SECTORAL INFLATION AND THE PHILLIPS CURVE: WHAT HAS CHANGED SINCE THE GREAT RECESSION?

María José Luengo-Prado[a]
Federal Reserve Bank of Boston

Nikhil Rao[b]
Federal Reserve Bank of Boston

Viacheslav Sheremirov[c]
Federal Reserve Bank of Boston

July 2018

Abstract: Using sectoral data at a medium level of aggregation, we find that disaggregated inflation rates became less responsive to aggregate unemployment around 2009–2010. The slopes of the disaggregated Phillips curves diminished in many sectors, including housing and some services. We also document a decrease in sectoral inflation persistence, suggesting an increase in the weight of the forward-looking inflation expectation component and a decrease in the weight of the backward-looking component.

Keywords: disaggregated price indices, inflation persistence, Phillips curve
JEL Classification: E24, E31, E32

[a]Federal Reserve Bank of Boston, Research Department T-9, 600 Atlantic Ave, Boston, MA 02210, United States. Email: Maria.Luengo-Prado@bos.frb.org. **Corresponding author.**

[b]Federal Reserve Bank of Boston, Research Department T-9, 600 Atlantic Ave, Boston, MA 02210, United States. Email: Nikhil.Rao@bos.frb.org.

[c]Federal Reserve Bank of Boston, Research Department T-9, 600 Atlantic Ave, Boston, MA 02210, United States. Email: Viacheslav.Sheremirov@bos.frb.org.

1 Introduction

Policymakers have struggled recently to understand why inflation dynamics differ from the predictions of workhorse models. One of the pillars of such models, the aggregate Phillips curve, predicts that as the labor market tightens, prices eventually face an upward pressure, and inflation takes off. Yet, while the unemployment rate decreased from a peak of 9.9 percent in the fourth quarter of 2009 to 4.3 percent in the third quarter of 2017, core inflation remained stubbornly below the Federal Open Market Committee target of 2 percent.

This stark deviation from the norm has sparked a debate about the usefulness of the aggregate Phillips curve framework for policy analysis.[1] Many economists, though, continue to believe that the Phillips curve remains informative, and that the recent changes in inflation dynamics can be reconciled with the models by allowing for a structural break in model parameters.

Using sectoral consumer price index (CPI) data from the Bureau of Labor Statistics (BLS) at a medium level of disaggregation, we find robust evidence of a structural break in the Phillips curve slope around 2009–2010. The co-movement of sectoral inflation rates and labor market slack has weakened, and it is now almost negligible. This change occurred broadly across sectors, although some sectors, including housing, apparel, and certain services such as food away from home, stand out in a statistical sense. We also document a decline in inflation persistence in many sectors, using a hybrid (sectoral) Phillips curve. Specifically, the relative weight of inflation lags decreased and the relative weight of the forward-looking expectation component increased.

Our results have broad implications for economic models and monetary policy. On the theory side, a structural break in key reduced-form parameters of the Phillips curve calls for a search of the break's origin in terms of micro-founded structural parameters. On the policy side, a Phillips curve flattening implies a slower response of inflation to policy tools, potentially requiring an adjustment of the policy-rate path obtained from historical observations. Also, decreasing persistence implies that low inflation in the past may not necessarily be indicative of low inflation in the future, and that managing expectations should play a crucial role in policy decisions.

2 Disaggregated CPI Data

The BLS produces CPI series for different levels of disaggregation. For example, their Level I series comprise eight categories with weights ranging from 42.6 percent for "Housing" to 3.0 percent for "Apparel." These categories are problematic when estimating a sectoral Phillips curve. Regressions with CPI weights are dominated by only a few series, while unweighted regressions amplify relationships pertinent to small sectors. We therefore conduct our analysis at a lower level of aggregation by compiling a list of sectors that satisfy the following criteria: (*i*) a size that is not too large and not too small; (*ii*) reasonable homogeneity in the size distribution; and (*iii*) a definition that roughly corresponds to a market.

[1] For example, the *Economist* of November 1, 2017, highlights this debate in its daily chart; see economist.com/blogs/graphicdetail/2017/11/daily-chart.

To meet these criteria, we compose a list of 16 categories comprising seven Level III BLS categories, seven Level II categories, and two Level I categories. These 16 categories cover 94.1 percent of the CPI. Each category has an individual CPI weight in the range of 2 percent to 8 percent, with the exception of "Owners' Equivalent Rent of Residence." This category has a weight of 24.6 percent, and cannot be split further in a meaningful way. Figure 1 depicts sectoral inflation dynamics for these selected categories during our sample period, showing rich variation in disaggregated inflation rates.[2]

3 Structural Break in Sectoral Phillips Curves

To test for a structural break in the Phillips curve slope using sectoral data, we estimate the model:

$$\pi_{jt} = \mu_j + \widetilde{\mu}_j \mathbb{1}_{t>\tau} + \sum_{i=1}^{5} \alpha_i \pi_{j,t-i} + \beta u_t + \widetilde{\beta} \mathbb{1}_{t>\tau} u_t + \gamma z_{jt} + \varepsilon_{jt}, \tag{1}$$

where π_{jt} is sector j inflation in quarter t; u_t is the aggregate unemployment gap with the natural rate estimated by the Congressional Budget Office; z_{jt} is a vector of controls; $\mathbb{1}_{t>\tau}$ is an indicator function for a period after a structural break τ; ε_{jt} is the error term; and the rest are parameters to be estimated. z_{jt} includes relative prices in sector j (the CPI in sector j normalized by the aggregate CPI) as well as the percentage change in the oil price. It is important to control for relative prices, because relative prices often appear in sectoral Phillips curves derived from multi-sector New Keynesian models.[3] Oil prices control for supply-side shifters that may be important in practice but are often omitted from the models. We allow for sectoral fixed effects μ_j to account for differences in productivity levels across industries, which may lead to different levels of steady-state inflation. The test results, however, are robust to excluding the intercept break or allowing for the effect of oil-price changes to vary by sector.

This model is estimated for a range of possible structural breaks τ. For each τ, the structural break test computes an F-statistic for a null hypothesis $\widetilde{\beta} = 0$ (no structural break). The identified break τ^* is chosen to maximize the F-statistic of this test. The statistical significance of the break is determined by comparing $F(\tau^*)$ with the test's critical values.

Panel A of Figure 2 shows the F-statistic of the test across possible structural breaks. The test picks 2010q3 as a structural break, but the F-statistic is elevated throughout the Great Recession. When we drop highly volatile categories such as "Motor Fuel" and "Fuels and Utilities" (Panel A of Figure 3) or keep only core CPI categories (Panel B), the break estimate occurs just one quarter later: 2010q4. The identification of a break is also robust to using a real-time unemployment gap measure (Panel C), the unemployment *rate* as the driving variable (Panel D), or controlling for five lags of aggregate inflation (Panel E). The test is also robust to using alternate controls for relative

[2]More information on these categories, including CPI weights, can be found in Luengo-Prado, Rao, and Sheremirov (2017).

[3]Examples of multi-sector New Keynesian models include Carvalho (2006) and Reis and Watson (2010).

prices: (1) controlling for the second lag of our baseline measure of relative prices, which does not have any common terms with sectoral inflation (Panel F); (2) controlling for inflation in other sectors by excluding sector j and re-normalizing the weights for other sectors (Panel G). We also show the test result when we exclude relative prices from the vector of controls (Panel H).[4]

To determine which sectors contributed the most to the structural break in the Phillips curve slope, we estimate Equation (1) for each sector j separately using $\tau^* = 2010q3$ as the structural break date (as identified from the panel regression).[5] Note that whereas in the baseline panel setup all parameters except the fixed effects are the same across sectors, in this exercise all parameters are sector specific. Panel B of Figure 2 shows p-values for the test that $\widetilde{\beta}_j = 0$. The null hypothesis that the slope is the same before and after the break is rejected at the 1 percent significance level in three sectors: "Food Away from Home," "Owners' Equivalent Rent," and "Rent of Primary Residence." The null of equal slopes is also rejected for "Apparel" at the 5 percent level, and for several other categories at the 10 percent level, suggesting that a change took place across the board.

We reach this conclusion also when we look at individual slopes before the break β_j and at the difference in slopes before and after the break $\widetilde{\beta}_j$ (Panel A of Table 1). In most cases, $\widetilde{\beta}_j > 0$, and often $\widetilde{\beta}_j \approx -\beta_j$. These estimates indicate a flattening of the Phillips curve after the Great Recession to essentially a horizontal line (i.e., inflation does not respond to the unemployment gap). Panel B compares these sectoral coefficients with those from the panel regression in Equation (1), which were obtained with and without sectoral weights. It also excludes volatile fuel and food prices. Whether we pool the sectors or not, we obtain a consistent picture of a flattening Phillips curve.

Our estimates of sectoral slope changes suggest that, given the path of the unemployment gap after the break, inflation would have been 0.63 to 1.40 percentage points higher in 2017q3 if the break had not occurred, depending on the aggregation method (Figure 4). Aggregate inflation would have been 0.49 percentage points higher if the break had occurred only in the four sectors with a statistically significant flattening (at the 5 percent level).

4 Sectoral Inflation Persistence and Forward-Looking Expectations

The reduced-form version of the sectoral Phillips curve has an important omission: inflation expectations. At the aggregate level, the reduced-form equation can be augmented with survey measures such as consumer expectations from the Michigan Survey of Consumers. However, no measure of *sectoral* inflation expectations exists. Everything else being equal, a change in the relative weights of forward-looking and backward-looking inflation expectation components should be tightly linked to a change in inflation persistence (Fuhrer 2010).

To differentiate between intrinsic and inherited inflation persistence, we need theory that disciplines structural parameters. Byrne, Kontonikas, and Montagnoli (2013) provide one solution:

[4] The remaining two panels of Figure 3 show specifications that trim the sample differently (Panel I) or exclude imputed rents (Panel J). All specifications produce a similar estimate of the break.

[5] To conserve degrees of freedom in time-series specifications, we allow for a break only in the slope (and not in the intercept). This modeling choice does not have a material effect on our conclusions.

the reduced-form Equation (1) can be derived from a hybrid New Keynesian Phillips curve (NKPC), and one can retrieve the coefficients of the model using a stationary rational-expectations solution. To solve for such an equilibrium, we need to impose stationarity on the driving variable (the unemployment gap).

To illustrate this principle, assume that the driving variable follows a stationary AR(1) process (i.e., the unemployment gap persistence is less than unity). Combining the unemployment gap process with the hybrid NKPC leads to the following model:

$$\pi_{jt} = \lambda_j^b \, \pi_{j,t-1} + \lambda_j^f \, \mathbb{E}_t \, \pi_{j,t+1} + \theta_j \, u_t + e_{jt}, \tag{2}$$

$$u_t = \rho^u \, u_{t-1} + \eta_t, \tag{3}$$

where λ_j^b and λ_j^f are the relative weights of the backward-looking and forward-looking expectation components ($\lambda_j^b + \lambda_j^f = 1$), θ_j is the slope of the hybrid NKPC, and ρ^u is the persistence of the unemployment gap.[6] The stationary rational-expectations solution of this model is

$$\pi_{jt} = \mu_j + \rho_j^\pi \, \pi_{j,t-1} + \beta_j \, u_t + \varepsilon_{jt}, \tag{4}$$

where the parameters are related as follows:

$$\begin{aligned} \lambda_j^b &= \frac{\rho_j^\pi}{1+\rho_j^\pi}, \\ \lambda_j^f &= \frac{1}{1+\rho_j^\pi}, \\ \theta_j &= \beta_j \, (1-\rho^u)\left(1-\lambda_j^b\right). \end{aligned} \tag{5}$$

The parameters governing the relative weights of the backward-looking and forward-looking expectation components (λ_j^b, λ_j^f) depend only on parameters from the reduced-form Phillips curve and do not require estimating the process on the driving variable. In contrast, estimates of the structural slope θ_j also depend on the unemployment gap persistence (the higher the persistence, the lower the slope). We do not attempt to estimate the persistence of the unemployment process, and focus on λ_j^b, which mechanically also informs us about $\lambda_j^f = 1 - \lambda_j^b$.

Our estimation algorithm is as follows: (1) estimate ρ_j^π from Equation (4) for each j; (2) compute λ_j^b using the mapping from Equations (5); (3) obtain the mean group estimate $\hat{\lambda}_{\mathrm{MG}}^b$ and its standard error (Pesaran and Smith 1995, Pesaran 2006).[7]

[6] Imbs, Jondeau, and Pelgrin (2011), Byrne, Kontonikas, and Montagnoli (2013), and others use sectoral marginal cost as the driving variable, allowing it to vary across j. Since we use the unemployment gap instead and data on sectoral unemployment are limited, we deviate from this practice and use a sector-invariant measure u_t. Additionally, we do not control for relative prices, because doing so would complicate the solution and make the results difficult to interpret. Since relative prices have only a small effect on the results in the previous section, this omission likely does not alter our conclusions.

[7] In practice, we allow for five lags of sectoral inflation and compute ρ_j^π as the sum of autoregressive coefficients. We also include additional controls such as the percentage change in the oil price. Furthermore, the baseline mean group

Panel A of Table 2 presents the mean group estimate of the backward-looking component weight (λ^b) in the hybrid NKPC for the entire sample period, as well as for the subsamples split around the Great Recession (2009q3).[8] The weight of the backward-looking (forward-looking) component is lower (higher) after the Great Recession. This result also holds for core sectors. Panel B of Table 2 shows that an increase (decrease) in the weight of the forward-looking (backward-looking) component is observed in many sectors. In summary, it appears that intrinsic inflation persistence has declined.[9]

5 Conclusions

Using disaggregated CPI data, we document a significant decrease in the Phillips curve slope and an increase in the relative weight of forward-looking expectations around 2009–2010. These changes are observed in a wide range of product categories. However, it is worth paying special attention to housing and services such as food away from home, because these sectors contributed the most to the break in a statistical sense. Such categories did not necessarily experience a significant technological revolution, nor were they affected by global trends in prices for tradeable goods. Other sector-specific explanations could be a promising subject of future research.

Acknowledgements

We are grateful to Michelle Barnes, Falk Bräuning, Daniel Cooper, José Fillat, Jeff Fuhrer, Chris Foote, Giovanni Olivei, Ali Ozdagli, Joe Peek, Eric Rosengren, Jenny Tang, Geoff Tootell, and an anonymous referee for insightful comments. Katrina Truebenbach helped us understand disaggregated employment data, and Chloe Lee provided superb research assistance. We also thank Lawrence Bean for excellent editorial assistance. The views in this paper are our own and do not necessarily reflect the views of the Federal Reserve Bank of Boston or the Federal Reserve System. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

Byrne, Joseph P., Alexandros Kontonikas, and Alberto Montagnoli. 2013. "International Evidence on the New Keynesian Phillips Curve Using Aggregate and Disaggregate Data." *Journal of Money, Credit and Banking* 45(5): 913–932.

estimator uses sectoral weights, but our conclusions also hold with equal weights.

[8]In Section 3, we identified a structural break in the slope of the Phillips curve β in 2010q3. However, the F-test used to select the date was elevated throughout the Great Recession. Since we split the sample into two periods for the analysis in this section, we start the second period earlier, in 2009q3, to gain a few observations. This slight inconsistency does not affect our conclusions.

[9]Luengo-Prado, Rao, and Sheremirov (2017) show that unconditional persistence, obtained by estimating univariate autoregressive models of sectoral inflation, also declined.

Carvalho, Carlos. 2006. "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks." *The B.E. Journal of Macroeconomics* 6(3): 1–58.

Fuhrer, Jeffrey C. 2010. "Inflation Persistence." In *Handbook of Monetary Economics*, vol. 3, edited by Benjamin M. Friedman and Michael Woodford, chap. 9, pp. 423–486. Amsterdam: Elsevier.

Imbs, Jean, Eric Jondeau, and Florian Pelgrin. 2011. "Sectoral Phillips Curves and the Aggregate Phillips Curve." *Journal of Monetary Economics* 58(4): 328–344.

Luengo-Prado, María José, Nikhil Rao, and Viacheslav Sheremirov. 2017. "Sectoral Inflation and the Phillips Curve: What Has Changed since the Great Recession?" Current Policy Perspectives 17-5, Federal Reserve Bank of Boston.

Pesaran, M. Hashem. 2006. "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure." *Econometrica* 74(4): 967–1012.

Pesaran, M. Hashem, and Ronald P. Smith. 1995. "Estimating Long-Run Relationships from Dynamic Heterogeneous Panels." *Journal of Econometrics* 68(1): 79–113.

Reis, Ricardo, and Mark W. Watson. 2010. "Relative Goods' Prices, Pure Inflation, and the Phillips Correlation." *American Economic Journal: Macroeconomics* 2(3): 128–157.

Figure 1. Disaggregated Inflation Rates: Categories Selected for Our Analysis

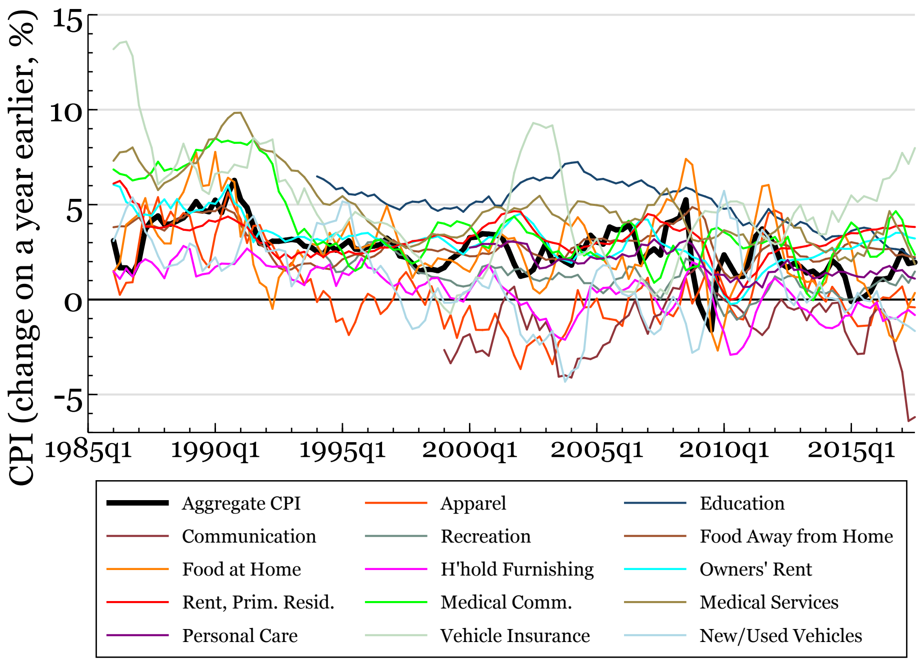

Source: BLS.
Note: To enhance visibility, “Motor Fuel” (Transportation) and “Fuels and Utilities” (Housing) are excluded from the graph (but included in our analysis).

Figure 2. Structural Break Test

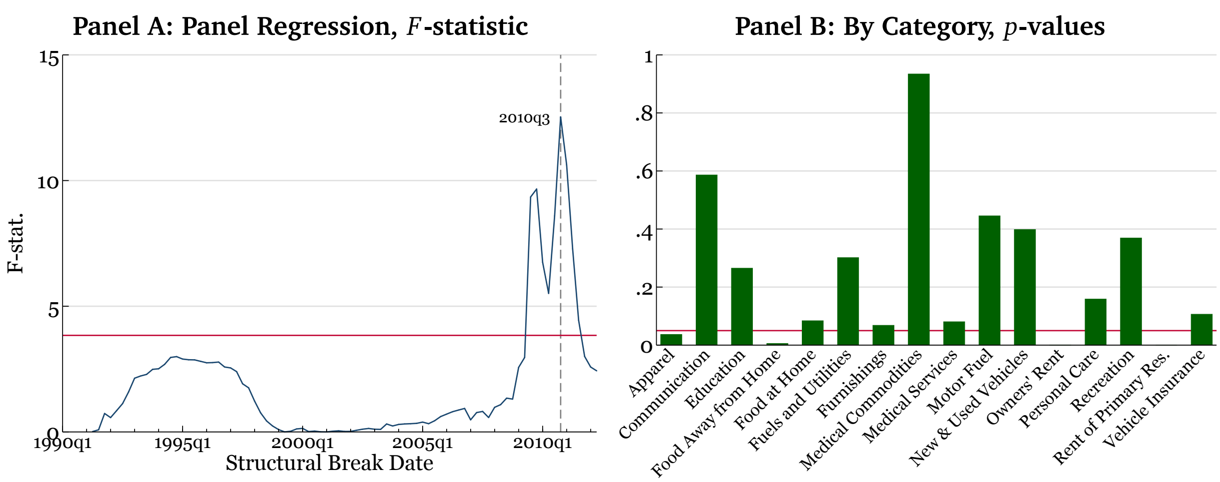

Source: Authors’ calculations.
Notes: Panel A shows the F-statistic (vertical axis) from the test that the change in the unemployment gap coefficient is zero for a range of potential break points (horizontal axis) obtained from panel regressions of sectors. High values of F indicate a change. Panel B shows p-values for a similar test run for each sector separately and a break date that maximizes the F-statistic in the left panel (2010q3). Sample period: 1986q1–2017q3. Dependent variable: the change in the quarterly seasonally adjusted sectoral CPI from the previous year. Controls: sectoral fixed effects, five lags of dependent variable, percentage change in the price of oil. The left panel allows for a break in the sectoral fixed effect and applies sectoral weights. Newey–West standard errors are used to calculate p-values. The horizonal red lines correspond to a 5 percent significance level.

Figure 3. Panel Regression, F-statistic: Robustness Checks

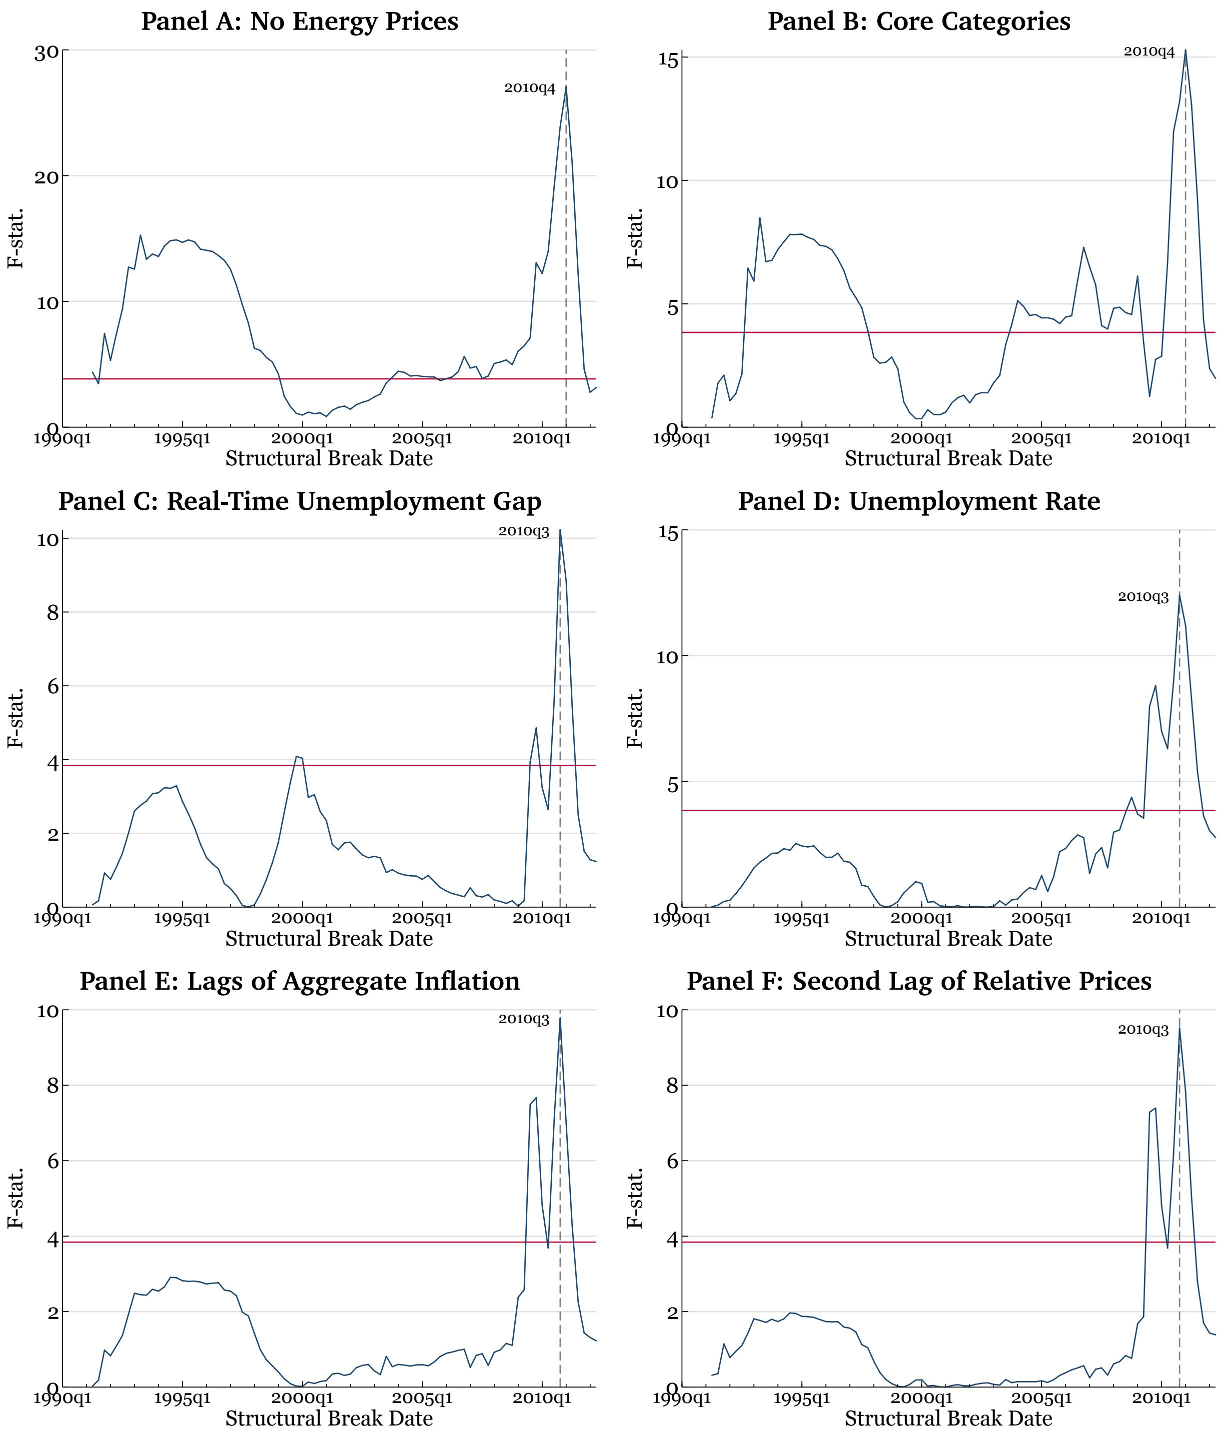

Figure 3. ***(cntd.)*** **Panel Regression, F-statistic: Robustness Checks**

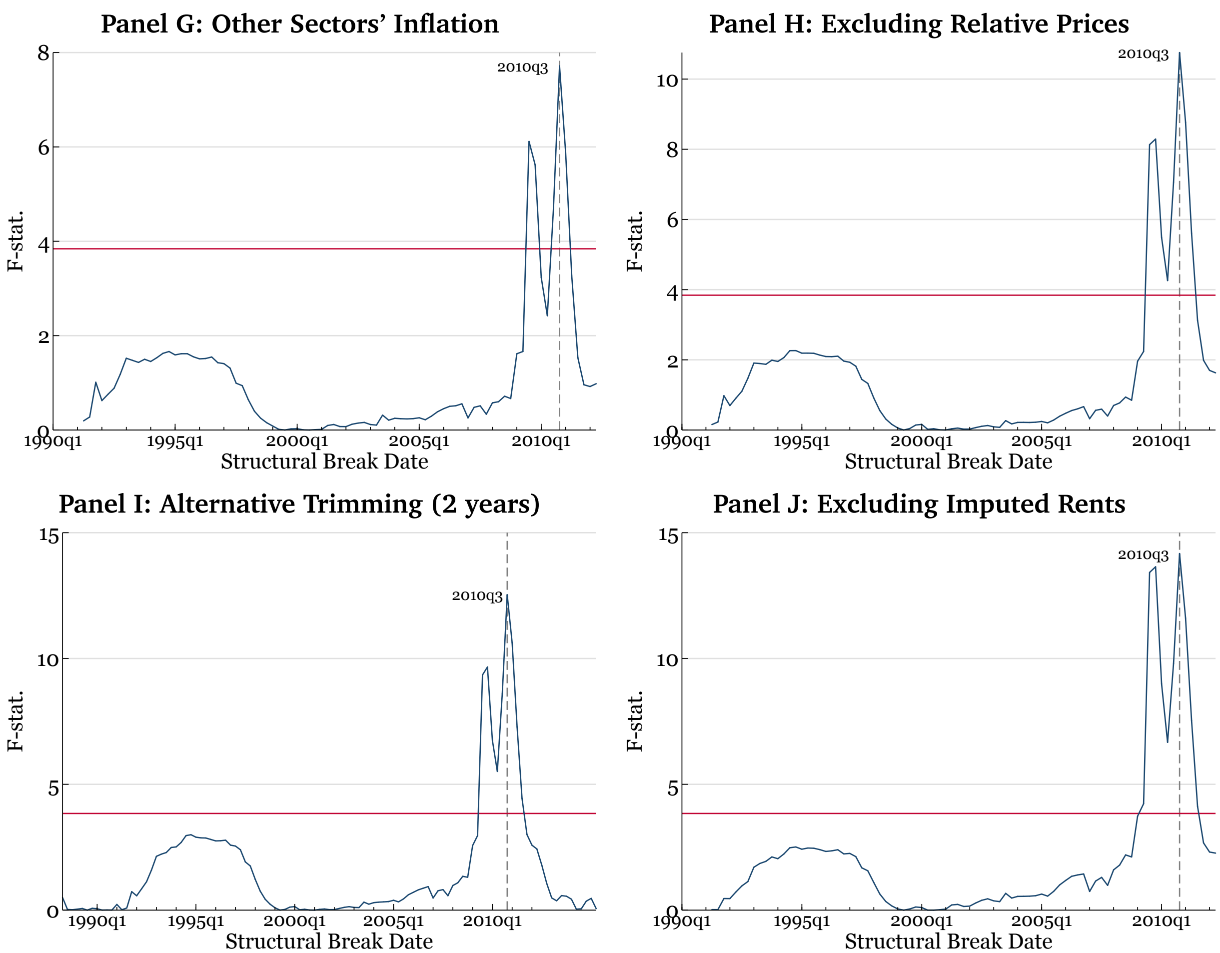

Source: Authors' calculations.

Figure 4. How Much Higher Sectoral Inflation Would Be without a Phillips Curve Flattening?

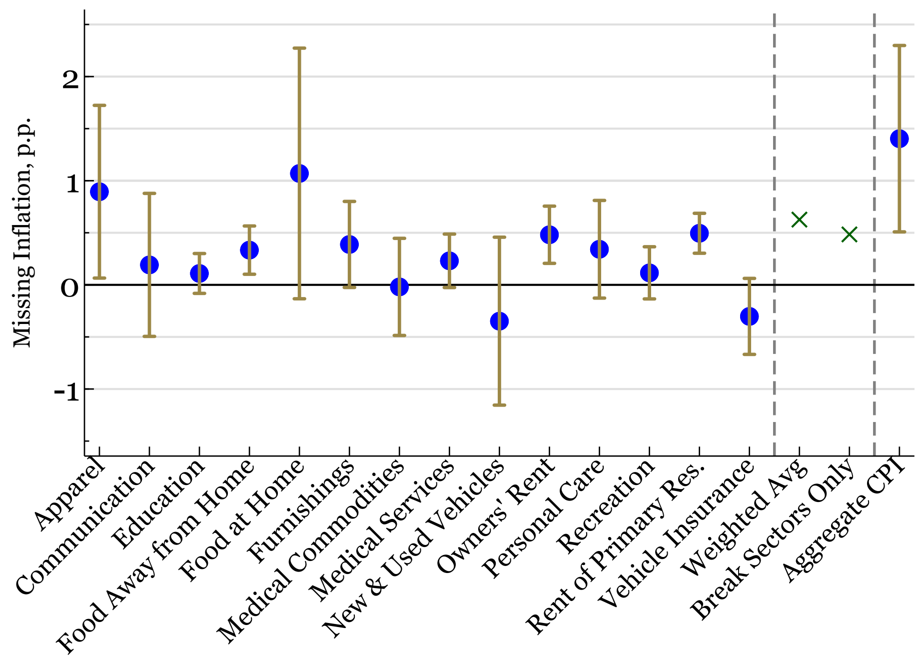

Source: Authors' calculations.

Notes: Blue dots show a missing increase in 2017q3 sectoral inflation due to a Phillips curve flattening, calculated based on sectoral Phillips curves using time-series (all but the last dot) and panel (last dot for aggregate CPI) analyses, and the bars represent the corresponding 95 percent confidence bands. The × mark is for the weighted average of sectoral inflation rates with current CPI weights. The panel and weighted average estimates also include "Motor Fuel" (Transportation) and "Fuels and Utilities" (Housing), which were dropped from the graph for visibility. "Break Sectors Only" include "Food Away from Home," "Owners' Rent," and "Rent of Primary Residence."

Table 1. Structural Break Tests: Slopes of Sectoral Phillips Curve

	Slope before 2010q3, β_j (1)	Change in slope, $\widetilde{\beta}_j$ (2)	Obs., N (3)
Panel A: Sector-Specific Slopes from Time-Series Regressions			
Apparel	0.011	0.185**	122
	(0.052)	(0.088)	
Education	−0.016	0.023	90
	(0.015)	(0.020)	
Communication	−0.017	0.040	70
	(0.076)	(0.073)	
Recreation	−0.053**	0.024	90
	(0.024)	(0.026)	
Food Away from Home	−0.071***	0.069***	122
	(0.019)	(0.024)	
Food at Home	−0.165**	0.221*	122
	(0.077)	(0.127)	
Fuels and Utilities	−0.353**	0.204	122
	(0.177)	(0.196)	
Household Furnishings and Operations	−0.056**	0.080*	122
	(0.027)	(0.044)	
Owners' Equivalent Rent	−0.106***	0.100***	122
	(0.023)	(0.029)	
Rent of Primary Residence	−0.136***	0.103***	122
	(0.019)	(0.020)	
Medical Commodities	−0.013	−0.004	122
	(0.035)	(0.049)	
Medical Services	−0.034	0.048*	122
	(0.024)	(0.027)	
Personal Care	−0.101***	0.071	66
	(0.022)	(0.050)	
Motor Fuel	−1.194	1.332	122
	(0.875)	(1.740)	
Vehicle Insurance	0.028	−0.063	122
	(0.029)	(0.039)	
New and Used Vehicles	0.189**	−0.072	122
	(0.077)	(0.085)	
Panel B: Common Slope from Panel Regressions			
All categories (sectoral weights)	−0.179***	0.291***	1,780
	(0.043)	(0.082)	
All categories (equal weights)	−0.278***	0.394***	1,780
	(0.083)	(0.109)	
All categories except fuel (sectoral weights)	−0.067***	0.104***	1,536
	(0.010)	(0.020)	
Core categories only (sectoral weights)	−0.053***	0.077***	1,292
	(0.010)	(0.020)	

Source: Authors' calculations.
Notes: The table shows the unemployment gap coefficient before the break (defined as 2010q3) and the change in this coefficient after the break obtained from time-series regressions estimated separately for each sector (Panel A) and in a panel of sectors (Panel B). A positive change indicates a Phillips curve flattening. Sample period: 1986q1–2017q3. Dependent variable: change in the quarterly seasonally adjusted sectoral CPI from the previous year. Controls: five lags of the dependent variable and the percentage change in the price of oil. Newey–West standard errors in parentheses. ***, **, and * represent 1 percent, 5 percent, and 10 percent significance levels, respectively.

Table 2. Intrinsic Inflation Persistence: Backward- vs. Forward-Looking Expectations in Hybrid NKPC

Sector	Full sample (1)	'86q1–'09q3 (2)	'09q4–'17q3 (3)
Panel A: Mean-Group Aggregate Estimates, $\hat{\lambda}^b_{\mathrm{MG}}$			
All categories	0.464^{***}	0.469^{***}	0.386^{***}
	(0.006)	(0.008)	(0.015)
All categories except fuel	0.468^{***}	0.475^{***}	0.388^{***}
	(0.005)	(0.006)	(0.016)
Core categories only	0.473^{***}	0.481^{***}	0.390^{***}
	(0.004)	(0.003)	(0.019)
Panel B: Sectoral Weight of Backward-Looking Expectations, $\hat{\lambda}^b_j$			
Owners' Equivalent Rent	0.470	0.481	0.435
Rent of Primary Residence	0.460	0.469	0.331
Food at Home	0.422	0.417	0.394
Medical Services	0.494	0.504	0.272
New and Used Vehicles	0.471	0.480	0.450
Food Away from Home	0.474	0.481	0.350
Recreation	0.460	0.472	0.346
Fuels and Utilities	0.437	0.444	0.317
Household Furnishings and Operations	0.464	0.460	0.330
Communication	0.480	0.494	0.395
Motor Fuel	0.398	0.356	0.437
Education	0.504	0.493	0.492
Apparel	0.466	0.476	0.334
Vehicle Insurance	0.488	0.486	0.452
Personal Care	0.465	0.466	0.330
Medical Commodities	0.491	0.495	0.422

Source: Authors' calculations.
Notes: Panel A presents mean group estimates (with sectoral weights) of the relative weight of the backward-looking expectation component in the hybrid NKPC described in Equation (2). Panel B presents coefficients by sector. The estimated regressions include five lags of sectoral inflation, the contemporaneous change in the oil price, and consumer inflation expectations from the Michigan survey. Rows are ordered by sectoral weights. ***, **, and * represent 1 percent, 5 percent, and 10 percent significance levels, respectively.