THE ECONOMIC JOURNAL

The Economic Journal, 1–34 https://doi.org/10.1093/ej/ueae006
Advance Access Publication Date: 27 January 2024

JOB POLARISATION, LABOUR MARKET FLUIDITY AND THE FLATTENING OF THE PHILLIPS CURVE*

Daniele Siena and Riccardo Zago

This paper shows that job polarisation—i.e., the disappearance of routine jobs—is changing the characteristics of the labour market. This has structural implications for the relationship between inflation and unemployment, the price Phillips curve. Using data from the European Monetary Union and exploiting the fact that job polarisation accelerates during recessions, we obtain two empirical results. First, countries experiencing a bigger shift in the occupational structure during a downturn exhibit a flatter Phillips curve afterwards. Second, the occupational shifts experienced during the Great Recession and the Sovereign Debt Crisis explain more than a fourth of the flattening of the curve in the 2002–18 period. Then, using a New Keynesian model with unemployment and search and matching frictions, we highlight a channel through which labour market characteristics operate on the slope of the Phillips curve. Increasing labour market *fluidity*—i.e., a higher separation and hiring rate—decreases the slope of the Phillips curve. Using micro-data, we find that in the European Monetary Union non-routine jobs are more *fluid*. We conclude that job polarisation flattened the Phillips curve.

In the European Monetary Union (EMU), the negative relationship between price inflation and unemployment—the price Phillips curve (PC)—was weak before the Great Recession and further flattened after 2009 (see Table 1). Contemporaneously, the share of routine employment has declined (see Figure 1). This phenomenon, called job polarisation, is mostly explained by technological change that led to employment relocation from routine to non-routine tasks (see Autor *et al.*, 2003; Goos *et al.*, 2009 and Firpo *et al.*, 2011, among others). The contribution of this paper is to combine these two facts in order to show if and how job polarisation has played a role in the flattening of the PC in the EMU. In other words, this paper points out that changes in the occupational composition—due to polarisation—affect the overall characteristics of the labour market, with direct implications for the structural relationship between prices and unemployment. This goes beyond the simple idea that polarisation—if interpreted as a result of technological change—can affect the level of prices through a reduction in marginal costs.

* Corresponding author: Daniele Siena, Department of Management, Economics and Industrial Engineering (DIG), Politecnico di Milano, via Lambruschini 4/b, 20156 Milano, Italy. Email: dansiena@gmail.com

This paper was received on 16 February 2022 and accepted on 23 January 2024. The Editor was Franck Portier.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.10546032.

We thank Cristiano Cantore, Andrea Colciago, Jordi Galí, Gunes Gokmen, Basile Grassi, Jonathon Hazell, Nir Jaimovich, Marco Lombardi, Guido Menzio, Tommaso Monacelli, Nicolas Petrosky-Nadeau, Franck Portier, Morten Ravn, Marianna Riggi, Andrej Sokol, Antonella Trigari, Gianluca Violante and participants at the PSE-BDF workshop, AMSE-BDF workshop, Canadian EA Annual Meeting 2021, CEF 2021, HEC Lausanne, ICES 2021, IAAE 2021, NDIF Workshop, ECSB Cluster II 2021, ICMAIF 2021, Padova MW 2022, Bocconi brown bag, T2M 2022, Bicocca University, Università Statale, CBS and Lund University for their helpful comments and suggestions. This paper previously circulated with the title ‘Job Polarization and the Flattening of the Phillips Curve’. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Banque de France. This project also benefited from the HumanTech Project. This project is financed by the Italian Ministry of University and Research for the 2023–7 period as part of the ministerial initiative ‘Departments of Excellence’ (L. 232/2016). The initiative rewards departments that stand out for the quality of the research produced and funds specific development projects. All the remaining errors are ours.

Table 1. *The Phillips Curve Correlation.*

The Stock–Watson Phillips curve $\Delta\pi^{Core}_{i,t} = \alpha_i + \kappa_1 \hat{u}_{i,t} + \kappa_2 \hat{u}_{i,t} \times \mathbb{I}(year > 2009) + \mathbb{I}(year > 2009) + \varepsilon_{i,t}$	
	π^{Core}
$\hat{u}$	−0.0133***
	(0.0029)
$\hat{u} \times \mathbb{I}(year > 2009)$	0.0131***
	(0.0038)
Window	[2002–18]

Note: This table reports the Phillips curve correlation estimated using a panel composed of countries that joined the EMU before 2002 (Luxembourg excluded). The estimating equation and variables definition come from Stock and Watson, 2019 (see Online Appendix C.1 for more details). *** Significance at the 99% levels.

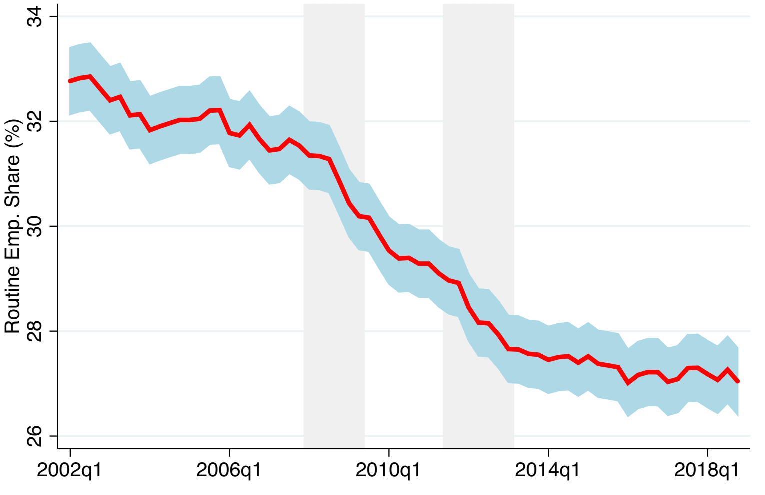

Fig. 1. *Routine Share.*
Note: This figure plots the evolution of the average routine employment share from 2002q1 to 2018q4 across a panel composed of countries that joined the EMU before 2002 (Luxembourg excluded) along with the 95% confidence interval. The routine employment share is defined as the sum of employment in clerical, craft and plant occupations over total employment. The two vertical shaded areas respectively indicate the periods of the Great Recession and of the Sovereign Debt Crisis, as defined by the CEPR Business Cycle Committee. Data are at quarterly frequency.

Both in Europe and in the United States, economists broadly agree that the price PC has weakened in the aftermath of the Great Recession of 2008. This fact has become of great concern among central bankers since a flatter PC prevents monetary policy being effective when trying to stabilise prices (as in the EMU), unemployment or both (as in the United States). For this reason, a good deal of research was conducted to properly assess to what extent the slope of the PC has decreased, and why this has happened. The literature has so far proposed explanations that can be grouped into two not mutually exclusive categories. The first category focuses on inflation expectations and the stronger ability/commitment of central banks to keep inflation low (e.g., Blanchard, 2016). The second category studies the impact of structural changes in the economy, like demographic transition, globalisation and labour market transformations (see, among others, Guerrieri *et al.,* 2010 and Faccini and Melosi, 2023). We contribute to the latter by providing empirical evidence that changes in the occupational structure of the labour market have critical relevance for the slope of the PC.

To do this, we leverage on recent developments in the job-polarisation literature, which documents that the disappearance of routine jobs (clerical, craft and plant occupations) is, not only a long-run phenomenon, but also has cyclical features. In fact, as demonstrated in Jaimovich and Siu (2020) for the United States, job polarisation accelerates during downturns. In other words, the cycle leads to (out-of-trend) shifts in the occupational composition of the labour market in favour of non-routine jobs (professional, managerial, services and sales occupations). Given this, first we provide evidence that the long-run and cyclical properties of job polarisation also hold in the EMU.

In particular, we show that, in normal times, the decline of the routine employment share is very homogeneous across EMU members. Conversely, the Great Recession (GR) and the following Sovereign Debt Crisis (SDC) operate on the common long-run trend of job polarisation through occupational shifts, which are very heterogeneous across countries and recessions. More importantly, these occupational shifts depend on the depth and length of the downturn rather than on pre-recession (i.e., labour or product market) country characteristics. Hence, we exploit these exogenous and heterogeneous compositional changes to assess if and by how much the disappearance of routine jobs affected the relationship between prices and unemployment.

Our main finding is that countries experiencing a bigger change in the composition of the job ladder during a recession exhibit a flatter PC afterwards. In particular, the occupational shifts witnessed during the last two recessions in the EMU explain more than a fourth of the flattening of the price PC observed in the last ten years. These results are robust (*i*) to three specifications of the Phillips curve: the New Keynesian, the regional (Hazell *et al.*, 2022) and the neoclassical; (*ii*) to controlling for other structural breaks; (*iii*) to controlling for changes in the sectoral composition of the economy (i.e., for the transition towards a service economy). Therefore, we conclude that the composition of the labour market matters for the slope of the price PC.

If occupational composition matters for the price PC, we should also find similar results for the wage PC. By applying the same identification strategy, we see that changes in the occupational structural—coming from job polarisation—also flattened the relationship between wages and unemployment.

But, why is this the case? The answer lies in the differences between the surviving and disappearing jobs, i.e., between non-routine and routine occupations. As suggested by the polarisation literature, these jobs are very different in several dimensions. For example, routine workers can be easily substituted by automation and ICT technology (Acemoglu, 2002), and routine jobs are more affected by trade shocks (Autor *et al.*, 2013). Here, we highlight another important difference: *labour market fluidity*. The rate at which workers separate from the current employer and find another job is higher in non-routine occupations.

Does this dimension matter for the slope of the PC? To show that it does, we take the standard New Keynesian model with unemployment and the search and matching frictions of Blanchard and Galí (2010) and we derive the analytical relationship between the slope of the PC and labour market characteristics. We prove that increasing the *fluidity* of the labour market indeed flattens the price Phillips curve. Hence, relocation of workers from less to more fluid markets—due to job polarisation—can indeed weaken the relationship between prices and unemployment. The intuition behind this result is simple: higher fluidity reduces the elasticity of marginal costs to economic conditions (e.g., market tightness) such that employers adjust more the stock of employment rather than wages. This happens because the labour demand becomes more elastic as employers can fire and hire workers more easily. In other words, the labour demand becomes flatter.

Finally, we provide micro-evidence showing that the non-routine labour market is indeed more *fluid*: non-routine jobs exhibit higher separation and hiring rate; non-routine employees are more likely to be offered temporary contracts and to have multiple jobs contemporaneously. As a result, an increase in *fluidity*—due to polarisation—reduces the elasticity of inflation to unemployment.

Literature review This paper relates to two strands of the literature. The first one is on job polarisation, which documents the long-run falling of employment in jobs with high content of routine tasks (among the many, see Acemoglu, 2002; Autor *et al.,* 2006; Acemoglu and Autor, 2011). In this literature, this phenomenon is typically explained by technological change: throughout time, new and cheaper technologies allow the substitution of workers with machines in performing routine tasks, whereas it complements workers in performing non-routine tasks (see, for example, Autor *et al.*, 2003; Autor, 2008 and Acemoglu and Restrepo, 2020). Other sources of polarisation, usually cited in the literature, are international trade and globalisation. In fact, trade and offshoring respectively allow one to substitute home routine productions with imports and to move routine activities in countries with lower labour costs (see Autor *et al.*, 2013; 2015) so as to trigger the decline of home routine employment.

Instead of looking at polarisation as a result of technological change, we focus on its implications on labour market characteristics. To do so, we lean on the cyclical properties of polarisation. As explained in Gaggl and Kaufmann (2020) and Jaimovich and Siu (2020) for the United States, the long-run trend of job polarisation accelerates during downturns, with routine jobs being permanently destroyed. As mentioned above, we show that this property also holds for EMU countries and we leverage on it to study its implications for the PC.

The second strand of the literature is on the flattening of the Phillips curve, and it is both empirical and theoretical. On the empirical side, the work is abundant on both shores of the Atlantic. For the United States, Blanchard (2016), Murphy (2018) and Powell (2018) stated that the PC is still alive, but its slope became flatter from the 1980s on, while inflation expectations have become more anchored. Mavroeidis *et al.* (2014), Hooper *et al.* (2020) and Fitzgerald *et al.* (2022) continued in this line of work. McLeay and Tenreyro (2020) showed that this evolution over time of the PC is also true at the state and city levels, although the correlation between unemployment and inflation is stronger than in the aggregate time series. Fitzgerald *et al.* (2022) showed, using state-level data, that the price Phillips curve flattened only marginally due to structural changes. Hazell *et al.* (2022), using a multi-region model, showed that the slope of the Phillips curve is small and was small even during the early 1980s. Similarly, Portier *et al.* (2020; 2023) showed that the PC has been quite flat in the last two decades. On the other hand, Del Negro *et al.* (2020) offered evidence that the flattening started in the 1990s along with a progressive flattening of the aggregate supply curve. All these papers point out that the flattening is therefore less recent than we thought, and not entirely explained by the Great Recession. Additionally, Bergholt *et al.* (2023) and Ascari *et al.* (2023) have respectively attributed the estimated flattening to a change in the slope of the demand curve and to non-linearities.

For the EMU, Deroose and Stevens (2017), Berson *et al.* (2018), Moretti *et al.* (2019) and Ball and Mazumder (2021) showed that the price PC has flattened from the financial crisis of 2008, but that the structural relation between price dynamics and unemployment and other variables of economic slack still exists. In disagreement with this view is the work of Giannone *et al.* (2014), who showed that the PC was actually steeper during the GR, whereas Ciccarelli *et al.* (2017)

showed that the disconnection between prices and unemployment started after 2012 due to both structural and cyclical factors that have affected aggregate demand.

In both continents, one of the main explanations for the flattening of the PC is based on inflation expectations (e.g., Coibion and Gorodnichenko, 2015). In fact, expectations have become more firmly anchored as the Fed and the ECB have more clearly committed to their inflation objectives. This view has been analysed and expressed in a wide range of research, Fed and ECB communications, including, among others, Roberts (2006), Bernanke (2007), Mishkin (2011), Kiley (2015), Yellen (2015), Ng *et al.* (2018), Pfajfar and Roberts (2022) for the United States, and Draghi (2015), Ciccarelli *et al.* (2017), Dovern and Kenny (2017), Natoli and Sigalotti (2017), Speck (2017) and Bobeica and Jarociski (2019) for the euro area.

Other studies have focused more on structural changes in economic fundamentals, for example, due to demographic dynamics. Daly *et al.* (2016) showed that the shifting composition of the labour force—due to the retirement of baby boomers—has imparted a downward bias to aggregate wage inflation, thus affecting the PC. The importance of demographic dynamics for inflation and inflation expectations is also documented in Pfajfar and Santoro (2008), Bruine de Bruin *et al.* (2010) and Yoon *et al.* (2018), showing—in summary—that an ageing population is deflationary. Other papers focus on the role of technology for the level of inflation. Akerlof *et al.* (1996), Mincer and Danninger (2000), Jorgenson (2001), Ciccarelli *et al.* (2017) and many others showed that technological innovation, digitalisation, automation and ICT contribute to the long-run downward trend of inflation. For what concerns the level of inflation, our paper is in line with this literature: polarisation, as a product of technological change, can have a deflationary effect. Additionally, our paper relates to the work of Wolf and Fornaro (2021) and Basso and Rachedi (2023), who respectively emphasised the role of robots' adoption for the transmission of monetary policy and, vice versa, the effects of monetary policy on automation decisions.

Additionally, our paper is related to a growing literature emphasising the role of labour market dynamics and characteristics. Ravenna and Walsh (2008) estimated a New Keynesian PC with a frictional labour market and showed that search and matching frictions are important for a better fit of the price PC with the data. In a similar theoretical framework, Trigari (2009) showed that search frictions in the labour market generate a lower elasticity of marginal costs with respect to output. Ravenna and Walsh (2012) showed that labour market composition is important for the unemployment-inflation trade-off faced by the monetary authority. Moscarini and Postel-Vinay (2023) introduced on-the-job search in a New Keynesian model and showed that cyclical labour misallocation leads to deflation. Through a similar theoretical set-up, Faccini and Melosi (2023) stressed the importance of the mobility of workers on the job ladder to rationalise the missing inflation and a flatter PC in the post-GR era.[1] Cantore *et al.* (2020) highlighted how it is important to look at labour market dynamics, as the labour market share, to correctly model the relationship between the real economy and prices. Pace and Hertweck (2019) relied on labour search and matching frictions with internal habit formation to explain the co-movement across durable and non-durable good prices after a monetary contraction. Finally, Lombardi *et al.* (2023) showed that the decline in bargaining power of workers has weakened the inflation-output gap relationship. We contribute to this literature by showing that labour market *fluidity* is another important channel to explain the recent evolution of the price PC. Moreover, our paper relates to that niche in the literature showing how polarisation can also explain the de-unionisation of the

[1] Despite its importance, our theoretical framework does not include on-the-job search since no data are available on job-to-job transitions by occupation for the sample of countries under consideration.

workforce (see, for example, Acemoglu *et al.*, 2001; Açkgöz and Kaymak, 2014; Dinlersoz and Greenwood, 2016 and Foster *et al.*, 2016).

Along with the literature on the price PC, there is another strand focusing on the wage PC. For example, Leduc and Wilson (2017) and Galí and Gambetti (2019) showed that, in the United States, the relationship between wages and unemployment also flattened. As explained in Benigno and Ricci (2011), Schmitt-Grohé and Uribe (2013) and Daly and Hobijn (2014), downward wage rigidities are important to rationalise this fact. Conversely, Petrosky-Nadeau *et al.* (2021) showed that efficient rent sharing between consumers and producers in the goods market drives down wage bargaining and causes the flattening of the wage PC. While these papers provide evidence for a flattening of the wage PC in the United States, evidence for the euro area is less clear. For example, Bulligan and Viviano (2017) showed that the wage PC has been steepening, while Nickel *et al.* (2019) showed that the GR flattened the PC.

Our paper is organised as follows. Section 1 presents the data and facts on job polarisation in the EMU. Section 2 estimates an augmented price PC for the EMU that takes into account changes in the occupational structure of the labour market. Section 3 repeats the exercise for the wage PC. Section 4 uses an analytical theoretical model to highlight one channel through which job polarisation can affect the slope of the price PC and gives micro-evidence that differences among surviving and non-surviving jobs is key for our result. Section 6 concludes.

1. Data and Labour Market Dynamics in the EMU

1.1. *Data, Definitions and Sample Selection*

Our focus is on the European Monetary Union. In particular, on countries that joined the EMU from the beginning: Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal and Spain. We do this for two reasons. First, countries that joined the EMU more recently have very peculiar convergence trajectories in terms of inflation and employment. Additionally, their entrance in the EMU in some cases coincides with the beginning of a downturn (e.g., see Estonia). Therefore, it would be inappropriate to use our identification strategy for these countries. Second, late entrants have very unreliable employment and unemployment series. Unfortunately, Luxembourg suffers the same problem with employment data. Therefore, despite being a member of the EMU since the beginning, we exclude it from the analysis. As a result, we end up with 11 countries (EMU11) and consider data from 2002q1 until 2018q4.

Data come from five sources: Eurostat, the ECB Data Warehouse, national statistical institutes, the ECB Survey of Professional Forecasters and the OECD. Eurostat gives information on employment by occupation, according to the International 2008 Standard Classification of Occupations (ISCO-08). We consider country-specific employment series for workers in the 15–74 age bracket. Once these series are corrected from statistical breaks and changes in occupation classification, we follow Jaimovich and Siu (2020) and group these jobs into three major categories based on their task content: (*i*) managers, professionals, technical and associate professionals, armed force employees as *abstract* workers; (*ii*) clerical, craft and plant employees as *routine* workers; (*iii*) elementary, skilled agricultural, forestry and fishery employees, sales and service workers as *manual* workers. Finally, under this grouping, we build employment share series for each category. We use country-level data from Eurostat to also build (*i*) the unemployment rate for the population in the 15–74 age bracket; (*ii*) the price inflation rate as a year-on-year percentage change of the Harmonised Index of Consumer Prices (HICP); (*iii*) the wage inflation rate as a year-on-year percentage change of the labour-cost index (salaries and wages) for the business economy; (*iv*) the non-tradable wage inflation rate as a year-on-year percentage change of the labour-cost index for non-tradable sectors; (*v*) the energy price index. Datastream provides the import price index for each country in the sample taken from national statistics. Data for the natural level of unemployment (NAIRU) are from the OECD economic outlook, which we use to construct the unemployment gap. The ECB Survey of Professional Forecasters provides information on the expected HICP inflation by country. All series are at quarterly frequency.

We use the ECB Data Warehouse to extrapolate country-specific business cycle dates from the time series of real GDP. For each country, we define a recession as a period in which GDP falls for at least two consecutive quarters. The peak of the recession is identified as the last quarter before which real GDP starts falling and the trough of the recession as the last quarter after which real GDP starts increasing again for at least two consecutive quarters. This allows us to punctually identify in which phase of the business cycle every country is in every quarter. The business cycle dates extracted using our methodology do not differ from those obtained using the algorithm of Harding and Pagan (2002) without imposing a minimum length of a phase. Finally, we use the ECB Data Warehouse to build a non-tradable inflation series. In order to do so, following Siena (2021), we use the year-on-year change in the GDP deflator of non-tradable sectors. Online Appendices A.2–A.5 report details on the construction of each variable along with figures for all series and countries in the sample.

1.2. *Job Polarisation in the EMU11*

As mentioned, the large literature on job polarisation documents the long-run falling of employment in jobs with a high content of routine tasks. Yet, this long-run trend has a short-run counterpart. As shown in Jaimovich and Siu (2020) for the United States, job polarisation accelerates during recessions with job losses concentrated the most in routine occupations. We leverage on these results for the United States and show that both the long- and short-run properties of job polarisation hold for countries in the EMU11 as well. As Figure 1 displays, the share of routine employment across EMU11 countries has been following a downward trend. From the beginning of the GR until the end of the SDC, the process of polarisation accelerated and the downward trend became steeper. Afterwards, the routine share returned to the pre-GR trend. However, this first-sight analysis is confounding since it looks like the trend accelerated without interruption until the end of the SDC, and without any effect of the economic expansion between the two crises. But, if 'job polarisation follows the cycle'—as explained in Jaimovich and Siu (2020)—we should observe a change in the trend in every single phase of the business cycle, i.e., in every single contractionary and expansionary phase. In order to check this point, we estimate

$$Share^R_{i,t} = \alpha_i + \beta_1 time + \beta_2 phase_{i,t} + \beta_3 phase_{i,t} \times time + \varepsilon_{i,t}, \tag{1}$$

where $Share^R_{i,t}$ is the routine employment share at time t in country i, α_i is the country fixed effect, *time* is the number of quarters, $phase_{i,t}$ = [Before GR, GR, Between GR and SDC, SDC, After SDC] is a vector of mutually exclusive dummies taking value one if, at time t, country i is currently in that cyclical phase and $\varepsilon_{i,t}$ is the error term. To construct this variable, we use country-specific business cycle dates, as explained in Section 1.1.

Table 2 shows results from this panel regression. From column (1), we see that—between 2002q1 and 2018q4—the routine employment share follows a downward trend across countries, with an average decline of 0.10 percentage points (pps) every quarter. In column (2), we investigate

Table 2. *Trend Decomposition of Occupational Shares across the EMU11.*

	(1) $Share^R$	(2) $Share^R$	(3) $Share^A$	(4) $Share^A$	(5) $Share^M$	(6) $Share^M$
time	-0.103^{***}	-0.062^{**}	0.087^{***}	0.047^{**}	0.016	0.016
	(0.012)	(0.020)	(0.018)	(0.020)	(0.015)	(0.015)
GR × *time*		-0.332^{***}		0.272^{**}		0.060
		(0.058)		(0.092)		(0.063)
Between GR and SDC × *time*		−0.061		0.045		0.016
		(0.035)		(0.056)		(0.029)
SDC × *time*		-0.128^{**}		0.091		0.037
		(0.049)		(0.054)		(0.034)
After SDC × *time*		0.048		0.003		-0.051^{**}
		(0.030)		(0.028)		(0.017)
Observations	748	748	748	748	748	748
R^2	0.764	0.813	0.565	0.588	0.070	0.118
Country FEs	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	No	No	No	No

Note: SEs are reported in parentheses, clustered at the country level. The unit of observation in columns (1) and (2) is the share of routine employment, i.e., clerical, craft, plant employment. The unit of observation in columns (3) and (4) is the share of abstract employment, i.e., managers, professionals, technical and associate professionals, armed force employment. The unit of observation in columns (5) and (6) is the share of manual employment, i.e., elementary, skilled agricultural, forestry, fishery, sales and service employment. The variable *time* is the number of quarters. GR is a country-specific dummy variable taking value one for periods in which a country is experiencing the Great Recession. Between GR and SDC is a country-specific dummy variable taking value one for periods in which a country is in between the GR and the Sovereign Debt Crisis. SDC is a country-specific dummy variable taking value one for periods in which a country is experiencing the Sovereign Debt Crisis. After SDC is a country-specific dummy taking value one for periods after the Sovereign Debt Crisis. **, *** Significance at the 95% and 99% levels.

how much of this decline is imputable to each different phase of the business cycle. Before the GR, the routine share is decreasing by 0.06 pps every quarter. When entering the GR, the trend accelerates by five times across all countries. Between the two recessions, the slope of the trend is not significantly different from the slope estimated in pre-GR periods. When EMU11 economies enter the SDC, the trend of job polarisation returns to accelerating at a pace two times larger than before the GR. Once EMU11 countries are out of the SDC, again there is no statistical difference between pre-GR and post-SDC trends. This empirical evidence generalises the results of Jaimovich and Siu (2020) and proves that, also in the EMU, job polarisation has, not only a long-run driver, but also a cyclical component: there is a negative trend in routine employment, which temporarily accelerates every time the economy enters a recession.

What about the dynamics of other jobs? Figures 2(a) and 2(b) respectively show the evolution of the employment share of abstract and manual jobs. While both present strong seasonal fluctuations, the abstract share follows a clear upward trend, whereas the manual share looks more stationary. This is also confirmed in columns (3) and (5) of Table 2 where we perform the same analysis of (1), but now with the abstract ($Share^A$) and manual share ($Share^M$) as dependent variables. We see that, across all countries, the long-run fall in the routine share is almost entirely compensated by an expansion of the abstract share. When looking at the decomposition of this trend across different business cycle phases, we find that the abstract share increases by 0.047 pps every quarter in pre-recession periods (column (4)). When the EMU11 enters the GR, this positive trend accelerates by five times before going back to the pre-GR trend from there afterwards. On the other hand, in pre-GR periods, the share of manual employment is stationary (column (6)), and neither the GR nor the SDC significantly affect this behaviour. Only in periods

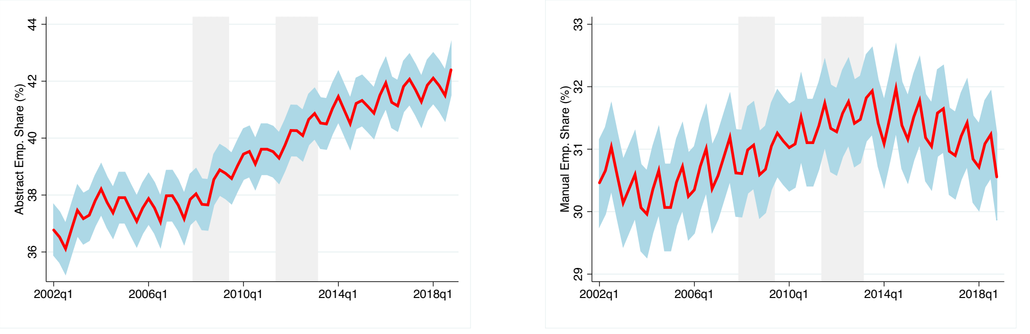

(a) Abstract Share (b) Manual Share

Fig. 2. *Abstract and Manual Share across the EMU11.*
Note: Panel (a) plots the evolution of the mean abstract employment share from 2002q1 to 2018q4 for those countries that joined the EMU before 2002 (Luxembourg excluded) along with the 95% confidence interval. The abstract employment share is defined as the sum of employment in managerial, professional, technical and associate professional, and armed force occupations over total employment. Panel (b) plots the evolution of the mean manual employment share from 2002q1 to 2018q4 for the same 11 EMU countries. The manual employment share is defined as the sum of employment in elementary, skilled agricultural, forestry and fishery, sales and service occupations over total employment. The two vertical shaded areas respectively indicate the periods of the Great Recession and of the Sovereign Debt Crisis as defined by the CEPR Business Cycle Committee. Data are at quarterly frequency.

after the SDC, the share of manual jobs starts to significantly decline by 0.05 pps each quarter. The estimated coefficients for the variable $phase_i$ are reported in Online Appendix B1.

2. The Slope of the Phillips Curve and the Occupational Composition of the Labour Market

Does employment composition matter for the relationship between price dynamics and unemployment? In order to check this, we start by plotting the cross-country correlation between the long-run change of the slope of the PC and the long-run change of the routine share. Country-specific slopes are estimated following the PC specification of Stock and Watson (2019) for periods before and after the GR (see Online Appendix C.1 for all the details). Then we take the change between the post- and pre-recession estimates. Figure 3 plots this over the long-run percentage change of the routine share between 2002q1 and 2018q4. The correlation is −0.76 and is significant at the 99% level.

In light of this evidence, we now want to understand if there is a causal relationship between changes in labour market composition and the flattening of the PC. This is challenging as there are several sources of endogeneity. First of all, there is a potential spurious correlation between the decline in routine employment and the flattening of the PC over time. Second, the estimates of the PC could be biased as there are factors that can influence both unemployment and inflation at the same time (e.g., supply shocks). Third, there are other long-run factors that might have influenced the slope of the PC and employment composition, such as changes in the sectoral composition.

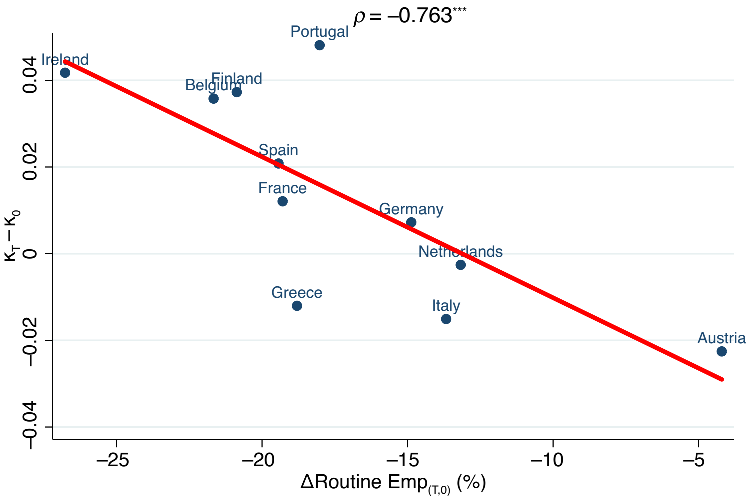

Fig. 3. *The Slope of the Phillips Curve and the Composition of the Labour Market.*
Note: This figure plots the long-run change of the slope of the Phillips curve over the long-run change of the share of routine employment by country. To compute the change in slope, we estimate for each country the Phillips curve defined in Stock and Watson (2019) (see Online Appendix C.1 for details) after (κ_T) and before the GR (κ_0), and then we take the difference between the two estimates. The long-run change in routine employment is the percentage change of the routine share between 2018q4 and 2002q1. The sample is composed of those countries that joined the EMU before 2002 (Luxembourg excluded). At the top of the graph, the correlation (ρ) between variables is reported along with its significance level. *** Significance at the 99% levels.

2.1. *Cyclical Occupational Shifts*

To address the endogeneity driven by the co-movement of job polarisation and the slope of the PC, we need to find a variation of the occupational composition that is ex ante orthogonal (*i*) to past country-specific characteristics and (*ii*) to past price dynamics. In this section we provide evidence that cyclical movements in job polarisation respect these criteria.

As we know from Section 1.2, recessions operate on the long-run trend of job polarisation through (level) shifts in the occupational composition of the labour market. This translates into a bigger destruction of routine jobs in favour of non-routine ones (most of all abstract jobs) during recessions. To quantify these structural shifts that occurred within each country, we consider the change in the level of the routine employment share matured during each recession. In other words, the extent to which the composition of the labour market of country i changed due to the cycle $c = \{GR, SDC\}$ can be measured as the percentage change of the routine employment share between the peak and trough of the recession c, according to the business cycle dates specific to country i. Formally,

$$Shift^R_{i,c} = \frac{Share^R_{peak_{i,c}} - Share^R_{trough_{i,c}}}{Share^R_{peak_{i,c}}}. \tag{2}$$

Figure 4 plots the levels of our measure of occupational shift ($Shift^R_{i,c}$) for each country of the EMU11 and each recession. Despite the fact that all countries were following the same job polarisation trend (as shown in Section 1.2), the cyclical rate of routine job destruction varies

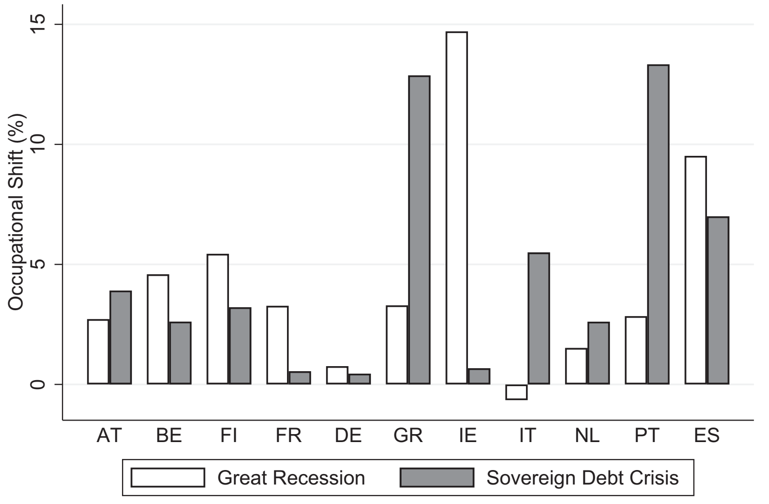

Fig. 4. *Occupational Shifts across EMU11 Countries, by Recession.*
Note: This figure plots the occupational shifts experienced by each country that joined the EMU before 2002 (Luxembourg excluded) during the Great Recession and the Sovereign Debt Crisis. Each occupational shift is defined as the percentage change in routine employment share between the peak and trough of each recession, according to country-specific business cycle dates.

substantially across EMU11 members and across recessions. The mean occupational shift is 4.6% (4.4% for the GR and 4.8% for the SDC). The correlation between the shifts matured during the two recessions is 0.15 and it is not significantly different from zero.

Yet, it is important to test whether the heterogeneity in $Shift_{i,c}$ observed across countries and recessions is due to pre-recession country-specific characteristics. In fact, as shown in Autor *et al.* (2013), the exposure of the labour market to a 'polarisation shock' can be explained by the employment composition of that market before the downturn. In light of this, we look at the correlation between the routine employment share—measured one-quarter before the beginning of the recession—and our measure of occupational shift. As can be seen from Figure 5(a), we find only a mild correlation not significantly different from zero. After having checked that the pre-recession employment composition does not matter, in Figure 5(b) we check if the pace at which each economy has polarised until the start of each recession matters for the size of the cyclical occupational shift. This is mostly to control if the polarisation trend is not explaining the size of the shift. Again, we find a correlation not significantly different from zero. This confirms that the cyclical occupational shifts are heterogeneous across countries and recessions and that they are not path dependent.

Since the occupational composition also varies across sectors in the economy (e.g., manufacturing and construction have a larger share of routine workers), we study whether the sectoral composition could be related to the cyclical shift. To do so, in Figure 5(c) we plot the share of value added from manufacturing and construction, measured one-quarter before the beginning of the recession, on the shift. Once again, we do not find any significant correlation. Finally, we check if pre-recession inflation correlates with the shift. Figure 5(d) shows no significant relationship between pre-recession inflation and our measure of occupational shift. In Online Appendix B2, we show that the occupational shift is also not significantly correlated with other features of the

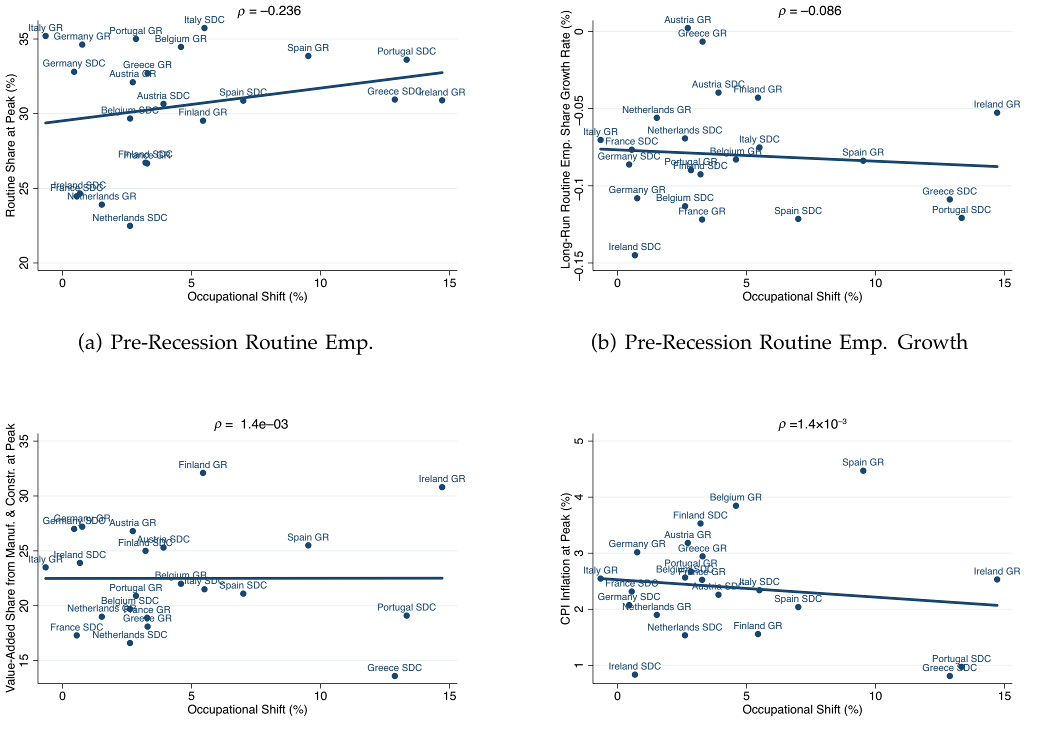

(a) Pre-Recession Routine Emp.

(b) Pre-Recession Routine Emp. Growth

(c) Pre-Recession Sectoral Composition

(d) Pre-Recession Inflation

Fig. 5. *Pre-Recession Characteristics versus the Occupational Shift.*
Note: For all subplots, the x axis shows the occupational shift (in percentages) experienced during both downturns by each country that joined the EMU before 2002 (Luxembourg excluded). Each occupational shift is defined as the percentage change in routine employment share between the peak and trough of each recession, according to country-specific business cycle dates. In panel (a) the y axis is the routine employment share, measured at the peaks of the GR and SDC. In panel (b), the y axis is the long-run growth rate of the routine employment share, measured as the slope of the linear trend fitting the routine share series until the GR and SDC, respectively. For panel (c), the y axis is the value added from manufacturing and construction (as the percentage of GDP), measured at the peaks of the GR and SDC. For panel (d), the y axis is the level of core inflation, measured at the peaks of the GR and SDC. At the top of each graph, the correlation (ρ) between variables is reported along with its significance level.

labour market (such as the pre-recession level of unemployment or the separation rate) and the slope of the PC.

It is now important to show that the variation in the occupational shift is really imputable to the economic downturn only. As shown in Figure 6(a), our measure of occupational shift is significantly correlated at the 99% level with the percentage change in GDP matured between the peak and trough of each recession. Similarly, Figure 6(b) displays a strong correlation—significant at the 99% level—between the duration of each downturn (expressed in quarters) and our measure of occupational change.

This evidence suggests that the labour market transformation experienced by each country during each recession is orthogonal to country-specific characteristics, but its magnitude depends

(a) GDP Contraction

(b) Recession Duration

Fig. 6. *Duration and Size of the Recession versus the Occupational Shift.*
Note: For both subplots, the *x* axis shows the occupational shift (in percentages) experienced during both downturns by each country that joined the EMU before 2002 (Luxembourg excluded). Each occupational shift is defined as the percentage change in routine employment share between the peak and trough of each recession, according to country-specific business cycle dates. In panel (a), the *y* axis is the GDP percentage change measured between the peak and trough of the GR and SDC. In panel (b), the *y*-axis is the duration (in quarters) of each recession, measured as the number of quarters between the peak and trough. At the top of each graph, the correlation (ρ) between variables is reported along with its significance level. *** Significance at the 99% levels.

on the size and persistence of the downturn only. As explained in depth in Online Appendix B.3, this is due to the specific nature and causes behind the GR and SDC. In fact, these are respectively a financial and a government debt crisis. Therefore, although labour markets are important for economic dynamics, the severity and length of these two specific recessions do not seem to primarily depend on labour market features, but rather on other factors such as the solvency of the financial system and debt-to-GDP ratio. This ensures that countries do not self-select into larger occupational shifts through the characteristics of their labour market, and that aggregate shocks hit some economies more than others for reasons unrelated to the labour market.

In light of this, we can conclude that (*i*) recessions operate on the process of polarisation through out-of-trend shifts; (*ii*) countries do not self-select into these shifts nor a deeper recession based on the characteristics of the labour market.

2.2. *Sectoral Dynamics behind Job Polarisation*

Although Figure 5(c) tells us the that the sectoral composition does not matter ex ante for the cyclical shift in the occupational structure of the labour market, we know that some specific sectors are much more concentrated of routine workers and more cyclical. This is particularly true for the manufacturing and construction sector. Figure 7(a) plots the cross-country employment share in manufacturing and construction. Clearly, the sectoral employment dynamic mimics the routine employment share dynamic of Figure 1. This implies that—although job polarisation remains a fact across all sectors—its long-run and cyclical behaviour is intertwined with manufacturing and construction.

This raises a red flag, in particular, for the implication that specific sectoral dynamics can have on both price dynamics and employment composition. In other words, sectoral dynamics

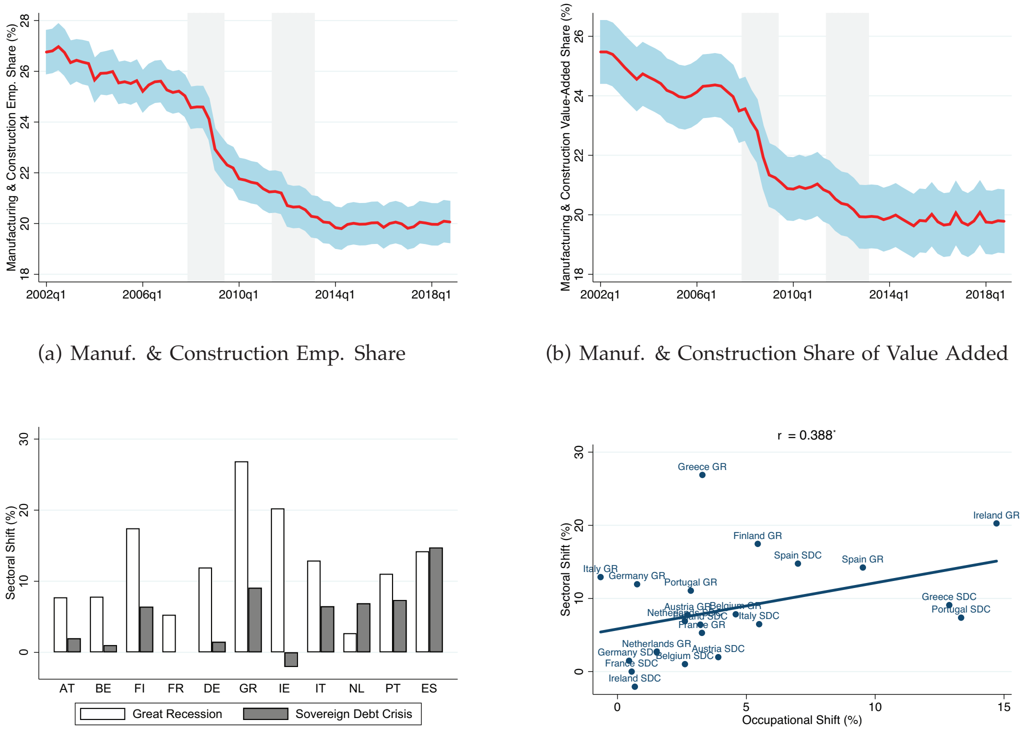

(a) Manuf. & Construction Emp. Share

(b) Manuf. & Construction Share of Value Added

(c) Sectoral Shifts

(d) Sectoral versus Occupational Shifts

Fig. 7. *Manufacturing and Construction Employment and Relative Weight in the Economy.*
Note: Panel (a) plots the evolution of the mean employment share in manufacturing and construction across those countries that joined the EMU before 2002 (Luxembourg excluded) along with the 95% confidence interval. Panel (b) plots the evolution of the mean share of value added from manufacturing and construction across the same 11 EMU countries along with the 95% confidence interval. The two vertical shaded areas respectively indicate the periods of the Great Recession and of the Sovereign Debt Crisis as defined by the CEPR Business Cycle Committee. Data are quarterly and span from 2002q1 to 2018q4. Panel (c) plots the sectoral shifts experienced by each country that joined the EMU before 2002 (Luxembourg excluded) during the Great Recession and the Sovereign Debt Crisis. Each sectoral shift is defined as the percentage change in the share of value added from manufacturing and construction measured between the peak and trough of each recession, according to country-specific business cycle dates. For panel (d), the *y* axis is the sectoral shift and the *x* axis is the occupational shift, defined as the percentage change in routine employment share measured between the peak and trough of each recession, according to country-specific business cycle dates. At the top of the graph, the correlation (ρ) between variables is reported along with its significance level. * Significance at the 90% levels.

could be confounding factors when trying to address the role of occupational composition on the slope of the PC. Therefore, it is convenient to analyse how the sectoral structure of EMU11 countries has evolved over time and over the cycle. As shown in Figure 7(b), the contribution of manufacturing and construction (in terms of value added) has moved across countries roughly from 25% to 20% in the last two decades. This trend has accelerated in both recessions. Given this, we build a measure of sectoral shift isomorphic to (2), i.e.,

$$Shift^{Manuf}_{i,c} = \frac{VA.Share^{Manuf}_{peak_{i,c}} - VA.Share^{Manuf}_{trough_{i,c}}}{VA.Share^{Manuf}_{peak_{i,c}}}, \tag{3}$$

which captures the percentage change of the share of value added from manufacturing and construction in country i, as measured between the peak and trough of each recession $c = \{GR, SDC\}$, according to the business cycle dates of country i. Figure 7(c) plots the levels of our measure of sectoral shift ($Shift^{Manuf}_{i,t}$) for each country of the EMU11 and each recession. The mean sectoral shift is 8.7% (12.6% for the GR and 4.8% for the SDC). Finally, Figure 7(d) plots the sectoral shift on the occupational shift. As expected, there is a positive correlation equal to 0.38, although it is only significant at the 90% level. Further details and analysis on sectoral dynamics are available in Online Appendix B.4.

2.3. *Occupational Structural Changes and the Flattening of the Price PC*

If the composition of the labour market matters for the slope of the PC, out-of-trend changes in employment composition should affect the structural relationship between prices and unemployment. In light of this argument, in this section we exploit the cross-country variation in occupational shifts that occurred during the GR and SDC to study the flattening of the PC in periods following each downturn. Doing so requires two things. First, we need a rigorous definition and estimation of the Phillips curve. This alone is quite an empirical challenge and the literature offers different approaches. We start by considering a baseline New Keynesian Phillips curve (NKPC). Second, we need to include the occupational shifts in the model in order to test if employment composition really matters for the slope. Formally, we consider the equation

$$\begin{aligned}
\pi_{i,t} = {} & \alpha_i + \kappa \hat{u}_{i,t} + \gamma_1 \mathbb{E}(\pi_{i,t+4}) + X'_{i,t}\gamma_2 \\
& + \sum_{c=\{GR,SDC\}} \{\delta_{1,c} After_{i,c} + \kappa_c After_{i,c} \times \hat{u}_{i,t}\} \\
& + \sum_{c=\{GR,SDC\}} \{\delta_{2,c} After_{i,c} \times Shift^{R}_{i,c} + \tilde{\kappa}_c After_{i,c} \times Shift^{R}_{i,c} \times \hat{u}_{i,t}\} \\
& + \sum_{c=\{GR,SDC\}} \{\delta_{3,c} After_{i,c} \times Shift^{Manuf}_{i,c} + \tilde{\tilde{\kappa}}_c After_{i,c} \times Shift^{Manuf}_{i,c} \times \hat{u}_{i,t}\} + \varepsilon_{i,t},
\end{aligned} \tag{4}$$

where $\pi_{i,t}$ is the HICP inflation in country i at time t, measured as the year-on-year percentage change of the harmonised consumer price index, α_i is the country fixed effect, $\hat{u}_{i,t}$ is the unemployment gap measured as the percentage deviation of the unemployment series from the country-specific NAIRU (OECD), $\mathbb{E}(\pi_{i,t+4})$ is the current expectation of HICP inflation one year from now, as provided by the ECB survey of professional forecasters[2] and $X'_{i,t}$ is a vector of controls. As energy and trade price fluctuations can influence current inflation, this vector includes the import and energy price index. We add to this set of controls three time dummies: a time dummy $phase_{i,t}$ to take into account changes in the level of inflation, and $After_{i,GR}$ and $After_{i,SDC}$ that take value one for all periods after the GR and the SDC, respectively. We denote

[2] An alternative would be to use households' expectations that, as shown by Coibion and Gorodnichenko (2015) and Gorodnichenko and Candia (2021), are closer to firms' expectations (i.e., the price setters). However, this measure is not available for all countries and periods under consideration here.

by $Shift^{R}_{i,GR}$ ($Shift^{R}_{i,SDC}$) the shift in the occupational structure that occurred during the GR (the SDC), as defined in (2) of Section 2.1, and by $Shift^{Manuf}_{i,GR}$ ($Shift^{Manuf}_{i,SDC}$) the shift in the sectoral structure that occurred during the GR (the SDC), as defined in (3) of Section 2.2. Here $\varepsilon_{i,t}$ is the error term.

In other words, (4) augments the baseline New Keynesian Phillips curve (i.e., the first line of the equation) by taking into account all potential structural changes that occurred during each recession that might have affected both the relationship between unemployment and inflation and the level of inflation in post-recession periods (i.e., the second line of the equation). On top of this, the third line of (4) takes into account how much of the flattening in post-recession periods can be attributed to structural changes in the occupational composition of the labour market that occurred during each recession. Line four controls for contemporaneous changes in the sectoral composition. Therefore, we use the augmented NKPC of (4) to test whether shifts in the occupational composition matter for the slope of the PC. Formally, we want to test

$$H_0 \colon \tilde{\kappa}_c = 0 \quad \text{for all } c = \{GR, SDC\},$$

once netting out the effect of all other possible structural changes that might have flattened the curve after the GR and the SDC. Yet, this hypothesis cannot be tested via an ordinary least square regression. In fact, these estimates would be biased as supply shocks can contemporaneously affect the unemployment gap, inflation and inflation expectations. Therefore, all regressors ($\hat{u}_{i,t}$, $After_{i,GR} \times \hat{u}_{i,t}, \ldots, After_{i,SDC} \times Shift^{Manuf}_{i,SDC} \times \hat{u}_{i,t}$ and $\mathbb{E}(\pi_{i,t+4})$) should be instrumented.

To do so, we rely on both external and internal instrumental variables (IVs). As external instruments for the unemployment gap and all its interactions, we use aggregate off-the-shelf high-frequency monetary policy shocks for the euro area (mps_t) from Altavilla *et al.* (2019). In Altavilla *et al.* (2019) monetary policy surprises are identified as exogenous/unexpected changes in the three-month overnight index swap that occurred during the monetary policy communication window. We select those shocks that are not correlated with the stock market to separate them from information shocks (see Jarociński and Karadi, 2020). We sum these shocks at a quarterly frequency and use mps_{t-k}, $mps_{t-k} \times After_{i,c}$, $mps_{t-k} \times Shift^{R}_{i,c}$, $mps_{t-k} \times Shift^{Manuf}_{i,c}$ for $c \in \{GR, SDC\}$ and $k \in \{0, \ldots, 4\}$ as instruments for the unemployment gap and all its interaction terms. On the other hand, we instrument country-specific inflation expectations with the lag of the aggregate inflation expectations for the EMU (i.e., the average of lagged inflation expectations across countries). The fact that these instruments are common across all countries could potentially be a threat for our identification. Despite this, these instruments are sufficiently relevant (Wald F-statistic $= 12.82$). In particular, the monetary shocks explain well movements in country-specific unemployment and therefore represent a good measure for (common) aggregate demand shocks at the EMU level (see Table C.2 in Online Appendix C.2 for first-stage statistics).

Columns (1)–(3) of Table 3 show two-stage least square (2SLS) results under this instrumentation. Note that we keep constant the set of instruments across columns to facilitate comparison. In column (1), we start by considering the baseline PC and we include all controls along with all time dummies and all shifts (in level) that might have affected the level of inflation over time. We find that, over the entire 2002–18 window, the slope of the PC is -0.006. In column (2), we study the flattening of the PC in post-GR years by adding the interaction $After_{GR} \times \hat{u}$. Here we find that the slope of the PC was -0.022 in pre-GR periods, but it significantly flattened and became equal to $-0.022 + 0.016 = -0.006$ after the financial crisis. This suggests that the economy has

Table 3. *The Flattening of the Price Phillips Curve across the EMU11.*

	New Keynesian Phillips curve						Regional Phillips curve of Hazell *et al.* (2022)					
	(1) π^{CPI}	(2) π^{CPI}	(3) π^{CPI}	(4) π^{CPI}	(5) π^{CPI}	(6) π^{CPI}	(7) π^{NT}	(8) π^{NT}	(9) π^{NT}	(10) π^{NT}	(11) π^{NT}	(12) π^{NT}
$\hat{u}$	−0.0061***	−0.0221**	−0.0252**	−0.0054***	−0.0105***	−0.0139***	−0.0117***	−0.0231***	−0.0253***	−0.0148***	−0.0199***	−0.0173***
	(0.0023)	(0.0088)	(0.0102)	(0.0017)	(0.0033)	(0.0037)	(0.0012)	(0.0030)	(0.0031)	(0.0012)	(0.0024)	(0.0026)
$After_{GR} \times \hat{u}$		0.0157*	−0.0115		0.0062*	0.0240***		0.0139***	0.0035		0.0067**	0.0111***
		(0.0083)	(0.0174)		(0.0034)	(0.0074)		(0.0033)	(0.0058)		(0.0029)	(0.0041)
$After_{GR} \times Shift^{R}_{GR} \times \hat{u}$			0.0029***			0.0011*			0.0029***			0.0027***
			(0.0008)			(0.0006)			(0.0003)			(0.0003)
$After_{SDC} \times \hat{u}$			0.0001			0.0002			0.0184***			0.0024
			(0.0072)			(0.0052)			(0.0036)			(0.0040)
$After_{SDC} \times Shift^{R}_{SDC} \times \hat{u}$			0.0017**			0.0010*			0.0009**			0.0010**
			(0.0008)			(0.0006)			(0.0005)			(0.0004)
$After_{GR} \times Shift^{Manuf}_{GR} \times \hat{u}$			0.0001			−0.0014**			−0.0011***			−0.0017***
			(0.0005)			(0.0006)			(0.0003)			(0.0002)
$After_{SDC} \times Shift^{Manuf}_{SDC} \times \hat{u}$			−0.0002			−0.0019***			−0.0015***			−0.0011***
			(0.0009)			(0.0005)			(0.0004)			(0.0004)
Observations	748	748	748	748	748	748	748	748	748	748	748	748
FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes
IV	Ext.	Ext.	Ext.	Int.	Int.	Int.	Int.	Int.	Int.	Int.	Int.	Int.

Note: In columns (1) to (6), the unit of observation is consumer price inflation, measured as the year-on-year percentage change of the harmonised consumer price index and $\hat{u}$ is the unemployment gap measured as the deviation of the unemployment series from the NAIRU. In columns (7) to (12), the unit of observation is non-tradable inflation, measured as the year-on-year percentage change of the GDP deflator for non-tradable sectors and $\hat{u}$ is the one-year discounted sum of future unemployment in deviation from the NAIRU. $After_{GR}$ ($After_{SDC}$) is a dummy taking value one for periods after the GR (SDC) according to country-specific business cycle dates; $Shift^{R}_{GR}$ ($Shift^{R}_{SDC}$) is the shift in the occupational structure that occurred during the GR (SDC), i.e., the percentage change in the routine employment share between the peak and trough of the recession according to country-specific business cycle dates; $Shift^{Manuf}_{GR}$ ($Shift^{Manuf}_{SDC}$) is the shift in the share of value added from manufacturing and construction that occurred during the GR (SDC), i.e., the percentage change in the share of value added from these two sectors between the peak and trough of the recession according to country-specific business cycle dates. Online Appendix C.2 provides the description of instruments and first-stage statistics. The sample is composed of all countries that joined the EMU before 2002 (Luxembourg excluded). Data are quarterly. SEs are reported in parentheses. *, **, *** Significance at the 90%, 95% and 99% levels.

indeed undergone some structural change that has permanently affected the relationship between prices and unemployment.

Therefore, we now study to which extent this flattening is a consequence of the permanent change in the occupational composition of the labour market that occurred during each downturn, while taking into account all other potential structural changes that could have influenced the slope (i.e., the SDC, the permanent change in sectoral composition following each recession). For this reason, in column (3) we add all other interaction terms. Under this complete specification, we find that the occupational shift that occurred during each recession significantly explains part of the flattening of the PC. In detail, for the same level of unemployment gap, a 1% shift in the occupational structure that occurred during the GR flattens the PC by 0.003 just after the GR. However, the PC flattens out even further in countries where there was a bigger shift in the occupational structure during the SDC. In particular, a 1% occupational shift during the SDC leads to a further flattening by 0.002.

In light of this, we can reject H_0 for all $c = \{GR, SDC\}$ and state that the structural change that occurred in the labour market during the last two recessions had a role in the recent flattening of the PC. By how much? Now, we can evaluate the aggregate contribution of both occupational shifts on the flattening through a back-of-the-envelope calculation. By simply using our estimates from columns (2) and (3), we can say that the occupational shifts can jointly explain $(0.003 + 0.002)/0.016 \approx 30\%$ of the overall flattening of the Phillips curve from the end of the GR onward.

To corroborate these results, we also repeat our analysis with internal instruments. In fact, we know that the transmission of monetary policy on unemployment and inflation expectations is heterogeneous across countries. However, the instruments used so far are common to all countries. Therefore, it is important to check if our results also hold when using country-specific IVs. In particular, we instrument the current unemployment gap and all its interactions with the previous six-month realisation ($\hat{u}_{i,t-2}$, $\hat{u}_{i,t-2} \times After_{i,c}$, $\hat{u}_{i,t-2} \times Shift^{R}_{i,c}$, $\hat{u}_{i,t-2} \times Shift^{Manuf}_{i,c}$ for $c \in \{GR, SDC\}$ and $i \in \{\text{Austria}, \dots, \text{Spain}\}$), and expectations with the one-year moving average of past inflation expectations. Hence, our complete model is now perfectly identified and instruments—now being country specific—are stronger (F-test $= 36.94$). As from columns (4) and (5), also under this instrumentation, we find that the PC has flattened just after the GR. Column (6) confirms that the occupational shifts during the GR and SDC jointly explain up to $(0.001 + 0.001)/0.006 \approx 30\%$ of the estimated flattening.

Although these results show that there was a flattening and the occupational shifts contributed to it, this approach has several limitations. As explained in Hazell *et al.* (2022), HICP inflation can spillover across countries such that unemployment is less correlated with inflation at the country level. At the same time, inflation expectations (in particular in the long run) should be constant and common across all members of a monetary union as the ECB operates under an aggregate inflation mandate. The estimation of our NKPC is exposed to these shortfalls. In light of this, it is now important to assess the validity of our results by applying the methodology of Hazell *et al.* (2022). In fact, there are three main gains from this approach: (*i*) non-tradable (NT) prices are more sensitive to regional unemployment than the aggregate HICP is to aggregate unemployment; (*ii*) variation in long-run inflation expectations (due to the behaviour of the Central Bank) can be controlled for by using time fixed effects; (*iii*) other differences across regions, as long as these differences are constant over time, will be absorbed by country fixed effects.

In light of this, we modify (4) accordingly. In particular, now the dependent variable is NT inflation. As in Hazell *et al.* (2022), unemployment is now measured as the expected discounted sum of future unemployment in deviation from its long-run equilibrium level. Given the short time length of our data, we use a four-quarter discounted sum of future unemployment in deviation from the NAIRU. Similarly, the discounted four-quarter sum of future NT prices replaces inflation expectation. Additionally, we drop the energy and import price index from controls. Now, endogenous variables are instrumented with the fourth lag of past NT prices, the fourth lag of unemployment and its interaction with occupational/sectoral shifts and post-recession dummies (F-test $= 25.05$).

Column (7) shows that, in the 2002–18 period, the slope of the PC is -0.012. Column (8) shows that there was a significant flattening after the GR, with the slope moving from -0.023 to $-0.023 + 0.014 = -0.009$. When looking at the decomposition of the flattening in column (9), we again find a significant contribution of each occupational shift: they can explain it up to $(0.003 + 0.001)/0.014 \approx 28\%$. In light of this, also under this strategy, we conclude again that (*i*) there was a flattening and (*ii*) the occupational shifts during both recessions can explain part of the weakening of the PC.

It is important to note that, so far, we have not been using time fixed effects, as in Hazell *et al.* (2022), but business cycle phases ($phase_{i,t}$). To make sure that this difference does not drive our result, we re-estimate the model by replacing our phase dummy with time fixed effects. Results are shown in columns (10)–(12). Also under this specification (F-test $= 39.70$), we find the flattening to be significant in post-GR periods. When decomposing the change in slope, we find that most of the flattening actually already occurred after the GR and the occupational shifts in both recessions significantly contributed to it by $(0.003 + 0.001)/0.007 \approx 55\%$.[3]

In the Online Appendix, we perform a series of robustness checks. By using the same instrumentation of this section, in Online Appendix C.4, we estimate model (4) with quarter-on-quarter HICP and quarter-on-quarter NT inflation. Results hold for the NKPC estimation (although the estimate for $\tilde{\kappa}_{SDC}$ is not significant). In Online Appendix C.5, we also check another specification often used in the empirical literature: we estimate the neoclassical PC, with past expectations on current inflation used as (exogenous) control. Also in this case, results hold under external instruments (the estimate for $\tilde{\kappa}_{SDC}$ is not significant when using internal instruments). Finally, in Online Appendix C.7, we adopt a narrative-identifying approach, as in Siena and Zago (2022). We cross-validate our results by exploiting reforms of employment protection liberalisations. In fact, this type of reform generates exogenous variation in the routine employment share and it therefore represents a good instrument to study the (endogenous) relationship between the slope of the PC and polarisation. Also in this case we find that, when the routine employment share declines as result of the reform, the slope of the PC becomes smaller.

2.4. *Any Implication for the Wage PC?*

If the occupational structure matters for prices, it should also matter for wages. By applying the exact same two models and IV strategy of the previous section, we check whether the results for the price PC also hold when looking at the wage PC. With respect to the price PC, results are weaker and the slope of the wage PC appears more unstable over time, as found by the previous literature. Yet, evidence suggests that the occupational shift that

[3] In Online Appendix C.6, we also implement the methodology of Hazell *et al.* (2022) with monetary policy shocks as instruments, but without time fixed effects (as the instrument is time varying and common across countries). Our results are robust.

Table 4. *The Behaviour of the Wage Phillips Curve across the EMU11.*

	New Keynesian Phillips curve		Regional Phillips curve	
	(1) $\pi^{W,BE}$	(2) $\pi^{W,BE}$	(3) $\pi^{W,NT}$	(4) $\pi^{W,NT}$
$\hat{u}$	−0.0215	0.0328*	−0.0098*	−0.0028
	(0.0374)	(0.0174)	(0.0053)	(0.0048)
$After_{GR} \times \hat{u}$	0.0183	−0.1079***	−0.0335***	−0.0219***
	(0.0641)	(0.0288)	(0.0099)	(0.0075)
$After_{GR} \times Shift^{R}_{GR} \times \hat{u}$	0.0052*	0.0037*	0.0011**	0.0011**
	(0.0028)	(0.0022)	(0.0006)	(0.0005)
$After_{SDC} \times \hat{u}$	0.0018	0.0268*	0.0150**	−0.0076
	(0.0266)	(0.0138)	(0.0062)	(0.0073)
$After_{SDC} \times Shift^{R}_{SDC} \times \hat{u}$	−0.0008	−0.0037	−0.0009	−0.0008
	(0.0029)	(0.0032)	(0.0008)	(0.0008)
$After_{GR} \times Shift^{Manuf}_{GR} \times \hat{u}$	−0.0045**	−0.0010	0.0008*	0.0002
	(0.0019)	(0.0023)	(0.0005)	(0.0004)
$After_{SDC} \times Shift^{Manuf}_{SDC} \times \hat{u}$	−0.0012	0.0047**	0.0005	0.0008
	(0.0032)	(0.0023)	(0.0007)	(0.0007)
Observations	748	748	748	748
FEs	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Time FEs	No	No	No	Yes
IV	Ext.	Int.	Int.	Int.

Note: In columns (1) and (2), the unit of observation is wage inflation, measured as the year-on-year percentage change of the labour-cost index (salary and wages) for the business economy and $\hat{u}$ is the unemployment gap measured as the deviation of the unemployment series from the NAIRU. In columns (3) and (4), the unit of observation is non-tradable wage inflation, measured as the year-on-year percentage change of the labour-cost index (salary and wages) for non-tradable sectors, and $\hat{u}$ is the one-year discounted sum of future unemployment in deviation from the NAIRU. $After_{GR}$ ($After_{SDC}$) is a dummy taking value one for periods after the GR (SDC) according to country-specific business cycle dates; $Shift^{R}_{GR}$ ($Shift^{R}_{SDC}$) is the shift in the occupational structure that occurred during the GR (SDC), i.e., the percentage change in the routine employment share between the peak and trough of the recession according to country-specific business cycle dates; $Shift^{Manuf}_{GR}$ ($Shift^{Manuf}_{SDC}$) is the shift in the share of value added from manufacturing and construction that occurred during the GR (SDC), i.e., the percentage change in the share of value added from these two sectors between the peak and trough of the recession according to country-specific business cycle dates. We refer the reader to Online Appendix C.2 for the description of the instruments and first-stage statistics. The sample is composed of all countries that joined the EMU before 2002 (Luxembourg excluded). Data are quarterly. SEs are reported in parentheses. *, **, *** Significance at the 90%, 95% and 99% levels.

occurred during the GR has contributed to a weakening of the relationship between wages and unemployment.

To show this, we begin with the estimation of the NKPC with the year-on-year percentage change of the labour-cost index (wages and salaries for the business economy) available on Eurostat. Column (1) of Table 4 shows the results when using external instruments. In this case, the slope of the wage PC is negative, but not significant. Yet, we still find a significant effect of the occupational shift that occurred during the GR. In column (2), we repeat our estimation with internal instruments. Although the slope moves from positive to negative between periods before and after the GR, countries that experienced larger shifts during the GR witnessed a weakening of the wage PC after this downturn. Then, we apply the methodology of Hazell *et al.* (2022). We use non-tradable wages, constructed using wages from non-tradable sectors. Columns (3) and (4) respectively show results without and with time fixed effects. Using these two models, the wage PC is more well behaved and we do find that the occupational shift of the GR matters.

3. Theoretical Framework and Micro-Foundation

Why does job polarisation flatten the price Phillips curve? We now move to the theoretical analysis with three objectives. First, to show that labour market characteristics matter for the slope of the price PC. Second, to illustrate how changing the composition of the labour market can flatten the price PC through an increase in labour market *fluidity*. Third, to provide micro-evidence on how job polarisation increased overall *fluidity*. Finally, in light of the theoretical model, we show how technological change (the main driver of polarisation) and *fluidity* interact and affect the level of inflation.

3.1. *The Model*

We start by introducing a simple New Keynesian model with unemployment and search-and-matching frictions. We take the model proposed by Blanchard and Galí (2010) and here present the main features key for our purpose and refer to the original paper for all additional details. There is a continuum of members in a representative household that consumes a differentiated basket of imperfectly substitutable goods, supplies labour $0 \leq N_t \leq 1$ and discounts the future at rate β. The household maximises the expected utility

$$E_0 \sum \beta^t \left(\log C_t - \chi \frac{N_t^{1+\phi}}{1+\phi} \right),$$

where $C_t = (\int_0^1 C_t(z)^{(\varepsilon-1)/\varepsilon} dz)^{\varepsilon/(\varepsilon-1)}$ and ϕ is the inverse Frisch labour supply elasticity. There is a continuum of firms $i \in [0, 1]$ producing a differentiated final good $Y_t(i)$:

$$Y_t(i) = X_t(i).$$

Here $X_t(i)$ is the quantity of the intermediate good bought by firm i from the large number of identically and perfectly competitive intermediate firm producers $j \in [0, 1]$. Intermediate firms produce the homogeneous good X with a linear production function $X_t(j) = A_t N_t$, where A_t is an exogenous process depicting technology. Employment decisions are taken by the intermediate firm j and are described by the following labour demand accumulation equation:

$$N_t(j) = (1 - \delta)N_{t-1}(j) + H_t(j).$$

Here $\delta \in (0, 1)$, a crucial parameter for our analysis, determines the exogenous separation rate and $H_t(j)$ measures the workers hired in period t. Parameter δ can be interpreted as the fraction of workers that had a job at $t - 1$, but are not working any longer at the beginning of period t and need to find a job. Therefore, δN_{t-1} will be the increase in the stock of people unemployed between periods $t - 1$ and t. This drives the necessity to define two 'types' of unemployment: U_t, ex ante unemployment (i.e., unemployment at the beginning of the period) and u_t, ex post unemployment (i.e., unemployment, after hiring, at the end of period t). Therefore, $U_t = u_{t-1} + \delta N_{t-1}$.

As long as our parametrisation guarantees that the benefit from an extra hour of work is higher than its marginal rate of substitution at full employment (i.e., $W_t > \chi C_t$), then the labour market is characterised by full participation. This condition implies that $u_t = 1 - N_t$ and $U_t = 1 - (1 - \delta)N_{t-1}$. As a consequence, the flow of newly hired workers in period t can be rewritten as $H_t = \int_0^1 H_t(j)dj = N_t - (1 - \delta)N_{t-1}$.

We now define labour market tightness x_t (or the job-finding rate). This measures the ratio of aggregate hires to unemployment $x_t = H_t/U_t \in [0, 1]$, capturing the probability of being hired in period t. Hiring is costly and the cost is a positive function of market tightness and vacancies are filled any time the hiring cost is paid:

$$G_t = A_t B x_t^{\alpha}.$$

Here $\alpha \geq 0$ and $B > 0$, where B is the parameter governing matching efficiency. We also follow Blanchard and Galí (2010) to introduce real wage rigidities in a simple manner, assuming a wage schedule of the form

$$W_t = \left[\frac{1}{\mu} - (1 - \beta(1-\delta))Bx^{\alpha}A^{\gamma}\right]A_t^{1-\gamma} = \Theta A_t^{1-\gamma},$$

where x and A are respectively the steady state value of market tightness and the unconditional mean of productivity A_t. This implies that, when $\gamma = 0$, our wage will correspond to Nash bargaining, while, when $\gamma = 1$, we have rigid wages, as in Hall (2005). Here μ is the gross desired markup of the final good producer.

To complete the model, we need to introduce the final firm's price setting behaviour. Prices are rigid and follow the pricing formulation of Calvo (1983): each period the final good producer has probability $(1-\theta)$ to reset prices, while the remaining producers θ keep their prices unchanged. The optimal price setting rule turns to be the standard

$$E_t\left\{\sum_{k=0}^{\infty}\theta^k Q_{t,t+k}Y_{t+k,t}(P_t^* - \mu P_{t+k}MC_{t+k})\right\} = 0,$$

where $P_t = [(1-\theta)(P_t^*)^{1-\varepsilon} + \theta(P_{t-1})^{1-\varepsilon}]^{1/(1-\varepsilon)}$, P_t^* denotes the price picked by the firm able to reset prices, $Y_{t+k,t}$ is the level of output in period $t+k$ for a firm last able to reset prices in period t, $Q_{t.t+k}$ is the stochastic discount factor common across all households $\equiv \beta^k(C_t/C_{t+k})(P_t/P_{t+k})$ and, finally, MC_{t+k} is the real marginal cost. The latter is given by the relative prices of intermediate good producers $P_t^I = MC_t^n$, given the assumption of perfect competition, and the aggregate consumption price level P_t:

$$\begin{aligned} MC_t &= \frac{P_t^I}{P_t} \\ &= \frac{1}{A_t}\left[W_t + G_t - \beta(1-\delta)E_t\left\{\frac{C_t}{C_{t+1}}G_{t+1}\right\}\right] \\ &= \Theta A_t^{-\gamma} + Bx_t^{\alpha} - \beta(1-\delta)E_t\left\{\frac{C_t}{C_{t+1}}\frac{A_{t+1}}{A_t}Bx_{t+1}^{\alpha}\right\}. \end{aligned}$$

In this formulation of the marginal cost function lies the crucial difference with respect to the standard NK model. It is immediate to see how labour market frictions and the real wage rigidity appear and affect the marginal cost function.[4]

3.2. *The Price Phillips Curve*

Now, we focus on the relationship between inflation and unemployment, the NK Phillips curve, in log deviations from a zero inflation steady state. We denote with lower-case letters with

[4] Interestingly, this feature of the model highlights how the relationship between inflation and the marginal cost is not affected by labour market characteristics, as it depends only on the frequency of price adjustments. In Online Appendix C.1 we show, in fact, that the elasticity of inflation to the marginal cost did not change over the period considered.

hats the log deviations of the corresponding upper-case variables from their steady-state values. Steady-state values are displayed with their corresponding letter without time subscripts (e.g., g is the steady-state value of g_t). See Online Appendix D.1 for a full derivation. Note that, to be consistent with our empirical formulation, we denote by κ_0 the negative elasticity of inflation to current output gap fluctuations. We have

$$\pi_t = \beta E_t\{\pi_{t+1}\} + \kappa_0 \hat{u}_t + \kappa_L \hat{u}_{t-1} + \kappa_F E_t\{\hat{u}_{t+1}\} - \lambda\Phi\gamma\hat{a}_t, \tag{5}$$

where

$$\begin{aligned}
\lambda &= \frac{(1-\beta\theta)(1-\theta)}{\theta},\\
\Phi &\equiv \frac{\mu W}{A} = 1-(1-\beta(1-\delta))g\mu,\\
\kappa_0 &\equiv -\frac{\lambda h_0}{1-u}, \qquad \kappa_1 \equiv \frac{\lambda h_L}{1-u}, \qquad \kappa_F \equiv \frac{\lambda h_F}{1-u},\\
h_0 &\equiv \frac{\alpha g\mu}{\delta}[1+\beta(1-\delta)^2(1-x)] + \beta(1-\delta)g\mu(\xi_1-\xi_0),\\
h_1 &\equiv \frac{\alpha g\mu}{\delta}(1-\delta)(1-x) - \beta(1-\delta)g\mu\xi_1,\\
h_F &\equiv \beta(1-\delta)g\mu\left(\frac{\alpha}{\delta}-\xi_0\right),\\
\xi_0 &\equiv \frac{1-g(1+\alpha)}{1-\delta g}, \qquad \xi_1 \equiv \frac{g(1-\delta)(1+\alpha(1-x))}{1-\delta g}.
\end{aligned}$$

What comes out clearly from this formulation is that the slope of the price Phillips curve, κ_0, depends on labour market characteristics, in particular on the separation rate δ, market tightness x and the curvature of the cost function α.

Before analysing the characteristics of this formulation in detail, we derive a simplified version of it. As in Blanchard and Galí (2010), we assume that both the hiring cost with respect to output g and the separation rate δ are small. We also assume that technology follows a stationary first-order autoregressive (AR(1)) process, with auto-regressive parameter $\rho_a \in [0, 1)$. These assumptions allow us to simplify the approximated Phillips curve as

$$\widehat{\pi}_t = \kappa\widehat{u}_t - \kappa(1-\delta)(1-x)\widehat{u}_{t-1} - \Psi\gamma\widehat{a}_t, \tag{6}$$

where

$$\kappa \equiv -\frac{\alpha g\mu\lambda}{\delta(1-u)} \quad \text{and} \quad \Psi \equiv \frac{\lambda\Phi}{1-\beta\rho_a},$$

which is easy to study analytically. We start by noting that the slope of the Phillips curve can be written as a function of standard Calvo parameters λ, the markup μ and labour market characteristics (i.e., the level of equilibrium unemployment (u), the separation rate (δ), the market tightness condition (x) and the parameters of the hiring cost function (B, α)):

$$\frac{\alpha g\mu\lambda}{\delta(1-u)} = \frac{\alpha Bx^{\alpha}\mu\lambda}{\delta(1-u)} = \frac{\alpha B(H/U)^{\alpha}\mu\lambda}{\delta(1-u)} = B\mu\lambda\frac{\alpha(\delta N/[1-(1-\delta)N])^{\alpha}}{\delta N}.$$

In order to find an analytical expression of the effects of the labour market composition on the slope of the Phillips curve, first we need to take a stand on the effect of polarisation on the

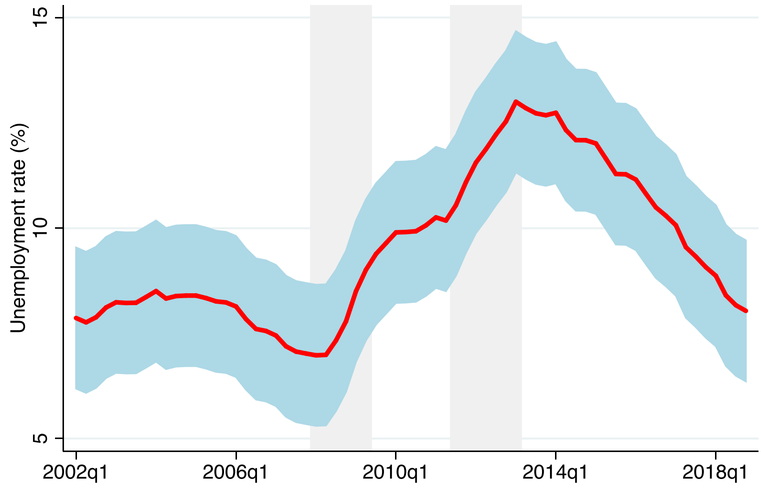

Fig. 8. *EMU Unemployment Rate.*
Note: This figure plots the evolution of the unemployment rate (for the labour force in the 15–74 age bracket) across those countries that joined the EMU before 2002 (Luxembourg excluded) along with the 95% confidence interval. The two vertical shaded areas respectively indicate the periods of the Great Recession and of the Sovereign Debt Crisis as defined by the CEPR Business Cycle Committee. Data are quarterly and span from 2002q1 to 2018q4.

equilibrium level of unemployment. Supported by the overall dynamic of unemployment in the European Monetary Union (Figure 8 shows that the unemployment level converged back to its pre-recession level), we maintain steady-state unemployment $u = 1 - N$ constant. This implies that every movement in the separation rate will result in an adjustment of the job-finding rate to maintain constant the equilibrium level of unemployment. In particular, an increase in the separation rate implies an increase in tightness. Given that $x = \delta N/[1 - (1-\delta)N]$, we have $\partial x/\partial\delta = N(1-N)/[1-(1-\delta)N]^2 > 0$.

Let us define, as in Blanchard and Galí (2010), a *fluid* labour market as one characterised by high separation and a high job-finding rate. How does the Phillips curve slope change in response to an increase in the fluidity of the labour market? How does an increase in the separation rate affect the slope of the Phillips curve? We can easily show that

$$\frac{\partial\kappa}{\partial\delta} = -\mu\lambda B\frac{\alpha(\delta N)^{(\alpha-2)}N}{(1-(1-\delta)N)^{1+\alpha}}[(\alpha-1)(1-(1-\delta)N) - \alpha\delta N]$$

is positive (i.e., an increase in δ makes the slope greater, and therefore flatter, given that $\kappa < 0$) when $\alpha < (1 - N + \delta N)/(1-N) = U/u$. Note that this condition is always satisfied when $\alpha < 1$, an empirical (see Pissarides and Petrongolo, 2001 and Barnichon and Figura, 2015) and theoretical regularity (Shimer, 2005).[5] Note that this result does not depend on the assumption of a constant steady-state unemployment rate. Actually, keeping this rate constant to match equilibrium unemployment data reduces the effect of δ on the slope of the PC. As shown in

[5] Equation (6) also shows that the lagged unemployment gap could play an important role for the relationship between inflation and unemployment. Under reasonable parametrisation—such as that presented below—the steady-state relationship between inflation and unemployment $\kappa[1-(1-x)(1-\delta)]$ is also dampened as fluidity increases. In fact, it can be easily shown numerically that $\partial\kappa[(1-\delta)(1-x)-1]/\partial\delta > \kappa(1-x) + \partial x\kappa(1-\delta)/\partial\delta$.

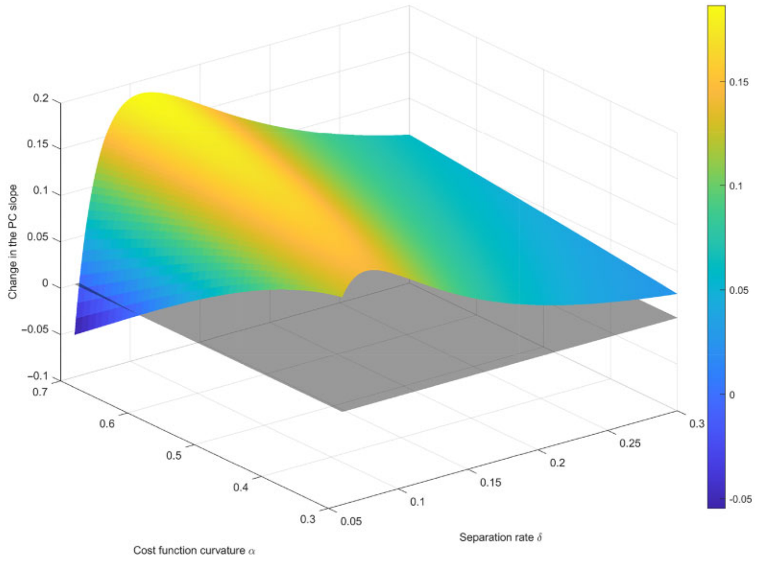

Fig. 9. *Derivative of the Slope of the Price Phillips Curve—Full Model.*
Note: This figure plots the partial derivative of the price Phillips curve with respect to the separation rate $\partial\kappa_0/\partial\delta$, computed from (5), for different starting values of the separation rate δ and the curvature of the cost function α.

Online Appendix D.4, allowing steady-state unemployment to vary in response to an increase in fluidity flattens the Phillips curve even further.

We now need to investigate whether this result generalises to the full extended model. We proceed with calibrating the model, considering each period a quarter. For preferences, price setting and wage rigidity, we take the parameters used in Blanchard and Galí (2010): $\beta = 0.99$, $\phi = 1$, $\epsilon = 6$ and $\theta = 0.7457$. For the other parameters, we look at evidence from the EMU. We estimate the hiring cost (labour costs other than wages and salaries, from Eurostat) as a fraction of GDP to be 1.512%, implying a matching efficiency B equal to 0.3297. The equilibrium level of unemployment is set to 8%, the average value for the EMU11 in the pre-recession period. Regarding labour market parameters, we take a large range of possible values, considering an aggregate $\delta \in [0.05, 0.3]$, implying that $x \in [0.36, 0.78]$, and a curvature parameter $\alpha \in [0.3, 0.7]$. Figure 9 shows the sign of the derivative of the slope of the Phillips curve for different values of δ and α. The relationship is highly non-linear, but positive for most of the considered subset. The exception is when the separation rate is extremely low and α quite high, which are unrealistic values for these parameters. Therefore, even in the extended model, a higher separation rate leads to a flatter price PC.

What is the economic intuition behind this result? Higher fluidity reduces the elasticity of marginal costs to economic conditions (e.g., market tightness) such that employers adjust more the stock of employment rather than wages. This happens because the labour demand becomes more elastic as employers can lay off and hire workers more easily. In other words, the labour demand becomes flatter. Therefore, for a higher δ, the elasticity of wages to an aggregate shock will be smaller: the marginal cost will be more stable and therefore also prices. As a result, the relationship between prices and unemployment will be weaker.

3.3. *Job Polarisation and Labour Market Fluidity*

3.3.1. *Labour market fluidity versus price stickiness*

As shown in the previous section, a higher separation rate δ leads to a more *fluid* labour market and a flattening of the PC. If job polarisation affects the slope of the PC through this channel, in the data we should observe some heterogeneity in the separation rate across jobs such that employment relocation into non-routine occupations would lead to higher fluidity. This would corroborate the idea that the transition to a more fluid labour market—through the disappearance of routine jobs—has implication for the observed flattening of the PC. In order to analyse if this is the case, first we build a measure for δ by job. In particular, in line with the methodology of Shimer (2012) and Hobijn and Şahin (2009), we collect country-level (Eurostat) data on unemployment composition by duration and last occupation. This allows us to approximate the timing and size of flows from each occupation to unemployment. Then, we normalise each job-specific flow from employment to unemployment by the level of job-specific employment in the previous period and make minor corrections for the potential measurement errors rising from the fact that employment and unemployment compositions are trendy. Hence, we obtain three job-specific separation rates such that their sum—weighted by the employment share of each occupation—equates the aggregate separation rate in the economy (see Online Appendix E.1 for details).

Figure 10(a) shows the cross-country mean separation rate conditional on the previous job (along with the 95% confidence interval).[6] The average separation rate of non-routine jobs is significantly higher than the rate of routine ones. In particular, the average separation rates for abstract and manual workers are respectively 4% and 6%, whereas the routine market exhibits a separation rate equal to 3%.

As shown in Online Appendix E.3, each job-specific separation rate—at the net of recession periods—is very stable over time. Consequently, since the trend of job polarisation implies employment relocation towards non-routine occupations, the aggregate separation rate should also increase due to the increasing weight of abstract and manual jobs in the labour market. In Figure 10(b) we show that this is indeed the case. On average, the aggregate separation rate across EMU11 members moved from 3.5% in 2002 to 4%.[7] Using the calibrated model of Section 3.2, this change in δ would imply a flattening of the PC of 13.2%.

If a higher separation rate (as driven by the process of polarisation) is important for the slope of the PC, we should observe that countries that experienced a bigger increase in separation in the long run are the same countries that had a bigger flattening of the PC. We check this fact in Figure 10(c) where we plot the long-run change of the coefficient of the price PC—estimated following Stock and Watson (2019) in periods before and after the GR—on the long-run change in the separation rate. The correlation is 0.24, i.e., when a labour market becomes relatively more fluid, the relationship between price inflation and unemployment becomes weaker.

[6] The mean is computed by considering only periods before the GR and after the SDC (according to country-specific business cycle dates). See Online Appendix E.1 for details and Online Appendix E.3 for the cross-country dynamic of each job-specific separation rate.

[7] As a theoretical consequence of a higher separation rate for abstract and manual jobs, the hiring rate x should also be higher for these occupations. In Online Appendices E.2 and E.3, we show that this is the case, even if the aggregate hiring rate has slightly declined, in particular after the SDC. For an unemployment rate close to pre-GR levels in post-SDC periods, this suggests an aggregate decline in the matching efficiency parameter B, which we have instead assumed to be constant. An indirect way to test such deterioration in the matching efficiency would be to check whether the Beveridge curve has shifted out in recent years. We know that this is the case. The latest data on the Beveridge curve are available here.

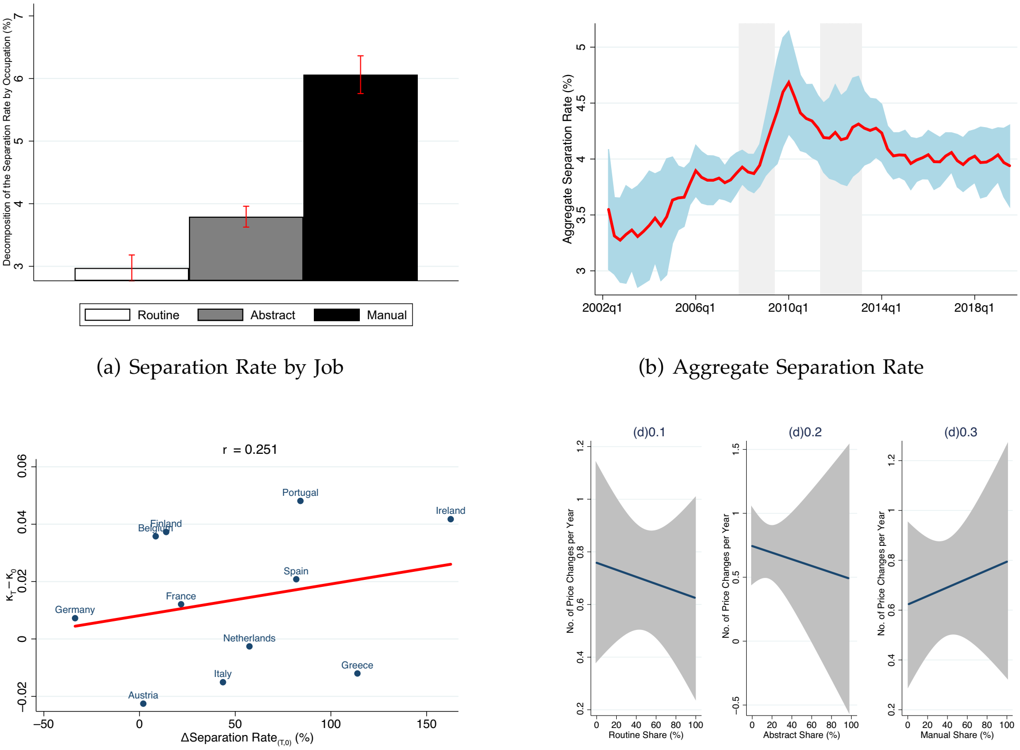

(a) Separation Rate by Job

(b) Aggregate Separation Rate

(c) Slope of the PC versus Separation Rate

(d) Calvo Parameter

Fig. 10. *Separation Rates and Calvo Parameter by Occupation across the EMU11.*
Note: Panel (a) plots the mean separation rate (with the 95% confidence interval) by occupation (routine, abstract or manual) across countries that joined the EMU before 2002 (Luxembourg excluded). Each job-specific separation rate is built by studying the (last) job composition and duration of the unemployment pool in each year and country in order to correctly identify the timing and size of flows from each job to unemployment (see Online Appendix E.1 for details). Each cross-country mean is computed considering only periods before the GR and after the SDC, according to country-specific business cycle dates. Panel (b) plots the aggregate separation rate along with the 95% confidence interval. Panel (c) plots the long-run change of slope of the Phillips curve—estimated following Stock and Watson (2019) in pre- and post-GR periods—on the long-run percentage change of the separation rate of each country. At the top of the graph, the correlation (ρ) between variables is reported along with its significance level. Panels (d)0.1 to (d)0.3 plot the linear relationship between the number of product price updates and the occupational workforce composition at the firm level. Data come from the three waves of the Wage Dynamics Survey of the ECB, which includes responses from firms in all countries that joined the EMU before 2002, but Finland. The three waves were conducted in 2008, 2009, 2014.

If these arguments point directly at labour market *fluidity* as a plausible explanation of the flattening of the PC, it is important to check if the traditional variable controlling price-update behaviour—the Calvo parameter λ—is somehow influenced by the composition of the labour market. In fact, we know that the process of job polarisation goes hand in hand with technological adoption, automation and offshoring, which can ultimately influence pricing behaviour in the product market (see, for example, Fueki and Maehashi, 2019; Aghion *et al.*, 2020 and Fujiwara and Zhu, 2020). For this reason, we exploit the Wage Dynamic Survey from the ECB to relate the frequency of the final good price update to the workforce composition of a sample of firms across the EMU11.[8] Figure 10(d) plots the linear relationship (and the 95% confidence interval) between the number of price changes (per year)—as reported by the management of the firm—and the workforce job composition of the firm. In Figure 10(d)0.1, the share of routine workers is on the x axis. Although the relationship is negative, there is no statistical difference in the frequency of the price update between firms fully composed of routine workers and firms fully composed of non-routine workers. If, on the contrary, such heterogeneity would be true then the decline of non-routine workers in the economy should lead to lower price stickiness, higher λ and, all else equal, a steeper slope of the PC, which is not the case in the data. We obtain similar (non-significant) evidence when considering the share of abstract workers on the x axis, as in Figure 10(d)0.2. The relationships turn positive when considering the share of manual workers on the x axis, as in Figure 10(d)0.3, but, also in this case, there is no significant difference between firms rich or poor of manual workers.

All in all, this evidence proves that employment relocation from less to more fluid occupations—as triggered by the process of job polarisation—is indeed an important channel to rationalise the observed flattening of the PC in recent years.

3.3.2. *Further anecdotal evidence on labour market fluidity*

The heterogeneity in the separation rate can also be explained by different labour market regulations and working arrangements across jobs. In this section, we introduce further evidence to corroborate the idea that non-routine occupations are more *fluid*.

The measure under consideration is the probability for a worker to have more than one job (i.e., more than one employer) when employed in a specific occupation. As shown in Figure 11, abstract workers exhibit a higher (average) probability of having multiple contractual arrangements with multiple employers (around 4.7%), whereas the probability for all other workers is statistically smaller (manual 3.8% and routine 2.8%). This evidence suggests that abstract employment is more uncertain since it depends on short-term contracts and multiple employers. This is particularly true for those abstract workers—e.g., designers, architects, lawyers, etc.—whose work arrangement depends more on the delivery of a specific service (service-based employment) or project (project-based employment) rather than on a continuous and binding relationship with a single employer. As explained in Blanchard and Landier (2002), all of this increases the turnover rate and dynamism of the labour market, but at the expense of more frequent unemployment spells.

3.4. *Cross-Validation: Fluidity versus Technological Change*

The literature on polarisation is mostly grounded on the role of technological change for the relocation of workers from less productive (routine) to more productive (non-routine) jobs. In

[8] This survey was conducted in three ways (2008, 2009, 2014) and asks the representative manager of a company some price-related questions. For example, if the management has recently changed prices of the final product and how many times prices were changed in a year. Moreover, the survey asks also the share of workers employed in a routine, abstract or manual occupation. Considering only those responses to questions on both price updates and workforce composition, we end up with a sample of 3,325 firms spread over ten of the EMU11 countries considered (no data are available for Finland). Click here for more information on the survey and variable constructions.

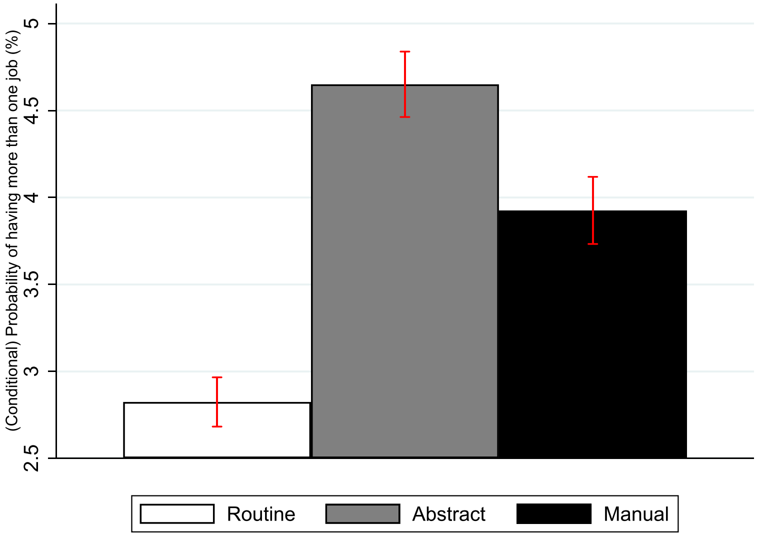

Fig. 11. *Probability of Multiple Jobs Worked.*
Note: This figure plots the mean (conditional) probability of having multiple jobs (with the 95% confidence interval) for workers currently employed in a specific occupation (routine, abstract or manual). See Online Appendix E.4 for details.

particular, as discussed in Acemoglu and Restrepo (2020), technological change and adoption typically co-move with the process of polarisation since investment in new technologies and inventions better complement non-routine workers. This is what we also observe for the EMU11. In Figure 12, we plot the dynamic of the share of investment in ICT technology and innovation, a proxy for non-routine-biased technology adoption and productivity enhancement efforts. As evident, just after the GR, (non-routine) productivity enhancement investment deviated from its pre-recession trend.

How does this phenomenon affect inflation? In the context of the model of Blanchard and Galí (2010), an increase in technology should have temporary disinflationary effects. This is captured by $-\Psi\gamma\hat{a}_t$ in the analytical PC of (6). However, if job polarisation comes contemporaneously with higher productivity (i.e., $\hat{a}_t > 0$), the acceleration of polarisation would also imply an increase in fluidity ($\Delta\delta > 0$). The latter dampens the disinflationary effect of technological change. In fact, in the theoretical PC, the level of inflation is also governed by the parameter Ψ, which is a decreasing function of δ. Hence, in a cross-country comparison, those with similar increases in productivity, but larger increases in δ, should have relatively higher inflation.

We can check this in the data. This will corroborate the idea that the main forces explaining the level of inflation (fluidity versus technology) are actually in place and operate as the theory predicts. To do so, we compare the cross-country level of inflation in the long run: before the GR and after the SDC. Then, we study how the change in the level of inflation can be explained by a heterogeneous increase in technology adoption and a separation rate in the long run. Formally,

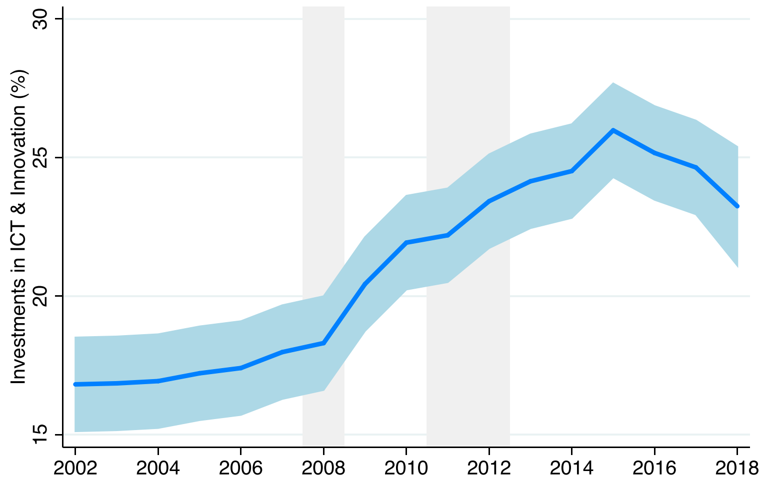

Fig. 12. *Productivity Enhancement Investment.*
Note: This figure plots the evolution of investment in ICT and innovation (as a share of total investment) across those countries that joined the EMU before 2002 (Luxembourg excluded) along with the 95% confidence interval. This investment share is measured as the sum of the investment in ICT equipment, computer and software database and intellectual property products. The two vertical shaded areas respectively indicate the periods of the Great Recession and of the Sovereign Debt Crisis as defined by the CEPR Business Cycle Committee. Data are from Eurostat, they are at annual frequency and span from 2002q1 to 2018q4.

we estimate the regression

$$\pi_{i,t}^{Core} = \underset{(0.055)}{2.033^{***}} - \underset{(0.112)}{0.905^{***}} After_{i,SDC} + \underset{(0.002)}{0.004^{***}} After_{i,SDC} \times \Delta\delta_i - \underset{(0.001)}{0.006^{***}} After_{i,SDC} \times \Delta Tech_i,$$

$$R^2 = 0.326, \quad n = 517.$$

In periods before the GR, the average inflation across EMU11 countries was 2.03%. On the other hand, the level of inflation decreased on average by almost 1 percentage point in periods after the SDC, i.e., there was disinflation. Such a phenomenon is mitigated in countries experiencing a larger percentage increase in the aggregate separation rate between 2002 and 2018 ($\Delta\delta_i > 0$), whereas it exacerbates in countries experiencing a larger percentage increase in technology adoption in the same years ($\Delta Tech_i > 0$).

In light of this, we conclude that the role of technological change and fluidity in the data operates on the level of inflation as the theory predicts. This cross-validates the role of both channels (fluidity and technology) to explain how job polarisation can differently affect the slope of the PC (through fluidity only) and the level of inflation (through both fluidity and technology).

4. Conclusions

In the last 20 years, labour markets across the EMU have dramatically changed composition: the share of routine employment (clerical, craft and plant occupations) has shrunk in favour of abstract employment (professional, managerial occupations). At the same time, the same economies experienced a flattening of the price PC. This paper combines these two events and proves that occupational composition and differences across jobs have important implications for the structural relationship between unemployment and inflation.

In the empirical part of the paper, we demonstrate that countries experiencing bigger changes in the occupational structure exhibit a flatter price (and wage) PC. By exploiting the exogenous acceleration of polarisation induced by recessions, we show that changes in job composition that occurred during the Great Recession and Sovereign Debt Crisis are responsible for more than a fourth of the flattening of the PC observed between 2002 and 2018.

In the theoretical part of the paper, we propose a possible explanation: the transformation of labour market characteristics induced by job polarisation. Using the analytical properties of the model of Blanchard and Galí (2010), we prove that a key factor affecting the slope of the PC is the *fluidity* of the labour market, i.e., the rate at which workers separate from employers and find other jobs. Hence, we show that higher *fluidity* leads to a flatter PC as the labour demand becomes more elastic to wages.

We conclude by providing micro-evidence supporting the implications of our theoretical mechanism. The market of abstract jobs is on average more fluid than the market of routine jobs: it has a higher separation and hiring rate, and it makes more frequent use of temporary contracts and multiple-job arrangements. Therefore, the overall transition from routine to non-routine occupations has increased the overall fluidity of the labour market in the EMU. This has decreased the elasticity of prices to unemployment.

Politecnico di Milano, Italy
Banque de France, France

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Acemoglu, D. (2002). 'Directed technical change', *The Review of Economic Studies*, vol. 69(4), pp. 781–809.
Acemoglu, D., Aghion, P. and Violante, G.L. (2001). 'Deunionization, technical change and inequality', *Carnegie-Rochester Conference Series on Public Policy*, vol. 55(1), pp. 229–64.
Acemoglu, D. and Autor, D. (2011). 'Skills, tasks and technologies: Implications for employment and earnings', in (D. Card and O. Ashenfelter, eds.), *Handbook of Labor Economics*, vol. 4, pp. 1043–171, Amsterdam: Elsevier.
Acemoglu, D. and Restrepo, P. (2020). 'Robots and jobs: Evidence from US labor markets', *Journal of Political Economy*, vol. 128(6), pp. 2188–244.
Açkgöz, O.T. and Kaymak, B. (2014). 'The rising skill premium and deunionization', *Journal of Monetary Economics*, vol. 63(C), pp. 37–50.
Aghion, P., Antonin, C., Bunel, S. and Jaravel, X. (2020). 'What are the labor and product market effects of automation? New evidence from France', Discussion paper, Centre for Economic Policy Research.
Akerlof, G.A., Dickens, W.T., Perry, G.L., Gordon, R.J. and Mankiw, N.G. (1996). 'The macroeconomics of low inflation', *Brookings Papers on Economic Activity*, vol. 1996(1), pp. 1–76.
Altavilla, C., Brugnolini, L., Gürkaynak, R.S., Motto, R. and Ragusa, G. (2019). 'Measuring euro area monetary policy', *Journal of Monetary Economics*, vol. 108, pp. 162–79.
Ascari, G., Bonomolo, P. and Haque, Q. (2023). 'The long-run Phillips curve is ... a curve', 2023-37 CAMA Working papers.
Autor, D. (2008). 'Technological change and earnings polarization: Implications for skill demand and economic growth', Working paper, The Conference Board.

Autor, D.H., Dorn, D. and Hanson, G.H. (2013). 'The China syndrome: Local labor market effects of import competition in the United States', *The American Economic Review*, vol. 103(6), pp. 2121–68.
Autor, D.H., Dorn, D. and Hanson, G.H. (2015). 'Untangling trade and technology: Evidence from local labour markets', Economic Journal, vol. 125(584), pp. 621–46.
Autor, D.H., Katz, L.F. and Kearney, M.S. (2006). 'The polarization of the U.S. labor market', *American Economic Review*, vol. 96(2), pp. 189–94.
Autor, D.H., Levy, F. and Murnane, R.J. (2003). 'The skill content of recent technological change: An empirical exploration', *The Quarterly Journal of Economics*, vol. 118(4), pp. 1279–333.
Ball, L. and Mazumder, S. (2021). 'A Phillips curve for the euro area', *International Finance*, vol. 24(1), pp. 2–17.
Barnichon, R. and Figura, A. (2015). 'Labor market heterogeneity and the aggregate matching function', *American Economic Journal: Macroeconomics*, vol. 7(4), pp. 222–49.
Basso, H. and Rachedi, O. (2023). 'Robot adoption and inflation dynamics', Mimeo.
Benigno, P. and Ricci, L.A. (2011). 'The inflation-output trade-off with downward wage rigidities', *American Economic Review*, vol. 101(4), pp. 1436–66.
Bergholt, D., Furlanetto, F. and Vaccaro-Grange, E. (2023). 'Did monetary policy kill the Phillips curve? Some simple arithmetics', Discussion paper, Centre for Economic Policy Research.
Bernanke, B. (2007). 'Inflation expectations and inflation forecasting', Speech delivered at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute.
Berson, C., de Charsonville, L., Diev, P., Faubert, V., Ferrara, L., Guilloux-Nefussi, S., Kalantzis, Y., Lalliard, A., Matheron, J. and Mogliani, M. (2018). 'Does the Phillips curve still exist?', *Rue de la Banque*, vol. 56, pp. 1–5.
Blanchard, O. (2016). 'The Phillips curve: Back to the '60s?', *American Economic Review*, vol. 106(5), pp. 31–4.
Blanchard, O. and Galí, J. (2010). 'Labor markets and monetary policy: A New Keynesian model with unemployment', *American Economic Journal: Macroeconomics*, vol. 2(2), pp. 1–30.
Blanchard, O. and Landier, A. (2002). 'The perverse effects of partial labour market reform: Fixed-term contracts in France', Economic Journal, vol. 112(480), pp. F214–44.
Bobeica, E. and Jarociski, M. (2019). 'Missing disinflation and missing inflation: A VAR perspective', *International Journal of Central Banking*, vol. 15(1), pp. 199–232.
Bruine de Bruin, W., Vanderklaauw, W., Downs, J.S., Fischhoff, B., Topa, G. and Armantier, O. (2010). 'Expectations of inflation: The role of demographic variables, expectation formation, and financial literacy', *Journal of Consumer Affairs*, vol. 44(2), pp. 381–402.
Bulligan, G. and Viviano, E. (2017). 'Has the wage Phillips curve changed in the euro area?', *IZA Journal of Labor Policy*, vol. 6(1), pp. 1–22.
Calvo, G.A. (1983). 'Staggered prices in a utility-maximizing framework', *Journal of Monetary Economics*, vol. 12(3), pp. 383–98.
Cantore, C., Ferroni, F. and Len-Ledesma, M. (2020). 'The missing link: Monetary policy and the labor share', *Journal of the European Economic Association*, vol. 19(3), pp. 1592–620.
Ciccarelli, M., Garca, J.A. and Montes-Galdn, C. (2017). 'Unconventional monetary policy and the anchoring of inflation expectations', Working paper, European Central Bank.
Coibion, O. and Gorodnichenko, Y. (2015). 'Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation', *American Economic Journal: Macroeconomics*, vol. 7(1), pp. 197–232.
Daly, M.C. and Hobijn, B. (2014). 'Downward nominal wage rigidities bend the Phillips curve', *Journal of Money, Credit and Banking*, vol. 46(S2), pp. 51–93.
Daly, M.C., Hobijn, B. and Pyle, B. (2016). 'What's up with wage growth', FRBSF Economic Letter 2016-07.
Del Negro, M., Lenza, M., Primiceri, G.E. and Tambalotti, A. (2020). 'What's up with the Phillips curve?', *Brookings Papers on Economic Activity*, vol. 2020, pp. 301–57.
Deroose, M. and Stevens, A. (2017). 'Low inflation in the euro area: Causes and consequences', *Economic Review*, vol. I, pp. 111–25.
Dinlersoz, E. and Greenwood, J. (2016). 'The rise and fall of unions in the United States', *Journal of Monetary Economics*, vol. 83, pp. 129–46.
Dovern, J. and Kenny, G. (2017). 'The long-term distribution of expected inflation in the euro area: What has changed since the Great Recession?', Working paper, European Central Bank.
Draghi, M. (2015). 'Structural reforms, inflation and monetary policy', BIS central bankers speeches.
Faccini, R. and Melosi, L. (2023). 'Job-to-job mobility and inflation', Working paper, Federal Reserve Bank of Chicago.
Firpo, S., Fortin, N.M. and Lemieux, T. (2011). 'Occupational tasks and changes in the wage structure', Discussion paper, Institute of Labor Economics.
Fitzgerald, T., Jones, C., Kulish, M. and Nicolini, J.P. (2022). 'Is there a stable relationship between unemployment and future inflation?', Working paper, Red Nacional de Investigadores en Economa.
Foster, L., Grim, C. and Haltiwanger, J. (2016). 'Reallocation in the Great Recession: Cleansing or not?', *Journal of Labor Economics*, vol. 34(S1), pp. S293–331.
Fueki, T. and Maehashi, K. (2019). 'Inflation dynamics in the age of robots: Evidence and some theory', Working paper, Bank of Japan.
Fujiwara, I. and Zhu, F. (2020). 'Robots and labour: Implications for inflation dynamics', *BIS Papers*, vol. 11, pp. 41–9.

Gaggl, P. and Kaufmann, S. (2020). 'The cyclical component of labor market polarization and jobless recoveries in the US', *Journal of Monetary Economics*, vol. 116, pp. 334–47.

Galí, J. and Gambetti, L. (2019). 'Has the U.S. wage Phillips curve flattened? A semi-structural exploration', Discussion paper, Centre for Economic Policy Research.

Giannone, D., Lenza, M., Momferatou, D. and Onorante, L. (2014). 'Short-term inflation projections: A Bayesian vector autoregressive approach', *International Journal of Forecasting*, vol. 30(3), pp. 635–44.

Goos, M., Manning, A. and Salomons, A. (2009). 'Job polarization in Europe', *American Economic Review*, vol. 99(2), pp. 58–63.

Gorodnichenko, Y. and Candia, B. (2021). 'The inflation expectations of U.S. firms: Evidence from a new survey', Discussion paper, Centre for Economic Policy Research.

Guerrieri, L., Gust, C. and López-Salido, J.D. (2010). 'International competition and inflation: A New Keynesian perspective', *American Economic Journal: Macroeconomics*, vol. 2(4), pp. 247–80.

Hall, R.E. (2005). 'Employment fluctuations with equilibrium wage stickiness', *American Economic Review*, vol. 95(1), pp. 50–65.

Harding, D. and Pagan, A. (2002). 'Dissecting the cycle: A methodological investigation', *Journal of Monetary Economics*, vol. 49(2), pp. 365–81.

Hazell, J., Herreño, J., Nakamura, E. and Steinsson, J. (2022). 'The slope of the Phillips curve: Evidence from U.S. states', *The Quarterly Journal of Economics*, vol. 137(3), pp. 1299–344.

Hobijn, B. and Şahin, A. (2009). 'Job-finding and separation rates in the OECD', *Economics Letters*, vol. 104(3), pp. 107–11.

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

Jaimovich, N. and Siu, H.E. (2020). 'Job polarization and jobless recoveries', *Review of Economics and Statistics*, vol. 102(1), pp. 129–47.

Jarociński, M. and Karadi, P. (2020). 'Deconstructing monetary policy surprises—the role of information shocks', *American Economic Journal: Macroeconomics*, vol. 12(2), pp. 1–43.

Jorgenson, D.W. (2001). 'Information technology and the US economy', *American Economic Review*, vol. 91(1), pp. 1–32.

Kiley, M.T. (2015). 'Low inflation in the United States: A summary of recent research', FEDS Notes, Board of Governors of the Federal Reserve System (U.S.).

Leduc, S. and Wilson, D.J. (2017). 'Has the wage Phillips curve gone dormant?', FRBSF Economic Letter 2017-30.

Lombardi, M.J., Riggi, M. and Viviano, E. (2023). 'Workers' bargaining power and the Phillips curve: A micro-macro analysis', *Journal of the European Economic Association*, vol. 21, pp. 1905–43.

Mavroeidis, S., Plagborg-Møller, M. and Stock, J.H. (2014). 'Empirical evidence on inflation expectations in the New Keynesian Phillips curve', *Journal of Economic Literature*, vol. 52(1), pp. 124–88.

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

Mincer, J. and Danninger, S. (2000). 'Technology, unemployment, and inflation', Working paper, National Bureau of Economic Research.

Mishkin, F.S. (2011). 'Monetary policy strategy: Lessons from the crisis', Working paper, National Bureau of Economic Research.

Moretti, L., Onorante, L. and Zakipour Saber, S. (2019). 'Phillips curves in the euro area', Working paper, European Central Bank.

Moscarini, G. and Postel-Vinay, F. (2023). 'The job ladder: Inflation vs. reallocation', NBER Working Papers 31466.

Murphy, A. (2018). 'The death of the Phillips curve?', Working Paper 1801, Federal Reserve Bank of Dallas.

Natoli, F. and Sigalotti, L. (2017). 'A new indicator of inflation expectations anchoring', Working paper, European Central Bank.

Ng, M., Wessel, D. and Sheiner, L. (2018). 'The Hutchins center explains: The Phillips curve', *Brookings*, 21 August.

Nickel, C., Bobeica, E., Koester, G., Lis, E. and Porqueddu, M. (2019). 'Understanding low wage growth in the euro area and European countries', Working paper, European Central Bank.

Pace, F.D. and Hertweck, M. (2019). 'Labor market frictions, monetary policy, and durable goods', *Review of Economic Dynamics*, vol. 32, pp. 274–304.

Petrosky-Nadeau, N., Wasmer, E. and Weil, P. (2021). 'When Hosios meets Phillips: Connecting efficiency and stability to demand shocks', Discussion paper, Centre for Economic Policy Research.

Pfajfar, D. and Roberts, J.M. (2022). 'The role of expectations in changed inflation dynamics', *International Journal of Central Banking*, vol. 18(1), pp. 199–238.

Pfajfar, D. and Santoro, E. (2008). 'Asymmetries in inflation expectation formation across demographic groups', Working paper, University of Cambridge.

Pissarides, C.A. and Petrongolo, B. (2001). 'Looking into the black box: A survey of the matching function', *Journal of Economic Literature*, vol. 39(2), pp. 390–431.

Portier, F., Beaudry, P. and Hou, C. (2020). 'Monetary policy when the Phillips curve is locally quite flat', Discussion paper, Centre for Economic Policy Research.

Portier, F., Beaudry, P. and Preston, A. (2023). 'Some inference perils of imposing a Taylor rule', Discussion paper, Centre for Economic Policy Research.
Powell, J.H. (2018). 'Monetary policy and risk management at a time of low inflation and low unemployment', *Business Economics*, vol. 53(4), pp. 173–83.
Ravenna, F. and Walsh, C.E. (2008). 'Vacancies, unemployment, and the Phillips curve', *European Economic Review*, vol. 52(8), pp. 1494–521.
Ravenna, F. and Walsh, C.E. (2012). 'Screening and labor market flows in a model with heterogeneous workers', *Journal of Money, Credit and Banking*, vol. 44(s2), pp. 31–71.
Roberts, J.M. (2006). 'Monetary policy and inflation dynamics', *International Journal of Central Banking*, vol. 2(3), pp. 193–230.
Schmitt-Grohé, S. and Uribe, M. (2013). 'Downward nominal wage rigidity and the case for temporary inflation in the eurozone', *Journal of Economic Perspectives*, vol. 27(3), pp. 193–212.
Shimer, R. (2005). 'The cyclical behavior of equilibrium unemployment and vacancies', *American Economic Review*, vol. 95(1), pp. 25–49.
Shimer, R. (2012). 'Reassessing the ins and outs of unemployment', *Review of Economic Dynamics*, vol. 15(2), pp. 127–48.
Siena, D. (2021). 'The euro area periphery and imbalances: Is it an anticipation story?', *Review of Economic Dynamics*, vol. 40, pp. 278–308.
Siena, D. and Zago, R. (2022). 'Employment protection legislation matters for the Phillips curve', *Economics Letters*, vol. 220, 110883.
Speck, C. (2017). 'Inflation anchoring in the euro area', Working paper, European Central Bank.
Stock, J.H. and Watson, M.W. (2019). 'Slack and cyclically sensitive inflation', Working paper, National Bureau of Economic Research.
Trigari, A. (2009). 'Equilibrium unemployment, job flows, and inflation dynamics', *Journal of Money, Credit and Banking*, vol. 41(1), pp. 1–33.
Wolf, M. and Fornaro, L. (2021). 'Monetary policy in the age of automation', Discussion paper, Centre for Economic Policy Research.
Yellen, J.L. (2015). 'Inflation dynamics and monetary policy', Speech delivered at the Philip Gamble Memorial Lecture, University of Massachusetts, Amherst.
Yoon, J.W., Kim, J. and Lee, J. (2018). 'Impact of demographic changes on inflation and the macroeconomy', *KDI Journal of Economic Policy*, vol. 40(1), pp. 1–30.

The Economic Journal, 1–34 https://doi.org/10.1093/ej/ueae006 © The Author(s) 2024. Published by Oxford University Press on behalf of Royal Economic Society. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com.
Advance Access Publication Date: 27 January 2024