Labor Market Shortages and the Phillips Curve*

Cristian Arcidiacono†

March 11, 2026

Abstract

This paper introduces a *granular* measure of labor market slack to estimate the slope of the Phillips curve. The measure is based on business surveys of perceived labor shortages and is available at the sub-sector-country level. This granularity enables panel estimation with interactive fixed effects, which flexibly controls for unobserved common factors related to inflation expectations and supply-side determinants. Crucially, common components are allowed to impact differently both sub-sectors and countries. Estimates for the euro area result in a significantly steeper Phillips curve than those obtained using conventional measures and standard empirical specifications.

Keywords: Phillips Curve; Inflation; Labor Markets; Non-Linearities; Monetary Policy
JEL Codes: E24, E31, J64

*I am deeply indebted to Pierpaolo Benigno and Ricardo Reis for their support all along. I thank Matthias Gnewuch, Douglas Araujo, Dante Gaviano and various seminar participants for helpful comments and discussions. I also thank Gabriele Iannizzi for insightful discussions and for his help at the early stage of the project.

†University of Bern, cristian.arcidiacono@unibe.ch.

1 Introduction

Ever since the influential work by Phillips (1958), economists have tied low unemployment to wage inflation. A textbook version of the New Keynesian Phillips Curve (NKPC) can be written as

$$\pi_t = \alpha \hat{x}_t + \theta \pi_t^e + v_t + \epsilon_t \tag{1}$$

where π_t is inflation, $\hat{x}_t$ is the relevant measure of economic activity (slack), π_t^e captures inflation expectations, v_t represents cost-push shocks and ϵ_t is the error term.

Although the original contribution of Phillips (1958) dates back nearly seventy years, consistent empirical estimation of the Phillips curve remains elusive. The literature has converged on two fundamental obstacles. First, commonly used indicators of economic slack – such as unemployment gaps or output gaps – are noisy, weakly informative and often poorly aligned with underlying labor market conditions. Second, inflation expectations and cost-push shocks are only imperfectly observed, giving rise to omitted variable bias that may severely distort estimates of the Phillips curve slope.

These limitations have important implications, as the slope of the Phillips curve is key for monetary policy, governing the transmission from real activity to inflation and shaping central banks' assessments of inflationary pressures. When slack is mismeasured and both inflation expectations and supply-side forces are inadequately controlled for, the empirical Phillips curve appears highly sensitive to specification choices or flat, undermining its usefulness for policy analysis.

This paper makes two contributions. First, we introduce a *granular* measure of labor market tightness that serves as a clearer proxy for slack than conventional aggregate indicators. Second, we apply an interactive fixed effects regression that enforces a low-rank factor structure on unobserved components, flexibly capturing inflation expectations and cost-push shocks without depending on strong parametric assumptions. Crucially, by allowing the impact of these common factors to vary heterogeneously across sectors and countries, we show that conventional estimation approaches implicitly restrict how common shocks propagate across units and lead to systematically biased estimates of the Phillips curve slope.

Traditional indicators of slack, such as the unemployment rate or the output gap, have long been employed in empirical Phillips curve estimations as proxies for firms' real marginal costs, in line with textbook New Keynesian models (e.g., Galí, 2015). Yet they suffer from

important limitations, as they often exclude workers outside the labor force or focus only on average labor costs, which may not coincide with the marginal cost of labor[1].

The first contribution of this paper is to introduce a *granular* proxy for labor market slack in the euro area, which we term the *Perception of Labor Shortages* (PLS). Available from 1990Q1 at the level of 23 manufacturing sub-sectors, PLS is obtained from the European Commission's Business and Consumer Surveys and measures the share of firms reporting labor shortages as a key constraint on production. We argue that this measure is naturally interpreted as a proxy for firms' marginal hiring costs – and, ultimately, for labor market slack. This interpretation is grounded in New Keynesian models with search and matching frictions, in which a key determinant of firms' real marginal cost is the vacancy filling cost (e.g., Pissarides, 2000). In such frameworks, inflation responds to the cost of hiring the *marginal* worker rather than to *average* labor costs. PLS directly captures firms' perceived difficulty in filling vacancies and thus provides a sub-sector-level measure of labor market tightness that maps closely into the marginal cost channel emphasized by textbook-type Phillips curves. To validate this interpretation, we document a strong empirical correlation between PLS and standard measures of slack commonly used in the literature.

Importantly, the granularity of PLS – not yet used in literature – allows us to exploit both time-series variation and cross-sectional variation across countries and manufacturing sub-sectors when estimating the Phillips curve. This feature substantially strengthens identification in the spirit of recent works emphasizing the importance of disaggregated data for recovering aggregate relationships[2]. By moving beyond traditional aggregate indicators of slack, our approach not only provides a more informative empirical basis for assessing inflation dynamics but also alleviates important data limitations present in euro area data, which include the relatively short post-euro sample and the limited availability of potentially important variables, such as job vacancies.

The second major challenge in estimating the Phillips curve is controlling for unobserved factors, such as inflation expectations and cost-push shocks, that vary over time and may affect countries and sub-sectors differently. A common approach in the recent literature is to include time fixed effects, which amounts to absorbing these shocks through a set of time dummies and thus implicitly assumes that they load uniformly across all units[3]. This

[1] Barnichon and Shapiro (2024); Albuquerque and Baumann (2017).
[2] Fitzgerald et al. (2024); McLeay and Tenreyro (2020); Hazell et al. (2022).
[3] E.g. Gitti and Cerrato (2025) and Hazell et al. (2022).

restriction can be overly strong; if unobserved aggregate shocks interact with sub-sector- or country- specific characteristics, time fixed effects may fail to adequately capture this heterogeneity and can bias the estimated Phillips curve's slope.

We address this second issue by employing an interactive fixed effects (IFE) framework, which explicitly accounts for latent common components that vary over time and allows their impact to differ across sub-sectors and countries. Formally, we estimate

$$\pi_{ijt} = \alpha x_{ijt} + \lambda_{ij} F_t + \gamma_i + \tau_j + \delta \Gamma_{ijt} + \eta_{ijt}, \tag{2}$$

where γ_i and τ_j are sub-sector and country fixed effects, λ_{ij} is a vector of factor loadings, and F_t is a vector of common factors; both λ_{ij} and F_t are unobserved. Γ_{ijt} is a vector of sub-sector-country specific controls and includes deviations of the sub-sector-country relative producer price index from its steady-state value of one.

Interactive fixed effects flexibly capture latent common shocks – such as energy price shocks, global demand or inflation expectations – by estimating a factor structure rather than simply absorbing them non-parametrically as time fixed effects would. An interesting aspect of this approach is that it naturally relates to the literature using principal components to extract common inflation dynamics across countries, such as Ciccarelli and Mojon (2010) and Auer et al. (2024). Importantly, this framework also allows the influence of the common shocks to vary across units, providing heterogeneous exposure. These features turn out to be important, as that conventional estimation approaches – including time fixed effects, widely used in this literature – implicitly restrict how common shocks propagate across units and lead to systematically biased estimates of the Phillips curve slope. Furthermore, this bias is not incidental: it is amplified precisely by the granularity of the data, suggesting that studies exploiting disaggregated series may be particularly exposed to that if heterogeneous factor loadings are left unaccounted for. In contrast, the proposed approach delivers substantially more robust estimates, revealing labor market tightness to be a quantitatively important determinant of inflation and restoring a statistically and economically significant Phillips curve relationship in the data.

By combining a granular measure of slack with interactive fixed effects, our framework addresses both major sources of identification bias in the Phillips curve: measurement of marginal labor costs and unobserved common shocks. This approach builds on and advances recent contributions highlighting that the Phillips curve may appear flat in aggregate data

but remains responsive at more disaggregated levels (McLeay and Tenreyro, 2020), while also leveraging cross-sectional variation across sectors and countries to strengthen inference (Mavroeidis et al., 2014).

In the second part of the article the novelties introduced in the paper are deployed in three applications, which are relevant for policy questions and that recently sparked new discussions in the academic literature. The first one investigates a competing mechanism proposed in Afrouzi et al. (2024). The authors construct a theoretical framework in which rigidity in nominal wages motivates workers to switch directly from one job to another following an increase in the overall price level. This mechanism generates an expansion in aggregate vacancies, thereby producing the appearance of a tight labor market. To engage with this newly proposed mechanism, we conduct an additional empirical exercise that corroborates the original mechanism put forward by Phillips (1958). Specifically, we examine the degree of pass-through from firms' perceived labor shortages to wages, by showing that firms' labor shortage perceptions translate systematically into higher wages, and confirm the main findings of this paper. While this evidence does not settle this ongoing debate, it provides suggestive evidence on the likely mechanism at play.

The second application revisits the framework proposed for the US by Benigno and Eggertsson (2023) in the European context. We contribute to that literature by accounting explicitly for non-linearities in the NKPC via a piece-wise linear relation using a data-driven, machine learning approach to estimate its kink. In particular, we exploit the unique granularity of PLS to estimate sub-sector-country specific kinks using a random forest algorithm. By leveraging this machine-learning code, we detect levels of labor market tightness at which inflation dynamics change, without ever imposing any *a priori* functional form or specific threshold. Unlike traditional approaches that assume pre-defined break points, our random forest framework lets the data speak for itself, flexibly detecting heterogeneity in kinks both across countries and sub-sectors.

A key strength of this identification strategy relies on the fact that multiple euro area countries and manufacturing sub-sectors experienced concurrent increases in both labor market tightness and inflation even before the pandemic, notably in 2007 and 2018-2019. These pre-Covid episodes offer rare but valuable variation – along with the post-pandemic period – that help identifying the Phillips curve, especially given the limited time span of the euro area's monetary union. Importantly, though, if both labor market and inflation dynamics were uniformly synchronized across members and sub-sectors, panel data would offer little

advantage over aggregate analysis. Instead, the data shows substantial idiosyncratic variation as well – not just in the *levels*, but also in the *timing* and *intensity* of tightness-inflation episodes. This heterogeneity transforms the panel approach from merely "useful" to methodologically essential, as it provides the identifying variation needed to isolate slope coefficients and non-linearities that would otherwise be unobservable in short country-specific series.

This analysis confirms a non-linear sectoral Phillips curve that experiences a slope increase by more than two times when the labor market is tight. Indeed, failing to account for this non-linearity may lead the central bank to underestimate the inflationary pressures that arise when labor markets become sufficiently constrained, leading to a larger-than-expected surge in inflation.

In the third application, we extend the analysis by estimating aggregate coefficients that map micro-level fluctuations in perceived labor shortages into macro inflation dynamics. This allows us to assess whether differences in labor market tightness across sectors and countries can quantitatively account for the euro area's aggregate inflation behavior. The results confirm the sectoral findings and show that our specification captures major inflation episodes, including the decline during the Great Financial Crisis, the prolonged disinflation following the European sovereign debt crisis and the sharp rise in inflation during the post-pandemic recovery.

Overall, these findings have important policy implications for the conduct of monetary policy. We show that firms' perception about labor shortages are a useful proxy for labor market slack, underscoring the importance for policymakers to keep a vigilant eye on companies' reported hiring challenges as a warning sign of inflationary pressures. While some central banks already take into account qualitative information from business surveys[4], these results suggest that a more systematic, quantitative and sector-disaggregated approach, especially related to hiring difficulties, could substantially improve the monitoring and assessment of inflationary pressures.

Moreover, the significant non-linearity we uncover in the Phillips curve implies that inflation becomes much more sensitive to labor shortages when the labor market is tight. This finding suggests that relying on linear models may lead to an underestimation of inflation risks during tight labor market phases and an overestimation of the costs of disinflation.

The rest of this paper is structured as follows. Section 2 summarizes the related litera-

[4]E.g., the US Federal Reserves' Beige Book.

ture and highlights our contributions and Section 3 discusses the data used throughout the analysis. Section 4 presents the identification and the estimation strategies along with the main empirical findings. Section 5 explores some recent debates in the literature through the lenses of our data and empirical framework. Section 6 concludes.

2 Related Literature

We do not aim to survey the large and ever growing literature on the Phillips curve[5], we focus on those papers that are more closely related to ours.

First and foremost, this paper contributes to the literature that uses cress-sectional regional panel data to better identify the structural relationship between slack and inflation. Prior studies show that exploiting cross-sectional variation across U.S. states or subnational units yields more robust Phillips curve estimates than aggregate time series approaches (Fitzgerald et al., 2024; Hazell et al., 2022; McLeay and Tenreyro, 2020; Mavroeidis et al., 2014), which often suffer from endogeneity and weak-instrument issues.

While most research focuses on the U.S. (Beaudry et al., 2025; Gitti and Cerrato, 2025; Gitti, 2024; Hooper et al., 2020), we focus on the euro area, where distinct challenges arise. Similar attempts (Smith et al., 2025; Florio et al., 2025; Eser et al., 2020)[6] face important data constraints: short time spans, low frequency and lack of harmonized proxies for slack. To address these, we introduce the Perception of Labor Shortages (PLS) as a proxy for labor market slack at the granular level[7]. We construct a novel panel dataset spanning euro area countries and manufacturing sub-sectors, exploiting both cross-national and cross-sectoral variation, which enables us to overcome data limitations and to get to a robust identification of Phillips curve coefficients through idiosyncratic variation in both labor market tightness and producer prices.

This paper also engages with three recent debates of the literature. The first relates to an alternative mechanism proposed, relative to the conventional one, which poses that higher inflation can sometimes create the appearance of labor market tightness (Afrouzi et al., 2024).

[5]For a recent review see Furlanetto and Lepetit (2025).

[6]Mostly with country-level proxies, except Levy (2019) and Beschin et al. (2025) who used NUTS-2 data, which however is limited to annual frequency.

[7]Some studies use the survey measures to approximate vacancies or analyze aggregate wage-setting (Kiss et al., 2022; Bonthuis et al., 2016; Bonam et al., 2021); our approach exploits sub-sector-country granularity to link labor shortages to inflation dynamics.

While we do not aim to address this issue thoroughly, our wage Phillips curve estimates show a positive, significant link between labor market tightness and wages, supporting the view that tightness drives inflation rather than the reverse[8]. The second debate concerns how to properly control for inflation expectations. Recent work stresses the importance of short-term expectations in Phillips curve estimation (Reis, 2025; Beaudry et al., 2025). We find that the Phillips curve estimated under our preferred specification is largely unchanged when controlling for inflation expectations, in contrast to traditional regression approaches for which the estimated relationship is markedly weakened and not significant. Finally, we contribute to the discussion on Phillips curve non-linearities. U.S. evidence is mixed, with some studies finding strong non-linearities (Gitti and Cerrato, 2025; Gitti, 2024; Benigno and Eggertsson, 2023; Hooper et al., 2020; Nalewaik, 2016), while others reporting mild or no non-linear effects (Beaudry et al., 2025; Ball et al., 2022; Hazell et al., 2022; Murphy, 2017; Babb and Detmeister, 2017). For the euro area, Faber and Züllig (2025) and Byrne and Zekaite (2018) find heightened sensitivity of wage and price inflation under tight labor conditions, whereas Lo Bello and Viviano (2024) detect no non-linear wage effects but still report that tightness measures improve forecasting accuracy.

Our contribution is twofold: (i) we employ the PLS indicator as a better proxy for labor market slack in the euro area and (ii) we provide robust evidence of non-linear Phillips curve relationships using granular cross-country, cross-sector panel data. This approach improves identification and mitigates endogeneity – critical advantages given the euro area's limited (historical) data.

3 Data and Measurement

Our primary data source is the Joint Harmonised EU Programme of Business and Consumer Surveys (BCS), which is administered by the European Commission's Directorate-General for Economic and Financial Affairs each month, with additional quarterly questions. The surveys encompass a large sample of firms: 19,380 firms for the 23 manufacturing sub-sectors for the countries in our sample. We use data for most euro area founding members (Austria, Belgium, Germany, Spain, Finland, France, Italy, the Netherlands and Portugal[9]) and the

[8] Ampudia et al. (2024) document strong wage-to-price pass-through in Europe.

[9] We exclude Luxembourg, for which inflation (PPI) is missing, and Ireland, for which PLS data are only available for the most recent period.

baseline sample period spans from 2002 to 2024[10].

Participation is mandatory in several of the countries we consider (Spain, France, Italy, the Netherlands and Portugal), ensuring high response rates and data quality. According to the European Commission's methodological guide, the sample of firms does not differ much from one period to another, as a panel of companies is established and surveyed each month and quarter.

Among the questions asked in the surveys, we focus on responses to the quarterly *Question 8*, which asks firms "*What main factors are currently limiting your production?*". Respondents are requested to select one or several factors: (i) none; (ii) insufficient demand; (iii) shortage of labor force; (iv) shortage of material and/or equipment; (v) financial constraints; (vi) other factors. The BCS dataset then reports, at the NACE-2 level (i.e. at the sub-sector level for each country)[11], the share of firms that respond that their production is currently constrained by each of those factors.

Table 1 (top panel) reports the correlation between PLS and the vacancy-to-unemployment (V/U) ratio – commonly used as proxy for labor market slack in recent Phillips curve studies – as well as with its individual components (vacancies and unemployment), for the countries for which those data are available from Eurostat. While these series are reported at the country level, we additionally compute correlations at the aggregate euro area level[12]. The main takeaway is that PLS correlates strongly and positively with the V/U ratio, with a correlation coefficient very close to one. As expected, PLS correlates positively with vacancies and negatively with unemployment, reinforcing its validity as a proxy for labor market tightness. Figure 1a visually shows the co-movement of two the time series.

We validate this correlation by showing that a survey measure close to PLS for the United States also displays a strong positive correlation with the V/U ratio. Results are reported

[10]While data are available for some countries and sub-sectors from 1990Q1, in the baseline results presented we restrict the sample starting from the first quarter of 2002 to avoid capturing co-movements due to the convergence period before the introduction of the euro (see Benati (2008) on the importance of a change in the monetary regime for inflation). Results are robust to this choice and the table for the full 1990-2024 sample is reported in the Appendix.

[11]NACE is the French acronym for "*Nomenclature statistique des Activités économiques dans la Communauté Européenne*" and is the statistical classification of economic activities used by the European Union to ensure comparability of data across member states. See Table 6 in the Appendix for more details on the sub-sectors considered in this study.

[12]Throughout, "euro area" refers to the group of countries included in our sample, as detailed earlier and in the Appendix.

in Table 1 (bottom panel) and in Figure 1b. Details on the data sources are relegated to the Appendix[13].

TABLE 1: Correlation Coefficients with PLS

Euro Area			
Aggregation Level	V/U Ratio	Unemployment	Vacancies
Country	0.877	-0.349	0.518
Euro Area	0.957	-0.888	0.868
United States			
U.S.	0.912	-0.531	0.960

Notes: the table shows the correlation of the vacancy-to-unemployment ratio, number of unemployed people and number of vacancies with PLS. When aggregation is performed, we compute weighted averages based on official weights reported by the EU Commission.

This positive correlation further supports interpreting PLS as a meaningful proxy for firms' marginal costs. In New Keynesian models with search and matching frictions (e.g., Pissarides, 2000), firms' pricing decisions are driven by their *marginal*, rather than average, labor costs. When labor markets tighten – reflected in higher PLS – firms face greater difficulty filling vacancies. The marginal cost of production therefore reflects the costs associated with unfilled positions. PLS captures precisely this friction, as it reflects the resources and effort required to hire an *additional* worker – the so-called "vacancy-filling" costs – which constitute a key component of firms' real marginal cost and, consequently, influence their price-setting behavior.

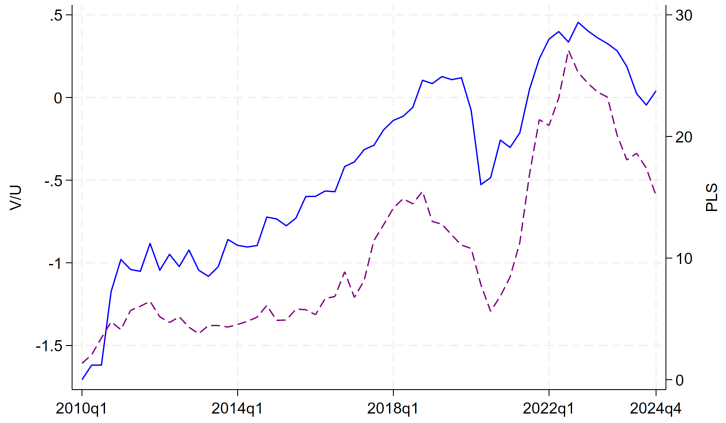
This interpretation bridges the microfoundations of modern New Keynesian Phillips curves (e.g., Galí, 2015) with labor market search frictions (e.g., Pissarides, 2000). In standard New Keynesian models, inflation is driven by real marginal costs; once search frictions are introduced, these costs also incorporate the expected duration and expense of recruiting. A tighter labor market raises vacancy-filling costs, thereby increasing marginal costs even conditional on wages, and ultimately generating inflationary pressures.

Unlike some studies for the U.S., our dataset is based on harmonized questionnaires and a common methodology, which ensures comparability across countries and sub-sectors, pro-

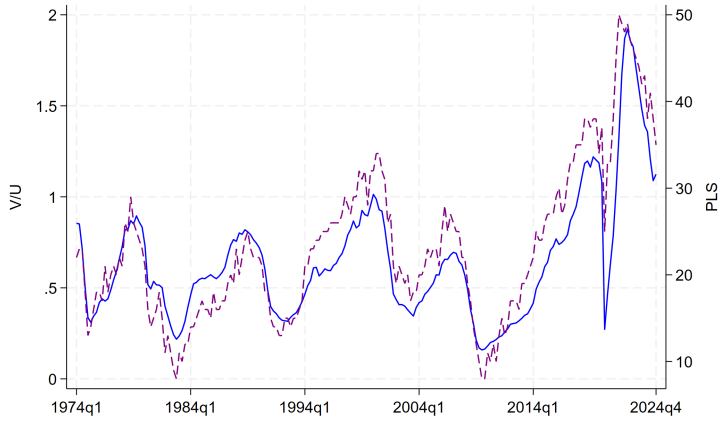
[13]We thank Pierpaolo Benigno for kindly sharing his data on vacancies and unemployment for the U.S.

FIGURE 1: Correlations of V/U and PLS

(A) Euro Area Average V/U ratio and PLS



(B) V/U ratio and PLS for U.S.



Notes: the figures show the time-series of aggregate PLS – in purple (dashed), right scale – and the natural logarithm of the V/U ratio – in blue, left scale. The plots start at different points in time due to data availability of the V/U series. When aggregating the euro area data, we compute weighted averages based on official weights reported by the EU commission. U.S. data are downloaded directly at the country level.

viding a more robust framework for cross-country analysis.

Another important dataset for our purposes is Eurostat's Short-Term Statistics database for the Producer Price Index (PPI), which serves as our primary measure of inflation. The PPI is particularly well-suited for our analysis as it aligns directly with BCS variables – particularly our measure of labor market slack – at the NACE-2 (sub-sector) level, allowing for a precise examination of the relationship between labor shortages (together with the other variables)

and inflation.

Appendix A.2 provides more detailed information on the data sources and the sample selection procedures adopted.

4 Empirical Strategy

We now outline the empirical methodology for estimating the Phillips curve in our panel of euro area countries.

4.1 Identification

A major challenge in studying inflation dynamics is to correctly disentangle the impact of (labor market) slack from supply shocks and inflation expectations. The recent literature[14] has proposed using disaggregated cross-sectional data to produce additional empirical evidence in an effort to obtain better estimates of the Phillips curve, a central building block for economists' understanding of inflation dynamics.

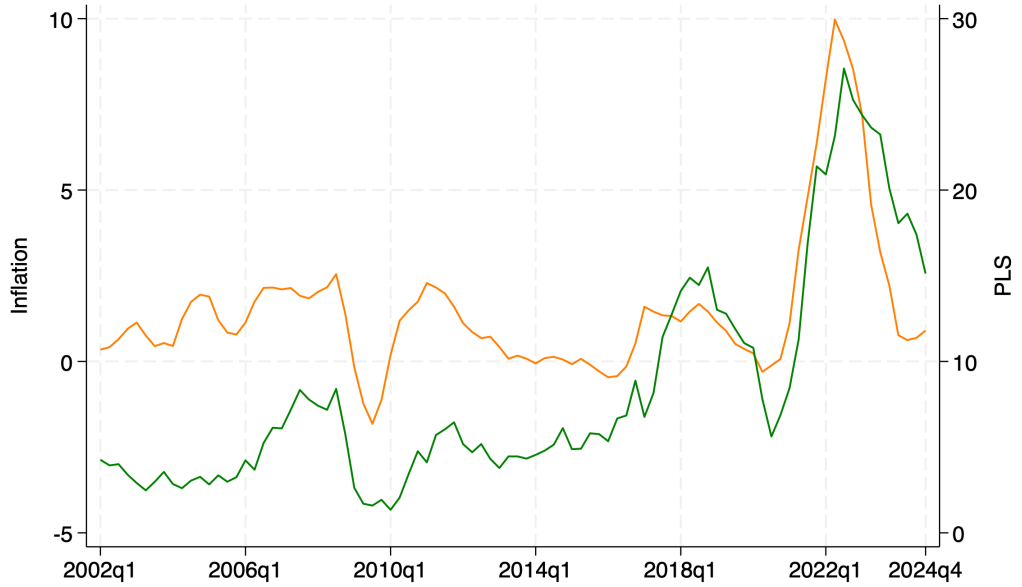
Our newly assembled dataset overcomes the limitations of single-country time series by exploiting cross-country and cross-sector variation. This feature helps mitigate the simultaneity bias that arises from the endogenous response of monetary policy to demand shocks. Intuitively, inflation and PLS are influenced both by idiosyncratic shocks specific to each sub-sector–country pair and by aggregate conditions. While monetary policy reacts to the latter, it does not respond to idiosyncratic deviations at the sub-sector–country level. As a result, these local variations in PLS act as quasi-exogenous shocks in the Phillips curve regression, allowing us to more cleanly identify the effect of labor market tightness on inflation. This logic underlies the approach of Fitzgerald et al. (2024), who use regional U.S. data to address the same endogeneity concerns.

Another important aspect of our panel structure is that it allows to estimate common factors that capture aggregate supply shocks and inflation expectations. In particular, long-run inflation expectations are largely independent of the business cycle, reflecting shared beliefs about the monetary regime (Hazell et al., 2022; Benati, 2008). As a result, these expectations are constant across sub-sector-country pairs in the long run and are absorbed by the

[14]For excellent and more detailed overviews the reader may want to refer to Mavroeidis et al. (2014) and McLeay and Tenreyro (2020).

estimated common factors, helping to isolate the effect of labor market tightness on inflation.

FIGURE 2: Euro Area Average Perception of Labor Shortages (PLS) and Inflation



Notes: the figure shows the time-series of aggregate PLS – in green, right scale – and (PPI) inflation – in orange, left scale. When aggregating, we compute weighted averages based on official weights reported by the EU commission.

Importantly, though, recent studies (Reis, 2025; Beaudry et al., 2025) argue that inadequately controlling for inflation expectations can lead to biased Phillips curve's coefficients and, more specifically, they emphasize the importance of properly accounting for short-term inflation expectations. Therefore, in our analysis we will augment our most preferred specification with information on short-term selling price expectations from the BCS database. In particular, the BCS provides answers to the following question at the sub-sector-country level: "How do you expect your selling prices to change over the next 3 months?". Firms can answer to the question in three ways: (i) increase, (ii) stay the same or (iii) decrease. We will include the responses in the regression analysis.

Descriptive evidence. The weighted averages of both inflation and PLS are plotted in Figure 2 and show how the time series evolve over the sample period. It highlights two aspects of the data: (i) there is a clear co-movement between PLS and inflation and (ii) post-COVID, inflation and perceived labor shortages both rose sharply, pointing to heightened inflationary pressures and tighter labor markets in the wake of the pandemic.

FIGURE 3: Percentage of Sub-Sector-Countries with Tight Labor Markets



Notes: the figure shows the percentage of sub-sector-country pairs that are above the estimated threshold.

Importantly, however, such dynamics are not exclusive to the post-COVID period. The percentage of sub-sector-country pairs with tight labor markets[15] plotted in Figure 3 show that, already in 2007 and even more so in 2018-2019, several euro area countries and manufacturing sub-sectors experienced simultaneous tight labor markets, indicating that such periods of tightness and price pressures are also present in pre-pandemic data. Crucially though, these patterns are not necessarily uniform across countries or sectors. As shown in Online Appendix B.1, there is considerable heterogeneity in both the *timing* and the *intensity* of these episodes. This cross-sectional variation is particularly valuable in the context of the euro area. Unlike in the United States, long historical series for unemployment and vacancies are limited, and the euro area itself is a relatively recent monetary union. As a result, clear identification of the Phillips curve from national time series alone is difficult. Exploiting cross-country and sectoral heterogeneity provides a richer basis for estimating the slope of the Phillips curve in the euro area and allows us to make better use of the few observable episodes of labor market tightness.

[15]In the next section we will define what characterize a slack or tight labor market. As an anticipation, labor markets are tight when the perception of labor shortages surpasses the respective threshold (i.e. the kink of the curve) and in the next section we will clarify how we identify and estimate these thresholds.

4.2 Estimation

We begin our empirical analysis with a simple regression

$$\pi_{ijt} = \alpha x_{ijt} + \epsilon_{ijt}, \tag{3}$$

which corresponds to a stripped-down version of the standard NKPC, including only labor market slack and abstracting from inflation expectations and cost-push shocks. Any correlation between the error term, ϵ_{ijt}, and our measure of labor market slack, x_{ijt}, can bias the estimate of α, the coefficient of interest. The baseline regression in equation 3 does not account for such correlations, providing a raw estimate of the link between inflation and labor market tightness. In doing so, it omits a number of important factors known to affect inflation, including inflation expectations, cost-push shocks and persistent sub-sector and country characteristics. Since these forces potentially account for a large share of inflation variation, attributing all inflation movements to changes in labor market tightness alone will likely overstate the true α.

A first, admittedly crude, attempt to mitigate this omitted variable bias is to include a constant term in the regression. Doing so decomposes the residual as $\epsilon_{ijt} = c + \eta_{ijt}$, where c captures the average level of inflation and η_t reflects deviations from that mean. By absorbing the unconditional mean of inflation, the constant already removes a non-trivial source of variation that would otherwise be mechanically attributed to labor market slack.

This adjustment provides a useful benchmark and sets the stage for understanding how progressively richer treatments of heterogeneity and unobserved components – through fixed effects, controls and, ultimately, interactive fixed effects – shape the estimated Phillips curve relationship.

Fixed effects. A common way to address the potential correlation between the error term and labor market slack is to include fixed effects, which control for persistent heterogeneity across sub-sectors and countries.[16] We begin by including two fixed effects, one for sub-sectors and one for countries, such that $\epsilon_{ijt} = \gamma_i + \tau_j + \eta_{ijt}$, where γ_i and τ_j capture differences across sub-sectors and countries, respectively, and η_{ijt} is the remaining idiosyncratic error.

Sectors and countries differ in many respects, including labor productivity, wage-setting institutions, bargaining power and labor market regulations. By absorbing these time-invariant

[16] The constant term is included in all subsequent regressions.

differences, the two-way fixed effects specification removes an important source of omitted variable bias. Relative to the baseline, we expect the estimated slope α to decline as it is likely that part of the raw correlation between labor market tightness and inflation reflects structural differences across sub-sectors or countries rather than genuine variation in slack.

We then include a time fixed effect, such that $\epsilon_{ijt} = \gamma_i + \tau_j + \omega_t + \eta_{ijt}$, where ω_t captures shocks common to all units in a given period. Including time fixed effects has been a common strategy in recent contributions (e.g., Fitzgerald et al., 2024; Hazell et al., 2022; McLeay and Tenreyro, 2020; Mavroeidis et al., 2014) to absorb unobserved, time-varying factors that affect inflation simultaneously across sectors and countries, such as shifts in inflation expectations or changes in global supply costs. By controlling for these common components, the researcher aims to reduce bias in the estimated slope of the Phillips curve.

However, conventional time fixed effects absorb common changes in a purely non-parametric way, treating each period as a separate dummy rather than modeling an underlying latent factor that summarizes the shared dynamics across units. In contrast, the interactive fixed effects framework explicitly *estimates* latent common factors from the data and allows their impact to vary across units, rather than imposing the restrictive assumption that all sub-sectors and countries respond identically. Conventional time fixed effects, by contrast, enforce this uniformity, assuming that every unit reacts in the same way to the same common factor. In practice, however, the transmission of these influences is likely heterogeneous. Energy price increases may disproportionately affect energy-intensive manufacturing sectors, while monetary policy may have different effects across euro area countries. When such heterogeneity is present, time fixed effects may absorb not only common shocks but also meaningful variation in inflation that is due to labor market tightness. Therefore, relative to the two-way fixed effects specification, we expect the estimated slope α to be further attenuated, and potentially flatter than the "true" relationship would imply. These potential challenges – carefully controlling for common shocks and allowing for heterogeneous responses – provide a central motivation for the interactive fixed effects framework introduced next.

Controls. Before introducing interactive fixed effects, we estimate specifications that include observable control variables at both the disaggregated level, Γ_{ijt}, and the country level, Θ_{jt}. The residual can be written as $\epsilon_{ijt} = \gamma_i + \tau_j + \delta\Gamma_{ijt} + \sigma\Theta_{jt} + \eta_{ijt}$. These controls[17]

[17]Sub-sector-country specific controls available from BCS are: insufficient demand; shortage of material and/or equipment; financial constraints; other factors and a proxy for short-term inflation expectations. We also include deviations of the sub-sector-country relative producer price index from its steady-state value

are designed to capture time-varying forces that are plausibly correlated with inflation and labor market tightness. By excluding time fixed effects, this specification avoids imposing the restrictive assumption that all sub-sectors and countries respond identically to aggregate shocks. This flexibility is particularly important in a highly disaggregated setting, where the transmission of common disturbances can differ markedly across sectors and countries. At the same time, this approach remains inherently limited. Observable controls are imperfect proxies for the underlying unobserved components that drive inflation dynamics, such as inflation expectations or broad cost-push pressures. If these forces are only partially captured by the available controls, the residual may still contain unobserved common shocks that are correlated with labor market tightness, leaving scope for bias in the estimated Phillips curve slope.

Relative to the three-way fixed effects specification, the estimated slope may increase if time fixed effects previously absorbed variation in inflation that reflects heterogeneous responses to common factors and is, instead, genuinely related to labor market tightness. However, the estimated slope may also decline if the added controls absorb variation that was previously – and spuriously – attributed to slack in the three-way fixed effects regression. These controls are inherently imperfect though, because, as observable proxies for unobserved components of the NKPC, they may absorb variation that instead reflects labor market slack. For this reason, the direction of change in terms of magnitude is *a priori* ambiguous.

Interactive Fixed Effects. To overcome the limitations discussed above, we propose an interactive fixed effects (IFE) regression as in equation (2)

$$\pi_{ijt} = \alpha x_{ijt} + \lambda_{ij} F_t + \gamma_i + \tau_j + \delta \Gamma_{ijt} + \eta_{ijt}.$$

Here, the residual can be decomposed as: $\epsilon_{ijt} = \gamma_i + \tau_j + \lambda_{ij} F_t + \eta_{ijt}$, where $\lambda_{ij} F_t$ captures latent common shocks that vary over time but have heterogeneous effects across sub-sectors and countries. λ_{ij} is a vector of factor loadings and F_t is a vector of unobserved common factors. Both λ_{ij} and F_t are estimated from the data using a principal components approach, providing a parsimonious representation of unobservable inflation dynamics (Mumtaz et al., 2011; Miranda et al., 2022)[18].

of one and, for panel regressions, lagged inflation over the previous four quarters. At the country level we include real effective exchange rates and energy prices, at the euro area level.

[18]This aspect is rather important as one could in principle allow for fully interacted fixed effects at the sub-sector-country-time level. However, in a panel of this dimensionality such a specification would effectively

In more detail, this component of the variance decomposition captures the co-movements of inflation across countries by extracting $F \geq 1$ unobservable common shocks, each weighted by sub-sector-country-specific factor loadings. Put differently, all countries face the same underlying shocks, but their impact on inflation varies across sectors and countries. The number of common factors, F, is chosen using data-driven information criteria (Ahn and Horenstein, 2013; Onatski, 2010; Gagliardini et al., 2019). A more detailed discussion on the estimation procedure is provided in Appendix A.8.

In this framework, the variation that is common across units is captured by F_t and, while persistent sub-sector-country differences in inflation can still exist (for example due to structural productivity differentials), as long as these differences are time-invariant, they are absorbed by the two-way fixed effects (Hazell et al., 2022). The interactive fixed effects approach thus decomposes inflation variation into two components: (i) variation explained by observable factors, such as PLS; (ii) variation driven by unobserved common inflation dynamics, such as long-run expectations and common cost-push shocks.

This decomposition offers several advantages over conventional time fixed effects. First, it provides a parsimonious representation of unobserved common components (Mumtaz et al., 2011). Second, it is data-driven, relying on principal component analysis to estimate the latent factors, which reduces the risk of misspecification (Miranda et al., 2022) while allowing for heterogeneous responses across sub-sectors and countries. Furthermore, this regression framework provides a natural bridge to the literature on global inflation dynamics, in which common inflation components are identified via principal component analysis (Ciccarelli and Mojon, 2010; Auer et al., 2024). To our knowledge, this represents the first application of interactive fixed effects to disaggregated (euro area) data in the Phillips curve literature.

As a result, the slope may be smaller compared to two-way fixed effects – reflecting proper control for aggregate dynamics – but should be larger than under three-way fixed effects, in which case it would be consistent with the view that the apparent flattening of the Phillips curve is largely an identification artifact rather than an economic one.

saturate the model, requiring the estimation of an extremely large number of parameters and absorbing much of the variation needed to identify the Phillips curve slopes.

4.3 Empirical Findings

This section presents the empirical results and evaluates whether the hypotheses developed in the previous section are supported by the data. Tables 2 and 3 report estimates of the Phillips curve slope using alternative specifications that differ in both how they decompose the error term and how they account for unobserved components.

TABLE 2: Phillips Curve Estimates

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.151***	0.107***	0.133***	0.031***	0.042***	0.002
	(0.004)	(0.005)	(0.011)	(0.008)	(0.006)	(0.007)
Constant	-	1.042***	1.513***	0.120	1.192***	-0.449
		(0.052)	(0.261)	(0.381)	(0.162)	(0.353)
Observations	14399	14399	14399	14399	9095	9094

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Baseline. Column (1) reports the baseline regression corresponding to equation (3), which relates inflation directly to labor market tightness without inflation expectations, the constant term, cost-push shocks or any fixed effects. While this specification provides a useful benchmark, it is likely biased due to omitted variables. Hence, column (2) introduces a constant term and this simple adjustment already significantly reduces the estimated slope. This decline is consistent with our prior that part of the raw correlation in column (1) reflects average inflation levels rather than within-unit comovement between slack and inflation. Although highly stylized, this specification illustrates that even minimal structure can absorb a nontrivial share of variation previously attributed to labor market tightness, setting the stage for more refined decompositions of the error term.

Fixed effects. Column (3) adds two-way fixed effects at the sub-sector and country level, accounting for persistent structural differences across units. Once this heterogeneity is controlled for, the estimated slope remains sizable and statistically significant – and notably increases relative to the constant-only specification, differently from what we hypothesized.

This indicates that the Phillips curve relationship is driven primarily by within sub-sector-country fluctuations in PLS over time, rather than by cross-sectional differences in average inflation or tightness.

Column (4) augments this specification with time fixed effects, absorbing shocks common to all sub-sectors and countries in a given period. This approach has been widely used in the recent Phillips curve literature to control for unobserved time-varying factors such as global energy prices and inflation expectations. Consistent with our prior, the estimated slope falls sharply.

This substantial attenuation partly reflects the restrictive nature of time fixed effects. By construction, not only they capture common factors only through time dummies rather than estimating their underlying structure but they also impose homogeneous responses to them across all units. In a disaggregated setting, the latter assumption is unlikely to hold. When heterogeneous responses are present, time fixed effects can absorb variation that is genuinely related to labor market tightness. As a result, while three-way fixed effects reduce bias from common factors, they may also render the Phillips curve excessively flat relative to the true underlying relationship.

Controls. Columns (5) and (6) replace time fixed effects with observable control variables, while retaining sub-sector and country fixed effects. In column (5), which includes controls excluding short-term inflation expectations, the estimated Phillips curve slope is smaller than in the two-way fixed effects specification in column (3), suggesting that part of the variation previously attributed to labor market tightness may reflect omitted factors correlated with PLS. At the same time, the coefficient remains larger than in the time fixed effects specification in column (4), indicating that time fixed effects may be overly restrictive and absorb variation related to labor market conditions.

In column (6), we additionally control for a proxy of short-term inflation expectations. Once this variable is included, the estimated slope declines sharply and becomes statistically insignificant. This result admits two interpretations. On the one hand, it is consistent with the argument advanced by Reis (2025) and Beaudry et al. (2025) that conditioning on short-term inflation expectations is critical and can substantially flatten – or even eliminate – the Phillips curve relationship. On the other hand, it may reflect the limitations of control-based approaches: observable proxies, including those for inflation expectations, are imperfectly measured and may absorb variation that is instead attributable to labor market slack.

Interactive fixed effects. Columns (7)-(13) report estimates from interactive fixed effects (IFE) regressions, which explicitly estimate latent common factors from the data and allow their effects to vary across sub-sectors and countries.

In column (7), the baseline IFE specification without additional fixed effects or controls delivers an estimated slope that is twice as large as the corresponding time fixed effects estimate. This increase is consistent with our hypothesis that by extracting the latent common factor and permitting heterogeneous transmission across units, the IFE approach recovers variation in labor market slack that would otherwise be absorbed by time fixed effects.

Column (8) augments the interactive fixed effects specification by adding sub-sector and country fixed effects, which absorb persistent differences in average inflation across units, such as structural productivity gaps or institutional features. The interactive fixed effects then isolate time-varying common shocks whose impact differs across sectors and countries. The estimated slope increases, indicating that time-invariant heterogeneity and time-varying common components capture distinct sources of inflation variation. Once persistent level differences are netted out, the remaining variation in labor market tightness is more cleanly identified. This result shows that failing to separately account for structural heterogeneity can attenuate the estimated response of inflation to slack.

Table 3: Phillips Curve Estimates

	(7)	(8)	**Factor Loadings**				(13)
	IFE	FE	Same	Country	Sector	Sec. × Country	Infl. Exp.
PLS	0.062***	0.070***	0.039***	0.038***	0.083***	0.088***	0.074***
	(0.003)	(0.003)	(0.007)	(0.007)	(0.005)	(0.005)	(0.004)
Constant	0.688***	1.359***	2.807***	2.864***	1.716***	1.706***	1.225***
	(0.040)	(0.038)	(0.158)	(0.157)	(0.101)	(0.097)	(0.098)
Observations	14399	14399	9445	9445	9445	9445	9444

Notes: Column (7) includes interacted fixed effects (IFE). Column (8) adds fixed effects (FE). Columns (9)–(12) report factor loadings under different specifications: homogeneous for both countries and sub-sectors (9), country (10), sub-sector (11) and sub-sector × country (12). Column (13) includes a proxy for short-term inflation expectations. All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Columns (9)–(12) systematically explore the importance of heterogeneous factor loadings,

isolating which specific aspect of heterogeneity matters most for identification. Column (9) imposes *homogeneous* factor loadings – meaning that common shocks are constrained to affect all units equally. The resulting slope estimate is very similar to the time fixed effects estimate reported in column (5) of Table 2. This is informative: our results are not driven by the presence of a factor structure *per se*, but by the flexibility to allow for heterogeneous propagation of shocks. When factor loadings are forced to be identical, the IFE specification effectively converges to the time fixed effects case, as intuition would suggest. This column therefore offers a key sanity check for the identification strategy of our paper.

Columns (10) and (11) then relax this homogeneity assumption along two distinct dimensions. Column (10) permits factor loadings to differ across countries, whereas column (11) allows them to differ across sub-sectors. Strikingly, introducing country-specific loadings leaves the estimate almost unchanged relative to the homogeneous benchmark, while introducing sub-sector-specific loadings nearly doubles it. This contrast indicates that heterogeneity across sub-sectors, rather than across countries, is the primary force behind the results. It implies that broad macroeconomic shocks – for example, movements in energy prices or shifts in global demand – may exert relatively similar aggregate effects across countries, but their transmission to prices at the sub-sector level can vary markedly, owing to the granularity and heterogeneity of the underlying data. At this level of disaggregation, the factor loadings allowed by the IFE specification capture meaningful differences in how common shocks propagate through various parts of the economy – differences that time fixed effects, which impose a uniform response, are unable to reflect.

Finally, column (12) combines interactive fixed effects with sub-sector-country specific controls, yielding the most thorough specification. This specification delivers the largest estimate in Table 3 for the Phillips curve's slope while remaining precisely estimated. Interestingly, while the interactive fixed effects capture broad common components the controls account for more granular, sub-sector-country specific variation. By absorbing these additional sources of confounding variation, the controls sharpen the identifying variation in labor market tightness rather than crowding it out. As a result, the estimated response of inflation to labor market slack strengthens, reinforcing the interpretation that previous specifications understated the true slope of the Phillips curve.

Column (13) further augments this specification by adding a proxy for short-term inflation expectations. We include this variable as a robustness check to further stress the flexibility of our framework, by allowing short-term expectations to deviate from the long-run ones and,

therefore, to vary across sub-sectors and countries. Even after conditioning on short-term inflation expectations, the estimated Phillips curve's slope remains large and statistically significant. This result confirms that the strengthening observed in column (12) is not driven by potentially omitted expectations dynamics and reinforces the robustness of the interactive fixed effects specification. For this reason, we adopt column (12) as our baseline specification for the remainder of the paper.

Taken together, the results strongly support the priors outlined in the previous section. Baseline and fixed effects specifications tend to overstate or understate the Phillips curve slope. Time fixed effects, while effective at absorbing common shocks, impose restrictive homogeneity assumptions that flatten the Phillips curve excessively. Control-based approaches relax these assumptions but may suffer from measurement error and incomplete coverage of unobserved inflation drivers.

Therefore, the interactive fixed effects framework offers a more credible compromise. By explicitly modeling common dynamics and allowing their impact to vary across sub-sectors and countries, IFE reduces bias without relying on restrictive assumptions. The results in columns (9)–(12) carry an additional implication for the broader empirical literature: studies that exploit highly disaggregated data – as is increasingly common in the recent literature on inflation dynamics – may be particularly susceptible to this form of bias. When data are granular, common shocks can transmit very differently across units, and imposing homogeneous factor loadings, as time fixed effects implicitly do, may systematically distort the estimated relationship between labor market slack and inflation. Failing to account for this heterogeneity is therefore not a minor specification choice, but a potentially consequential source of attenuation bias – one that becomes more severe precisely as the data become more disaggregated.

We therefore view the IFE specifications as providing the most reliable estimates of the relationship between labor market tightness and inflation. Importantly, the estimated Phillips curve slope is both economically and statistically significant. This finding challenges the view that the Phillips curve is intrinsically flat and instead suggests that much of the apparent flattening found in the literature reflects identification choices rather than structural changes in inflation dynamics.

5 Wages, Non-Linearity and Aggregate Implications

Having highlighted the value of granular labor market data and interactive fixed effects, we now use these tools to address a series of questions that play a central role in current policy and academic discussions. In particular, we use our data and methodology to (i) study the transmission of labor market tightness to wages, (ii) the presence of a non-linearity in the Phillips curve and (iii) the implications of sectoral inflation dynamics for aggregate inflation[19].

5.1 From Tight Labor Markets to Wages

The canonical view on slack and inflation, originating with Phillips (1958) and reinforced by a large body of theoretical and empirical work, holds that tight labor markets exert upward pressure on wages and, ultimately, on prices. When labor demand exceeds the available supply of workers, firms must offer higher wages to attract and retain employees, and these cost increases are – at least partially – passed through into output prices (Ampudia et al., 2024).

TABLE 4: Phillips Curve Estimates

	Wages
PLS	0.036**
	(0.015)
Constant	2.943***
	(0.421)
Observations	827

Notes: fixed effects (country level) and controls outlined in Section 4.2 are included. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A recent contribution by Afrouzi et al. (2024) proposes a complementary mechanism whereby nominal wage rigidity implies that increases in the general price level temporarily reduce real wages, inducing workers to search for better-paying jobs. A rise in job-to-job transitions can mechanically increase posted vacancies, generating the *appearance* of tighter labor markets even without a commensurate strengthening in underlying labor demand. This mechanism

[19] Unless otherwise noted, all subsequent results are based on the specification reported in column (9).

highlights a potential identification concern: if inflation itself triggers higher turnover and vacancy posting, standard tightness indicators may partly reflect the consequences of inflation rather than its causes.

Our setting is well suited to shed light on this distinction because it relies on firms' reported hiring difficulties rather than solely on vacancy posting behavior. To contribute to this discussion, we conduct an additional empirical exercise examining the pass-through from perceived labor shortages to wage growth. While this test is not definitive, it is informative: if perceived shortages mainly reflected inflation-induced search behavior, one might expect a weaker or less systematic link between shortages and wage growth. Instead, the results in Table 4 indicate that increases in PLS are associated with statistically significant wage increases.[20] Overall, this pattern is consistent with the interpretation that firms' perceived labor shortages capture economically meaningful constraints in the labor market, and it suggests that the Afrouzi et al. (2024)'s mechanism is unlikely to be the sole driver of the tightness signals observed in the data. Further work in this area offers a promising avenue for future research.

5.2 Non-Linearity

Recent work has highlighted that the Phillips curve may not be well approximated by a linear specification. In particular, Benigno and Eggertsson (2023) provide theoretical and empirical evidence – at the aggregate level and for the U.S. – that the slope of the Phillips curve is not constant but rather follows an inverse-L shape, with inflation responding very little to slack when labor markets are loose, but becomes highly sensitive once tightness surpasses a certain threshold. An contribution of this paper is to test this framework with disaggregated data for the euro area and a novel granular proxy for labor market slack. We consider the piecewise linear version of (2):

$$\pi_{ijt} = \alpha_o \, x_{ijt} + \alpha_d \times (x_{ijt} - x^*_{ij}) \times \mathbb{1}\{x_{ijt} > x^*_{ij}\} + \delta \Gamma_{ijt} + \lambda_{ij} F_t + \alpha_i + \tau_j + \eta_{ijt}, \qquad (4)$$

where $\mathbb{1}\{x_{ijt} > x^*_{ij}\}$ is an indicator variable that takes value of 1 when labor market tightness is greater than a certain threshold and 0 otherwise.

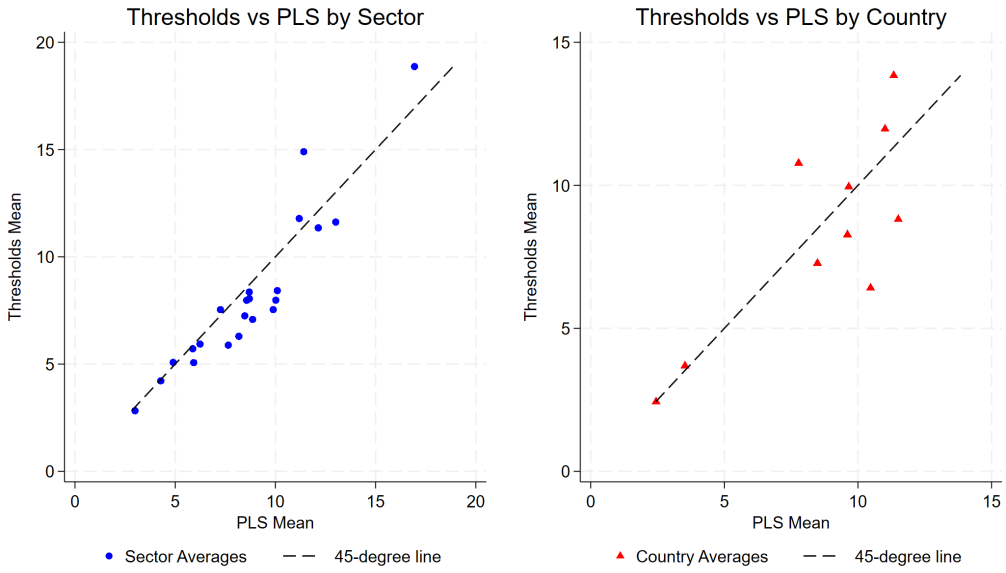
The difference between x_{ijt} and x^*_{ij} generates the non-linearity. When PLS is less than PLS*,

[20] The analysis is conducted at the country level, as wage indices are not available at the NACE-2 level.

the slope of the Phillips curve is α_o (the constant is omitted for simplicity). When PLS is greater than PLS*, the slope of the Phillips curve is $\alpha_o + \alpha_d$.

Thresholds' Estimation. To identify the potential threshold in the relationship between labor shortages and inflation, we use a model-free, data-driven approach. Specifically, we estimate thresholds using a random forest regressor[21] trained separately on each sub-sector-country pair. Each forest consists of 500 decision trees, where every tree is fit to a bootstrap resample of the data. Within a single decision tree, the algorithm recursively searches for the best split points in the predictor space. When considering the relation between labor shortages and inflation, the algorithm tests candidate splits at which the relation may change and, at each step, the algorithm selects the split that yields the greatest improvement in prediction accuracy[22]. This process is repeated recursively for each branch of the tree.

FIGURE 4: Scatter Plot of Thresholds vs PLS Averages by Sector and Country



Notes: the figure shows the scatter plots of by sector and country of thresholds against PLS averages. The figure serves as anecdotal evidence and in support of the economic significance of the random forest procedure adopted. When aggregating, we compute weighted averages based on official weights reported by the EU commission.

[21] Random forests are widely used in modern machine learning because they handle high-dimensional data, reduce overfitting and work well with minimal tuning – unlike e.g. neural networks, which are highly sensitive to hyperparameter choices.

[22] Equivalently, the largest reduction in mean squared error or increase in explained variance.

Each split on PLS can be interpreted as a candidate threshold, which is the point at which the slope of the inflation-slack relationship may change. Because random forests consist of hundreds of such trees, each fit on a slightly different sample and, by allowing different branches, they generate a distribution of split points. From this distribution, we take the median PLS split as our estimate of the threshold for that sub-sector-country combination. Intuitively, this aggregates the "votes" of many trees about where the relationship between labor shortages and inflation changes most sharply.

This procedure has a natural economic interpretation within our framework. When PLS is below the threshold, the Phillips curve is relatively flat, indicating that inflation responds little to variation in labor market slack. Once PLS rises above the threshold, the Phillips curve steepens and inflation becomes highly sensitive to a tighter labor market.

The advantage of this method is that it does not impose any a priori functional form, enabling the data to inform the choice of the threshold, which we then use in equation (4). To the best of our knowledge, the use of random forests to estimate thresholds, at least in a Phillips curve setting, is new in the literature[23].

In the Appendix we show an example of this procedure on simulated data that share similar characteristics with our empirical sample, and we also report summary statistics of the estimated thresholds by sub-sector and by country. Interestingly, the estimated thresholds highly correlate[24] with the average perceived labor shortages (PLS) reported over the sample period, as can be seen in Figure 4. Sectors and countries with persistently high perceived labor shortages tend to display higher thresholds, while those with looser labor markets feature lower thresholds. This tight relationship suggests that the estimated thresholds are not mere statistical artifacts but capture meaningful heterogeneity in the inflation process.

Non-Linear Results. Results are reported in Table 5 and show that non-linear component of the curve ($\alpha_d + \alpha_o$), which captures the effect of labor shortages when they exceed the sub-sector-country specific threshold, is large and highly significant across both methods. This suggests that inflation becomes much more sensitive to labor shortages in periods of acute labor market tightness, consistent with some of the recent evidence for the U.S.

The coefficient on the "linear" part of the curve (α_o), which captures the effect of labor

[23]For robustness, we report in the Appendix the results using an alternative approach employed by Gitti (2024).

[24]The correlation coefficients are equal to 0.93 at the sub-sector level and to 0.80 at the country level.

TABLE 5: Phillips Curve Estimates

	Inflation
PLS	0.047***
	(0.008)
PLS > PLS*	0.060***
	(0.012)
Constant	1.852***
	(0.096)
Observations	9445

Notes: fixed effects and controls outlined in Section 4.2 are included. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

shortages under normal labor market conditions, is positive and statistically significant. This indicates that even moderate labor tightness exerts upward pressure on inflation. At the same time, the magnitude of this effect is rather small and smaller than in periods of high shortages, suggesting a non-linear relationship such that inflationary pressures escalate as labor scarcity intensifies, while the response to increases in labor market tightness during slack periods remains more muted.

Viewed in a broader perspective, the estimates suggest that labor market conditions accounted for roughly 52% of the cumulative inflation observed between 2021Q1 and 2024Q4. A simple back-of-the-envelope calculation suggests that a linear model would have underpredicted inflation by about 0.58 percentage points *per quarter* over the 2021-2024 period.

Overall, these findings confirm that the relationship between labor shortages and inflation is not constant and varies depending on the tightness of the labor market. The presence of a significant non-linearity suggests that standard linear specifications of the Phillips curve may underestimate inflationary pressures when labor shortages become particularly pronounced. This has important implications for policymakers, as it indicates that inflationary risks may escalate more rapidly once labor market constraints reach critical levels. Pushing results further, they suggest that central bankers may be able to bring back inflation closer to the their target, without necessarily causing a recession.

5.3 Aggregate Implications

Until this point, the analysis focused on a sub-sector-country level, through the use of disaggregated data. However, one may question if the differences in labor market tightness across various sectors and countries could explain the aggregate inflation rate in the euro area. Before doing that, we have to take a few steps back and be more explicit about equation (4.2).

We follow the same logic of Hazell et al. (2022) and begin with a standard textbook NKPC, such as Eq. (1) in the introduction, and iterate it forward

$$\pi_{ijt} = \mathbb{E}_t \sum_{n=0}^{\infty} \beta^n \left(\kappa \, \hat{x}_{ij,t+n}\right) + \mathbb{E}_t \, \pi_{ij,t+\infty} + \omega_{ijt}, \tag{5}$$

While short-run inflation expectations vary across sub-sector-country pairs due to differing economic conditions, long-run expectations are shaped by shared beliefs about the monetary regime and thus do not vary for each pair in a monetary union. In panel regressions with common estimated components, long-run expectations will thus be captured.

These observation imply that that we can rewrite the previous equation as

$$\pi_{ijt} = \mathbb{E}_t \sum_{n=0}^{\infty} \beta^n \left(\kappa \, x_{ij,t+n}\right) + \lambda_{ij} F_t + \alpha_i + \tau_j + \omega_{ijt}, \tag{6}$$

Variation in $\mathbb{E}_t \, \pi_{ij,t+\infty}$ that is common across sub-sector-country pairs is captured by the vector of common factors F_t, while constant differences are absorbed by sub-sector and country fixed effects.

Building on the intuition above, we use the theoretical form of Eq. (4), borrowed from Gitti (2024) and iterated forward[25]

$$\begin{aligned} \pi_{ijt} = \mathbb{E}_t \, \pi_{ij,t+\infty} + \mathbb{E}_t \sum_{n=0}^{T} \beta^n \left(\kappa \, \hat{x}_{ij,t+n}\right) + \mathbb{E}_t \sum_{n=0}^{T} \beta^n \left[\kappa' \times \left(\hat{x}_{ij,t+n} - x^*_{ij,t+n}\right) \times \mathbb{1}\{x_{ijt} > x^*_{ij}\}\right] + \\ + \mathbb{E}_t \sum_{n=0}^{T} \beta^n \left(\delta \hat{\Gamma}_{ijt}\right) + \omega_{ijt}. \end{aligned} \tag{7}$$

[25] To clarify, this is just the texbook NKPC iterated forward in its piecewise linear version.

The current formulation is instrumental in showing that we have concentrated on estimating the coefficients α_o and α_d in equation (4), rather than focusing on κ and κ' in equation (7). However, to connect micro-level variations in labor shortages with the macroeconomic behavior of inflation, it is necessary to estimate these latter coefficients, as it will become clear below.

In fact, moving from estimating Eq. (4) to Eq. (7) poses certain empirical difficulties. In particular, the key difficulty arises from the two summations within the equation. Fortunately, Hazell et al. (2022) offers a solution by suggesting to truncate of the discounted sums on the right side to $T = 20$ (quarters)[26]. A major shortcoming of this approach is that the forward-looking sums in the previous equation extend 20 periods ahead and, therefore, using $T = 20$ would result in the loss of the final five years of our sample. To mitigate this issue, they suggest implementing a two-sample two-stage least squares (2SLS) approach. Specifically, the first-stage regression is estimated on a truncated sample that excludes the last five years, while the second stage is carried out using the full sample. We apply the two-sample 2SLS correction to the standard errors, as proposed by Chodorow-Reich and Wieland (2020), to ensure valid inference. Results are reported in Table 22 in the Appendix together with a discussion around the coefficients, the calibration and connected robustness checks.

Once κ and κ' are estimated, we can now understand whether the sub-sector-country results of the slope we found previously can help account for aggregate inflation dynamics in the euro area the sample period.

We begin with the Phillips curve expressed in its solved-forward form in equation (6). Assuming that PLS follows an AR(1) process enables a convenient derivation linking current PLS to the discounted sum of future PLS. However, this assumption may not accurately capture the true dynamics of PLS. To relax that, we now estimate scaling factors ζ and ζ' that empirically relate current PLS to its expected discounted path, using the following regressions

$$\sum_{n=0}^{T} \beta^n \, \hat{x}_{t+n} = \zeta \hat{x}_t + c + \nu_t \tag{8}$$

$$\sum_{n=0}^{T} \beta^n \, (\hat{x}_{t+n} - x^*_{t+n}) = \zeta'(\hat{x}_t - x^*_t) + c + \nu_t \tag{9}$$

[26] Hazell et al. (2022), with simulated data, show that this empirical procedure is able to accurately estimate the true value of the slope in this setting for a very wide range of its true values.

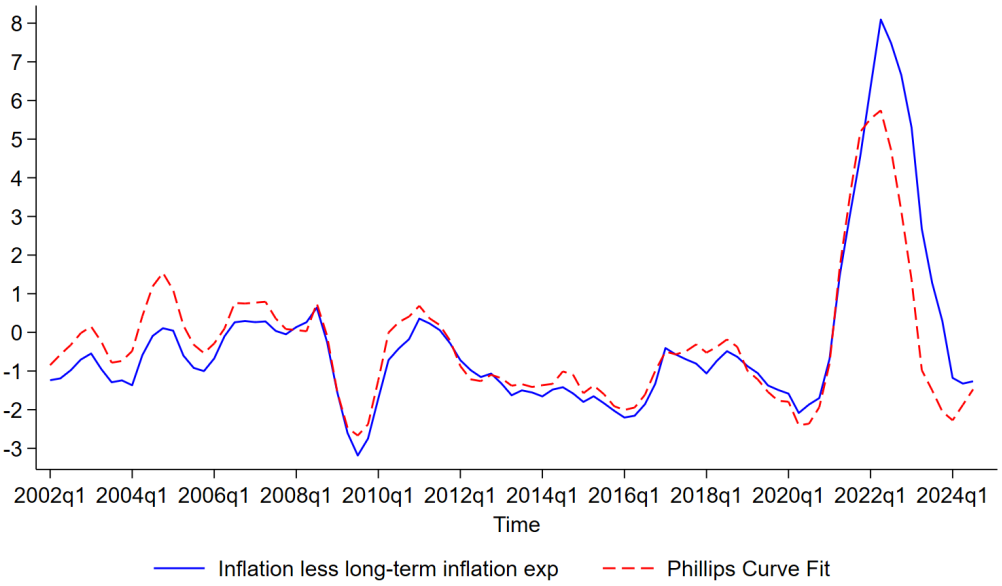
This approach yields an estimated ζ of 3.05 for the linear part of the Phillips curve and a larger ζ' of 8.16 for the non-linear part.

Using equations (8) and (9), we can rewrite equation (7) as

$$\pi_t - \mathbb{E}_t \, \pi_{t+\infty} = \kappa\zeta \, \hat{x}_t + \kappa'\zeta' \, (\hat{x}_t - x_t^*) \times \mathbb{1}\{x_t > x^*\} + \delta\Gamma_t + \lambda F_t + \epsilon_t, \tag{10}$$

where for Γ_t we applied the same logic of equation (8), and F_t is the common component[27]. At this point, it is evident why κ and κ' were estimated earlier; they are necessary, together with the scaling factors, to determine the aggregate coefficients in the previous equation. Note that equation (10) imposes a single kink at the aggregate level and is therefore a simplified representation of the sectoral piecewise linear relationships, which would otherwise depend on the distribution of sectoral slack.

FIGURE 5: Aggregate Phillips Curve



Notes: the figure shows the fit of the aggregate Phillips curve for PPI inflation. The blue line is the difference between PPI inflation and the 5-year ahead IMF's WEO forecasts which proxy inflation expectations. The red (dashed) line plots the fitted Phillips curve, based on the right hand side of equation (10), which is the variation in inflation predicted by our estimates.

[27] While for Table 22 we used the estimated common factor as explained previously, before estimating equation (10) we subtracted inflation expectations.

Figure 5 plots realized inflation, corresponding to the left-hand side of equation (10) (blue line), against the fitted values implied by its right-hand side terms (red dashed line). The figure shows that our Phillips curve regression tracks key inflation episodes, including the decline during the Great Financial Crisis, the prolonged disinflation following the European sovereign debt crisis and the sharp rise in inflation during the post-pandemic recovery. The full specification accounts around 80 percent of the variation in inflation ($R^2 = 0.78$), with a correlation coefficient of 0.88. Notably, PLS alone accounts for 55 percent of the variation in inflation ($R^2 = 0.55$), suggesting that labor market slack is an important determinant of inflation dynamics, together with supply, demand, financial and global factors, which are associated with an additional 23 percentage points to the explained variation.

5.4 Robustness Checks

We present a range of additional robustness checks to further assess the strength of our results, complementing those discussed in the main text. The main results remain valid across different variable selections and estimation strategies. This overall robustness is comforting, given the notorious sensitivity of Phillips curve estimates to different specifications and variable choices.

Sample. We examine whether the results hold when including the earlier part of the sample, prior to the introduction of the euro. In the Appendix, we report results using the full 1990Q1-2024Q4 sample and confirm that our main findings remain robust even with this extended period. We further show that the core findings of the paper continue to hold even when the post-pandemic period is controlled for.

Alternative Estimation. As a further effort to avoid potential endogeneity issues we use the one-quarter[28] lagged PLS as independent variable – a widely adopted approach in the Phillips curve's literature.

Supply Shortages. We employ an alternative indicator to capture supply chain disruptions. In the main estimation we used the percentage of firms reporting shortages of material and/or equipment as our proxy variable. However, from the BCS it is readily available an-

[28] Alternatively, two- and four- quarter lags lead to similar results and are reported in the Online Appendix B.2.

other measure of supply bottlenecks: capacity constraints[29]. These have been particularly pronounced in the European context, as recently underlined by Lane (2024).

Outliers. In an effort to mitigate the possibility that only a few sub-sector-country pairs drive the results, we re-estimate the main tables after having winsorized the observations at the 1% and 99% levels.

Thresholds. In order to confirm the robustness of the results regarding the choice of the kinks, instead of relying on the random forest approach to detect them, we follow a more traditional approach as outlined by Gitti (2024). In particular, for values of PLS between 10 and 90 and equally spaced by 5 percentage points, we regress inflation on PLS and its interaction with a dummy that takes the value of 1 when PLS exceeds the value we are considering. We then compute the root mean squared error (RMSE) for each regression and select the value of PLS that minimizes the RMSE in each sub-sector-country pair. Despite this different identification procedure, which imposes some structure to find the kinks, we confirm the baseline results.

6 Conclusion

This paper introduces the *Perception of Labor Shortages* (PLS), a new *granular* proxy for labor market tightness in the euro area. Using an interactive fixed effects regression analysis, we show that explicitly estimating latent common factors and allowing their effects to vary across sub-sectors and countries is crucial, as conventional specifications that ignore these points can lead to substantially biased estimates of the slope of the Phillips curve.

In line with theoretical predictions, PLS is positively associated with wages, supporting its interpretation as reflecting firms' marginal labor costs. Beyond linear effects, we document pronounced non-linearities in the Phillips curve. A machine-learning approach reveals kinks in the relationship between labor shortages and inflation, with inflation responding only modestly under normal conditions but becoming sharply more sensitive once shortages exceed sector-specific thresholds.

Aggregating these effects, we find that our specification accounts for the main euro area

[29] Data come as time series for each country-sub-sector pair and it is the percentage from the answer to the following question: "*At what capacity is your company currently operating (as a percentage of full capacity)? The company is currently operating at X % of full capacity*".

inflation episodes at the macro level, including the fall during the Great Financial Crisis, the extended disinflation after the EU sovereign debt crisis and the sharp post-pandemic surge. This confirms that sub-sector-country movements in perceived labor shortages are also quantitatively important for explaining aggregate inflation in the euro area.

Overall, our results highlight the importance of granular labor market indicators and flexible empirical methods for capturing both the level and the state-dependent effects of slack and they suggest that firms' hiring difficulties provide a meaningful signal of inflationary pressures.

References

Afrouzi, H., A. Blanco, A. Drenik, and E. Hurst (2024). A theory of how workers keep up with inflation. Technical report, National Bureau of Economic Research.

Ahn, S. C. and A. R. Horenstein (2013). Eigenvalue ratio test for the number of factors. *Econometrica 81*(3), 1203–1227.

Albuquerque, B. and U. Baumann (2017). Will us inflation awake from the dead? the role of slack and non-linearities in the phillips curve. *Journal of Policy Modeling 39*(2), 247–271.

Ampudia, M., M. Lombardi, and T. Renault (2024). *The wage-price pass-through across sectors: evidence from the euro area.* Number 2948. ECB Working Paper.

Auer, R., M. Pedemonte, and R. Schoenle (2024). Sixty years of global inflation: a post-gfc update.

Babb, N. R. and A. K. Detmeister (2017). Nonlinearities in the phillips curve for the united states: Evidence using metropolitan data.

Ball, L., D. Leigh, and P. Mishra (2022). Understanding us inflation during the covid-19 era. *Brookings Papers on Economic Activity 2022*(2), 1–80.

Barnichon, R. and A. H. Shapiro (2024). Phillips meets beveridge. *Journal of Monetary Economics 148*, 103660.

Beaudry, P., C. Hou, and F. Portier (2025). On the fragility of the nonlinear phillips curve view of recent inflation. Technical report, National Bureau of Economic Research.

Benati, L. (2008). Investigating inflation persistence across monetary regimes. *The Quarterly Journal of Economics 123*(3), 1005–1060.

Benigno, P. and G. B. Eggertsson (2023). It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve. Technical report, National Bureau of Economic Research.

Beschin, A., J. Paredes, G. Polichetti, and T. Renault (2025). The slope of the euro area price phillips curve: evidence from regional data.

Bonam, D., J. de Haan, and D. Van Limbergen (2021). Time-varying wage phillips curves

in the euro area with a new measure for labor market slack. *Economic Modelling 96*, 157–171.

Bonthuis, B., V. Jarvis, and J. Vanhala (2016). Shifts in euro area beveridge curves and their determinants. *IZA Journal of Labor Policy 5*(1), 20.

Byrne, D. and Z. Zekaite (2018). Missing wage growth in the euro area: is the wage philips curve non-linear? Technical report, Central Bank of Ireland.

Chodorow-Reich, G. and J. Wieland (2020). Secular labor reallocation and business cycles. *Journal of Political Economy 128*(6), 2245–2287.

Ciccarelli, M. and B. Mojon (2010). Global inflation. *The Review of Economics and Statistics 92*(3), 524–535.

Ditzen, J. and S. Reese (2023). xtnumfac: A battery of estimators for the number of common factors in time series and panel-data models. *The Stata Journal 23*(2), 438–454.

Eser, F., P. Karadi, P. R. Lane, L. Moretti, and C. Osbat (2020). The phillips curve at the ecb. *The Manchester School 88*, 50–85.

Faber, M. and G. Züllig (2025). Kinky europe: Evidence from the regional phillips curve in the euro area.

Fitzgerald, T., C. Jones, M. Kulish, and J. P. Nicolini (2024). Is there a stable relationship between unemployment and future inflation? *American Economic Journal: Macroeconomics 16*(4), 114–142.

Florio, A., D. Siena, and R. Zago (2025). Global value chains and the phillips curve: A challenge for monetary policy. *European Economic Review 174*, 104966.

Forbes, K. J. (2019). Has globalization changed the inflation process?

Furlanetto, F. and A. Lepetit (2025). The slope of the phillips curve. In *Research Handbook on Inflation*, pp. 167–185. Edward Elgar Publishing.

Gagliardini, P., E. Ossola, and O. Scaillet (2019). A diagnostic criterion for approximate factor structure. *Journal of Econometrics 212*(2), 503–521.

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

Gitti, G. (2024). Nonlinearities in the regional phillips curve with labor market tightness.

Gitti, G. and A. Cerrato (2025). The return of the phillips curve: Evidence from us cities. *Kilts Center at Chicago Booth Marketing Data Center Paper*.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2022). The slope of the phillips curve: evidence from us states. *The Quarterly Journal of Economics 137*(3), 1299–1344.

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

Im, K. S., M. H. Pesaran, and Y. Shin (2003). Testing for unit roots in heterogeneous panels. *Journal of econometrics 115*(1), 53–74.

Kiss, A., A. Turrinil, and A. Vandeplas (2022). Ii. slack vs tightness in euro area labour markets: growing mismatch after covid-19? *Quarterly Report on the Euro Area 2*(19), 19–28.

Lane, P. (2024). The 2021-2022 inflation surges and the monetary policy response through the lens of macroeconomic models. *SUERF Policy Note 364*.

Levy, A. (2019). The euro area phillips curve: identification with regional data. *Unpublished manuscript, Department of Economics, Massachusetts Institute of Technology 9*.

Lo Bello, S. and E. Viviano (2024). Some considerations on the phillips curve after the pandemic. *Bank of Italy Occasional Paper* (842).

Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *American Economic Journal: Journal of Economic Literature 52*(1), 124–188.

McLeay, M. and S. Tenreyro (2020). Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual 34*(1), 199–255.

Miranda, K., P. Poncela, and E. Ruiz (2022). Dynamic factor models: Does the specification matter? *SERIEs 13*(1), 397–428.

Mumtaz, H., S. Simonelli, and P. Surico (2011). International comovements, business cycle and inflation: A historical perspective. *Review of Economic Dynamics 14*(1), 176–198.

Murphy, A. (2017). Is the us phillips curve convex? some metro level evidence. *Manuscript*.

Nalewaik, J. (2016). Non-linear phillips curves with inflation regime-switching.

Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *The Review of Economics and Statistics 92*(4), 1004–1016.

Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861-1957. *economica 25*(100), 283–299.

Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press.

Reis, R. (2025). Why did inflation rise and fall in 2021-24? channels and evidence from expectations. *LSE manuscript*.

Smith, S. C., A. Timmermann, and J. H. Wright (2025). Breaks in the phillips curve: Evidence from panel data. *Journal of Applied Econometrics 40*(2), 131–148.

A Appendix

A.1 List of Manufacturing Sub-Sectors

Table 6: Manufacturing Sub-Sectors

Sector definition	**NACE 2-digit Codes**	**Sub-sectors**
Food products, beverages, and tobacco products	10 – 12	Food products, Beverages, Tobacco products
Textiles, apparel, leather, and related products	13 – 15	Textiles, Wearing apparel, Leather and related products
Wood and paper products, and printing	16 – 18	Wood products, Paper and paper products, Printing and recorded media
Chemicals and chemical products	20	Chemicals and chemical products
Rubber and plastics products, and other non-metallic mineral products	22 – 23	Rubber and plastic products, Other non-metallic mineral products
Basic metals and fabricated metal products	24 – 25	Basic metals, Fabricated metal products
Computer, electronic, and optical products	26	Computer, electronic and optical products
Electrical equipment	27	Electrical equipment
Machinery and equipment n.e.c.	28	Machinery and equipment n.e.c.
Transport equipment	29 – 30	Motor vehicles, trailers and semi-trailers, Other transport equipment
Other	19, 21, 31, 32, 33	Manufacture of coke and refined petroleum products, Manufacture of basic pharmaceutical products and pharmaceutical preparations, Manufacture of furniture, Other manufacturing, Repair and installation of machinery and equipment

A.2 Data Appendix

In this section we describe the sources of our data and any transformation applied.

Our primary data source is the Joint Harmonised EU Programme of Business and Consumer Surveys, administered by the Directorate-General for Economic and Financial Affairs of the European Commission[30].

These surveys are conducted at the national level by various EU Commission's partner institutes, including ministries, statistical offices, central banks, research institutes, business associations and private companies. The program follows a harmonized methodology to ensure comparability across countries.

The survey samples are taken out from national frames intended to capture the entire relevant population. These frames are derived from official or statistical registers, as well as membership lists from business associations and chambers of commerce. Given the dynamic nature of business populations – due to mergers, bankruptcies and the entry of new firms – the sample frames are regularly updated to maintain representativeness.

Survey samples are designed to ensure efficiency and precision. Many national institutes employ stratified random sampling, which divides the population into non-overlapping subgroups (strata) with similar variance concerning key variables. Business survey stratification is typically based on firm size and sector of activity.

The Business and Consumer Surveys operate under two fundamental principles to achieve harmonization across countries:

- all national institutes use the same harmonized questionnaires;
- surveys are conducted and results transmitted according to a common timetable.

Survey timing is consistent across countries. Questionnaires are administered in the first two to three weeks of each month for the monthly surveys and quarterly ones are carried out during the first month[31] of each quarter (January, April, July, and October).

The panel structure of business surveys ensures consistency over time. In most cases, companies included in the sample remain in the survey for multiple periods, with periodic updates

[30] The following description is based on the official methodological user guide of the Commission, available here bcs_user_guide.pdf – for a more detailed overview.

[31] Always around the first two to three weeks.

or partial replacements (rotating panels). This approach improves response speed, reduces variability between successive surveys and optimizes sample size.

We apply the following data-cleaning steps to address missing values and outliers. First, we focus on those series for which we have at least (i) half of the sample time frame available and (ii) that are also different from zero. This ensures that our sample includes countries and sub-sectors for which we have a reasonable amount of data and that experienced hiring constraints. Second, we exclude values below zero and above one-hundred, as these are clearly reporting issues, as these variables are in percentage terms and cannot be negative or above one-hundred.

Due to the relatively high correlation between PLS and supply shortages, before estimating the main results we decided to orthogonalize supply shortages having it regressed on PLS and using the residuals in the regressions. This ensures that only the variation in supply shortages, independent of PLS, is used. We use capacity constraint as an alternative measure to supply shortages and confirm the main results.

The second set of variables we employ come from different sources. The producer price index is downloaded directly from Eurostat at quarterly frequency both in percentage change compared to same period in previous year[32] and as an index, with base year 2021. The number of vacancies and unemployed people are also download from Eurostat. In order to adjust for the difference in the measurement scale of the two variables, we divide vacancies by one-thousand. Once the ratio is computed, we apply the natural logarithm transformation. Long-term inflation expectations are 5-year ahead forecast of inflation by the IMF. These forecasts are released twice a year, hence we follow Forbes (2019) and treat those released in the spring's World Economic Outlook (WEO) as Q1, and those released in the fall's WEO as Q3. The remaining two quarters (Q2 and Q4) are linearly interpolated between the nearest spring and fall forecasts. Energy prices come also from the IMF and are available through their commodity data portal, while the nominal effective exchange rate are download from the BIS data portal.

Variables for the U.S. come from two sources. The first one is the National Federation of Independent Business' *Small Business Economic Trends* indicators, which are publicly available. We downloaded and used the series with current job openings, which gives the percent of respondents with a position they are currently not able to fill (at the national level). This

[32]The same applies for wage inflation.

measure is indeed close to the PLS one we employed from the EU Commission's Business and Consumer Surveys. Regarding the vacancies and unemployment data for the U.S. (necessary to compute the V/U ratio), we use the same data as in Benigno and Eggertsson (2023). We thank the authors for kindly sharing their data.

A.3 Summary Statistics

Table 7 presents summary statistics for the key variables over the full sample period, as well as separately for the pre-COVID and post-COVID periods, given the scale of the pandemic shock and to reflect the sharp discontinuity it introduced.

TABLE 7: Summary Statistics

	Full Sample		Pre-Covid		Post-Covid	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Inflation	1.483	2.120	0.891	0.915	3.760	3.533
PLS	8.369	6.414	5.954	3.476	17.649	6.715
Supply Shortages	11.353	11.760	7.045	3.421	27.905	16.968
Insufficient Demand	29.967	9.049	30.584	8.829	27.598	9.733
Financial Constraints	3.893	1.094	3.989	1.166	3.524	0.656
WEO Inflation Expectations	1.815	0.158	1.805	0.171	1.856	0.070
Energy Prices	-0.000	57.086	-0.869	43.937	3.337	93.480
NEER	-0.000	0.108	0.001	0.099	-0.004	0.140
Other Constraints	7.148	3.925	5.860	0.986	12.098	6.434
Wage Inflation	2.526	1.225	2.394	0.817	3.035	2.138

Notes: the table reports basic summary statics (mean and standard deviation) for the full sample (2002Q1-2024Q4), the pre-Covid period (2002Q1-2020Q1) and the post-Covid period (2020Q1-2024Q4). When aggregating, we compute weighted averages based on official weights reported by the EU commission.

A.4 Illustration of the Random Forest Threshold Detection Procedure

This appendix gives a step-by-step explanation of the schematic example reported in the Figure 6 below. The purpose of the illustration is pedagogical: it shows that the proposed data-driven procedure recovers a meaningful threshold in a simple data generating process that mimics relevant features of the empirical panels.

For this exercise, we generate a right-skewed distribution for the perceived labor shortages using a Gamma draw in order to have a distribution similar to the one observed in the empirical data. Moreover, to mimic the empirical feature that a non-negligible fraction of panels report zero or near-zero PLS, we force 10% of observations to be exactly zero. The true threshold is set to be equal to 8.5. Inflation is generated as the piecewise linear function of PLS

$$y = \begin{cases} 2.0 + 0.04 \cdot \text{PLS} + \varepsilon, & \text{if PLS} < 8.5 \\ 2.0 + 0.04 \cdot 8.5 + 0.107 \cdot (\text{PLS} - 8.5) + \varepsilon, & \text{if PLS} \geq 8.5 \end{cases}$$

where $\varepsilon \sim N(0, 0.5^2)$. We fit a random forest similar to what we estimate in the data with each tree fitted to a bootstrap sample of the data and grows by recursive partitioning. A single decision tree searches over candidate splits of the form PLS $<$ PLS* (for many candidate PLS* values drawn from observed PLS) and selects the split that maximizes the reduction in predictive error. The chosen split is applied and the procedure repeats within each resulting node. Every internal node that splits on PLS provides a candidate split point (PLS*) that marks a potential change in the slope of the inflation–PLS relationship.

After fitting the forest we iterate over every tree and every internal node. For nodes that split on the PLS feature we collect the corresponding threshold value. We will have a vector therefore that contains all split points on PLS found across all trees in the forest. Because each tree is grown on a different bootstrap sample and can split the predictor space differently, the vector contains a sample of plausible threshold candidates.

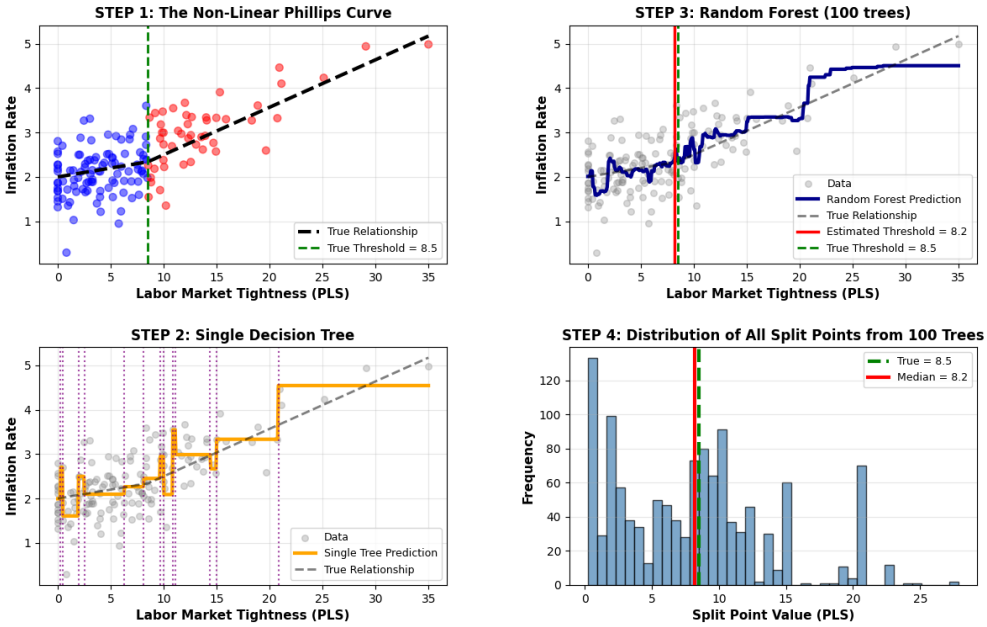
Finally, we take the median of that vector of candidate thresholds as the point estimate of the threshold for the example panel. The median is robust to outlying splits (for instance splits near zero that arise because many observations concentrate there) and corresponds to

the majority "vote" of the ensemble about where the relationship changes.

The four panels in Figure 6 visualize and schematize this process. In this illustrative example the forest yields a median split close to the true break (median close to 8.2 vs true 8.5) and the ensemble prediction reproduces the piecewise pattern. The single tree is noisier and its many candidate split points (purple lines) illustrate why an ensemble vote is preferable. The histogram in Panel 4 provides a direct, data-driven way to communicate both the point estimate (median). The exact same extraction and aggregation steps are applied to each empirical sub-sector-country panel in the dataset.

Results and summary statistics for the full empirical sample (summary stats by sub-sector and by country for thresholds) are provided in Appendix Table 8 and plotted in the two scatter plots in the main text.

FIGURE 6: Step-by-Step Random Forest Procedure on Simulated Data



Notes: the figure show the procedure used to estimate thresholds for each seub-sector-country pair.

A.5 Thresholds' Summary Stats by Country and Sub-Sector

TABLE 8: Summary Statistics - Thresholds

	Mean	SD	Min	Max	N
Austria	7.278	3.267	3.40	16.75	1564
Belgium	11.970	7.666	1.50	34.60	1840
Germany	8.812	2.896	4.10	14.45	1840
Greece	10.775	9.516	1.85	26.55	368
Spain	3.690	2.783	1.05	11.25	1380
Finland	13.840	7.927	4.45	38.00	1380
France	8.275	5.129	2.30	25.95	2024
Italy	2.435	0.878	1.15	4.25	1840
Netherlands	6.412	2.284	3.45	12.60	1840
Portugal	9.947	6.126	3.40	28.75	1564
Food prod.	7.540	4.633	1.15	16.05	920
Beverages	2.825	1.000	1.50	4.10	368
Textiles	8.050	4.179	2.00	13.00	828
Wearing apparel	5.069	2.630	1.55	11.05	736
Leather prod.	8.425	4.659	2.45	16.65	552
Wood products	7.250	1.961	3.50	10.25	828
Paper prod.	5.883	2.283	2.75	10.15	552
Printing	7.542	4.344	2.20	14.15	552
Chemicals	4.219	1.613	1.40	6.50	736
Pharma	5.079	2.209	1.40	8.75	644
Plastics	7.083	4.425	1.85	15.50	828
Non-metallic min.	6.294	3.769	1.05	14.55	828
Basic metals	5.717	2.818	2.40	10.50	828
Fabricated metals	11.620	5.065	3.75	21.10	920
Electronics	8.357	3.102	2.00	11.40	644
Electrical eq.	7.972	3.844	1.65	14.50	828
Machinery	11.786	6.186	3.15	24.50	828
Motor vehicles	5.933	5.158	1.15	19.00	828
Transport eq.	14.900	13.226	1.85	38.00	552
Furniture	11.350	7.731	3.00	28.75	736
Other manuf.	7.979	4.684	2.50	17.45	644
Machinery repair	18.870	10.525	4.25	34.60	460

Notes: the table reports basic summary statics (mean, standard deviation, minimum and maximum values, and number of observations) for the thresholds, aggregated in two different ways: the upper part of tables shows the stats by country, while the bottom part by sub-sector.

A.6 Im, Pesaran and Shin (2002) Unit Root Test

Table 9: Unit Root Test Results

Variable	Z-stat	p-value
Inflation	-13.418	0.000
PLS	-34.163	0.000
Supply Shortages	-48.012	0.000
Insufficient Demand	-37.850	0.000
Other Constraints	-54.299	0.000
Financial Constraints	-47.626	0.000
WEO Inflation Expectations	-5.315	0.000
Energy Prices	-27.042	0.000
NEER	-99.050	0.000

Notes: In the Im et al. (2003) unit root test the null hypothesis is that all panels contain unit roots, while the alternative hypothesis is that at least some panels are stationary.

A.7 Collinearity Diagnostics

Table 10: Collinearity Diagnostics

Variable	VIF	SQRT VIF	Tolerance	R-Squared
Inflation	1.23	1.11	0.8131	0.1869
WEO Inflation Expectations	1.08	1.04	0.9278	0.0722
PLS	1.14	1.07	0.8807	0.1193
Energy Prices	1.19	1.09	0.8439	0.1561
Supply Shortages	1.01	1.01	0.9875	0.0125
Insufficient Demand	1.12	1.06	0.8966	0.1034
Other Constraints	1.02	1.01	0.9825	0.0175
Financial Constraints	1.07	1.03	0.9365	0.0635
NEER	1.00	1.00	0.9985	0.0015
Relative Prices	1.01	1.00	0.9925	0.0075
Mean VIF	1.09			

Notes: Collinearity diagnostics report variance inflation factors (VIF), their square roots, tolerance values, and the corresponding R^2. The mean VIF is 1.09, indicating the absence of severe multicollinearity in our regressions.

A.8 Common Component Estimation

Consider a model similar to the one estimated with Method 1[33]:

$$\pi_{ijt} = \alpha\, x_{ijt} + \lambda_{ij} F_t + w_{ijt}.$$

Ahn and Horenstein (2013) suggest two estimators to estimate the number of common factors F: the eigenvalue ratio (ER) and growth rate (GR) estimators. The estimators for a given number of common factors k

$$\begin{aligned} \mathrm{ER}(k) &= \frac{\widetilde{\mu}_{NT,k}}{\widetilde{\mu}_{NT,k+1}}, \quad k = 1, 2, \ldots, k_{\max} \\ \mathrm{GR}(k) &= \frac{\ln\left(1 + \widetilde{\mu}^*_{NT,k}\right)}{\ln\left(1 + \widetilde{\mu}^*_{NT,k+1}\right)} \end{aligned}$$

with $\widetilde{\mu}_{NT,k}$ being the kth largest eigenvalue of $XX/(NT)$ and $\widetilde{\mu}^*_{NT,k} = \widetilde{\mu}_{NT,k}/V(k)$, where $V(k)$ is the mean of the squared residuals of a regression of X on the first k principal components of $XX/(NT)$. The number of common factors is then

$$\begin{aligned} \widetilde{k}_{\mathrm{ER}} &= \max_{1 \le k \le k_{\max}} \mathrm{ER}(k) \\ \widetilde{k}_{\mathrm{GR}} &= \max_{1 \le k \le k_{\max}} \mathrm{GR}(k). \end{aligned}$$

Onatski (2010) proposes identifying the number of factors by examining the gaps between the ordered eigenvalues of $X'X/T$. He notes that eigenvalues beyond the true number of factors tend to cluster around a finite value, creating a threshold δ. The method selects the smallest eigenvalue among the top $k_{\max}$ whose gap with the next smaller eigenvalue exceeds δ, as this indicates the presence of a genuine factor.

The following algorithm is used:

1. Obtain the eigenvalues $\mu_1, \mu_2, \ldots, \mu_N$ of $T^{-1}\sum_{t=1}^{T} X_t X_t'$. Set $j = k_{\max} + 1$.
2. Estimate the slope coefficient $\hat{\beta}$ in a linear regression of $\mu_j, \ldots, \mu_{j+4}$ on $(j-1)^{2/3}, \ldots, (j+3)^{2/3}$ and a constant, and set the threshold δ to $\delta = 2|\hat{\beta}|$.

[33] This section is based on Ditzen and Reese (2023).

3. Obtain the estimated number of factors as $\hat{r}(\delta) = \max(i \leq r_{\max}^N : \mu_i - \mu_{i+1} \geq \delta)$ if there exists some $i \leq r_{\max}^N$ subject to $\mu_i - \mu_{i+1} \geq \delta$. Otherwise, set $\hat{r}(\delta) = 0$.

4. Set $j = \hat{r}(\delta)$. Repeat steps 2 and 3 until $\hat{r}(\delta)$ remains unchanged in two subsequent repetitions. Let this final estimator be denoted $\tilde{r}_{\text{ED}}$.

Gagliardini et al. (2019) propose a postestimation method based on residuals from linear models to distinguish between weak cross-sectional dependence and a factor structure. The approach compares eigenvalues of the residual covariance matrix to a penalty term that depends on N and T. If the penalty exceeds zero, at least one factor is present. The number of factors is determined as the largest k such that the difference between the k-th largest eigenvalue and the penalty becomes negative for the first time.

The difference between the k-th largest eigenvalue, $\mu(k)$, and the penalty $g(N,T)$ is

$$\begin{aligned} \xi(k) &= \mu(k) - g(N,T) \\ g(N,T) &= \frac{(\sqrt{N}+\sqrt{T})^2}{NT} \ln\left\{\frac{NT}{(\sqrt{N}+\sqrt{T})^2}\right\} \end{aligned} \tag{11}$$

The number of factors is then selected according to

- no factors: $\xi(1) < 0$
- $\widehat{k}$ factors: $\widehat{k} = \min\{0, \ldots, T-1 : \xi(k) < 0\}$

If $\xi(k_i) > 0 \; \forall \; k_i = 0, \ldots, T-1$, then $\widehat{k} = T$.

Across all methods considered, we find strong evidence in favour of a single component (F, in our baseline equation). This result aligns with the literature that emphasizes the presence of inflation co-movements and which suggests that one factor alone explains a large proportion of cross-country inflation variation (Ciccarelli and Mojon, 2010; Auer et al., 2024).

This result lends empirical support to the view that inflation is shaped by common shocks, which we interpret as potentially related to (i) shared inflation expectations dynamics within the monetary union and (ii) aggregate shocks.

A.9 Robustness Checks

Sample Period: 1990-2024

TABLE 11: Full Sample Phillips Curve Estimates: 1990-2024

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.149***	0.102***	0.127***	0.026***	0.041***	0.003
Constant		1.092***	1.450***	2.677***	1.222***	-0.337
		(0.052)	(0.229)	(0.674)	(0.159)	(0.545)
Observations	16950	16950	16950	16950	10378	10377

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 12: Full Sample Phillips Curve Estimates: 1990-2024

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.060***	0.066***	0.079***	0.067***
	(0.003)	(0.003)	(0.004)	(0.004)
Constant	0.755***	1.384***	1.623***	1.162***
	(0.044)	(0.036)	(0.092)	(0.092)
Observations	16950	16950	10643	10642

Notes: Column (7) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (8), control variables 98), and a proxy for short-term inflation expectations (10). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Controlling for Post-Covid Inflation

TABLE 13: Phillips Curve Estimates: 2002-2024

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.151***	0.107***	0.072***	0.031***	0.030***	0.006
Constant		1.042***	1.276***	0.120	1.145***	-0.458
		(0.052)	(0.219)	(0.381)	(0.155)	(0.364)
Observations	14399	14399	14399	14399	9095	9094

Notes: A post-2021Q1 dummy variable is included to account for inflation dynamics associated with the period following the Covid pandemic. Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 14: Phillips Curve Estimates: 2002-2024

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.039***	0.043***	0.056***	0.050***
	(0.003)	(0.003)	(0.004)	(0.004)
Constant	1.627***	1.590***	1.978***	1.581***
	(0.036)	(0.038)	(0.095)	(0.101)
Observations	14399	14399	9445	9444

Notes: A post-2021Q1 dummy variable is included to account for inflation dynamics associated with the period following the Covid pandemic. Column (7) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (8), control variables (9), and a proxy for short-term inflation expectations (10). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Lagged PLS

TABLE 15: Phillips Curve Estimates: Lagged PLS

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.149***	0.101***	0.125***	0.037***	0.035***	0.002
	(0.004)	(0.005)	(0.010)	(0.008)	(0.006)	(0.007)
Constant		1.122***	1.610***	0.426	1.257***	-0.445
		(0.053)	(0.254)	(0.339)	(0.165)	(0.355)
Observations	14245	14245	14245	14245	9035	9034

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 16: Phillips Curve Estimates: Lagged PLS

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.067***	0.075***	0.091***	0.078***
	(0.003)	(0.004)	(0.005)	(0.005)
Constant	0.641***	1.342***	1.783***	1.261***
	(0.041)	(0.039)	(0.096)	(0.097)
Observations	14245	14245	9327	9326

Notes: Column (7) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (8), control variables (9), and a proxy for short-term inflation expectations (10). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Capacity Constraints

TABLE 17: Phillips Curve Estimates: Capacity Constraints

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.151***	0.107***	0.133***	0.031***	0.038***	-0.001
Constant		1.042***	1.513***	0.120	-2.797***	-2.959***
		(0.052)	(0.261)	(0.381)	(0.954)	(0.892)
Observations	14399	14399	14399	14399	9224	9223

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 18: Phillips Curve Estimates: Capacity Constraints

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.062***	0.070***	0.089***	0.074***
	(0.003)	(0.003)	(0.005)	(0.004)
Constant	0.688***	1.359***	0.535	0.850
	(0.040)	(0.038)	(0.582)	(0.553)
Observations	14399	14399	9586	9585

Notes: Column (7) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (8), control variables (9), and a proxy for short-term inflation expectations (10). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Outliers

Table 19: Phillips Curve Estimates: Outliers

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.160***	0.116***	0.145***	0.036***	0.043***	0.003
	(0.004)	(0.005)	(0.011)	(0.008)	(0.006)	(0.007)
Constant		0.978***	1.410***	0.095	1.188***	-0.449
		(0.053)	(0.265)	(0.382)	(0.162)	(0.353)
Observations	14399	14399	14399	14399	9095	9094

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 20: Phillips Curve Estimates: Outliers

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.067***	0.076***	0.089***	0.075***
	(0.003)	(0.004)	(0.005)	(0.004)
Constant	0.660***	1.315***	1.700***	1.221***
	(0.041)	(0.039)	(0.097)	(0.098)
Observations	14399	14399	9445	9444

Notes: Column (7) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (8), control variables (9), and a proxy for short-term inflation expectations (10). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Thresholds

TABLE 21: Phillips Curve Estimates: RMSE

	Inflation
PLS	0.079***
	(0.005)
PLS > PLS*	0.038**
	(0.015)
Constant	1.747***
	(0.096)
Observations	9445

Notes: fixed effects and controls outlined in Section 4.2 are included. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.10 Estimates of κ and κ'

First of all, to isolate cyclical fluctuations in our PLS measure ($\hat{x}$), we follow Barnichon and Shapiro (2024) and remove panel-specific trends by fitting a linear time trend for each sub-sector-country pair and taking the residuals as our cyclical PLS. Intuitively, it serves as a gap measure similar to the unemployment gap, reflecting transitory departures driven by business cycle dynamics[34].

TABLE 22: Phillips Curve Estimates: Aggregate

	(1)
κ	0.035^{***}
	(0.002)
κ'	0.006^{***}
	(0.001)

Notes: the outcome variable is PPI inflation. The regressors are discounted future sums of quarterly sub-sector-country PLS. We estimate the coefficients by two-sample two stage least squares and apply the correction to our standard errors from Chodorow-Reich and Wieland (2020). We include sub-sector-country fixed effects and include the estimated common factor. Corrected standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

At first glance, we observe that the absolute size of the sub-sector-country coefficients are (much) larger than the ones just estimated. This reflects the fact that in equation (4) the contemporaneous value for each variable is standing in for the entire future sum in equation (7). Since PLS is quite persistent (the same line of reasoning goes for the other variables), time variation in the future sum is much larger than time variation in PLS, which results in a much larger coefficient in equation (4) than in equation (7). It is perhaps important to recall that the slope of the non-linear part of the curve is equal to the sum of κ and κ' (i.e., the sum of the two coefficients in the first two rows of the table).

In our baseline specification, we calibrate the discount factor at $\beta = 0.99$. This calibration presumes a high degree of forward-looking behavior in firms' pricing decisions. However, a growing body of research suggests that firms may be substantially less forward-looking than this benchmark implies. To assess the sensitivity of our results, Table 23 reports also estimates of PLS for both the linear and the non-linear part of the curve under alternative

[34] Unlike unemployment, PLS has no well-defined "natural" level, so this residual captures deviations from each panel's long-run trajectory.

calibrations of β; results are barely affected by this choice.

TABLE 23: Phillips Curve Estimates: Structural

	$\beta = 0.90$	$\beta = 0.95$	$\beta = 0.99$
κ	0.035^{***}	0.034^{***}	0.035^{***}
	(0.002)	(0.002)	(0.002)
κ'	0.005^{***}	0.006^{***}	0.006^{***}
	(0.001)	(0.001)	(0.001)

Notes: the outcome variable is PPI inflation. The regressors are discounted future sums of quarterly sub-sector-country PLS. We estimate both coefficients by two-sample two stage least squares, and apply the correction to our standard errors from Chodorow-Reich and Wieland (2020). We include sub-sector-country fixed effects and include the estimated common factor. Corrected standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

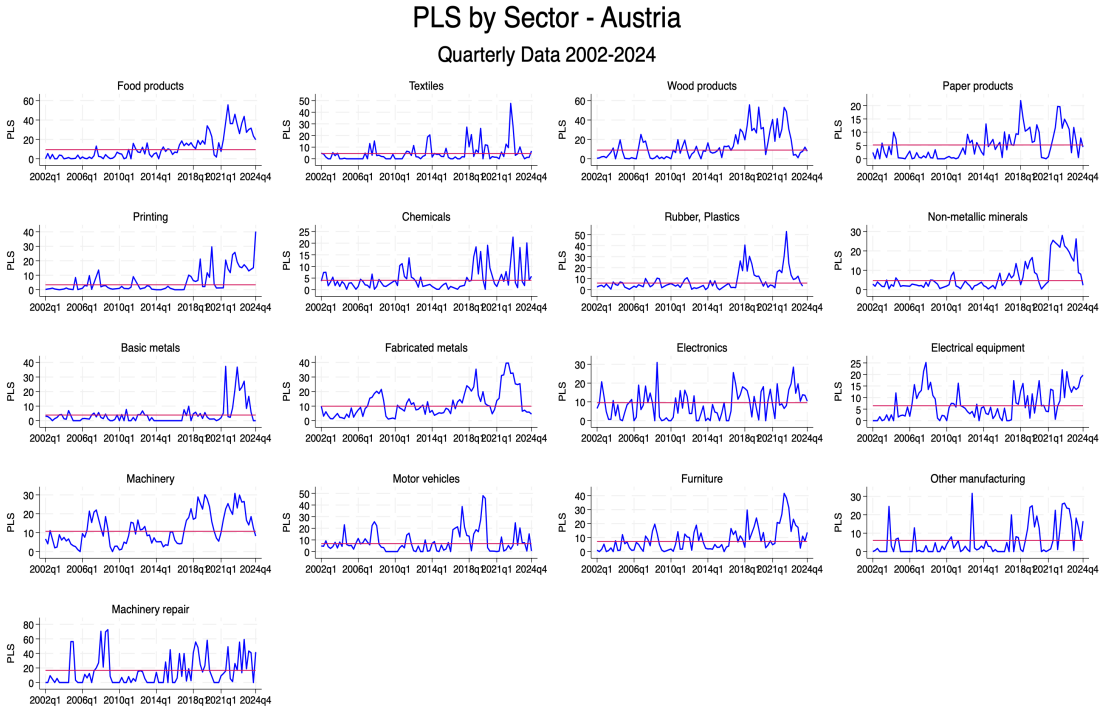
Online Appendix to "Labor Market Shortages and the Phillips Curve"

Cristian Arcidiacono

B Online Appendix

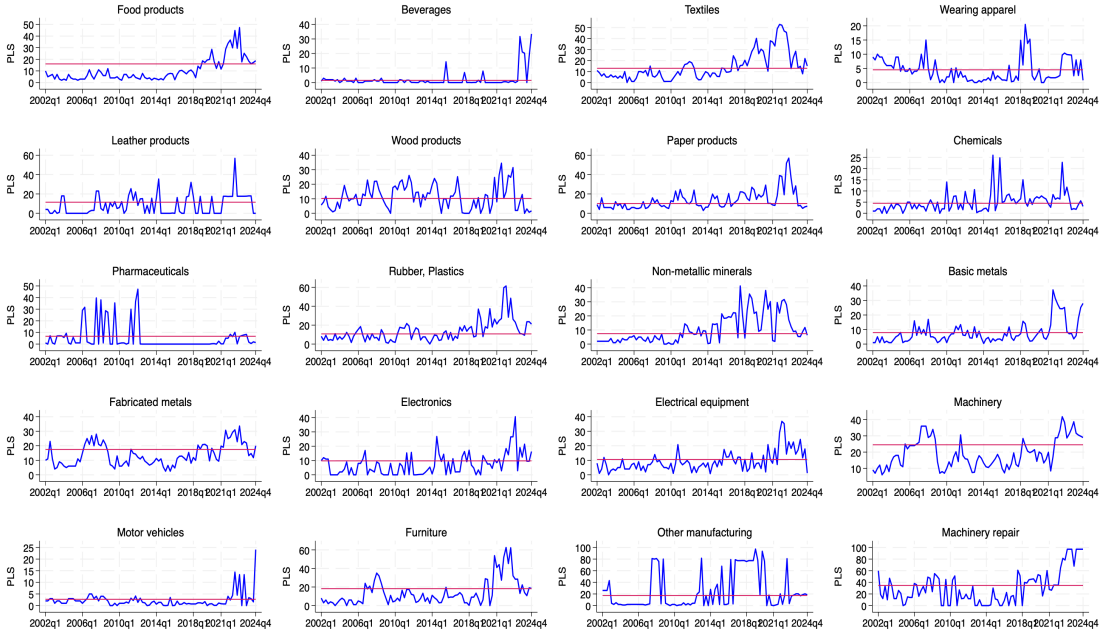
B.1 Time Series of PLS by Country and Sub-Sector

Together with the time series of PLS we plot the corresponding threshold.



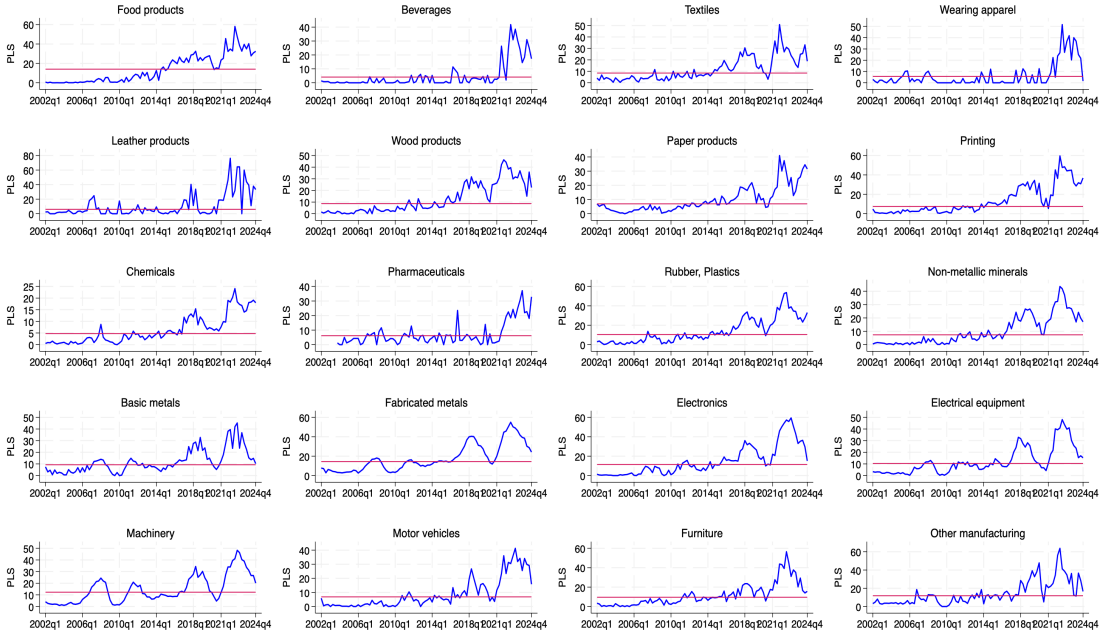
PLS by Sector - Belgium

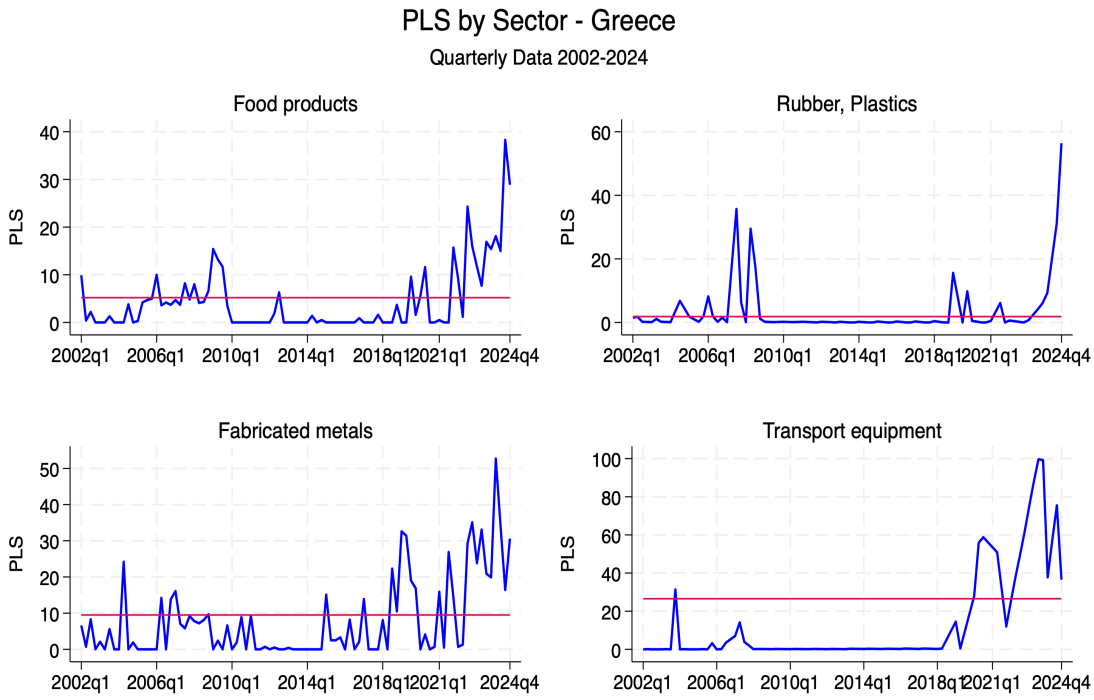
Quarterly Data 2002-2024

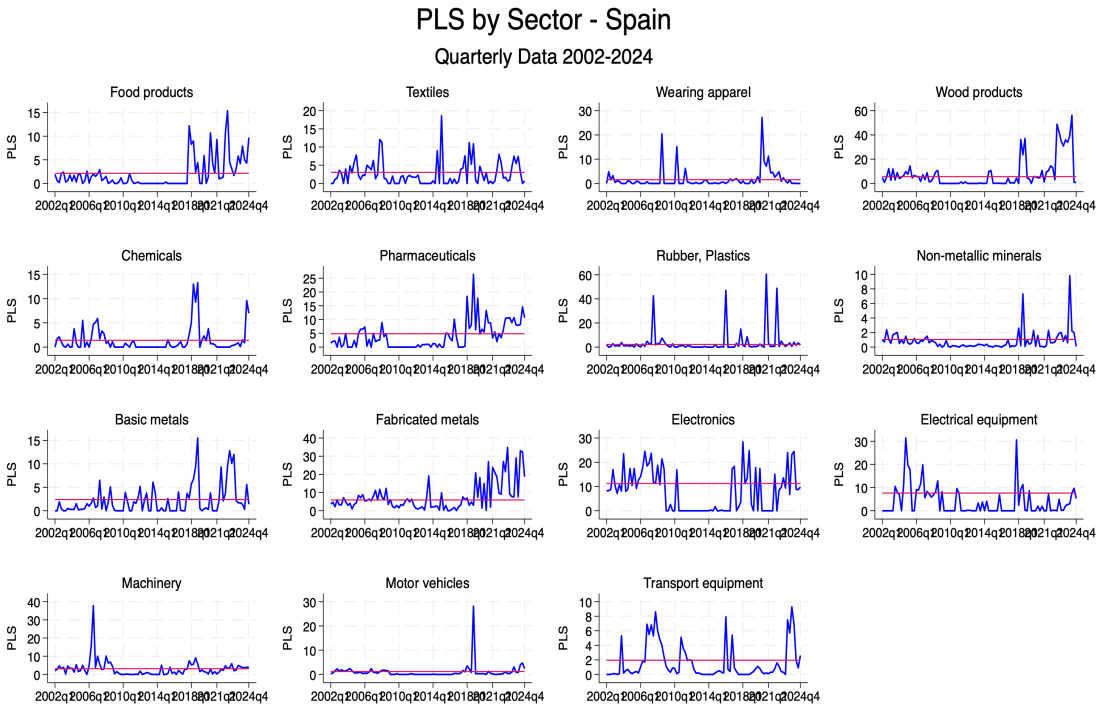


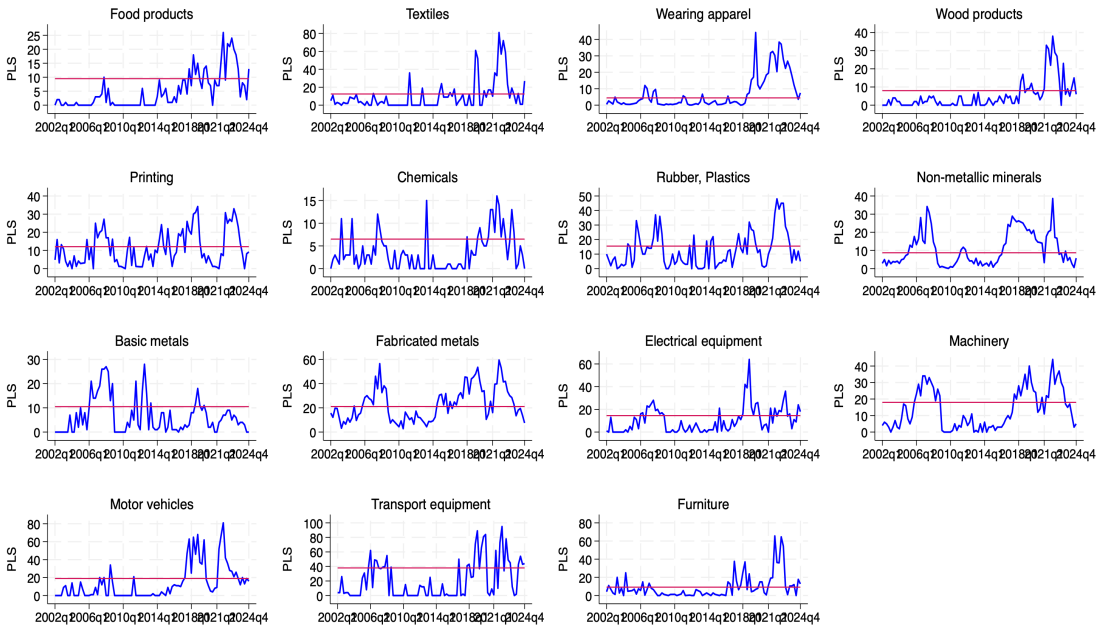
PLS by Sector - Germany

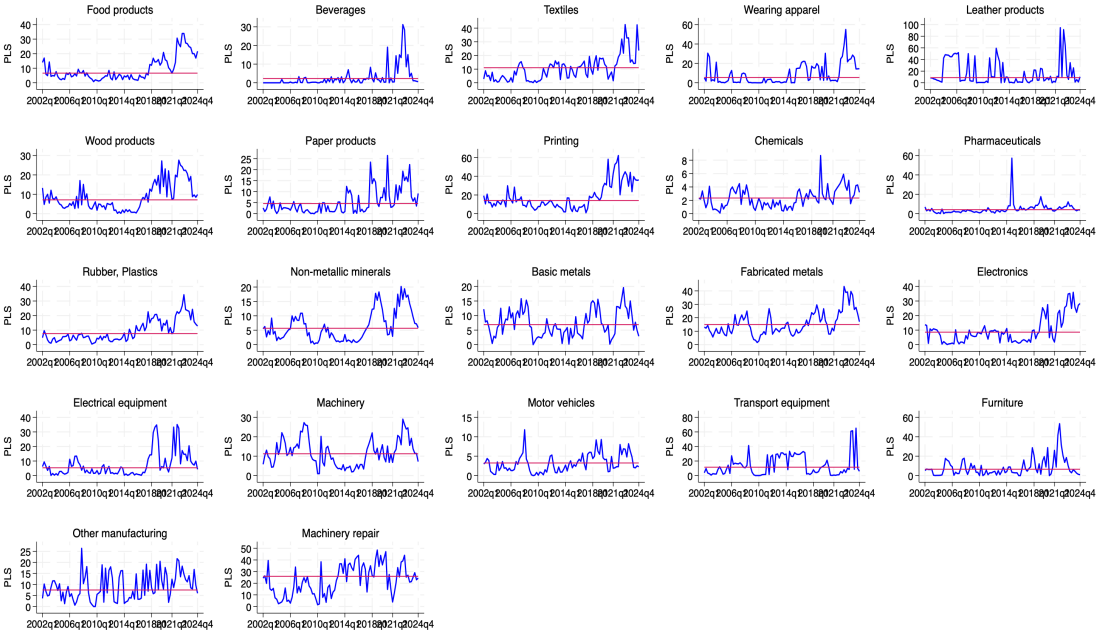
Quarterly Data 2002-2024




PLS by Sector - Greece
Quarterly Data 2002-2024
Food products
Rubber, Plastics
Fabricated metals
Transport equipment
PLS
2002q1
2006q1
2010q1
2014q1
2018q1
2021q1
2024q4

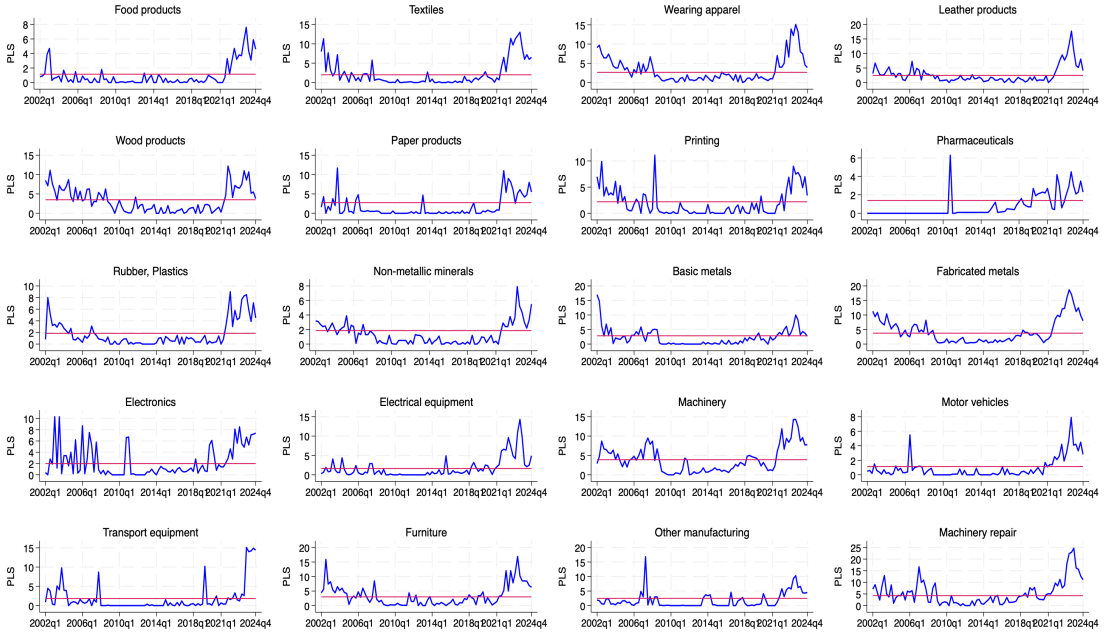

PLS by Sector - Spain
Quarterly Data 2002-2024
Food products
Textiles
Wearing apparel
Wood products
Chemicals
Pharmaceuticals
Rubber, Plastics
Non-metallic minerals
Basic metals
Fabricated metals
Electronics
Electrical equipment
Machinery
Motor vehicles
Transport equipment
PLS


PLS by Sector - Finland
Quarterly Data 2002-2024
Food products
Textiles
Wearing apparel
Wood products
Printing
Chemicals
Rubber, Plastics
Non-metallic minerals
Basic metals
Fabricated metals
Electrical equipment
Machinery
Motor vehicles
Transport equipment
Furniture
PLS
2002q1 2006q1 2010q1 2014q1 2018q1 2021q1 2024q4


PLS by Sector - France
Quarterly Data 2002-2024
Food products
Beverages
Textiles
Wearing apparel
Leather products
Wood products
Paper products
Printing
Chemicals
Pharmaceuticals
Rubber, Plastics
Non-metallic minerals
Basic metals
Fabricated metals
Electronics
Electrical equipment
Machinery
Motor vehicles
Transport equipment
Furniture
Other manufacturing
Machinery repair
PLS
2002q1 2006q1 2010q1 2014q1 2018q1 2021q1 2024q4

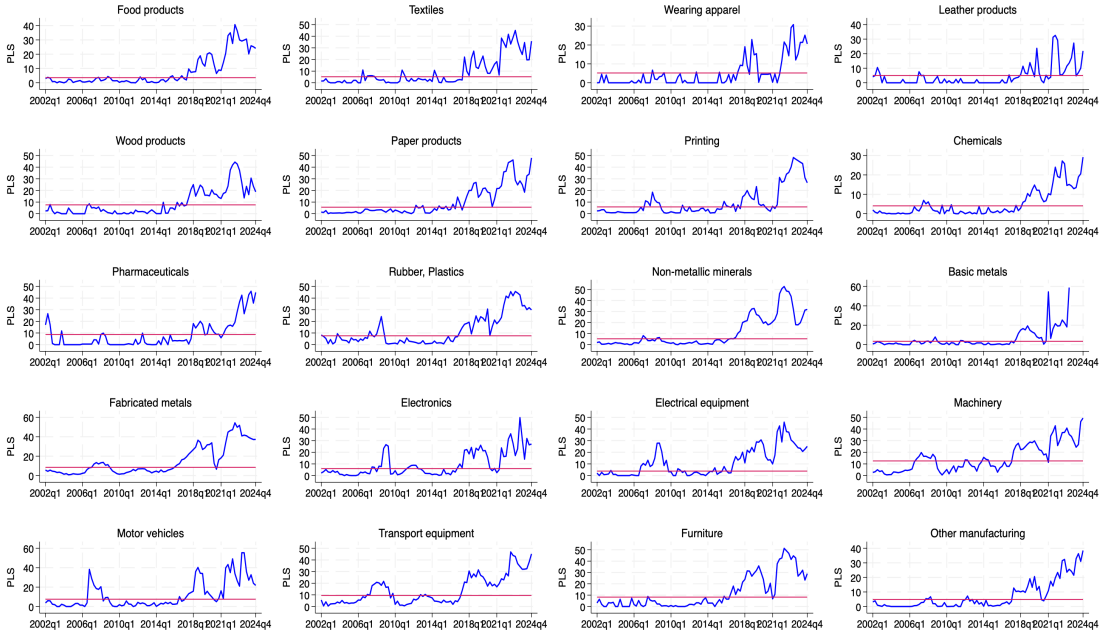
PLS by Sector - Italy

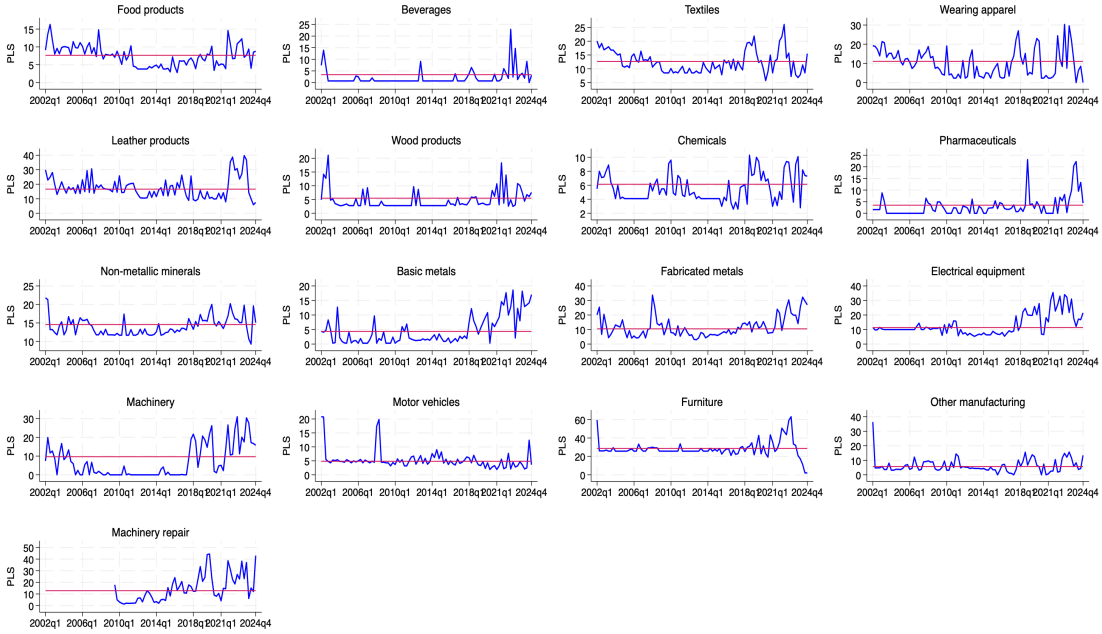
Quarterly Data 2002-2024



PLS by Sector - Netherlands

Quarterly Data 2002-2024




PLS by Sector - Portugal
Quarterly Data 2002-2024
Food products
Beverages
Textiles
Wearing apparel
Leather products
Wood products
Chemicals
Pharmaceuticals
Non-metallic minerals
Basic metals
Fabricated metals
Electrical equipment
Machinery
Motor vehicles
Furniture
Other manufacturing
Machinery repair
PLS
2002q1
2006q1
2010q1
2014q1
2018q1
2021q1
2024q4

B.2 Additional Robustness Checks

TABLE OA.1: Phillips Curve Estimates: Lagged PLS – Two Lags

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.141***	0.088***	0.108***	0.041***	0.022***	-0.003
	(0.004)	(0.005)	(0.009)	(0.007)	(0.006)	(0.006)
Constant		1.255***	1.735***	0.782**	1.325***	-0.435
		(0.054)	(0.234)	(0.328)	(0.167)	(0.356)
Observations	14090	14090	14090	14090	9022	9021

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE OA.2: Phillips Curve Estimates: Lagged PLS – Two Lags

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.068***	0.077***	0.093***	0.081***
	(0.003)	(0.004)	(0.005)	(0.005)
Constant	0.624***	1.346***	1.838***	1.296***
	(0.042)	(0.038)	(0.096)	(0.096)
Observations	14090	14090	9257	9256

Notes: Column (6) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (7), control variables (8), and a proxy for short-term inflation expectations (9). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE OA.3: Phillips Curve Estimates: Lagged PLS – Four Lags

	(1) Baseline	(2) Constant	(3) FE	(4) Time FE	(5) Controls	(6) Infl. Exp.
PLS	0.116***	0.046***	0.055***	0.034***	-0.011*	-0.021***
	(0.004)	(0.004)	(0.008)	(0.007)	(0.006)	(0.006)
Constant		1.632***	2.124***	1.716***	1.418***	-0.404
		(0.056)	(0.191)	(0.305)	(0.161)	(0.355)
Observations	13778	13778	13778	13778	9033	9032

Notes: Column (1) presents the basic model of Eq. (3). Subsequent columns sequentially add a constant (2), fixed effects (3), time fixed effects (4), control variables (5), and a proxy for short-term inflation expectations (6). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE OA.4: Phillips Curve Estimates: Lagged PLS – Four Lags

	(7) IFE	(8) FE	(9) Sec. × Country	(10) Infl. Exp.
PLS	0.061***	0.068***	0.083***	0.071***
	(0.003)	(0.004)	(0.005)	(0.005)
Constant	0.689***	1.448***	2.052***	1.471***
	(0.044)	(0.039)	(0.095)	(0.095)
Observations	13778	13778	9151	9150

Notes: Column (6) includes interacted fixed effects (IFE). Subsequent columns sequentially add (sub-sector and country) fixed effects (7), control variables (8), and a proxy for short-term inflation expectations (9). All additions are cumulative. Robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.