Kinky Europe: Evidence from the regional Phillips curve in the euro area

Marius Faber, Gabriel Züllig

SNB Working Papers
12/2025

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK

ISSN 1660-7716 (printed version)
ISSN 1660-7724 (online version)

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

Marius Faber† Gabriel Züllig‡

August 19, 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 steeper when the labor market is sufficiently tight. This kink is more pronounced in the euro area than in the United States, potentially due to more rigid labor markets. Our estimates suggest, however, that despite this nonlinearity, most of the post-pandemic inflation surge is attributable to factors other than labor market tightness.

JEL classification: E30

Keywords: Phillips curve, inflation, nonlinearities

*We would like to thank Emmanuel De Veirman, Andreas Freitag, Jonathon Hazell, Christian Hepenstrick, Peter Karadi, Sarah Lein, Evi Pappa, Maxime Phillot, Raphael Schoenle, and Michael Weber for helpful discussions. We also thank seminar participants at the BIS-SNB Research Workshop, the SSES Annual Congress 2025, and the 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—a core element of macroeconomic models and a relationship 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, discussion has centered on whether the Phillips curve had flattened, which has led to new approaches for identifying its slope. Inflation appeared to have become less and less sensitive to changes in unemployment in the period between 1960 and 2020 in the United States, prompting economists to debate the potential flattening of the Phillips curve (Blanchard, 2016, Stock and Watson, 2020). It became clear from this literature that, in normal times, economic slack does impact inflation, but only to a relatively limited extent (Hazell et al., 2022).[1] 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.[2] For example, it became clear that using cross-sectional variation from a monetary union is useful to rule out endogeneity bias from aggregate shocks, monetary policy responses or changes in long-run inflation expectations.

Despite a seemingly flat Phillips curve, inflation rose far above central bank targets in the post-COVID period. Since the inflation surge was accompanied by tight labor markets in most advanced economies, the focus shifted toward understanding if the Phillips curve is nonlinear, that is, flat in normal times but steep in times of a booming economy with tight labor markets.[3] Notably, Phillips (1958) already had a nonlinear relationship in mind, asserting that the "relation between unemployment and the rate of change of wage rates is (...) likely to be highly non-linear" (p. 1). While euro area data offer clear suggestive evidence of a nonlinear Phillips curve, a split of the data into pre- and post-COVID samples highlights the difficulty of disentangling the role of supply shocks or changes in long-run inflation expectations from a nonlinear slope using 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 begin by justifying our empirical approach, laying out the main theoretical insights from Hazell et al. (2022) with regard to inferring 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,

[1] The finding that the Phillips curve has not, in fact, substantially flattened is confirmed by Bergholt et al. (forthcoming) using a complementary approach based on aggregate data.

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

[3] Another closely related debate concerns 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

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 from endogeneity stemming from reverse causality, contemporaneous regional supply shocks, and measurement error, we instrument the contemporaneous unemployment gap with the four-quarter-lagged unemployment gap.

The 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 that this is robust to several alternative specifications and address various concerns related to measurement and identification. For example, we vary the definition of a "tight" labor market and use multiple subsamples. Across all specifications, the Phillips curve is at least three times and up to fifteen times steeper in tight labor markets.

The kink that we document for the euro area is more pronounced than that reported for the United States, potentially due to the relatively high rigidity of labor markets in the euro area. In a heterogeneity exercise, we show that, within the euro area, countries with more rigid labor markets—measured by the degree of employment protection—have a more pronounced kink, indicating that the difference in labor market rigidity between the two currency areas could account for the difference in the shape of their Phillips curves. Finally, we consider the extent to which the post-COVID inflation surge has been driven by the labor market given the strong nonlinearity documented throughout this paper. Our estimates imply that, despite the nonlinearity, most of the inflation surge was due to factors other than labor market tightness.

This paper contributes to the literature on a potentially nonlinear Phillips curve, as it is being the first to provide solid empirical evidence of such a nonlinearity in the unique setting that is the euro area. While there are various theoretical explanations 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.[4] For the United States, Hooper et al. (2020) and Smith et al.

[4]Moretti et al. (2019) and Byrne and Zekaite (2020) investigate the potential nonlinearity of the euro area wage Phillips curve and reach different conclusions. However, both studies only use aggregate data; thus, they are subject to the various criticisms noted above and discussed in greater detail by Furlanetto and Lepetit (2024).

(2024) find evidence of a relatively steep Phillips curve at low levels of unemployment using cross-sectional variation from MSA-level data.[5] Gitti (2025) confirms these results, improving both the measurement of labor market tightness and the identification approach over previous studies. It should be 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 complete set of time fixed effects in our regressions. On Europe, the evidence is particularly scarce. While Smith et al. (2024) also include some tests for the nonlinearity in the European Union (EU, incl. the UK), their focus is on investigating structural breaks in the (linear) Phillips curve in the United States.

Our paper expands on this literature along several dimensions: First, in contrast to most of the above mentioned empirical research on regional Phillips curves, we focus on the euro area. The euro area lends itself naturally to estimating the Phillips curve with regional data. Given the region's relatively low level of labor-market integration (and low geographic mobility), countries in the euro area represent a reasonable proxy for local labor markets. This is critical to inferring the slope of the aggregate Phillips curve from regional estimates. Our results are therefore novel even in terms of linear Phillips curve estimates. Second, by focusing on the identification of kinks in the euro area Phillips curve, we expand substantially on the findings in Smith et al. (2024). In addition to the methodological advantages of our approach, which improve measurement and identification, we show that our main result is robust against various alternative specifications. We also explore the role of labor market rigidity as a mechanism for the nonlinearity, and quantify the extent to which labor market tightness has contributed to inflation in the euro area over time.[6]

The remainder of this 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 that we use to identify its slope(s). Section 4 presents our main results and Section 5 various robustness tests. Section 6 explores the role of labor market rigidity as it pertains to kinks in the Phillips curve. Section 7 quantifies the extent to which labor market tightness has historically contributed to euro area inflation given the documented nonlinearity in the Phillips curve. Section 8 concludes.

[5] There is also evidence of 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).

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

2 The regional Phillips curve from theory to empirics

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

2.1 Inferring the aggregate Phillips curve slope from regional data

To fix ideas, consider a linear New Keynesian Phillips curve 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:

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

Estimating κ—the slope of the Phillips curve—from aggregate data is subject to a number of econometric challenges. First, if unemployment is persistent—as it is very much in the euro area—shocks to $\hat{u}_t$ will have an impact on π_t for several time periods. We will explicitly distinguish between the contemporaneous (model-based) slope κ and a dynamic (empirical) counterpart. Second, inflation expectations, π_{t+1}, may be correlated with both current inflation and labor market slack. For example, a positive shock to inflation expectations may raise actual inflation through Equation (1) but simultaneously reduce real rates and labor market slack. Third, cost-push shocks may raise inflation and induce the central bank to raise policy rates, increasing labor market slack. In principle, these challenges could be alleviated if properly controlled for or through the use of valid instruments. However, it was shown that this solution is suboptimal, as results are sensitive to particular choices of instruments, control variables or time periods (Mavroeidis et al., 2014).

A recent literature proposes overcoming such endogeneity concerns through the use of a panel of regions in a monetary union, allowing for the inclusion of time fixed effects to absorb inflation expectations and aggregate shocks (McLeay and Tenreyro, 2019, Hazell et al., 2022, Fitzgerald et al., 2024). Hazell et al. (2022) develop a two-country New Keynesian model with separated labor markets,[7] showing that the structural parameter κ in Equation (1) can be inferred from the region-level (or, in our case, country-level) equation

$$\pi_{ct}^N = \kappa\hat{u}_{ct} + \beta E_t\pi_{c,t+1}^N + \lambda\hat{p}_{ct}^N + \nu_{ct}^N. \tag{2}$$

There are two differences between Equations (1) and (2). First, the dependent variable is inflation of non-tradables at the country level, π_{ct}^N. This is important, as prices of tradables are

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

uniform across the monetary union, meaning that they remain unaffected by regional slack. The second difference is the presence of the term $\hat{p}^N_{ct} = p^N_{ct}/p_{ct} - 1$, which controls for the deviation of relative prices of non-tradables from their steady state level of 1.

To bring Equation (2) to the data, first solve forward to obtain

$$\pi^N_{ct} = E_t \sum_{j=0}^{\infty} \beta^j \left(\kappa \tilde{u}_{c,t+j} + \lambda \hat{p}^N_{ct} \right) + E_t \pi^N_{c,t+\infty} + \omega^N_{ct}, \tag{3}$$

where we have introduced $\tilde{u}_{ct} = \hat{u}_{ct} - E_t \hat{u}_{c,t+\infty}$ as a decomposition of the unemployment gap into a transitory and a permanent 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 (2). $\omega_{ct} = E_t \sum_{j=0}^{\infty} \beta^j \nu_{c,t+j}$ is the discounted sum of all current and future cost-push shocks.

The Phillips curve is now expressed in terms of long-run inflation expectations. That means that we can replace them with time fixed effects γ_t, as these will, at least to a large extent, absorb non-observable long-run beliefs about inflation, which should co-move across countries in the same monetary union. These time fixed effects also absorb aggregate supply (and demand) shocks as well as the endogenous monetary policy response to them, which is, by definition, the same across countries in a monetary union. In addition to time fixed effects, we also include country fixed effects α_c to account for permanent differences in inflation (expectations) and labor market slack across countries.

Finally, we simplify Equation (3) by assuming that $\tilde{u}_{ct}$ and $\hat{p}^N_{ct}$ follow AR(1) processes with autocorrelation coefficients ρ_u and ρ_{pN}, respectively. This is a surprisingly accurate description of movements in the unemployment gap in the euro area, as discussed in Section 3.2. With these assumptions, Equation (3) can be rewritten as

$$\pi^N_{ct} = \psi \tilde{u}_{ct} + \delta \hat{p}^N_{ct} + \alpha_c + \gamma_t + \omega^N_{ct}, \tag{4}$$

where the estimable coefficients ψ and δ directly map to their structural counterparts as follows: $\psi = \frac{\kappa}{1-\beta\rho_u}$ and $\delta = \frac{\lambda}{1-\beta\rho_{pN}}$.

2.2 Incorporating potential nonlinearities

In the next and final step, we introduce potential nonlinearities into the Phillips curve. In other words, we let κ depend on labor market slack itself or any other state of the economy. The theoretical literature establishes various mechanisms that may endogenously give way 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, Schmitt-Grohé and Uribe, 2023 or Gitti, 2025). Where there is considerable slack in the labor market, there is little downward pressure on wages and thus prices. Conversely, wages in tight labor markets are more flexible, rising quickly with the utilization of labor, and are then (in part) 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 increase when demand is high (and labor markets are tight), as in Gasteiger and Grimaud (2023), Blanco et al. (2024), and Ascari et al. (2025). Finally, nonlinearities in the Phillips curve can also be micro-founded via quasi-kinked demand for goods, as in Kimball (1997), see e.g. Harding et al. (2022), or via firms' maximum capacity constraints, see e.g. Boehm and Pandalai-Nayar (2022).

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

$$\pi_t = \kappa \hat{u}_t + \tau 1\{\text{tight}_t\}\hat{u}_t + \beta E_t \pi_{t+1} + \nu_t, \tag{5}$$

where $1\{\text{tight}\}$ is a dummy variable equal to 1 when a labor market is tight and 0 otherwise. The slope of the structural Phillips curve is then equal to κ in normal times (i.e., when the labor market is not tight) and $\kappa + \tau$ when the labor market is tight.

Similar to the linear case, we can derive its equivalent which is estimable with panel data:

$$\pi_{ct}^N = \psi \tilde{u}_{ct} + \phi 1\{\text{tight}_{ct}\}\tilde{u}_{ct} + \delta \hat{p}_{ct}^N + \alpha_c + \gamma_t + \omega_{ct}^N, \tag{6}$$

with $\psi = \frac{\kappa}{1-\beta\rho_u}$, $\phi = \frac{\tau}{1-\beta\rho_u}$ and $\delta = \frac{\lambda}{1-\beta\rho_{pN}}$. This derivation requires two additional assumptions, which we validate in Section 3.2. First, the autoregressive coefficient ρ_u does not vary by tightness regime. We find that this assumption is satisfied in the euro area data. The second assumption is that, at each point in time, the labor market is expected to stay in its current tightness regime, that is, $E_t 1\{\text{tight}_{c,t+j}\} = E_t 1\{\text{tight}_{c,t}\}$. In this regard, we show that labor market slack is highly persistent in the euro area data (in fact, close to a unit root). Moreover, it is well known that agents over-extrapolate from current economic states and underestimate the degree of mean reversion (Fuster et al., 2010). Given these facts, we find that this assumption is also justifiable.

Notice that the slopes of the *contemporaneous* structural Phillips curve are κ and τ, whereas the coefficients that 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 the unemployment rate in subsequent periods on account 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 and the data employed to estimate Equation (6). We then discuss the main threats to identification in this setting and the measures taken to alleviate these.

3.1 Estimating equation

Our baseline regression specification is:

$$\pi_{ct}^S = \psi \hat{u}_{ct} + \phi 1\{\text{tight}_{ct}\}\hat{u}_{ct} + \eta 1\{\text{tight}_{ct}\} + \delta \hat{p}_{c,t-1}^S + \alpha_c + \gamma_t + \epsilon_{ct}, \tag{7}$$

where π_{ct}^S is the 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. Note that much of the literature on regional Phillips curves uses headline inflation; this leads to an underestimation of the slope of the aggregate Phillips curve, as it contains tradable goods that are insensitive to slack in the local labor market. $\hat{u}_{ct}$ is the unemployment gap, $1\{\text{tight}_{ct}\}$ is the dummy for tight labor markets, $\hat{p}_{c,t-1}^S$ is the relative price of services (lagged by one quarter to avoid a mechanical correlation), and α_c and γ_t are country- and period (quarter-year) fixed effects.[8] In our preferred specification, we estimate this equation with two-stage least squares (2SLS), using the four-quarter lags of the unemployment gap and labor market tightness as instruments. We weight the regressions by the country's share of euro area employment in the year 2000. Standard errors allow for arbitrary clustering at the time period level.

Our empirical strategy thus exploits differences in the relationship between inflation and labor market slack among countries with tight and slack labor markets at the same point in time and accounting for permanent differences in (observable and unobservable) country characteristics. For example, the estimation exploits the fact that in before the global financial crisis, some countries (e.g., France, Italy and Spain) had tight labor markets while others did not (e.g., Germany and the Netherlands). It compares the relationship between inflation and labor market slack across these sets of countries, while recognizing the permanent differences among them.

[8]Our empirical specification has two minor deviations from the structural equation (6). First, we use the raw measure of the unemployment gap rather than extracting its transitory component, because the permanent component is absorbed in fixed effects. Additionally, the dummy representing tight labor markets is included both in the interaction with the unemployment gap and as a separate control variable (with coefficient η), as is standard in empirical work.

3.2 Data and descriptive statistics

We draw on multiple data sources to estimate Equation (7) at a quarterly frequency for the period between 2001Q1 and 2024Q4 for a panel of 18 euro area countries.[9] We use the "HICP: Services" indices from Eurostat as a proxy for non-tradable prices. As these data 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 represents a deviation from much of the literature, which typically employs 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 is likely to be underestimated. Thus, we consider sequential growth rates to be 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), interpolating the latter from an annual to a quarterly frequency. Detrending the unemployment rates is important in the case of the euro area because the "raw" unemployment rates exhibit highly heterogeneous and nonlinear time trends. These do not only reflect economic slack, but also stem from supply-side factors, such as demographic trends, migration patterns, and changes in labor market policy, one example being the Hartz reforms passed in Germany in 2005. Failing to account for such trends in u^n_{ct} would bias our results. We go on to establish robustness to the measurement of labor market slack in Section 5.[10]

The mapping between estimated and structural parameters introduced in the previous section assume *i)* that the unemployment gap follows an AR(1) process and *ii)* that the parameter of this process is not dependent on the tightness regime. We present estimates of the first autoregressive coefficient in Table A2, using both the euro area time series and the panel with all 18 countries. In both cases, ρ_u is estimated to be around 0.98—very close to a unit root—with an R^2 of this simple AR(1) of 0.97. Table A2 also shows the very small (and statistically insignificant) estimates of an interaction term between lagged unemployment gaps and tightness indicators. We conclude that the two additional assumptions necessary for a simple mapping between our estimated coefficients and the structural parameters are likely to hold for the euro area.

We define labor markets as "tight" in all quarters for which the 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. We choose a

[9] These include all 20 current euro area members, except for Croatia, which only joined in 2023, and Malta, for which the data are incomplete. Notably, we include the newer member countries that have been part of the European Exchange Rate Mechanism (with their currencies pegged to the euro) since the beginning of our sample period.

[10] See Figure A1 and Table A1 for descriptive statistics, including the average levels and standard deviations of the unemployment gaps across euro area members.

negative cutoff value because euro area labor markets have generally been loose over the sample period; however, we establish the robustness of our results to the choice of this value in Section 5.

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”, subtracting a linear country-level trend.[11]

3.3 Identification

Various identification issues arise when estimating the effect of economic slack on inflation using the contemporaneous correlation between the two variables. First, aggregate 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 detailed in Section 2 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, labor market institutions, or norms) may 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.

Finally, even with period- and country fixed effects, a few sources of bias remain. We address these by using the four-quarter lags of the unemployment gap and labor market tightness indicator as instruments and estimating Equation (7) with 2SLS. OLS estimates may be biased due to reverse causality, with inflation affecting labor market slack, in addition to the latter affecting the former. Using lagged variables as instruments assuages this concern. Moreover, 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, these 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 the extent that the impact of such regional supply shocks on quarter-on-quarter services inflation and labor market tightness vanishes after about one year, using these instruments should ease concerns about simultaneity. Lastly, given the measurement error inherent in an unobservable object like the unemployment gap, using lagged variables as instruments should reduce any bias towards zero resulting from that.

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

4 Results

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

We start by estimating the linear Phillips curve, excluding the interaction with the indicator of labor market tightness in Equation (7), as shown in Columns (1)–(4) of Table 1. Panel A presents OLS estimates, Panel B presents those from a reduced-form estimation in which the instruments are used directly as the explanatory variables, and Panel C presents those from the 2SLS estimation, in which the contemporaneous unemployment gap and labor market tightness 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) presents the estimates from a version without any fixed effects. These coefficients may pick up observable or unobservable differences in country characteristics (especially expected inflation) or aggregate supply shocks that affect labor markets and inflation. Columns (2) and (3) include country- and time fixed effects, respectively, while Column (4) includes them jointly. Notably, the point estimate becomes only somewhat smaller (in absolute terms) upon 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.

Our 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 toward 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 1 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 across the sample, translatin—according to this linear estimate—to an effect on services inflation of less than 0.3 percentage points, annualized. This is approximately one-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 nonetheless 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.[12]

[12] $\psi = -0.22$ 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 (see Table A2), we estimate a κ of -0.0052, which is very similar to the direct estimates of κ provided by Hazell et al. (2022) for the United States.

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 (7), showing the estimates for the 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 1 in periods where the difference between a country's unemployment rate and the natural rate of unemployment is at least one standard deviation below its 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 at the time level. 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%.

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 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.

Across 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.[13] 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 both country and time fixed effects (Column (8) of Table 1.A), ϕ is estimated to be -0.34. This points to a point estimate of the slope of the Phillips curve of -0.53 when labor markets are tight, which is almost three times the slope in normal times. In the 2SLS specification in Panel C, this wedge increases even further, pointing to a Phillips curve that is nearly 15 times as steep in tight labor markets as it is in normal times (with $\psi = -0.07$ and $\phi = -1.01$).[14] The estimate of ϕ does, however, have a relatively large standard error, implying a 95% confidence interval of -0.5 to -1.4.

Under the conditions outlined in Section 2, our estimates of ψ and ϕ directly map to the structural parameters κ and τ. Given that the autocorrelation of the unemployment gap ρ_u barely differs across labor market regimes (see Table A2), the nonlinearity implied by the empirical estimates entails a corresponding nonlinearity of the structural Phillips curve.

5 Robustness

In this section, we conduct several robustness checks to assess the sensitivity of our results to various important choices made in our baseline specification. Among other things, we vary the definition of tight labor markets, use alternative measures of labor market slack, and estimate coefficients based on different subsamples.

Definition of tight labor markets. In our baseline result, the $1\{\text{tight}_{ct}\}$ dummy is equal to 1 when the unemployment gap is at least a country-specific standard deviation below the long-term average. In Figure 2(a), we show estimates for different threshold values. The estimate of ϕ is negative regardless of the chosen value, and statistically significantly so for values -0.5 or below. The estimates for a value of -0.5 are also laid out in Column (1) of Table 2.

Additionally, we estimate a version of the nonlinear Phillips curve with a smooth convexity

[13]The 2SLS results are unlikely to suffer from substantial bias stemming 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 the 5% significance level (see Stock and Yogo, 2002).

[14]Figure A2 depicts the steepening implied by our preferred estimates.

Figure 2: Estimates of ψ and ϕ with various thresholds for the labor market tightness regime

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

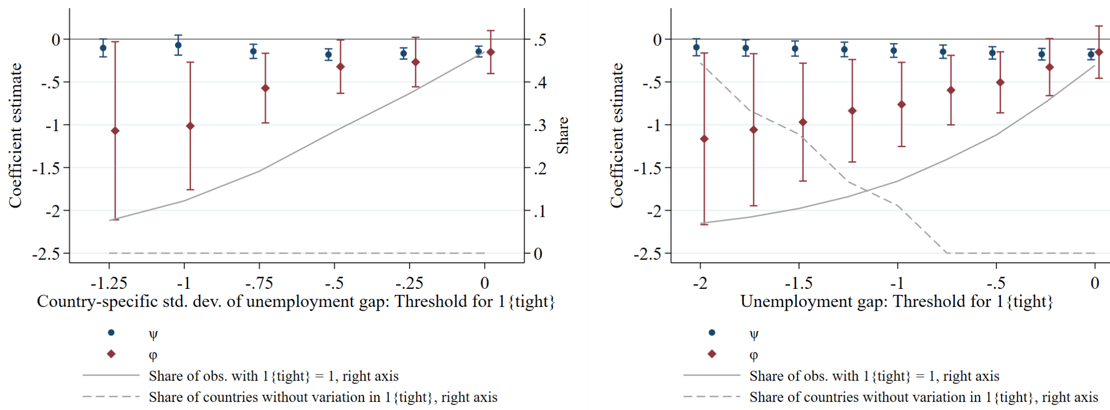

Notes: The depicted estimates of Equation (7) are the slope of the Phillips curve in normal times (ψ) and the additional effect on the slope under a labor market defined as tight (ϕ). We show variations in the definition of a tight labor market on the x-axes. In Panel (a), the underlying variable for this definition is the unemployment gap standardized at the country level; Panel (b) is based on the raw unemployment gap.

rather than a piece-wise linear shape. Figure A2 depicts the estimate of a restricted cubic spline with three knots at a standardized unemployment gap of 0, -1 and -2. This flexibly allows for (potentially non-monotonic) changes in the slope of the Phillips curve at any point between 0 and -2 standard deviations while imposing linearity below and above. The resulting shape confirms that the steepening occurs within a relatively narrow interval around -1; thus, our main version with a piece-wise linear slope around a kink is a good approximation of this shape.

The next set of robustness tests, laid out in Columns (2) and (3) of Table 2, again use piece-wise linearity but allow $1\{\text{tight}_{ct}\}$ to be a function of the non-standardized unemployment gap. While we think that country-level standardization sensibly accounts for different variances of unemployment rates across member countries, this alternative specification may be easier to interpret, as it is measured directly in percentage point deviations. A disadvantage of this alternative is that it does not necessarily use variation from all countries, as some with a particularly stable unemployment gap did not experience a tight labor market when using the lower-threshold definition. Results for a wide range of these thresholds are presented in Figure 2(b). The results confirm the view that the Phillips curve steepens considerably when the unemployment gap is -0.5 or lower, that is, when labor markets are sufficiently tight.

Measurement of slack. Our baseline measure of labor market slack is the unemployment gap. Subtracting the natural rate from the unemployment rate is important, as failing to account for trends in the natural rate can lead to the underestimation of the slope of the Phillips curve (Hall and Kudlyak, 2023). This is particularly relevant in the context of the euro area, where structural forces (e.g., demographic trends, changes in labor market policy) have led to heterogeneous and nonlinear trends in unemployment rates that cannot be absorbed by time- or country fixed effects.

Nevertheless, we see two potential threats to our measure of labor market slack. First, the natural rate is an estimate and is, therefore, subject to uncertainty; this, by extension, also concerns $\hat{u}_{ct}$ and $1\{\text{tight}_{ct}\}$. This could lead to measurement error, biasing the results toward zero. Second, our baseline measure of the natural rate is the NAWRU estimated by the European Commission using wage inflation to help extract the component of unemployment related to labor market slack (see Hristov et al., 2017 for methodological details). This might bias our results toward finding a relationship between slack and unemployment. Notice, however, that the European Commission's estimation is linear in labor market slack. Therefore, while our estimates of ψ might be biased toward finding a significant relationship, our ϕ would likely be biased toward zero, i.e., finding no nonlinearity.

To address these concerns, in Column (4) of Table 2, we first present results using the vacancy–unemployment ratio θ as an alternative measure of labor market slack. Barnichon and Shapiro (2024) show that, in the United States, this has more predictive power for inflation than other measures of labor market conditions. Moreover, related studies (e.g., Gitti, 2025) use MSA-level (log) tightness measures to estimate a regional Phillips curve. The data on vacancies for the euro area are unbalanced, often not seasonally adjusted, and exhibit strongly increasing trends over time.[15] Due to these measurement issues, we prefer to use the unemployment gap as our baseline measure of labor market slack, but we consider using the vacancy–unemployment ratio as a useful robustness check. Column (4) of Table 2 shows, first, that a higher ratio of vacancies to unemployment leads to higher inflation, as implied by the Phillips curve relationship. Second, this relationship is substantially stronger when this measure of tightness is high (i.e., 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 has no mechanical relationship with wage inflation: an HP filter with a smoothing parameter of 16,000. Note that even though the smoothing value is relatively high, the resulting unemployment gaps have a lower amplitude than the European Commission estimates, with a standard deviation of 1.7 (compared to 2.6). Reassuringly, these results are very similar to those of our baseline specification.

In Column (6), we use the unemployment rate directly to avoid relying on any particular estimate of the NAWRU. This is a very challenging specification, as it effectively only allows for permanent cross-country differences in the level of the NAWRU (absorbed in country fixed effects) and for a common euro area time trend (in time fixed effects). Given the sometimes-strong nonlinear trends in unemployment rates unrelated to labor market slack, this likely leads to a substantial measurement error in the unemployment gap. Nevertheless, the point estimate of the interaction coefficient ϕ remains negative and sizable (-0.27), but is only statistically significant at the 10% level.

[15] 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 exhibits a clear upward trend in many countries (see Barlevy et al., 2024 for potential reasons).

Table 2: Main robustness tests for the 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 (7), 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 (Column (8) of Table 1.C) are as follows: Columns (1)–(3) vary the definitions of $1\{\text{tight}_{ct}\}$, meaning the classification of periods as having a normal or tight labor market. 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 onward, (8) excludes observations in which the country has followed an independent monetary policy, (9) weights all observations equally instead of doing so based on their share of overall euro area employment in the year 2000. (10) uses year-on-year instead of quarter-on-quarter services inflation as the dependent variable, while (11) uses quarter-on-quarter headline inflation (with our own seasonal adjustment). Standard errors allow for arbitrary clustering at the time period level, except for Column (12), which allows for clustering at the time period and country levels. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Subsamples. Column (7) of Table 2 shows the results of our baseline specification but only with data up to 2019, meaning that the brief period of disinflation during the COVID-19 pandemic and the subsequent surge in inflation are excluded. The results are very similar and, if anything, stronger, which rules out the idea 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 exclude observations from EU accession countries in periods in which they technically had an independent monetary policy. This is an important robustness check, as the identification of ψ and ϕ relies on the absorption of the monetary regime in the common time fixed effects. A possible threat to this identification method may be that, in the early days of the monetary union and for countries that joined the euro area later, long-term inflation expectations might still have been exhibiting heterogeneous trends. To account for this possibility differently, in a separate exercise in Column (1) of Table A3 in the Appendix, we control for long-term inflation expectations by including the average 5y5y inflation forecast of professional forecasters surveyed by Consensus Economics. In both cases, the estimate of ϕ becomes slightly smaller but remains statistically significantly negative and large compared to the estimate of ψ.[16]

Other specifications. Finally, in Columns (9)–(12) of Table 2, we make other adjustments to our baseline specification to rule out that our results are driven by specific choices on single properties. In Column (9), we present the unweighted version of our baseline specification.[17] In Column (10), we measure services inflation in year-on-year terms, which does not require any seasonal adjustments, to ensure that the results are not driven by our own seasonal adjustment. In Column (11), we demonstrate that the nonlinearity persists even when using headline CPI inflation rates, which matches specifications of other contributions 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 using a price measure driven predominantly by regional factors and would be biased toward zero otherwise, estimates become smaller and less significant in this specification. However, the estimate of ϕ remains statistically significant at the 10% level. Finally, in Column (12), when calculating standard errors, we allow for arbitrary clustering, including serial correlation, at the time period *and* country level. While this increases the standard errors somewhat, ϕ remains negative at the 5% level of statistical significance.[18]

Overall, we conclude that there is robust evidence of the Phillips curve in the euro area becoming substantially steeper when labor markets are sufficiently tight.

[16] In Appendix Figure A3, we separately drop each of the 18 countries in the sample to show that our finding of a kink in the Phillips curve does not rely on observations from a single country.

[17] Appendix Table A3 also includes a version with consumption rather than employment weights.

[18] We prefer clustering only at the time period level in our baseline specification, given that there are only 18 units (countries) in our data. As discussed in Cameron et al. (2008) , there is a considerable risk with two-way clustering of biased standard errors in such a case.

6 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 those 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 (2025) documents a slope that is almost three times as large in tight labor markets. Beaudry et al. (2025) show that, depending on the specification, the nonlinearity nay vanish entirely. In the results presented in the previous two sections, the size of this wedge is around 5, on average, across the main results and robustness tests, and it 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 (as documented in Schoefer, 2025). The main idea is that in rigid labor markets, labor supply is not as flexible, such that when labor markets are tight, firms need to offer workers larger wage increases for them to switch occupations, industries, locations, or enter the labor market. For example, in labor markets with stricter job protection, job turnover is lower, resulting in fewer positions opening up at any point in time. In such an environment, it is riskier for workers to accept new jobs, as they know that there will be fewer vacancies to choose from should the accepted job turn out to be a bad fit. As a result, labor supply is less flexible, and firms need to offer larger wage increases to fill vacancies. Similarly, in countries with more generous unemployment benefits, 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. In principle, these mechanisms may steepen the Phillips curve even when labor markets are slack. However, it is likely that, in such times, downward nominal wage rigidity substantially flattens the Phillips curve amid both rigid and flexible labor markets, implying a more pronounced kink (as opposed to a steeper Phillips curve more generally) in regions with rigid labor markets, such as the euro area.

To test this hypothesis, in Table 3, we exploit the euro area's country-level variation in labor market rigidity. In addition to the interaction of labor market slack with the tightness indicator $1\{\text{tight}_{ct}\}$, we include an interaction with an indicator of labor market rigidity, $1\{\text{rigid}_c\}$. Labor market rigidity is difficult to measure and compare across countries. As a first pass, we compare countries with traditionally high versus low average unemployment rates, using this as a coarse proxy for rigidity. Furthermore, we explore two more specific dimensions of labor market rigidity limiting both unemployment-employment and job-to-job transitions, namely *i)* workers' opportunity cost of being unemployed and *ii)* firms' difficulties in laying off workers. To disentangle the roles played by these two dimensions, we use as additional indicators for rigid labor markets *i)* unemployment benefits as a share of previous income (OECD, 2023a) and *ii)* the strictness of employment protections legislation (OECD, 2023b).[19]

Results are presented in Table 3, and our coefficient of interest in this exercise is the one on the

[19] Classifications of countries based on all three measures are listed in the summary statistics in Appendix Table A1. While there is some overlap between these measures of rigidity, there are also notable differences.

Table 3: Phillips curve estimates with labor market rigidities

	(1)	(2)	(3)
	Unemployment rate (avg.)	Unemployment benefits	Employment protection
Unemployment gap$_{c,t-4}$	-0.40***	-0.22***	-0.10**
	(0.07)	(0.02)	(0.04)
— × 1{rigid$_c$}	0.25***	0.12***	-0.15***
	(0.07)	(0.04)	(0.04)
— × 1{tight$_{c,t-4}$}	0.35	-0.48***	-0.47***
	(0.24)	(0.15)	(0.10)
— × 1{tight$_{c,t-4}$} × 1{rigid$_c$}	-0.41**	0.19*	-0.40***
	(0.18)	(0.10)	(0.09)
Country and time fixed effects	✓	✓	✓
Observations	1,727	1,727	1,727

Notes: Estimation of Equation (7), including interaction terms with three different proxies for labor market rigidity: Column (1) uses the country-level average unemployment rate; Column (2) uses unemployment benefits as a share of previous income (from the OECD); Column (3) uses the employment protection index (from the OECD). The sample consists of 18 countries—half of which are classified as rigid—and 96 periods (spanning over 24 years). Apart from the additional interaction terms (and the main effect of labor market rigidity), the specifications are identical tose ones in Panel B and Column (8) of Table 1). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

triple interaction in the last row.[20] Column (1) shows that, when classifying labor markets as rigid based on the average unemployment rate, the steepening of the Phillips curve appears to be driven by the set of countries with relatively rigid labor markets. Beyond this rather crude measure, Column (2) suggests that the generosity of unemployment benefits is not the relevant dimension driving this heterogeneity, as the coefficient on the triple interaction becomes positive (albeit barely statistically significantly different from zero). Instead, Column (3) suggests that the strictness of employment protection is a relevant dimension driving the steepening of the Phillips curve. This suggests that the additional inflationary pressure in tight labor markets that results from more rigid institutions stems rather from the difficulty of inducing employed workers to switch, rather than inducing unemployed workers to become employed.

These results support our conjecture that a large share of the very steep Phillips curve in tight labor markets, which we estimate for the euro area, is driven by their relatively rigid labor markets. In these markets, demand shocks are less easily met with additional labor supply, resulting in stronger inflation pressure. As euro area labor markets are generally more rigid than those in the United States, this finding may explain why our results show a more pronounced kink for the euro area Phillips curve than what the literature has found for the United States. Our results also suggest that, in this regard, the more relevant institutional

[20]In this table, we focus on reduced form specifications, as they allow us to identify potential heterogeneity without worrying about differences in the strength of the first-stage relationships driving the results.

difference between the euro area and the United States lies in the former's stricter employment protection legislation rather than its more generous unemployment benefits.

7 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 quantify the extent to which the strong kink we find in the euro area Phillips curve can explain historical variation in (services) inflation. We are particularly interested in quantifying the contribution of the labor market to services inflation amid the post-COVID inflation surge.

According to the European Commission's 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 lies 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, which is more than two orders of magnitude lower than the actual rate of services inflation at the time.

However, there is considerable uncertainty regarding 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 were underestimated and thus the unemployment gap overestimated, the contribution of the labor market to inflation could be significantly higher, due to 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.[21] Figure 3(a) shows the distribution of the resulting unemployment gaps over time. Over the last two years of the time series, the 90% interval of plausible values for 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 or 2023 in which the unemployment gap is at least one standard deviation below its historical mean, meaning that the nonlinear component of the Phillips curve—$\hat{\phi} \cdot \hat{u}_t$—added 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 applying them to each path of the unemployment gap. The resulting 100,000 draws of the contribution of the labor market to services inflation are depicted in Figure 3(b), transforming the quarter-on-quarter

[21] First, they follow a random walk, just like the mean estimate measured by the European 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.

contributions back to year-on-year rates for the sake of readability.

Figure 3: Historical contribution of the Phillips curve to euro area services inflation

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

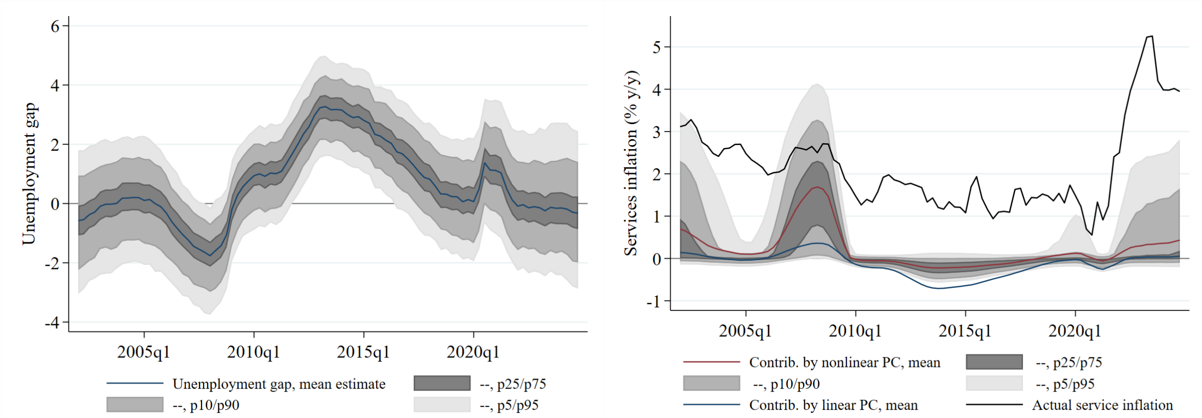

Notes: Panel (a) shows the mean estimate of the unemployment gap published by the European Commission; Gray areas illustrate the distribution of 1,000 simulations of the unemployment gap accounting for the uncertainty around the mean estimate, which we take to be 1 percentage point. Under this and parameter uncertainty in the Phillips curve estimation, Panel (b) depicts the contribution of labor market slack to historical services inflation in the euro area, that is, $\hat{\psi}^{(d)}\hat{u}_t^{(d)} + \hat{\phi}^{(d)}\hat{u}_t^{(d)}1\{\text{tight}_t^{(d)}\}$ 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, as indicated by the darkest shade in Panel (b), which shows 50% of all resulting contributions of the unemployment gap to services inflation. However, the distribution of the contributions in 2022 and 2023 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.

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 fairly small. Instead, these calculations suggest that the inflation surge in the euro area was predominantly due to factors other than labor market tightness, namely supply shocks (such as supply bottlenecks and surging natural gas prices) and their second-round effects on service prices through a catch-up of nominal wages). We confirm this by showing estimates of the time fixed effects in Figure A4 in the Appendix. They imply a contribution to the inflation surge of at least 3 percentage points.

The contributions based on our nonlinear Phillips curve estimates also track the historical dynamics of services inflation significantly better than the linear slope. For example, the effects of the labor market cooling in the early 2000s, the tightening in the lead-up 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 rates pf services inflation. This suggests that, despite its limited role in the post-COVID inflation surge, labor market slack and the nonlinear Phillips curve were significant drivers of inflation fluctuations in the first 20 years of the euro area.

8 Conclusion

This paper provides robust evidence of a nonlinear Phillips curve in the euro area. In normal times, the relationship between labor market slack and inflation is likely present but quite weak. In contrast, when the labor market runs hot, that is, 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 in the steepening of the Phillips curve in tight labor markets, with the steepening being particularly pronounced in more rigid labor markets. This may explain our finding that steepening is more pronounced in the euro area than it is, according to recent studies, in 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 and 2023. Thus, the clear steepening of the negative correlation between inflation and the unemployment rate depicted in Figure 1 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

Ascari, Guido, Alexandre Carrier, Emanuel Gasteiger, Alex Grimaud, and Gauthier Vermandel, "Monetary Policy in the Euro Area, when Phillipe Curves... Are Curves," *CEPR Discussion Paper*, 2025, *DP20489.*

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.

Bergholt, Drago, Francesco Furlanetto, and Etienne Vaccaro-Grange, "Did Monetary Policy Kill the Phillips Curve? Some Simple Arithmetics," *Review of Economics and Statistics*, forthcoming.

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.*

Boehm, Christoph E. and Nitya Pandalai-Nayar, "Convex Supply Curves," *American Economic Review*, 2022, *112* (12), 3941–3969.

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.

Fuster, Andreas, David Laibson, and Brock Mendel, "Natural Expectations and Macroeconomic Fluctuations," *Journal of Economic Perspectives*, 2010, *24* (4), 67–84.

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," 2025.

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.

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).

OECD, "Benefits in unemployment, share of previous income," `https://www.oecd.org/en/data/indicators/benefits-in-unemployment-share-of-previous-income.html` 2023. Accessed: 2025-07-25.

_ , "OECD Employment Protection Legislation (EPL) Index Dataset," `https://www.oecd.org/employment/emp/oecdindicatorsofemploymentprotection.html` 2023. Accessed: 2025-07-25.

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.

Schmitt-Grohé, Stephanie and Martin Uribe, "Heterogeneous Downward Nominal Wage Rigidity: Foundations of a Nonlinear Phillips Curve," *NBER Working Paper*, 2023, *30774.*

Schoefer, Benjamin, "Eurosclerosis at 40: Labor Market Institutions, Dynamism, and European Competitiveness," *NBER Working Paper*, 2025, *33975.*

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.

Appendix

A Additional tables and figures

Figure A1: Labor market indicators for the euro area, 2001–2024

(a) Unemployment rate

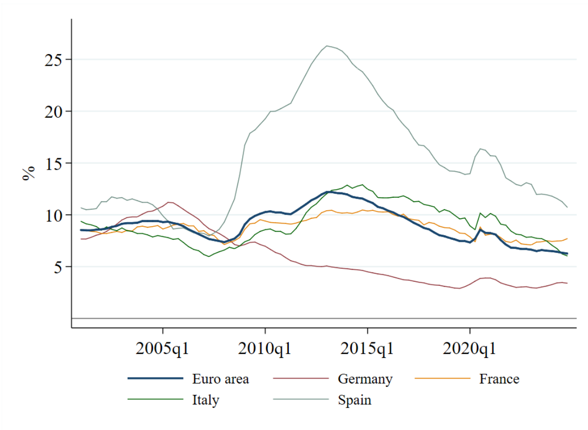

(b) NAWRU

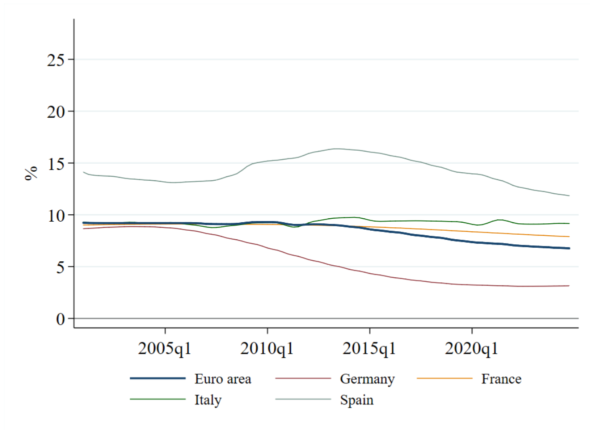

(c) Unemployment gap

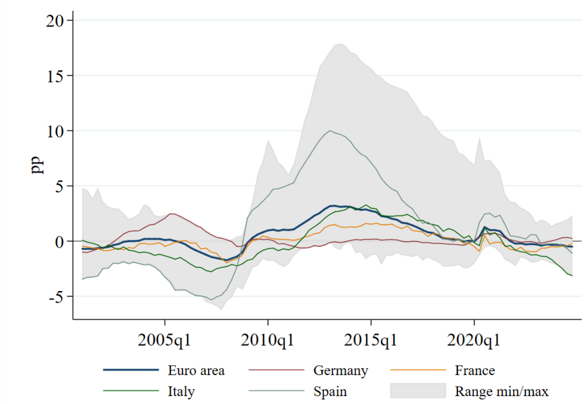

(d) Share of countries with tight labor market

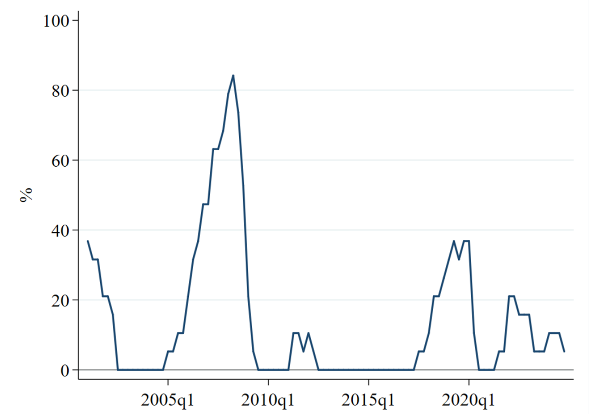

Notes: Panel (a) presents unemployment rates for the euro area aggregate as well as the four largest member states. Panel (b) shows the European Commission's estimates of the NAWRU. Panel (c) features the difference between (a) and (b), our main measure of labor market slack. Detrending is important here due to the heterogeneous and nonlinear trends in the unemployment rates, as shown in (b), which are not driven by labor market slack, but by labor market policies and demographic trends. The gray area encompasses the range of unemployment gaps across the 18 countries in the sample. Panel (d) shows the share of these countries that have a standardized (at the country level) unemployment gap of -1 or below, which serves as our baseline definition of a tight labor market.

Table A1: 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})	Mean(u)	UB	EPL
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	2.63	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	.	.	.

Notes: All numbers in percent. The first column contains the country weights we use in our regression, which are based on the share of each country's employment in 2000. The last three columns contain time-invariant dummies at the country level describing whether the labor market is more rigid than the median. To measure labor market rigidity, we use the average unemployment rate over the 24-year sample, unemployment benefits as a share of previous income ("UB", OECD, 2023a) and the employment protection legislation index ("EPL", OECD, 2023b).

Table A2: Autocorrelation of unemployment gaps

	(1)	(2)	(3)	(4)
	Euro area level $\hat{u}_t$		Country level $\hat{u}_{ct}$	
$\hat{u}_{(c),t-1}$	0.98***	0.99***	0.97***	0.98***
	(0.02)	(0.02)	(0.01)	(0.01)
$\hat{u}_{(c),t-1} \times 1\{\text{tight}_{(c),t-1}\}$		-0.04		-0.03
		(0.06)		(0.02)
Country FE and time FE			✓	✓
R^2	0.97	0.97	0.97	0.97
Observations	96	96	1,727	1,727

Notes: We regress quarterly unemployment gaps on their lags. Columns (3)–(4) do so for the set of 18 euro area countries in our data. Depicted are standard errors clustered by time period.

Figure A2: Estimates of the Phillips curve with various functional forms

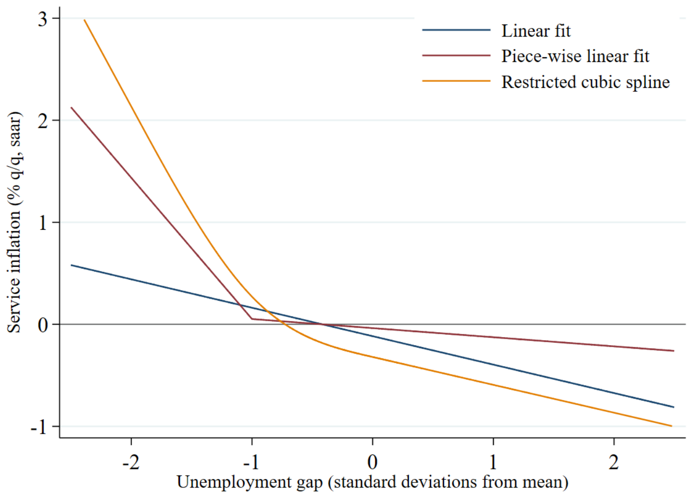

Notes: The orange line depicts the Phillips curve when estimated with a restricted cubic spline with three knots at -2, -1, and 0. Additionally, we show estimates of the Phillips curve with a linear slope of $\psi = -0.22$ (in blue) and the nonlinear version (in red) with $\psi = -0.07$ and $\phi = -1.01$. Both of them are estimated using Equation (7), which is a function in the unemployment gap.

Recent SNB Working Papers

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK