The Nonlinear Unemployment-Inflation Relationship and the Factors that Define It*

Andrew Keinsley† — Weber State University

Sandeep Kumar Rangaraju‡ — Weber State University

Abstract

In this paper, local projections are used to explore the unemployment-price inflation relationship in a more generalized, empirical fashion. We find that inflation's reaction to changes in the unemployment rate varies across states of the economy, with timing as the primary difference. In low-unemployment environments, inflation reacts immediately and persistently. In high-unemployment environments, the same reaction manifests after a one-year lag. We then use industry-level data and a two-stage feasible generalized least squares (FGLS) method to explore the factors that drive this relationship. We find that increased reliance on labor and intermediate inputs adds to the inflationary pressures of low unemployment, while increased market concentration (among others) dampens this effect.

JEL Classification: E24, E31

Keywords: Unemployment, Inflation, Local Projections, Nonlinear Responses

*We are thankful to Olivier Blanchard, David Blanchflower, Ana María Herrera, John Keating, Sandeep Mazumder, and A. Lee Smith for comments on earlier drafts. We have also benefitted from the many helpful comments provided by participants at economic conferences and seminars. All remaining errors are ours.

† *Corresponding Author* – Email: andrewkeinsley@weber.edu – Address: Department of Economics WB 226, 1337 Edvalson St, Ogden, UT 84408-3807

‡ Email: srangaraju@weber.edu – Address: Department of Economics WB 226, 1337 Edvalson St, Ogden, UT 84408-3807

1 Introduction

The unemployment-wage inflation relationship described by Phillips (1958) and its unemployment-price inflation adaptation have faced an existential crisis in recent years. The price inflation version especially, which serves as a pillar of macroeconomic theory, seems to no longer hold. After the 2007-2009 recession, for instance, neither core nor headline inflation settled around the targeted rate on a sustained basis until 2018, despite a four percentage point reduction in the unemployment rate since 2012.[1] The accumulation of low-unemployment/low-inflation observations during the 2014-2019 period has led many in the economics community to part with the Phillips curve, but what has been driving this apparent change?

There are multiple schools of thought when it comes to the evolution of the Phillips curve. Some argue that expected inflation has become extremely well-anchored in recent decades (i.e. Blanchard, 2016; IMF, 2013), which reduces the weight on the real variable of choice.[2] In a similar light, others have considered the impact of nominal wage rigidities on the shape of the Phillips curve (i.e. Daly and Hobijn, 2014). Still others cite shortcomings in the econometric techniques (i.e. Clark and McCracken, 2006) or even that simply exchanging input variables can go a long way to salvaging this theory (i.e. Coibion and Gorodnichenko, 2015). Across all of these potential solutions, however, few have ventured far from some version of the standard Phillips curve

$$\pi_t = \omega + \alpha \mathbb{E}_t(\pi_{t+1}) - \beta(x_t - \bar{x}) + \varepsilon_t,$$

where π_t is inflation and $(x_t - \bar{x})$ is the gap in the real variable of choice, despite the fact

[1] As measured by the personal consumption expenditures chain-type price index, which is believed to be the preferred metric of the Federal Reserve based on Humphrey-Hawkins reports.

[2] Blanchard (2016) considers the unemployment rate, though many empirical estimates are of models derived from theory, which considers the output gap as the real variable of choice. As they are linked by Okun's Law, the discrepancies are often ignored.

that empirical macroeconomics is rooted in Sims' (1980) vector autoregression paradigm. Restricting analysis in this manner hamstrings broader understanding of the underlying connections.

We use Jordà (2005) local projections to analyze the unemployment-inflation relationship; allowing for richer dynamics, nonlinearity, and a broad set of control variables that lead us to three key results. First, in low-unemployment environments, further unanticipated reductions in the unemployment rate produce immediate inflation that is persistent for 8-12 months after the shock. Second, the same persistent effect takes a full year to materialize in high-unemployment environments. These two findings reveal that unemployment and inflation are still connected; though this relationship is nonlinear, potentially lagged, and subject to multiple outside factors as well. Lastly, we show that controlling for broader unemployment metrics is important for extracting this result, though we do not find a substantial direct effect of these broader metrics on inflation. This serves as a reminder that the headline unemployment rate is a simple quantity metric and that there are other aspects of the labor market that can influence inflation as well. In total, these aggregate-level results shed new light on the dynamics of the Phillips-curve relationship.

Exploring the roots of our aggregate findings, we use industry-level data and a two-stage feasible generalized least squares (FGLS) method to pinpoint which factors define this unemployment-inflation connection. First, a greater reliance on labor and intermediate inputs increases the sensitivity of inflation to changes in the national unemployment rate in a low-unemployment environment. Both of these channels are found to have immediate effects and are operable for nearly a year after the shock. This implies that, in tight labor markets, the wage-price mechanism is well-established and that these pressures build along the supply chain. Second, increased sectoral output and market concentration dampen the unemployment-inflation passthrough in the same low-unemployment environments. These channels are also quite persistent, though the market concentration channel is less so and operates with a lag. Combining the negative market concentration response with the positive labor input response suggests that monopsony (not monopoly) power is a contributing factor in this mechanism. If the wage-price connection is strong (labor input), and rising market concentration reduces the response of inflation to changes in the unemployment rate, then there must be a breakdown in the unemployment-wage component of the mechanism. This conclusion matches some findings in the literature, such as those of Webber (2018), who outlines an intuitive source of nonlinearity in his finding that hiring tendencies differ between competitive and monopsonistic firms.

Our analysis contributes to the extensive literature and our general understanding of this macroeconomic dynamic in two key ways. First, we show that the primary difference between the inflationary response to movements in unemployment across states of the economy is timing. There are only a few works that consider the potential for nonlinearity or a more dynamic model structure (e.g. Debelle and Laxton, 1997; Doser *et al.*, 2017; Gagnon and Collins, 2019). Thus, this is one of a small number of papers to apply what has become the rule of empirical macroeconomics to this theoretical cornerstone. Second, to our knowledge, we are the first to analyze the impact of industry-specific characteristics on unemployment-inflation pass-through.

The primary policy implication that stems from this study focuses on the use of the aggregate relationship versus the underlying sectoral factors. Across our sample, we find that changes in the unemployment rate pass through to inflation in one way or another, but more recent trends in the industry-level factors suggest that they may be of more use to policy makers. These factors seem to provide more insight into what to expect as labor markets tighten. For instance, declining labor income share, increasing market concentration, and further drift towards a services- and information-based economy would result in an erosion of the aggregate-level channel.

The remainder of this paper is structured in a straight-forward manner. In Section 2 we introduce our identification strategy and use of Jordà (2005) local projections to estimate the impulse response functions of focus. The results of these local projections are presented in Section 3. Section 4 outlines the methodology and results of the sectoral analysis, which we use to explain our baseline, aggregate result. Robustness checks can be found in Section 5. Section 6 concludes.

2 Methodology

To investigate whether the unemployment shock exhibits a differential impact during periods of low and high unemployment, we follow Jordà's (2005) local projection method to compute impulse responses of our macro variables. This still provides us with a single equation model, but offers more in terms of dynamics, flexibility, nonlinearities, and rigor.

There are several advantages in using the local projections versus conventional vector autoregressive (VAR) models.[3] First, impulse responses are estimated directly by linear least squares regressions. Second, the local projection model is more robust to misspecification biases than the conventional VAR estimates, where we would impose implicit dynamic restrictions on the variables. Third, point-wise and joint inference for the impulse responses is very straightforward and it does not require the delta method. Fourth, unlike a regime-switching VAR model, the non-linear local projection method does not require assumptions on the duration of a given state or on the mechanism triggering the transition between regimes. Hence, the estimated coefficients here will represent the

[3] Plagborg-Møller and Wolf (2019) shows that the impulse responses estimated in VAR and linear local projection methodology are same for population.

average effects of unemployment shocks depending on the initial state, and capture the possible change in the projected horizon (see Jordà, 2005; Ramey and Zubairy, 2018; Alpanda *et al.*, 2019; Kilian and Lütkepohl, 2018).

2.1 Identification of Unanticipated Unemployment Rate Shocks

The first step in this analysis is identifying unanticipated shocks to the variable of choice: the unemployment rate (u_3). Using a simple forecasting model

$$u_{3,t} = \Phi(L)X_t + \varepsilon_t,$$

where X_t consists of twenty monthly variables and ε_t are the forecast errors, this identification strategy generally follows with that of Auerbach and Gorodnichenko (2013).[4] In this setup, we consider three lags of all twenty macroeconomic as well as contemporaneous values of all variables outside of the unemployment rate itself. The variables considered can be found in Section A of the online appendix. This model is estimated over the 1970:03–2018:12 period. Controlling for both lagged and contemporaneous values of the variables means that the forecast errors represent unanticipated shocks to the unemployment rate, which can be seen in Figure 1. While not the only way to identify shocks to the unemployment rate, we show in Section 3 that our strategy seems to do the job well.

A couple comments must be noted before we continue. First, our identification scheme is not state-dependent, though we do use these shocks in state-dependent models. Second, our unanticipated shocks are not separated in to structural supply-side or demand-side innovations. If we were trying to test the slope of the Phillips curve directly, we

[4] While there are one-step ahead forecasts of u_3, the longest-running series of relevant forecasts (Survey of Professional Forecasters) is a quarterly series, while we are attempting to work with monthly data here.

would need to see how inflation responds contemporaneously to a demand-side shock alone. However, we are interested in the dynamic response of inflation to an unanticipated fall in the unemployment rate. Thus, the source of the shock is not as important for our research purposes.[5] Even when estimating a structural vector autoregression model, identification of shocks to the included variables is more likely to be done via dynamic exclusion restrictions (e.g. Cholesky decompositions) than in taking a stand on the source of the shocks. Many in the literature use sign-restrictions to separate supply from demand shocks, though these tend to be less informative due to the small-scale models used (see Uhlig, 2005; Fujita, 2011; Fry and Pagan, 2011; Kilian and Murphy, 2014; Kilian and Lütkepohl, 2018). In addition, the conventional priors used in estimating a sign-identified model may unduly influence the posterior of the structural impulse response. The use of sign-identified models comes at cost of imposing sign restrictions on responses and bounds on parameter estimates. Hence, we might gain in one dimension but may be lose in other dimensions (see Kilian and Lütkepohl, 2018; Herrera and Rangaraju, 2020). Thus, we do not take a stand on where these shocks originate, but merely consider how the economy responds to a random decrease in the unemployment rate.

2.2 Linear Model

In our initial examination of the unemployment-inflation relationship, we consider the linear model

$$z_{t+h} = c_t + \Phi_h(L)y_t + \beta_h \varepsilon_t + \epsilon_{t+h}. \tag{1}$$

Estimating via local projections allows us to incorporate additional control variables y_t and provides additional dynamics that are often missing from traditional Phillips curve

[5] Our identification scheme eliminates as much endogeneity as possible, but comes at a cost of structural identity.

analysis. The variable of interest z_{t+h} is forecasted at horizon $t+h$ for some $h \geq 0$. The variable-specific innovations are represented by ε_t, ϵ_{t+h} depicts the aggregated model error term, and the estimated coefficients β_h will form the impulse response functions. Lastly, the term c_t is a vector that includes a constant, linear trend, and quadratic trend.[6] The results of this model will be used as a point of reference when evaluating the nonlinear model below.

For our baseline analysis, we are interested in what drives the personal consumption expenditures (PCE) inflation rate excluding food and energy. Going forward, we simply refer to this as "core PCE inflation." The control variables consist of lags of the unemployment rate u_3, the spread between the u_6 and u_3 unemployment measures, the core PCE inflation rate, the Congressional Budget Office's estimate of the natural rate of unemployment, the five-year expected inflation rate derived from the University of Michigan's Survey of Consumers, the percent change in WTI oil prices, and the annualized growth rate of the end-use import price index (excluding food and energy) adjusted for imports share of gross domestic product.[7] While we consider six lags of each of these variables, we also control for the contemporaneous percent change in oil prices, given its general importance and volatility. All series in this baseline analysis are monthly and span the 1994:01–2018:10 period.[8] Together, these variables constitute a majority of what is believed to drive inflation, allowing us to isolate the unemployment-inflation channel.

A couple of explanations are needed regarding our initial choice of control variables.

[6] Under this specification the deterministic term in (1) is $c_t = \begin{bmatrix} c_0 & c_1 & c_2 \end{bmatrix} \begin{bmatrix} 1 & t & t^2 \end{bmatrix}'$

[7] While other measures of the expected inflation rate exist, the University of Michigan's Survey of Consumers has the longest history at the needed frequency. Additionally, household expectations have been shown to be the better metric for this type of analysis (Coibion and Gorodnichenko, 2015), explaining the "missing deflation" episode during the financial crisis. Lastly, imports share of GDP is linearly extrapolated from its quarterly publication to the monthly frequency needed.

[8] Note that the "marginally attached to the labor force" variable was first added to the Current Population Survey in 1994. However, this does limit the scope of our analysis to the period in which the literature claims the Phillips curve has been flattening, making our results stand out that much more.

First, Coibion and Gorodnichenko (2015) find that 12-month household expectations are highly correlated with movements in oil prices. Since controlling for both is generally desired, we use the 5-year inflation expectations variable instead of its 12-month counterpart. This should allow us to separate out the effects of movements in energy prices from other factors influencing these expectations as well while simultaneously reducing any potential multicollinearity problems. Second, our inclusion of the spread between the broadest measure of unemployment provided by the Bureau of Labor Statistics u_6 and the headline u_3 rate is meant to capture non-quantity movements in the labor market that may influence inflation. For example, movements from part-time to full-time labor do not show up in u_3, but can influence the pricing decisions of firms. We test the robustness of these initial assumptions across Sections 3.2 and 5 herein, as well as Section B in the online appendix; providing additional insight along the way.

2.3 Nonlinear Model

We follow the approach taken by Ramey and Zubairy (2018) to modify (1) to a state-dependent model in order to investigate whether the unemployment-inflation relationship exhibits different dynamics during periods of low and high unemployment environment. The non-linear model can be written in a single equation

$$z_{t+h} = c_t + d_{t-1}[\alpha_{b,h} + \Phi_{b,h}(L)y_t + \beta_{b,h}\varepsilon_t] + (1 - d_{t-1})[\alpha_{g,h} + \Phi_{g,h}(L)y_t + \beta_{g,h}\varepsilon_t] + \epsilon_{t+h}, \quad (2)$$

where the definitions and control variables in (1) still hold. In addition, the dummy variable d_t takes the value of one when the $u_{3,t}$ is above the threshold $\bar{u}$ and zero when the civilian unemployment rate $u_{3,t}$ is below the threshold $\bar{u}$. All the coefficient estimates except the constant, linear trend, and quadratic trend terms vary depending on the state of the economy. The coefficient estimates are labeled with a g subscript when $u_{3,t} < \bar{u}$ to denote the low-unemployment state, or periods of relative non-slack, and a b subscript when $u_{3,t} > \bar{u}$ to denote the high-unemployment state, or periods of relative slack. Again, the estimated coefficients $\beta_{n,h}$ for $n \in \{b, g\}$ will form the impulse response functions of interest going forward.

Threshold Choice An intuitive choice of threshold $\bar{u}$ would be the natural rate of unemployment, thus allowing the threshold to vary over time. However, doing so results in 63% of the observations lying in the high-unemployment subsample and only 37% in the low-unemployment subsample. In an attempt to balance the observations, we fix the threshold value to 5.1%. This is approximately equal to both the mean and median values of unemployment during the 1994:01-2007:12 subsample. The severity of the 2008-2009 recession increases the average unemployment rate to 5.8% and median to 5.4% for the full sample. At 5.1%, this threshold unemployment rate provides a good cutoff for an analysis of what is historically considered a "tight" labor market and falls generally in the middle of the natural rate of unemployment, which decreases almost linearly from 5.5% to 4.5% during our sample. Running the analysis with the time-varying threshold does not alter the general dynamics of our results, but the small sample size does significantly increase the standard errors in our low-unemployment results.

3 Empirical Results

In this section, we report the results of the models outlined in Section 2. First, we check our identification strategy by subjecting the unemployment rate to its own, unanticipated shocks. We then estimate both the linear and state-dependent impulse response functions (IRFs) represented by (1) and (2), respectively. All estimated IRFs are shown for a horizon of 20 months after the shock and include both 68% and 95% confidence intervals, which are computed via Newey-West standard errors to correct for serial correlation. In all cases, we use negative shocks to unemployment since that should—theoretically—be inflationary. Together, this forms the basis of our analysis, allows us to evaluate a more generalized version of the Phillips curve relationship from multiple perspectives, and paints a picture that is consistent with both the classical theory and the contemporary puzzles in the process.

Since our identification scheme is not state-dependent, we use the unemployment rate as the dependent variable in the linear model (1) and analyze the impact of its own, identified shocks. This is simply an exercise in making sure we are properly identifying exogenous, unanticipated shocks. Figure 2 shows the reaction of unemployment to its own, negative shock. As can be seen, this is a temporary shock, though there is a slight oscillation to it initially. Thus, it implies that our identification process was generally successful in generating unanticipated, exogenous shocks.

3.1 Linear Effects of Unemployment Shocks

Figure 3 presents the response of core PCE inflation to a decrease in the headline unemployment rate in the linear model. Even with the additional dynamics and control variables, our findings generally follow the growing narrative of a flat Phillips curve. That said, we do find a positive initial impact at lower confidence levels. Generally speaking, though, it is understandable that the standard, linear models used in Phillips curve analyses would find little correlation between the unemployment rate and inflation.

3.2 State-Dependent Effects of Unemployment Shocks

The estimated IRFs for the nonlinear model (2) reveal two important dynamics. First, a decrease in the unemployment rate in a low-unemployment environment (Figure 4, right panel) has a positive, persistent, and statistically significant effect on inflation. Though core PCE inflation is meant to capture more trend-like movements in inflation, we find that these innovations to unemployment in tight labor markets produce immediate inflation that persists for 8-12 months thereafter. Second, the same shock still produces a vigorous and persistent inflationary response in the high-unemployment environment, but the effect takes twelve months to materialize (Figure 4, left panel). Together, these results show that the pass-through from unemployment to inflation persists, but the timing of the mechanism depends on the state of the economy at the time of the shock.

Ball and Mazumder (2019) find that the Phillips curve relationship becomes more clear as the focus moves towards more stable, less noisy measures of inflation. To this end, we replace the core PCE inflation rate with the trimmed-mean PCE inflation rate published by the Federal Reserve Bank of Dallas.[9] Instead of simply removing the individual food and energy price components from the index, this measure truncates the price distribution, removing the more extreme sectoral movements found in the tails. As can be seen in Figure 5, the general narrative matches that of the baseline analysis (Figure 4), but there are some key differences. First, in the high-unemployment state, the negative effect from the u_3 shock is now statistically significant, suggesting that the disinflationary effects of recessions takes are more persistent. This matches the timing seen in the data where unemployment begins rising before inflation falls at the start of a recession. Thus, our high-unemployment analysis may simply be picking up the initial dynamics of a recessionary event.[10] Second, in a low-unemployment environment, the positive inflationary effect does not manifest itself for up to a quarter after the shock, which is to be expected from a more trend-oriented variable. In general, the results for this inflation metric promotes our finding that the Phillips curve alive and well in a

[9] Alternatively, Ball and Mazumder (2019) consider a weighted-median PCE inflation rate.

[10] Sahm (2019) notes that the fast-moving dynamics of the unemployment rate at the beginning of US recessions make it a good recession indicator and trigger mechanism for a proposed automatic stabilizer policy.

nonlinear sense, but also adds the possibility that a large swath of industries may see a quarter-long lag in their inflation response.

4 Which Industry-Level Factors Drive This Relationship?

In this section we explore which characteristics drive the variation in PCE inflation rate responses. To begin, we consider 30, three-digit NAICS industries and the PCE inflation components that best map to those sectors directly. The pairings can be found in Table 1.[11] We estimate each of the sectoral PCE inflation reactions via the nonlinear model (2) and collect the cumulative impulse responses for use as the dependent variables in this analysis.[12] Then, following the methodology of Carlino and DeFina (1998, 1999), Bachmeier and Cha (2011), and Herrera and Rangaraju (2019), we regress these responses against industry-level data in an attempt to identify how the unemployment-inflation passthrough occurs. Our results shed light on where we are more likely to see higher sensitivities and perhaps why the Phillips curve has seemed to flatten over time.

4.1 Empirical Methodology

We conduct multiple iterations of this analysis, with the dependent variable considered being the respective month's cumulative response of the sector-specific PCE inflation rates, respectively. These cumulative responses are estimated using the nonlinear model represented in (2). To account for estimation uncertainty and control for the fact that sampling uncertainty may not be constant across sectors, we follow Lewis and Linzer

[11] The matches we propose are a best-fit attempt, but since we are combining industry-level data with the price indices of the resulting consumer goods/services, the categories cannot be expected to match perfectly.

[12] For space considerations, the cumulative impulse response functions for the 30 nonlinear model estimations are not reported here, but are available upon request. These are then used as the dependent variables in our sectoral estimation (3).

(2005) and Guisinger, Hernandez-Murillo, Owyang and Sinclair (2018) in constructing the two-stage feasible generalized least squares estimator (FGLS).

Let y_i be a dependent variable that is not directly observable. Instead we observe an estimate

$$y_i^* = y_i + e_i,$$

where $E(e_i) = 0$ and $Var(e_i) = w_i^2$ for each sector i. Our objective is to estimate the cross-sector equation

$$y_i = \beta x_i + \epsilon_i,$$

where x_i is a vector of sector-level covariates and ϵ_i is a homoskedastic error term. The covariates of interest include sectoral output, labor input, capital intensity, intermediate input, multi-factor productivity and 4-firm concentration across all 30 industries. Since y_i is unobservable, it is only feasible to estimate

$$y_i^* = \beta x_i + v_i, \tag{3}$$

where $v_i = \epsilon_i + e_i$ and y_i^* is the estimated cumulative response of the sectoral-level PCE inflation rate for each sector i in each month of the IRF analysis. $Var(e_i) = \omega_i^2$ is obtained from the nonlinear model (2). The unbiased estimator of the variance of β is computed as

$$\sigma^2 = \frac{\Sigma\widehat{v}_i^2 - \Sigma\omega_i^2 + tr((\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{G}\mathbf{X})}{N-K}.$$

where $\mathbf{X}$ is the matrix of covariates, $tr(\cdot)$ is trace operator, and $\mathbf{G}$ is the diagonal matrix with ω_i^2 as the i^{th} diagonal element.[13] Estimation of (3) requires weighted least squares

[13] As is tradition, the $\widehat{a}$ designation represents an estimated value of a given a.

with weights θ_i given by

$$\theta_i = \frac{1}{\sqrt{\widehat{\sigma^2} + \omega_i^2}}.$$

In Section 3.2, the computed impulse responses represented the average responses over the sample. To control for sectoral composition in the PCE inflation rate, we use the annual average of sectoral output, labor input, capital intensity, intermediate input, and multi-factor productivity between 1994 and 2018. These averages are computed using industry-level multi-factor productivity tables (KLEMS) from the Bureau of Labor Statistics (BLS). For clarity, the annual average of labor input index reflects the combination of labor composition index and hours at work. The intermediate input index represents the goods and services that are used in the production of other goods and services and not sold in the final demand markets. Additionally, we include the four-firm concentration metric from the 2012 Census. This is computed as the precent of total industry revenue generated by the industry's four largest firms, providing us with a measure of market concentration in each of 30 three-digit NAICS industries.

4.2 Estimation Results

The period-by-period sectoral results are reported in Figures 6 and 7. Each shows the estimated coefficients (circles) and 95% confidence intervals (gray bars) for each of the six industry-level variables and a constant.[14]

In the low-unemployment environment (Figure 6), we find both positive and negative influences as well as a general trend in the dynamics of the process. First, as the labor market tightens, the increased use of labor and intermediate inputs in production results in immediate, positive, and sustained sensitivity to changes in the national unemploy-

[14] To save space and make the figures easier to read, we do not report the results pertaining to the constant in the regressions.

ment rate. This suggests that the wage-inflation channel is strong in labor-dependent industries, but also that these tightening labor market pressures accumulate along the supply chain, leading to price pass-through to consumers. Second, increased exposure to outside-industry demand and higher degrees of market concentration have negative impacts on the inflationary response to a tightening labor market. The former could be due to a situation in which, when selling to a wide audience of consumers, a company faces increased competitive pressure, which would hurt their ability to pass along the pressures of a tight labor market to their prices. This channel of inflationary pressure, however, requires more exploration, which we leave to future research. Lastly, we find that, of those variables that have a statistically significant impact, the effects start out quite small and build over time, peaking in the 8- to 9-month range. The only variable for which this does not hold is the 4-firm concentration ratio, which has a more temporary impact than the others, and comes with a 5-month lag. Individually, these results provide key insights into what makes inflation more or less sensitive to changes in unemployment at the aggregate level.

Combining some of these results provides some insight into the mechanism at play. Take the negative association between the inflationary response and firm concentration, for example. This result is somewhat surprising, as one would expect that increased concentration would lead to more price-making power. But when combined with the positive labor input result, this negative concentration effect suggests a monopsony dynamic rather than a monopoly dynamic. That is, the breakdown in this channel seems to not lie necessarily in the wage-inflation part of the mechanism, but rather in the unemployment-wage portion. This would align with the findings of Webber (2018) and others who find that the market concentration leads to more rigid wage dynamics. Another potential explanation for this result could be regulation. Natural monopolies—such as utilities—are, by definition, highly concentrated, but they are also highly regulated.

This adds another rigidity in the unemployment-inflation mechanism. That said, given that we only have one definitive sector that may attribute to this, we suspect this regulatory channel plays a relatively minor role, with monopsony as the primary driver.

Recall that for the high-unemployment environment, a decrease in the unemployment rate had no statistically significant impact for the first 12 months (Figure 4). One might expect that this is simply due to combating factors within these industries. Our results, however, do not reflect such dynamics, finding a lack of high-degree statistical significance throughout most of the corresponding sectoral analysis (Figure 7). This is especially surprising with regards to labor input, as wage dynamics are generally the baseline channel considered in the unemployment-inflation relationship. The only statistical significance we find in this high-unemployment environment comes from sectoral output, labor input, and capital intensity. The statistical significance of these results is only at the point of impact and is very small in magnitude. Also, it should be noted that the signs are reversed when moving from one unemployment environment to another, which we believe picks up the odd, negative, though non-significant effect seen at impact in Figure 4. From this, we conclude that our analysis does not provide much information for this type of shock in this type of unemployment environment.[15]

5 Robustness Checks

In this section, we present a number of robustness checks. These include exploring the model's sensitivity to changes in the threshold unemployment rate, changes in our control variables, and others. For space considerations, we only present robustness checks to the core PCE inflation results.

[15] This is also reflected in the low R-squared values for the models in this high-unemployment environment. Those of the other scenarios are substantially larger.

5.1 Testing the Threshold

The first robustness check we consider regards the threshold value of the unemployment rate used. In the baseline model, we set the threshold to 5.1% for two key reasons. First, it represents both the median and mean of unemployment for the 2007-2009 crisis period, which makes it a natural point of analysis. Second, unemployment rates below this threshold level are generally considered to indicate high degrees of labor-market tightness, which is the overall purpose of this analysis.

Narrowing our definition of a tight labor market does not change the baseline results substantially. A 4.9% threshold keeps approximately 35% of our observations (101 of 289 total observations) in the low-unemployment state, and the IRFs under this assumption can be found in Figure 8. Here, we find that the unemployment-inflation relationship to manifest with a one-month lag, and be slightly less persistent. Overall, however, the general dynamics are intact.

Increasing the threshold to 5.4%, which is the median value for the entire 1994:01–2018:10 period, also yields similar dynamics. The IRFs for the baseline model under the new threshold can be seen in Figure 9 and are nearly identical to the baseline analysis. The only difference here is a reduction from an 8–12 month persistence to five months of persistence in the low-unemployment environment. Together, these suggest that our initial assumption for the unemployment rate threshold is robust.

5.2 Alternative Measure of Inflation Expectations

The next robustness check we need to consider is how we define inflation expectations. In the baseline model, the five-year ahead expected inflation rate (UM Survey of Consumers) was considered due to the derivation of our chosen inflation measure and due to the persistent correlation of the 12-month reading with changes in oil prices, for which we already control. For completeness, Figure 10 presents the IRFs when we control for 12-month ahead inflation expectations. Just as with the previous robustness checks, the general dynamics are very similar to the baseline results (Figure 4), though they are not as persistent. Thus, our choice of expected inflation control is generally robust to other options as well.

5.3 Subsample Analysis

In an attempt to address the possibility of a time-varying relationship (e.g. Blanchard, 2016), we conduct the same test over the pre-2007-2009 crisis subsample. Recall that the original threshold unemployment rate was set to the median of this time period, ensuring that we have a balanced analysis in this robustness check. As can be seen in Figure 11, the general qualitative nature of our baseline results remain intact. There are some key differences to report, however. First the pre-crisis result in the low-unemployment environment is less persistent than in the full sample. Second, the contemporaneous impact of the shock in this subsample is larger than the full sample model at lower levels of confidence. The same cannot be said at higher confidence levels, but it is still worth noting. While not definitive, these results suggest that there has been a "flattening" of the Phillips curve in some respects, but there also seems to be a tradeoff. That is, while the impact may not be as vigorous, there has been an increased persistence in the effect. This could stem from any number of sources, such as increased monopsony power or even just added rigidities in the mechanism. We leave further exploration of this to future research, but this simplified analysis—as well as the results of our sectoral analysis—leaves the door open for time-variation.

We do not present a post-crisis analysis due to a sample size issue. While the post-crisis period contains more than enough observations for the total number of coefficients being estimated, maintaining a consistent 5.1% threshold unemployment rate leaves only 33 observations below this threshold to estimate 37 coefficients. The severity of the crisis pushed the post-crisis median unemployment rate to 6.55%. Simply put, there have not been enough observations since the crisis that fall under our definition of a "tight" labor market.

6 Conclusion

In this paper, we take a more generalized approach to exploring the Phillips-curve, unemployment-inflation relationship. Using local projections to estimate the inflation response to negative unemployment rate shocks, we find that the key difference across states of the economy is the timing of the response. While inflation has a positive response in both environments, the response in low-unemployment environments is immediate, the same positive response only manifests after a twelve-month lag. This state-dependent timing reveals that the pass-through from unemployment to inflation still exists despite not passing the eye test in recent years.

We then analyze industry-level data to show that the low-unemployment environment result is driven by elevated levels of labor and intermediate inputs, but diminished by higher levels of market concentration and sectoral output. While interesting in their own right, the combination of these results suggests that monopsony power may be a significant contributing factor. Recent trends in these underlying factors also leave the door open to the "flattening" narrative, providing an explanation to the diminishing reliability of the Phillips curve. That is, further decreases in labor income share, drift towards a service- and information-oriented economy, and increased market concentration will continue to erode the connection between unemployment and inflation. Therefore, it may behoove policy-makers to look to the changes in the factors that define it going forward rather than to the aggregate result more generally.

While we find that our model is quite robust, there are additional avenues for future research. For instance, our identification of the unanticipated unemployment rate shocks does not distinguish between supply- and demand-side innovations. Thus, our analysis is more from a general equilibrium point of view than a Phillips-curve-specific one. The goal here would be a way to maintain the information contained in our unanticipated shocks, but filter them in such a way that eliminates the supply-side innovations. Additionally, more work needs to be done on the sectoral composition and its effect on aggregate relationships like these. Lastly, we mentioned that an increase in sectoral output diminishes the unemployment-inflation pass through, but we do not specifically explore why that is. So while this study reveals new insights into this storied theory, there is still much work to be done.

References

Alpanda, S., Granziera, E. and Zubairy, S. (2019) State dependence of monetary policy across business, credit and interest rate cycles, Research Discussion Papers 16/2019, Bank of Finland.

Auerbach, A. J. and Gorodnichenko, Y. (2013) Output Spillovers from Fiscal Policy, *American Economic Review*, **103**, 141–146.

Bachmeier, L. J. and Cha, I. (2011) Why Don't Oil Shocks Cause Inflation? Evidence from Disaggregate Inflation Data, *Journal of Money, Credit and Banking*, **43**, 1165–1183.

Ball, L. M. and Mazumder, S. (2019) The Nonpuzzling Behavior of Median Inflation, NBER Working Papers 25512, National Bureau of Economic Research, Inc.

Blanchard, O. (2016) The Phillips Curve: Back to the '60s?, *American Economic Review*, **106**, 31–34.

Carlino, G. and DeFina, R. (1998) The Differential Regional Effects Of Monetary Policy, *The Review of Economics and Statistics*, **80**, 572–587.

Carlino, G. A. and DeFina, R. H. (1999) Do states respond differently to changes in monetary policy?, *Business Review*, pp. 17–27.

Clark, T. E. and McCracken, M. W. (2006) The Predictive Content of the Output Gap for Inflation: Resolving In-Sample and Out-of-Sample Evidence, *Journal of Money, Credit and Banking*, **38**, 1127–1148.

Coibion, O. and Gorodnichenko, Y. (2015) Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation, *American Economic Journal: Macroeconomics*, **7**, 197–232.

Daly, M. C. and Hobijn, B. (2014) Downward Nominal Wage Rigidities Bend the Phillips Curve, *Journal of Money, Credit and Banking*, **46**, 51–93.

Debelle, G. and Laxton, D. (1997) Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom, and the United States, *IMF Staff Papers*, **44**, 249–282.

Doser, A., Nunes, R., Rao, N. and Sheremirov, V. (2017) Inflation expectations and nonlinearities in the Phillips curve, Working Papers 17-11, Federal Reserve Bank of Boston.

Fry, R. and Pagan, A. (2011) Sign Restrictions in Structural Vector Autoregressions: A Critical Review, *Journal of Economic Literature*, **49**, 938–960.

Fujita, S. (2011) Dynamics of worker flows and vacancies: evidence from the sign restriction approach, *Journal of Applied Econometrics*, **26**, 89–121.

Gagnon, J. E. and Collins, C. G. (2019) Low Inflation Bends the Phillips Curve, Working Paper Series WP19-6, Peterson Institute for International Economics.

Guisinger, A. Y., Hernandez-Murillo, R., Owyang, M. T. and Sinclair, T. M. (2018) A state-level analysis of Okun's law, *Regional Science and Urban Economics*, **68**, 239–248.

Herrera, A. M. and Rangaraju, S. K. (2019) The quantitative effects of tax foresight: Not all states are equal, *Journal of Economic Dynamics and Control*, **107**, 103726.

Herrera, A. M. and Rangaraju, S. K. (2020) The effect of oil supply shocks on US economic activity: What have we learned?, *Journal of Applied Econometrics*, **35**, 141–159.

International Monetary Fund (2013) The dog that didn't bark: Has inflation been muzzled or was it just sleeping?, in *World Economic Outlook: Hopes, Realities, and Risks*, International Monetary Fund, Washington D.C., chap. 3, pp. 79–96.

Jordà, O. (2005) Estimation and Inference of Impulse Responses by Local Projections, *American Economic Review*, **95**, 161–182.

Kilian, L. and Lütkepohl, H. (2018) *Structural Vector Autoregressive Analysis*, no. 9781107196575 in Cambridge Books, Cambridge University Press.

Kilian, L. and Murphy, D. P. (2014) The role of inventories and speculative trading in the global market for crude oil, *Journal of Applied Econometrics*, **29**, 454–478.

Lewis, J. B. and Linzer, D. A. (2005) Estimating Regression Models in Which the Dependent Variable Is Based on Estimates, *Political Analysis*, **13**, 345–364.

Phillips, A. W. (1958) The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861–1957, *Economica*, **25**, 283–299.

Plagborg-Møller, M. and Wolf, C. K. (2019) Local projections and vars estimate the same impulse responses, *Unpublished paper: Department of Economics, Princeton University*, p. 1.

Ramey, V. A. and Zubairy, S. (2018) Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data, *Journal of Political Economy*, **126**, 850–901.

Sahm, C. (2019) Direct stimulus payments to individuals, in *Recession Ready: Fiscal Policies to Stabilize the American Economy* (Eds.) H. Boushey, R. Nunn and J. Shambaugh, The Hamilton Project and the Washington Center on Equitable Growth, Washington, DC, pp. 67–92.

Sims, C. A. (1980) Macroeconomics and Reality, *Econometrica*, **48**, 1–48.

Uhlig, H. (2005) What are the effects of monetary policy on output? results from an agnostic identification procedure, *Journal of Monetary Economics*, **52**, 381 – 419.

Webber, D. A. (2018) Employment Adjustment over the Business Cycle: The Impact of Competition in the Labor Market, IZA Discussion Papers 11887, Institute for the Study of Labor (IZA).

Figure 1: Identified Unemployment Shocks

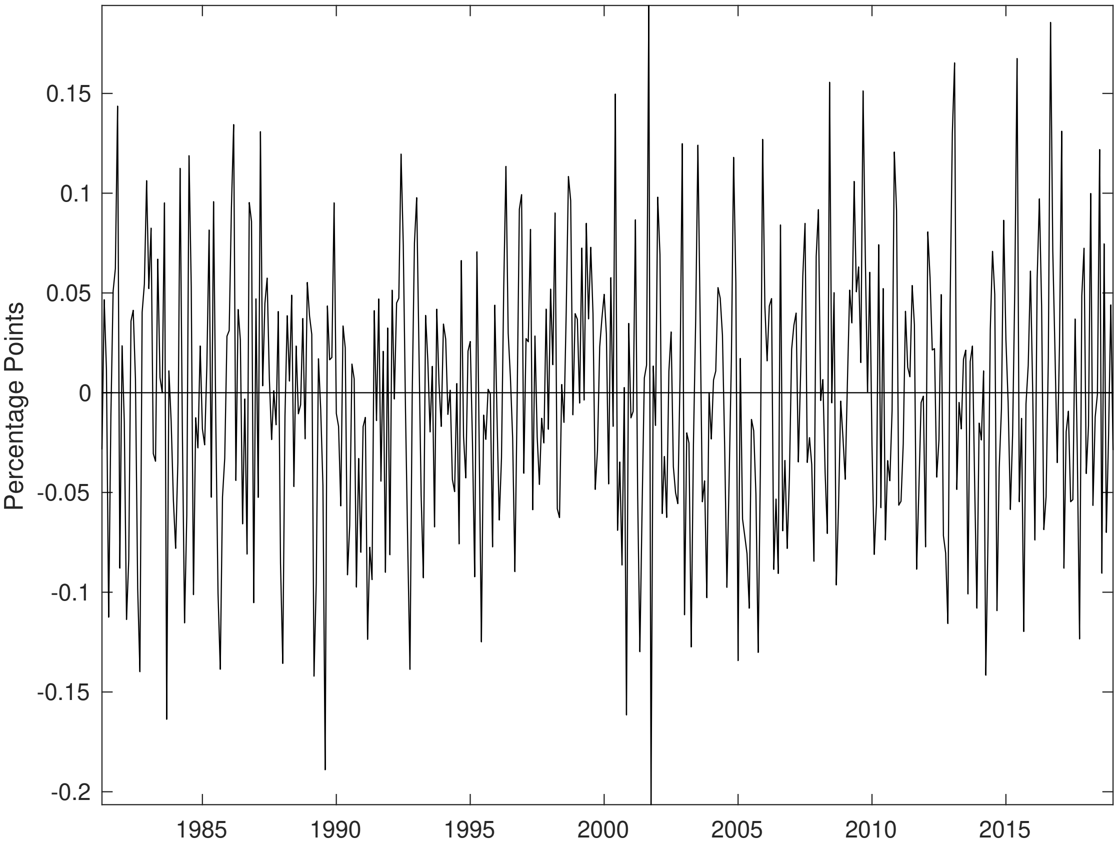

Figure 2: Identified Shock Impact on Unemployment Rate

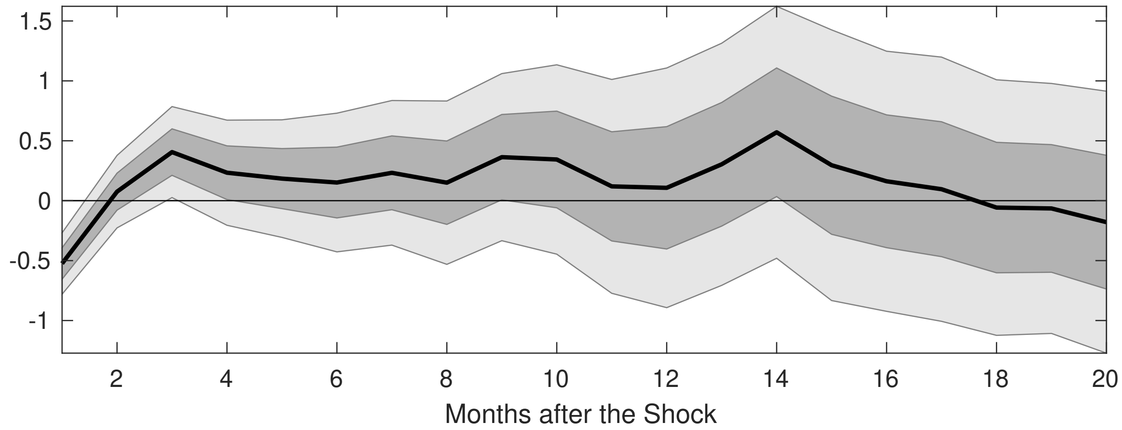

Note: The impulse response function represents the unemployment rate's reaction to its own, negative shock. The point estimate is denoted by the solid, black line while the shaded regions represent the 68% and 95% confidence intervals.

Figure 3: Core PCE Inflation: Linear Model Specification

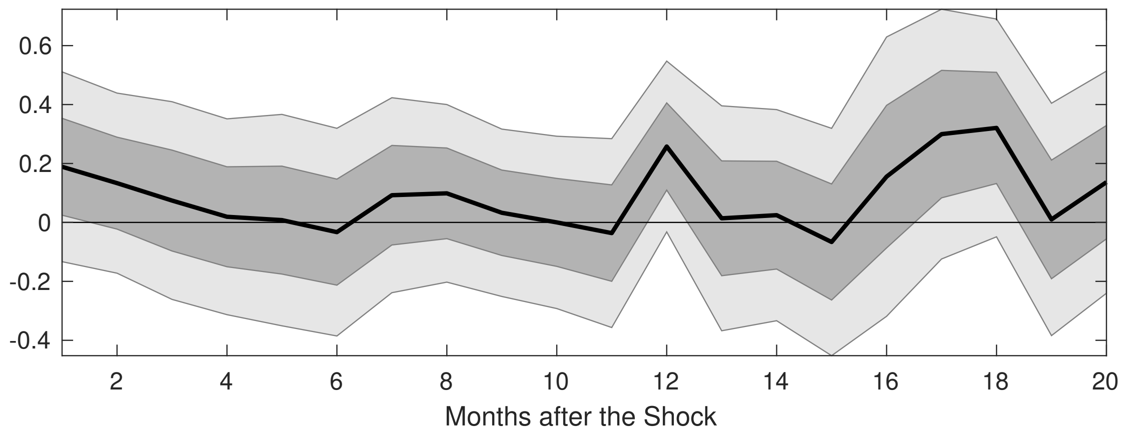

Note: The point estimate for this impulse response function represents the core PCE inflation rate and is denoted by the solid, black line. Shaded regions represent the 68% and 95% confidence intervals. The shock considered is a *negative* innovation to the unemployment rate.

Figure 4: Core PCE Inflation: Nonlinear Model Specification

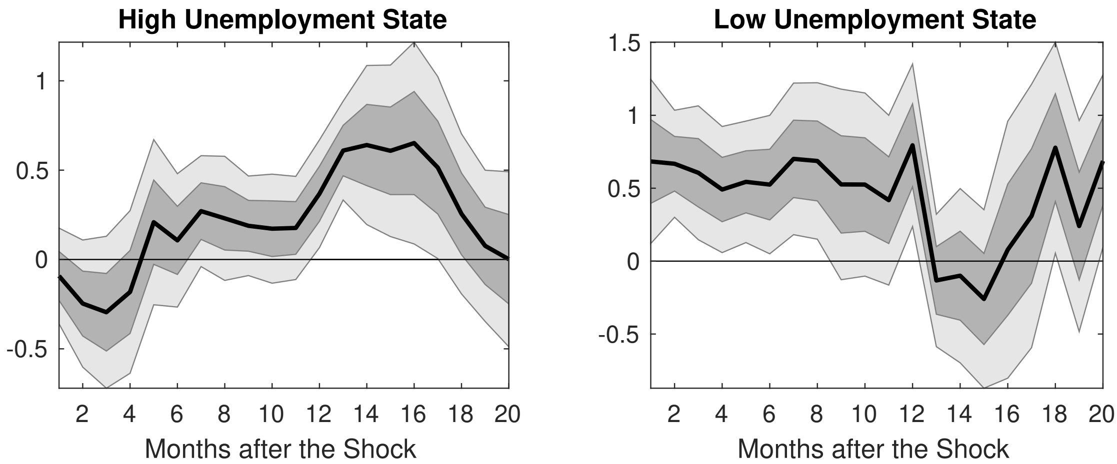

Note: The impulse response functions represent the core PCE inflation rate and are denoted by the solid, black lines. Shaded regions represent the 68% and 95% confidence intervals. Both shocks represent unanticipated *decreases* in the unemployment rate.

Figure 5: Trimmed-Mean PCE Inflation: Nonlinear Model Specification

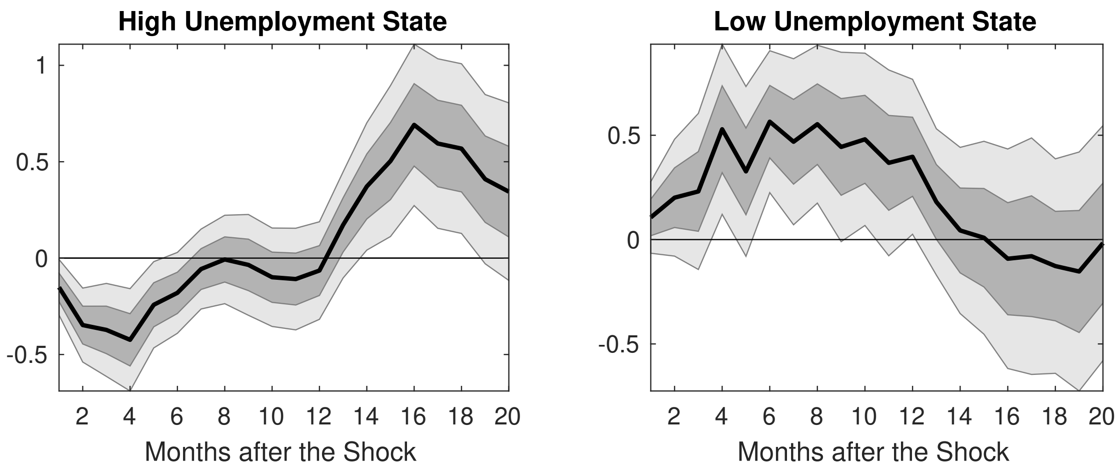

Note: The impulse response functions represent the trimmed-mean PCE inflation rate published by the Federal Reserve Bank of Dallas and are denoted by the solid, black lines. Shaded regions represent the 68% and 95% confidence intervals. Both shocks represent unanticipated *decreases* in the unemployment rate.

Figure 6: Sectoral Analysis of Cumulative PCE Response: Low-Unemployment Environment

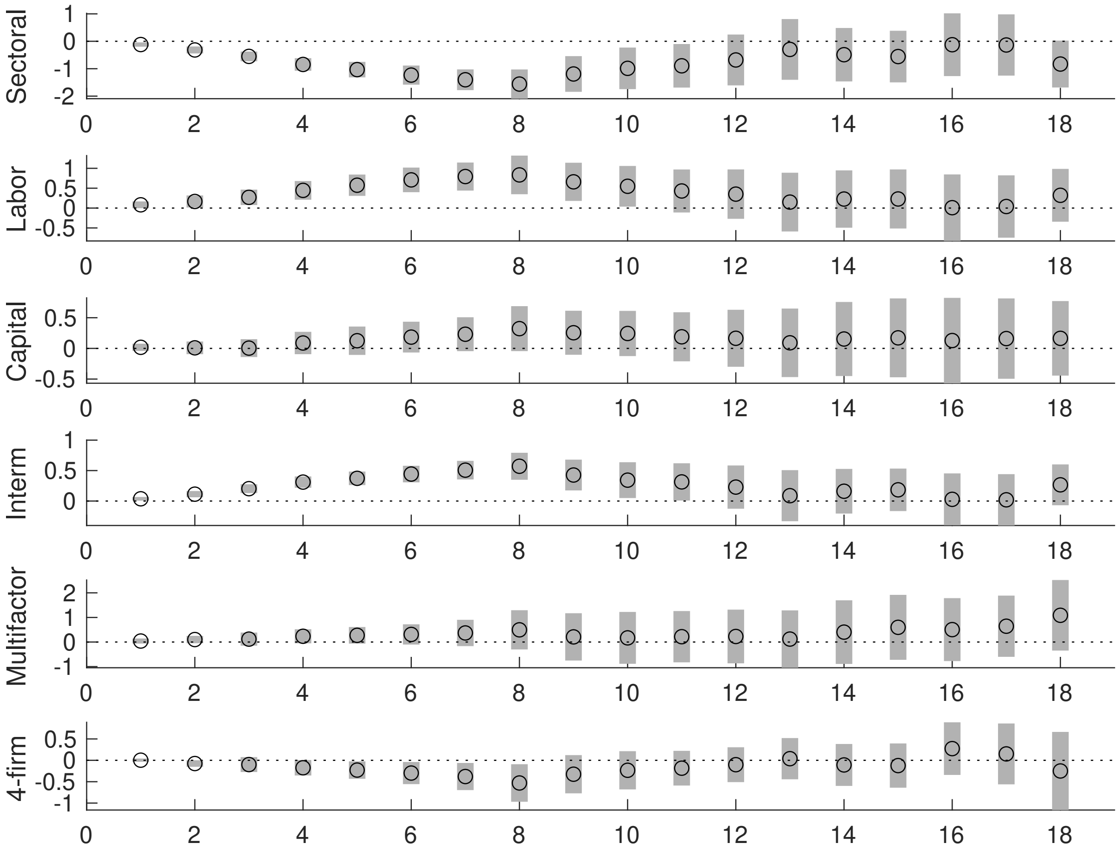

Note: The results above represent period-by-period FGLS regressions of the cumulative impulse responses against the six sector-level variables and a constant (not shown). The labels on the y-axes are short for sectoral output, labor input, capital intensity, intermediate input, multifactor productivity, and the 4-firm concentration ratio, respectively. The circles designate the point estimates, while the gray bars correspond to the 95% confidence intervals.

Figure 7: Sectoral Analysis of Cumulative PCE Response: High-Unemployment Environment

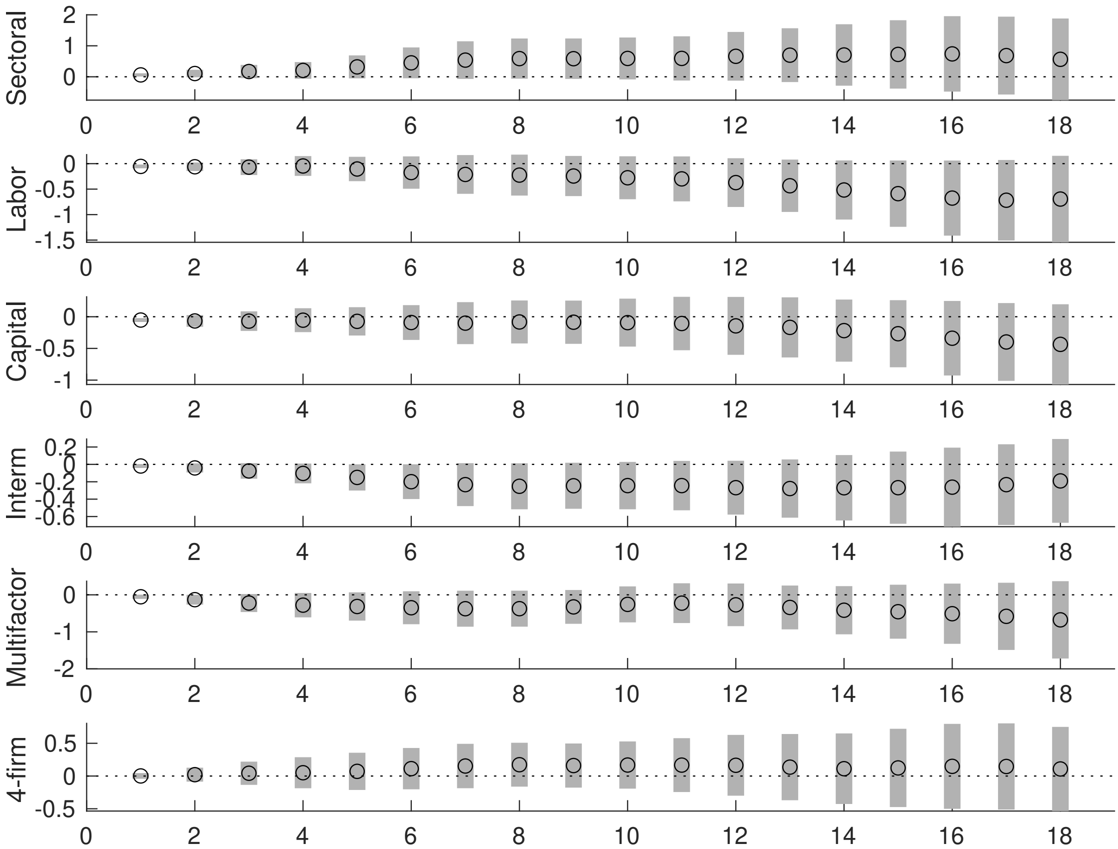

Note: The results above represent period-by-period FGLS regressions of the cumulative impulse responses against the six sector-level variables and a constant (not shown). The labels on the y-axes are short for sectoral output, labor input, capital intensity, intermediate input, multifactor productivity, and the 4-firm concentration ratio, respectively. The circles designate the point estimates, while the gray bars correspond to the 95% confidence intervals.

Figure 8: Core PCE Inflation: Threshold = 4.9%

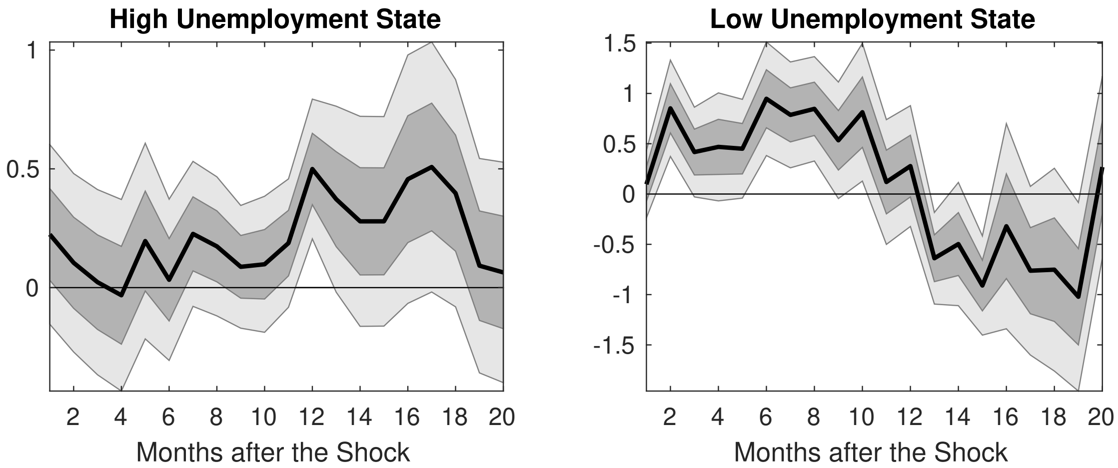

Note: The impulse response functions represent the core PCE inflation rate and are denoted by the solid, black lines. The threshold is set at 4.9% unemployment rate versus the 5.1% baseline case. Shaded regions represent the 68% and 95% confidence intervals. Both shocks represent *decreases* in the unemployment rate.

Figure 9: Core PCE Inflation: Threshold = 5.4%

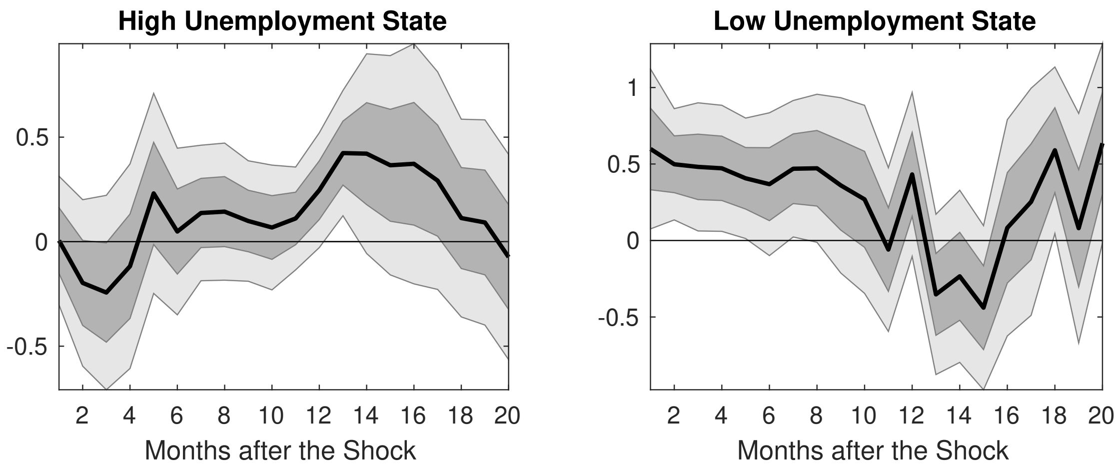

Note: The impulse response functions represent the core PCE inflation rate and are denoted by the solid, black lines. The threshold is set at 5.4% unemployment rate versus the 5.1% baseline case. Shaded regions represent the 68% and 95% confidence intervals. Both shocks represent *decreases* in the unemployment rate.

Figure 10: Core PCE Inflation: 12-month Inflation Expectations

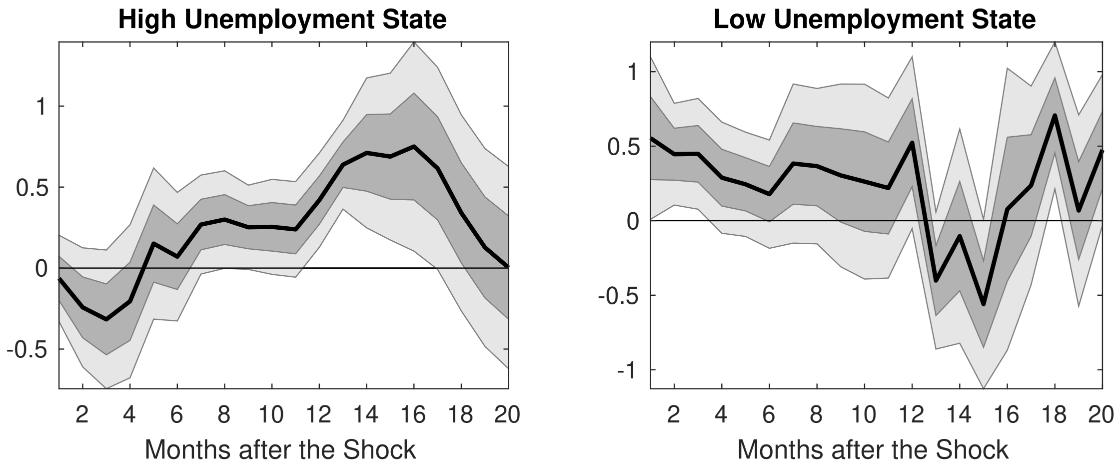

Note: The impulse response functions represent the core PCE inflation rate and are denoted by the solid, black lines. 12-month inflation expectations (UM Survey of Consumers) used versus 5-year baseline case. Shaded regions represent the 68% and 95% confidence intervals. Both shocks represent *decreases* in the unemployment rate.

Figure 11: Core PCE Inflation: Pre-crisis subsample

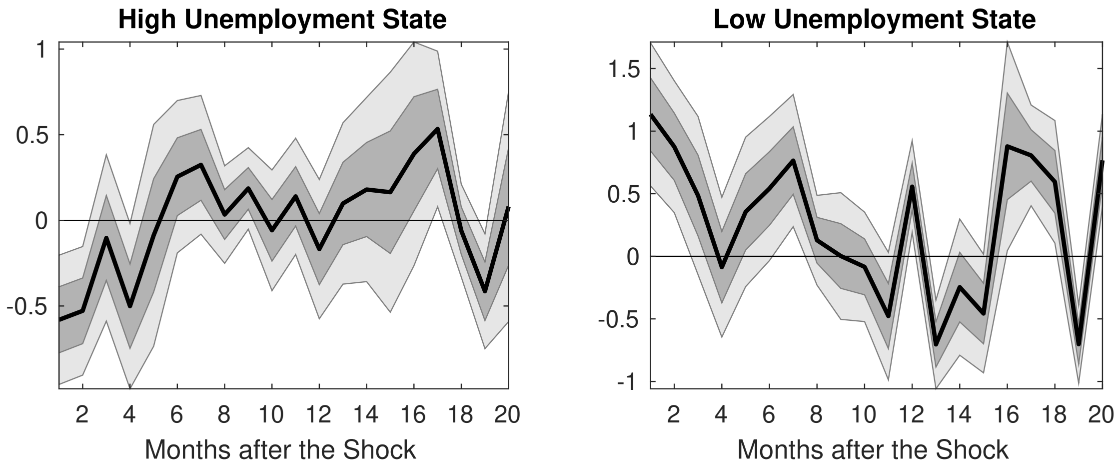

Note: The impulse response functions represent the core PCE inflation rate and are denoted by the solid, black lines. The subsample studied here spans the 1994:01–2007:12 period. Shaded regions represent the 68% and 95% confidence intervals. Both shocks represent *decreases* in the unemployment rate.

Table 1: NAICS/PCE Categorical Pairing

NAICS Number	NAICS Category	PCE Line	PCE Category
22	Utilities	161	Household Utilities
334	Computer and Electronic Products	46	Information Processing Equipment
335	Electrical Equipment, Appliances, and Components	27	Household Appliances
3361	Motor Vehicles, Bodies and Trailers, and Parts	4	Motor Vehicles and Parts
337	Furniture and Related Products	22	Furniture and Furnishings
311FT	Food and Beverage and Tobacco Products	71	Food and Beverages Purchased for Off-Premises Consumption
315AL	Apparel and Leather and Allied Products	102	Clothing and Footwear
322	Paper Products	131	Household Paper Products
324	Petroleum and Coal Products	111	Gasoline and Other Energy Goods
481	Air Transportation	203	Air Transportation
482	Rail Transportation	197	Railway Transportation
483	Water Transportation	204	Water Transportation
485	Transit and Ground Passenger Transportation	198	Road Transportation
512	Motion Picture and Sound Recording Industries	210	Motion Picture Theaters
513	Broadcasting and Telecommunications	278	Telecommunication Services
523	Securities, Commodity Contracts, and Investments	245	Financial Service Charges, Fee, and Commissions
524	Insurance Carriers and Related Activities	266	Insurance
HS	Housing	151	Housing
532RL	Rental and Leasing Services and Lessors of Intangible Assets	152	Rental of Tenant-Occupied Nonfarm Housing
5411	Legal Services	295	Legal services
5412OP	Miscellaneous Professional, Scientific, and Technical Services	294	Professional and Other Services
61	Educational Services	286	Educational Services
621	Ambulatory Health Care Services	169	Outpatient Services
622	Hospitals	179	Hospitals
623	Nursing and Residential Care Facilities	183	Nursing Homes
624	Social Assistance	311	Social Services and Religious Activities
711AS	Performing Arts, Spectator Sports, Museums, and Related Activities	209	Admissions to Specified Spectator Amusements
713	Amusements, Gambling, and Recreation Industries	205	Recreation Services
721	Accommodation	245	Accommodations
722	Food Services and Drinking Places	231	Food Services