficient on the aggregate will be equal.

Period	Dependent Variable	Short Term		Long Term	
1985–2013	City price inflation	−0.22	(0.05)	−0.27	(0.07)
1994–2013	State CPS mean hourly real wage inflation	−0.503	(0.143)	−0.008	(0.137)
1994–2013	State CPS median hourly real wage inflation	−0.335	(0.153)	−0.353	(0.127)
1994–2013	State CES average manufacturing real wage inflation	−0.544	(0.263)	0.018	(0.603)
1994–2013	State QCEW average weekly real wage inflation	−0.254	(0.120)	−0.069	(0.107)

Table 1: Slope Estimates of Short Term and Long Term Unemployment in the Literature. Kiley's estimate in row 1 come from his Table 3, in the column labeled Metropolitan data 1985–2013. Kumar and Orrenius findings in the other rows come from their Table 9.

The unequal coefficients in the average–wage equations only suggest but do not demand a decomposition.

The estimated effect of unemployment on inflation varies over time. Leduc and Wilson (2017, see their figure 2) run a set of 7 year rolling regressions between 1997 and 2015. There is a period between 2000 and 2006 where the panel slope is steeper (about −1.0) than in the period before or after (about −0.3 before 2000, about 0 after 2009). Despite this apparent association with the mid-2000s housing boom, in unpublished results from a rolling regression that goes back to 1990, I find roughly stable slope values before the housing boom.

These studies control for trend inflation in different ways. Kiley relies on inflation expectations being anchored and constant during the Great Moderation. Kumar and Orrenius use a real wage specification and rely on wages and prices being cointegrated over the long run. Leduc and Wilson use a rolling regressions design so that expectations are assumed constant over a seven year period. Fitzgerald et al. focus on cross section deviation and rely on the cointegration of expectations across places. This paper relies on Kiley's constant expectations assumption but I have experimented with all of these techniques and arrive at similar conclusions.

3. Methods

$$\pi_{i,t} = \sum_i \mu_i I_i + \kappa u_{i,t} + \epsilon_{i,t} \tag{1}$$

This paper uses the standard panel estimator for the effect of local unemployment on local inflation κ. The variable $u_{i,t}$ is the unemployment rate in a given CMSA i at time t. The inflation measure is a CMSA–specific component of the year–on–year growth rate in the CPI. Following Kiley's work, there is no explicit expectations term in the model. Constant expectations are estimated as a CMSA specific intercept μ_i. I_i is a location specific indicator variable.

$$\pi_{i,t} = \sum_i \mu_i I_i + \kappa_i u_{i,t} + \epsilon_{i,t} \tag{2}$$

$$\hat{\kappa}_{i,shelter} = \beta_0 + \beta_1 L_i + \epsilon_{i,t} \tag{3}$$

I then estimate a panel variable coefficient model to examine slope variation (eq. 2). The city–specific slope estimates (weighted by the reciprocal of each estimates' own standard error) are regressed in equation 3 on a measure of that CMSAs undevelopable land, L_i. This regression examines how local housing supply constraints affect shelter inflation.

4. Data

The Bureau of Labor Statistics' CMSA specific inflation series are available through the CPI database on the BLS website. This study focuses on all items (SA0), all items less shelter (SA0L2), and shelter (SAH1) inflation. Shelter inflation is used rather than housing (SAH) which includes utilities and so is affected by fuel price fluctuations.

The definitions for the 16 CMSAs in this study are taken from the second appendix of the 1995 Statistical Abstract of the United States (Census, 1995). For the purposes of this study, a CMSA is a set of counties. The aim was to best match the geography described by the CPI series documentation. The central cities of the CMSAs used in this study are: New York, Philadelphia, Boston, Pittsburgh, Chicago, Detroit, St. Louis, Cleveland, Washington DC, Dallas, Houston, Atlanta, Miami, Los Angeles, San Francisco, and Seattle.

The counties in the CMSA definition can be matched to county–level activity statistics from the Local Area Unemployment (LAU) database. The LAU contains monthly unemployment and employment estimates for counties going back to 1990. It is straightforward to sum the LAU data over counties in a CMSA to calculate a CMSA unemployment rate.

Saiz (2010) reports a measure of undevelopable land for 100 metropolitan statistical areas. It is the share of land within a 50 kilometer radius of the city center that is either under water or has a slope greater than or equal to 15%. Many of the MSAs in his table are the central cities of the CMSAs in this study. The most constrained MSA in his table is Ventura, CA at 79.64% undevelopable which is one county in the Los Angeles-Riverside-Orange County, CA CMSA. I combine this number with the Los Angeles–Long Beach MSA (52.47% undevelopable) and Riverside-San Bernardino, CA (37.9% undevelopable) by taking an average. The mapping from MSA share to CMSA share for some cities like Miami (76.63%) is straightforward with no other MSAs in Saiz' table. Saiz' measure gives an imperfect sense of land constraints but it has more face validity than using a measure of the housing stock.

5. Results

5.1. CPI Component Specific Variation

Table 2 presents the estimated relationship between unemployment and inflation in high level components of the CPI. There are significant slopes at the $p < .01$ level for all of these components save for SAN and SA0L2. What differs across items is the model fit. The association that stands out is the housing component of CPI. The subcomponents (SAH1–SAH3) show that it is rent and not fuel shocks that are driving the association between CPI and unemployment. Nor is the correlation between SAH1 and unemployment an artifact caused by the imputation of Owner Equivalent Rent (OER), as both the observed Rent and the imputed OER sub-sub-components are correlated with unemployment. Finally, the correlation in the services component (SAS), which at first appears to suggest that non-tradables are at issue, disappears when rent is removed (SASL2RS). It is rent that is following the business cycle.

The evidence here is that shelter inflation contributes the most to the Phillips correlation. For this reason, I focus upon shelter (SAH1) versus non-shelter (SA0L2) indices of the CPI.

5.2. Regional Variation

Table 3 shows the estimated panel slopes from figure 1. The slopes for the all items inflation model here are approximately -0.2 and have poor model fit with an adjusted r^2 at 0.074. Controlling for time fixed effects

Inflation Component		Unemployment	Adj. r^2
SA0	All items	-0.220 (0.017)	0.077
SA0L2	All items less shelter	-0.007 (0.024)	-0.003
SAA	Apparel	0.522 (0.060)	0.024
SAC	Commodities	0.146 (0.033)	-0.006
SAD	Durables	0.388 (0.028)	0.051
SAE	Education & communication	-0.137 (0.026)	-0.004
SAF	Food and beverages	-0.175 (0.019)	0.021
SAG	Other goods and services	-0.450 (0.038)	-0.004
SAH	Housing	-0.582 (0.020)	0.275
SAH1	Shelter	-0.637 (0.017)	0.284
SAH2	Fuels and utilities	-0.562 (0.099)	0.022
SAH3	Household furnishings & ops.	-0.312 (0.035)	0.010
SAM	Medical care	-0.121 (0.031)	-0.005
SAN	Nondurables	0.037 (0.041)	-0.002
SAR	Recreation	-0.474 (0.037)	0.074
SAS	Services	-0.451 (0.016)	0.199
SASL2RS	Services less rent of shelter	-0.222 (0.024)	0.006
SAT	Transportation	0.566 (0.068)	-0.002
SETB	Motor fuel	0.602 (0.196)	-0.002

Table 2: Panel Regressions of Component Inflation Rates on the Unemployment Rate. Data is for the 1990M1–2017M6 period except for SAE and SAR which begin on 1998M12.

Inflation Component		Unemployment		Adj. r^2	Effects
SA0	All items	-0.224	(0.014)	0.074	
SA0L2	All items less shelter	-0.067	(0.020)	-0.001	city
SAH1	Shelter	-0.543	(0.014)	0.276	
SA0	All items	-0.246	(0.018)	-0.063	
SA0L2	All items less shelter	-0.011	(0.020)	-0.140	twoway
SAH1	Shelter	-0.697	(0.029)	0.055	

Table 3: Panel Regressions Of All items (SA0), Shelter (SAH1), and All items less Shelter (SA0L2) Component Inflation Rates On The Unemployment Rate. Effects are fixed effects models applied to either city or twoway (both city and time period).

gives nearly the same slope but with worse model fit (as one would expect).

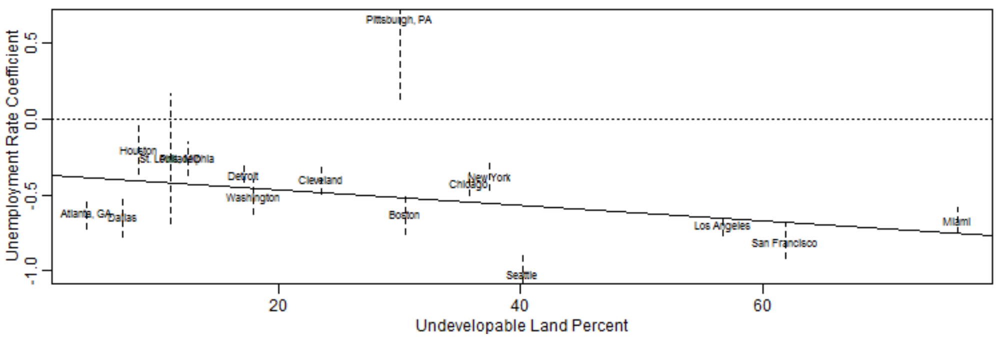

Figure 2: Phillips Coefficient versus Undevelopable Land. Dashed vertical lines indicate 1.96 standard error range on CMSA specific coefficients. Regression line excludes Pittsburgh CMSA (see equation 4).

The time fixed effects demonstrate that unemployment is associated with lower inflation not only across time periods but within time periods. Shelter inflation exhibits the same pattern but with steeper slopes and better fit for the individual effects model. Excising shelter from all items creates a slope that is essentially flat.

As was implied by the city–specific dashed lines in figure 1, there is variation in the slopes and intercepts of these models. Poolability tests for the individual effects models of all items (F-stat. 3.209, df1:15, df2:2790) and shelter inflation (F-stat. 16.747, df1:15, df2:4028) are statistically significant and reject stability of the model across cities.

The slopes of the shelter inflation models are correlated with the undevelopable land share. The only exception to the pattern is that Pittsburgh's Phillips coefficient is positive. This is likely the result of the short time series (1990–1997). Excluding Pittsburgh produces a statistically significant regression between the CMSA slopes and Saiz' measure plotted in figure 2 and equation 4. The estimated slope of the shelter inflation Phillips curve is steeper for cities that are more land constrained.

$$\hat{\kappa}_i = -0.367(0.084) - 0.005(0.002)L_i \qquad adj.r^2 = 0.248 \tag{4}$$

5.3. National Data: All Items Less Food Energy and Shelter SA0L12E

The correlation of shelter inflation with unemployment carries over into the aggregate data. The correlation is complicated by the increase in trend inflation during the 1970s. To address this, I subtract out a measure of core inflation and produce the same downward sloping relationship described by the city panel. The question then becomes which measure of core inflation one should use: the standard core all items less food and energy (CPILFESL) or all items less food shelter and energy (SA0L12E).

Figure 3 shows all items CPI minus a measure of core inflation versus the unemployment rate. The excess inflation on the vertical axis is chiefly the category omitted from the core measure. For example, the vertical axis in the right panel is CPI minus CPI less food and energy which approximately equals food and energy inflation.

There is a negative relationship between both measures of excess inflation and unemployment. The slope is much steeper and the model fit is stronger when excess inflation includes shelter. The food and fuel component offers a more modest downward sloping relationship.

A substantial shelter inflation occured in 1980. This appears as a loop that peaks at six percent excess inflation while unemployment is between six and eight percent.

The housing component of the CPI was measured differently – using home prices and mortgage interest rates – before 1984 (Gillingham and Lane, 1982). The potential confusion between measuring the investment value of owning a home, with measuring the user cost of shelter, lead the BLS to switch to the current owner–equivalent rent method.

I correct the all items CPI series by subtracting out the difference between CPI and CPI-U-RS; a research series that uses the rental equivalence method (CPI-U-RS) which goes back to 1977 (Stewart and Reed, 1999). This produces the bold lines in the figure 3 shown against the unadjusted data. The reduction in the 1980 peak is substantial and improves the model fit without altering the slope estimate. The same results hold for the CPILFESL measure.

The correlation between unemployment and inflation is better when using CPI less food, fuel, and shelter, rather than the standard measure of core inflation. This supports the hypothesis that shelter inflation varies with the unemployment rate. It also shows that the correlation matters not only in some land constrained cities, but in the aggregate CPI index.

5.3.1. Comparison of Core Measures

The following criteria have been suggested for a measure of core inflation. Core inflation (1) ought to measure trend price changes, (2) ought not reflect

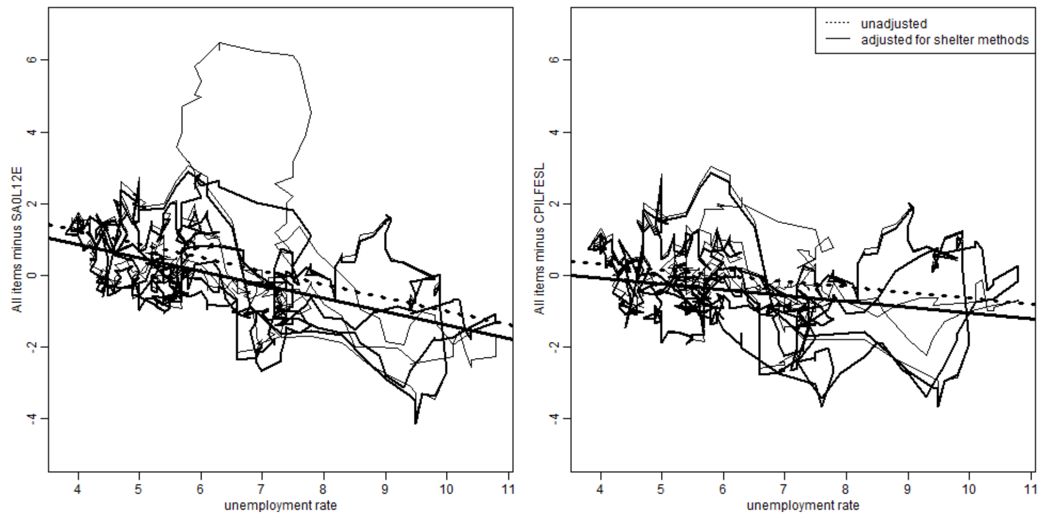

Figure 3: Phillips correlations contrasting measures of core inflation with CPI corrected for rental equivalence method 1978–2017. See table 4 for estimates and fit.

Core Measure	**Intercept**		**Unemployment**		r^2	**Adj.**
CPILFESL	0.938	(0.191)	−0.158	(0.029)	0.057	
CPILFESL	0.570	(0.206)	−0.160	(0.032)	0.050	X
SA0L12E	2.749	(0.26)	−0.375	(0.04)	0.157	
SA0L12E	2.381	(0.198)	−0.377	(0.03)	0.246	X

Table 4: Regression Estimates contrasting Core. All items inflation minus core measures as a function of the unemployment rate. The Adj. column shows that the dependent variable is adjusted for rental equivalent shelter inflation. Data is limited to 1978 M12 to 2017 M12. All estimates are made with the dynlm R package.

changing relative prices, and (3) should be useful for forecasting headline inflation (Clark, 2001).

The standard core includes shelter and reflects the increase in the relative price of shelter. In 2017, the relative importance of all items less food and energy is 79% of the headline index of which 32% is shelter; implying a 40% shelter share of CPILFESL. The relative price of shelter changes along with the business cycle. Failing to excise shelter from core leaves in an important relative price change that can be excised.

A measure of core inflation ought to measure trend and implicitly should be less volatile than headline inflation. Figure 4 compares headline infla-

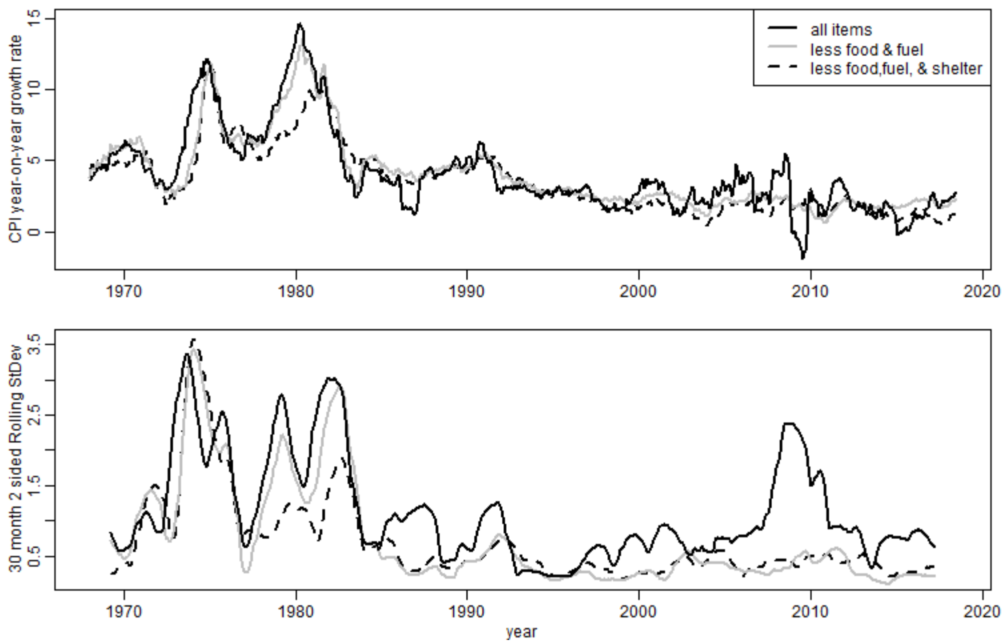

Figure 4: All items CPI and Core CPI measures. Data (upper) and rolling standard deviations (lower).

tion, the standard measure of core, and SA0L12E. The upper panel shows the series themselves. The lower panel shows a 30 month two sided rolling standard deviation. Much of the volatility in headline inflation is reduced by taking out fuel shocks. The two measures of core have approximately the same standard deviation. Where they differ is in the late 1970s and early 1980s when shelter inflation was high. Based on volatility, all items less food shelter and energy is about as good a measure of core infaltion as all items less food and energy.

A measure of core inflation ought to be useful for forecasting oncoming price changes. Clark (2001) suggests the following regression to assess within sample fit. The model fit is compared over 3, 6, 12, 24, month forecast horizons, h. Table 5 shows that SA0L12E is about as good as CPILFESL at forecasting changes in headline inflation rates.

$$\pi_{t+h} - \pi_t = \alpha + \beta(\pi_t^{core} - \pi_t) + \epsilon_t \tag{5}$$

$$\pi_t = \alpha + \beta\pi_{t-h}^{core} + \epsilon_t \tag{6}$$

Core Measure	Time Period		Horizon	r^2	coefficient	
CPILFESL	1985	2000	3	0.051	0.182	(0.059)
CPILFESL	1985	2018	3	0.088	0.235	(0.038)
CPILFESL	1985	2000	6	0.118	0.400	(0.083)
CPILFESL	1985	2018	6	0.181	0.475	(0.051)
CPILFESL	1985	2000	12	0.297	0.930	(0.111)
CPILFESL	1985	2018	12	0.355	0.914	(0.063)
CPILFESL	1985	2000	24	0.443	1.222	(0.11)
CPILFESL	1985	2018	24	0.399	0.992	(0.063)
SA0L12E	1985	2000	3	0.043	0.169	(0.06)
SA0L12E	1985	2018	3	0.069	0.200	(0.037)
SA0L12E	1985	2000	6	0.093	0.360	(0.085)
SA0L12E	1985	2018	6	0.142	0.404	(0.05)
SA0L12E	1985	2000	12	0.253	0.881	(0.117)
SA0L12E	1985	2018	12	0.293	0.803	(0.064)
SA0L12E	1985	2000	24	0.296	1.018	(0.126)
SA0L12E	1985	2018	24	0.346	0.886	(0.063)

Table 5: Predicting the change in inflation with All items less Food Shelter and Energy (SA0L12E) versus All items less Food and Energy (CPILFESL). Comparison of core inflation measures ability to forecast oncoming changes in inflation. All estimates are made with the dynlm R package.

Core Measure	Horizon				
	3 months	6 months	1 year	2 year	
CPIAUCSL	0.936	1.259	1.543	1.593	
SA0L12E	1.437	1.094	0.899	0.828	**
CPILFESL	1.392	1.062	0.881	0.850	**

Table 6: Relative Root Mean Squared Error of forecast using different measures of core inflation. The baseline model in rows labelled CPIAUCSL uses a single lag of all items CPI. Rows labelled CPILFESL and SA0L12E are relative RMSE. The models are trained on data from 1985–2000 and the forecast is made for post 2000 data. Diebold-Mariano test of forecast difference from baseline model implemented in the forecast R package. Significance Codes: ** $p < .01$.

To assess out–of–sample fit, I train the regression model in equation 6 on data before 2000, and examine how it affects the root mean squared forecast error on data after 2000. The baseline model uses a lag of the all items CPI in place of a measure of core. The alternative models use CPILFESL and SA0L12E as measures of core inflation. Significant difference from the baseline model is assessed with the Diebold–Mariano test implemented in R's forecast package (Hyndman et al., 2019).

The results in table 6 show that both measures of core inflation improve forecasts at the two year horizon. The two measures of core are competitive with one another. Neither stands out as obviously better.

6. Conclusion

This article shows that the shelter component of the CPI is correlated with the unemployment rate. The aggregate CPI inherits the correlation found in the panel of cities. The evidence here suggests that relative price increases of shelter rents are an important part of the inflation process.

The correlation should be viewed cautiously. This panel uses 16 CMSAs; of which Pittsburgh does not have the expected sign relating shelter inflation to unemployment. Whatever the precise cause and effect relationship may be, it seems clear that rent is a crucial part of inflation in the short–run. More research is needed to pin down the link between unemployment and shelter inflation.

Core inflation should take out shelter inflation to remove components that are correlated with the business cycle. Doing so offers little benefit for forecasting, but neither is there any substantial cost. SA0L12E is more in line with our intuition that core inflation should remove cyclical inflation.

Whether Federal Reserve policy should target SA0L12E is debatable. The core PCE index used by the Federal Reserve has a lower weight on shelter (16–20%) than CPILFESL (20–30%). Federal Reserve policy does not seem to be hampered by including shelter inflation, but they should consider removing it to avoid confusing transitory inflation pressures with long run inflation.

Acknowledgements

Many thanks to Dr. Xuguang Simon Sheng for reading early versions of this article.

Funding Statement

This research did not receive any specific grant from funding agencies in the public, commercial, or not–for–profit sectors.

References

Bureau of Labor Statistics, U.S. Department of Labor, a. CPI Research Series Using Current Methods (CPI-U-RS). https://www.bls.gov/cpi/research-series/home.htm.

Bureau of Labor Statistics, U.S. Department of Labor, b. Local Area Unemployment Statistics Home Page. https://www.bls.gov/lau/home.htm.

Bureau of Labor Statistics, U.S. Department of Labor, 2017. Consumer Price Index (CPI) Databases: U.S. Bureau of Labor Statistics. https://www.bls.gov/cpi/data.htm.

Bureau of Labor Statistics, U.S. Department of Labor, 2018. Archived Relative Importance of Components in the Consumer Price Indexes. https://www.bls.gov/cpi/tables/relative-importance/home.htm.

Census, 1995. Statistical Abstract of the United States. 115 ed., US Government Printing Office. https://www.census.gov/library/publications/1995/compendia/statab/115ed.html.

Clark, T.E., 2001. Comparing Measures Of Core Inflation. Economic Review-Federal Reserve Bank of Kansas City 86, 5.

Croissant, Y., Millo, G., 2008. Panel data econometrics in R: The plm package. Journal of Statistical Software 27.

116th Congress House of Representatives Committee on Financial Services, U.S., 2019. Monetary policy and the state of the economy. URL: https://financialservices.house.gov/calendar/eventsingle.aspx?EventID=403999#Wbcast03222017.

Fitzgerald, T.J., Holtemeyer, B., Nicolini, J.P., 2013. Is There A Stable Phillips Curve After All? Economic Policy Paper 13.

Fitzgerald, T.J., Nicolini, J.P., 2014. Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from U.S. Cities , forthcoming.

Gillingham, R., Lane, W., 1982. Changing the Treatment of Shelter Costs for Homeowners in the CPI. Monthly Labor Review 105, 9.

Hyndman, R., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., O'Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E., Yasmeen, F., 2019. forecast: Forecasting functions for time series and linear models. URL: `http://pkg.robjhyndman.com/forecast`. R package version 8.5.

Kiley, M.T., 2015. An Evaluation Of The Inflationary Pressure Associated With Short-and Long-term Unemployment. Economics Letters 137, 5–9.

Kumar, A., Orrenius, P.M., 2016. A Closer Look At The Phillips Curve Using State-level Data. Journal of Macroeconomics 47, 84–102.

Leduc, S., Wilson, D., 2017. Has the Wage Phillips Curve Gone Dormant? FRBSF Economic Letter 2017, 30.

Moretti, E., 2011. Local Labor Markets, in: Handbook of Labor Economics. Elsevier. volume 4, pp. 1237–1313.

Saiz, A., 2010. The Geographic Determinants Of Housing Supply. The Quarterly Journal of Economics 125, 1253–1296.

Stewart, K.J., Reed, S.B., 1999. Consumer Price Index Research Series Using Current Methods, 1978–1998. Monthly Labor Review , 29–38.

Zeileis, A., 2016. dynlm: Dynamic Linear Regression. URL: `http://CRAN.R-project.org/package=dynlm`. R package version 0.3-5.