Bank of England

Revisiting the effects of long-term unemployment on inflation: the role of non-linearities

Staff Working Paper No. 1,018
March 2023

Vania Esady, Bradley Speigner and Boromeus Wanengkirtyo

Bank of England

Staff Working Paper No. 1,018

Revisiting the effects of long-term unemployment on inflation: the role of non-linearities

Vania Esady,[1] Bradley Speigner[2] and Boromeus Wanengkirtyo[3]

Abstract

We demonstrate that it is necessary to control for state dependence in the Phillips curve in order to be able to appropriately identify separate slopes for short and long-term unemployment rates. Whereas several existing studies have typically concluded that long-term unemployment is largely immaterial for price pressures, our evidence suggests that the effect of long-term unemployment on inflation is highly state dependent. In particular, reductions in long-term unemployment are found to be significantly inflationary when aggregate unemployment is low, displaying a larger and more immediate peak effect on inflation than short-term unemployment. The explanation for our finding is a direct consequence of allowing for non-linearity in the Phillips curve together with short and long-term unemployment gaps that enter the specification separately. Variation in long-term unemployment typically arises following large recessionary shocks and the Phillips curve also tends to be flatter in deep recessions. It therefore follows that the comovement between long-term unemployment and inflation will be understated in linear regressions. In order to address this, we adopt a flexible methodology that combines non-linearity and heterogeneity in the unemployment duration distribution, enabling us to control for this confounding effect of state dependence on the identification of separate Phillips-curve slopes for short and long-term unemployment. Our results would caution against underweighting long-term unemployment in the inflation-relevant measure of economic slack, especially when unemployment is low.

Key words: Phillips curve, non-linearity, long-term unemployment, cross-sectional identification.

JEL classification: C22, E24, E32, J6.

(1) Bank of England. Email: vania.esady@bankofengland.co.uk
(2) Bank of England. Email: bradley.speigner@bankofengland.co.uk
(3) Bank of England. Email: boromeus.wanengkirtyo@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful for useful comments from participants at Bank of England seminars, IAAE 2022 and New Directions for Inflation Forecasting.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1. Introduction

The headline unemployment rate is the most widely used indicator of economic slack. A large empirical literature on Phillips curve estimation has explored whether more general definitions of labour utilisation are more informative than this simple measure (Hornstein et al., 2014; Faberman et al., 2020). In this paper, we investigate whether the duration distribution of unemployment contains useful information for modelling inflation dynamics. Put differently, do short- and long-term unemployment play separate roles in the Phillips curve? A consensus seems to have formed in much of the preceding literature that suggests that the long-term unemployment rate (LTU) plays no role at all in determining inflation. Only short-term unemployment (STU) appears to matter for wage competition and therefore inflationary pressure. Information about the duration distribution of unemployment matters for inflation prediction only insofar as it seems to rule out a role for LTU. This question is important in the context of large increases – and recoveries – in the long-term share of total unemployment that has occurred during the downturns following the 2008 recession and the more recent pandemic.

In this paper, we propose instead that convexity in the Phillips curve – that it is flatter in recessions – can explain these empirical results, building on the observation of Speigner (2014). When we account for this state-dependence, LTU becomes a key variable for explaining the behaviour of inflation, contributing even more acutely to inflationary pressure than STU during tight labour market conditions. The intuition behind this result is simple. LTU is typically a deep recession phenomenon. Moreover, there is a phase difference between LTU and STU over the business cycle in the sense that LTU lags STU, almost by definition. By the time LTU actually rises, the economy is much more likely to have already transitioned to a relatively flat region of the Phillips curve. At that point, however, labour market slack *in general* exerts a smaller drag on inflation at the margin, making it appear as though there is a disconnect between LTU and inflation when in fact there may not be.

For this reason, it is imperative to control for state-dependence in the Phillips curve when assessing whether the partial correlation of LTU with inflation is significant. Crucially, we also demonstrate that this intuition holds in reverse when the economy is in the steep part of the Phillips curve. At that stage, reductions in the LTU rate can be highly inflationary when the aggregate unemployment rate is low. We draw from this that there is in fact significant informational content in the duration distribution of unemployment that can be used to help predict inflation conditional on the right empirical strategy being adopted.

The empirical strategy we adopt incorporates two major methodological improvements from the recent literature. Firstly, whereas Speigner (2014) imposes convexity through a restrictive parametric modification to an otherwise standard time-varying Phillips curve, we apply a much more flexible estimation method based on local projections. Such techniques can easily be adapted to handle rich unemployment dynamics in the presence of state-dependencies in the Phillips curve.[1] Secondly, we leverage cross-sectional information to further aid the identification of the Phillips curve, making use of state-level rather than aggregate data (McLeay and Tenreyro, 2020; Hazell et al., 2020). Our econometric setup therefore allows us to jointly estimate the curvature of the Phillips curve as well as identify separate slopes for short- and long-term unemployment rates.

Our main empirical findings are two-fold. First, for a linear baseline specification we confirm the results of Kiley (2015) who finds that STU and LTU have the same Phillips curve slope when estimated on regional US data. In addition to this baseline result, our estimation methodology also enables us to model the inflation rate's *dynamic* behaviour in a rich way. This allows us to further demonstrate that the effect of LTU on inflation is sharper but less persistent compared to STU when a linear model is used. This additional dynamic dimension is important and illustrates one a key benefit of adopting a local projections approach as it enables us to highlight key differences in the time profile of inflation following changes in the different unemployment measures.

The picture changes however when a non-linear model is adopted. As discussed above, the non-linear model specification reveals a key interaction between state-dependence and unemployment duration that is hidden in the linear estimates. Once we allow the slope of the Phillips curve to change with the level of unemployment, it remains the case that STU and LTU have similar (small) effects on inflation during periods of economic slack when unemployment is high. But we also find that the impact of LTU on inflation during low unemployment periods is much larger than the effect of STU. Another way to express this is that the state dependence of the Phillips curve is more acute with respect to LTU than STU.

This paper is related to the literature which argues that the long-term unemployed are less relevant for wage determination than the short-term unemployed and therefore should be discounted or omitted entirely from the measurement of economic slack (Gordon, 2013; Krueger et al., 2014; Watson, 2014; Linder et al., 2014). This idea has its roots in insider-outsider theories of the labour market (Lindbeck and Snower, 2001). The central hypothesis is that either through human capital loss or an outright preference by firms to preferably rank newly unemployed individuals, the long-term unemployed become less substitutable for employed workers and therefore do not factor into the wage bargaining process as much.

Allowing for differential weights on short- and long-term unemployment in the Phillips curve, Llaudes (2005) finds that the weight on the long-term unemployed was often statistically smaller for a number of countries. Kiley (2015) raises an important econometric issue, arguing that the high degree of collinearity between STU and LTU in the aggregate data makes such an inference difficult. Using the extra variation contained in regional data to help alleviate the identification problem, he finds that STU and LTU actually exert equal downward pressure on inflation. Our reliance on

[1]Ever since the original study by Phillips (1958) it has commonly been accepted that the Phillips curve is possibly non-linear. Modern micro-foundations of convexity have included downward nominal wage rigidity (see Daly and Hobijn (2014)) and strategic complementarities (see Lindé and Trabandt (2019)).

cross-sectional identification of the Phillips curve is therefore based on the idea put forth by Kiley (2015). One difference however is that instead of metropolitan statistical areas (MSAs) we make use of more granular state-level data.

One of the aims of the current study is to bridge this body of work with the extensive empirical literature on non-linear Phillips curve estimation. Several of the more recent contributions to the latter also utilise cross-sectional information to help trace out non-linearities (Kumar and Orrenius (2016); Hooper et al. (2020); Byrne and Zekaite (2020)). From a technical standpoint, there are of course many ways to specify non-linear models. The methods that are commonly applied in the empirical literature are most often based on spline functions or Markov-switching models. We take a slightly different methodological route in this paper. Our econometric methodology instead uses local projections to characterize the state-dependence of the Phillips curve, in much the same vein as Ramey and Zubairy (2018) applied the technique to estimate the state dependence of fiscal multipliers.

The rest of the paper is organised as follows. The related literature is discussed in the next section. Then, in order to bring out the intuition more clearly, we construct a stylised model in Section 2 to illustrate in a simplified setting that our findings are a natural outcome when our two key ingredients are mixed together: (i) a phase difference, or lag, between LTU and STU cycles and (ii) a convex Phillips curve. This stylised model makes it very clear that a given degree of convexity in the aggregate Phillips curve will have a disproportionate effect on the degree of state-dependence of the LTU slope coefficient. The broader empirical methodology is outlined Section 3. Main empirical results are reported in Section 5. The final section concludes.

2. Stylised model

We now set out a stylised model to illustrate how state-dependence in the slope of the Phillips curve (PC) and duration dependence in unemployment, when combined, lead to testable predictions that can be assessed with the available data. Take a standard backward-looking Phillips curve augmented with a piecewise-linear slope:

$$(\pi_t - \pi^*) = \rho_\pi(\pi_{t-1} - \pi^*) - \kappa_t(u_t - u^N) \tag{1}$$

$$\kappa_t = \begin{cases} \kappa_1 & \text{if } u_t \leq u^N \text{ (``Expansion'')} \\ \kappa_2 & \text{if } u_t > u^N \text{ (``Recession'')} \end{cases} \tag{2}$$

where $\kappa_1 \gg \kappa_2$ and u^N denotes the natural rate of unemployment.

A simple mechanism is then specified to capture the empirical observation that the LTU share varies systematically with the level of aggregate unemployment due to the phase difference between

LTU and STU.[2] By definition, a negative shock that precipitates a rise in unemployment initially transmits predominantly to STU. LTU follows with a phase difference, rising only subsequently as the recession deepens. A standard logistic function for the LTU share is a reasonably good approximation for these mechanics:[3]

$$\frac{\text{ltu}_t}{u_t} = \frac{b_1}{1 + \exp(-b_2(u_t - b_3))} \tag{3}$$

where the sum of STU and LTU is aggregate unemployment ($\text{stu}_t + \text{ltu}_t = u_t$).[4] Under these assumptions, the STU and LTU Phillips curve slopes are constrained to both be the same. We then simulate a path for unemployment around the natural rate using an AR(1) process:

$$(u_t - u^N) = \rho_u(u_{t-1} - u^N) + \varepsilon_t^u, \quad \text{where } \varepsilon_t^u \sim N(0, \sigma^2_{\varepsilon^u}) \tag{4}$$

The model is purely for the purposes of illustration, and is calibrated as follows. The persistence of π (ρ_π) and u (ρ_u) is set to 0.5 and 0.9, respectively. The natural rate of unemployment and trend inflation are 5% and 2%. The standard deviation of unemployment shocks is 0.1. The PC slope is calibrated to $\kappa_1 = 0.2$ and $\kappa_2 = 0.3\kappa_1$. This 70% flattening is roughly what Lindé and Trabandt (2019) find. Recall that κ_t is the *aggregate* PC slope.

This simple model can then be used to construct a simulation as shown in Figure 1. The scatter points are blue when the economy is in an expansion ($u_t \leq u^N$) and red in a recessionary state ($u_t > u^N$). By design, the simulation produces a convex PC slope when plotting inflation against aggregate unemployment u_t. However, the interaction of non-linearity with the state-dependence of the LTU share results in an interesting implication for the degree of convexity of the Phillips curve with respect to STU and LTU. Recall that, by construction, the marginal impact of STU and LTU on inflation is the same, equal to κ_t for a given state of the economy.

However, it can clearly be seen in the third panel of Figure 1 that the Phillips curve with respect to LTU is much more convex than the corresponding curve for STU. This accentuated non-linearity will make it more likely for LTU to appear statistically insignificant during recessions because the slope is very flat, while at the same time amplifying the effect that LTU has on inflation during booms. In contrast, the opposite effect occurs for the Phillips curve plotted in STU space, for which the non-linearity is attenuated. The aggregate curve lies somewhere in between.

[2] Figure 2c in Section 3 illustrates the relationship between aggregate unemployment and the LTU share.

[3] In theory, any monotonically increasing function is capable of capturing the story. The results are also robust to using a simpler linear function. Tenreyro and Thwaites (2016) also uses a similar logistic function to denote economic states.

[4] The parameters b_1, b_2, b_3 are estimated on the state-level unemployment rates (shown in the next section) with non-linear least squares. The full estimation results can be found in the Appendix (Table 3).

Figure 1: *Simulated Phillips curves*

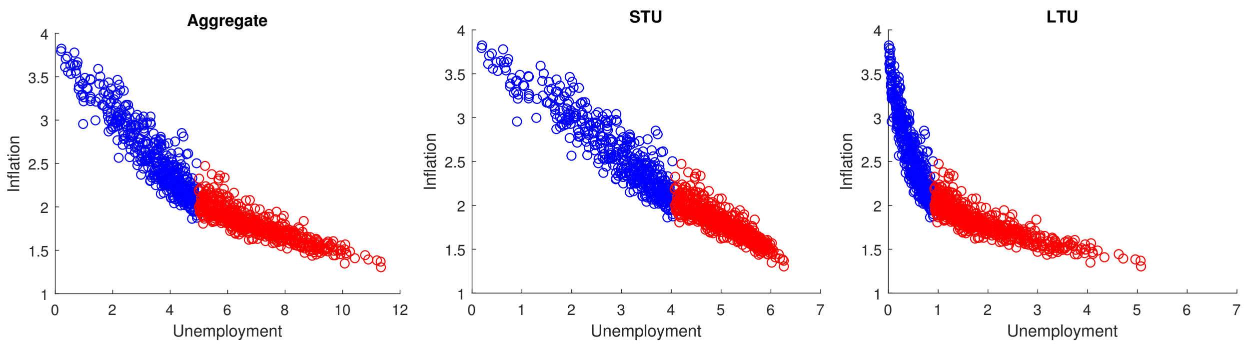

Table 1 summarises the empirical predictions of our framework. If the true Phillips curve is convex (the 'convexity' hypothesis), then LTU will have a larger average effect on inflation during expansions and a roughly similar (or smaller) effect during booms. In contrast, if LTU does not matter for inflation as the 'detached worker' hypothesis would suggests, then LTU should have a smaller (possibly zero) effect on inflation than STU in all business cycle conditions. Finally, it is also possible that unemployment duration is not important for Phillips curve estimation, in which case LTU and STU have 'equal effects'.

What this simple exercise clearly illustrates is that statistical inferences about the degree of state-dependence in the PC and the relative importance of LTU should be carried out in a setup which considers these issues jointly. We turn to this in the next section.

Table 1: *Theoretical predictions of the PC slope from stylised model*

Hypotheses	Expansion	Recession	*Non-Linearity*
Convexity	$LTU > STU$	$LTU <\approx STU^*$	Stronger in LTU
Detached worker	$LTU < STU$	$LTU < STU$	None
Equal effects	$LTU \approx STU$	$LTU \approx STU$	N/A

* The theory predicts a smaller difference in slack periods, but whether it is roughly the same depends on the calibration.

3. Methodology

3.1. A Cross-Sectional Approach to Identifying the Phillips Curve

McLeay and Tenreyro (2020) highlight a fundamental endogeneity problem in Phillips curve identification. The macroeconomic influence of an inflation targeting central bank that trades off inflation and output gap stabilisation results in a reduced form PC that is substantially flatter than the structural PC or perhaps even has the 'wrong' sign. A regional approach to PC identification exploits the fact that business cycles across states are not perfectly aligned and that the central

Figure 2: *Unemployment rates at the state level in the US*

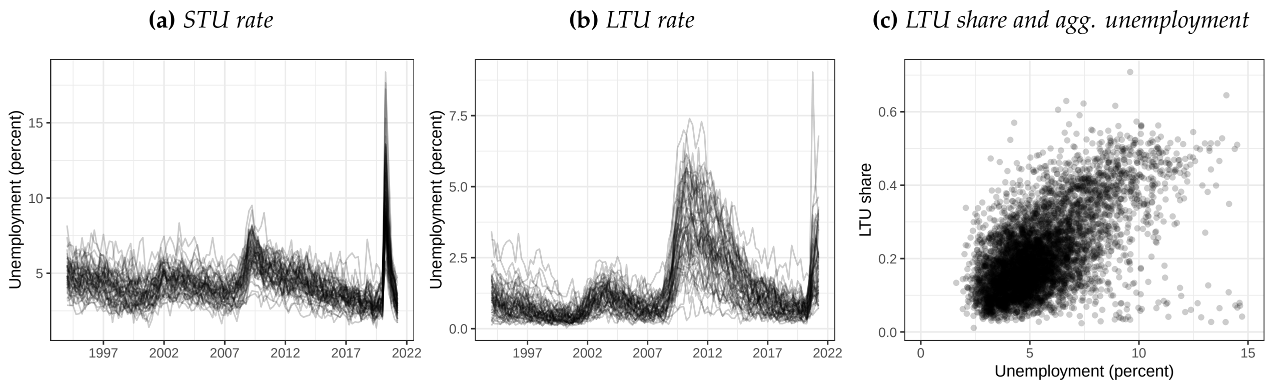

Note: The left and middle charts plots *state-level* short-term and long-term unemployment rates, respectively, calculated from the CPS micro-data. The right chart plots the state-level LTU share (of aggregate unemployment) against aggregate unemployment, illustrating the positive association between the two variables discussed in Section 2. We follow the Bureau of Labor Statistics' definition of long-term unemployment as in excess of 27 weeks. The sample is from 1996Q1 to 2021Q2. The summary statistics can be found in the Appendix (Table 2).

bank does not target specific states' inflation or unemployment rates. Figures 2a and 2b depict the variation in short- and long-term unemployment rates at the state level for the US.

In recognition of this, and to exploit the significant amount of variation in the cross-section, we estimate Phillips curves using state-level price and unemployment data, also adding time fixed effects. Hazell et al. (2020) point out that time fixed effects adjust for the presence of inflation expectations in a forward-looking Phillips curve. Furthermore, regional (state-level) fixed effects are included to control for potentially different trend inflation rates as well as heterogeneous natural rates of unemployment, to the extent that they are time-invariant. The latter implies that for the regression analysis we can simply use levels of unemployment rather than having to enter unemployment in gap form. Results for local projections without time fixed effects are reported in the Appendix. As predicted by McLeay and Tenreyro (2020), in this case the response of inflation to a rise of unemployment often ends up with a 'wrong' sign.

3.2. Data

Kumar and Orrenius (2016) and Hooper et al. (2020) also investigate non-linearities using disaggregated US state-level data, approximating state-level prices with CPI-U data (an index of prices for urban consumers). Our preferred data source is the state-level CPI data available from Hazell et al. (2020).

We construct state-level unemployment rates (split by duration) using the micro-data from the Current Population Survey (CPS) (Flood et al., 2020). We follow the definition of the Bureau of Labor Statistics (BLS) of long-term unemployment being that of a continuous unemployment spell of more than 27 weeks. Data are seasonally-adjusted with X-13 ARIMA.[5]

The state-level unemployment data is then merged with the state-level inflation rates constructed by Hazell et al. (2020) from the BLS CPI micro-data.[6] Following Hazell et al. (2020), in addition to the headline inflation rate we also report results for non-tradable goods inflation. Combining the state-level unemployment and inflation data, our sample period is 1994Q1 to 2017Q4.

Figure 2c plots the state-level LTU share against aggregate unemployment. A general upward sloping relationship can be clearly seen. When economic conditions deteriorate, a larger fraction of the unemployment pool is long-term unemployed. This observation forms the empirical basis for the assumption used in our illustrative model in Section 2.

3.3. Local Projections Specification

The estimation method we use is based on the local projections approach of Jordà (2005):[7]

$$\pi_{r,t+h} = \alpha_h + \psi_h(L)\pi_{r,t-1} + \beta_h \text{stu}_{r,t} + \delta_h \text{ltu}_{r,t} + \gamma_t + \gamma_r + \varepsilon_{r,t+h} \quad \text{for} \quad h = 0, 1, 2, \ldots, H \tag{5}$$

where π refers to inflation (aggregate or non-tradable, as discussed further below), $\psi_{h(L)}$ is a lag polynomial of order 3 (so the maximum lag is 4), γ_t are the time fixed effects that control for variation in average unemployment rates across U.S. states and γ_r are the region (state) fixed effects. ε_{t+h} is the residual term. The key parameters are β_h and δ_h, which capture the Phillips curve slope estimates h quarters ahead corresponding to STU and LTU, respectively. It is in this sense that the PC model is dynamic; a slope parameter is estimated for each horizon h, separately for STU and LTU. This specification pools the Phillips curves together across regions. Using Bayesian panel breakpoint methods, Smith et al. (2021) find no cross-sectional heterogeneity in regional (MSA) PC slopes, providing some justification for a pooled panel approach.

State-dependence is introduced into the above setup using the approach of Ramey and Zubairy (2018) as follows:

$$\begin{aligned}\pi_{r,t+h} = F_{r,t-1} & \left[\alpha_h^{\text{HU}} + \psi_h^{\text{HU}}(L)\pi_{r,t-1} + \beta_h^{\text{HU}} \text{stu}_{r,t} + \delta_h^{\text{HU}} \text{ltu}_{r,t}\right] \\ & + (1 - F_{r,t-1}) \left[\alpha_h^{\text{LU}} + \psi_h^{\text{LU}}(L)\pi_{r,t-1} + + \beta_h^{\text{LU}} \text{stu}_{r,t} + \delta_h^{\text{LU}} \text{ltu}_{r,t}\right] \\ & + \gamma_t + \gamma_r + \varepsilon_{r,t+h} \qquad \text{for} \quad h = 0, 1, 2, \ldots, H\end{aligned} \tag{6}$$

The indicator variable F equals 1 when the economy is in high unemployment (HU) and 0 when it is in low unemployment (LU) periods. HU periods are defined as when *total state-level* unemployment

[5] In some very small states, the sample of the CPS is limited resulting in occasional missing values, even though the CPS samples more than 100,000 individuals each month. For these missing state-month observations, missing long-term unemployment values are linearly interpolated.

[6] Note that the dataset by Hazell et al. (2020) only has 34 states. However, this would still more than 50% greater than the number of MSAs (typically used by other studies).

[7] Angeli et al. (2021) also takes a similar local projections approach for the UK but do not focus on the interaction between non-linearities and the duration distribution of unemployment.

is higher than the state's median unemployment rate over the sample period. The interaction with the indicator variable allows *all* coefficients to vary according to the regime of the economy.[8] We chose this particular scheme to define the two regimes in the absence of state-specific recession indicators such as NBER recession dates, which are only available nationally.

4. Results

Results from the linear specification are reported in Figures 3a and 3b. For the first 12 quarters, the impact on inflation of a one percentage point fall in STU and LTU is broadly similar, yielding a peak inflationary impulse of approximately 0.1 percentage points (pp). This is qualitatively similar to the results reported by Kiley (2015) who also found similar effects on inflation from STU and LTU. We therefore confirm Kiley's MSA-level results in a linear setup using our state-level data but we note that we are able to achieve somewhat greater estimation precision. This is most likely the result of using state-level data and the greater cross-sectional variation it provides.[9]

However, the effect of LTU on inflation is observed to be significantly less persistent than for STU. It is no longer statistically significant from zero at the end of the impulse response horizon. This stands in notable contrast to the inflation response to STU which tends to build over time rather than become weaker. This difference in the dynamic path of inflation in response to STU and LTU illustrates an added insight of adopting a local projections approach. Therefore, if anything, the linear model would therefore appear to (weakly) favour the detached-worker hypothesis that LTU workers are less active market participants and consequently have a smaller overall effect on inflation (Table 1).

In contrast, the results from the non-linear model in Figures 3c and 3d reveal a key state-dependence that is hidden in the linear estimates. As predicted by the illustrative model in Section 2, the peak impact of LTU on inflation in low unemployment periods is much higher than the effect of STU (0.4pp compared to 0.2pp). This difference almost entirely disappears during high unemployment states, with the effects of a 1pp rise in STU and LTU in high unemployment periods being broadly similar at less than 0.1pp. Although even then the impact of LTU is, if anything, slightly stronger than STU.

Another qualitative aspect of the results is that inflationary pressure emerges much more quickly for a reduction in LTU than it does for STU. It is of course a common feature of Phillips curve specifications to include lagged dynamics. What is novel about our finding is that the time profile of inflation is clear a function of unemployment duration. Differences therefore arise not just with respect to the peak impact of STU and LTU on inflation but also the timing of such impacts.

It is possible to speculate why this might be the case. Typically, people who are short-term

[8] The results are also robust using a smooth-transition logistic function instead, available upon request.

[9] On the sample period closest to ours (post-1998), Kiley (2015) finds the coefficients on STU and LTU to be -0.17 and -0.14, respectively, with standard errors of 0.15 for each.

Figure 3: *Response of aggregate inflation to unemployment*

(a) *Inflation response to STU (linear model)*

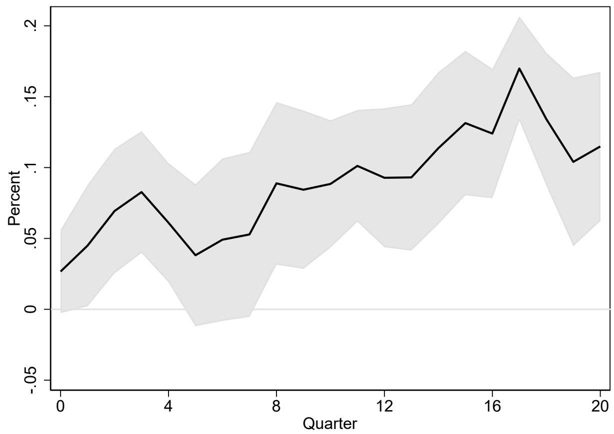

(b) *Inflation response to LTU (linear model)*

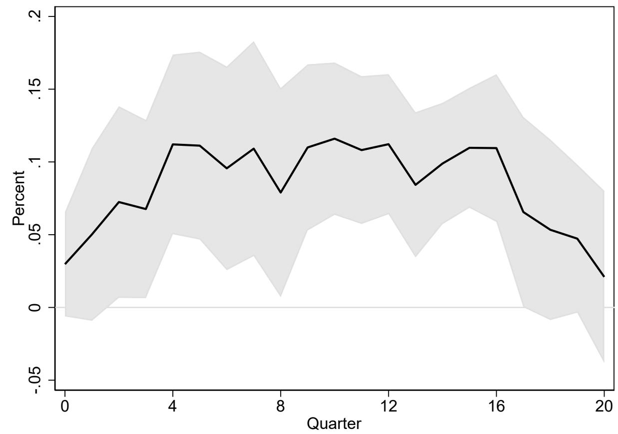

(c) *Inflation response to STU (non-linear model)*

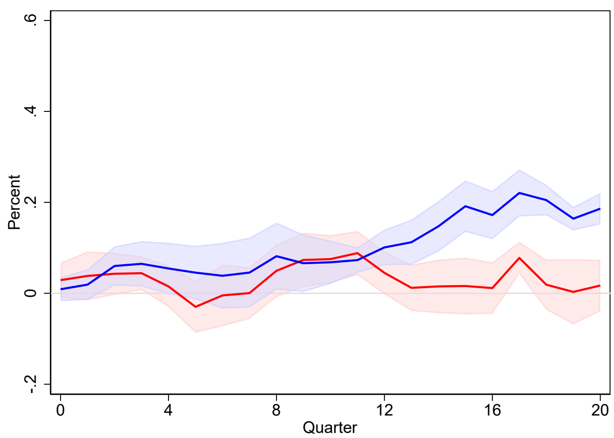

(d) *Inflation response to LTU (non-linear model)*

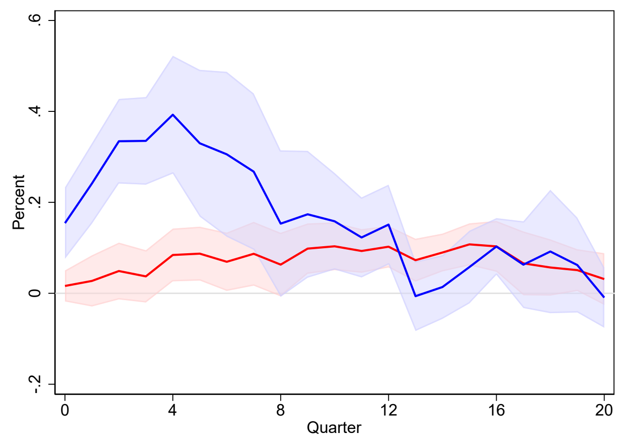

Note: This shows the impulse responses to a 1pp fall in the short-/long-term unemployment rate. The shaded areas show the 68% Driscoll-Kraay confidence interval. The first row is the linear model and the second row is the state-dependent model (where blue is low unemployment, and red is high unemployment periods).

unemployed have higher transition rates into employment and are therefore likely to become employed sooner than the long-term unemployed during a period in which general employment levels are rising. Figures 3c and 3d indicate that the economy can begin to absorb excess slack through a reduction in STU with relatively little consequence for inflation until, with a delay, the hiring process becomes more congested. It is possible that reductions in STU can be made without necessitating a burst of wage pressure to attract more hires, which would explain why initial reductions in STU take more time to filter through to inflation. In contrast, a reduction in LTU precipitates a more immediate rise in inflation, signifying that the labour market cannot absorb a change in this form of slack easily.

Figure 4: *Non-tradable Inflation*

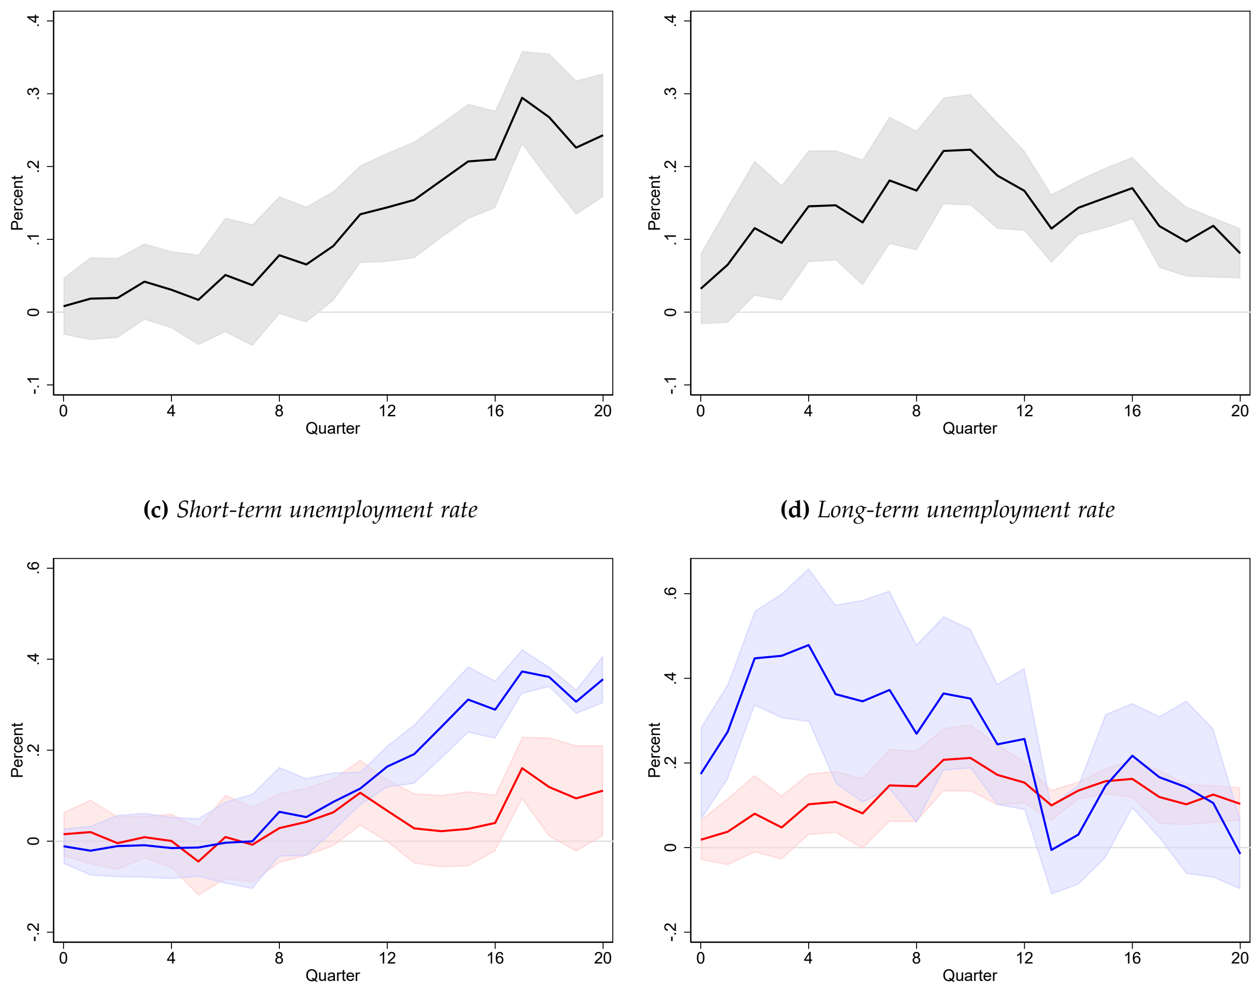

(a) *Short-term unemployment rate*

(b) *Long-term unemployment rate*

(c) *Short-term unemployment rate*

(d) *Long-term unemployment rate*

Note: Impulse responses to a 1pp fall in the short-/long-term unemployment rate. The shaded areas show the 68% Driscoll-Kraay confidence interval. The first row is the linear model and the second row is the state-dependent model (where blue is low unemployment, and red is high unemployment periods).

The results for non-tradable inflation are reported in Figure 4 and suggest a similar story.[10] The qualitative features of the exercise remain more or less unchanged. Quantitatively, in the linear specification both STU and LTU exert greater pressure on inflation, as can be seen by the larger impact in Figure 4 compared with Figure 3. The delayed response for STU becomes particularly more accentuated.

Turning to the non-linear specification, the peak impact in low unemployment periods is higher

[10]With tradables inflation, we do not find statistically significant results for both the linear and non-linear specification. We also do not find non-linearities as we do for aggregate and non-tradable inflation. This is in-line with Hazell et al. (2020) who noted that for prices set at the national level – as is more likely for tradables – the slope of the regional PC will be zero.

for both LTU and STU at approximately 0.5pp and 0.35pp, respectively. Therefore the difference between the peak effect of LTU and STU is now less stark. Inflationary pressure still builds relatively more gradually following a reduction in STU, but the pace of momentum quickens more notably after 8 quarters in Figure 4c compared to figure Figure 3c. The key result nevertheless remains that the PC is more non-linear with respect to LTU than with STU.

For both the aggregate inflation and non-tradables specifications, the difference between high and low unemployment states is statistically significant for both STU and LTU, as Figure 5 in the appendix shows.

In summary, the Phillips curve slopes with respect to STU and LTU diverge most significantly during expansionary phases when the labour market is tight – which fits the convexity hypothesis the closest (Table 1).[11] We find that, in stark contrast to the prevailing literature, LTU has a highly significant impact on inflation with the transmission coming through with less of a delay than is the case for STU. We do not find much of a meaningful difference between LTU and STU during high unemployment periods when the effect of both measures of unemployment on inflation is estimated to be relatively weak. Non-linear PC patterns are most acute with respect to LTU rather than STU. The ratio of the peak inflationary impact of LTU in the two regimes is around four (0.4pp vs. 0.1pp), whereas it is only a factor of two with STU (0.2pp vs. 0.1pp).

4.1. Further Discussion and Policy Implications

What are the possible macroeconomic channels that could explain the result that a reduction in LTU has significant inflationary effects? We offer two ideas: (i) a supply-side effect through the labour market and (ii) a demand-side effect through household demand. At this stage, both of these hypotheses are speculative but relate closely to previously established literatures.

First, to the extent that the LTU are relatively more detached from the labour market, the search and matching process is likely to be more difficult and less efficient than for the STU pool, effectively making hiring more costly for firms. If such recruitment difficulties are met with higher wage offers from firms, then this would add to inflationary pressure. Krueger et al. (2014) find that the job finding rate of the LTU pool is lower than the STU pool, consistent with labour market detachment rising with unemployment duration. This could be due to a combination of human capital depreciation and discouragement by the LTU pool, alongside potential discrimination on the part of employers. However, Abraham (2014) highlights various measurement issues with the CPS that may mean that job finding rates of LTU searchers are actually not very different to the STU searchers. Furthermore, while Krueger et al. (2014) find that the LTU job finding rate is less cyclical than STU, Abraham (2014) argues that this could be because the most attractive candidates in the LTU pool are the first ones hired as labour market conditions improve in expansions. This

[11] However, it is noteworthy that various labour market frictions could also be the underpinnings of the convexity. We discuss it further in the the next subsection 4.1.

compositional effect could cause a worsening in the average employability of the remaining LTU pool, masking the effects of the tighter labour market on LTU exit rates.

Either way, whether due to detachment or compositional changes, if marginal hiring costs are greater for the LTU pool compared to the STU pool, then that could provide one rationalisation for the main results in this paper. The labour market is effectively bumping up against a capacity constraint once the STU pool starts to become depleted. This insight contains useful information for policymakers as it implies that the LTU rate is an important gauge to monitor when assessing inflation risks, particularly during low unemployment periods.

There is also potentially a different demand-side channel that could provide an underpinning for why long-term unemployment is an important barometer of inflationary pressure. Becoming unemployed is typically a negative income shock and household consumption responds accordingly. There is a well-established literature, theoretically and empirically, showing that households' consumption responds by more the more persistent the income shock is (see Jappelli and Pistaferri (2010) for a review). The majority of unemployment in the US over recent decades has been short-term (the average STU share in our sample is 77.5%) and it is easier for households to smooth their consumption over a relatively short period of unemployment. However, the longer individuals are unemployed, the more likely they are to perceive the income shock to be more persistent (potentially amplified by the reduced job-finding rates of becoming longer-term unemployed, if that is the case). Additionally, the longer-term unemployed are also more likely to drain their liquid assets, making consumption more sensitive to income shocks. Jappelli and Pistaferri (2010) highlight the importance of liquidity constraints. These effects combined could lead to a significant fall in household consumption in the face of widespread long-term unemployment, amplifying the shortfall in aggregate demand. In other words, in general equilibrium, fluctuations in LTU could be associated with larger fluctuations in aggregate demand. Firms' prices would then be likely to respond by more to a change in LTU compared to STU. This behavior is consistent with Nekarda and Ramey (2020), who find that mark-ups are procyclical in response to demand shocks.

5. Conclusion

Long-term unemployment typically rises only when there is substantial slack in the labour market and aggregate unemployment is relatively high. These are the same macroeconomic conditions that also tend to be associated with a relatively flat Phillips curve. Most of the previous literature that has attempted to investigate whether short- and long-term unemployment should have a different weight in the Phillips curve has tended to concentrate on linear inflation-unemployment specifications, typically finding that long-term unemployment has relatively small effects on the behaviour of inflation. We have attempted to shed some new light on this issue by using state-level consumer price data coupled with a non-linear local projections estimation methodology that enables a richer characterisation of the dynamic elements of the Phillips curve compared with prior work in this area. We find evidence that short- and long-term unemployment have roughly similar Phillips curve slopes in recessions when the aggregate Phillips curve is relatively flat. But we also find evidence of a *larger* Phillips curve slope with respect to long-term unemployment during expansions when unemployment is low. Our results would therefore suggest that policymakers and econometricians stand to gain from including long-term unemployment as part of the inflation-relevant measure of economic slack.

References

Abraham, K. (2014): "Discussion of Krueger, Cramer, and Cho, 'Are the Long-Term Unemployed on the Margins of the Labor Market?'," *Brookings Papers on Economic Activity*, 45, 281–87.

Angeli, M., M. McLeay, S. Tenreyro, and B. Wanengkirtyo (2021): "The Phillips Curve in the UK," *mimeo*.

Byrne, D. and Z. Zekaite (2020): "Non-linearity in the wage Phillips curve: Euro area analysis," *Economics Letters*, 186, 108521.

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

Faberman, R. J., A. I. Mueller, A. Şahin, and G. Topa (2020): "The Shadow Margins of Labor Market Slack," Tech. rep., National Bureau of Economic Research.

Flood, S., M. King, R. Rodgers, S. Ruggles, and J. R. Warren (2020): "Integrated Public Use Microdata Series, Current Population Survey: Version 8.0 [dataset]," *Minneapolis, MN: IPUMS, 2020*, 34.

Gordon, R. J. (2013): "The Phillips curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery," Tech. rep., National Bureau of Economic Research.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2020): "The Slope of the Phillips Curve: Evidence from U.S. States," NBER Working Papers 28005, National Bureau of Economic Research, Inc.

Hooper, P., F. S. Mishkin, and A. Sufi (2020): "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is it Just Hibernating?" *Research in Economics*, 74, 26–62.

Hornstein, A., M. Kudlyak, and F. Lange (2014): "Measuring Resource Utilization in the Labor Market," *Economic Quarterly*, 1–21.

Jappelli, T. and L. Pistaferri (2010): "The Consumption Response to Income Changes," *Annual Review of Economics*, 2, 479–506.

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

Kiley, M. T. (2015): "An Evaluation of the Inflationary Pressure Associated with Short- and Long-term Unemployment," *Economics Letters*, 137, 5–9.

Krueger, A. B., J. Cramer, and D. Cho (2014): "Are the Long-term Unemployed on the Margins of the Labor Market?" *Brookings papers on economic activity*, 2014, 229–299.

Kumar, A. and P. M. Orrenius (2016): "A closer look at the Phillips curve using state-level data," *Journal of Macroeconomics*, 47, 84–102, what Monetary Policy Can and Cannot Do.

Lindé, J. and M. Trabandt (2019): "Resolving the Missing Deflation Puzzle," CEPR Discussion Papers 13690, C.E.P.R. Discussion Papers.

Lindbeck, A. and D. J. Snower (2001): "Insiders versus Outsiders," *Journal of Economic Perspectives*, 15, 165–188.

Linder, M. H., R. Peach, and R. W. Rich (2014): "The Long and Short of it: the Impact of Unemployment Duration on Compensation Growth," Tech. rep., Federal Reserve Bank of New York.

Llaudes, R. (2005): "The Phillips Curve and Long-term Unemployment," *ECB working paper*.

McLeay, M. and S. Tenreyro (2020): "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 34, 199–255.

Nekarda, C. J. and V. A. Ramey (2020): "The Cyclical Behavior of the Price-Cost Markup," *Journal of Money, Credit and Banking*, 52, 319–353.

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.

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

Smith, S. C., A. Timmermann, and J. H. Wright (2021): "The Phillips Curve: Heterogeneity Across Space and Time," *mimeo*.

Speigner, B. (2014): "Long-term Unemployment and Convexity in the Phillips Curve," Bank of England working papers 519, Bank of England.

Tenreyro, S. and G. Thwaites (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8, 43–74.

Watson, M. W. (2014): "Inflation Persistence, the NAIRU, and the Great Recession," *American Economic Review*, 104, 31–36.

A. Appendix

Table 2: *Summary Statistics of Unemployment Data*

					Quantile			
Variable	n	Mean	S.D.	Min	.25	Med	.75	Max
Short-term Unemployment	3264	4.41	1.05	1.41	3.68	4.32	5.03	9.51
Long-term Unemployment	3264	1.47	1.18	0.10	0.66	1.09	1.83	7.00
Total Unemployment	3264	5.88	1.91	1.63	4.55	5.51	6.88	14.09

Table 3: *Non-linear least squares regression of LTU share*

	b_1	b_2	b_3
	0.585***	0.308***	7.500***
	(0.071)	(0.051)	(1.035)
Observations		3,264	
R-squared		0.898	

Standard errors clustered around states in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 5: *t-tests for Coefficient Equivalence in High and Low Unemployment Periods*

(a) *Aggregate Inflation (STU)*

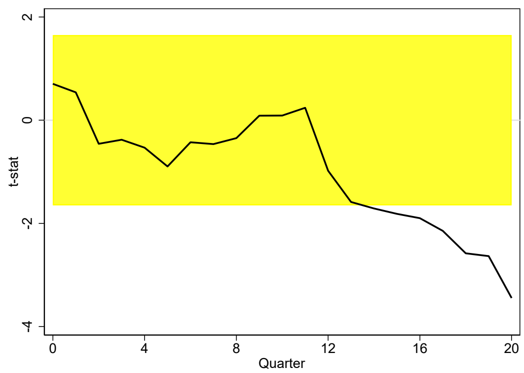

(b) *Aggregate Inflation (LTU)*

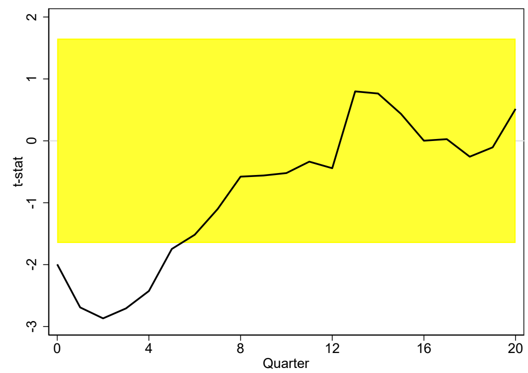

(c) *Non-tradable Inflation (STU)*

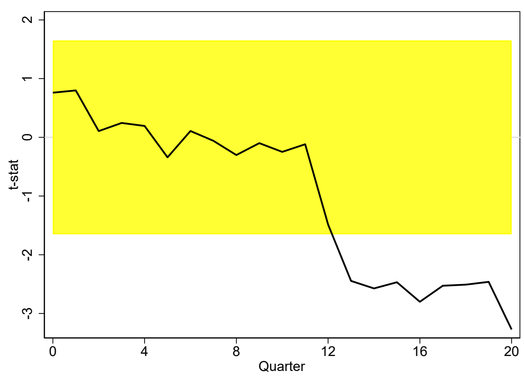

(d) *Non-tradable Inflation (LTU)*

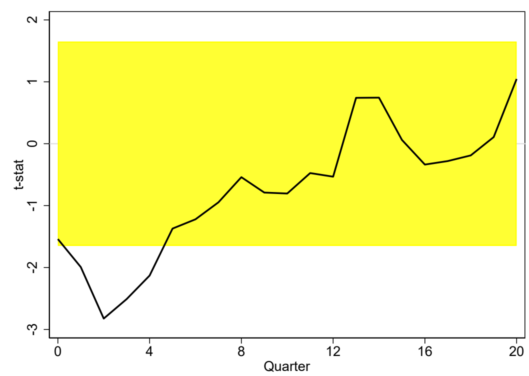

Note: This shows the t-statistics for the impulse responses reported in Figures 3 and 4, for the IRF for expansions minus the IRF on slack periods, for a given horizon h. The yellow shaded area indicates ± 1.65 (a 90% confidence interval).

Figure 6: *Model without time fixed effects*

(a) *Short-term unemployment rate*

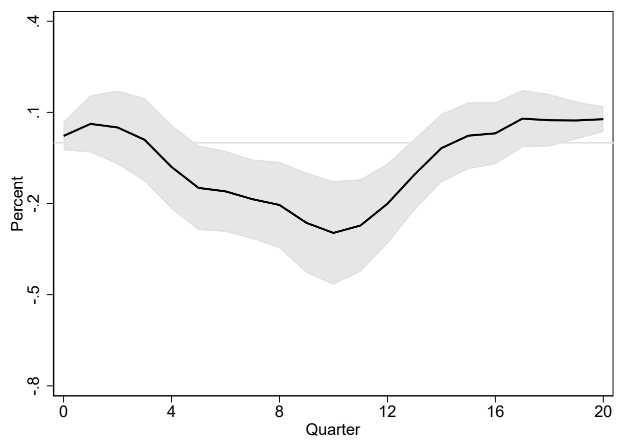

(b) *Long-term unemployment rate*

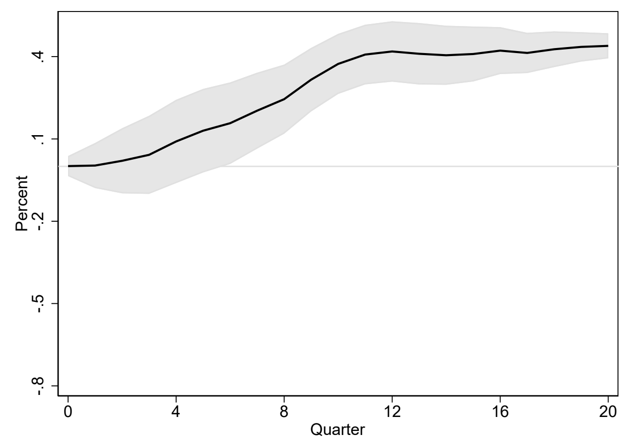

(c) *Short-term unemployment rate*

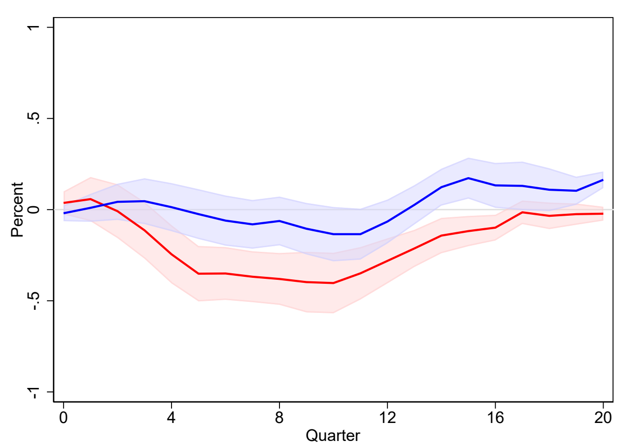

(d) *Long-term unemployment rate*

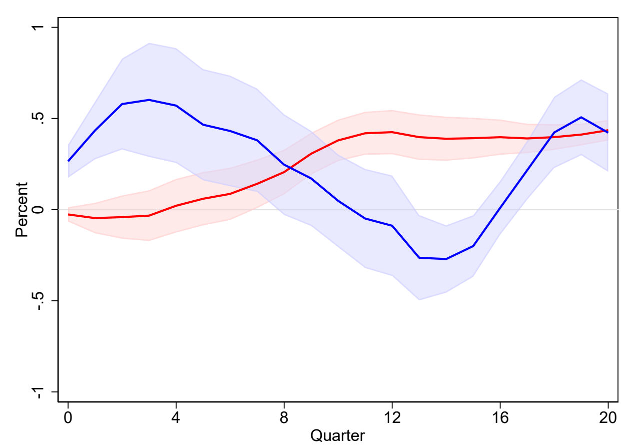

Note: This shows the impulse responses to a 1pp fall in the short-/long-term unemployment rate. The shaded areas show the 68% Driscoll-Kraay confidence interval. The first row is the linear model, and the second row is the state-dependant model (where blue is low unemployment, and red is high unemployment periods).