What do wages in online job postings tell us about wage growth?

Pawel Adrjan*

Reamonn Lydon†

January 11, 2024

Abstract

We use data from millions of online job postings to construct monthly estimates of annual growth in advertised wages in the US, UK, France, Germany, Ireland, Italy, the Netherlands, and Spain from 2019 to 2022. The resulting wage growth tracker is a source of timely and forward-looking data on the wages of the marginal worker and serves as a useful leading indicator of momentum in wage dynamics. We show that the online job postings data benchmarks well against official sources on job vacancies, new hires, and wage levels. In both Europe and the US, growth in advertised wages accelerated sharply after the pandemic. Granular data shows a heterogeneous pattern of post-pandemic wage growth, with lower-paid jobs experiencing stronger wage growth in most of the countries in our sample. We attribute this to stronger labor demand for lower paid jobs coming out of the pandemic.

JEL No.: J31, J30, J63, J23, J20, E24, E31

Keywords: wages, labor cost, vacancies, online job postings, labor demand, labor supply, employment, inflation.

*Indeed Hiring Lab and Regent's Park College, University of Oxford; padrjan@indeed.com

†Central Bank of Ireland, Monetary Policy Division; reamonn.lydon@centralbank.ie (corresponding author). We thank the editors, the anonymous referees, Yusuke Aoki, Nick Bunker, Sorcha Foster, Chris Glynn, Robert Goodhead, Svenja Gudell, Annina Hering, Enda Keenan, Gerrit Koester, Laura Moretti, Gerard O'Reilly, Gillian Phelan, Tara Sinclair, Zivile Zekaite, and participants in ECB seminars, the Irish Economic Association Annual Conference, the Royal Economic Society Annual Conference, ECONDAT Spring Meetings, ESCoE Annual Conference on Economic Measurement, The Federal Reserve Board/ECONDAT/CEBRA 5th conference on Nontraditional Data, Machine Learning, and Natural Language Processing in Macroeconomics, and the Annual Conference of the French Economic Association for helpful comments. The views in this paper are those of the authors only. Earlier work on the wage tracker appeared as policy notes in the Central Bank of Ireland Economic Letter Series and on the Indeed Hiring Lab blog.

1. INTRODUCTION

A combination of high inflation and a tight labor market emerging from the COVID-19 pandemic led to concerns about the emergence of a potential wage-price spiral for the first time in decades. Central bankers in particular have been paying close attention to wage developments, looking for signs of 'second round' effects generated by the staggered catchup of nominal wages to inflation (see, for example, Lane (2022) and Powell (2022)).

Despite the importance of wages for inflation dynamics, timely and reliable data on wage growth is hard to come by, especially in Europe. Measures such as compensation of employees from National Accounts tend to be released at long lags and can be more backward- than forward-looking. Changes in the composition of workers and hours worked also complicate the interpretation of aggregate estimates of wage growth (Daly and Hobijn (2017), Christodoulopoulou and Kouvavas (2022)).

To address this data gap, this paper presents a new wage growth indicator based on advertised wages extracted from millions of online job postings in eight countries, over five years from January 2018 to December 2022. The job posting data comes from the Indeed job site. The countries are the US, UK, and, in the euro area, France, Germany, Ireland, Italy, the Netherlands, and Spain. The six euro area countries account for over 80% of employment in the bloc.

We extract a wide range of information from each job posting to construct a measure of annual wage growth, which we call the Indeed Wage Tracker (IWT). Our approach is similar to the Atlanta Fed's Wage Growth Tracker, but instead of tracking the wages of individuals, we track wage growth within narrowly defined jobs. We first calculate the median advertised wage for each country, month, job title, region, and wage or salary type (hourly, monthly, or annual). Within each country, we then calculate year-on-year wage growth for each job title-region-salary type combination, generating a monthly distribution. Our headline estimate of annual wage growth at the country level is the median of that distribution.

Our tracker provides a timely signal of the momentum in wage growth dynamics, as well as serves as a leading indicator of the state of the labor market more generally. In this paper, we present the output from the tracker to December 2022, although the output continues to be updated online monthly.[1] A wide variety of users have drawn on the monthly data to assess wage developments, including central banks, multilateral institutions, and investors.[2]

[1]See the GitHub site for the latest Indeed Wage Tracker data [link].

[2]Organizations that have cited the tracker in policy publications include the Bank of England (Bank of England (2023)), the European Central Bank (Lane (2022)), the Federal Reserve (Powell (2023)), and

This paper presents the tracker and focuses on the following questions. First, how representative are the job ads we use to construct the wage tracker? The answer to this question is important for the validation of our wage growth estimates and gives us confidence that they are informative about labor market developments. Second, what were the trends in advertised wages before, during, and after the pandemic? Analyzing these trends helps us better understand the evolution of the labor market in recent years when many survey-based measures of aggregate wage growth were distorted by compositional shifts and government support schemes. Third, how do trends in advertised wages vary across occupational groups, and what explains any differences we find?

We have three main results. First, we find that online job postings from Indeed benchmark well against official statistics on vacancies and new hires. Specifically, the distribution of online job postings across occupations and sectors and the distribution of wage levels in job postings are comparable to official data from sources such as national Labor Force Surveys, job vacancy statistics, labor cost surveys of wages across sectors, and survey data on the wages of new hires. While there are some differences in the occupation and sector mix of online postings and other labor market data, we show, through re-weighting, that these differences do not substantially alter the results of our wage tracker.

Second, we show that wage growth accelerated as economies re-opened after the pandemic, growing at rates far above pre-pandemic levels. In some countries, such as the US, the IWT is a leading indicator of wage growth from other sources. This is particularly valuable around turning points in economic activity and wage growth momentum. In other countries – notably continental Europe – the IWT is more concurrent relative to other wage growth indicators, such as those based on negotiated wages from collective bargaining agreements (see Koester et al. (2020)). Nonetheless, the IWT retains some important advantages over these indicators, including being more timely – it is available within a matter of days after month-end compared to a two- to three-month lag for negotiated wages – and it is more forward-looking – wage negotiations generally take place over several months and can be backward-looking in nature as they often include clauses for past inflation surprises.

Third, we show that wage growth after the pandemic was generally stronger for lower-paid jobs. This gap between wage growth for low- and high-paid jobs is largest in the US and UK, but is also evident in euro area countries. Exploiting the granular information in the online job posting data, we show that wage growth in recent years was stronger for jobs that saw a larger increase in job postings.

We are not the first to propose a new way to measure wage growth that complements official statistics. In the US, for instance, the Atlanta Fed Wage Growth Tracker measures

the OECD (Araki et al. (2023)), among others.

average wage growth across a sample of individuals surveyed 12 months apart, drawing on the methodology in Daly and Hobijn (2017). The ECB's Negotiated Wage Tracker provides another source of timely data by measuring trends in wage settlements in a subset of euro area countries (Koester et al. (2020)). In some countries, private-sector researchers have also begun to publish complementary measures of wage growth, such as CaixaBank's indicator based on salaries deposited by employers into Spanish workers' bank accounts (Mestres Domènech (2022)).

However, all these indicators have limitations that make our tracker a valuable contribution. The Atlanta Fed's tracker, while widely followed, is based on relatively small sample sizes from the monthly Current Population Survey. This means that robust estimates are limited to a small number of groups so that analysis of the heterogeneity of wage growth – across geographies, wage quantiles or occupations – is limited. The ECB's Negotiated Wage Tracker is based on union bargaining agreements, which often adjust wages for past inflation, and, as such, the indicator may be overly backward-looking (Lane (2022)). Furthermore, agreed changes in wages only feed through slowly to average wage growth as old agreements run out and new agreements are concluded. Existing private-sector measures of wage growth are timely and use large samples but are usually limited to single countries.

In contrast, the data from job postings on Indeed has several unique advantages. They include timeliness of delivery, large sample sizes that enable detailed breakdowns by occupation and other variables, and coverage of multiple countries with consistent data from a single source, which facilitates cross-country comparisons. Another useful feature of wages in job postings is that they can be viewed as reflecting employers' medium-term expectations about the value of the worker-firm match (Adrjan et al. (2023a)). This gives our tracker a forward-looking aspect.

Another advantage of our data is that we only use *actual* posted wages from job descriptions, not imputed values. According to Lafontaine et al. (2023), in other sources of wages from online job postings, such as Lightcast (formerly Burning Glass), as much as three-quarters of job ads 'with wages' are an algorithmic imputation by third-party data providers.

We make two main contributions to the labor economics literature and to policy. First, we collect wage data from online job postings for eight countries over five years from 2018 to 2022, with the dataset growing as more data is added each month. In contrast, most studies that use online job posting data to analyze wage dynamics focus on single countries. Examples are Djumalieva et al. (2020) for the UK and Hazell et al. (2022) and Adrjan et al. (2023b) for the US. One exception is Pham et al. (2023), which uses online job postings to analyze labor demand in Ukraine and Poland during wartime. However, to our knowledge, ours is the first paper to collect comparable and consistent data on wages posted in job ads in multiple countries to track wage growth over time. Although the timeframe we cover does not represent a full economic cycle, it does capture periods of extreme swings in labor demand – the shuttering and reopening of the economy during and after the pandemic – as well as labor supply shocks in some countries and high rates of inflation. All these factors should impact nominal wages, which is exactly what we find.

This points to our second contribution: analysis of the heterogeneity of wage growth. The dataset is very rich, allowing us to highlight both cross-country wage growth dynamics and differences across occupations. As noted earlier, one of our main findings is that wage growth was generally higher for lower-paid jobs after the pandemic.

Our paper also fits into a growing literature that uses online job postings to study a wide range of labor market topics, including skills demand (e.g., Modestino et al. (2020)), digitalization (Bellatin et al. (2023), Soh et al. (2022)), remote work (Adrjan et al. (2023a), Hansen et al. (2023)), wage heterogeneity (Marinescu and Wolthoff (2020)), labor market concentration (Azar et al. (2020)), and matching (Turrell et al. (2021), Bhuller et al. (2023)). Some of this literature uses data from Indeed. We also add to an emerging body of policy-relevant research on the impact of the pandemic on the labor market. Pandemic-related labor market upheaval boosted demand from policymakers for timely and granular information on the heterogeneous nature of the shock. Large datasets of online job postings are well suited to meet this demand, as many papers have shown (see, for example, Chen (2020), Anderton et al. (2020), Adrjan and Lydon (2020), Chetty et al. (2020), Forsythe et al. (2020) and Hensvik et al. (2021)).

The remainder of the paper is structured as follows. In Section 2, we present the data, including the results from our validation of the Indeed job posting dataset. Section 3 presents the wage tracker methodology and results, including comparisons with other wage growth indicators. Section 4 goes below the national level and looks at the heterogeneity of wage growth across occupations. Section 5 concludes.

2. DATA AND VALIDATION

2.1. *Wage data in job postings*

A key innovation of our tracker is using data on wages from 70 million job postings on Indeed, the largest job site in the world by number of visitors. Indeed lists jobs that are posted directly to the site by employers, as well as those aggregated from other online sources. Postings are de-duplicated so that when the same job is collected from multiple sources it is shown only once.[3]

To develop the tracker, we use data from eight countries: the US, UK, Germany, France, Italy, Spain, the Netherlands, and Ireland between January 2018 and December 2022. These countries were selected based on Indeed having an established market presence for many years, while the period was determined by data availability. Although we would have preferred to use a longer time series for validation purposes, data before 2018 is not available. Nevertheless, we have access to 26 months of data before March 2020, which lets us establish a pre-pandemic baseline for the subsequent analysis of trends during and after the pandemic in 2020, 2021, and 2022.

The data we use to construct our wage growth tracker is described in Table 1. From the text of each job posting that includes a wage or salary, we extract the following information:

- **Wage or salary**, and whether it is hourly, monthly, or annual. We exclude any additional benefits, focusing on basic pay. For cell size reasons, we drop the small proportion of postings with daily or weekly wages (less than 0.5% of jobs, in just a handful of countries). Advertised wages can be posted as a single value (e.g., "€50,000"), a range ("€50,000-€60,000"), a floor ("from €20 per hour"), or a ceiling ("up to €3,000 per month"). We use the midpoint of any ranges and drop the floors and ceilings because the expected wage is not clear-cut in those situations. All of our wage growth estimates are robust to different treatments, for example, using the bottom or the top of a range, or including wages expressed as floors or ceilings.

- **Job title**, the most granular level of information we use to describe the job, which is more detailed than an occupation. The job title is normalized to be consistent across all job postings within each country and over time. We discuss job titles further below.

- **Occupational category** from the Indeed taxonomy of 54 occupational groups. This is similar to two-digit occupational codes used by statistical agencies, with some wider groupings in certain occupations.

- **Location of the job** within the country, such as the state, county, or administrative region. Appendix Table A1 lists the geographic variables that we use to construct the wage tracker for each country.

[3]More information on Indeed job posting data is available online [link]. The Indeed Job Postings Index, aggregated by country and occupational category, is available in the FRED database [link].

- **Date** when the posting first appeared on Indeed. In the tracker, each posting is included only once in the month it first appeared on the site, meaning that we use the flow of postings rather than the stock. That is because wage growth based on the stock of postings could be biased by differential changes in posting duration by wage level. In addition, the pandemic experience shows that the flow of postings reacts more rapidly to labor market conditions.[4]

Table 1. Data used to construct the wage tracker

	Number of postings with wage data (millions)	Share of postings with wage data (%)	Hourly (%)	Monthly (%)	Yearly (%)	Number of job titles
France	6.4	30%	34%	34%	32%	5,725
Germany	1.9	11%	58%	28%	14%	9,165
Ireland	0.4	32%	38%	0%	61%	4,691
Italy	0.3	9%	4%	68%	28%	3,219
Netherlands	2.3	35%	21%	77%	3%	6,977
Spain	0.3	12%	16%	30%	54%	5,420
EA6 Total	11.6	21%	34%	41%	24%	
UK	16.0	49%	45%	0%	54%	7,647
US	42.6	25%	72%	1%	27%	10,161
Total	70.2	27%	59%	8%	33%	

Notes: Data is from Indeed job postings from January 2018 to December 2022. Hourly, monthly, and yearly column headings refer to how the wage or salary is expressed in job postings, shown as a share of all postings with wage or salary information. The 'EA6 total' reflects the six euro area countries listed in the table.

2.2. *Benchmarking online job posting data*

An important question for online job posting data is how representative it is of the labor market in each country. The first step in analyzing representativeness is defining the relevant population of interest. Several papers have benchmarked the sectoral and occupational distribution of online job postings against total employment, including Cammeraat and Squicciarini (2021), Hazell et al. (2022), and Bellatin et al. (2023).[5] The pattern that

[4]See Figure 3 in Kennedy (2020), which plots the flow of new job postings on Indeed in France, Germany, Italy, the UK, and the US in the days following the start of the Covid-19 pandemic.

[5]Cammeraat and Squicciarini (2021) compared postings from Burning Glass to employment by occupation in Australia, Canada, New Zealand, Singapore, UK, and US between 2010 and 2019. They found that higher-skilled occupational categories, such as managers and professionals, were better represented than other occupational categories. Similarly, Hazell et al. (2022) benchmarked Burning Glass postings with wages from 2010 to 2019 against US employment by occupation and sector. They found that the data over-represented occupations in computing, transportation, and management and under-represented food preparation and construction. Bellatin et al. (2023) compared the distribution of Canadian job postings on Indeed to the distribution of employment across provinces and ten broad occupational groups. They

emerges from these analyses is that online job postings sometimes under- or over-represent employment in certain sectors or occupations, but the distributions generally match fairly well. This is consistent with what we find.

While detailed data on employment is readily available in most countries, it is not the ideal benchmark for online job postings, because worker turnover and job vacancy rates vary across sectors and occupations (Hershbein and Kahn (2018)). As a result, the sectoral and occupational composition of the demand for new workers – as expressed in job postings or vacancies – can diverge from the composition of employment.

In this paper, we are interested in tracking posted wages, which relate specifically to the market for new workers. We, therefore, focus our benchmarking and validation exercise on comparing job postings to vacancies and new hires, rather than employment. This is similar to the approaches of Hershbein and Kahn (2018), who compare US postings from Burning Glass with job openings by sector and with new jobs by occupation, and Soh et al. (2022), who compare US postings from Indeed with job openings by sector and state.

Crucially for our estimates of wage growth, we also compare wages advertised in job postings with the wages of new hires and the wages of all workers from representative surveys. We show that the Indeed data benchmarks well on all the dimensions we consider. Finally, we also show that reweighting our wage tracker using the distribution of all postings or new hires by occupation does not substantially alter the aggregate trends. The results of these validation exercises make us confident that our wage tracker captures broad trends in the labor market.

2.3. *Online job postings versus statistics on job vacancies*

We start with aggregate trends. Figure 1 shows that online job postings correlate well with trends in job vacancies from official sources between 2018 and 2022.[6] Not only is the dip during the initial stages of the pandemic in the first half of 2020 and during subsequent lockdowns evident in the online data but so is the rapid and strong rebound after economies reopened. This is consistent with the patterns shown by Bellatin et al. (2023) for Indeed postings in Canada and Chen (2020) for Indeed postings in 22 countries. These patterns give us confidence that our online job posting dataset captures overall

concluded that the distributions looked fairly similar, despite some over-representation of higher-skilled occupations, such as managers.

[6] In this paper, "job postings", "job openings" and "job vacancies" are broadly synonymous. We use the terminology that reflects the original data sources. For example, both Eurostat and the UK ONS publish data on "job vacancies", whereas the JOLTS data from BLS in the US refers to "job openings". When we refer to Indeed data, we use "online job postings".

trends in job openings. While one job posting does not necessarily equal one job opening, both series approximate the demand for new workers and trend closely together.

Next, we compare the composition of Indeed online job postings and published job vacancy statistics over the same period. One challenge is that the main Indeed taxonomies are job titles and occupational categories, but vacancies are only available by sector. We follow Soh et al. (2022), who also benchmark the sectoral distribution of Indeed job postings against JOLTS data for the United States, by mapping occupations to one-digit sectors. For most occupations, such as those in construction, transport or retail, there is an obvious mapping from the occupation or job title to a sector. For others – such as professional or managerial roles – the allocation is less clear. Like Soh et al. (2022), where there is no obvious direct mapping, we assign occupations to sectors based on occupation-sector shares in labor force surveys.

Figure 2 shows the results. Whilst not a perfect match, the sectoral composition of Indeed job postings and job vacancies from national statistics is very similar. In all countries, the Indeed data is over- or under-represented in certain sectors, and there are some patterns, but they do not apply to all countries equally. Construction is the only sector that appears to be over-represented in online job postings in all eight countries. Accommodation and food is consistently under-represented in all countries, perhaps due to a high share of informal hiring methods or because it is a sector where each online posting relates to more than one vacancy to a greater extent than in other sectors.[7] Transport and storage, wholesale/retail, administrative, and professional services tend to be over-represented in online job postings in most – but not all – countries, although to a different extent in each one.[8] In contrast, manufacturing, finance, and healthcare tend to be under-represented.

Overall, these benchmarking exercises suggests that our online dataset captures a slightly different sectoral mix of jobs than survey-based vacancy estimates. However, some of those differences could be driven by our imperfect ability to map postings to sectors. Our mapping is derived from occupations. So, for instance, if a high share of vacancies in the healthcare sector is for administrative roles, then we will underestimate the number of postings in the healthcare sector and overestimate the number of postings in the administrative services sector. The impact is difficult to test since vacancies in official surveys are not broken down by occupation. Hence, in the next section, we improve on this by mapping postings to standard occupational categories and comparing the resulting

[7]We are not able to test this hypothesis with our data.

[8]The over-representation of professional services and IT in the EU countries in our sample is consistent with Napierała et al. (2022), who find that managerial, professional, and technical roles tend to be over-represented in postings compared with Labor Force Surveys, as well as with the trends observed by Cammeraat and Squicciarini (2021). For Germany, this is also consistent with Carrillo-Tudela et al. (2023), who find that German firms that use the job posting channel for hiring tend to recruit for higher-skilled positions than firms that use other channels like networks or the public employment service.

distribution to that of new hires.

In the meantime, we note, however, that despite some apparent differences in the sectoral mixes of postings and vacancies in Figure 2, the quarterly sector shares from the two data sources are strongly positively correlated. Panel A in Table 2 shows that the correlation between the quarterly sector shares in the two datasets ranges from 0.65 in Spain to 0.85 in the UK when we consider all postings and from 0.62 in Italy to 0.90 in the US when we consider only postings with wages.[9] This suggests postings and vacancies evolve similarly over time.

Table 2. Correlation of the sector and occupation shares of Indeed job postings with official statistics on job vacancies and new hires, 2018-22

	Germany	France	Spain	Italy	Netherlands	Ireland	US	UK
A. Sector shares: job postings vs. vacancies								
All job postings	0.74	0.67	0.65	0.68	0.80	0.78	0.84	0.85
Postings with wages	0.80	0.63	0.75	0.62	0.72	0.64	0.90	0.85
N	320	285	260	279	320	180	300	260
B. Occupation shares: job postings vs. new hires								
All job postings	0.72	0.85	0.53	0.65	0.72	0.83	n/a	0.86
Postings with wages	0.66	0.83	0.67	0.68	0.75	0.90	n/a	0.86
N	414	567	575	561	573	536	n/a	640

Source: Quarterly job vacancy and new hire data is from Eurostat for the euro area countries, ONS for the UK, and BLS-JOLTS for the US. *Notes:* US data represents non-farm job openings. Job vacancies for France are for firms with ten employees or more. Job vacancy shares for Italy are derived from the job vacancy rate because estimates of vacancy numbers are not available. Sample sizes vary depending on the number of sectors and occupations for which job vacancy and new hire statistics are available in each country. Figures 2 and 3 compare the shares graphically, for the total population of Indeed job postings. Figures A1 and A4 in the appendix show the shares for the subset of job postings that include wage data.

[9] Appendix Figure A1 replicates the bar charts in Figure 2 using only postings with a wage.

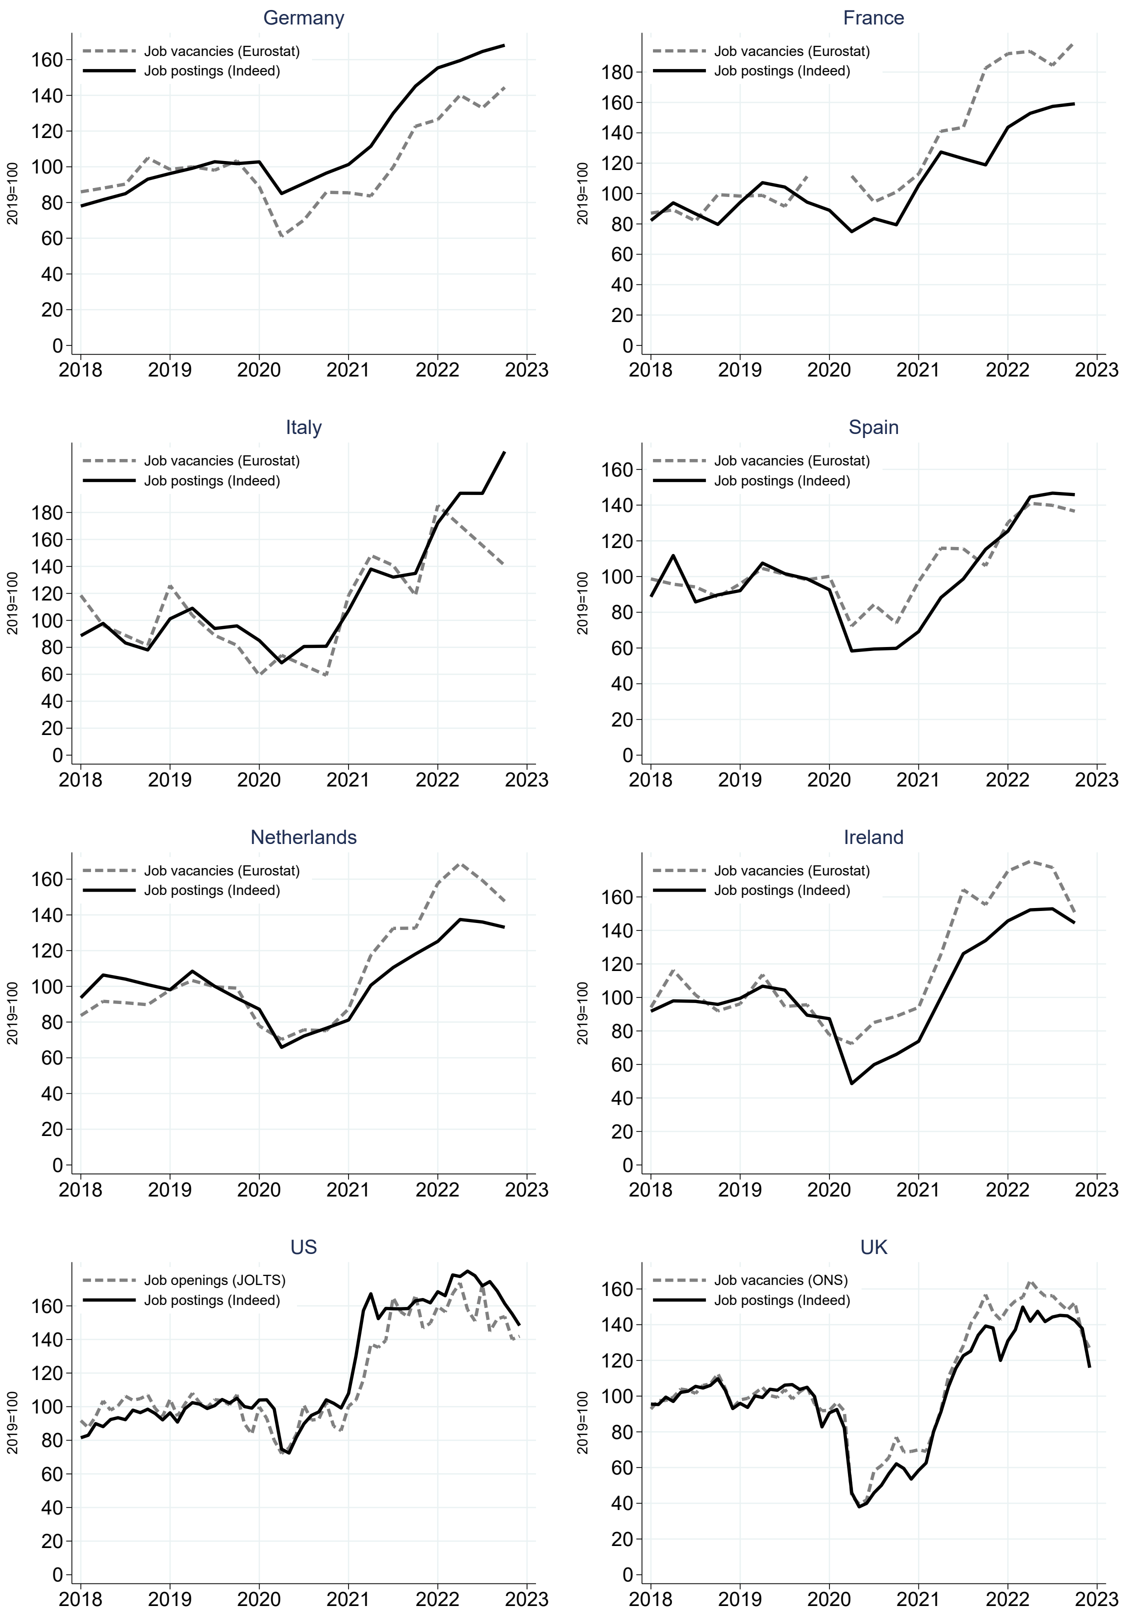

Fig. 1. Trends in Indeed online job postings and job vacancies from statistical agencies, 2018-22 (2019=100). *Source:* Quarterly Eurostat for euro area countries, monthly ONS for UK, and monthly JOLTS for US. *Notes:* Data for France is for firms with ten employees or more, no French data for Q1 2020. Data for Italy is derived from the job vacancy rate because estimates of vacancy numbers are not available. US data is non-farm job openings. Series are not seasonally adjusted.

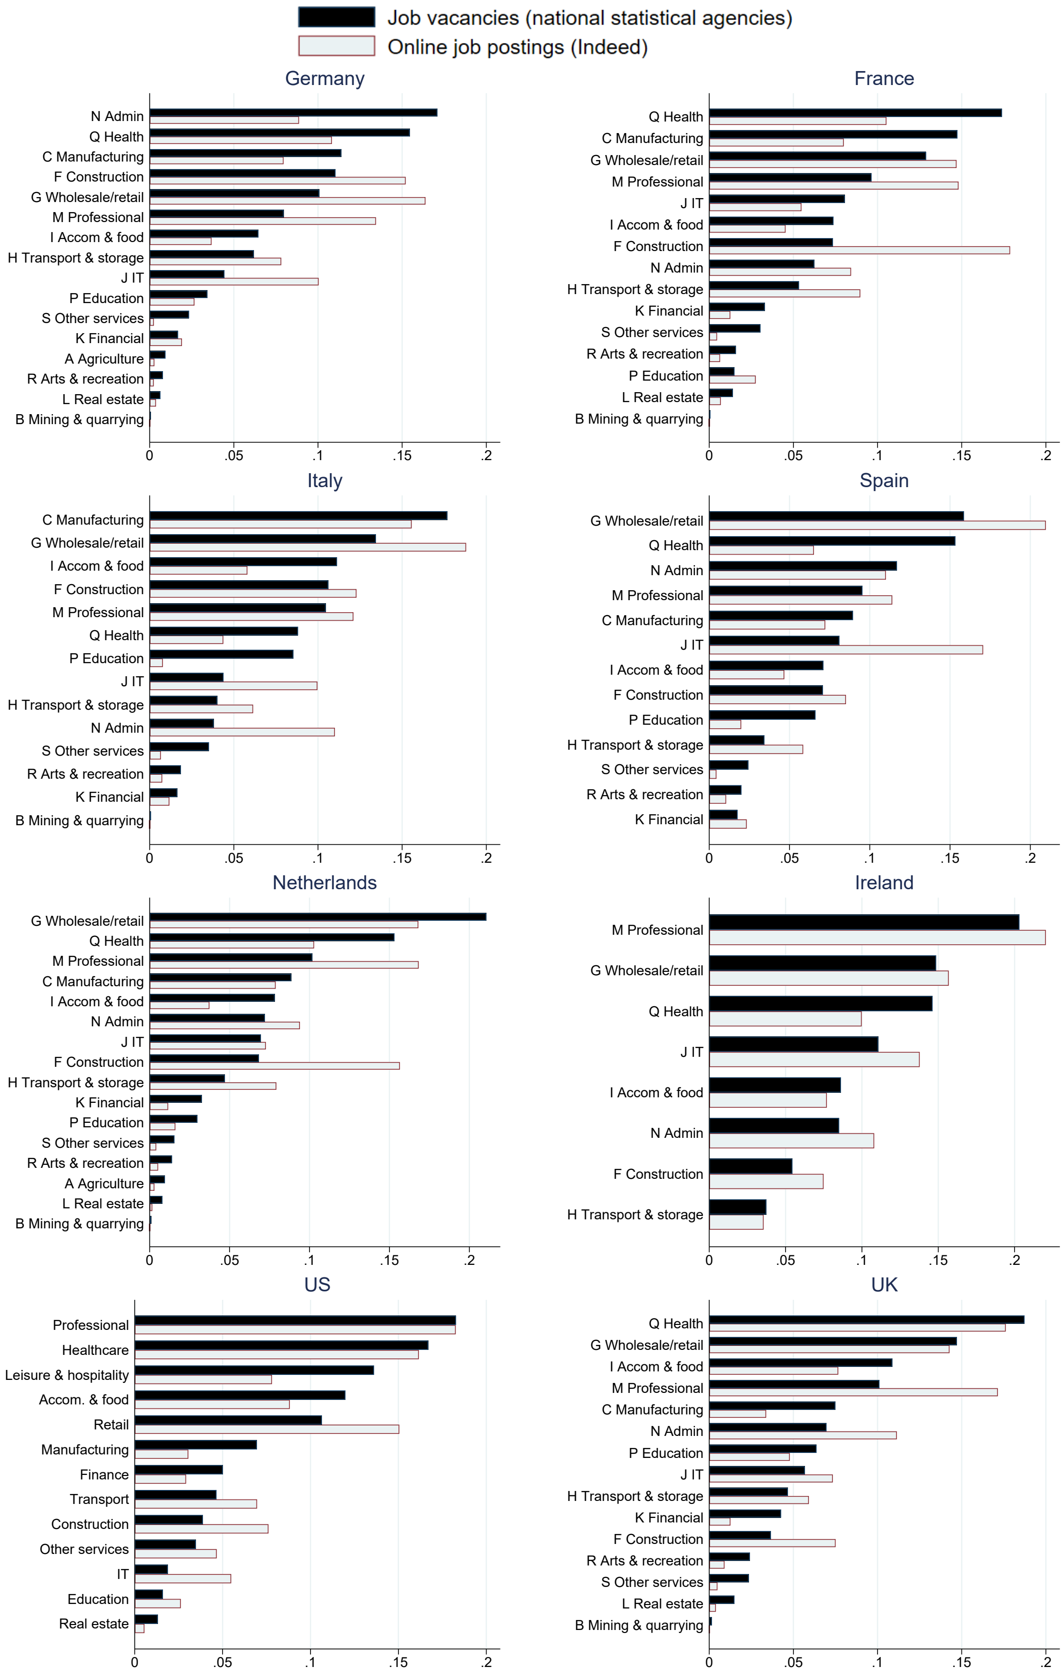

Fig. 2. Sector shares of Indeed online job postings and job vacancies from national statistical agencies (2018-22) *Source:* Eurostat for euro area countries, ONS for UK, and JOLTS for US. *Notes:* Data for France is for firms with ten employees or more. In the absence of job vacancy numbers for Italy, the chart data is estimated from sector vacancy rates and employment from Eurostat. US data is non-farm job openings.

2.4. *Job postings versus new hires*

Next, we compare online job postings with data on new hires. We focus on occupations, rather than sectors, for two main reasons. First, the Indeed occupational categories map closely to occupation codes used by statistical agencies (i.e., ISCO and SOC), providing a potentially more accurate match to published statistics than sectors.[10] Second, while not every vacancy will lead to a hire (Azar et al. (2020)), comparisons with new hires provide a benchmark for weighting our wage tracker to a group of *actual* workers. We return to this below when we discuss weights in Section 3.2.

Data on the occupational shares of new hires comes from quarterly Labor Force Surveys from 2018 to 2021 for the euro area countries and to 2022 for the UK. The US is excluded from the occupational benchmarking because JOLTS data on hires is only available by sector (in Figure A2 in the Appendix, we show the comparisons by sector, including for the US). We define a new hire as an employee with job tenure of three months or less. Figure 3 shows the average occupational shares in the two datasets for occupations with a job posting share of 1% or higher for legibility, while Figure A4 does this for postings with wages. Panel B of Table 2 shows the correlations of the quarterly shares using all postings and all occupations for each country.

Overall, the occupational mix of the online job posting data is very similar to that of new hires. For most countries, such as the UK, Ireland, France, Germany, and the Netherlands, the comparison is very good, with correlation coefficients ranging from 0.72 to 0.86. For Italy and Spain, the correlation is weaker but still positive, at 0.65 and 0.53, respectively. In both countries, a relatively high concentration of online job postings in a small number of occupations – such as *'Retail & Sales'* and *'Software Development'* – drives this result.

Focusing only on postings with wages, we find that the correlation coefficients range from 0.66 in Germany (fairly high despite the low share of postings with wage information) to 0.90 in Ireland. While these correlations are high, they are not perfect. For this reason, as noted above, when we present the wage growth tracker at the country level, we show the results with different weightings, including a re-weighting of the wage tracker data using the occupational shares of new hires.

[10]For comparison with Section 2.3, which shows job vacancies by sector, Appendix Figure A2 and Table A2 show data on new hires by sector. Figures A3 and A4 compare the sector and occupation shares in the subset of job postings *with wages* with the share shares of new hires.

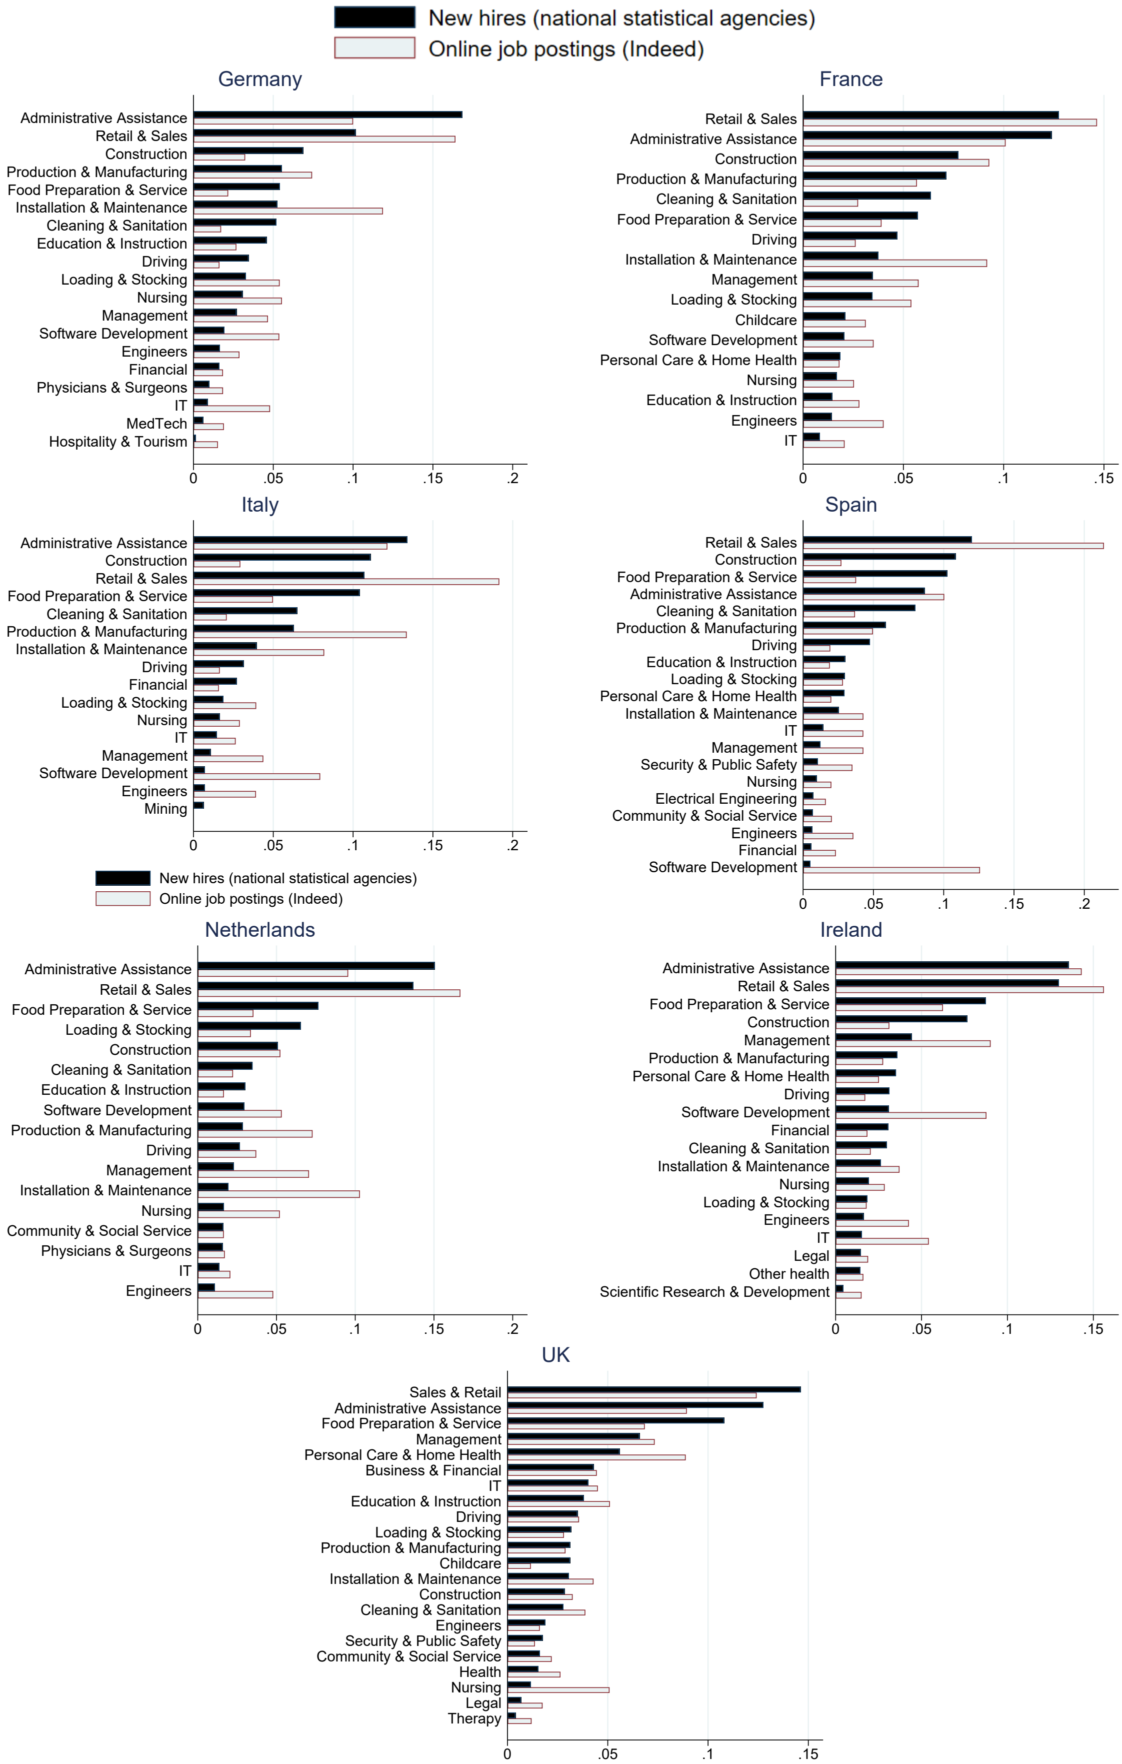

Fig. 3. Occupational shares of Indeed job postings and new hires from national statistical agencies. *Source:* Eurostat for euro area countries (2018-21), ONS for UK (2018-22). *Notes:* Chart shows occupations with share postings share of at least 1%.

The re-weighting potentially matters because, as Table 1 shows, only 27% of postings in our eight-country dataset include wage or salary information. One concern is that our sample becomes less representative when we only use the subset of postings with a wage to construct our tracker. This could happen if some categories of jobs were systematically more or less likely to include wages in the posting. To test whether this is a problem, Appendix Figure A5 plots the occupational shares of the population of job postings and the sub-sample of postings with a wage from 2018 to 2022. The occupational shares are strongly positively correlated – with a correlation coefficient that ranges from 0.86 in the US to 0.98 in the UK – and there is little evidence of systematic over- or under-representation on either side of the 45-degree line.

These correlations tell us that the occupational shares are broadly similar in both datasets. However, when there are differences, we might still be concerned that this is correlated with our variable of interest, wages. To test this, we construct a 'gap' measure, which is the difference between the occupation shares in the sub-sample of job postings that include wages and the shares in the full population of Indeed job ads. We then plot this gap against average log wages by occupation and year in Figure A6. In all countries, the vast majority of observations are on the horizontal axis, which suggests no correlation between average wage levels and the propensity to include wages in the job posting. In some countries – notably Germany, Spain, the Netherlands, and Ireland – there is a pattern of some lower-paid jobs or occupations being more likely to include wages in the posting and some higher-paid jobs being less likely to do so. We can control for any potential bias by re-weighting our wage growth tracker to be representative of occupations in the population of job postings. However, as we discuss below, this has only a small impact on country-level estimates of wage growth, given the small differences in the shares and the fact that this gap arises for a few occupations.

2.5. *Wages in online job postings versus other wage data*

Our final data validation exercise compares the level and composition of posted wages with other available wage data. For the euro area countries and the UK, we compare wages in job postings with the wages of new hires from the 2019 EU-SILC survey. We use 2019 data because it is available for all European countries, and it represents a year of relative stability before the pandemic. In EU-SILC, a new hire is any worker who answers yes to the question: "Did you start a new job in the last 12 months?". For the US, we use the wages of new hires from the Current Population Survey (CPS 2019). Following the methodology of the Atlanta Fed Wage Growth Tracker, new hires are respondents who say they have moved to a new employer between survey waves.

First, we compare the percentiles of the monthly wage distribution. For Indeed postings with hourly wages, we assume that weekly hours are in line with the country averages in national surveys and 4.3 weeks worked per month. Annual salaries from postings are divided by 12. Figure 4 presents the results. It suggests that our wage data from job postings captures very well the percentiles of the new hire wage distribution.

If there is a part of the Indeed wage distribution that is comparably lower than in the survey data, it is the highest-paid workers, but only in certain countries. In Ireland, France, the Netherlands, and the UK, the top decile of wages from the population is above the top decile in the Indeed data. There are two possible reasons for this. First, higher-paid jobs may simply be less likely to include advertised wages in the job ad. We saw some evidence of this in the previous section. Second, advertised pay in some higher-paying jobs may be more negotiable. While we are unable to test this hypothesis in our data, the advertised wage may be akin to a floor, with the effective starting wage for new hires in these jobs tending to be higher on average. There are also some deviations in the lowest decile, likely driven by the fact that the conversion of advertised wages to monthly figures assumes full-time work.

Table 3 combines the graphical assessment in Figure 4 into a single statistic. For each country, we calculate the absolute difference between the survey estimate of new hires' earnings at each percentile, and the estimate from the Indeed data, and take the average across percentiles. For example, we see from this metric that, relative to the population estimate from the survey, the Indeed data tends to be closest in Italy, France and Germany, where the statistic ranges from 0.035 to 0.055). For the likes of the US, UK and Ireland, where it ranges from 0.070 to 0.092, it is a little less aligned. The second column, which shows the same statistic excluding the bottom and top deciles, confirms that the Indeed data is least aligned in the tails.

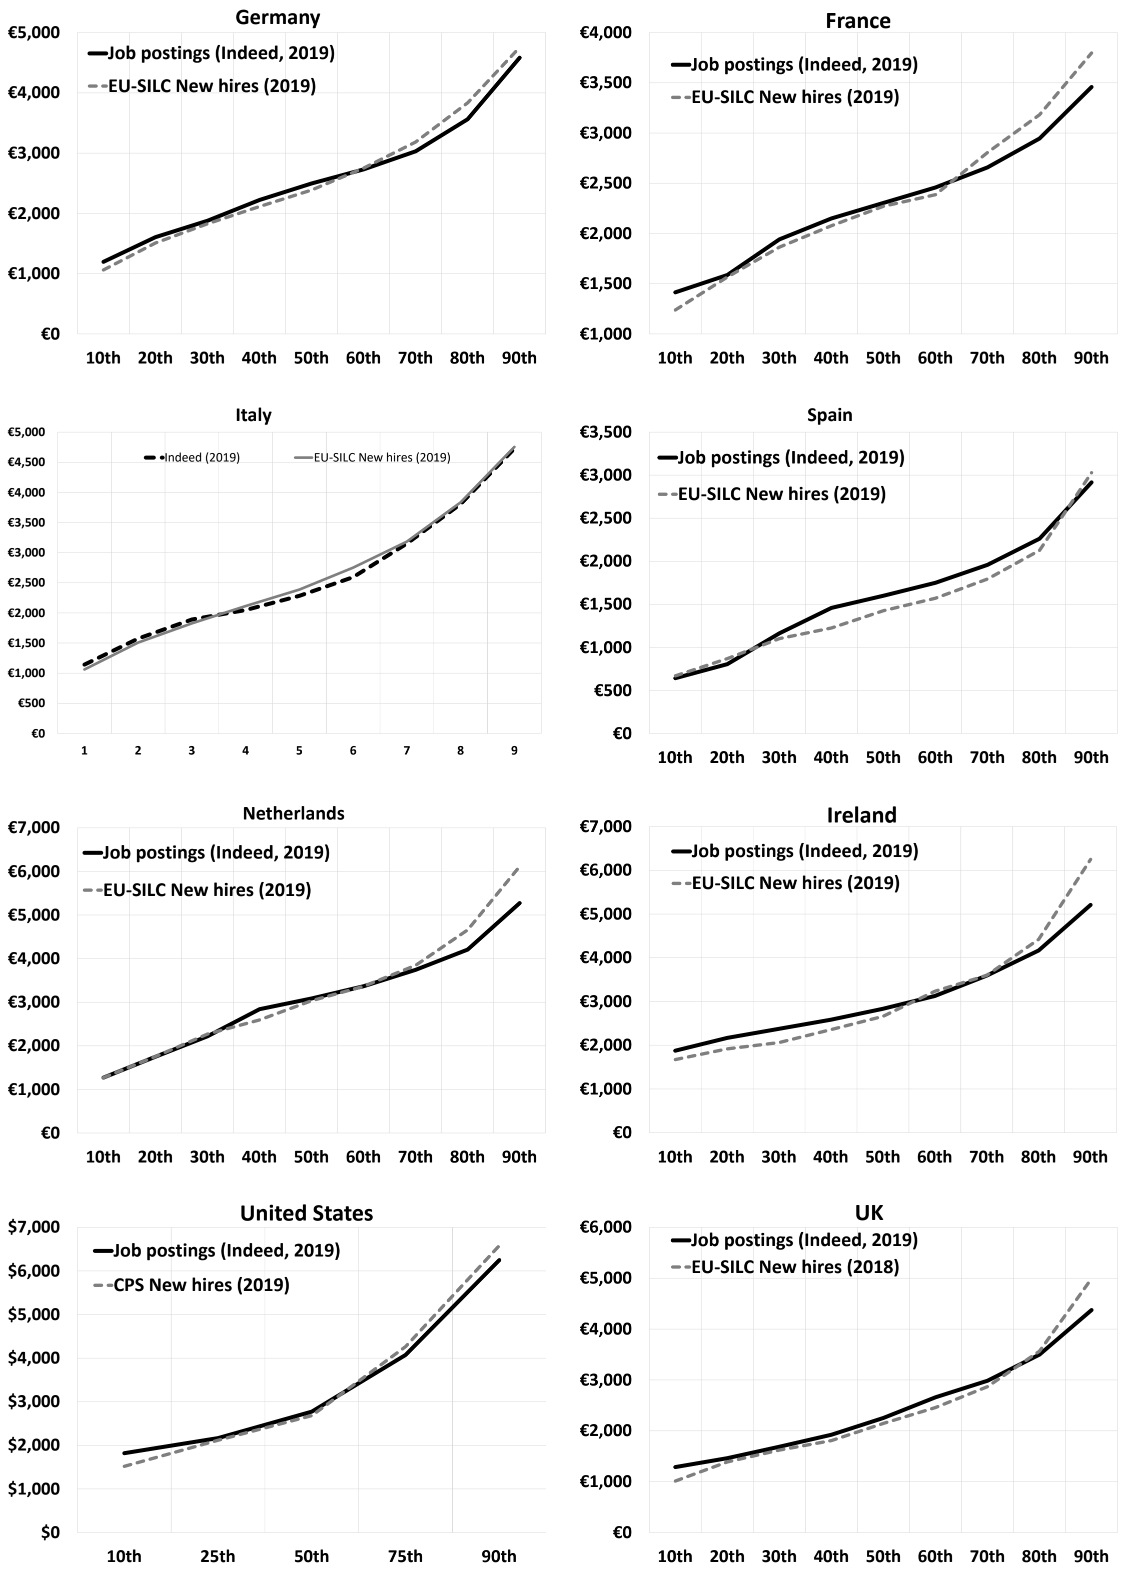

Fig. 4. Percentiles of monthly pay in Indeed job postings and statistical surveys. *Source:* Indeed data and Eurostat for EU-SILC, NBER for Basic Monthly CPS datasets for the US. *Notes:* EU-SILC microdata for the UK is only available up to 2018. In the Indeed data, hourly wages from job postings are converted to monthly equivalents assuming that weekly hours are in line with the country averages in national surveys and there are 4.3 weeks worked per month. Annual salaries from job postings are divided by 12. Survey data is population-weighted. Indeed data is unweighted.

Table 3. Average absolute deviation of monthly pay in the survey data and Indeed job postings (as a percentage of the survey value), averaged across percentiles

	All percentiles	Excluding bottom and top percentiles
Germany	0.053	0.045
France	0.055	0.037
Italy	0.035	0.033
Spain	0.087	0.101
Netherlands	0.044	0.037
Ireland	0.092	0.077
UK	0.083	0.050
US	0.070	0.034

Notes: This table shows the absolute differences between monthly pay in the survey data (e.g. SILC or CPS) at each earnings percentile and wage in the Indeed job postings at the same earnings percentile, as a percentage of the survey data and averaged across all percentiles. For example, the UK value of 0.083 indicates an average 8.3% difference in earnings across all percentiles. The second column is the same metric, excluding the bottom and top deciles.

Next, we compare average wages by sector from the online job posting data with average monthly wages by sector from national statistical agencies. For euro area countries we use the Annual Labor Cost Survey published by Eurostat from 2020-22; for the UK, data is from the Annual Survey of Hours and Earnings; and for the US, data is from the Current Population Survey. Figure 5 shows this comparison for each country.

As the charts show, average wages by sector extracted from online job postings are a good match for average wage levels by sector from official statistics. In most countries workers in the IT (sector J) or Professional (M) tend to have the highest average wages, a pattern replicated in the online job postings data. In both data sources, average wages tend to be lowest in the Accommodation and Food sector (I), Other Services (S), Transport and Storage (H), and Administration (N).

Our overall conclusion from the validation exercises in Sections 2.3, 2.4, and 2.5 is that the Indeed job posting data aligns well with job vacancies and new hires across sectors and occupations. Whilst we explore a re-weighting of the wage tracker to ensure closer alignment to the populations of job postings and new hires in Section 3.2, we conclude that wages extracted from online job postings are a fairly representative indicator of the wages of new hires for the purpose of tracking aggregate wage growth.

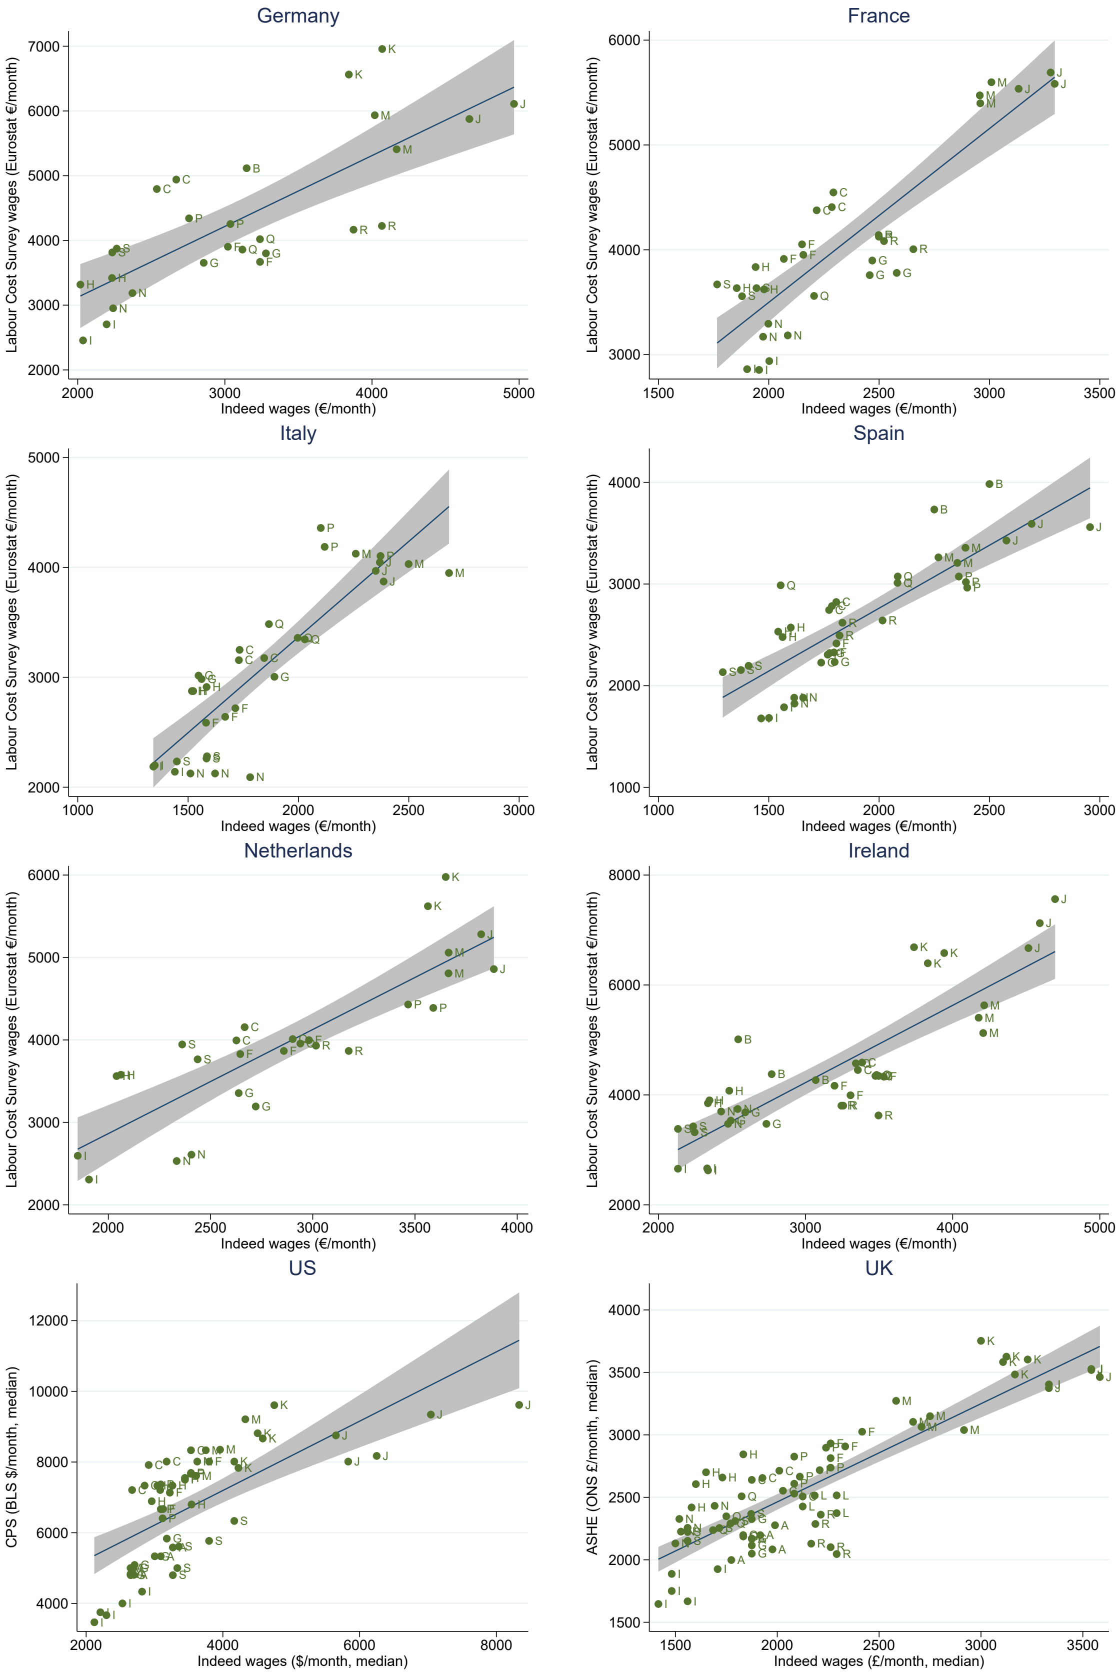

Fig. 5. Average wage by sector in Indeed job postings and survey data (employees 2018(20)-22, monthly wage). *Source:* Indeed data; Eurostat Annual Labor Cost Survey (2020-22); ONS Annual Survey of Hours and Earnings (ASHE, 2018-22); US wages by sector from the Current Population Survey (CPS, 2018-22). *Notes:* Survey data for Germany and Netherlands is available for 2021-22 only. Sectors are one-digit NACE codes B-S. UK and US survey data are medians.

3. WAGE TRACKER METHODOLOGY AND RESULTS

We now present our wage tracker methodology and estimates. Our methodology proceeds in three steps:

1. For each country and month, we group job postings into cells defined by job title, region (see Table A1 in the appendix), and salary type (hourly, monthly, or annual), and calculate the median wage in each cell.
2. We then calculate the year-on-year change in the median wage for each job title-region-salary type cell, generating a monthly distribution.
3. Our monthly estimate of wage growth at the country level is the median of this distribution. For our baseline country-level estimate, we use an unweighted median. As we show below, our overall results tend not to be sensitive to alternative weightings.

We use medians because it makes our estimates of wage growth less susceptible to outliers, especially for cells with fewer observations.[11] As Table 4 shows, whilst we have a large number of cells, on average the number of job postings in each cell tends to be small, in single digits for most countries. An alternative would be to set a minimum cell size or drop one of the categories. However, as this would likely lead to a bias towards capturing wage trends for large job titles, we opt not to do this. Furthermore, we prefer a larger number of cells to generate a distribution from which to take the median for the country-level estimate of wage growth.

Our approach is similar to the Atlanta Fed's Wage Growth Tracker, but instead of tracking the wages of individuals, we track wage growth within narrowly defined jobs. Calculating wage growth within job titles allows us to control for changes in the composition of jobs which could have a substantial impact on aggregate advertised wage growth. As Daly and Hobijn (2017) show, aggregate wage dynamics are heavily influenced by changes in the composition of the employed over the business cycle, and we expect wages in job ads to exhibit similar composition effects. The US Bureau of Labor Statistics' *Employment Cost Index* – a closely-tracked quarterly metric of wage growth – controls for compositional shifts in a similar way to us, tracking wages withing occupation-industry and geographic cells; see Ruser (2001).

[11] We have also estimated wage growth using alternative methodologies, such as the regression-based approaches in Marinescu and Wolthoff (2020) and Haefke et al. (2013). These methods effectively track *average* wage growth over time. Nonetheless, in terms of the overall wage growth trends at the country level, they give similar results to the median-median method we report here.

Table 4. Cell sizes in the Indeed Wage Tracker (January 2019- December 2022)

	Average number of cells per month	Average number of job postings per cell-month
France	14,085	5
Germany	2,378	10
Ireland	827	5
Italy	680	4
Netherlands	4,366	7
Spain	772	3
UK	28,363	7
US	83,735	7

Notes: Data is from Indeed job postings. A cell refers to a group of postings with the same job title, salary type, and region in a given country.

The job title is extracted from the text of the job posting itself. When an employer posts a job on the platform they input a 'job posting title'. This text is cleaned algorithmically to remove information not related to the title of the role, such as the location, contract type, employer name, or common descriptive words that do not denote specific roles or tasks, such as "job" or "position". These cleaned titles are mapped to canonical Indeed job titles that group similar jobs. To give a sense of this process, Table 5 lists some examples.

Table 5. Job title mapping examples

Original job posting title	Indeed job title
Java engineer	Java developer
Java developer (Based in Surrey or remote)	Java developer
Anticipated middle school math teaching position	Mathematics teacher
Trace counter and showroom sales	Sales professional
Manage external supply operations	Supply manager
Site manager – residential developer	Site manager
Hiring hr mgr	Human resources manager
HR manager (12 month contract, hybrid)	Human resources manager

Job titles are a key input for search optimization, returning a list of relevant jobs for a given job search. Importantly for our purposes, they also help explain much of the cross-sectional variation in wages in the data. Table 6 shows the adjusted R-squared from two log wage regressions: one that controls for 54 broader occupational categories plus salary type and region, and one that replaces the occupational controls with job titles. The adjusted R-squared is much higher in the second regression, and the difference between the adjusted and unadjusted R-squared is very small, meaning that this is not simply an over-fitting of the data. This shows that Indeed job titles are a good predictor of wages, in line with the results in Marinescu and Wolthoff (2020), who use job titles to explain the cross-sectional dispersion of wages in online job postings.

Table 6. Adjusted R-squared from a log wage regression with different controls for occupation and job title

	Log wage regression			
	54 occupations	*Job titles*		
Country	Adj. R^2	Adj. R^2	R^2	N (millions)
Germany	0.44	0.70	0.70	1.4
France	0.32	0.61	0.61	5.5
Spain	0.26	0.46	0.46	0.3
Italy	0.20	0.58	0.59	0.3
Netherlands	0.39	0.61	0.61	2.0
Ireland	0.35	0.74	0.75	0.3
UK	0.41	0.67	0.67	14.3
US	0.34	0.69	0.69	36.3

Notes: R-squared from regressions of the log wage on occupation or job title dummies using Indeed data from Jan 2018 to Sep 2022. All specifications include fixed effects for year-month, region, and salary type (hourly, monthly, annual).

3.1. *Country-level results*

Using the methodology outlined above, Figure 6 shows the wage growth tracker for each of the eight countries in our dataset. For comparison, we also include selected other available sources of wage growth data. For five of the six euro area countries, we include the ECB's Negotiated Wage Tracker, constructed from wages in collective bargaining agreements (Koester et al. (2020), Górnicka and Koester (2023)).[12] For the US, we include the Atlanta Fed Wage Growth Tracker for job switchers, which is constructed from the Current Population Survey. For Ireland, we include wage growth for new hires from the Labor Force Survey, and for the UK, we include average wage growth for all workers from the Office for National Statistics.

Wage growth in online job postings accelerated sharply as countries reopened after the pandemic in 2022. In the US, growth peaked at over 9% year-on-year in the spring of 2022 before slowing to 6.4% by the end of 2022. In euro area countries and the UK, wage growth peaked at a lower level – closer to 6-7% – and later than in the US, reflecting, in part, the later reopening of European economies after the pandemic.

[12]We thank Górnicka and Koester (2023) for providing us with their monthly Negotiated Wage Growth Tracker series.

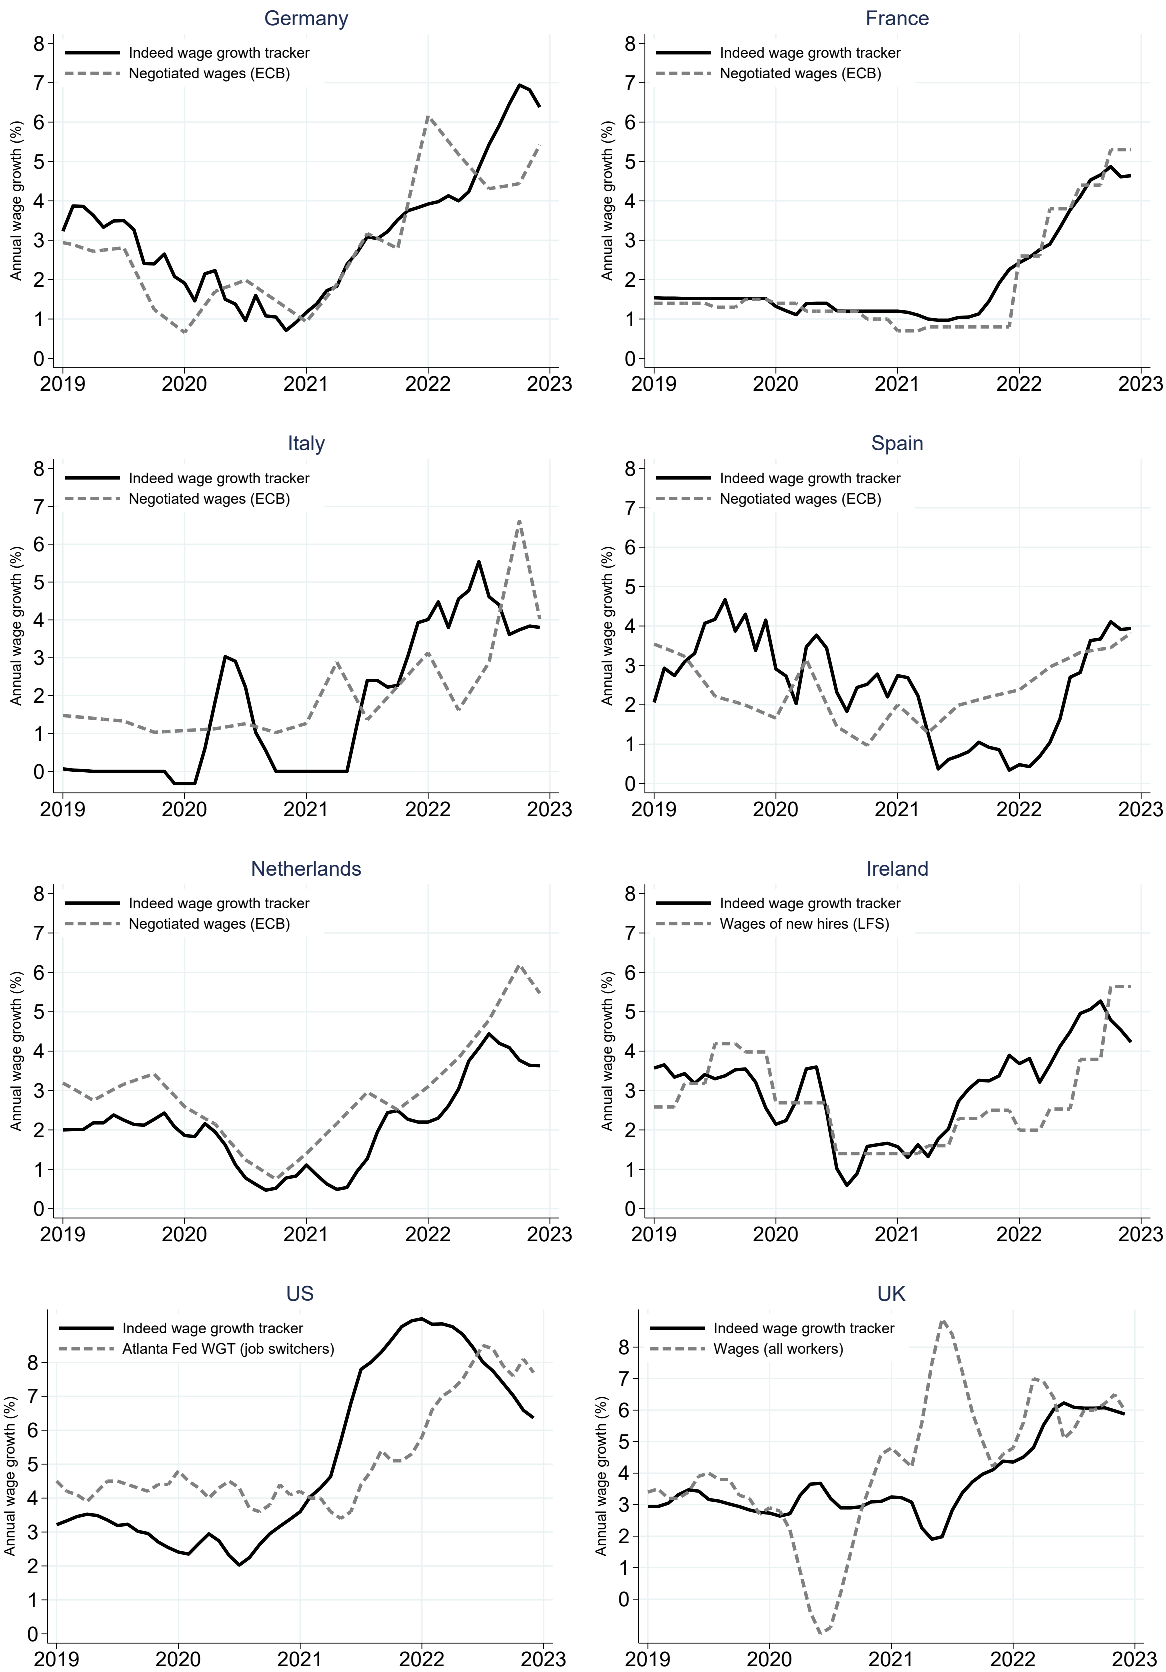

Fig. 6. Indeed Wage Growth Tracker and selected other wage trackers, Jan 2019- Dec 2022 (annual percentage change). *Source:* Monthly negotiated wages from the ECB (Górnicka and Koester (2023)). The Atlanta Fed Wage Growth Tracker is available online [link]. UK data from the ONS. *Notes:* See Section 3 for details of the IWT methodology, the chart shows 3-month averages and the last observation is December 2022.

We find that our estimates track negotiated wage growth closely in the five euro area countries where we have the ECB data – Germany, France, Italy, Spain, and the Netherlands. While the two series exhibit similar trends, the Indeed data has the advantage of being more timely, as it is available at least two months before the ECB estimates.

For countries where collective bargaining coverage is lower and where we have comparable data on the marginal worker – be it new hires in Ireland or job switchers in the US – the Indeed data on advertised wages tends to lead data on actual wage growth by several months.[13] In the UK, our tracker has tended to move contemporaneously with the wages of all workers, at least outside the pandemic period in 2020 and 2021, when changes in the wages of all workers were heavily impacted by compositional effects and the furlough scheme (Athow (2021)). The timeliness of the data supports the use of the Indeed Wage Tracker as an early or leading indicator of wage growth momentum for policy purposes.

3.2. *Robustness: sensitivity of wage growth tracker to alternative weights*

The country-level wage growth estimates presented in Figure 6 are unweighted. This means that in the third step in our methodology, where we take the median wage growth across all cells, every job title-region-salary type cell receives equal weight. Given our findings in Sections 2.3 and 2.4 that: a) some sectors and occupations are over- or under-represented in online job postings relative to vacancies and new hires, and b) the extent of over- and under-representation varies by country, in this section we look at the sensitivity of the wage tracker to four alternative weightings, calculated separately for each country:

1. Re-weighting to the *occupational shares* in the population of Indeed job postings;
2. Re-weighting to the *sector shares* of job vacancies in national statistics;
3. Re-weighting to the *occupational shares* of new hires from national surveys (the US is excluded from this comparison due to a lack of data);
4. Re-weighting to the *sector shares* of new hires from national surveys.

Re-weighting has been used in the context of online job posting data, e.g., by Turrell et al. (2021), who re-weight aggregate trends in UK job postings from Reed based on the occupation shares of employment. In the case of our wage tracker, re-weighting entails some changes to our previous methodology outlined above. Instead of taking the median

[13] Granger causality tests show that, statistically, the Indeed Wage Tracker is useful for predicting wage growth in the economy, as we show in Section 3.3

wage growth across cells at the *country level* in step 3, we calculate the median for either occupations or sectors, depending on the re-weighting in question. We then apply occupation or sector weights to these medians before averaging them to the country level, and these are re-weighted country estimates.

Figure 7 shows the results. The unweighted (i.e. baseline) and re-weighted series follow a broadly similar trend in all countries, with wage growth tending to weaken at the onset of the pandemic before recovering around re-opening. For the larger countries, there is very little difference between the different series. This is unsurprising, given the strong positive correlations we reported in the validation section.

For some countries – Ireland, Italy and Spain are the main examples – we do observe some differences. The largest disparities occur around the pandemic, during 2020 and the first half of 2021. The sector data on new hires confirms that pandemic-specific hiring (and layoff) effects are the key driver of this volatility. For example, in Ireland in the first half of 2020, the share of new hires in health sector increases three-fold, from 10 to 30%, whereas the shares in sectors that were shut due to the pandemic – consumer facing and construction are the main examples – collapses. This pattern then reverses when sectors reopen.

The un-weighted IWT estimates effectively smooth this pandemic-related volatility, much of which depends on how individual countries responded to the pandemic. From the perspective of monitoring underlying wage pressures, the unweighted estimates, we argue, are more relevant, especially around this period. However, from a validation perspective, the key conclusion we draw from the comparison of un-weighted and weighted estimates is that the overall trends in all series are very similar.

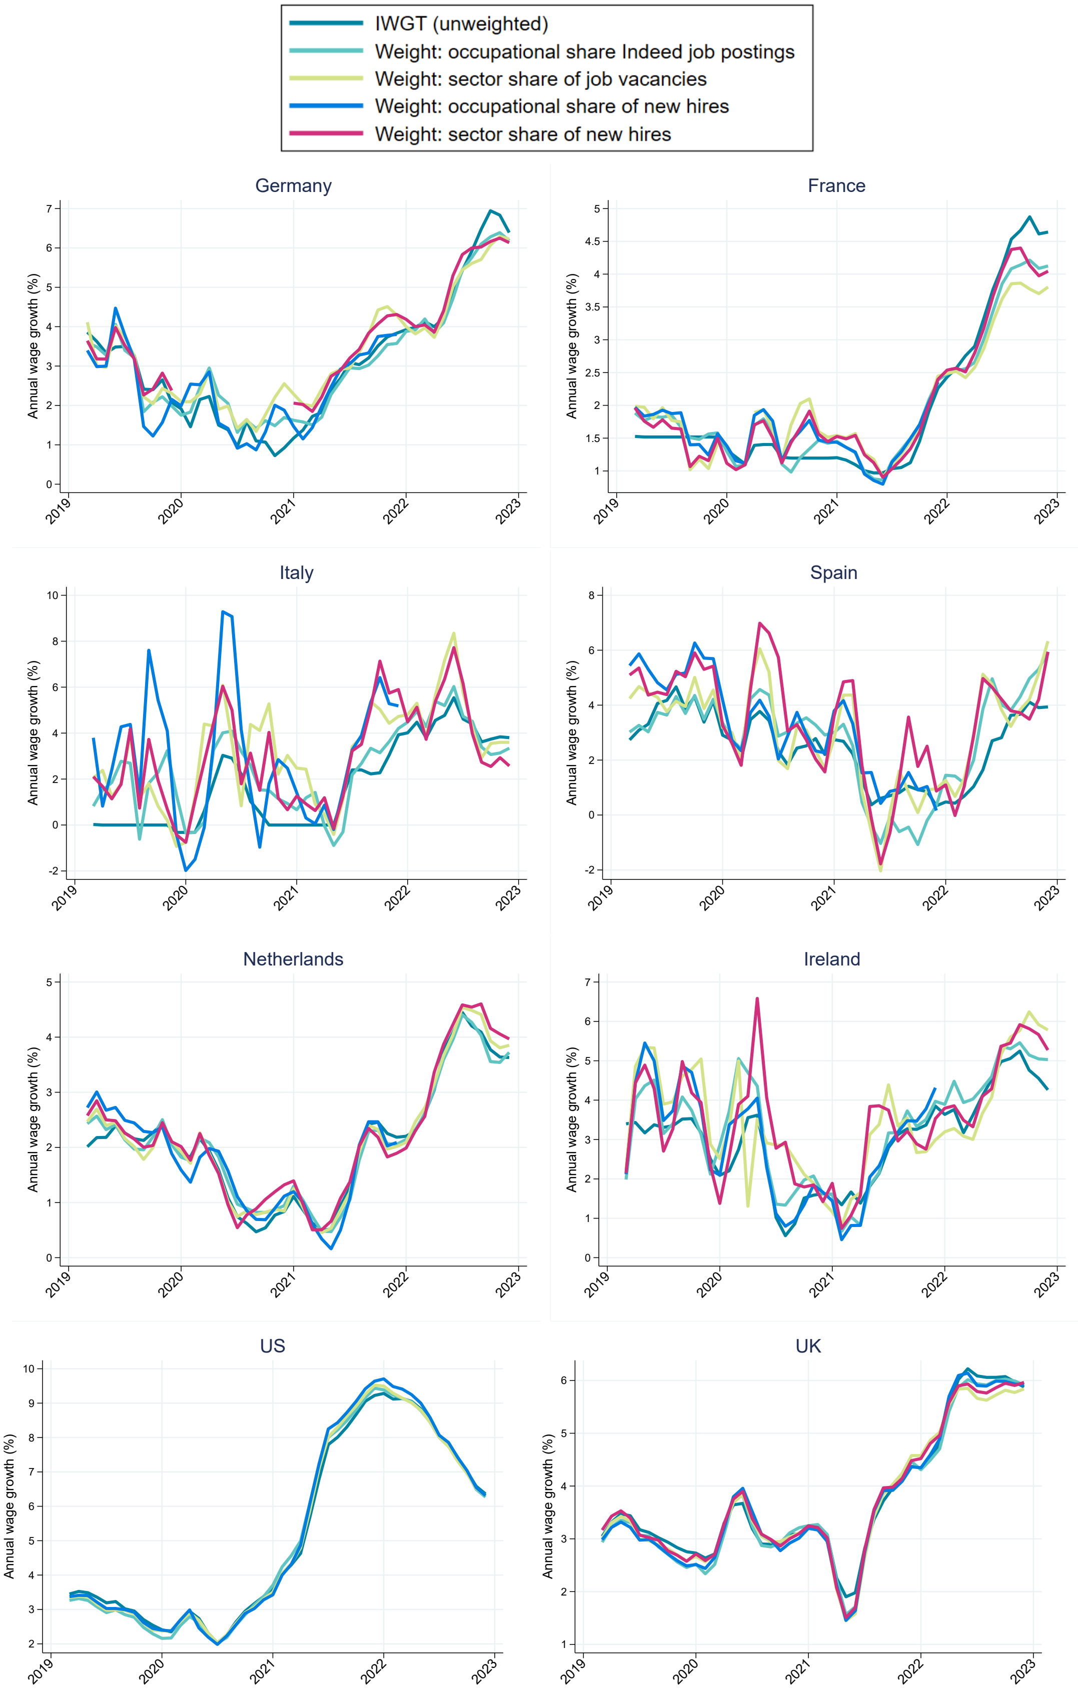

Fig. 7. Wage growth tracker with alternative weights. *Notes:* For data availability reasons, the following are not included: occupational shares of new hires for the US, all years; sector share of new hires for Germany, all of 2020; sector shares of job vacancies for France, January-March 2020; occupational shares of new hires for all six euro area countries, all of 2022.

3.3. *Wages from online job postings as a leading indicator of wage growth in the economy*

Forecasts of wage growth are a key output of forecasting exercises at all central banks. In this section, we ask what the Indeed Wage Tracker (IWT) can tell us about the future direction of wage growth in the economy.

The first thing to note is there are fundamentally two aspects to the potential leading indicator properties of the IWT. One is the *timeliness* of the data itself, which is available within a matter of days of the end of each month. This compares with lags of two months for some official estimates of wage growth, usually based on national accounts, such as compensation per employee in Europe.[14] Negotiated wage trackers have a similar lag, depending on when data on wage agreements is released. In the US and UK, certain measures of wage growth such as the quarterly Employment Cost Index (ECI), Average Labour Compensation per Hour (ALCH), and Average Weekly Earnings (AWE) also have a one- to two-month lag.

The second aspect, and the one we focus on here, is whether growth in advertised wages is informative about future wage dynamics, *irrespective of the superior timeliness of the data itself*. An employer's expectation of the wage required to make a new hire reflects forward-looking judgements around future demand, output, productivity, and labor market tightness.[15] Furthermore, as advertised wages relate to the marginal worker, it takes time for the factors that influence new hire wage dynamics to impact economy-wide average wages. We would therefore expect to see some relationship between lagged growth in our wage tracker and average wage growth more generally.

The rather short wage growth time series available to us, from January 2019 to December 2022, precludes a comprehensive forecasting assessment, such as comparisons of performance from rolling regressions over multiple forecast horizons.[16] Instead, we adopt a simpler approach consisting first of an assessment of Granger causality, followed by a comparison of the root mean squared errors from separate regressions of compensation growth on the IWT and the alternative wage growth trackers presented in Figure 6.

We test for Granger causality from our wage growth tracker to average wage growth to see whether, statistically, the IWT is informative about future rates of wage growth in the economy. For euro area countries, we use average annual growth in the wages and

[14]As noted on the Eurostat website, a flash estimate for GDP is available at t+30 days from quarter-end, but the data on GDP main aggregates, which includes the relevant data on compensation of employees, is only available at t+65 days.

[15]In search and matching frameworks this is reflected in the expected surplus from the job match, and the sharing of the surplus between the employer and the new hire.

[16]See Consolo et al. (2023) for an example of these approaches for forecasting wages in the euro area.

salary component of Eurostat's quarterly Labour Cost Index as our measure of economy-wide wage growth. For the US, we use the annual change in the quarterly Employment Cost Index (ECI), which is a composition adjusted version of growth in compensation of employees, and therefore closer to the IWT. For the UK, we use growth in compensation of employees from National Accounts.

The wage growth data is available quarterly for all countries (albeit at a significant lag, as noted above) and is a key forecast target for central banks. To match the monthly frequency of the Indeed Wage Tracker, we interpolate the data to monthly frequency using a centered three-month moving average. The underlying data is shown in Figure A8 in the appendix. The results, shown in Table 7, suggest that there is statistical causation from the IWT to wage growth, except for Italy, which is statistically insignificant.

Table 7. Granger causality tests:
VAR compensation per employee, Indeed Wage Tracker

Country	Chi^2	p-value	Lags	N
Pooled	13.45	0.04	6	336
Germany	22.05	0.00	4	44
France	56.53	0.00	3	45
Italy	5.69	0.13	3	45
Spain	12.34	0.01	3	45
Netherlands	21.57	0.00	3	45
Ireland	15.03	0.04	7	41
UK	13.48	0.01	4	44
US	49.71	0.00	6	42

Notes: Granger causality test results from a VAR of growth in compensation on the Indeed Wage Tracker. Lags are chosen according to the Akaike Information Criterion (AIC).

Next, we look at how well the IWT explains growth in compensation of employees in a regression, compared to using either (a) the negotiated wage trackers for the five euro area countries for which we have the data; and (b) the Atlanta Fed Wage Growth Tracker for the US. With no alternative wage growth trackers, both the UK and Ireland are excluded from this comparison.

To compare the explanatory variables, We summarise the regression adjusted R-squared and the ratio of the root mean squared errors (RMSE) from each regression. This is the ratio of the RMSE from the regression with the Indeed Wage Tracker only to the RMSE with the other wage growth trackers only. A value below one indicates that that the IWT can explain more of the data, relative to the other wage trackers.

Table 8 shows the results. For Germany, France, Netherlands and the US, the IWT performs better than alternative wage trackers, and this is without taking account of the more timely IWT data. In the case of Italy, the IWT performs only marginally better than the negotiated wage tracker, though both regressions leave more than half the data unexplained. For Spain, the IWT compares poorly with negotiated wages.

In Table 8, the US IWT stands out as performing much more strongly than the alternative tracker – here, the Atlanta Fed Wage Growth Tracker (AFWGT) for job switchers. Not only is the R-squared in the ECI-IWT regression higher (0.97, compared with 0.77), but in absolute error terms it the IWT is superior, as we see from the ratio comparison of the RMSE (0.34). Graphically, we can also see from A8 in the Appendix, that the IWT is much closer to the ECI. The most likely explanation lies with the differences in how the series are constructed, and in particular the underlying data that is used. The aim of the ECI is understand the hourly labor cost *to employers*, and, as such, the data and methods are similar to online job postings data we use to construct the IWT: it is establishment-level data that is composition adjusted by tracking wages within occupation-industry and geographic cells.[17] In contrast, the AFWGT is based on surveys of *individual employees*, from the monthly Current Population Survey, and is composition adjusted by comparing the same individuals twelve months apart.[18] The conclusion we draw is that if the aim of a given wage growth tracker is to understand in a timely manner the underlying labor cost pressures facing *firms*, and therefore pipeline pressures on prices and inflation, the IWT is likely to be more informative.

Table 8. Comparison of prediction errors from wage growth regressions

Country	R^2 (IWT)	R^2 (Other trackers)	RMSE ratio	Lags	N
Germany	0.66	0.40	**0.75**	4	44
France	0.78	0.70	**0.86**	4	45
Italy	0.49	0.46	**0.97**	4	44
Spain	0.43	0.77	1.59	3	45
Netherlands	0.90	0.84	**0.79**	3	45
US	0.97	0.77	**0.34**	6	42

Notes: R-squared from a regression of growth in compensation of employees (or ECI in the case of the US) on either the IWT or other wage trackers. For the euro area, the other wage trackers are negotiated wage growth, published by the ECB, and for the US we use the Atlanta Fed Wage Growth Tracker. The RMSE ratio is the ratio of the root mean squared error of the IWT regression to the other trackers regression. A value below one favours the IWT.

[17]See Ruser (2001) and the methodology section of the BLS-ECI website for details.
[18]See the methodology section of the AFWGT website for details.

4. HETEROGENEITY IN THE GROWTH OF ADVERTISED WAGES

One of the strengths of the online job postings data is the large number of observations, which means we can also assess wage developments below the country level. In this section, we exploit this data granularity to analyse the difference in advertised wage growth across the wage distribution.

To summarise the heterogeneity in wage growth trends, we show wage growth for different jobs depending on the wage level. To do this, we divide occupational categories into terciles depending on the level of their median advertised wage before the pandemic in 2018-19. We label the three groups low-, middle- and high-wage jobs. We then calculate wage growth trackers using the methodology described above, but separately for each country-tercile group. Figure 8 shows the results.

In almost all countries, there is a pattern of stronger wage growth for lower-paid jobs after the pandemic. The one exception is Spain, where the wage growth tracker for all three groups follows a broadly similar pattern, suggesting post-pandemic acceleration in wage growth was fairly general. In other countries, wages for lower-paid jobs either started to rise earlier and/or increased by more, compared to wages of middle- and higher-paid jobs. The differences can be very large. For example, in Germany, France, the Netherlands, the US, and UK – the gap in the wage growth rate between low- and high-paid jobs was over five percentage points at its peak.

The differential patterns of wage growth we observe are driven by a combination of supply, demand and other institutional or country-specific factors. Differences in the rate of inflation is one obvious candidate for explaining cross-country differences. Within countries, differences across occupations or wage terciles will also reflect the balance of labour demand and supply, in other words, labour market tightness.

Combining the Indeed data on wages and job postings can provide some insights into the role played by labor demand, which increased rapidly as economies re-opened after the pandemic (Figure 1). The question is whether the increases were larger for lower-paid jobs, thereby potentially contributing to the wage growth patterns we see in Figure 8. Figure 9 offers suggestive evidence for this narrative by showing the job posting trends for occupations by wage tercile. In several countries – notably Germany, France, Italy, Ireland, and the UK – we observe stronger growth in job postings for lower-paid jobs in recent years.

Fig. 8. Wage growth tracker for low-, middle- and high-wage jobs (annual percentage change). *Notes:* Data from job postings on Indeed. 'Low-wage' jobs are occupations that are in the bottom-third of wages posted in job ads in 2018-19, 'Middle-wage' are the middle third, and 'High-wage' are the top third of jobs by wage level. Hourly and monthly salary types are normalised to annual using average hours worked per week from country survey, multiplied by 52 or 12.

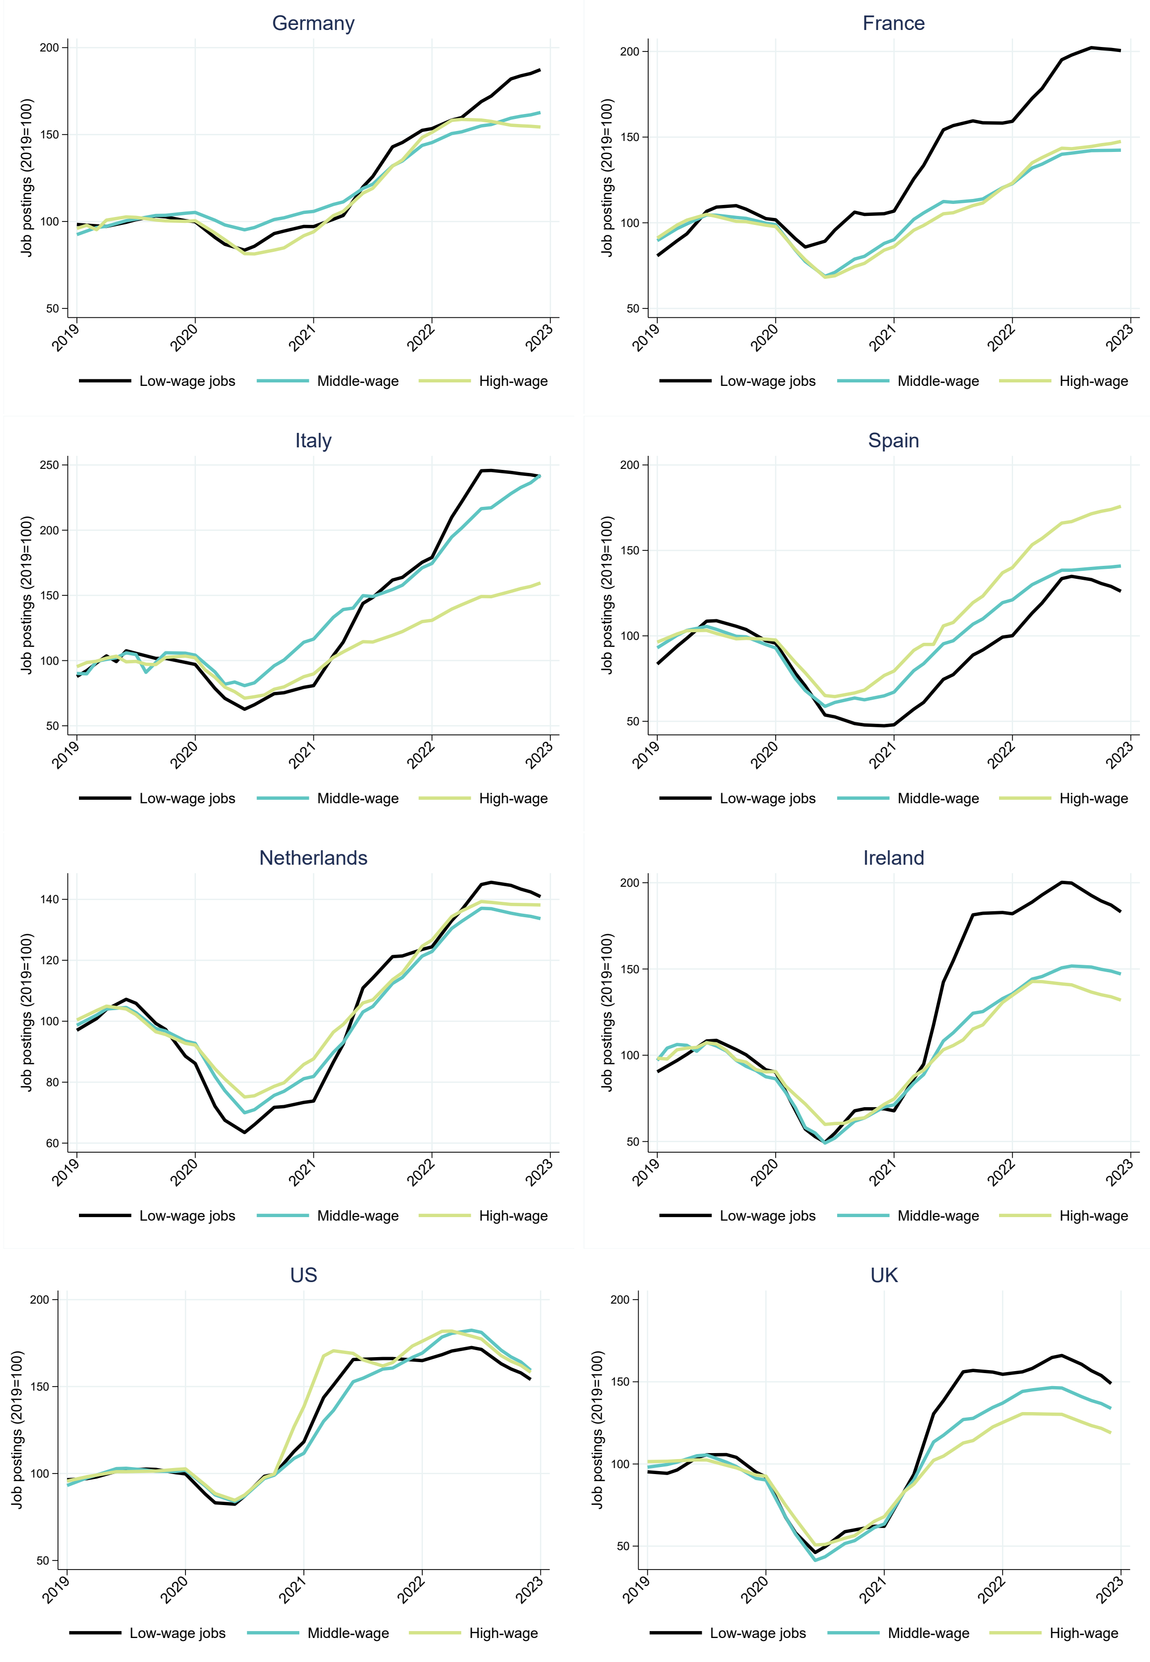

Fig. 9. Job postings for low-, middle- and high-wage jobs (2019=100). *Notes:* Data from job postings on Indeed. 'Low-wage' jobs are occupations that are in the bottom-third of wages posted in job ads in 2018-19, 'Middle-wage' are the middle third, and 'High-wage' are the top third of jobs by wage level.

In the online appendix, we combine the data in Figures 8 and 9 in a scatter plot of wage growth on job postings. The pattern for many countries is clear: higher-wage growth for lower-paid jobs after the pandemic tends to go hand-in-hand with the growth in job postings (see Figure A7), Spain being the obvious exception given similar trends for all wage terciles. Regression analysis of wage growth by tercile (Table A3) and occupation (Table A4) controlling for inflation and job postings confirms the scatter-plot relationship.

5. CONCLUSION

In this paper, we use data on wages advertised in tens of millions of job postings to construct a monthly wage growth tracker for eight countries from 2018 to 2022. The dataset has unique advantages, which include timeliness of delivery, availability of large sample sizes that allow for detailed breakdowns by occupation and sector, and the ability to compare trends across countries using a single source of information.

We show that the online data benchmarks well against national statistics and representative surveys on the distribution of job openings and new hires across sectors and occupations, as well as against the level of wages across sectors in each country. We then use the data to develop an indicator of wage growth. To the best of our knowledge, this is the first attempt in the literature to use online job postings to measure wage growth in multiple countries and to assess the relationship between this indicator and other labor market variables.

We show that trends in our wage growth tracker are similar to trends from other sources, including the Atlanta Fed Wage Growth Tracker for the US and the ECB's Negotiated Wages tracker for euro-area countries. As well as being more timely, our wage tracker appears to be a leading indicator of wage growth momentum in many countries.

Using the granular information in the online data, and looking underneath the country-level trends, we show wage growth after the pandemic was significantly stronger for lower-paid jobs in many countries. We show that stronger wage growth tends to be correlated with the large increase in job openings we see in the data, especially for lower-paid jobs, even after controlling for cross-country differences in inflation rates. This is suggestive of strong demand being one of the key factors driving stronger wage growth in recent years.

References

Adrjan, P., Ciminelli, G., Judes, A., Koelle, M., Schwellnus, C., and Sinclair, T. M. (2023a). Unlocked Potential: Work-from-Home Job Postings in 20 OECD Countries. *AEA Papers and Proceedings*, (113).

Adrjan, P., Gudell, S., Nix, E., Shrivastava, A., Sockin, J., and Starr, E. (2023b). We've Got You Covered: Employer and Employee Responses to Dobbs v. Jackson. IZA Discussion Paper No. 16360.

Adrjan, P. and Lydon, R. (2020). COVID-19 and the Global Labour Market: Impact on Job Postings. *Central Bank of Ireland Economics Letters Series*, 2020(3).

Anderton, R., Botelho, V., Consolo, A., da Silva, A. D., Foroni, C., Mohr, M., and Vivian, L. (2020). The Impact of the COVID-19 Pandemic on the Euro Area Labour Market. *ECB Economic Bulletin*, 8(2020).

Araki, S., Cazes, S., Garnero, A., and Salvatori, A. (2023). Under pressure: Labour Market and Wage Developments in OECD Countries. Chapter in OECD Employment Outlook 2023.

Athow, J. (2021). Far from average: How COVID-19 has Impacted the Average Weekly Earnings data. `https://blog.ons.gov.uk/2021/07/15/far-from-average-how-covid-19-has-impacted-the-average-weekly-earnings-data/`.

Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2020). Concentration in US Labor Markets: Evidence from Online Vacancy Data. *Labour Economics*, 66:101886.

Bank of England (2023). Monetary policy report. `https://www.bankofengland.co.uk/monetary-policy-report/2023/august-2023`.

Bellatin, A., Chu, V., and Galassi, G. (2023). What Can Job Posting Data Tell Us About the Labor Market? An Application to Recent Trends in Digitalization. Bank of Canada Staff Working Paper 2022-17.

Bhuller, M., Ferraro, D., Kostøl, A., and Vigtel, T. (2023). The Internet, Search Frictions and Aggregate Unemployment. NBER Working Paper No. 30911.

Cammeraat, E. and Squicciarini, M. (2021). Burning Glass Technologies' Data Use in Policy-relevant Analysis: An Occupation-level Assessment. OECD Science, Technology and Industry Working Paper No. 2021/05.

Carrillo-Tudela, C., Kaas, L., and Lochner, B. (2023). Matching Through Search Channels. IZA Discussion Paper No. 16583.

Chen, W. (2020). Disparities in Real Time. *Finance & Development (IMF)*, 57(004).

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., and The Opportunity Insights Team (2020). Real-time Economics: A New Platform to Track the Impacts of COVID-19 on People, Businesses, and Communities Using Private Sector Data. NBER Working Paper No. 27431.

Christodoulopoulou, S. and Kouvavas, O. (2022). Wages, Compositional Effects and the Business Cycle. ECB Working Paper No. 2022/2653.

Consolo, A., Foroni, C., and Martínez Hernández, C. (2023). A Mixed Frequency BVAR for the Euro Area Labour Market. *Oxford Bulletin of Economics and Statistics*, 85(5):1048–1082.

Daly, M. C. and Hobijn, B. (2017). Composition and Aggregate Real Wage Growth. *American Economic Review*, 107(5):349–52.

Djumalieva, J., Garasto, S., and Sleeman, C. (2020). Evaluating a new earnings indicator. Can we improve the timeliness of existing statistics on earnings by using salary information from online job adverts? ESCoE Discussion Paper 2020-19.

Forsythe, E., Kahn, L. B., Lange, F., and Wiczer, D. (2020). Labor Demand in the Time of COVID-19: Evidence From Vacancy Postings and UI Claims. *Journal of Public Economics*, 189:104238.

Górnicka, L. and Koester, G. (2023). A Forward-Looking Tracker of Negotiated Wages in the Euro Area. *ECB Occasional Paper*, forthcoming.

Haefke, C., Sonntag, M., and Van Rens, T. (2013). Wage Rigidity and Job Creation. *Journal of Monetary Economics*, 60(8):887–899.

Hansen, S., Lambert, P. J., Bloom, N., David, S. J., Sadun, R., and Taska, B. (2023). Remote Work across Jobs, Companies, and Space. NBER Working Paper No. 31007.

Hazell, J., Patterson, C., Sarsons, H., and Taska, B. (2022). National Wage Setting. NBER Working Paper No. 30623.

Hensvik, L., Le Barbanchon, T., and Rathelot, R. (2021). Job Search During the COVID-19 Crisis. *Journal of Public Economics*, 194:104349.

Hershbein, B. and Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review*, 108:1737–1772.

Kennedy, J. (2020). Coronavirus and UK Job Postings: Data from Indeed. `https://www.hiringlab.org/uk/blog/2020/04/14/coronavirus-and-uk-job-postings/`.

Koester, G., Benatti, N., and Vlad, A. (2020). Assessing wage dynamics during the COVID-19 pandemic: Can data on negotiated wages help? *ECB Economic Bulletin*, 8(2020).

Lafontaine, F., Saattvic, S., and Slade, M. (2023). No-Poaching Clauses in Franchise Contracts, Anticompetitive or Efficiency Enhancing? Working paper,SSRN.

Lane, P. R. (2022). Inflation Diagnostics. ECB Blog.

Marinescu, I. and Wolthoff, R. (2020). Opening the Black Box of the Matching Function: The Power of Words. *Journal of Labor Economics*, 38(2):535–568.

Mestres Domènech, J. (2022). Dinámicas salariales en España ¿Qué nos dice el indicador de salarios de Caixabank Research? *Informe Mensual - La Caixa*, (467):35–36.

Modestino, A. S., Shoag, D., and Ballance, J. (2020). Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful? *Review of Economics and Statistics*, 102(4):793–805.

Napierała, J., Kvetan, V., and Branka, J. (2022). Assessing the Representativeness of Online Job Advertisements. CEDEFOP Working Paper No. 17.

Pham, T., Talavera, O., and Wu, Z. (2023). Labor Markets During War Time: Evidence From Online Job Advertisements. *Journal of Comparative Economics*. In press.

Powell, J. H. (2022). Monetary policy and price stability. *Speech at the economic policy symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming.*

Powell, J. H. (2023). Inflation: Progress and the path ahead. *Speech at the economic policy symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming.*

Ruser, J. W. (2001). The Employment Cost Index: what is it? *Monthly Labor Review*, 2001(September):3–16.

Soh, J. S., Oikonomou, M., Pizzinelli, C., Shibata, I., and Mendes Tavares, M. (2022). Did the Covid-19 Recession Increase the Demand for Digital Occupations in the United States? Evidence from Employment and Vacancies Data. IMF Working Paper 2022-195.

Turrell, A., Speigner, B., Copple, D., Djumalieva, J., and Thurgood, J. (2021). Is the UK's Productivity Puzzle Mostly Driven by Occupational Mismatch? An Analysis Using Big Data on Job Vacancies. *Labour Economics*, 71:102013.

Online Appendix for: What do wages in online job postings tell us about wage growth?

A. BENCHMARKING ONLINE JOB POSTING DATA - SUPPORTING INFORMATION

Table A1. Geographic level of detail used to define cells in the wage tracker

Country	**Geographic information**	**N**
US	States	52
UK	Counties/Shires	92
Germany	States	16
France	Departments	96
Italy	Administrative Regions	20
Spain	Autonomous Regions	17
Netherlands	Provinces	12
Ireland	Counties	26

Table A2. Correlation of the sector shares of Indeed job postings with official statistics on new hires, 2018-22

	Germany	France	Spain	Italy	Netherlands	Ireland	US	UK
All Indeed job postings	0.83	0.83	0.45	0.63	0.70	0.67	0.74	0.61
Postings with wages	0.84	0.81	0.59	0.67	0.60	0.67	0.64	0.65
N	256	320	320	320	320	320	260	320

Notes: Data is quarterly from 2018 Q1 to 2022 Q4, with the exception of Germany, which excludes 2020 data. New hire data is from Eurostat for the euro area countries, ONS for the UK, and JOLTS for the US. Sample sizes vary depending on the number of sectors and occupations for which new hire statistics are available in each country.

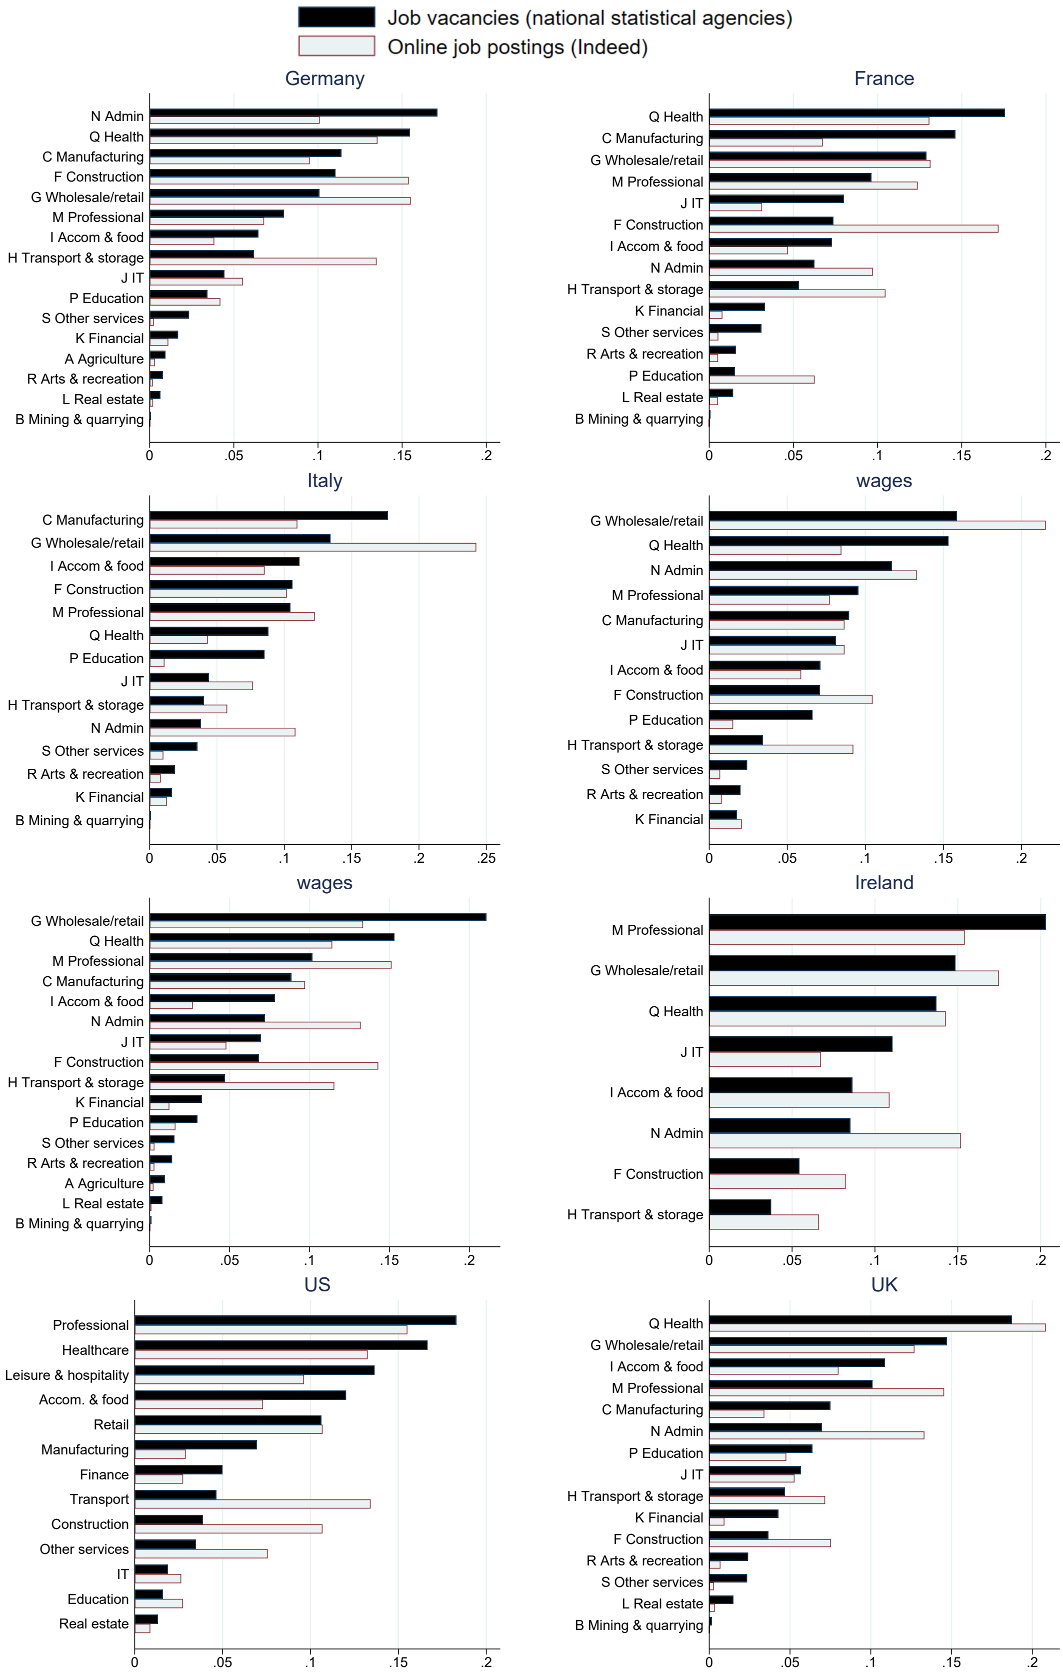

Fig. A1. Sector shares of Indeed online job postings *with wages* and job vacancies from national statistical agencies (2018-22) *Source:* Eurostat for euro area countries, ONS for UK, and JOLTS for US. *Notes:* Data for France is for firms with ten employees or more. In the absence of job vacancy numbers for Italy, the chart data is estimated from sector vacancy rates and employment from Eurostat. US data is non-farm job openings.

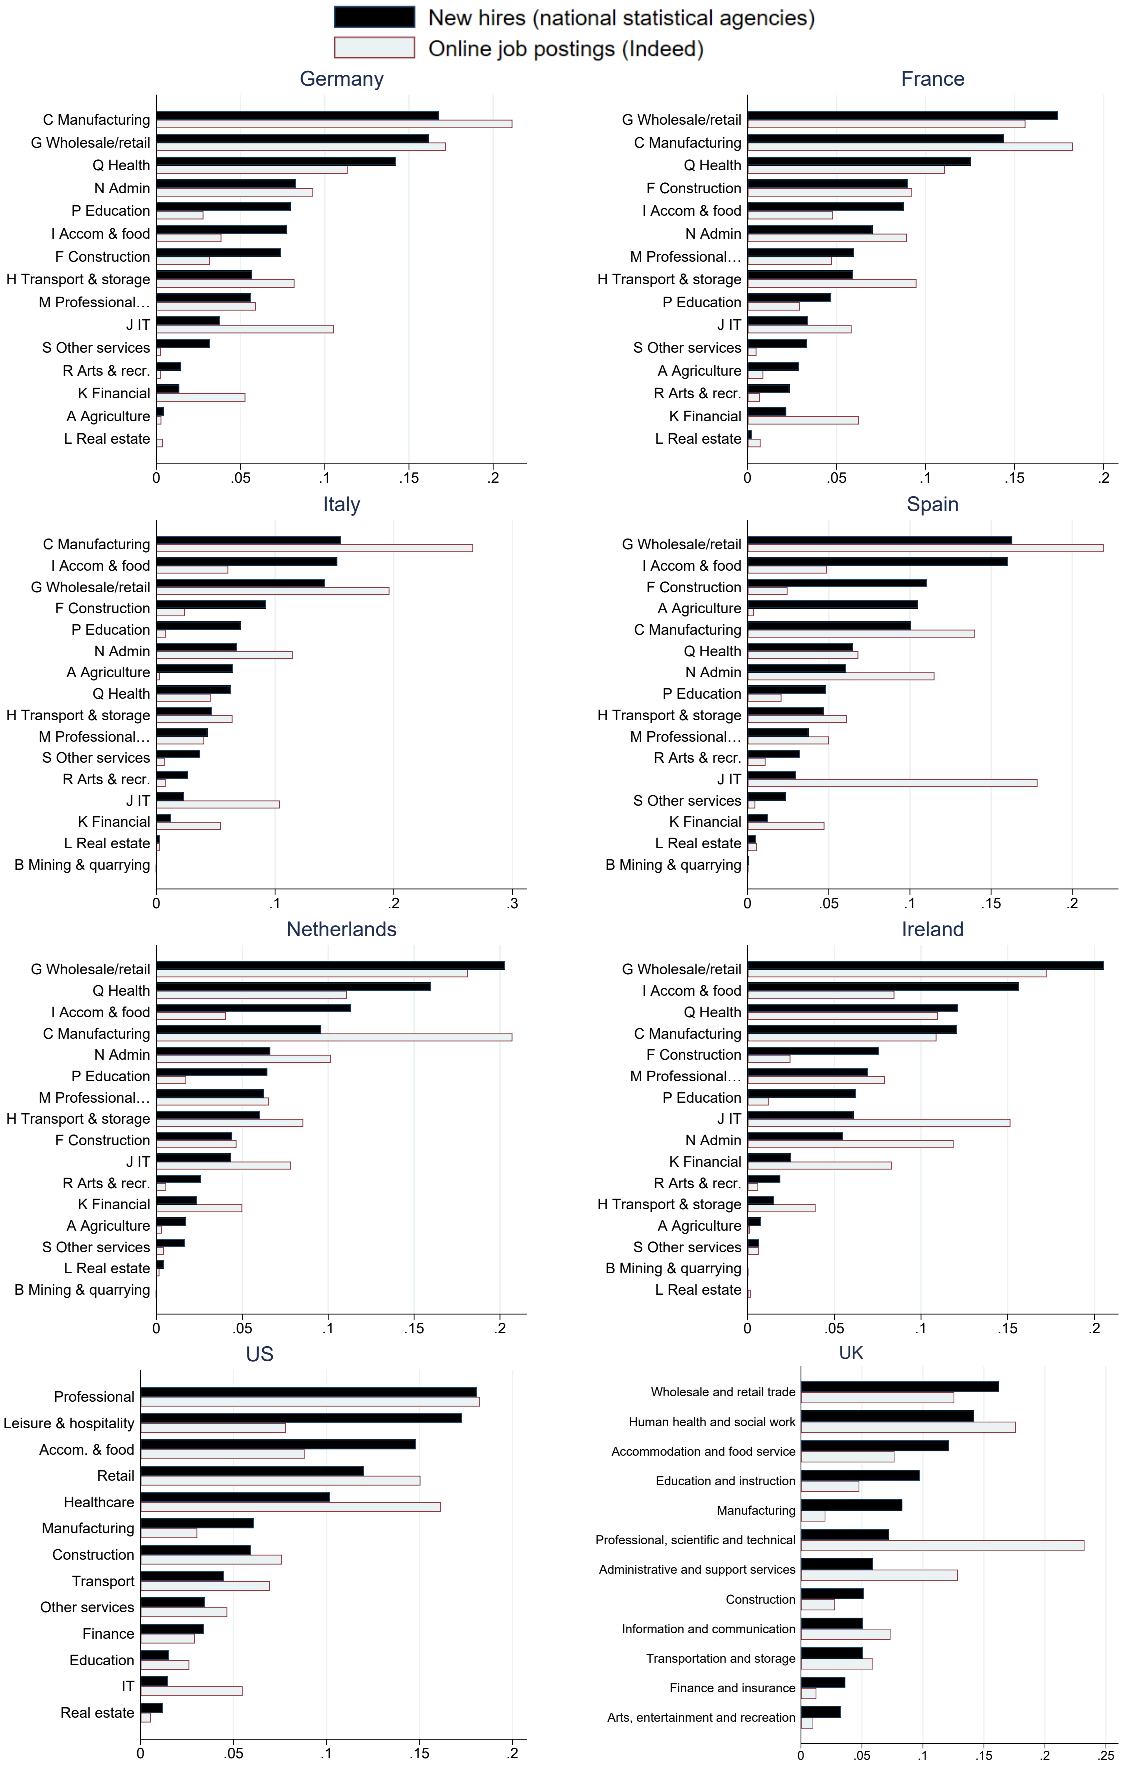

Fig. A2. Sector shares of Indeed job postings and new hires from national statistical agencies. *Source:* Eurostat for euro area countries (2018-21), ONS for UK (2018-22), JOLTS for US (2018-22).

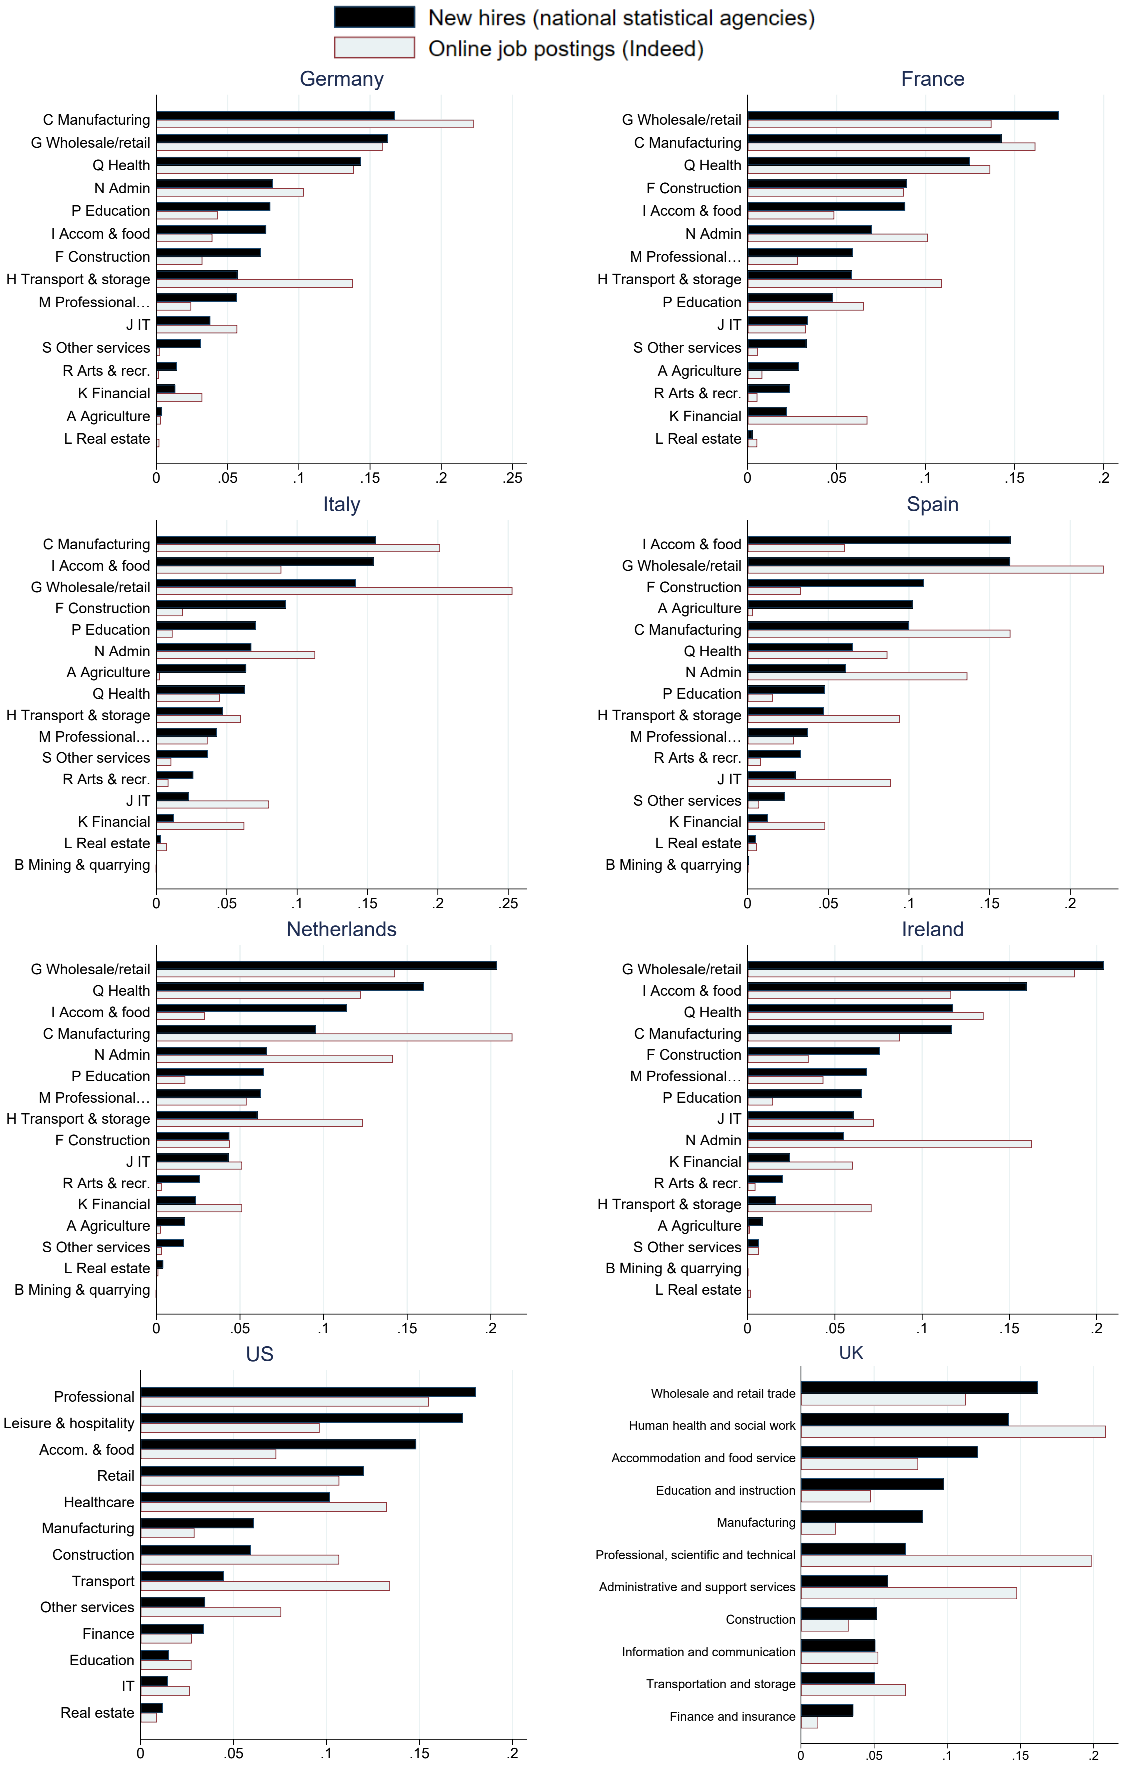

Fig. A3. Sector shares of Indeed job postings *with wages* and new hires from national statistical agencies. *Source:* Eurostat for euro area countries (2018-21), ONS for UK (2018-22), JOLTS for US (2018-22).

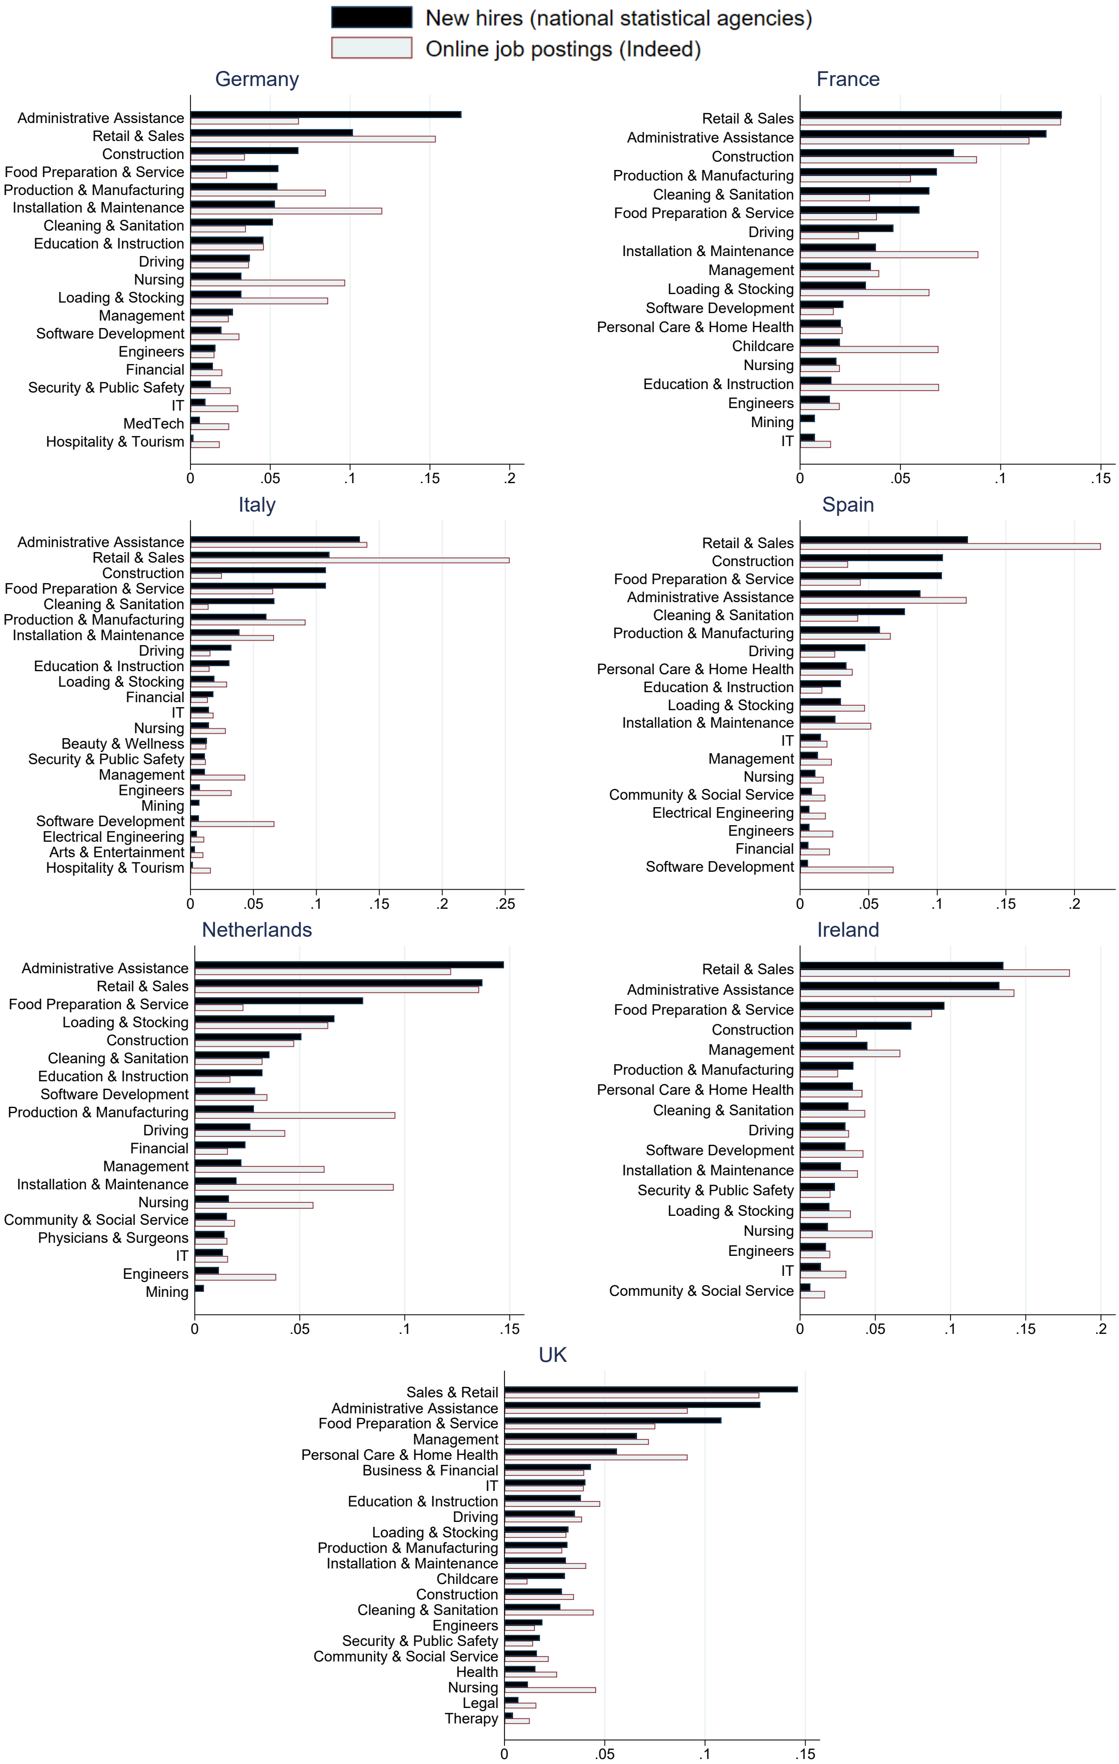

Fig. A4. Occupational shares of Indeed job postings *with wages* and new hires from national statistical agencies. *Source:* Eurostat for euro area countries (2018-21), ONS for UK (2018-22).

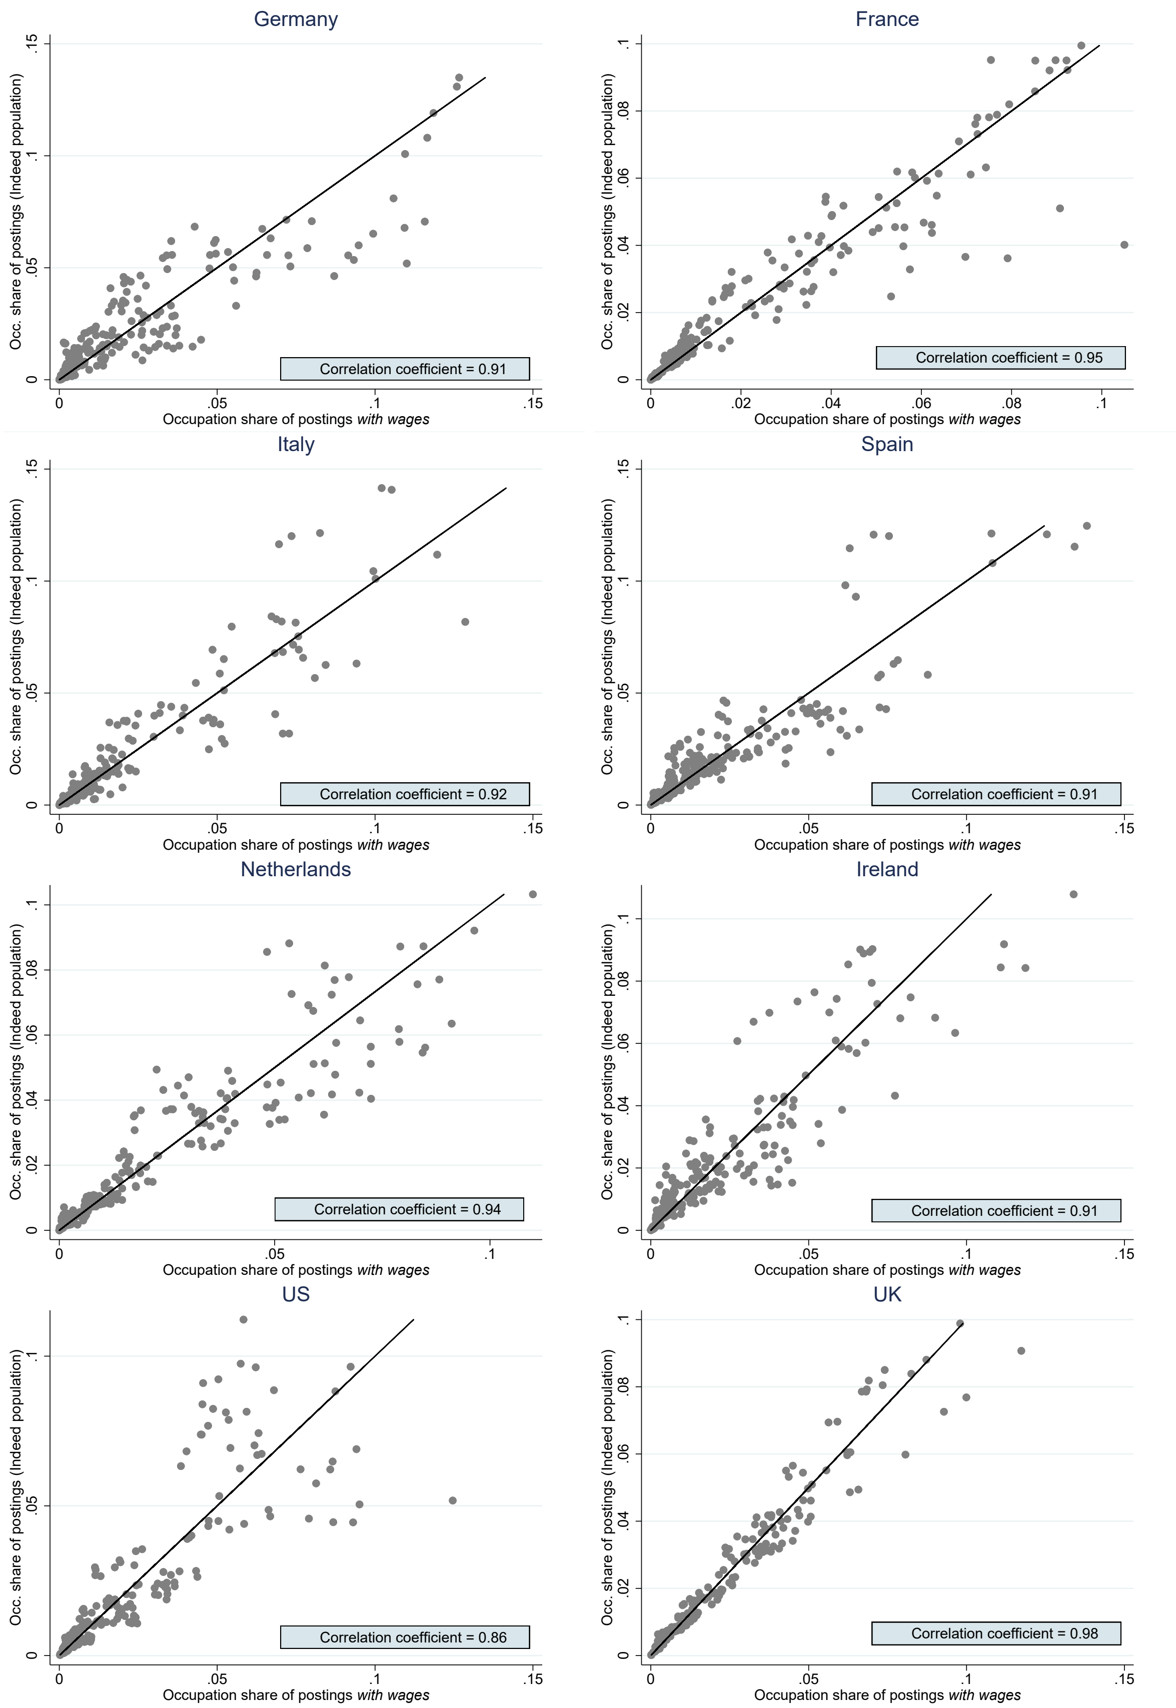

Fig. A5. Occupational shares of *all* Indeed Job postings (y-axis) and subset of postings that *include* wages (x-axis). *Source:* Indeed data. *Notes*: Each gray dot is the *occupational* share of job postings in a given year (2018-22) from either the full sample of all job postings (y-axis) or the sub-sample that just includes wages (x-axis). The black line is a 45-degree line.

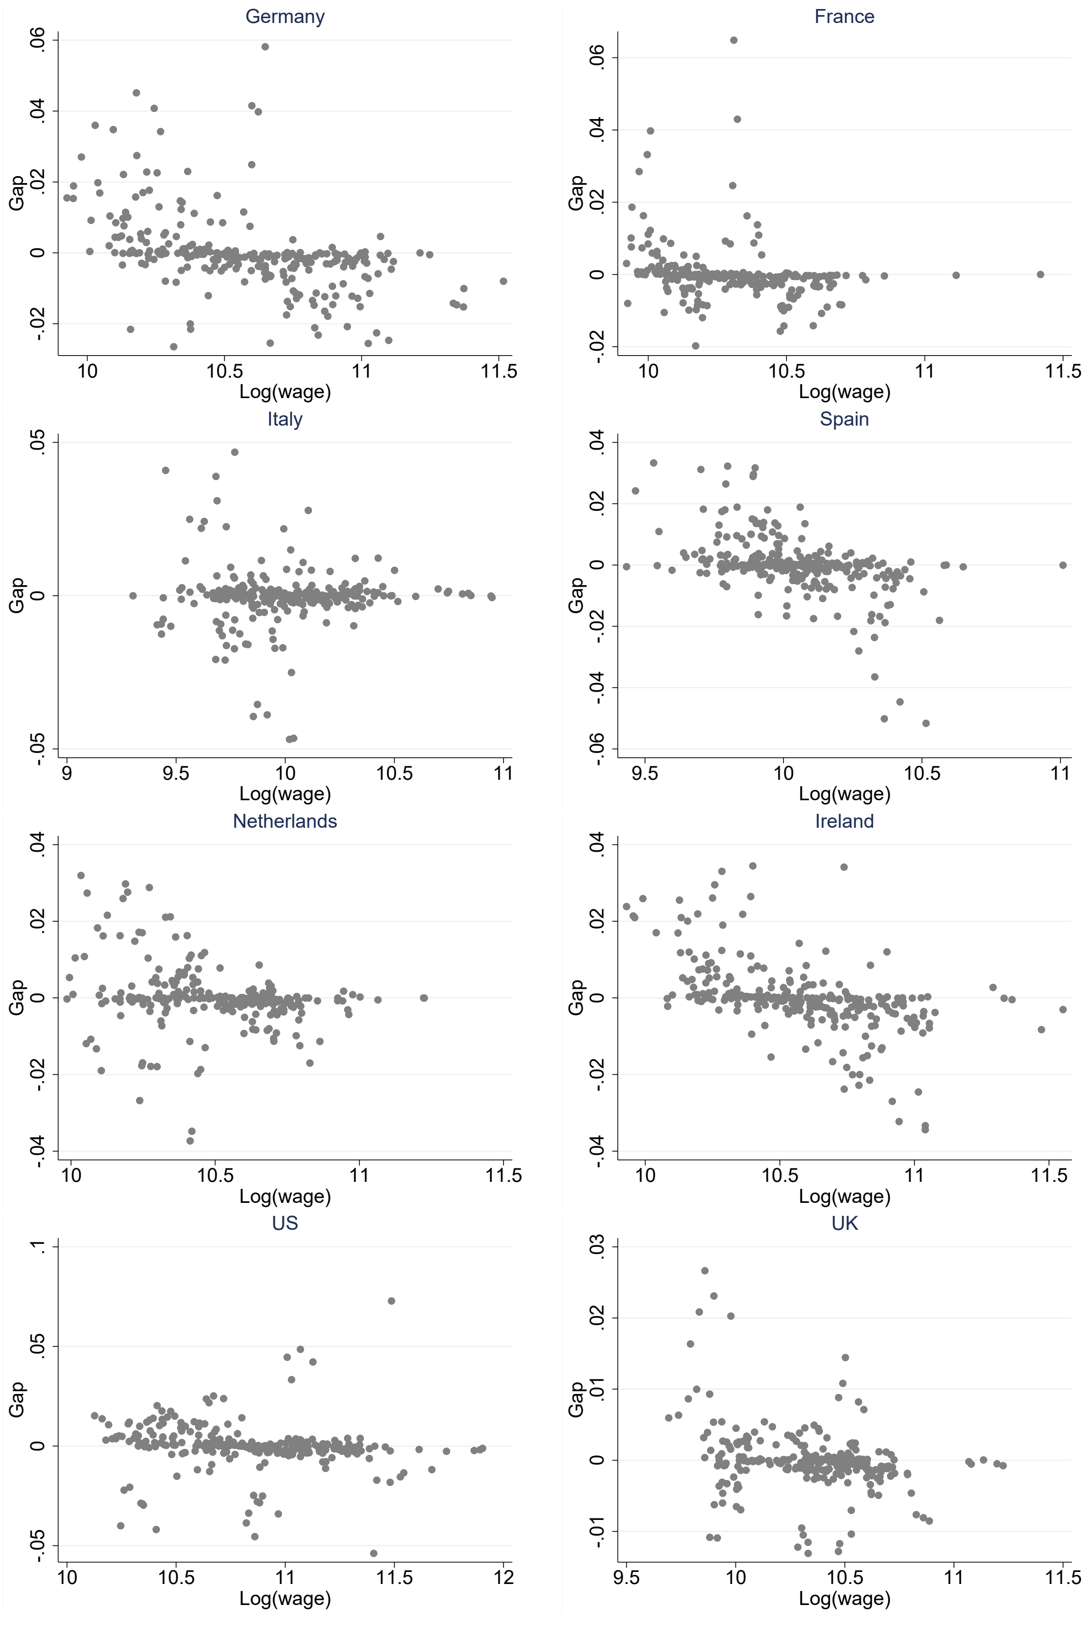

Fig. A6. Correlation between (1) the difference in occupational shares in the sub-sample of Indeed job and the population of Indeed job postings ('Gap'), and (2) average log wages by occupation in Indeed job postings. *Source:* Indeed data. *Notes*: Each dot is the 'gap' or difference between the occupational share in a given year in the sub-sample of Indeed job postings that include wages and the full population of Indeed job postings. The x-axis is the log of mean annual wage within occupations, by year. Hourly and weekly posted wages are annualised using average hours worked per week from country surveys and assuming 52 weeks worked per year. A downward slope means that, relative to higher paying jobs, lower wage job postings tend to be more likely to include wages in the job ad.

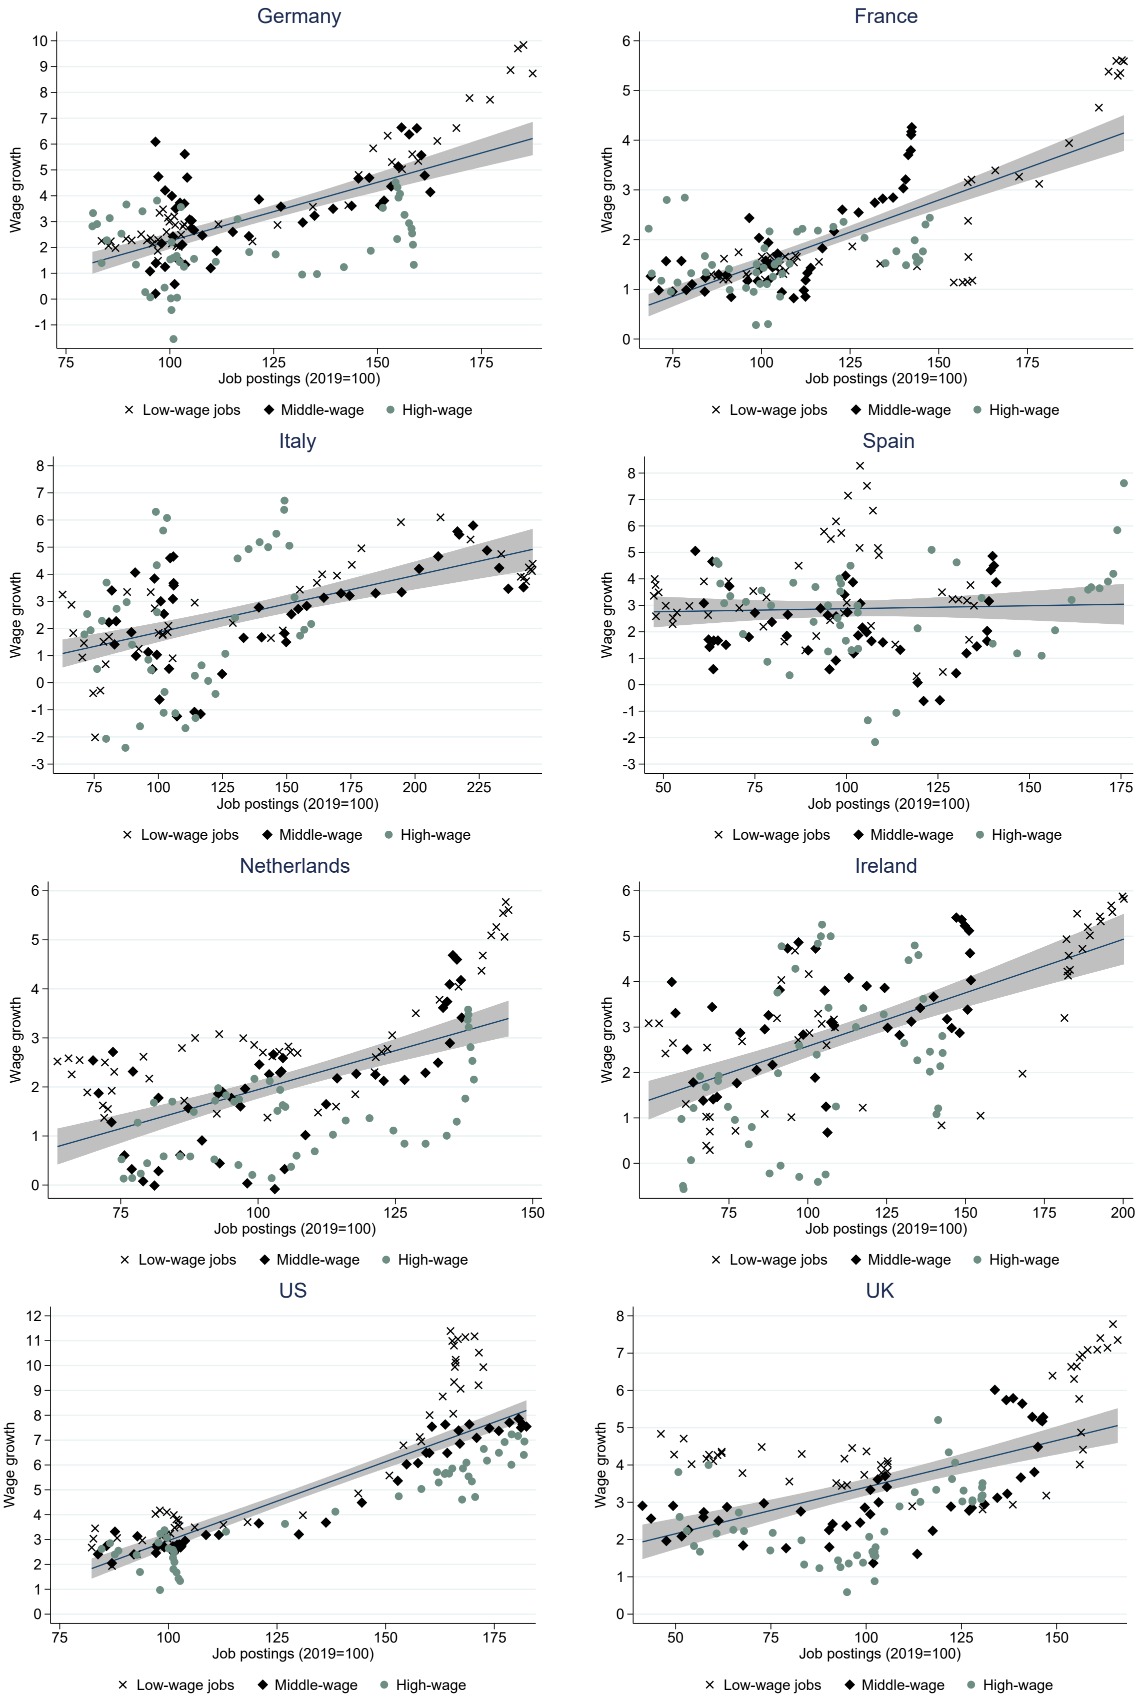

Fig. A7. Cross-plot of wage growth and growth in job postings for low-, middle- and high-wage jobs. *Notes:* 'Low-wage' jobs are occupations that are in the bottom-third of wages posted in job ads in 2018-19, 'Middle-wage' are the middle third, and 'Top-wage' are the top third of jobs by wage level. Hourly and monthly salary types converted to annual using average weekly hours worked from country surveys, multiplied by 52 or 12.

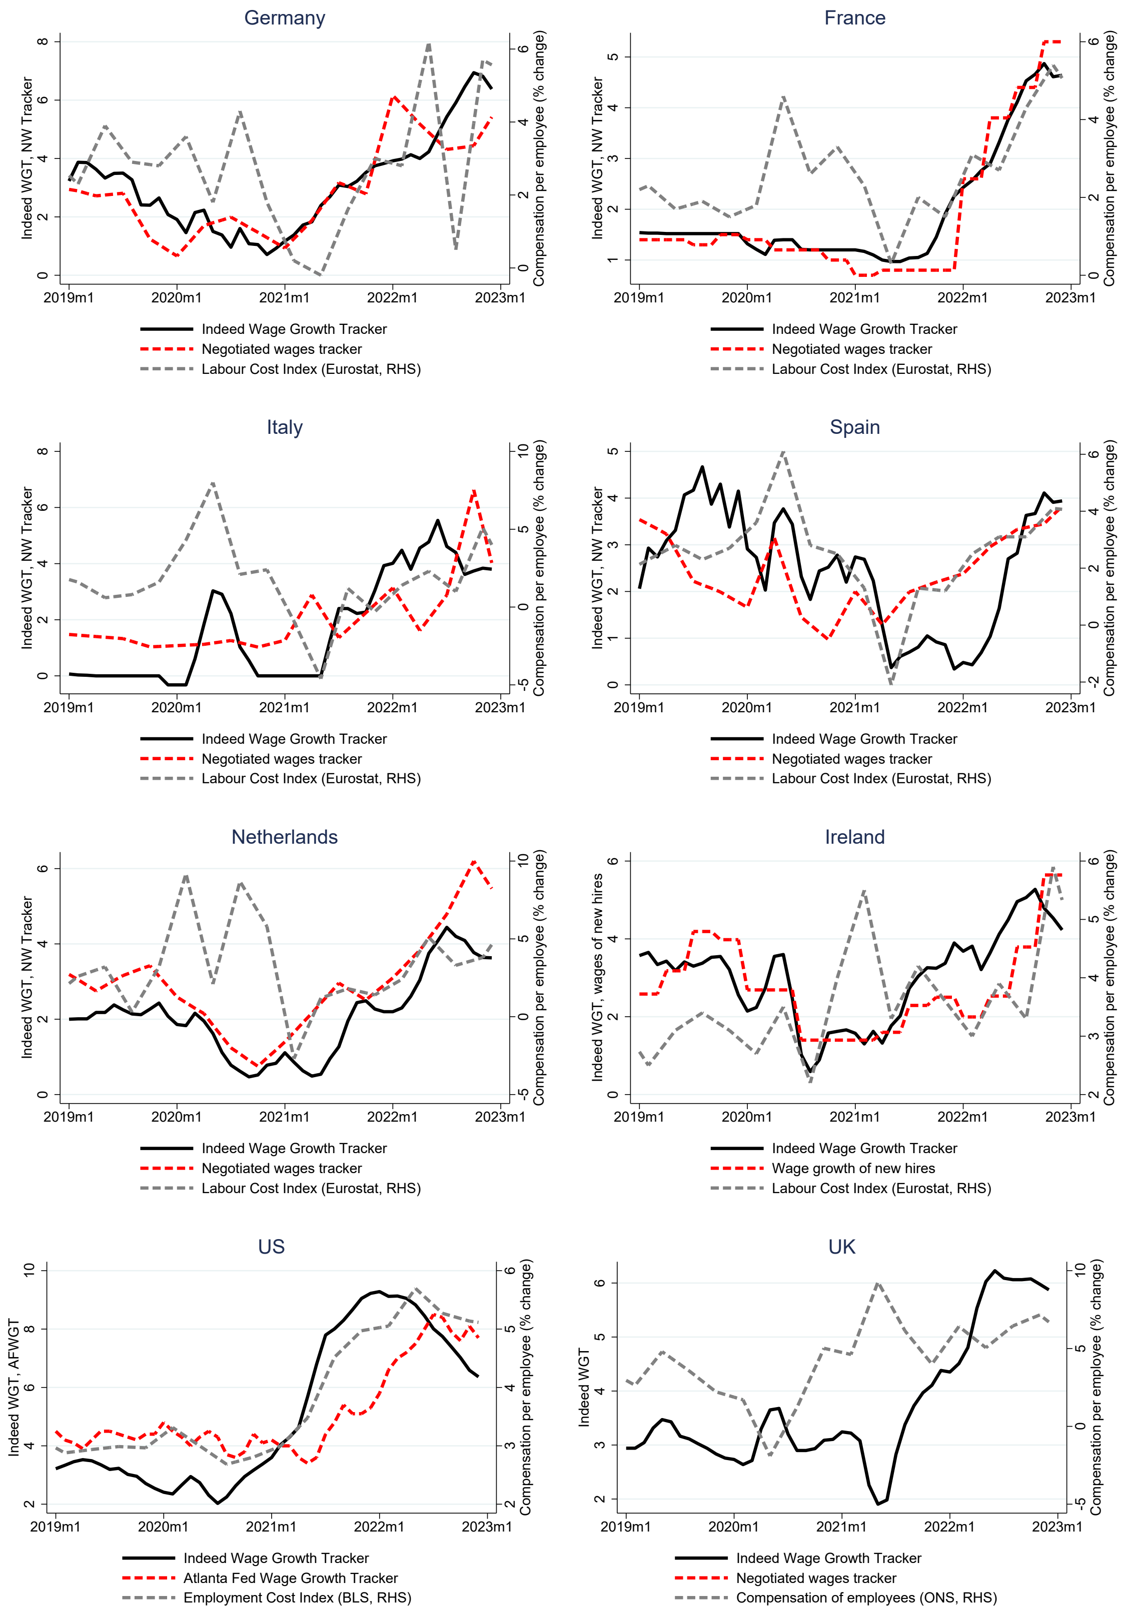

Fig. A8. Indeed Wage Tracker and economy-wide wage growth, 2019-22 (annual percentage change). *Notes:* Wage growth in the Indeed Wage Tracker in each month is the median of the distribution of annual changes in the median wage advertised in job postings on Indeed for each job title, region, and salary type cell. The chart shows 3-month averages. The last observation is December 2022. Labour Cost Index from Eurostat, Employee Cost Index from the BLS, Compensation of employees from the ONS.

Table A3. Regressions of wage growth (per cent) on job postings (index 2019=100) and inflation (per cent annual change) by **tercile** of the 2018-19 wage distribution from job postings

	Pooled	Germany	France	Italy	Spain	Netherlands	Ireland	US	UK
Job postings	0.017	0.036	0.009	0.019	0.005	0.026	0.009	0.059	0.005
	(6.84)	(4.31)	(2.48)	(2.81)	(0.35)	(6.87)	(1.19)	(13.24)	(1.62)
Job postings	-0.009	-0.005	-0.002	-0.001	-0.013	-0.008	0.002	-0.013	-0.015
x middle-wage	(3.02)	(2.76)	(1.48)	(0.23)	(2.77)	(5.93)	(0.60)	(10.24)	(11.82)
Job postings	-0.013	-0.016	-0.006	-0.000	-0.005	-0.015	-0.008	-0.018	-0.024
x top-wage	(3.90)	(7.73)	(3.91)	(0.08)	(1.02)	(10.45)	(2.39)	(14.30)	(18.01)
Inflation	0.464	0.290	0.49	0.09	0.240	0.400	0.314	0.349	0.866
	(9.11)	(1.75)	(5.09)	(0.42)	(0.93)	(7.28)	(1.91)	(4.40)	(14.49)
R^2	0.38	0.61	0.64	0.15	0.06	0.71	0.28	0.90	0.83
N	1,095	136	138	131	138	138	138	138	138
H2=0	0.02	0.00	0.15	0.00	0.58	0.00	0.25	0.00	0.00
H3=0	0.32	0.03	0.51	0.04	0.99	0.00	0.95	0.00	0.00

Notes: Coefficients from regressions of advertised wage growth by wage tercile (low-, middle- and high-paid jobs) from January 2019-December 2022, controlling for job postings (indexed to 2019=100) and core inflation (percentage change on previous year). H2 is the p-value for the null hypothesis that the sum of the coefficients on *job postings* and the interaction of *job-postings* and *middle-wage jobs* equals zero. H3 is the same for the high-wage jobs interaction. t-statistics in parentheses.

Table A4. Regressions of wage growth (per cent) on job postings (index 2019=100) and inflation (per cent annual change) by **occupation**

	Pooled	Germany	France	Italy	Spain	Netherlands	Ireland	US	UK
Job postings	0.021	0.017	0.009	0.032	0.015	0.030	0.006	0.051	0.011
	(14.82)	(4.88)	(6.81)	(7.13)	(2.79)	(14.10)	(1.76)	(31.54)	(7.93)
Job postings	-0.005	-0.007	-0.004	0.001	-0.009	-0.008	-0.002	-0.010	-0.011
x middle-wage	(2.62)	(4.73)	(4.10)	(0.39)	(2.68)	(7.00)	(0.76)	(13.65)	(12.47)
Job postings	-0.008	-0.018	-0.006	-0.000	-0.013	-0.011	-0.007	-0.016	-0.021
x top-wage	(3.77)	(10.66)	(7.21)	(0.19)	(3.97)	(10.38)	(3.04)	(21.00)	(20.95)
Inflation	0.480	0.733	0.52	0.06	0.477	0.360	0.471	0.416	0.700
	(16.21)	(7.05)	(10.15)	(0.31)	(3.76)	(7.42)	(6.13)	(11.71)	(18.02)
R^2	0.18	0.21	0.21	0.08	0.04	0.26	0.06	0.69	0.83
N	10,180	1,151	1,364	131	1,019	1,350	1,324	1,481	1,305
H2=0	0.00	0.01	0.00	0.00	0.15	0.00	0.29	0.00	0.73
H3=0	0.00	0.85	0.06	0.00	0.65	0.00	0.69	0.00	0.00

Notes: Coefficients from a regression of advertised wage growth by wage occupation from January 2019-December 2022, controlling for job postings (indexed to 2019=100) and core inflation (percentage change on previous year).H2 is the p-value for the null hypothesis that the sum of the coefficients on *job postings* and the interaction of *job-postings* and *middle-wage jobs* equals zero. H3 is the same for the high-wage jobs interaction. t-statistics in parentheses.