Kinky Europe: Evidence from the Regional Phillips Curve in the Euro Area*

Marius Faber[†] Gabriel Züllig[‡]

April 15, 2025

Abstract

We estimate the slope of the Phillips curve in the euro area, allowing for nonlinearities – or kinks – in the relationship between labor market slack and inflation. We exploit cross-country variation in labor market conditions in the period 2001-2024, absorbing aggregate shocks and endogenous monetary policy reactions with time fixed effects. We find that while the Phillips curve is usually quite flat, it becomes at least three times as steep if the labor market is sufficiently tight. This kink is more pronounced in the euro area than in the United States, potentially because of more rigid labor markets.

JEL classification: E30

Keywords: Phillips curve, inflation, nonlinearities

*We would like to thank Andreas Freitag, Jonathon Hazell, Christian Hepenstrick, Peter Karadi, Sarah Lein, Maxime Phillot, Raphael Schoenle and Michael Weber for helpful discussions. We also thank seminar participants at an SNB brown bag seminar for helpful comments and suggestions. Leonie Frei provided excellent research assistance. The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the authors. They do not necessarily reflect the views of the Swiss National Bank (SNB). The SNB takes no responsibility for any errors or omissions in, or for the correctness of, the information contained in this paper.

[†]Swiss National Bank, marius.faber@snb.ch

[‡]Swiss National Bank, gabriel.zuellig@snb.ch

1 Introduction

The Phillips curve, the standard model used by economists to analyze and forecast inflation, has been the subject of intense debate for over half a century. In his seminal study, Phillips (1958) documented that in the United Kingdom, wage growth tended to be higher when unemployment was low, and vice versa. This relationship between wages and unemployment was later extended to price inflation, given the link between wages and prices.

In recent decades, there has been an intense debate as to whether the Phillips curve had flattened, which has led to new approaches for identifying its slope. Inflation had appeared to have become less and less sensitive to changes in unemployment in the period between 1960 and 2020 in the United States, leading economists to debate about a potential "flattening" of the Phillips curve (Blanchard, 2016, Stock and Watson, 2020). From this literature, it became clear that in normal times, economic slack does impact inflation, however, only to a relatively limited extent (Hazell et al., 2022). As a byproduct of this debate, new estimation techniques, which address various identification issues when estimating the causal relationship between unemployment and inflation, have been introduced.[1] For example, it became clear that using cross-sectional variation from a monetary union is useful to rule out endogeneity bias from aggregate shocks, endogenous monetary policy responses or changes in long-run inflation expectations.

Despite a seemingly flat Phillips curve, inflation rose far above central bank targets during the post-Covid period. Since the inflation surge was accompanied by tight labor markets in most advanced economies, the focus shifted towards understanding if the Phillips curve is nonlinear, i.e., flat in normal times but steeper in times of a booming economy with tight labor markets.[2] Notably, Phillips (1958) already had a nonlinear relationship in mind, stating that "[the] relation between unemployment and the rate of change of wage rates is (...) likely to be highly nonlinear" (p.1). While euro area data offers clear suggestive evidence of a nonlinear Phillips curve, a split of the data by pre- and post-Covid highlights the difficulty of disentangling the role of supply shocks or changes in long-run inflation expectations from a nonlinear slope with aggregate data (see Figure 1).

In this paper, we therefore estimate the slope of the Phillips curve in the euro area using regional data, allowing for nonlinearities – or kinks – in the relationship between labor market slack and inflation. We first justify our empirical approach by laying out the main theoretical insights from Hazell et al. (2022) on how to infer the slope of the aggregate Phillips curve from regional estimates. We then augment this structural equation to allow for nonlinearities depending on the degree of labor market tightness. Next, we turn to the empirical analysis, where we exploit cross-country variation in labor market conditions in the period 2001-2024, absorbing aggregate shocks and endogenous monetary policy reactions with time fixed effects. To purge results

[1] See, for example, McLeay and Tenreyro (2019) or Fitzgerald et al. (2024). For a concise overview of recent developments, see Furlanetto and Lepetit (2024).

[2] Another closely related debate is about the role of supply and demand forces causing the inflation surge. See, for example, Guerrieri et al. (2022) and Bernanke and Blanchard (2023).

Figure 1: First look at the Phillips curve relationship in the euro area

from endogeneity resulting from contemporaneous regional supply shocks, we instrument the contemporaneous unemployment gap with the four-quarter lagged unemployment gap.

Results show that while the Phillips curve in the euro area is usually quite flat, it becomes substantially steeper if the labor market is sufficiently tight. We show this to be robust to various alternative specifications addressing measurement and identification concerns. For example, we vary the definition of "tight" labor markets and use various subsamples. Across all specifications, the Phillips curve in tight labor markets is at least three and up to fifteen times steeper when labor markets are tight.

The kink we document for the euro area is more pronounced than what has been found for the United States, potentially due to more rigid labor markets in the euro area. In a heterogeneity exercise, we show that, within the euro area, countries with more rigid labor markets have a more pronounced kink, suggesting that labor market rigidity could play an important role for the difference between the two currency areas. Finally, we ask to what extent the labor market has driven the post-Covid inflation surge, given the strong nonlinearity we document. Our estimates imply that despite this nonlinearity, the majority of the inflation surge was due to factors other than tight labor markets.

This paper contributes to the literature on a possibly nonlinear Phillips curve by being the first to provide solid empirical evidence of a nonlinearity, using the unique setting of the euro area. While there are various theoretical reasons for a nonlinear Phillips curve (see Section 2.2 for a brief overview), solid empirical evidence of this nonlinearity is scarce, especially for the euro area.[3] On the United States, Hooper et al. (2020) and Smith et al. (2024) find evidence of a steeper Phillips curve at low levels of unemployment, using cross-sectional variation from MSA-level data.[4] Gitti (2024) confirms these results, improving both the measurement of labor market tightness and the identification approach relative to previous papers. It should be

[3] Moretti et al. (2019) and Byrne and Zekaite (2020) investigate a potential nonlinearity of the euro area (wage) Phillips curve, reaching different conclusions. Both studies use only aggregate data, however, and are thus subject to the various criticisms mentioned above and discussed in detail in Furlanetto and Lepetit (2024).

[4] There is also evidence for a nonlinear *wage* Phillips curve in the United States, based on cross-sectional data. See, for example, Kumar and Orrenius (2016) and Donayre and Panovska (2016).

noted, however, that Beaudry et al. (2025) review the evidence of a nonlinear Phillips curve in the United States, concluding that it is "very fragile", particularly after fully accounting for time fixed effects, which are crucial to absorb inflation expectations (see also Doser et al., 2022). We do include a full set of time fixed effects in our regressions. On Europe, Smith et al. (2024) also include some tests with respect to the nonlinearity in the European Union (EU, incl. UK), but their focus is on investigating structural breaks in the (linear) Phillips curve in the United States.

Our paper expands on this literature in a few dimensions: First, in contrast to most of the empirical evidence of regional Phillips curves mentioned above, we focus on the euro area. The euro area lends itself naturally to estimating the Phillips curve with regional data. Given the relatively low level of integration of labor markets (and low geographic mobility), euro area member countries are a plausible proxy of local labor markets, which is crucial for inferring the slope of the aggregate Phillips curve from regional estimates. Our results are therefore novel even in terms of the linear Phillips curve estimates. Second, our main focus on identifying kinks in the euro area Phillips curve allows us to expand substantially on the findings in Smith et al. (2024). In addition to methodological advantages that improve measurement and identification, we show robustness of our main result to various alternative specifications, explore the role of labor market rigidities as a mechanism behind the nonlinearity, and quantify the extent to which tight labor markets have contributed to euro area inflation over time.[5]

The remainder of the paper is structured as follows. Section 2 outlines the theoretical framework that links the (potentially nonlinear) aggregate Phillips curve to cross-country variation. Section 3 details the empirical strategy we use to identify its slope(s). Section 4 presents our main results. Section 5 explores the role of labor market rigidity for kinks in the Phillips curve. Section 6 quantifies the extent to which tight labor markets have historically contributed to euro area inflation, given the nonlinearity in the Phillips curve we document. Section 7 concludes.

2 The regional Phillips curve from theory to empirics

In this section, we begin by linking the slope of the aggregate Phillips curve with the one from a regional Phillips curve, following Hazell et al. (2022). We then allow the slope of this regional Phillips curve to vary depending on the level of labor market slack, in order to test for the nonlinearity of the aggregate Phillips curve using regional data.

[5] Our paper has various methodological advantages relative to Smith et al. (2024): First, we use an instrumental variable approach to purge results from endogeneity bias due to contemporaneous, regional supply shocks. Second, we focus on a monetary union (the euro area) instead of a political one (the EU incl. UK), allowing for a cleaner identification. Third, following Hazell et al. (2022), we use services inflation as a proxy for non-tradable inflation instead of headline inflation. Finally, we use quarterly instead of yearly data.

2.1 Inferring the aggregate Phillips curve slope from regional data

To fix ideas, consider the standard linear New Keynesian Phillips curve (Galí, 2015)

$$\pi_t = \kappa \hat{u}_t + \beta E_t \pi_{t+1} + \nu_t, \tag{1}$$

according to which inflation π_t is determined by three factors: labor market slack $\hat{u}_t$ is the unemployment gap, defined as the unemployment rate u_t in deviation from its natural level u_t^n, expected inflation $E_t\pi_{t+1}$, and cost-push shocks ν_t.

Hazell et al. (2022) show that in their two-country New Keynesian model, the structural parameter κ in Equation (1) can be inferred from the region-level (or, in our case, country-level) equation

$$\pi_{ct}^N = \psi \tilde{u}_{ct} + \delta \hat{p}_{ct}^N + \alpha_c + \gamma_t + \omega_{ct}^N, \tag{2}$$

where π_{ct}^N is inflation of non-tradables in country c in period t, $\tilde{u}_{ct}$ is a transitory component of the unemployment gap, $\hat{p}_{ct}^N = p_{ct}^N/p_{ct} - 1$ is the deviation of relative prices of nontradables from their steady state level of 1, and α_c and γ_t are country and period fixed effects, respectively. The inclusion of γ_t is crucial in this setting in order to absorb long-run inflation expectations, which are closely correlated in a monetary union, but hard to measure. Furthermore, they absorb aggregate supply (and demand) shocks, including the endogenous response of monetary policy to them. Finally, ω_{ct}^N is the discounted sum of future cost-push shocks. The estimable coefficient $\psi = \frac{\kappa}{1-\beta\rho_u}$ directly maps to its structural counterpart κ, as we have outlined in Appendix A, following Hazell et al. (2022).

2.2 Incorporating potential nonlinearities

The theoretical literature has established various mechanisms that may endogenously lead to a nonlinear Phillips curve. The most widely used class of such models features some form of downward nominal wage rigidity (see, e.g., Benigno and Ricci, 2011, Daly and Hobijn, 2014, Burgert et al., 2021, Barnichon et al., 2022, Benigno and Eggertsson, 2023 or Gitti, 2024). With much slack in the labor market, there is little downward pressure on wages and thus prices. Conversely, in tight labor markets, wages are more flexible and rise quickly with the utilization of labor, and are then (partially) passed on to consumers in the form of higher prices.

Another mechanism that may lead to a nonlinear Phillips curve is state-dependent pricing. In a menu cost model (e.g., Golosov and Lucas, 2007), large shocks lead to a higher share of price-adjusting firms than small ones, which is a key determinant of the slope of the Phillips curve (de Veirman, 2023, Karadi et al., 2024). Additionally, an empirical finding in the literature is that prices are more flexible during expansions than during recessions (Wulfsberg, 2016), meaning that the frequency of price adjustment and thus the Phillips curve slope increases when demand is high (and labor markets are tight), as in Gasteiger and Grimaud (2023) or Blanco et al. (2024). Finally, another way to micro-found nonlinearities in the Phillips curve is via quasi-kinked demand for goods as in Kimball (1997), see e.g. Harding et al. (2022).

Motivated by these theoretical arguments, but without assuming a specific one of them, we incorporate the nonlinearity in labor market slack into the aggregate Phillips curve as follows:

$$\pi_t = \kappa \hat{u}_t + \tau 1\{\text{tight}_t\}\hat{u}_t + \beta E_t \pi_{t+1} + \nu_t, \tag{3}$$

where $1\{\text{tight}\}$ is a dummy variable equal to one when a labor market is tight and zero otherwise. The slope of the structural Phillips curve is then equal to κ in normal times (i.e., if the labor market is not tight) and to $\kappa + \tau$ when the labor market is tight. Similarly to the linear case, we can derive its equivalent which is estimable with panel data:

$$\pi_{ct}^N = \psi \tilde{u}_{ct} + \phi 1\{\text{tight}_{ct}\}\tilde{u}_{ct} + \delta \hat{p}_{ct}^N + \alpha_c + \gamma_t + \omega_{ct}^N, \tag{4}$$

with $\psi = \frac{\kappa}{1-\beta\rho_u}$, $\phi = \frac{\tau}{1-\beta\rho_u}$ and $\delta = \frac{\lambda}{1-\beta\rho_{pN}}$.[6]

Notice that the slopes of the *contemporaneous* structural Phillips curve are κ and τ, whereas the coefficients we estimate are ψ and ϕ. The latter sets of parameters take into account that a surprise in the unemployment rate in period t also affects unemployment rates in subsequent periods, because of their autocorrelated nature. This future change in labor market slack has its own effect on inflation pressure, which is captured in ψ and ϕ, but not in κ and τ.

3 Empirical strategy

We now describe the baseline specification as well as the data we use to estimate Equation (4). Moreover, we discuss the main threats to identification in this setting and describe what we do to alleviate them.

3.1 Estimating equation

Our baseline regression specification is:

$$\pi_{ct}^S = \psi \hat{u}_{ct} + \phi 1\{\text{tight}_{ct}\}\hat{u}_{ct} + \eta 1\{\text{tight}_{ct}\} + \delta \hat{p}_{c,t-1}^S + \alpha_c + \gamma_t + \epsilon_{ct}, \tag{5}$$

where π_{ct}^S is quarter-on-quarter annualized inflation of services in country c and period t. We use inflation of services as a proxy for non-tradable inflation because of the predominantly domestic nature of service prices. We want to highlight that much of the literature on regional Phillips curves uses headline inflation, which leads to an underestimation of the slope of the Phillips curve, because it is contaminated by shocks from all regions through the inflation on

[6] As is outlined in Appendix A, we need to make the additional assumption that $E_t 1\{\text{tight}_{c,t+j}\} = E_t 1\{\text{tight}_{c,t}\}$, i.e. that the labor market is expected to stay in its current tightness regime. Given the persistence of these regimes, this is a justifiable assumption even in light of transitory movements in $\tilde{u}_{ct}$.

tradable goods. $\hat{u}_{ct}$ is the unemployment gap;[7] $\hat{p}^S_{c,t-1}$ is the relative price of services, lagged by one quarter to avoid a mechanical correlation; α_c and γ_t are country and period (quarter-year) fixed effects. We weight regressions by the country's share of overall euro area employment in the year 2000. Standard errors allow for arbitrary clustering at the time period level.

3.2 Data

We use various data sources to estimate Equation (5) at a quarterly frequency for the period between 2001Q1 and 2024Q4 for a panel of 18 euro area countries.[8] We use the "HICP: Services" indices from Eurostat as a proxy for non-tradable prices. As these are not seasonally adjusted by Eurostat, we conduct our own seasonal adjustment using X-13ARIMA-SEATS before constructing annualized quarter-on-quarter growth rates in percent. The fact that our main dependent variable enters in sequential growth rates is a deviation from much of the literature, which typically uses year-on-year changes. As explained in Section 2, κ and τ identify the *contemporaneous* slope of the Phillips curve. If the dependent variable entered in year-on-year growth rates, i.e. the cumulated sum of our measure for the current and three previous quarters, the effect of labor market slack on inflation are likely to be underestimated. We therefore find sequential growth rates more appropriate.

To measure labor market slack, we take the difference between the unemployment rate (Eurostat) and the non-accelerating wage rate of unemployment (NAWRU, European Commission), which we interpolate to the quarterly frequency, as it is published only at the annual frequency. We define labor markets as "tight" in all quarters in which this unemployment gap was at least one standard deviation below its long-run average, measured separately for each country. This definition implies that, on average, labor markets were tight 15% of the time.[9]

For the relative price of services, we use the ratio of our seasonally adjusted "HICP: Services" and overall core prices "HICP: All items excluding energy, food, alcohol and tobacco", and subtract a linear country-level trend.[10]

3.3 Identification

There are three main identification issues when estimating the effect of economic slack on inflation using the contemporaneous correlation between the two variables. First, aggregate

[7] We use the raw measure of the unemployment gap instead of its transitory component, as in Equation (4). See Appendix A for details.

[8] These include all 20 current euro area member countries, except for Croatia, which only joined the euro area in 2023, and Malta, for which the data is incomplete. We do include the other newer member countries have all been part of the European Exchange Rate Mechanism (ERM) with their currencies pegged to the euro since the early part of our sample period.

[9] Table B1 and Figure B1 in the Appendix contain some descriptive statistics on labor market tightness across countries and over time. We choose a negative cutoff value because euro area labor markets have tended to be loose over the period we study. In Figure C1, we show that results are robust to various alternative cutoff values.

[10] In the theory, $\hat{p}^S_{ct}$ has a steady state level of 1, but in practice, core service prices grow at a higher rate than core goods prices on average, making the raw ratio non-stationary.

supply shocks may affect both inflation and economic slack at the same time. For example, a negative oil supply shock may increase inflation and economic slack, biasing the estimated coefficient upwards. We address this issue by including period fixed effects in our baseline specification. As discussed in more detail in Appendix A and shown in Hazell et al. (2022), this also assuages concerns about contemporaneous movements of area-wide inflation expectations.

Second, in our cross-country setting, it is possible that other differences across countries (e.g., in inflation expectations, labour market institutions, norms) lead to spurious correlation between inflation and economic slack. For this reason, we include country fixed effects in our baseline specification, which absorb constant differences across countries.

Third, while period fixed effects account for *aggregate* supply shocks affecting all countries in the same period, it remains possible that supply shocks hit, or transmit throughout, the various economies differentially. Therefore, OLS might be contaminated by *regional* supply shocks that affect unemployment and inflation simultaneously even in the presence of period fixed effects. To address this concern, we use the four-quarter lagged unemployment gap and labor market tightness indicator as instruments and estimate Equation (5) using two-stage least squares (2SLS). To the extent that the effect of such regional supply shocks on quarter-on-quarter services inflation and labor market tightness has vanished after about one year, using these instruments should assuage concerns about simultaneity.

4 Results

4.1 The slope of the linear Phillips curve in the euro area

We start by estimating the linear Phillips curve, i.e., excl. the interaction with the labor market tightness indicator in Equation (5), in columns (1)-(4) of Table 1. Panel A presents OLS estimates, Panel B those from a reduced-form estimation in which the instruments are used as the explanatory variable directly, and Panel C those from the 2SLS estimation, where the contemporaneous unemployment gap and the labor market tightness indicator are instrumented with their four-quarter lags.

The negative coefficients in all four columns and all three panels suggest that looser labor markets cause lower inflation, as implied by the theoretical Phillips curve relationship. Column (1) shows the estimates from a version without any fixed effects. These coefficients may pick up observable or unobservable differences in country characteristics (in particular expected inflation) or aggregate supply shocks affecting labor markets and inflation. Columns (2) and (3) separately include country and time fixed effects, whereas in column (4) they enter jointly. Notably, the point estimate becomes only somewhat smaller (in absolute terms) when including time fixed effects, suggesting that aggregate supply shocks have tended to cause only limited co-movement of inflation and unemployment in the euro area. Overall, point estimates across specifications with different fixed effects are remarkably stable.

Results are also remarkably stable across the three panels in Table 1. Compared to the (potentially endogenous) OLS estimates, 2SLS estimates tend to be only marginally larger (in absolute terms), perhaps due to measurement error biasing the OLS estimates towards zero, again suggesting that most fluctuations in the unemployment gap in the euro area are driven by demand-side fluctuations.

The point estimate of our preferred specification in column (4) of Panel C implies that a one percentage point increase (decrease) in the unemployment gap lowers (raises) services inflation (quarter-on-quarter, annualized) by about 0.22 percentage points, with a standard error of 0.03 and a 95% confidence interval between -0.17 and -0.28. The standard deviation of the unemployment gap in the euro area was 1.27 over the sample, translating – according to this linear estimate – to an effect on services inflation of less than 0.3 percentage points, annualized. This is around a fourth of the standard deviation in services inflation. We conclude from this exercise that while the slope of the linear Phillips curve in the euro area is clearly not zero, it is quite flat. Having said that, our point estimates are about twice as large in absolute terms as the equivalent estimate of the closest specification in Hazell et al. (2022), which estimates the Phillips curve with regional data for the United States.[11]

4.2 The slopes of the nonlinear Phillips curve in the euro area

Next, in columns (5)-(8) of Table 1, we allow for a kink in the Phillips curve estimation. The coefficient on the unemployment gap reflects the slope of the Phillips curve when labor markets are roughly balanced or have slack. The coefficient on the interaction between the unemployment gap and the tightness indicator reflects the *additional* slope when labor markets are tight.

In all our specifications, regardless of whether we include country and/or time fixed effects or estimate the relationship with OLS or 2SLS, the estimate of ϕ is negative and statistically significantly different from zero at all conventional levels.[12] This implies that the Phillips curve is steeper when labor markets are tight.

Our estimates imply that this steepening is quantitatively meaningful. In the OLS specification with country and time fixed effects (column (8) of Table 1.A), ϕ is estimated to be -0.34. This implies a point estimate of the slope of the Phillips curve of -0.53 when labor markets are tight, almost three times the magnitude compared to normal times. In the 2SLS specification in Panel C, this wedge increases even further, implying a Phillips curve that is almost 15 times as steep in tight labor markets as in normal times (with $\psi = -0.07$ and $\phi = -1.01$).[13] The estimate of ϕ does, however, have a relatively large standard error, implying a 95% confidence

[11] $\psi = -0.22$ directly maps to an estimate of the structural Phillips curve parameter κ. Given a β of 0.99 and a first-order auto-regressive coefficient of 0.99 for the unemployment gap, we estimate a κ of -0.0052, very similar to the direct estimates of κ provided by Hazell et al. (2022) for the United States.

[12] 2SLS results are unlikely to suffer from substantial bias from weak instruments, as the Kleibergen-Paap F statistics in all specifications are either above or just slightly below the Stock-Yogo critical values for a maximum relative bias of 10% at a 5% significance level (see Stock and Yogo, 2002).

[13] Figure B2 depicts the steepening implied by our preferred estimates.

Table 1: Baseline Phillips curve estimates with varying fixed effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dependent variable: Services inflation π^S_{ct}, q/q saar							
	Linear				Nonlinear			
	A. OLS							
Unemployment gap_{ct}	-0.17***	-0.22***	-0.13***	-0.18***	-0.14***	-0.21***	-0.11***	-0.18***
	(0.03)	(0.03)	(0.02)	(0.03)	(0.02)	(0.03)	(0.02)	(0.03)
— $\times$ $1\{\text{tight}_{ct}\}$					-0.44***	-0.28***	-0.50***	-0.34***
					(0.09)	(0.10)	(0.09)	(0.10)
	B. Reduced form							
Unemployment $\text{gap}_{c.t-4}$	-0.16***	-0.21***	-0.14***	-0.19***	-0.13***	-0.18***	-0.11***	-0.17***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
— $\times$ $1\{\text{tight}_{c,t-4}\}$					-0.66***	-0.54***	-0.51***	-0.34***
					(0.08)	(0.11)	(0.07)	(0.09)
	C. IV							
Unemployment gap_{ct}	-0.19***	-0.25***	-0.16***	-0.22***	-0.08***	-0.11**	-0.06*	-0.07
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.05)	(0.03)	(0.06)
— $\times$ $1\{\text{tight}_{ct}\}$					-0.95***	-0.95***	-0.95***	-1.01***
					(0.22)	(0.30)	(0.25)	(0.38)
Kleibergen-Paap F-stat.	1,037.7	766.2	1,741.3	1,326.3	15.1	10.2	11.7	6.0
Stock-Yogo crit. values	16.4	16.4	16.4	16.4	7.0	7.0	7.0	7.0
Country fixed effects		✓		✓		✓		✓
Time fixed effects			✓	✓			✓	✓
Observations	1,727	1,727	1,727	1,727	1,727	1.727	1,727	1,727

Notes: Estimation of Equation (5), showing the estimates for slope of the Phillips curve (ϕ) and – in columns (4) through (8) – the additional effect on the slope in tight labor markets (ψ). The dummy variable $1\{\text{tight}_{ct}\}$ is equal to one in time where the different between a country's unemployment rate and the natural rate of unemployment is at least one standard deviation below is full-sample average. All columns include the relative price level of services as a covariate. Country and time fixed effects (α_c and γ_t) are included as indicated. The sample consists of 18 countries and 92 periods (over 23 years). Standard errors are robust against heteroskedasticity and allow for arbitrary clustering by time. We weight regressions by the country's share of overall euro area employment in the year 2000. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$. The Stock-Yogo critical values are those for a test with approximately a 5% significance level, of the hypothesis that the maximum relative bias from weak instruments is at least 10%.

interval of between -0.5 and -1.4.

4.3 Robustness

In Appendix C, we report results from several robustness checks, which are briefly summarized here. First, results are insensitive to our preferred definition of "tight" labor markets (columns 1–3 of Table C1 and Figure C1). While we find that a kink in the Phillips curve around an unemployment gap of -1 is most plausible, our finding does not hinge on this exact threshold or even the functional form of our model. Second, results are robust to using alternative measures of slack such as the vacancy-to-unemployment ratio and an unemployment gap estimated using an HP filter (columns 4–6 of Table C1). This addresses concerns regarding potential mismeasurement and mechanical correlation resulting from the use of the European Commission's NAWRU. Third, results hold when using subsamples of the data (columns 7–8 of Table C1 and Figure C2), showing that our results do not rely on a specific period (most notably the post-pandemic inflation surge) or country. Fourth, results remain similar when changing parts of the specification to match previous contributions in the literature, such as using y/y inflation, headline inflation as the dependent variable or by explicitly including long-term inflation expectations of professional forecasters as a control variable (columns 9–12 of Table C1 and Table C2). Overall, we conclude that there is solid evidence of the Phillips curve in the euro area becoming substantially steeper when labor markets are sufficiently tight.

5 Labor market rigidity and kinks in the Phillips curve

While studies using panel data from the United States have also found evidence of kinks in the Phillips curve, they tend to be less pronounced than what we find for the euro area. For example, Smith et al. (2024) find that the Phillips curve is about twice as steep when the unemployment rate is below 5% and Gitti (2024) documents a slope that is almost three times as large in tight labor markets. Beaudry et al. (2025) show that the nonlinearity might vanish altogether, depending on the specification. In the results we have presented in the previous two sections, the size of this wedge is around five, on average across the main results and robustness tests, and can range up to a factor of 15 in the case of our main IV specification.

In this section, we test our conjecture that the more pronounced nonlinearity we document stems from the euro area's less flexible labor markets relative to the United States. The main idea is that in rigid labor markets, labor supply is not as flexible, such that firms need to offer workers a larger wage increase for them to switch occupations, industries or locations, or enter the labor market. For example, with a high degree of job protection, giving up a secure job for a new one in a tight sector is more risky as fewer opportunities tend to open up in normal times. Similarly, in countries with generous unemployment benefits, often those that also feature high levels of job protection, the same level of labor market tightness may require higher wage growth to induce unemployed workers to enter the labor force, leading to more inflation.

To test this hypothesis, in Table 2 we exploit the euro area's country-level variation in labor market rigidity. Besides the interaction of labor market slack with the tightness indicator $1\{\text{tight}_{ct}\}$, we include an additional interaction with a measure of labor market rigidity. Labor market rigidity is difficult to measure and compare across countries. As a unified proxy, we use the 2000 labor market regulation subindex of the Economic Freedom of the World (EFW) index, produced by the Fraser Institute (Gwartney et al., 2024). This composite index measures regulations and frictions around hiring, firing, working hours and wage-setting. We split the sample of 18 countries at the median into high-rigidity and low-rigidity countries. Alternatively, a more transparent but less direct measure of labor market rigidity is the average unemployment rate over the sample period. Countries with less flexible labor markets should on average have higher unemployment.[14]

Results are presented in Table 2. Columns (1) and (3) show how the linear slope of the Phillips curve depends on labor market rigidity. The evidence is inconclusive. While countries with a higher EFW index seem to have a substantially steeper Phillips curve on average, countries with a higher average unemployment rate do not. Turning to the impact of labor market rigidity on the extent of the steepening of the Phillips curve, columns (2) and (4) include a triple interaction between the unemployment gap, the tightness indicator, and the labor market rigidity indicators. The coefficient on the triple interaction term has a statistically significantly negative and large coefficient in both cases, indicating that indeed the steepening of the Phillips curve in tight labor markets is more pronounced in countries with relatively rigid labor market institutions.

These results support our conjecture that a large share of the very steep Phillips curve in tight labor markets we estimate for the euro area is driven by relatively rigid labor markets, where demand shocks are less easily met with additional labor supply, thus leading to stronger inflation pressure. Because euro area labor markets are generally more rigid than those in the United States, this finding may explain why we find a more pronounced kink for the euro area Phillips curve than the literature has found for the United States.

[14] The classification of countries based on these two measures is listed in the summary statistics in Table B1 in the Appendix.

Table 2: Phillips curve estimates with labor market rigidities

	(1)	(2)	(3)	(4)
	Labor market regulation index		Average unemployment rate	
Unemployment gap$_{ct}$	-0.18***	-0.18***	-0.30***	-0.41***
	(0.02)	(0.03)	(0.07)	(0.13)
— × 1{high rigidity$_c$}	-0.44***	0.37*	0.02	0.37***
	(0.13)	(0.21)	(0.07)	(0.13)
— × 1{tight$_{ct}$}		-0.77**		1.13
		(0.31)		(0.79)
— × 1{tight$_{ct}$} × 1{high rigidity$_c$}		-2.24***		-3.60***
		(0.63)		(0.93)
Country and time fixed effects	✓	✓	✓	✓
Observations	1,727	1,727	1,727	1,727

Notes: Estimation of Equation (5) with an additional interaction term for labor market rigidity, proxied by two different measures. (1) and (2) use the Fraser Institute's EFW subindex on labor market regulation, (3) and (4) use the country-level average unemployment rate as a sufficient statistic for labor market frictions. The sample consists of 18 countries – 9 of which are classified as rigid – and 92 periods (over 23 years). Contrary to the baseline results (see Table 1), observations are not weighted by country-level employment, to exploit the full variation in labor market rigidity across all countries weighted equally. Standard errors are robust against heteroskedasticity and allow for arbitrary clustering by time periods. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

6 To what extent did labor market tightness contribute to the post-pandemic inflation surge?

If the Phillips curve is quite flat, the contribution of labor market tightness to inflation ($\hat{\psi} \cdot \hat{u}_t$) is always small. However, with a nonlinear Phillips curve of the form we estimate, labor market tightness might explain a larger share of inflation, particularly of inflation surges. In this section, we thus quantify the extent to which the strong kink we find in the euro area Phillips curve can explain historical variation in (services) inflation. In particular, we are interested in quantifying the contribution of the labor market to services inflation in the post-Covid inflation surge.

According to the European Commission estimate, the unemployment gap at the euro area level averaged -0.3 in the years 2022 and 2023, the years with large price increases for services. This is less than a standard deviation below the euro area average and thus still in the flat region of our estimated Phillips curve, with a slope of roughly -0.1. Consequently, the contribution of the labor market to services inflation during this period would average 0.03 percentage points, more than two orders of magnitude lower than actual services inflation at the time.

However, there is considerable uncertainty around the estimate of the unemployment gap. The European Commission reports standard errors of its NAWRU estimate of roughly one percentage point. If the natural rate was underestimated and thus the unemployment gap overestimated, the contribution of the labor market to inflation could be significantly higher, because of the nonlinear nature of the relationship. But by how much?

To quantify the labor market's contribution to services inflation under this uncertainty, we simulate 1,000 alternative paths of the natural rate with similar properties.[15] Figure 2(a) shows the distribution of the resulting unemployment gaps over time. In the last two years of the time series, the 90% interval of plausible values of the unemployment gap lies between -3 and 2.5 percentage points. A share of 30% of all the simulated paths have at least one quarter in 2022/23 where the unemployment gap is at least one standard deviation below its historical mean and thus the nonlinear component of the Phillips curve – $\hat{\phi} \cdot \hat{u}_t$ – adds to the contribution.

Additionally, we account for parameter uncertainty in the Phillips curve estimation by drawing 100 random variables from the parameter distributions of $\hat{\psi}$ and $\hat{\phi}$ and apply them to each of the paths of the unemployment gap. The resulting 100,000 draws of the contribution of the labor market to services inflation are depicted in Figure 2(b), whereas we transform the quarter-on-quarter contributions back to year-on-year rates for the sake of readability.

Figure 2: Historical contribution of Phillips curve to euro area services inflation

(a) Unemployment gap with confidence intervals (b) Contribution to services inflation

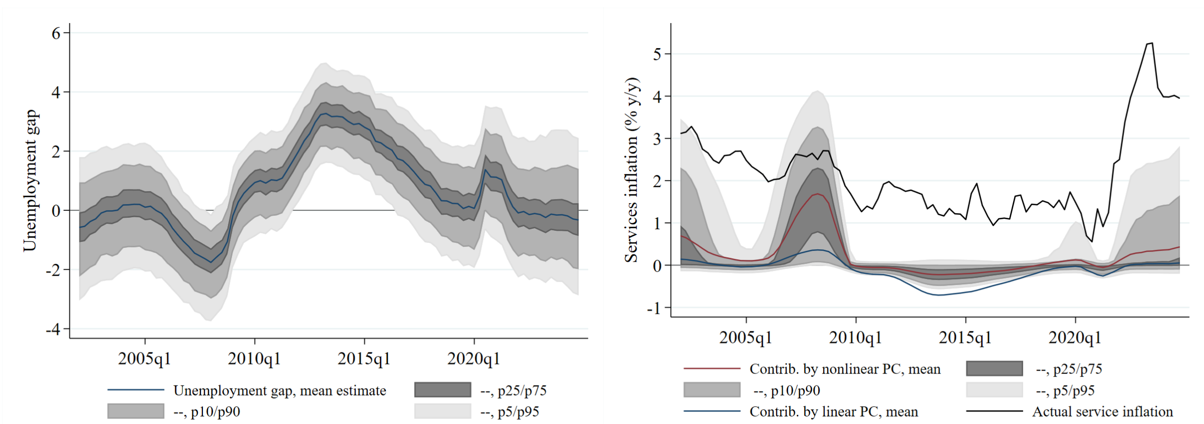

Notes: Panel (a) shows the mean estimate of the unemployment gap published by the European Commission, gray areas show the distribution of 1,000 simulations of the unemployment gap that accounts for the uncertainty around the mean estimate, which we take to be 1 percentage point. Under this and parameter uncertainty in the Phillips curve estimation, Panel (b) depicts contributions of labor market slack to historical services inflation in the euro area, i.e., $\hat{\psi}^{(d)}\hat{u}_t^{(d)} + \hat{\phi}^{(d)}\hat{u}_t^{(d)}1\{\text{tight}_t^{(d)}\}$ for each of the draws d. The blue line additionally shows the mean contribution using the linear Phillips curve estimate.

Under the likely paths of the unemployment gap in the post-pandemic period, its contribution to services inflation stays close to zero, which is indicated by the darkest shade in Panel (b), showing 50% of all resulting contributions of the unemployment gap to services inflation. However, the distribution of the contributions in 2022/23 is heavily skewed. The 90th percentile shows a contribution of the unemployment gap to services inflation of around 1.5 percentage points and the 95th percentile of almost 3 percentage points.

[15] First, they follow a random walk, just like the mean estimate measured by the Commission. Second, the distribution of the deviation between all our simulated paths and the published mean estimate has a mean of zero and a standard deviation of 1 percentage point.

Nevertheless, under the likely true values of the unemployment gap and the (nonlinear) Phillips curve coefficients, the contribution of labor market tightness to post-pandemic inflation was positive but quite small. Instead, these calculations suggest that the inflation surge in the euro area was predominantly due to factors other than tight labor markets (most notably, supply shocks such as supply bottlenecks and surging natural gas prices, and their second-round effects on services prices through a catch-up of nominal wages). We confirm this by showing estimates of the time fixed effects in Figure B3 in the Appendix. They imply a contribution to the inflation surge between 2021 and 2023 of at least 3 percentage points.

The contributions computed based on our nonlinear Phillips curve estimates also track the historical dynamics of services inflation significantly better than the linear slope. For example, effects from the labor market cooling in the early 2000s, the tightening prior to the Global Financial Crisis as well as the limited disinflation coming from the substantial labor market slack of the 2010s are all well-reflected in historical services inflation rates. This indicates that despite its limited role for the post-Covid inflation surge, labor market slack and the nonlinear Phillips curve were significant drivers of fluctuations in inflation in the first 20 years of the euro area.

7 Conclusion

This paper provides robust evidence of a nonlinear Phillips curve in the euro area. During normal times, the relationship between labor market slack and inflation is likely existent but quite weak. In contrast, when the labor market runs hot, i.e., when unemployment is sufficiently below its natural level, the Phillips curve becomes much steeper. This implies, first and foremost, that the central bank's inflation-unemployment trade-off is state-dependent and more 'favorable' (i.e., giving the central bank more leverage) in a well-utilized labor market. Second, evidence in this paper suggests that labor market institutions may play an important role for the steepening of the Phillips curve in tight labor markets. Our results suggest that the steepening is more pronounced in more rigid labor markets, and consequently, may explain the stronger steepening we find for the euro area compared to existing studies on the United States. A third important lesson of this paper is that a tight labor market can be a significant contributor to inflation developments, but that it likely was not during the post-pandemic inflation surge in 2022/23. The clear steepening of the negative correlation between inflation and the unemployment rate depicted in Figure 1 therefore likely reflects a common euro area supply shock at a time when unemployment rates were historically low, rather than movements on a (steep or kinky) Phillips curve.

References

Barlevy, Gadi, R. Jason Faberman, Bart Hobijn, and Ayşegül Şahin, "The Shifting Reasons for Beveridge Curve Shifts," *Journal of Economic Perspectives*, May 2024, *38* (2), 83–106.

Barnichon, Régis and Adam Hale Shapiro, "Phillips meets Beveridge," *Journal of Monetary Economics*, 2024, *148.*

Barnichon, Regis, Davide Debortoli, and Christian Matthes, "Understanding the Size of the Government Spending Multiplier: It's in the Sign," *Review of Economic Studies*, 2022, *89* (1), 87–117.

Beaudry, Paul, Chenyu Hou, and Franck Portier, "On the Fragility of the Nonlinear Phillips Curve View of Recent Inflation," *NBER Working Paper*, 2025, *33522.*

Benigno, Pierpaolo and Gauti B. Eggertsson, "It's baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve," *NBER Working Paper*, 2023, *31197.*

_ **and Luca Antonio Ricci**, "The Inflation-Output Trade-Off with Downward Wage Rigidities," *American Economic Review*, 2011, *101* (4), 1436–66.

Bernanke, Ben and Olivier Blanchard, "What caused the US pandemic-era inflation?," *Peterson Institute for International Economics Working Paper*, 2023, (23-4).

Beyer, Robert C. M. and Frank Smets, "Labour market adjustments and migration in Europe and the United States: how different?," *Economic Policy*, 2015, *30* (84), 643–682.

Blanchard, Olivier, "The Phillips curve: back to the'60s?," *American Economic Review: Papers and Proceedings*, 2016, *106* (5), 31–34.

Blanco, Andrés, Corina Boar, Callum J. Jones, and Virgiliu Midrigan, "The Inflation Accelerator," *NBER Working Paper*, 2024, *32531.*

Burgert, Matthias, Philipp Pfeiffer, and Werner Roeger, "Fiscal policy in a monetary union with downward nominal wage rigidity," *Swiss National Bank Working Paper*, 2021, *2021-16.*

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

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *The Review of Economics and Statistics*, 08 2008, *90* (3), 414–427.

Daly, Mary C. and Bart Hobijn, "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money, Credit and Banking*, 2014, *46* (S2), 51–93.

de Veirman, Emmanuel, "How Does the Phillips Curve Slope Vary with Repricing Rates?," *ECB Working Paper*, 2023, *2804*.

Decressin, Jörg and Antonio Fatás, "Regional labor market dynamics in Europe," *European Economic Review*, 1995, *39* (9), 1627–1655.

Donayre, Luiggi and Irina Panovska, "Nonlinearities in the U.S. wage Phillips curve," *Journal of Macroeconomics*, 2016, *48*, 19–43.

Doser, Alexander, Ricardo Nunes, Nikhil Rao, and Viacheslav Sheremirov, "Inflation expectations and nonlinearities in the Phillips curve," *Journal of Applied Econometrics*, 2022.

Fitzgerald, Terry, Callum Jones, Mariano Kulish, and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation?," *American Economic Journal: Macroeconomics*, 2024, *16* (4), 114–142.

Furlanetto, Francesco and Antoine Lepetit, "The Slope of the Phillips Curve," 2024.

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.

Gasteiger, Emanuel and Alex Grimaud, "Price setting frequency and the Phillips curve," *European Economic Review*, 2023, *158*, 104535.

Gitti, Giulia, "Nonlinearities in the Regional Phillips Curve with Labor Market Tightness," 2024.

Golosov, Mikhail and Robert E. Lucas, "Menu Costs and Phillips Curves," *Journal of Political Economy*, 2007, *115*, 171–199.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, "Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages?," *American Economic Review*, 2022, *112* (5), 1437–1474.

Gwartney, James, Robert A. Lawson, Ryan Murphy, Matthew D. Mitchell, Kevin Grier, Robin Grier, and Daniel J. Mitchell, "Economic Freedom of the World: 2024 Annual Report," *Fraser Institute*, 2024.

Hall, Robert E. and Marianna Kudlyak, "The Active Role of the Natural Rate of Unemployment During Cyclical Recoveries," *IZA Discussion Paper*, 2023.

Harding, Martín, Jesper Lindé, and Mathias Trabandt, "Resolving the missing deflation puzzle," *Journal of Monetary Economics*, 2022, *126*, 15–34.

Hazell, Jonathan, Juan Herreño, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from U.S. States," *Quarterly Journal of Economics*, 2022, *137* (3), 1299–1344.

Hooper, Peter, Frederic S Mishkin, and Amir Sufi, "Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating?," *Research in Economics*, 2020, *74* (1), 26–62.

Hristov, Atanas, Christophe Planas, Werner Roeger, and Alessandro Rossi, "NAWRU – Estimation using structural labour market indicators," *European Commission: Directorate-General for Economic and Financial Affairs*, 2017.

Karadi, Peter, Anton Nakov, Galo Nu no, Ernesto Pastén, and Dominik Thaler, "Strike while the Iron is Hot: Optimal Monetary Policy with a Nonlinear Phillips Curve," *CEPR Discussion paper*, 2024, *DP19339.*

Kimball, Miles S., "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit, and Banking*, 1997, *27* (4), 1241–1277.

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

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H Stock, "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve," *American Economic Journal: Journal of Economic Literature*, 2014, *52* (1), 124–188.

McLeay, Michael and Silvana Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 2019, *34*, 199–255.

Moretti, Laura, Luca Onorante, and Shayan Zakipour Saber, "Phillips curves in the euro area," 2019, (2295).

Phillips, A. W., "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957," *Economica*, 1958, *25* (100), 283–299.

Smith, Simon, Allan Timmermann, and Jonathan H. Wright, "Breaks in the Phillips Curve: Evidence from Panel Data," *Journal of Applied Econometrics*, 2024.

Stock, James H and Mark W Watson, "Slack and cyclically sensitive inflation," *Journal of Money, Credit and Banking*, 2020, *52* (S2), 393–428.

_ **and Motohiro Yogo**, "Testing for Weak Instruments in Linear IV Regression," Working Paper 284, National Bureau of Economic Research November 2002.

Wulfsberg, Fredrik, "Inflation and Price Adjustments: Micro Evidence from Norwegian Consumer Prices," *American Economic Journal: Macroeconomics*, 2016, *8* (3), 175–194.

Supplemental appendix (not for publication)

Table of Contents

A Theoretical background

To fix ideas, consider the standard linear New Keynesian Phillips curve (Galí, 2015)

$$\pi_t = \kappa \hat{u}_t + \beta E_t \pi_{t+1} + \nu_t, \tag{6}$$

according to which inflation π_t is determined by three factors: labor market slack $\hat{u}_t$, defined as the unemployment rate u_t in deviation from its natural level u_t^n, expected inflation $E_t\pi_{t+1}$, and cost-push shocks ν_t. Forward-solving Equation (6) and assuming that $\hat{u}_t$ follows an AR(1) process with an autocorrelation coefficient ρ_u yields the estimable equation:

$$\pi_t = \psi \hat{u}_t + E_t \pi_{t+\infty} + \omega_t, \tag{7}$$

where $\psi = \frac{\kappa}{1-\beta\rho_u}$ and $\omega_t = E_t \sum_{j=0}^{\infty} \beta^j \nu_{t+j}$ is the discounted sum of all current and future cost-push shocks.

The estimation of ψ (or implicitly, κ) in this equation from aggregate data is subject to a number of econometric challenges related to endogeneity. First, long-run inflation expectations $E_t\pi_{t+\infty}$ may be correlated with both current inflation and labor market slack. For example, a shock to long-run inflation expectations may raise actual inflation and simultaneously reduce real rates, reducing labor market slack. Second, cost-push shocks may raise inflation and induce the central bank to raise policy rates, raising labor market slack. Both issues could, in principle, be alleviated if properly controlled for or by the use of valid instruments. However, the literature has shown that results are sensitive to particular choices of instruments, control variables or time periods (Mavroeidis et al., 2014).

A recent literature proposes to overcome such endogeneity issues by using a panel of regions in a monetary union, allowing the inclusion of time fixed effects, which absorb both long-run expectations and aggregate shocks (McLeay and Tenreyro, 2019, Hazell et al., 2022, Fitzgerald et al., 2024). For example, Hazell et al. (2022) develop a two-country New Keynesian model in which there is no labor market mobility across regions or, as in our case, countries.[16] They show that the slope of a regional Phillips curve of the form

$$\pi_{ct}^N = \kappa \hat{u}_{ct} + \beta E_t \pi_{c,t+1}^N + \lambda \hat{p}_{ct}^N + \nu_{ct}^N \tag{8}$$

is identical to that of the aggregate Phillips curve, namely κ. There are two differences between Equations (6) and (8). First, the dependent variable is inflation of nontradables at the country level, π_{ct}^N. This is important because prices of tradables are uniform across the monetary union, i.e., remain unaffected by regional slack. The second difference is the presence of the term

[16] This assumption is particularly suitable in the case of the euro area, where labor market mobility between countries is much less prevalent than in the United States, at least in the short run (Decressin and Fatás, 1995). This manifests itself in a higher cross-sectional standard deviation of unemployment rates in Europe (Beyer and Smets, 2015).

$\hat{p}^N_{ct} = p^N_{ct}/p_{ct} - 1$, which controls for the deviation of relative prices of nontradables from their steady state level of 1. To bring Equation (8) to the data, first solve forward to obtain

$$\pi^N_{ct} = E_t \sum_{j=0}^{\infty} \beta^j \left(\kappa \tilde{u}_{c,t+j} + \lambda \hat{p}^N_{ct} \right) + E_t \pi^N_{c,t+\infty} + \omega^N_{ct}. \tag{9}$$

We have introduced $\tilde{u}_{ct} = \hat{u}_{ct} - E_t \hat{u}_{c,t+\infty}$ as a definition of the transitory component of labor market slack and have also used the fact that $E_t \pi_{c,t+\infty} = \frac{\kappa}{1-\beta} E_t \hat{u}_{c,t+\infty}$, which follows directly from the equilibrium version of Equation (7).[17]

From here on, the inclusion of time fixed effects γ_t has several benefits. First, they absorb at least to a large extent the non-observable long-run beliefs about inflation, which should co-move across countries in the same monetary union, and are therefore omitted in Equation (10).[18]

$$\pi^N_{ct} = E_t \sum_{j=0}^{\infty} \beta^j \left(\kappa \tilde{u}_{c,t+j} + \lambda \hat{p}^N_{c,t+j} \right) + \alpha_c + \gamma_t + \omega^N_{ct}. \tag{10}$$

Second, they capture the (endogenous) monetary policy reaction to cost-push shocks, which is by definition the same for all countries in a monetary policy union. Finally, they absorb also other aggregate shocks that may be correlated with inflation and labor market slack. In particular, they absorb all aggregate supply shocks, thus helping to identify the slope of the Phillips curve from variation in (local) demand only.

Equation (10) is not yet estimable because it contains unobserved forward-looking expectations of unemployment and relative prices of nontradables to the extent that they are not uniform across countries. However, with the assumption that $\tilde{u}_{ct}$ and $\hat{p}^N_{ct}$ follow AR(1) processes with autocorrelation coefficients ρ_u and ρ_{pN}, respectively, we can estimate

$$\pi^N_{ct} = \psi \tilde{u}_{ct} + \delta \hat{p}^N_{ct} + \alpha_c + \gamma_t + \omega^N_{ct}, \tag{11}$$

where the estimable coefficients $\psi = \frac{\kappa}{1-\beta\rho_u}$ and $\delta = \frac{\lambda}{1-\beta\rho_{pN}}$ directly map to their structural counterparts.[19]

[17] In our empirical specification, we will use the "raw" measure of the unemployment gap, rather than just the transitory component. This is equal to the assumption that the permanent component is absorbed in the country fixed effects.

[18] In addition to time fixed effects, we also include country fixed effects α_c in Equation (10), in order to account for potential time-invariant differences in inflation and labor market slack across countries.

[19] In the nonlinear case, we need to make the additional assumption that $E_t 1\{\text{tight}_{c,t+j}\} = E_t 1\{\text{tight}_{c,t}\}$, i.e. that the labor market is expected to stay in its current tightness regime. Given the persistence of these regimes, this is a justifiable assumption even in light of transitory movements in $\tilde{u}_{ct}$.

B Additional tables and figures

In this appendix, we show summary statistics and data visualizations as background information.

Table B1: Descriptive statistics

	Wt.	Services inflation		Slack				Rigid	
	Share	Mean(π^S)	Sd(π^S)	Mean(u)	Mean($\hat{u}$)	Sd($\hat{u}$)	Mean(1{tight})	EFW	Mean(u)
Austria	2.66	2.89	1.72	5.37	0.21	0.60	15.63		
Belgium	2.91	2.43	1.38	7.25	0.02	0.72	17.71	✓	
Cyprus	0.22	2.11	2.75	7.64	−0.19	2.78	7.29		
Estonia	0.41	4.25	4.03	8.27	−0.18	2.63	11.46	✓	✓
Finland	1.63	2.49	1.15	8.09	0.71	0.67	18.75		
France	18.20	2.05	0.98	8.78	0.01	0.85	13.54	✓	✓
Germany	28.32	1.90	1.74	5.93	0.21	0.73	8.33		
Greece	3.06	2.10	2.631	15.46	5.70	6.29	12.50		✓
Ireland	1.26	3.08	2.38	7.69	0.00	1.22	8.33	✓	
Italy	16.16	2.03	1.25	9.15	−0.10	1.73	15.63		✓
Latvia	0.65	4.15	4.88	10.38	0.05	2.67	11.46	✓	✓
Lithuania	0.99	4.41	4.28	9.75	0.26	2.92	13.54	✓	✓
Luxembourg	0.19	2.49	1.36	5.02	−0.06	0.62	13.54	✓	
Netherlands	5.82	2.75	1.87	5.56	0.35	1.04	16.67	✓	
Portugal	3.57	2.73	2.43	9.65	0.39	1.09	16.67		✓
Slovakia	1.43	4.83	4.39	11.61	0.25	1.95	20.83	✓	✓
Slovenia	0.65	3.57	3.03	6.33	0.59	1.79	16.67		
Spain	11.85	2.41	1.75	15.48	1.24	4.28	18.75		✓
Euro area	100.00	2.18	1.16	9.00	0.53	1.27	10.42	n/a	n/a

Notes: All numbers in percent. The first column contains country weights we use in our regression, for which we use the share of the country's employment in 2000. The last two columns are time-invariant dummies at the country level that describe whether the labor market is more rigid than the median of countries. To measure labor market rigidity, we use two proxies, namely the labor market regulation subindex of the Economic Freedom of the World (EFW) index, produced by the Fraser Institute, as well as the average unemployment rate over the sample period.

Figure B1: Baseline slack/labor market tightness measures over time

(a) Unemployment gap over time

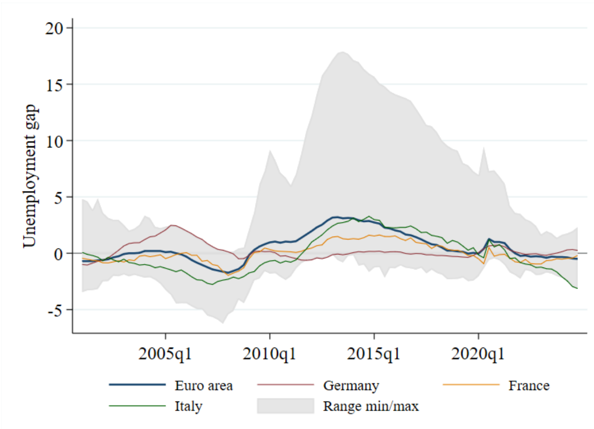

(b) Share of countries with tight labor market

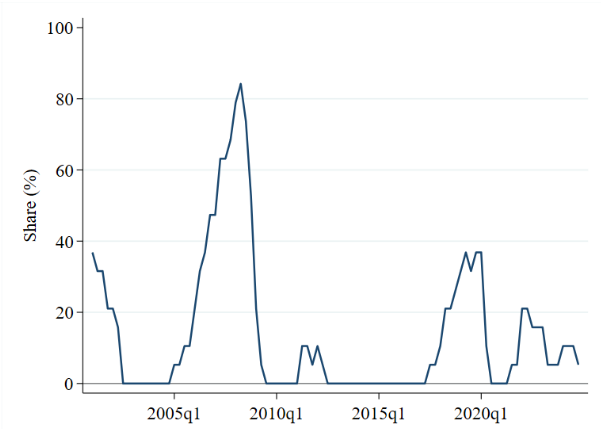

Notes: Panel (a) shows the range of unemployment gap – the difference between the quarterly average unemployment rate and the NAWRU estimate from the European Commission – across the 18 countries in the sample, as well as that of the euro area aggregate and the three largest member states. Panel (b) shows the share of countries that have a standardized (at the country level) unemployment gap of -1 or below, which is our baseline definition of a tight labor market.

Figure B2: Estimates of the Phillips curve with various functional forms

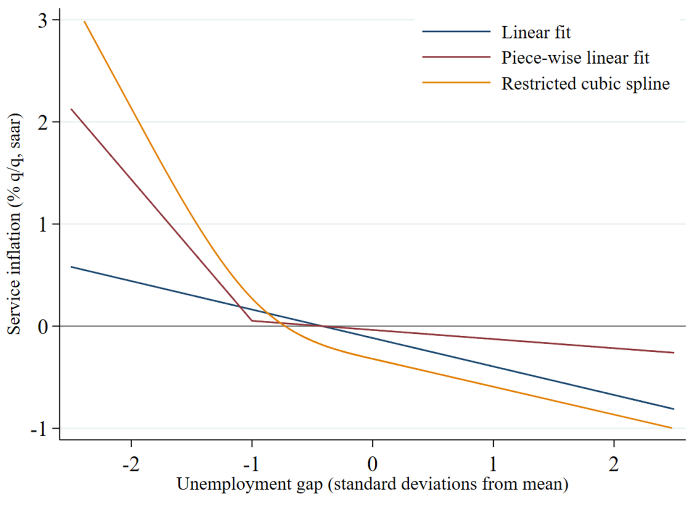

Notes: The orange line depicts the Phillips curve when estimated not with (piece-wise) linear slopes, but with a restricted cubic spline with two knots . Estimates of the Phillips curve with a linear slope of $\psi = -0.22$ (in blue) and the nonlinear version (in red) with $\psi = -0.07$ and $\phi = -1.01$. Both of them are estimated using Equation (5), which is a function in the unemployment gap.

Figure B3: Time fixed effects estimated in the nonlinear specification

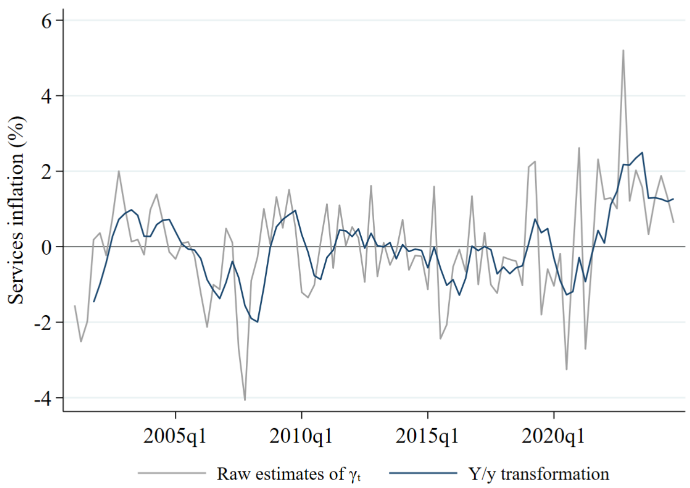

Notes: We depict the estimated γ_t, i.e. the time fixed effects from the estimation of the nonlinear regional Phillips curve (Equation (5)). These absorb aggregate demand and supply shocks, including the monetary policy response to these shocks, as well as the contribution of co-moving inflation expectations.

C Robustness

In this appendix, we show that results are robust to using a variety of alternative specifications, and rule out several alternative explanations.

Definition of tight labor markets. We begin by testing the sensitivity of our results to the definition of "tight" labor markets. In our main specification, the dummy variable $1\{\text{tight}_{ct}\}$ is defined as one for periods in which the unemployment gap was at least one country-specific standard deviation below the long-term average. As a first deviation from this, we consider a somewhat less extreme case, and half a standard deviation as the threshold instead. In that case, twice as many observations are classified as tight labor markets. Rerunning our baseline specification in column (1) of Table C1 with this alternative threshold yields an estimate of ϕ of -0.32, smaller in absolute terms than the corresponding baseline estimate of -1.01 in column (8) of Table 1.C, but still significant at the 5% level. This smaller coefficient still implies that the Phillips curve is substantially steeper in tight periods, in this case by a factor of around three.[20]

In a second set of robustness tests, we show that the kink in the Phillips curve is confirmed when $1\{\text{tight}_{ct}\}$ is a function of the non-standardized unemployment gap. While we think the country-level standardization is sensible to account for different variances of unemployment rates across member countries, this alternative specification may be somewhat easier to interpret as it is measured directly in percentage point deviations. A slight disadvantage of this alternative is, however, that it does not always use variation from all countries, because some never experienced tight labor markets for some thresholds. Columns (2) and (3) of Table C1 present results when using thresholds of -1 and -0.5 percentage points, respectively. Results remain qualitatively unchanged, again suggesting a substantial steepening of the euro area Phillips curve in tight labor markets.[21]

Overall, these results show that the Phillips curve in the euro area becomes substantially steeper if labor markets are sufficiently tight. They also suggest that in the euro area, labor markets can only be considered sufficiently tight if the unemployment rate is meaningfully (at least about half a percentage point) below the natural rate of unemployment.

Measurement of slack. Our baseline measure of labor market slack is the unemployment gap. Subtracting the natural rate from the unemployment rate is important, because failing to account for trends in the natural rate can lead to an underestimation of the slope of the Phillips curve (Hall and Kudlyak, 2023). This is likely to be particularly relevant in the euro area, where structural forces such as demographic trends have led to heterogeneous and nonlinear trends in unemployment rates. These are often unrelated to economic slack and cannot simply be absorbed by time or country fixed effects.

[20] In Figure C1(a), we include results for a broad range of thresholds. Our estimate of ϕ becomes statistically insignificant for values of -0.25 standard deviations or higher, reinforcing our view that labor markets have tended to be rather loose in the euro area in the sample period.

[21] Results for a wider range of thresholds are presented in Figure C1(b).

Table C1: Main robustness tests for nonlinearity of the Phillips curve

	(1)	(2)	(3)	(4)	(5)	(6)
	Kinks in the PC			Measurement of slack		
Definition: slack/tightness	$\hat{u}_{ct}$	$\hat{u}_{ct}$	$\hat{u}_{ct}$	$\hat{\theta}_{ct}$	$\dot{u}_{ct}$	u_{ct}
Definition: $1\{\text{tight}_{ct}\}$:	$< -\frac{1}{2}\sigma_{\hat{u},c}$	< -1	$< -\frac{1}{2}$	$> \sigma_{\hat{\theta},c}$	$< -\sigma_{\dot{u},c}$	$< -\sigma_{u,c}$
Slack/tightness measure$_{ct}$	-0.18***	-0.13***	-0.16***	0.64***	-0.09	-0.20***
	(0.03)	(0.04)	(0.04)	(0.21)	(0.14)	(0.03)
— × $1\{\text{tight}_{ct}\}$	-0.32**	-0.76***	-0.50***	2.51***	-0.87***	-0.27*
	(0.16)	(0.25)	(0.18)	(0.94)	(0.30)	(0.15)
Country FE, time FE	✓	✓	✓	✓	✓	✓
Observations	1,727	1,727	1,727	1,227	1,727	1,727

	(7)	(8)	(9)	(10)	(11)	(12)
	Subsamples		Specification			
Changes rel. to main spec.:	Excl. Covid	ERM only	Unwgt.	π^S_{ct} in y/y	$\pi^{\text{Headline.}}_{ct}$	2-way clust.
Unemployment gap$_{ct}$	-0.08	-0.11**	-0.06	-0.13***	-0.06	-0.07
	(0.07)	(0.05)	(0.05)	(0.04)	(0.05)	(0.12)
— × $1\{\text{tight}_{ct}\}$	-1.12**	-0.87**	-2.12***	-0.87***	-0.61*	-1.01**
	(0.45)	(0.33)	(0.52)	(0.32)	(0.31)	(0.44)
Country FE, time FE	✓	✓	✓	✓	✓	✓
Observations	1,367	1,620	1,727	1,719	1,727	1,727

Notes: Estimation of Equation (5), showing the estimates for slope of the Phillips curve (ϕ) and the additional effect on the slope in tight labor markets (ψ). The deviations from our main specification (results in column (8) of Table 1.C) are as follows: Columns (1)-(3) vary the definitions of $1\{\text{tight}_{ct}\}$, i.e., the classification of time periods to normal or tight labor markets. Columns (4)-(6) replace $\hat{u}_{ct}$ with three alternative measures of labor market slack, namely the vacancy-unemployment ratio and the raw unemployment rate. $\dot{u}$ denotes the unemployment gap calculated with an HP filter with a smoothing parameter of 16,000. Column (7) drops all observations from 2020 onwards, (8) excludes observations where the country has followed an independent monetary policy, (9) weights all observations equally instead of by their share of overall euro area employment in the year 2000. (10) uses y/y instead of q/q services inflation as the dependent variable and (11) q/q headline inflation (using our own seasonal adjustment). Standard errors allow for arbitrary clustering at the time period level, except for column (12), where we allow for clustering at the time period and country level. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Figure C1: Estimates of ψ and ϕ with various thresholds for the labor market tightness regime

(a) Threshold on standardized unemployment gap (b) Threshold on unemployment gap

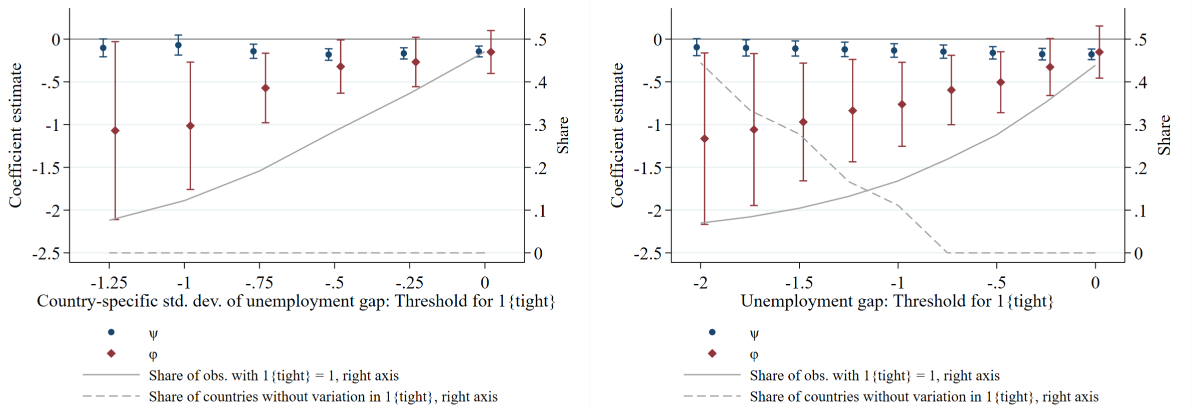

Notes: Estimates of Equation (5), depicted are the slope of the Phillips curve in normal times (ψ) and the additional effect on the slope when the labor market is defined as tight (ϕ). We show variations of the definition of tight labor markets on the x-axes. In panel (a), the underlying variable for this definition is the unemployment gap standardized at the country level, whereas it is the raw unemployment gap in Panel (b).

However, we see two potential threats to our measure of labor market slack. First, the natural rate is an estimate and therefore subject to uncertainty, which by extension also concerns $\hat{u}_{ct}$ and $1\{\text{tight}_{ct}\}$. This may lead to measurement error, biasing results towards zero. Second, our baseline measure of the natural rate is the non-accelerating wage rate of unemployment estimated by the European Commission. This estimation uses wage inflation to help extract the component of unemployment related to labor market slack (see Hristov et al., 2017 for methodological details), and might therefore bias our results towards finding a relationship between slack and unemployment. Notice, however, that their estimation is linear in labor market slack, so while our estimates of ψ might be biased towards finding a significant relationship, our ϕ would likely be biased towards zero.

To address these concerns, we first present results using the vacancy-unemployment ratio θ as an alternative measure of labor market slack in column (4). Barnichon and Shapiro (2024) show that in the United States, this measure has more predictive power for inflation than other measures of labor market conditions. Moreover, some related literature (for example, Gitti, 2024) has used MSA-level (log) tightness measures to estimate a regional Phillips curve. For the euro area, the data on vacancies is unbalanced, often not seasonally adjusted and exhibits strongly increasing trends over time. Therefore, we seasonally adjust and de-trend the data.[22] Because of these measurement issues, we prefer to use the unemployment gap as our baseline measure of labor market slack, but consider the vacancy-unemployment ratio a useful robustness check. Column (4) of Table C1 shows, first, that a higher ratio of vacancies to unemployment leads to higher inflation, as implied by the Phillips curve relationship. Second, this relationship

[22]We back out vacancies V from Eurostat's published data on vacancy rates ($v = V/(V+E)$) and employment (E), and seasonally adjust the data using X-13ARIMA-SEATS. Moreover, we de-trend the series, since it shows a clear upward trend in many countries (see Barlevy et al., 2024 for potential reasons).

is substantially stronger when this tightness measure is high, again defined as being at least one country-specific standard deviation above the trend.

In column (5), we address the potential issue with our baseline measure of the natural rate of unemployment by replacing it with a measure that is has no mechanical relationship with wage inflation, namely an HP filter with a smoothing parameter of 16,000. Note that even though the smoothing value is quite high, the resulting unemployment gaps have a lower amplitude than the European Commission estimates, with a standard deviation of 1.7 instead of 2.6. Reassuringly, results remain very similar to those in our baseline specification.

In column (6), we use the unemployment rate directly, which is also not subject to any of the measurement issues of the NAWRU mentioned above, as the measure of labor market slack. Consequently, we define the tightness indicator to be one if the unemployment rate is at least one country-specific standard deviation below the mean. This is a very challenging specification likely subject to substantial measurement error, given the sometimes strong nonlinear trends in the some countries' unemployment rates unrelated to labor market slack. Nevertheless, the point estimate of the interaction coefficient ϕ continues to be negative and sizable (-0.27), albeit only at the 10% significance level.

Subsamples. In columns (7) and (8), we adjust the estimation sample to rule out alternative explanations related to Covid-19 and the formation of the monetary union. Column (7) only considers data up to 2019, excluding both the brief disinflation during the Covid-19 pandemic and the subsequent inflation surge. This is important to rule out that special effects during this phase (i.e., supply shocks raising prices during a time when labor markets were tight) bias our results towards finding a stronger nonlinearity.

In column (8), we discard observations from EU accession countries during periods in which they technically had an independent monetary policy.[23] This is an important robustness check since the identification of ψ and ϕ relies on the absorption of the monetary regime in the common time fixed effects. One threat to this identification method, however, may be that in the early days of the monetary union and for countries which joined the euro zone later, expected long-term inflation might still have been on heterogeneous trends.[24]

Results in both columns show that our main findings are driven neither by contaminating effects from Covid-19 nor the harmonization of inflation expectations across member states in the early years of the euro area.

Other specifications. In columns (9)-(12), we adjust our baseline specification in various other ways to rule out that results are driven by specific choices on single properties. In column (9), we present the unweighted version of our baseline specification. In column (10), we measure services inflation in year-on-year terms, which does not require any seasonal adjustments, to

[23]In Figure C2, we separately drop each of the 18 countries in the sample. This shows that our result of a kink in the Phillips curve does rely on observations from a single country.

[24]To account for this possibility differently, in a separate exercise in column (1) of Table C2, we control for long-term inflation expectations by including the average 5y5y inflation forecast of professional forecasters surveyed by Consensus Economics.

Figure C2: Phillips curve estimates when dropping one country at a time

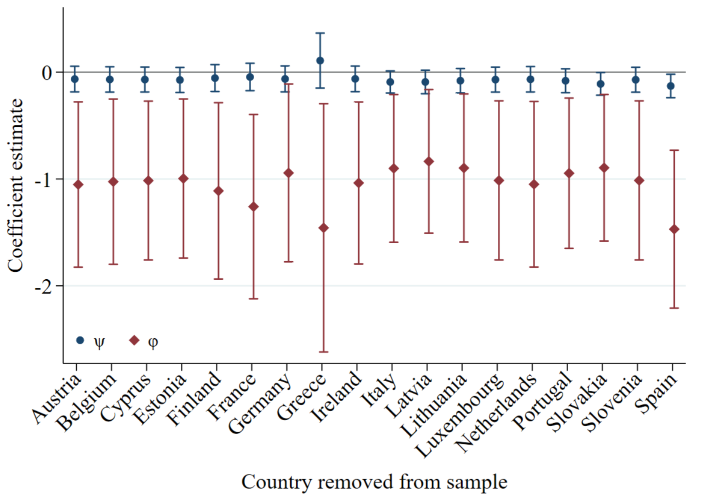

Notes: Estimates of Equation (5), depicted are the slope of the Phillips curve in normal times (ψ) and the additional effect on the slope when the labor market is defined as tight (ϕ). Each set of estimates is produced by dropping one country (out of 18) from the sample.

Table C2: Further robustness tests for nonlinear Phillips curve specification

	(1)	(2)	(3)	(4)	(5)
Changes rel. to main spec.:	Long-t. $E(\pi)$	Cons. wgt.	Winsor. π	$\hat{p}^S_{ct}$ not detr.	High infl.
Unemployment gap$_{ct}$	-0.05	-0.13**	-0.10*	-0.13***	-0.23***
	(0.06)	(0.06)	(0.05)	(0.05)	(0.04)
— × 1{tight$_{ct}$}	-0.82**	-0.65**	-0.87***	-0.71**	
	(0.31)	(0.31)	(0.32)	(0.29)	
— × 1{high inflation$_{ct}$}					0.13
					(0.44)
Country and time FE	✓	✓	✓	✓	✓
Observations	1,455	1,727	1,727	1,727	1,727

Notes: Estimation of Equation (5), showing the estimates for slope of the Phillips curve (ϕ) and, where applicable, the effect on the slope in tight labor markets (ψ). Relative to our baseline result, the specification changes as follows. In (1) we add 5y5y inflation expectations from the panel of professional forecasters surveyed by Consensus Economics. Column (2) uses country-level nominal consumption in 2000 as a share of the euro area total as weights instead of employment. Column (3) winsorizes inflation rates at the 5th and 95th percentile in each country to account for potential measurement error (potentially introduced by our own seasonal adjustment). Column (4) uses the raw ratio of services and goods prices as a control variable (instead of its deviation from a country-specific linear trend, as in the baseline). The sample consists of 18 countries and 92 time period (over 23 years). Column (5) includes the interaction with (and the main effect of) a dummy variable taking the value of one if a country's services inflation was more than one standard deviation above the country-specific mean in the respective time period. Standard errors are robust against heteroskedasticity and allow for arbitrary clustering at the time levels. Regressions are weighted by the country's share of overall euro area employment in the year 2000. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

make sure that our results are not driven by our own seasonal adjustment.[25] In column (11), we show that the nonlinearity is present even when using headline CPI inflation rates, which is used as the dependent variable in many papers in the regional Phillips curve estimation literature (e.g. Fitzgerald et al., 2024, Smith et al., 2024). In line with the reasoning in Hazell et al. (2022) that the regional Phillips curve should be estimated with a price measure that is driven predominantly by regional factors, estimates become closer to zero in this specification. Finally, in column (12), when calculating standard errors, we allow for arbitrary clustering at the time period *and* country level. While this increases the standard errors somewhat, results remain significant at the 5% level.[26]

Finally, in Table C2, we present further robustness checks, where we control for long-term inflation expectations, weight regressions by each country's share of euro area consumption instead of employment, winsorize inflation rates to account for potential measurement error, control for the non-detrended relative price of services instead of the detrended one, and include an interaction with a dummy for a "high inflation regime" instead of the one for "tight labor markets" in our baseline specification.[27]

[25] An alternative way to avoid potential measurement error in sequential services inflation is to winsorize the inflation rates. Results for this are included in Table C2.

[26] We prefer clustering only at the time period level in our baseline specification to avoid potentially biased standard errors resulting from the relatively small number of countries (18), as discussed in Cameron et al. (2008).

[27] The fact that the nonlinearity is not present in column (5) of Table C2 suggests to us that the nonlinearity is more likely to stem from downward nominal wage rigidity than a higher frequency of price adjustments during high-inflation periods.