Aggregation bias in wage rigidity estimation

Persyn, Damiaan

University of Göttingen

3 March 2021

Aggregation bias in wage rigidity estimation

Damiaan Persyn*
University of Göttingen

December 2020

Abstract

I argue in this paper that the estimation of wage rigidity using country level data suffers from aggregation bias. Using European data for the years 2000-2017, I find that wages respond less flexibly to changes in unemployment at the regional level, compared to estimation using the same data aggregated at the country level. A possible explanation is that in the European data changes in aggregate unemployment tend to be driven by regions with low unemployment rates, while unemployment in regions with high unemployment rates is less variable and less responsive to aggregate shocks. The relationship between unemployment and wages —the wage curve— is downward sloping and convex. Due to this nonlinearity, the higher variability in lower regional unemployment rates implies higher observed wage flexibility at the aggregate country level, and biased inference. The implication is that wages are even less responsive to changes in unemployment than is observed in aggregate data and commonly assumed in macro-economic models, such that for example fiscal stimulus would lead to less wage inflation than anticipated.

1 Introduction

The degree of wage rigidity in an economy plays a crucial role in determining how economic shocks affect employment and unemployment. This is the case both in the real world, and in macroeconomic models used to evaluate and steer fiscal and monetary policy. The vast majority of empirical and theoretical macro-economic models considering wage rigidities operate on the national level. This is intuitive since many fiscal and monetary policy questions are defined on the level of countries rather than at the regional level. Moreover, institutions such as labour unions or public unemployment insurance schemes that may shape wage rigidities operate at the national level or at

*Contact: Department of Agricultural Economics and Rural Development, University of Göttingen, Germany. damiaan.persyn@uni-goettingen.de I am greatly indebted to Ragnar Nymoen for many useful comments and insights that have significantly improved this paper. I also gratefully acknowledge helpful comments from Katja Heinisch, Javier Barbero, Enrique Lopez Bazo, Raul Ramos; participants at the Barcelona AQR group regional workshop, and the ERSA conference in Lyon. All remaining errors are mine.

least have an important national component, such that functional forms and parameters governing wage rigidity may be shared between regions within the same country. Another reason for performing analysis at the country-level is that the required data often is not available at the regional level, or models become intractable or difficult to handle computationally at a fine level of spatial disaggregation. Such reasons may explain why influential studies such as Hagedorn and Manovskii (2013) or Gertler *et al.* (2020) use thousands of observations on individual wages and worker characteristics, but consider the effect of the US-wide unemployment rate on these wages, rather than the local unemployment rate. A fundamental problem with this approach is that the relevant labour market for most workers is local, and changes in unemployment rates are unevenly distributed in space. If the relationship under investigation is nonlinear, estimation using country level data is biased if shocks are unevenly distributed between regions, even if the relation is identical in all regions.

Wage rigidities have been key to reconciling DSGE models encompassing search and matching with the high cyclicality of unemployment observed in the data. Most search and matching models imply a simple relationship between the level of unemployment and wages, i.e. a wage curve, reflecting the level of wage rigidity. There is a trend to perform reduced form estimation of this relationship, rather than attempting to jointly estimate or calibrate the deeper parameters underlying it together with the larger model. There are good reasons for this. The estimated wage curve elasticity may serve as a sufficient statistic, summarising all that is relevant about the labour supply and wage rigidities in the labour market. The estimated wage curve can be combined with labour demand to close the model. The separately estimated wage curve elasticity is a portable statistic, allowing to considering how identification strategies affect estimates, or -as in this paper- the level of aggregation. Such issues are much harder to track when jointly estimating of a large system of equations (see for example Chetty, 2009; Andrews *et al.*, 2017; Nakamura and Steinsson, 2018). Two recent examples of this approach are Beraja *et al.* (2019), who iterate between a DSGE model at the aggregate level and reduced form wage curve estimation using instrumentation at the regional level; and Koenig *et al.* (2020) who use reduced form wage curve estimation to investigate how introducing backward-looking reference wages in a search and matching model can reproduce a reduced-form estimated wage elasticity.

This paper uses reduced form wage curve estimation, to argue that estimation of wage rigidity using country level data suffers from spatial aggregation bias. Using European regional data from 2000 to 2017, considering a host of different specifications, wage curves are consistently found to be steeper at the country level compared to the regional level. This finding in itself is not novel, but has not received much attention: in their meta-study of 608 wage curve estimates Clar *et al.* (2007) note that wage curve estimates using national data on average find wages to be more cyclical compared to those using regional data. Also recent studies such as Koenig *et al.* (2020) report more rigid wages using regional data.[1] To the best of my knowledge, this is the first paper to

[1]Interestingly, Beraja *et al.* (2019) report more rigid wages at the country level compared to the state level as the motivating observation for their paper. But this result comes with some caveats since both the

further investigate the effect of aggregation on wage rigidity estimation, and to argue that aggregation in this context biases results.

Starting with the seminal work of Theil (1954), several authors[2] have emphasised that heterogeneity in slope parameters, or a shared but nonlinear relationship, implies that the slope parameters cannot be inferred from aggregate data without additional information or assumptions on the distribution of changes at the micro-level. This type of *distributional aggregation bias* is well known in the context of, for example, demand estimation, but has received less attention in macro-econometric analyses. Lewbel (1992) considers log-linear relationships without parameter heterogeneity, and shows that estimation using aggregate data is biased unless changes of the explanatory variables at the micro-level are proportional (mean-scaled). van Garderen *et al.* (2000) and Albuquerque (2003) consider log-linear aggregation with slope heterogeneity. I show that the bias described in this literature also may affect regressions using micro-level data with explanatory variables considered at a more aggregated level, such as Hagedorn and Manovskii (2013) or Gertler *et al.* (2020).

Pesaran and Smith (1995), Pesaran *et al.* (1999) and numerous more recent contributions consider aggregation in the context of dynamic heterogeneity. Dynamic heterogeneity leads to residual autocorrelation in the aggregate series, and bias in the presence of lagged dependent variables. Bias through heterogeneous dynamics has received a lot of attention in the macro-econometric literature, but mostly on the question whether pooled estimation can be used on such data, rather than on the consequences of aggregating data (see for example Canova, 2011, chapter 8). Pesaran *et al.* (1999) propose using the pooled mean group and mean group estimators on the micro-level data as strategies to avoid bias from pooling in the context of dynamic heterogeneity. An appendix verifies the robustness of the results presented in this paper to pooling under dynamic heterogeneity.

Although the parameters estimated using aggregate data correctly summarize the *observed* relationship between the macro-aggregates, there will be atypical changes in the micro-variables, for example policy-induced, leading to changes in macro-aggregates that deviate from the estimated relationship. Over the period considered, changes in European country-level unemployment are driven mainly by underlying changes in regions with low unemployment rates. Because the estimated wage curve in levels is downward sloping and convex, variation in unemployment in the regions with low unemployment rates causes large changes in wage pressure, both locally and at the country level. Using aggregate data then leads to overestimating the slope of the wage

---

data and methods used differ between their country and state level analysis. First, the reported weighted average of state level increase in nominal wages between 2007 and 2010 (their Figure 1, Panel A) deviates from the same variable at the country level (their Figure 2, Panel A), even in sign. Second, at the country level, a single observation of the ratio of changes in wages to employment is considered as a measure of wage flexibility. This ratio equals the slope of the line *between the origin* and this datapoint in ΔW-ΔE space, i.e. omitting a constant term. Their state level analysis, in contrast, is fundamentally different in using multiple observations and allowing for an intercept. Third, the country level wage changes are de-trended (which matters due to the absence of a constant term), but state-level wages are not.

[2] See for example Stoker (1986) for an overview.

curve, or overestimating wage flexibility. Policies, however, typically do not target low unemployment regions. A fiscal policy reducing unemployment in regions with an average or high unemployment rate will then lead to less aggregate wage pressure than a researcher or policy maker would have been led to believe from the country-level analysis.

Conditions for the distributional aggregation bias described by Stoker (1986) and Lewbel (1992) to occur are that (1) conditions at the regional level matter for local wage setting, and not just national variables; (2) the underlying relationship between regional wages and unemployment is non-linear; and (3) changes in regional unemployment rates are not 'mean scaled'. There is ample empirical evidence for these three conditions in the European data considered here.

First, the importance of local factors for wage determination has been attested by a vast empirical literature. The early Phillips curve literature related regional wage changes to regional unemployment. Lipsey (1960) considered nonlinearity, distributional effects and their link with regional and national Phillips curve and NAIRU estimates. The limited response of migration and labour mobility to labour demand shocks in Europe has been well documented (see for example Beyer and Smets, 2015; Arpaia *et al.*, 2016; Basso *et al.*, 2019) and contributes to the long lasting effects of local shocks. Other papers considering the relation between regional and country level variables in the context of wage setting are for example Roberts (1997), Jimeno and Bentolila (1998) and Campbell (2008). Kosfeld and Dreger (2018) model the effect of unemployment in neighboring regions on wage formation as spatial autocorrelation. In this paper, the formal derivation of the aggregation bias and empirical specifications model spatial autocorrelation with both regional and national unemployment rates affecting local wages and wage inflation.

Second, the nonlinearity of the relationship between unemployment rates and wages is well attested. It received attention in the early literature on the Phillips curve where log-linear and more convex relationships between wage inflation and unemployment rates were considered (see for example Lipsey, 1960). Also the wage curve literature (Blanchflower and Oswald, 1994; Card, 1994) has typically considered a log-linear rather than linear relationship between the unemployment rate and wages. Also in the data on European regions used here, I verify that the relationship between regional wages and unemployment rates is convex and approximately log-linear.

Third, regional unemployment differences in the EU are large, both between and within countries. I show that changes over time in the distribution of unemployment rates within countries are not mean-scaled. Increases in country level unemployment rates are on average accompanied by lower dispersion of unemployment rates. The ratio of lower quantiles of regional unemployment rates to the country level unemployment rate are on average positively correlated with the country level unemployment rate. This is only possible if the distribution of regional unemployment rates is compressing and expanding mostly on the left, i.e. if changes in aggregate unemployment rates are mainly driven by regions with relatively low levels of unemployment. Given the downward slope and convexity of the wage curve, this type of deviation from mean scaling leads to overestimating wage flexibility when aggregating.

Even if wage curves are log-linear and identical in all regions, asymmetric (non-mean-scaled) regional shocks may cause the aggregate wage curve to be highly non-loglinear and, in theory, have a vertical asymptote at some positive level of the country level unemployment rate. Aggregation may thus obfuscate the effect of the level of wages in restoring labour market equilibrium, and mislead econometricians into preferring a vertical Phillips curve as a representation of the long run labour market equilibrium, rather than a wage curve. Estimation of natural rates of unemployment or the NAIRU would then be based on misspecified models, and biased upward. Aggregation bias may therefore also explain the finding that the US is characterised by a Phillips curve and European countries by a wage curve, simply because the USA is a larger country and aggregation bias using country level data could be expected to be larger; or, for studies using regional data, because the analysis in the US is typically performed at a higher level of spatial aggregation (US states versus UK regions in Blanchard and Katz, 1997).

The implications of overestimating wage pressure or NAIRUs are significant. Overblown fears of inflation may have held back governments worldwide in using fiscal stimulus to fight crises in recent decades. This is quite clear in the case of the Eurozone, where the European Commission judges whether a EU member state has an excessive fiscal deficit using a set of rules that is explicitly based on econometric NAIRU estimates. Central banks may have wondered about the lack of aggregate (wage) inflation given record-low levels of interest rates and low levels of unemployment, while the level of unemployment at which inflation would pick up is lower than country-level analyses would suggest.

The remainder of this paper is organised as follows. Section 2 describes the European regional dataset. Section 3 shows that basic wage curve elasticity estimates are significantly lower when using regional data, compared to when estimating using the same data at the country level. That the estimated wage rigidity changes significantly with the level of aggregation of the data is a is a key result of this paper in itself. Section 4 derives an analytical expression for the aggregation bias under the assumption of log-linearity, showing how the bias depends on the behaviour of the underlying regional distribution of unemployment rates. It is verified for the European data that the relationship between unemployment rates and wages is approximately log-linear, and that the changes in the distribution of regional unemployment rates matches the observed upward bias in the slope estimate of the wage curve. Section 5 estimates wage curves with rich temporal dynamics and spatial auto-correlation, finding the same upward bias. Section 6 considers the effect of aggregation on NAIRU estimation and finds the same upward bias. Section 7 concludes. As a robustness check, an appendix presents the results of analyses controlling for dynamic parameter heterogeneity using mean-group and pooled mean group estimation, and finds the same results hold using these methods.

2 Data

The data used in the empirical analysis is of annual frequency, at the NUTS2 level of regional disaggregation, and freely available online from the Eurostat website (or from the author on request). The sample consists of 246 regions in 18 countries. For most countries the available data runs from 2000 to 2016 or 2017, resulting in 4138 region-year observations with an average of 16.82 yearly observations per region. Table 1 gives an overview. The NUTS2 regions vary greatly in size, and therefore the relative size of the region in the national aggregate hours worked is used as weights in the regressions.[3]

When regional regressions consider country-level explanatory variables as spatial lags, these are calculated excluding the region under consideration and therefore refer to 'the rest of the country'. Larger regions carry a larger weight in this country-level variable, which is not usually the case in spatial econometric analysis. The advantage of this approach is that the sum of the elasticities on the own-region and country level unemployment rates then reflects the effect on regional wages of a homogeneous increase of the unemployment rates of all regions of a country, without double-counting. The sum of these elasticities can be compared to the single elasticity obtained when using aggregate national data.

I take great care to ensure that the data in the country level analyses perfectly matches the data in the regional analyses. Only years in which all regions in a country have data are considered, and the country-level data is aggregated using these strictly balanced regional series. Small countries that consist of only one NUTS2 region cannot be considered[4] and these are therefore excluded from the analysis.

The variables used are the following[5]:

- w: nominal hourly cost of employees. Calculated as the total compensation of employees divided by the total hours supplied by employees, on the region or country level
- $gvap$: gross value added deflator, on the country level
- $rw = w/gvap$: real hourly wage cost
- $prod$: real value added per hour supplied by employees, on the region or country level
- u: unemployment rate, on the region or country level

[3] As a robustness check, I considered the higher-level NUTS-1 region in all NUTS-1 regions which contain a NUTS-2 region with less than 150,000 employees in any year, resulting in a sample of 216 regions regions; I also considered a cutoff of 500,000 employees resulting in 146 regions; and also repeated the entire analysis on the level of 86 NUTS-1 regions. In these analysis the results are qualitatively similar, with changes that are expected from using larger regions: e.g. larger own-region effects and smaller spatial lagged effects for higher levels of aggregation, and a smaller difference between national and regional estimates

[4] This excluded Lithuania, Latvia, Estland, Luxemburg, Malta, Cyprus and Slovenia from the analysis. For Croatia, the youngest EU member state, there are no time series available at the regional level. For Poland the most recent years were dropped due to clear coding errors in the data.

[5] The eurostat datasets used are *nama_10r_2coe* for the compensation of employees; *nama_10r_2emhrw* for total hours supplied by employees; *nama_10r_3gva* for real value added; *nama_10_a10* for the gross value added deflator; and lastly *lfst_r_lfu3rt* and *lfst_r_lfu3pers* for the number of unemployed, the size of the labour force and the unemployment rate.

Table 1: *The countries, the first and last year contained in the sample, the number of regions, and some summary statistics. All regional series within a country are strictly balanced. The reported smallest and largest unemployment rate and nominal wage are over all years and regions.*

	min(year)	max(year)	#regions	min(urate)	max(urate)	min(wage)	max(wage)
AT	2000	2016	9	2	11.3	12.2	28.3
BE	2000	2016	11	1.9	19.2	15.3	37.5
BG	2000	2017	6	2.9	24.6	0.7	5.4
CZ	1999	2016	8	1.9	15.2	2.2	11.3
DE	2000	2017	37	2	22.4	13.5	32.4
DK	2000	2017	5	3.2	8.2	20.1	40.4
EL	2000	2016	13	4.7	31.6	3.5	10.8
ES	2000	2017	17	4.1	36.2	8.1	20.2
FR	2003	2015	22	5.2	15	16.6	36.4
HU	2000	2015	7	3.7	16.4	2.1	7.3
IT	2000	2016	21	1.8	27.3	8.5	18.4
NL	2000	2016	12	1.2	11	14.7	29.7
PL	2000	2012	16	5.5	27.3	1.6	6.3
PT	1999	2016	5	1.9	18.5	4.8	12.3
RO	2000	2016	8	3	10.8	0.5	7.9
SE	2000	2016	8	3.2	10.3	15.3	32.8
SK	2000	2016	4	3.4	25	2.2	11.6
UK	2000	2017	37	1.8	13	12.9	35.8

A Harris-Tzavalis panel unit root test[6] does not reject the H_0 of unit roots in the series for any of these variables (even for the comparatively stable unemployment rate the p-value is 0.85). The same test strongly rejects the presence of a unit root in the variables in first differences. All variables are therefore assumed to be I(1).

Nominal wages are quite erratic compared to unemployment rates. A simple regression of wages on unemployment rates may lead to spurious inference although such regressions are frequently used on micro-data with a limited number of yearly observations. With real productivity defined as $prod = \mathrm{Y}/\mathrm{H}$ with Y aggregate real value added and H aggregate hours worked, the wage share in aggregate income is $wsh \equiv \frac{w \times \mathrm{H}}{gvap \times \mathrm{Y}} = \frac{w}{gvap \times prod}$. A stationary wage share in national income implies that nominal wages are co-integrated and homogeneous in prices and productivity. I therefore mostly consider prices and productivity as explanatory variables alongside unemployment to explain wages, and test for cointegration. The alternative is to use the log of the wage share in value added as the dependent variable which imposes homogeneity of wages in prices and productivity.

- $\ln(wsh) \equiv \ln(w) - \ln(prod) - \ln(gvap)$.

Using the unemployment rate to explain changes in the wage share amounts to a model

[6]This panel unit root test is appropriate because its asymptotic results are derived assuming T fixed, contrary to most panel unit root tests. This matches well with this dataset which has only up to 18 yearly observations. The HT test requires strongly balanced series, and therefore only the residuals corresponding to the 15 countries and 185 regions which have 17 yearly observations are used for the calculation of this test.

where labour market tightness influences the division of national income between capital and labour. This could happen when actors in a collective bargaining process take into account the unemployment situation, or when the unemployment affects the outside option of a labour union.

3 Exploring aggregation bias in wage rigidity estimation

3.1 The dynamic wage curve

I estimate wage rigidity using a dynamic wage curve based on Nymoen and Rødseth (2003). Write w_{rt} for the nominal hourly cost of employees in region r (belonging to country c) and year t, $prod_{rt}$ for the real output per hour worked and $gvap_{rt}$ for the value added deflator. Different versions of the following error-correction model will be estimated:

$$\begin{aligned}\Delta\ln(w_{rt}) = & \sum_{i=1}^{j} \gamma_i \Delta\ln(w_{r,t-i}) \\ & + \sum_{k=0}^{l} \left[\beta_{0k}\Delta\ln(prod_{r,t-k}) + \beta_{1k}\Delta\ln(prod_{c,t-k}) + \beta_{2k}\Delta\ln(gvap_{c,t-k})\right] \\ & + \alpha_{0r} + \alpha_1\ln(w_{r,t-1}) + \alpha_2\ln(prod_{r,t-1}) + \alpha_3\ln(prod_{c,t-1}) \\ & + \alpha_4\ln(gvap_{c,t-1}) + \alpha_5\ln(u_{r,t-1}) + \alpha_6\ln(u_{c,t-1}) + \nu_{rt}. \end{aligned} \tag{1}$$

This is a general framework that embeds both the case of a wage curve in the tradition of Blanchflower and Oswald (1994), e.g. a relationship between the level of the unemployment rate and the level of wages, and a wage Phillips curve which posits a relationship between the level of the unemployment rate and wage growth. Parameters are assumed to be shared between regions, apart from the region-specific level effects α_{0r}. Due to the limited number of annual observations per region at most $l = 1$ is considered, and often $\gamma = 0$ and or other constraints are imposed on the parameters. A first constraint considered is one of dynamic homogeneity where $\beta_{0k} + \beta_{1k} = \beta_{2k} = 1$, under which changes in prices and productivity fully translate into nominal wage changes.

Of key interest is the long-run equilibrium, which is found by setting all terms in first differences to 0 (or a constant c) in equation (1). If $\alpha_1 \neq 0$ the following log-linear relationship between wages and the unemployment rate, e.g. a wage curve, is obtained:

$$\begin{aligned}\ln(w_{rt}) = & -\frac{\alpha_{0r}}{\alpha_1} - \frac{\alpha_2}{\alpha_1}\ln(prod_{rt}) - \frac{\alpha_3}{\alpha_1}\ln(prod_{ct}) \\ & - \frac{\alpha_4}{\alpha_1}\ln(gvap_{ct}) - \frac{\alpha_5}{\alpha_1}\ln(u_{r,t-1}) - \frac{\alpha_6}{\alpha_1}\ln(u_{c,t-1}). \end{aligned} \tag{2}$$

For given national and regional unemployment rates, and assuming regional and national productivity are moving in line, a constant long-run labour share requires the coefficient on prices and the sum of the coefficients on local and country level productivity to equal 1. Variations of this long run equilibrium relationship are frequently estimated in the empirical literature on the wage curve. Typically spatial lags and productivity are ignored. Regressions considering real wages as the dependent variable amount to constraining the coefficient on prices to 1.

Considering equation (1) with $\alpha_1 = \alpha_2 + \alpha_3 = \alpha_4 = 0$ excluding the spatial lag of unemployment ($\alpha_6 = 0$), and setting all $\beta's$ and $\gamma's$ to 0, the long-run equilibrium rather corresponds to a vertical line in w-u space, the long run vertical wage Phillips curve defined by

$$\ln(u_r^*) = -\frac{\alpha_{0r}}{\alpha_5}. \tag{3}$$

Also here one can alternatively set some parameters to 1 in equation (1) and bring variables to the left hand side, to consider the level of unemployment rate at which there are no changes in real wages, or no changes in the wage share, rather than in nominal wages. I will do so in the empirical analysis. The unemployment rate at which wages are constant can be called the non increasing wage rate of unemployment, or NIWRU. It is a basic estimate of the natural rate of unemployment in the economy. With $\gamma = 1$ acceleration in nominal wages, real wages or wage shares is considered, rather than increases. This single level of unemployment u_r^* for which wage growth (wage inflation) is constant (wages are non-accelerating) can be called the NAWRU. I will join part of the literature in referring to all these levels of unemployment as the NAIRU or NAWRU.

3.2 Aggregation bias in wage curve estimation: basic regressions

To explore the basic properties of the data and illustrate possible aggregation bias, Table 2 compares the results of some basic wage curve regressions when using the original data at the regional level with the results obtained after aggregating the data at the country level. All specifications in this paper include cross-sectional and time dummies. Regions with a larger workforce carry a greater weight in the country level analysis. To ensure an apples-to-apples comparison between the region and country level, the regressions at the regional level use the regional share in the country level total hours worked as weights. The long run wage curve elasticity is reported separately in the row 'LR-elast.'. This corresponds simply to the coefficient on the unemployment rate in logs for the most basic regressions. It is calculated as the sum of the coefficients on regional and country level unemployment for specifications including a spatial lag. It is calculated following equation (2) for the error-correction specifications. Also the estimated unweighted long run wage curve elasticity is reported in the row 'LR-elast. (unw.)'. The coefficients underlying the calculation of the unweighted elasticities are not reported to preserve space.

The regression reported in column 1 of Table 2 considers the log real hourly wage cost at the regional level as the dependent variable, with the log regional unemployment rate as the sole explanatory variable. This amounts to estimating the long-run wage curve of equation (2) without spatial lags, imposing a coefficient of 1 on prices, and ignoring productivity. The estimated long run wage elasticity is -0.126. For the unweighted regression it is -0.086. Both are close to the value of -0.1 typically found in the literature. Taking the same specification, with the same data aggregated at the country level

Table 2: *Aggregation bias: The estimated long run wage curve elasticities (LR-elast.) are larger (more negative) when estimating on aggregated data (columns 2, 4 and 6) compared to regional data (columns 1, 3, 5, 7 and 9).*

	Specification 1		Specification 2		Specification 3		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\ln(rw_{\mathrm{rt}})$	$\ln(rw_{\mathrm{ct}})$	$\ln(wsh_{\mathrm{rt}})$	$\ln(wsh_{\mathrm{ct}})$	$\Delta\ln(wsh_{\mathrm{rt}})$	$\Delta\ln(wsh_{\mathrm{rt}})$	$\Delta\ln(wsh_{\mathrm{ct}})$
$\ln(u_{\mathrm{rt}})$	-0.126^{***}						
	(-13.00)						
$\ln(u_{\mathrm{ct}})$		-0.161^{***}					
		(-6.94)					
$\ln(u_{\mathrm{r,t-1}})$			-0.0193^{***}		-0.0184^{***}	-0.00521	
			(-5.28)		(-7.88)	(-1.02)	
$\ln(u_{\mathrm{c,t-1}})$				-0.0295^{***}		-0.0164^{***}	-0.0212^{***}
				(-3.76)		(-3.12)	(-4.71)
$\ln(wsh_{\mathrm{r,t-1}})$					-0.246^{***}	-0.249^{***}	
					(-10.94)	(-11.09)	
$\ln(wsh_{\mathrm{c,t-1}})$							-0.192^{***}
							(-3.62)
Constant	-14.18^{***}	-14.09^{***}	-0.708^{***}	-0.741^{***}	-0.218^{***}	-0.229^{***}	-0.192^{***}
	(-186.23)	(-86.19)	(-40.99)	(-23.63)	(-12.19)	(-12.70)	(-4.66)
LR-elast.	-0.126	-0.161	-0.0193	-0.0295	-0.0746	-0.0865	-0.111
	(-13.00)	(-6.939)	(-5.284)	(-3.755)	(-6.869)	(-7.674)	(-3.390)
LR-elast. (unw.)	-0.0859		-0.00820		-0.0510	-0.0618	
	(-17.39)		(-3.859)		(-8.517)	(-9.404)	
N.Obs.	4138	304	3892	286	3892	3892	286
Level	region	country	region	country	region	region	country
R-sq	0.971	0.970	0.933	0.949	0.260	0.264	0.316
Q AR(1) p	0	0	0	0	0	0	0.0294
Q AR(2) p	0	0	0	0	0.0112	0.0120	0.218
HT I(1) z	-1.522	-0.129	-1.205	-0.0480	-17.93	-17.73	-6.374

Robust t-statistics in parentheses. ***: p<0.01; **: p<0.05; *: p<0.1. Cross-sectional and year dummies are included. HT-test: reject H0 of no-cointegration for z < -1.65.

(column 2), the estimated elasticity is -0.161. Weighting the regional regression by employment (hours) closes about half of the gap between the unweighted regional and national estimates. This tends to be the case in many specifications which are considered below. This very basic specification is used frequently in the literature, especially in studies based on detailed micro-data with short time series. The bottom of the table shows a set of diagnostic tests which should make us sceptical of the results, however. The z-value of the Harris-Tzavalis (HT) test shows a lack of evidence[7] against the H0 of unit roots in the regression residuals, suggesting that the observed relationship may be spurious. Controlling for productivity is an obvious possible way to correct this. The p-values of the Ljung-Box portmanteau (Q) tests reported in the table strongly reject the absence of first and second order autocorrelation in the residuals. This suggests that a more elaborate dynamic specification may be required.

Prices and productivity are key stochastic drivers of wages. The specification in columns (3) and (4) replaces wages as the dependent variable by the labour share in value added. This amounts to estimating equation (2) without spatial lags, and with the restriction of long run homogeneity in prices and productivity imposed. The

[7]The HT tests rejects the H0 of presence of I(1) in the residuals which would imply no cointegration for z < -1.65, at the one-sided 5 percent significance level; this critical value is reported at the bottom of the tables.

unemployment rate is lagged by one year to allow for slower adjustment. Also here estimation on aggregated data leads to a higher elasticity. Still the HT test cannot reject unit roots in the residuals for the country level analysis. The absence of AR(1) and AR(2) in the residuals is strongly rejected.

Columns (5) to (7) no longer start from the long run equilibrium relation (2), but rather consider a simple version of the dynamic wage curve of equation (1). The diagnostic test statistics for this dynamic specification are more promising. Although the Q test rejects the absence of AR(1) in the residuals, it does so less strongly, with the z-value (not reported) decreasing from about 21 to 5 in absolute value for the region-level regression in column (5). The HT test now strongly rejects unit roots in the residuals. This suggests a cointegration relationship exists between wages, prices and productivity. The long run wage curve elasticity now is calculated as the ratio of the coefficients on the lagged log unemployment rate and the lagged log wage share as in equation (2). Column (6) considers the wage share at the regional level, while controlling for regional unemployment and country level unemployment. Again, the sum of the long run elasticities on the regional and country level unemployment rate is smaller than what is obtained using the same data aggregated at the national level (column 7).

In conclusion, all these relatively simple wage curve estimations with specifications that are commonly found in the literature find a lower elasticity when estimation is performed at the regional level, compared to analyses using the same data aggregated at the country level. The diagnostic tests suggest that an error correction model is preferred, and even more elaborate dynamics are needed. Section 5 considers such specifications. In the next section, distributional aggregation bias is first considered as a possible explanation of the observed difference between the regional and country level results, and whether the properties of the European regional data support this explanation.

4 Distributional aggregation bias in wage curve estimation

4.1 Formal derivation

This section formally derives the conditions under which distributional aggregation bias can explain the observed difference between the regional and country level wage curve elasticity estimation. Consider estimating the long run equilibrium relationship of equation (2) without spatial spillovers in productivity $\alpha_3 = 0$ while imposing long run homogeneity in productivity and prices $-\alpha_1 = \alpha_2 = \alpha_4$ such that for the wage share $\ln(wsh_{rt}) \equiv \ln(w_{rt}) - \ln(prod_{rt}) - \ln(gvap_{ct})$ it holds that

$$\ln(wsh_{rt}) = b_{0r} + b\ln(u_{rt}) + b_s\ln(u_{ct}) + \nu_{rt}, \tag{4}$$

To formally derive the bias that may occur when estimating equation (4) using national rather than regional data I follow Lewbel (1992) and van Garderen *et al.* (2000). Drop time indices for convenience. Define b_{0c} as the average over the region-specific deterministic constant terms b_{0r}. Exponentiating, multiplying by the regional share in aggregate hours worked ω_r, summing over regions and taking the expected value results in

$$E\left[\sum_r \omega_r wsh_r\right] = E\left[\sum_r \omega_r \exp\left(b_{0r}+b_{0c}-b_{0c}+b\ln\left(\frac{u_r}{u_c}\right)+b\ln(u_c)+b_s\ln(u_c)+\nu_r\right)\right].$$

The left hand side equals the expected value of the country level wage share. If the errors in the region-level wage equation are $\nu_r \sim N(0, \sigma^2_\nu)$ then $E[\exp(\nu_r)] = \sigma^2_\nu/2$ and

$$E[wsh_c] = \exp\left(b_{0c} + (b+b_s)\ln(u_c) + \frac{\sigma^2_\nu}{2}\right) E\left[\sum_r \omega_r \exp\left(b_{0r} - b_{0c} + b\ln\left(\frac{u_r}{u_c}\right)\right)\right].$$

An estimation equation at the country level then would be

$$wsh_c = \exp\left(b_{0c} + (b+b_s)\ln(u_c) + \frac{\sigma_\nu}{2}\right) E\left[\sum_r \omega_r \frac{e^{b_{0r}}}{e^{b_{0c}}}\left(\frac{u_r}{u_c}\right)^b\right] \exp(\epsilon_c),$$

with $E[\exp(\epsilon)] = 1$. Note that then in general $E[\epsilon] \neq 0$ (see for example van Garderen *et al.*, 2000; Silva and Tenreyro, 2006). If one is willing to assume that $\epsilon \sim N(E[\epsilon_c], \sigma^2_\epsilon)$ it holds that $E[\exp(\epsilon_c)] = \exp(E[\epsilon_c] + \sigma^2_\epsilon/2) = 1$ or $E[\epsilon_c] = -\sigma^2_\epsilon/2$. Defining $\xi_c = \epsilon_c + \sigma^2_\epsilon/2$ implies $E[\xi_c] = 0$. Define the new regional weight $\omega'_r = \omega_r e^{b_{0r}}/e^{b_{0c}}$. Moreover, exclude errors-in-variables and consider the sample equivalent $\sum_r \omega'_r(u_r/u_c)^b$ for its expected value, and the estimation equation becomes

$$\ln(wsh_c) = b_{0c} + (b+b_s)\ln(u_c) + \frac{\sigma_\nu}{2} - \frac{\sigma_\epsilon}{2} + \ln\sum_r \omega'_r\left(\frac{u_r}{u_c}\right)^b + \xi_c. \tag{5}$$

Now consider estimation of this equation including only a constant and $\ln(u_c)$ as the sole explanatory variable. Estimation of the intercept will in general be biased due to the presence of the various omitted terms. The presence of innocuous heteroskedasticity implies correlation between the observed $\ln(u_c)$ and the unobserved terms σ_ν and σ_ϵ, and biased estimation of the slope coefficient $b + b_s$. Assuming that the variances σ_ν and σ_ϵ are uncorrelated over time with the aggregate unemployment term, the expected value of the coefficient on $\ln(u_c)$ when estimating equation (5) while omitting the term $\ln E[(u_r/u_c)^b]$ equals

$$E[\widehat{b+b_s}] = b + b_s + \frac{cov\left(\ln\sum \omega'_r\left(\frac{u_r}{u_c}\right)^b, \ln(u_c)\right)}{var(\ln(u_c))}. \tag{6}$$

The bias is increasing in the covariance between $\ln E[(u_r/u_c)^b]$ and $\ln(u_c)$. If over time the individual u_r move proportionally with the country level average, the covariance is 0 and the distribution of u_r is said to be *mean-scaled* (Lewbel, 1992). Aggregate data can then be used to recover the underlying regional wage curve elasticity. The next section will show, however, that European regional unemployment rates are not mean-scaled.

Note that the term $\ln E[(u_r/u_c)^b]$ is a measure of dispersion of the u_r. This can also be seen by assuming log-normally distributed u_r. Assume that regions are equal in size and have equal wage curve intercepts. For $\ln(u_r) \sim N(\mu, \sigma_c^2)$ with μ and σ_c^2 the cross-regional mean and variance of regional unemployment rates within country c within a time period, it holds that $\ln(E[u_r^b]) = b\mu + b^2\sigma_c^2/2$ and therefore

$$\ln E\left[\left(\frac{u_r}{u_c}\right)^b\right] = \ln E\left[u_r^b E[u_r]^{-b}\right] = -b\ln(E[u_r]) + \ln E[u_r^b] = \frac{\sigma_c^2}{2}b(b-1).$$

The bias expressed as the percentage difference between the expectation of the slope parameter in a country level wage curve, and the underlying parameter value $b + b_s$ which shows how much wages would change assuming a homogeneous increase in unemployment rates in all regions of the country, then equals[8]

$$\frac{E[\widehat{b+b_s}] - (b+b_s)}{b+b_s} = \frac{b-1}{2}\frac{cov(\ln(u_c), \sigma_c^2)}{var(\ln(u_c))}\frac{b}{b+b_s}. \tag{7}$$

For the EU as a whole, regional unemployment rates in any given year seem approximately log-normally distributed. Log-normal distributions appear quite naturally: if rates of change of regional unemployment rates are independent of the level of the unemployment rate, the distribution of regional unemployment rates will tend to a log-normal distribution as a result of the central limit theorem. It may be risky to put much faith in equation (7) in this application, however, given that many countries contain just a handful of NUTS2 regions and given that the expression assumed equal sizes and wage curve intercepts between regions. Equation (6), in contrast, does not impose any distributional assumption on the regional unemployment rates, and allows for differences in the size of regions and their wage curve intercepts.

Analysis such as Hagedorn and Manovskii (2013) and Gertler *et al.* (2020) which consider data at the level of individuals i but use unemployment rates at a high level of aggregation c possibly suffer from a similar bias. Even when allowing for individual specific intercepts b_{0ir} and controlling for individual characteristics x_{irt} it similarly holds that

$$\ln(w_{irt}) = b_{0ir} + b_1\ln(x_{irt}) + b\ln(u_{ct}) + b\ln\frac{u_{rt}}{u_{ct}} + \epsilon_{irt}.$$

A regression controlling for unemployment at the level c rather than r which leaves out the term $\ln(u_{rt}/u_{ct})^b$ would then suffer from aggregation bias in presence of covariance between $\ln(u_{rt}/u_{ct})^b$ and $\ln(u_{ct})$, i.e. if regional unemployment rates are not mean-scaled. This bias does not disappear when considering a larger sample or more control variables at the individual level. With regressions at the micro-level the variances σ_v and σ_ϵ do not appear in the expression, however, and heteroskedasticity would not lead to biased inference.

[8] Lewbel (1992) reports $cov(\ln(u_c), \sigma_c)$ and not $cov(\ln(u_c), \sigma_c^2)$.

4.2 Graphical illustration

Equation (6) shows that assuming a log-linear relationship between unemployment rates and wages at the regional level, correlation between $\ln \mathrm{E}[(u_r/u_c)^b]$ and $\ln u_c$ (failure of mean-scaling) causes aggregation bias. Figure 1 illustrates how this works for the case of two regions of equal size. The regional wage curves are given by $\ln(w_{rt}) = -0.1\ln(u_{rt})$.

Figure 1: *A graphical illustration of the aggregation bias. Bold: overlapping wage curves in two regions. Thin line: average (country level) values. Left panels: levels. Right panels: logs. The unemployment rate in region* R1 *is fixed at 10 percent. For values* R2 < R1 *decreasing averages go with an increase in dispersion of unemployment rates and the country level wage curve is steeper than the regional one. For values of* R2 > R1 *the opposite holds. Bottom: if the wage curve in levels is convex and concave in logs (not log-linear), the aggregate wage curve in logs will still be convex due to aggregation bias.*

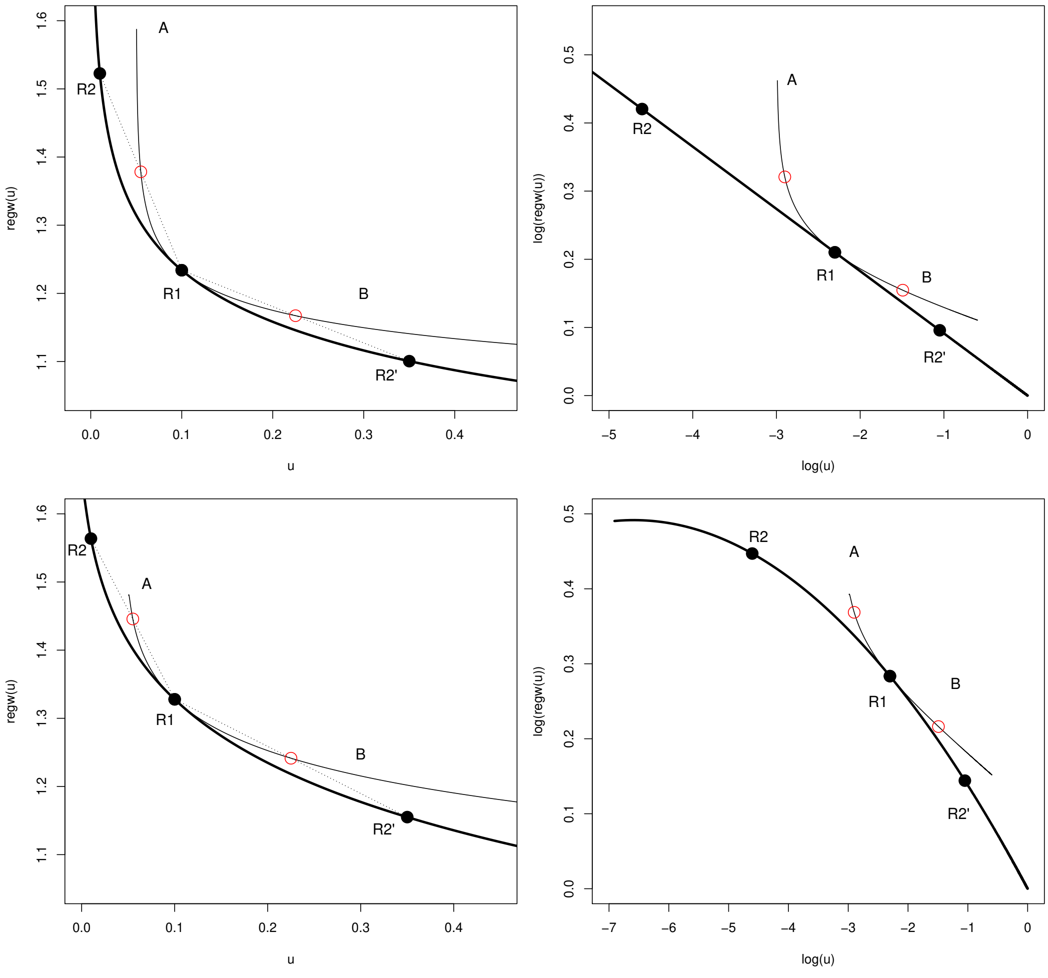

They are independent from variables outside of the region, overlap and are pictured in bold. The left panel shows the relationship in levels, and the wage curve therefore is a curve, with an asymptote at $u = 0$. The top right panel shows the log unemployment rate and log wages. The log-linear regional wage curve depicted in logs is obviously a straight line.

Assume that both regions initially have an unemployment rate of 10 percent. Consider changes in the unemployment rate in only one of the regions, R2, keeping the unemployment rate in R1 fixed. Such changes violate mean-scaling. The thinner line shows how the average wage and unemployment rates change in response to the changes in region R2. Two specific values for the unemployment in R2 are illustrated by a black dot, for unemployment rates of 0.01 and 0.35. For each, the average unemployment rate and wage levels are indicated by a red circle, which in the left panel lies at the midpoint of the line segment between the regions R1 and R2. If $R2 = R1 = 0.1$ then $\ln u_c = \ln(0.1) = -2.302$ and $\ln(\overline{(u_r/u_c)^b}) = 0$. If R2 changes to 0.01, then $\ln u_c = \ln(0.055) = -2.9$, whereas $\ln(\overline{(u_r/u_c)^b}) = \ln((0.01/0.055)^{-0.1} + (0.1/0.055)^{-0.1}) = 0.755$. As equation (6) shows, such negative covariance between $\ln E(u_r/u_c)^b$ and $\ln u_c$ leads to overestimation of the negative slope of the wage curve using aggregate data.

Consider the expression for log-normally distributed regional unemployment rates in equation (7). The term $\mathrm{cov}(\sigma^2_{ct}, \ln(u_{ct}))$ determines the sign of the bias. For the case $b < 0$ the estimated slope at the aggregate level will be steeper than the slope at the regional level if decreases in the aggregate unemployment rate $\ln(u_{ct})$ are accompanied by increases in the dispersion of regional unemployment rates as measured by σ^2_{ct}. This opposite movement in dispersion and the country average occurs at the left of point R1 in Figure 1. To the right of point R1, increases in the unemployment in R2 lead to increases in both $\ln(u_{ct})$ and σ^2_{ct}, and the bias is positive.

As shown in the top right panel, with log-linear regional wage curves, in absence of mean-scaling the relationship between the unemployment rate and wage at the country level is not log-linear and more convex. In this example the wage curve at the country level has an asymptote at an unemployment rate of 5 percent ($u = 0.05$), as the unemployment rate in R2 approaches 0. The steeper part A of the aggregate wage curve resembles a long run vertical Phillips curve near its asymptote at $u = 0.05$. Due to the distributional aggregation bias, there are conditions under which regional wage curves generate aggregate data that is observationally equivalent to a long run vertical Phillips curve.

4.3 Verifying the convexity of the wage curve and the absence of mean-scaling in European regional unemployment

This section considers whether the data on European regional unemployment rates and wages supports distributional aggregation bias as an explanation of the higher wage curve elasticities which are observed at the country level compared to the regional level. I first show that the basic relationship between regional wages and unemployment rates is convex and approximately log-linear, before showing how changes over time in the distribution of regional unemployment rates deviate from mean scaling.

Convexity

The regressions reported in Table 3 show that in Europe the relationship between wages and unemployment, in levels, is highly convex and reasonably approximated by a log-linear function, both at the regional and country level. Column (1) considers a

Table 3: *The relationship between regional wages and the unemployment rate is convex and approximately log-linear. The long run wage curve elasticity (LR-elast.), evaluated at the median level of unemployment of 7 percent, are higher at the country level.*

	(1) $\Delta\ln(wsh_{rt})$	(2) $\Delta\ln(wsh_{rt})$	(3) $\Delta\ln(wsh_{rt})$	(4) $\Delta\ln(wsh_{ct})$	(5) $\Delta\ln(wsh_{ct})$	(6) $\Delta\ln(wsh_{ct})$
$\ln(wsh_{r,t-1})$	-0.246^{***} (-11.00)	-0.246^{***} (-10.94)	-0.246^{***} (-10.97)			
$u_{r,t-1}$	-0.325^{***} (-5.36)					
$u^2_{r,t-1}$	0.553^{***} (3.21)					
$\ln(u_{r,t-1})$		-0.0184^{***} (-7.88)	-0.0286^{***} (-3.18)			
$\ln(u_{r,t-1})^2$			-0.00205 (-1.20)			
$\ln(wsh_{c,t-1})$				-0.190^{***} (-3.61)	-0.192^{***} (-3.62)	-0.192^{***} (-3.61)
$u_{c,t-1}$				-0.383^{***} (-3.32)		
$u^2_{c,t-1}$				0.709^{*} (1.87)		
$\ln(u_{c,t-1})$					-0.0212^{***} (-4.71)	-0.0280 (-1.28)
$\ln(u_{c,t-1})^2$						-0.00142 (-0.34)
Constant	-0.148^{***} (-9.74)	-0.218^{***} (-12.19)	-0.230^{***} (-11.05)	-0.109^{***} (-3.22)	-0.192^{***} (-4.66)	-0.199^{***} (-4.10)
LR-el.$\|_{u=0.07}$	-0.0704 (-5.840)	-0.0746 (-6.869)	-0.0718 (-6.636)	-0.104 (-3.086)	-0.111 (-3.390)	-0.107 (-3.448)
LR-el.$\|_{u=0.07}$ (unw.)	-0.0453 (-6.682)	-0.0510 (-8.517)	-0.0482 (-8.117)			
N.Obs.	3892	3892	3892	286	286	286
Level	region	region	region	country	country	country
AIC	-16737.0	-16738.8	-16738.9	-1376.9	-1378.9	-1377.0
BIC	-15070.1	-15078.1	-15072.0	-1237.9	-1243.6	-1238.1
AR-sq	0.206	0.206	0.207	0.213	0.217	0.214
Q AR(1) p	0	0	0	0.0275	0.0294	0.0279
Q AR(2) p	0.0137	0.0112	0.0131	0.230	0.218	0.226
HT I(1) z	-18.17	-17.93	-17.99	-6.639	-6.374	-6.429

Robust t-statistics in parentheses. ***: p<0.01; **: p<0.05; *: p<0.1. Cross-sectional and year dummies are included
HT-test: reject H0 of no-cointegration for z < -1.65.

second order polynomial in the unemployment rate. Column (2) uses the log of the unemployment rate. Column (3) uses both the log and the squared log unemployment rate. Columns (4) to (6) repeat this on the country level. The regional specification with the unemployment rate and its square (column (1)) shows that the relationship is highly convex. The wage share is estimated to change by 1.18 percent in relative terms for every 1 percentage point change in the unemployment rate at an unemployment rate of 3.7 percent (the lower 5th percentile of regional unemployment rates), and by 0.42 percent at an unemployment rate of 16 percent (the upper 95th percentile). So the regional wage curve is roughly three times as steep at the lower range of unemployment rates compared to higher levels of unemployment.

The derivations in section 4 assumed log-linearity of the wage curve at the regional level. Comparing specifications (1), (2) and (3) in Table 3 suggest that log-linearity is a reasonable approximation of the relationship between wages and unemployment in the data. The BIC prefers the log-linear form over a second order polynomial in the unemployment rate (and also over a linear and third order polynomial, not reported) and rejects the addition of a squared log term. The AIC (and the adjusted R2) only marginally prefer the addition of such a squared log term.

Noteworthy is that the relationship between unemployment and wages is more convex at the aggregate level (comparing the squared terms in columns 1 and 4). Finally, note that the main conclusion from before still holds: the estimated long run wage curve elasticity, now evaluated at the median unemployment rate of 7 percent, is about 50 percent higher when estimated using country level data.

Violation of mean-scaling

Lewbel (1992) shows that a necessary and sufficient condition for mean scaling and unbiased aggregation is that the ratio of the q'th quantile s_{qct} of the distribution of the regional unemployment rates to the aggregate unemployment rate, $rel_{qct} = \frac{s_{qct}}{u_{ct}}$ is independent from u_{ct}. Table 4 shows the coefficients on u_{ct} in a regression of rel_{qct} on u_{ct} for several choices of s_{qct}. The first row shows the results when pooling all countries while including country dummies to ensure that the reported regression coefficient is reflecting only within-country variation. There is a striking pattern: the lower quantiles

Table 4: *Regressing regional unemployment rate quantiles relative to the country level unemployment rate $rel_{qct} = \frac{s_{qct}}{u_{ct}}$, on the country level unemployment rate u_{ct}.*

	(1) rel1	(2) rel5	(3) rel10	(4) rel25	(5) rel50	(6) rel75	(7) rel90	(8) rel95	(9) rel99
u_c	0.484***	0.559***	0.416***	0.263*	−0.318**	−0.812***	−1.092***	−1.743***	−1.804***
	(3.78)	(4.47)	(3.45)	(2.30)	(−3.04)	(−5.14)	(−4.47)	(−7.29)	(−6.65)
_cons	0.549***	0.545***	0.552***	0.698***	0.862***	0.993***	1.805***	1.838***	1.841***
	(30.44)	(30.98)	(32.46)	(43.31)	(58.59)	(44.61)	(52.47)	(54.56)	(48.17)

expressed relative to the average tend to be positively correlated with the national average, and this is reversed for the higher quantiles. This violates mean scaling, and hints at the direction of the bias: The pattern found in the first row of Table 4 corresponds to section A of Figure 1, where the decrease (increase) in the average unemployment rate was driven by a decrease (increase) in the regions with the lowest unemployment rate.

The change in the lower quantiles of the regional unemployment rates is larger than the change of the average unemployment rate, leading s_{qct}/u_{ct} to move in line with u_{ct}. The higher quantiles decrease less (or in section A of the figure, not at all), leading s_{qct}/u_{ct} to move in the opposite direction of u_{ct}, leading to a negative correlation. In short, the changes in the distribution of regional unemployment rates within countries over time as summarised in Table 4 suggests that wage curve estimates at the national level should be steeper compared to the regional level.

More indications of distributional aggregation bias

As was illustrated in Figure 1, a log-linear regional wage curve is expected to become more convex after aggregation. This is even the case for a regional relationship which is concave in logs (as in the bottom row of the figure). In that sense, the fact that the quadratic terms on the unemployment rate in Table 3 was found to be more positive (in levels) or less negative (in logs) at the country level is suggestive of distributional aggregation bias (see van Garderen *et al.*, 2000, for more on aggregation of quadratic functions).

Another indication of distributional aggregation bias is obtained from adding the sample equivalent of the omitted term pertaining to the distribution of regional unemployment rates in equation (5) to the country level regressions while constraining the coefficient to 1. Equation (5) was derived for the case of a static estimation equation with a single explanatory variable, such as in columns (1) to (4) of Table 2. Doing this closes only about 18 percent of the gap between the elasticity estimates at the regional and country level. Part of this failure in fully reconciling the regional and country level results may be due to violation of some of the assumptions in the data, such as deviations of log-linearity or the fact that the expression is derived for a static equation whereas the data strongly suggests wages exhibit strong hysteresis. Nevertheless, this suggests that other mechanisms may be important in explaining part of the observed difference in wage elasticities between the regional and country level.

Lastly, there are other terms in equation (5) which are missing from our estimation equations: the variances of the wage curve at the regional and country level. Using standard residual based estimates of these variance terms, we find that controlling for these terms separately or jointly hardly changes the results, and their effects operate in offsetting directions. Whereas Silva and Tenreyro (2006) use simulations to argue that heteroskedasticity is a potentially important source of bias when log-linearising multiplicative relationships, we find that at least in the case of aggregation of the log-linear relationship between wages and unemployment, and for the data considered in this paper, heteroskedasticity is unlikely to be an important cause of the observed difference between the regional and country level results in wage curve elasticity estimates.

This section aimed to illustrate the mechanisms of distributional aggregation bias and to verify that some of the preconditions for it hold in European regional data. It was shown in turn that simple wage curve elasticities are higher at the aggregate country level compared to the regional level; that not only country level unemployment but also the regional unemployment rate matters for local wage settings (more elaborate wage curve specifications below will provide more evidence on this); that the relationship between regional and country level wages and unemployment rates is convex and approximately log-linear; and that changes in the distribution of regional unemployment rates are not mean-scaled. The deviation from mean-scaling is in line with the observation of higher wage curve elasticities at at the country level. The increase in convexity observed after aggregation is suggestive of aggregation bias. Adding the analytically derived terms related to the bias to the country-level regressions closes about 18 percent of the gap between the regional and country level analysis in the most basic static regressions.

5 Dynamics and spatial autocorrelation

The diagnostic tests for the basic regressions reported in Table 2 indicate significant residual autocorrelation. This section therefore considers specifications based on the error correction equation (1), but including more elaborate temporal dynamics and spatial lags. I continue assuming homogeneity of parameters other than the intercept across regions, such that (in absence of other complications) pooling and estimation by OLS is unbiased and efficient. All regressions at the regional level in Table 5 include the spatial lags of both productivity and the unemployment rate. Columns (1) and (2) consider only contemporaneous values for the variables in differences ($l = 0$) and exclude a lagged dependent variable ($\gamma = 0$). No restrictions are imposed on the parameters. The Q tests reject the absence of autocorrelation in the residuals. Columns (3) and (4) therefore add a lag of the differenced independent variables ($l = 1$), and of differenced wages, the dependent variable ($j = 1$). Again no parameter restrictions are imposed. The Q tests no longer reject the absence of residual autocorrelation. In this specification the effect of an increase in productivity growth in all regions (considering spatial and time lags) in the region level analysis equals (0.632+0.407+0.141+0.0984)/(1-(-0.165))=1.042; with a 95 percent confidence interval of [0.987,1.096]. Wages are close to dynamically homogeneous in productivity. For prices the elasticity in response to a change in inflation is (0.672+0.329)/(1+0.165)=0.86 ([.71,1.01]). The fact that the effect of prices is smaller than 1 and estimates are less precise compared to productivity may reflect the fact that only a national price deflator is used for lack of regional data. In the long-run equilibrium in levels, productivity and prices have estimated elasticities of 0.85 [0.69,1.01] and 1.10 [0.996,1.20].

The specifications shown in columns (5) and (6) impose dynamic and long run homogeneity in prices and productivity. The results are qualitatively not very different from the unconstrained ones in columns (3) and (4). Given their elaborate spatial and temporal lag structures, and given that they pass the specification tests, specifications (3)-(4) and (5)-(6) are the preferred dynamic wage curve estimates.

As a robustness check, the specifications in columns (7) and (8) impose strict homogeneity in prices and productivity within each period, by considering changes in the wage share rather than its individual components. The fact that the addition of a lag of the differenced wage share is needed to control for residual autocorrelation may be

Table 5: *Dynamic wage curve estimation assuming homogeneous slopes, including a spatial lag of productivity and the unemployment rate. Regressions using the regional share in aggregate hours worked as weights. LR-elast. (weigh.): the long run wage curve elasticity. The row LR-elast. (unw.) separately reports the elasticity for unweighted regressions.*

	(1) $\Delta\ln(w_{rt})$	(2) $\Delta\ln(w_{ct})$	(3) $\Delta\ln(w_{rt})$	(4) $\Delta\ln(w_{ct})$	(5) $\Delta\ln(w_{rt})$	(6) $\Delta\ln(w_{ct})$	(7) $\Delta\ln(wsh_{rt})$	(8) $\Delta\ln(wsh_{ct})$
$\Delta\ln(prod_{rt})$	0.639***		0.632***		0.628***			
	(16.98)		(16.68)		(16.55)			
$\Delta\ln(prod_{ct})$	0.394***	1.036***	0.407***	1.042***	0.377***	1.004***		
	(9.59)	(16.34)	(9.87)	(17.04)	(9.39)	(30.90)		
$\Delta\ln(gvap_{ct})$	0.831***	0.831***	0.672***	0.670***	0.783***	0.808***		
	(9.24)	(4.35)	(6.79)	(3.47)	(11.26)	(5.59)		
$\ln(w_{r,t-1})$	−0.242***		−0.218***		−0.231***			
	(−9.64)		(−8.66)		(−10.20)			
$\ln(prod_{r,t-1})$	0.173***		0.141***		0.145***			
	(5.95)		(4.32)		(4.59)			
$\ln(prod_{c,t-1})$	0.0898***	0.221***	0.0984***	0.196***	0.0854***	0.200***		
	(3.23)	(4.14)	(3.20)	(3.36)	(2.92)	(3.31)		
$\ln(gvap_{c,t-1})$	0.181***	0.146**	0.186***	0.152**	0.231***	0.200***		
	(6.13)	(2.29)	(6.21)	(2.10)	(10.20)	(3.31)		
$\ln(u_{r,t-1})$	−0.00695		−0.01000*		−0.00964*		−0.0223***	
	(−1.40)		(−1.92)		(−1.82)		(−9.04)	
$\ln(u_{c,t-1})$	−0.0162***	−0.0220***	−0.0149***	−0.0242***	−0.0140***	−0.0237***		−0.0245***
	(−2.99)	(−2.94)	(−2.67)	(−3.34)	(−2.62)	(−4.65)		(−5.13)
$\ln(w_{c,t-1})$		−0.195***		−0.172***		−0.200***		
		(−3.50)		(−2.82)		(−3.31)		
$\Delta\ln(w_{r,t-1})$			−0.165***		−0.144***			
			(−4.42)		(−3.79)			
$\Delta\ln(prod_{r,t-1})$			0.136***		0.122***			
			(3.32)		(2.97)			
$\Delta\ln(prod_{c,t-1})$			0.0381	0.172	0.0180	0.124		
			(0.99)	(1.44)	(0.48)	(1.07)		
$\Delta\ln(gvap_{c,t-1})$			0.329***	0.308**	0.362***	0.320**		
			(4.63)	(2.15)	(5.34)	(2.37)		
$\Delta\ln(w_{c,t-1})$				−0.163		−0.128		
				(−1.51)		(−1.16)		
$\Delta\ln(wsh_{r,t-1})$							−0.164***	
							(−4.25)	
$\ln(wsh_{r,t-1})$							−0.232***	
							(−9.99)	
$\Delta\ln(wsh_{c,t-1})$								−0.144
								(−1.50)
$\ln(wsh_{c,t-1})$								−0.193***
								(−3.24)
Constant	0.589***	0.593*	0.340*	0.305	−0.241***	−0.205***	−0.224***	−0.205***
	(3.37)	(1.81)	(1.81)	(0.84)	(−12.30)	(−4.60)	(−12.20)	(−4.57)
LR-elast.	−0.0959	−0.113	−0.114	−0.141	−0.102	−0.119	−0.0962	−0.127
	(−6.767)	(−2.676)	(−6.641)	(−2.438)	(−7.546)	(−2.881)	(−7.161)	(−3.064)
LR-elast. (unw.)	−0.0736		−0.0769		−0.0719		−0.0627	
	(−8.004)		(−8.581)		(−10.51)		(−8.921)	
N.Obs.	3892	286	3646	268	3646	268	3646	268
Level	region	country	region	country	region	country	region	country
R-sq	0.838	0.891	0.844	0.896			0.303	0.355
Q AR(1) p	0	0.0302	0.873	0.742	0.579	0.786	0.808	0.770
Q AR(2) p	0.0485	0.125	0.173	0.345	0.295	0.272	0.239	0.533
HT I(1) z	−15.44	−4.621	−12.33	−3.756	−13.73	−4.308	−13.65	−4.386

Robust t-statistics in parentheses. *: p<0.1; **: p<0.05; ***: p<0.01. Cross-sectional and year dummies are included. HT-test: reject H0 of no-cointegration for z < -1.65.

indicative of a specification error.

The wage curve at the country level is expected to be more convex in the aggregated data compared to the regional series (see van Garderen *et al.*, 2000, and Figure 1). This was observed to be the case in the more basic regressions reported in Table 3. It turns out that this also is (modestly) the case using the more elaborate specifications considered here. Adding a squared regional unemployment rate in the same specification (column 3) results in a coefficient of between -0.0496 and -0.057 depending on weighting and whether or not the country level unemployment rate is included. At the country level the coefficient on the squared unemployment rate is -0.0460. Repeating this exercise with the unemployment in levels rather than logs results in coefficients on the squared unemployment rate between 0.17 and 0.35 at the regional level, compared to 0.46 at the country level. While these differences are in line with the predictions, they seem small.

To summarise the findings of this section, the long-run wage curve is estimated to be steeper when using national data compared to regional data, also when considering more elaborate dynamics and spatial lags, and irrespective of whether or not homogeneity restrictions are imposed. The wage curve elasticity is estimated to be between 17 and 32 percent higher at the country level when comparing to regional regression while weighting regions by their share in national hours worked, and 50 to 100 percent without weighting.

An appendix repeats this analysis using mean group and pooled mean group estimation which are robust to pooling under dynamic heterogeneity, confirming larger wage elasticities at the country level also when using these methods.

6 NAIRU estimation

The lagged level of wages is highly significant in the error-correction based wage curve regressions, both at the regional and country level (see for example Table 5). This indicates that there is no single level of the unemployment rate at which wages are stable. Wages (or the wage share) rather stabilise at levels of unemployment and wages which are jointly described by the long-run wage curve. At least for European data and the time period under consideration, imposing $\alpha_1 = 0$ in equation (1) amounts to estimating a mis-specified model. Estimating NAIRUs imposing $\alpha_1 = 0$ without validating this assumption is widely practiced, however. In this section, I also assume $\alpha_1 = 0$ and investigate whether such NAIRU estimates also dependent on the level of aggregation of the data, as was found to be the case for the estimation of wage rigidities (the wage curve elasticity) before.

As can be seen in equation (5), and was argued in previous sections, correlation between the country level unemployment and a particular measure of dispersion of regional unemployment rates biases estimation of the wage curve elasticity in analyses using aggregate data. Such correlation implies failure of mean-scaling in the dynamics of regional unemployment rates. Bias in the estimation of the slope parameter also occurs in case of correlation between country level unemployment and the variances of the shocks to the regional and country level wage equations (heteroskedasticity). If country level unemployment is uncorrelated with these three terms, estimation of slope parameters is unbiased under the assumptions of the model.

Now consider expression (3). Both the slope parameter on unemployment and the intercept determine the NAIRU. As before, an unbiased estimate of the slope parameter on the country level unemployment requires mean-scaled changes in regional unemployment and homoskedasticity. But importantly, now also an unbiased estimate of the intercept of the wage curve is required. Under the assumptions of the model this is not possible, even assuming homoskedasticity and mean-scaled regional unemployment: As long as there is variation in regional unemployment rates and some residual variation in the wage curves at the regional and country level, estimates of the intercept that fail to control for these terms will be biased. The variance and dispersion terms would be captured in the constant term, and directly enter the country level NAIRU calculation. The NAIRU estimation based on regional data as in equation (4) does not contain these three terms and allows for unbiased NAIRU estimation.

Table 6 shows country level NAIRU estimates, comparing the result of analysis based on regional and country level data. The first two columns start from the unconstrained dynamic wage curves reported in columns (3) and (4) of Table 5. To estimate the NAIRU all the variables in levels are removed except for the regional unemployment rate and the region and country level fixed effects. Also the year dummies are removed. The dynamics are rich with a lagged dependent variable and a lag of changes in productivity and prices. Notice that also in this specification the estimated short run elasticity in productivity and prices is close to 1 (summing up the coefficients on the differenced terms for this variables and dividing by one minus the coefficient on the lagged differenced wages). The median of the region and country level estimated NAIRUs are reported near the bottom of the table. The difference between the region-level median NAIRU of 8.168 and the country level estimate of 8.632 is about 0.5 percentage points, or 6 percent in relative terms. Columns (3) and (4) impose dynamic homogeneity in prices and productivity. The difference between the region and country level estimate remains about the same. Whereas there was little proof of residual autocorrelation in specifications (3) to (6) of Table 5, removing the level variables introduces significant residual autocorrelation in the regional series. This residual autocorrelation does not readily disappear when considering additional lags.

Columns (5) and (6) in Table 6 consider the wage share as the dependent variable, imposing dynamic homogeneity within every time period. The results can be compared to the specifications in columns (7) and (8) of Table 5. The difference between the median of the NAIRUs estimated at the regional (weighted) and country level is less than 5 and 2 percent in these specifications. Adding two lags of the dependent variable is required to remove the autocorrelation in the residuals (columns 7 and 8). For this specification the difference between the national and regional estimate becomes very small.

It is not straightforward to pick a preferred specification in Table 6, given the omission of the highly significant variables in levels, given the appearance of residual autocorrelation and the multiple lags of the dependent variable that are needed to capture it. Overall, however, the difference between the regional and country level NAIRU is

Table 6: *NAIRU estimation. The regressions exclude all variables in levels (such as year dummies) except for the lagged unemployment rate and cross-sectional specific intercepts. The reported NAIRUs are the medians over the region and country specific NAIRU estimates.*

	(1) $\Delta\ln(w_{rt})$	(2) $\Delta\ln(w_{ct})$	(3) $\Delta\ln(w_{rt})$	(4) $\Delta\ln(w_{ct})$	(5) $\Delta\ln(wsh_{rt})$	(6) $\Delta\ln(wsh_{ct})$	(7) $\Delta\ln(wsh_{rt})$	(8) $\Delta\ln(wsh_{ct})$
$\Delta\ln(w_{r,t-1})$	-0.269^{***}		-0.254^{***}					
	(-6.62)		(-6.20)					
$\Delta\ln(prod_{rt})$	0.626^{***}		0.627^{***}					
	(15.79)		(16.00)					
$\Delta\ln(prod_{ct})$	0.342^{***}	0.969^{***}	0.340^{***}	0.966^{***}				
	(8.35)	(19.30)	(8.37)	(33.37)				
$\Delta\ln(gvap_{ct})$	0.552^{***}	0.551^{***}	0.741^{***}	0.758^{***}				
	(6.04)	(3.27)	(9.92)	(5.00)				
$\Delta\ln(prod_{r,t-1})$	0.211^{***}		0.200^{***}					
	(5.34)		(4.82)					
$\Delta\ln(prod_{c,t-1})$	0.0918^{**}	0.301^{***}	0.0864^{**}	0.266^{***}				
	(2.35)	(3.07)	(2.19)	(2.85)				
$\Delta\ln(gvap_{c,t-1})$	0.483^{***}	0.457^{***}	0.513^{***}	0.474^{***}				
	(6.25)	(3.26)	(6.54)	(3.21)				
$\ln(u_{r,t-1})$	-0.0274^{***}		-0.0227^{***}		-0.0237^{***}		-0.0248^{***}	
	(-11.24)		(-9.36)		(-12.15)		(-9.89)	
$\Delta\ln(w_{c,t-1})$		-0.264^{***}		-0.232^{**}				
		(-2.85)		(-2.53)				
$\ln(u_{c,t-1})$		-0.0298^{***}		-0.0234^{***}		-0.0254^{***}		-0.0261^{***}
		(-6.01)		(-4.78)		(-3.81)		(-3.34)
L.$\Delta\ln(wsh_{rt})$					-0.230^{***}		-0.245^{**}	
					(-3.41)		(-2.23)	
L2.$\Delta\ln(wsh_{rt})$							-0.0983	
							(-0.78)	
L.$\Delta\ln(wsh_{ct})$						-0.228		-0.239
						(-0.68)		(-1.44)
L2.$\Delta\ln(wsh_{ct})$								-0.0916
								(-0.72)
Constant					-2.585^{***}	-2.533^{***}	-2.573^{***}	-2.518^{***}
					(-72.83)	(-22.60)	(-74.14)	(-22.75)
NAWRU	8.168	8.632	6.676	7.143	6.883	7.203	7.129	7.257
NAWRU (unw.)	8.013		6.704		7.070		7.207	
N.Obs.	3646	268	3646	268	3646	268	3400	250
Level	region	country	region	country	region	country	region	country
R-sq	0.862	0.910			0.866	0.843	0.871	0.846
Q AR(1) p	0.0105	0.668	0.0700	0.655	0.780	0.853	0.433	0.932
Q AR(2) p	0	0.711	0	0.748	0	0.247	0.610	0.963
HT I(1) (z<-1.65)	-13.56	-3.728	-14.58	-4.211	-15.25	-4.373	-14.25	-4.120

smaller than the bias observed for the wage curve elasticity; ranging from 6 to 7 percent in the regressions allowing for rich dynamics, to less than 5 percent when imposing strict homogeneity in prices and productivity by considering wage shares.

An appendix repeats this analysis using mean group and pooled mean group estimation, confirming a moderately higher NAIRU estimate at the country level also when using these methods which are robust to pooling under dynamic heterogeneity.

6.1 The NAIRU and wage curve elasticity interpretation

The NIWRU or NAIRU corresponds to the unemployment rate at which wage growth or acceleration is 0. Being a specific level of the unemployment rate, at least its unit of measurement is easily interpretable. In comparison, the dimensionless slope of a wage curve is a more abstract concept. To aid interpretation, some back-of-the-envelope reasoning can be used to translate the wage curve elasticity into a statement on a specific level of the unemployment rate. The absolute value of the wage curve elasticity corresponds to the level of the unemployment rate with a specific wage pressure: it is the level at which a 1 percentage point change in the unemployment rate gives rise to a 1 percent change in the wage share. E.g. a wage elasticity of -0.1 implies that at an unemployment rate of $\text{abs}(-0.1)$ or 10 percent, a percentage point decrease to an unemployment rate of 9 percent leads to an increase in the wage share by approximately 1 percent. The wage share is relatively stable and bounded, such that a change in the wage share by 1 percent is actually relatively large. It turns out that the variation in the wage share and unemployment rates is such that taking the absolute value of the elasticity also corresponds to the level of the unemployment rate at which, as a rule of thumb, a one standard deviation in the unemployment rate approximately leads to a standard deviation in the wage share (or 78 percent of a standard deviation to be precise).[9]

If we call one standard deviation in the wage share 'significant', the unemployment rate corresponding to absolute value of the wage cure elasticity may be called the 'significant wage pressure rate of unemployment', or SWPRU. In this view the differences in the wage curve elasticities which are observed depending on the level of aggregation are quite large: In the preferred specification of columns (3) and (4) of Table 5 the estimated SWPRU is 11.4 percent on the regional level versus 14.1 percent using aggregated country level data. If one prefers two standard deviations of the wage share as a definition of what constitutes significant wage pressure, these values are halved, to unemployment rates of 5.7 percent and 7 percent respectively. These unemployment rates are not steady state or 'natural' rates to which the economy would return. They are levels of the unemployment rate at which a further tightening of the labour market would lead to a specific amount of upward pressure on wages.

7 Summary and conclusion

Using European regional data, this paper showed that estimates of wage curve elasticities and the NAWRU are consistently smaller when estimated at the regional level, compared to estimation using the same data aggregated at the country level. This result holds throughout a host of different specifications, including static wage curves

[9] $\ln(wsh) = \beta \ln(u_r)$ implies $\frac{\Delta wsh}{wsh} \approx \beta \frac{\Delta u_r}{u_r}$. Filling in a one percent increase in the wage share and a one percentage point decrease in the unemployment rate shows that $0.01 = \beta \frac{-0.01}{u_r}$ holds for $u_r = -\beta = \text{abs}(\beta)$ since $\beta < 0$. The interquartile range of the wage share is [0.43-0.56] with a median value of 0.525. The median of the country level standard deviations in the wage share over time is 0.0125, or 0.0238 relative to the median. For the unemployment rate, the standard deviation over time in absolute terms is 0.0186. Bravely rounding both 0.0238 and 0.0186 to 0.02 gives $0.02 = \beta \frac{-0.02}{u_r}$, showing that for European regional data, a wage curve elasticity of for example -0.1 implies that, crudely, a standard deviation in the unemployment rate is predicted to lead to a standard deviation in the wage share, when the unemployment rate is 0.1 or 10 percent (7.8 percent without rounding). A steeper wage curve estimate would imply that this amount of wage pressure for a given decrease in unemployment is reached at a higher level of unemployment.

and error-correction models, whether controlling for spatial lags of unemployment and productivity or not, including basic or elaborate dynamics, and also when using mean-group methods that are robust to pooling under dynamic parameter heterogeneity.

A possible cause is the combination of a downward sloping and convex relationship between unemployment and wages, and the failure of mean-scaling of regional unemployment rates, i.e. the fact that over time European regional unemployment rates do not change proportionally with the country level unemployment rate. Changes in aggregate unemployment are rather mainly driven by changes in low-unemployment regions, where the wage curve is steep and wages respond more strongly to changes in local unemployment. This leads to and over-estimation of wage flexibility using aggregate data. An analytically derived estimate of the size of the bias suggests that convexity and failure of mean-scaling explains about 20 percent of the gap between the estimated wage curve elasticity at the country and regional level, for the particular case of a static wage curve with homogeneous parameters and unemployment as a single explanatory variable. The observed increase in convexity of the wage-unemployment relationship at the aggregate level is also suggestive of this type of aggregation bias being at work.

The theoretical and empirical results also point to the problematic nature of the NAIRU. Empirically, lagged wages are found to be highly significant in regressions of wage growth including the unemployment rate as an explanatory variable. This points to a relationship between wages and the unemployment rate for which wages are stable, rather than a specific level of unemployment. Dropping the level of wages introduces residual autocorrelation which proves hard to control for, suggesting specification error. Whereas mean-scaled dynamics of regional unemployment rates allow for unbiased estimate of the wage curve elasticity, this is not the case for NAIRU estimation due to its dependence on the intercept of the wage equation. Due to the aggregation of nonlinear regional relationships, the country-level wage equation intercept contains terms related to the variances of the residuals of the regional and aggregate wage equations and a term related to the dispersion of regional unemployment rates. If these terms are not controlled for, estimation of the intercept will be biased, and therefore the NAIRU estimate will be biased. In spite of these problems, in the European data the country level estimates of the NAIRU are only moderately higher than the NAIRU estimates based on regional data. As a simple rule-of-thumb alternative to the NAIRU that is based on the econometrically preferred dynamic wage curve estimation, I suggest using the absolute value of the long-run wage curve elasticity as measure of the unemployment rate at which wage pressure is significant.

To base policies only on aggregate data which ignore the underlying regional structure in the relationship between wages and unemployment rates is hazardous. Although analysis at the national level would capture the observed relationship at the aggregate level, there will be changes, policy induced or otherwise, that deviate from the estimated relationship. Aggregate wage rigidity estimates are driven by the dynamics of low-unemployment regions, where wages respond more strongly to changing unemployment rates. Analysis at the aggregate level would overestimate the unemployment rate below which wage pressure builds up, whether estimated through a wage curve elasticity or NAIRU. For example fiscal policy affecting all regions or regions with average or high unemployment rates would lead to less wage pressure than predicted by country level analysis.

References

ALBUQUERQUE, P. H. (2003). A practical log-linear aggregation method with examples: heterogeneous income growth in the USA. *Journal of Applied Econometrics*, **18** (6), 665–678.

ANDREWS, I., GENTZKOW, M. and SHAPIRO, J. M. (2017). Measuring the Sensitivity of Parameter Estimates to Estimation Moments*. *The Quarterly Journal of Economics*, **132** (4), 1553–1592.

ARPAIA, A., KISS, A., PALVOLGYI, B. and TURRINI, A. (2016). Labour mobility and labour market adjustment in the EU. *IZA Journal of Migration*, **5** (1), 21.

BASSO, G., D'AMURI, F. and PERI, G. (2019). Immigrants, Labor Market Dynamics and Adjustment to Shocks in the Euro Area. *IMF Economic Review*, **67** (3), 528–572.

BERAJA, M., HURST, E. and OSPINA, J. (2019). The Aggregate Implications of Regional Business Cycles. *Econometrica*, **87** (6), 1789–1833.

BEYER, R. C. M. and SMETS, F. (2015). Labour market adjustments and migration in Europe and the United States: how different? *Economic Policy*, **30** (84), 643–682.

BLANCHARD, O. and KATZ, L. F. (1997). What We Know and Do Not Know About the Natural Rate of Unemployment. **11** (1), 33.

BLANCHFLOWER, D. G. and OSWALD, A. J. (1994). *The wage curve*. MIT press.

CAMPBELL, C. M. (2008). An efficiency wage approach to reconciling the wage curve and the Phillips curve. *Labour Economics*, **15** (6), 1388–1415.

CANOVA, F. (2011). *Methods for applied macroeconomic research*. Princeton university press.

CARD, D. (1994). *The wage curve*. Cambridge, Mass: MIT Press.

CHETTY, R. (2009). Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. *Annual Review of Economics*, **1** (1), 451–488, _eprint: https://doi.org/10.1146/annurev.economics.050708.142910.

CHUDIK, A. and PESARAN, M. H. (2015). Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. *Journal of Econometrics*, **188** (2), 393–420.

CLAR, M., DREGER, C. and RAMOS, R. (2007). Wage Flexibility and Labour Market Institutions: A Meta-Analysis. *Kyklos*, **60** (2), 145–163.

DITZEN, J. (2018). Estimating Dynamic Common-Correlated Effects in Stata. *The Stata Journal: Promoting communications on statistics and Stata*, **18** (3), 585–617.

GERTLER, M., HUCKFELDT, C. and TRIGARI, A. (2020). Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires. *The Review of Economic Studies*, **87** (4), 1876–1914.

HAGEDORN, M. and MANOVSKII, I. (2013). Job Selection and Wages over the Business Cycle. *American Economic Review*, **103** (2), 771–803.

HOLDEN, S. and NYMOEN, R. (2002). Measuring Structural Unemployment: NAWRU Estimates in the Nordic Countries. *Scandinavian Journal of Economics*, **104** (1), 87–104.

JIMENO, J. F. and BENTOLILA, S. (1998). Regional unemployment persistence (Spain, 1976–1994). *Labour Economics*, **5** (1), 25–51.

KOENIG, F., MANNING, A. and PETRONGOLO, B. (2020). Reservation wages and the wage flexibility puzzle. *mimeo*.

KOSFELD, R. and DREGER, C. (2018). Local and spatial cointegration in the wage curve – a spatial panel analysis for german regions. *Review of Regional Research*, **38** (1), 53–75.

LEWBEL, A. (1992). Aggregation with Log-Linear Models. *The Review of Economic Studies*, **59** (3), 635.

LIPSEY, R. G. (1960). The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1862-1957: A Further Analysis. *Economica*, **27** (105).

NAKAMURA, E. and STEINSSON, J. (2018). Identification in Macroeconomics. *Journal of Economic Perspectives*, **32** (3), 59–86.

NYMOEN, R. and RØDSETH, A. (2003). Explaining unemployment: Some lessons from Nordic wage formation. *Labour Economics*, **10** (1), 1–29.

PESARAN, M. and SMITH, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, **68** (1), 79–113.

PESARAN, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, **74** (4), 967–1012.

—, SHIN, Y. and SMITH, R. P. (1999). Pooled Mean Group Estimation of Dynamic Heterogeneous Panels. *Journal of the American Statistical Association*, **94** (446), 621–634.

ROBERTS, J. M. (1997). The wage curve and the Phillips curve. In *Finance and Economics Discussion Paper Series # 57, Federal Reserve Board of Governors*.

SILVA, J. S. and TENREYRO, S. (2006). The log of gravity. *The Review of Economics and statistics*, **88** (4), 641–658.

STOKER, T. M. (1986). Simple Tests of Distributional Effects on Macroeconomic Equations. *Journal of Political Economy*, **94** (4), 763–795.

THEIL, H. (1954). *Linear Aggregation of Economic Relations*. North-Holland.

VAN GARDEREN, K. J., LEE, K. and PESARAN, M. (2000). Cross-sectional aggregation of non-linear models. *Journal of Econometrics*, **95** (2), 285–331.

Appendix

A Robustness: dynamic heterogeneity

As argued by Pesaran and Smith (1995) and Pesaran *et al.* (1999) in the context of linear models, ignoring heterogeneity in the slope parameters between micro-units can induce residual autocorrelation when pooling or aggregating. This leads to bias in the presence of lagged dependent variables even for $T \to \infty$. Two alternatives to pooling are suggested: mean-group estimation and pooled mean group estimation. These methods start from the individual micro-series and cannot be used with aggregated data. van Garderen *et al.* (2000) show that the derivation of aggregation bias under parameter heterogeneity and nonlinearity is quite complicated. Unbiased estimation using aggregate data requires introducing additional higher order terms, even under restrictive assumption of mean-scaling which is clearly violated in the data. I will therefore not attempt to repair the aggregate regressions or try to uncover how much of the difference between the region and country level analysis is due to dynamic aggregation bias.

An important reason to consider these methods, however, is to make sure that the pooling between regions and countries which was used throughout is not the underlying cause of the observed difference in the country and region level analysis. The expected bias when pooling under dynamic heterogeneity is for the coefficients to tend to 0. A possible explanation for the lower slopes in the regional data could therefore be that the bias due to dynamic heterogeneity is larger for the pooled regional regressions. I show in this section that this is not the case: also using pooled mean group and mean group estimation the wage curve elasticities and the NAIRU are higher when using country level data compared to regional data.

The mean group estimator considers each underlying time series separately and averages over (functions of) parameter estimates. Given the limited length of the time series, it is likely that these estimates will be inefficient and quite noisy. A particular concern is the fact that both the wage curve elasticity in the error-correction specification and the NAIRU require taking ratio's, which leads to erratic estimates (see also Holden and Nymoen, 2002). Only parsimonious specifications are therefore considered to preserve sufficient degrees of freedom, and use the median as a more robust centrality estimate. I also consider the more efficient pooled mean group estimator of Pesaran *et al.* (1999), where only short run parameters are allowed to differ between the micro-units while some of the slope parameters for the long-run effects are assumed to be shared.

Given the lower efficiency of these methods and the fact that the main conclusions are not altered, the results of this section should be considered a robustness check, excluding pooling under dynamic parameter heterogeneity as an explanation of the observed difference between region and country level estimation of wage pressure in the economy.

A.1 Dynamic heterogeneity and the wage curve

Table 7 shows the result of estimating a simplified dynamic wage curve using the pooled mean group estimator and the mean group estimator.[10] To preserve degrees of freedom I exclude lags of the independent variable $l = 0$, assume contemporaneous dynamic homogeneity in differences and levels such that the differenced price and productivity levels can be brought to the left hand side to consider changes in the labour share of income. Using the pooled mean group estimator and constraining the coefficient on the

Table 7: *Dynamic wage curve: pooled mean group and mean group estimation. The table reports the unweighted regression results. The weighted long-run wage curve elasticity (LR-elast (weigh.)) is calculated by duplicating observations in proportion to their share in the national hours worked and therefore a standard error is omitted.*

	(1) $\Delta\ln(wsh_{rt})$	(2) $\Delta\ln(wsh_{rt})$	(3) $\Delta\ln(wsh_{ct})$	(4) $\Delta\ln(wsh_{rt})$	(5) $\Delta\ln(wsh_{rt})$	(6) $\Delta\ln(wsh_{ct})$	(7) $\Delta^2\ln(wsh_{rt})$	(8) $\Delta^2\ln(wsh_{ct})$
$\Delta\ln(wsh_{r,t-1})$	-0.00841	-0.000855		-0.0346^{**}	-0.0211			
	(-0.53)	(-0.05)		(-2.22)	(-1.37)			
$\ln(wsh_{r,t-1})$	-0.339^{***}	-0.339^{***}		-0.338^{***}	-0.309^{***}			
	(-23.01)	(-22.75)		(-10.10)	(-11.90)			
$\ln(u_{c,t-1})$	-0.0365		-0.0864	-0.0315^{*}		-0.127^{***}		
	(-0.22)		(-1.27)	(-1.83)		(-6.31)		
$\ln(u_{r,t-1})$	-0.0300	-0.0607^{*}		-0.0308^{*}	-0.0546^{***}			
	(-0.23)	(-1.94)		(-1.78)	(-8.13)			
$\Delta\ln(wsh_{c,t-1})$			-0.000313			-0.0303		
			(-0.01)			(-0.47)		
$\ln(wsh_{c,t-1})$			-0.280^{***}			-0.217^{**}		
			(-6.41)			(-2.21)		
$\Delta\ln(wsh_{r,t-1})$							-1.006^{***}	
							(-56.16)	
$\Delta\ln(u_{r,t-1})$							-0.00479	
							(-0.89)	
$\Delta\ln(u_{c,t-1})$							-0.0346^{***}	-0.0498^{***}
							(-5.25)	(-4.37)
$\Delta\ln(wsh_{c,t-1})$								-0.930^{***}
								(-16.38)
LR-elast. (weigh.)	-0.0690	-0.0630		-0.0760	-0.0640		-0.0390	
LR-elast. (unw.)	-0.0660	-0.0610	-0.0860	-0.0620	-0.0550	-0.127	-0.0390	-0.0500
	(-0.863)	(-1.941)	(-1.265)	(-7.631)	(-8.133)	(-6.315)	(-8.988)	(-4.373)
pooled:	llogurc llogur	llogur	llogurc	llogshare	llogshare	llogsharec		
N.Obs.	3646	3646	268	3646	3646	268	3646	268
Level	region	region	country	region	region	country	region	country
Q AR(1) p	0.529	0.337	0.748	0.0680	0.378	0.781	0.301	0.837
Q AR(2) p	0	0	0.270	0	0	0.140	0.230	0.940
HT I(1) z	-16.18	-16.31	-7.034	-17.19	-17.83	-7.358	-16.52	-7.094

Robust t-statistics in parentheses. *: p<0.1; **: p<0.05; ***: p<0.01. Cross-sectional and year dummies are included
HT-test: reject H0 of no-cointegration for z < -1.65.

unemployment rate to be shared among regions or countries gives a long run wage curve elasticity of -0.086 on the country level, compared to elasticities ranging from -0.061 to -.069 at the regional level, depending on whether regional weights are used and a spatial lag of unemployment is included. It is important to note that this specification allows for an idiosyncratic long-run wage curve elasticity through variation in the coefficient on the

[10]The xtdcce2 command from Ditzen (2018) in Stata is used to perform the analysis.

lagged wage share between the micro-units (see equation (2)). Since the insignificance of the unemployment rates in the first column is likely due to co-linearity, column (2) repeats the analysis at the regional level excluding the spatial lag of the unemployment rate, leading to similar estimates.

Columns (4) to (6) use pooled mean group estimation but rather restrict the coefficient on the lagged wage share in levels to be identical across the micro-units, and allow the coefficient on the unemployment rate to differ between regions and countries. This results in long run wage curve elasticities between -0.055 and -0.076 on the region level versus -0.127 at the country level.

There is substantial residual autocorrelation in these specifications, however, which does not disappear when including additional lags. Therefore now consider the equation in differences in columns (7) and (8).[11] Estimation using the mean group estimator allows all parameters to vary between regions and countries. The estimated wage curve elasticities are smaller using this specification, but remain substantially higher for the country level analysis.

The difference between the estimated wage curve elasticity when using regional and aggregated data remains substantial with the pooled mean group and mean group estimators, with differences of around 25 to 70 percent when comparing country level estimates to the regional estimates using the regional share in country level hours worked as weights.

A.2 Dynamic heterogeneity and the NAIRU

As emphasised in section 6, estimation of the constant term in the macro-level wage equation suffers from different biases due to the presence of terms related to the residual variances in the wage equations at the regional and country level, and related to the dispersion of regional unemployment rates. This bias remains present even under the restrictive assumption of mean-scaled regional unemployment rates which was sufficient to allow unbiased estimation of the wage curve elasticity. This implies that various biases are likely to affect NAIRU estimation using aggregate data, since the constant term directly enters the NAIRU estimation. Also here, dynamic heterogeneity and nonlinearity would require the addition of higher order terms for unbiased estimation under the assumption of mean-scaled regional unemployment rates. I will therefore not attempt to fix the aggregate regression or uncover the relative size of these biases, but stick to excluding that the observed lower NAIRU estimate using regional data is due to the pooling of dynamic heterogeneous series.

Table 8 considers mean group estimation of the NAIRU on the regional and country level. Columns (1) and (2) consider changes in the wage share, allowing for a lag of the dependent variable. The coefficients on this lag as well as on the lagged level of the unemployment rate is allowed to differ freely between the micro-units. The estimated difference between the NAIRUs estimated using regional and country level data is

[11] The lagged dependent variable is omitted in the differenced equations. Including it results in very similar regional elasticity estimates, but increases the country level elasticity to -0.15.

Table 8: Mean group NAIRU estimation

	(1) $\Delta\ln(wsh_{rt})$	(2) $\Delta\ln(wsh_{ct})$	(3) $\Delta\ln(wsh_{rt})$	(4) $\Delta\ln(wsh_{ct})$
L.$\Delta\ln(wsh_{rt})$	-0.156^{***}		-0.338^{***}	
	(-9.71)		(-10.43)	
L2.$\Delta\ln(wsh_{rt})$			-0.344^{***}	
			(-10.27)	
L3.$\Delta\ln(wsh_{rt})$			-0.182^{***}	
			(-5.64)	
L4.$\Delta\ln(wsh_{rt})$			-0.207^{***}	
			(-7.98)	
$\ln(u_{r,t-1})$	-0.0255^{***}		-0.0483^{***}	
	(-13.06)		(-9.25)	
L.$\Delta\ln(wsh_{ct})$		-0.136^{**}		-0.258^{**}
		(-2.16)		(-2.02)
L2.$\Delta\ln(wsh_{ct})$				-0.443^{***}
				(-3.91)
L3.$\Delta\ln(wsh_{ct})$				-0.187^{*}
				(-1.77)
L4.$\Delta\ln(wsh_{ct})$				-0.202^{**}
				(-1.98)
$\ln(u_{c,t-1})$		-0.0359^{***}		-0.0481^{***}
		(-5.50)		(-2.69)
Constant	-2.671^{***}	-2.525^{***}	-2.137^{***}	-1.334
	(-35.17)	(-22.86)	(-6.54)	(-1.16)
NAWRU	7.234	7.439	6.964	7.936
NAWRU (weigh.)	7.234		6.964	
N.Obs.	3646	268	2908	214
Level	region	country	region	country
R-sq	0.765	0.756	0.477	0.463
Q AR(1) p	0.174	0.960	0.00300	0.205
Q AR(2) p	0	0	0.0100	0.530
HT I(1) (z<-1.65)	−16.35	−6.465	−14.96	−6.204

quite small at just 0.2 percentage points. The residual autocorrelation does not readily disappear when adding more lags. Columns (3) and (4) show a specification with four lags, which is the maximum number of lags that can be considered while keeping the regional and country level samples identical. Even this number of lags is not sufficient to remove the residual autocorrelation. The difference between the country level and regional NAIRU estimate is about 1 percentage point, or 14 percent in relative terms.