

# VU Research Portal

## Agglomeration and Human Capital Spillovers

Verstraten, P.W.A.H.

2019

### **document version**

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

### **citation for published version (APA)**

Verstraten, P. W. A. H. (2019). *Agglomeration and Human Capital Spillovers: Identification and Spatial Scope*. [PhD-Thesis - Research and graduation internal, Vrije Universiteit Amsterdam].

### **General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

### **Take down policy**

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

### **E-mail address:**

[vuresearchportal.ub@vu.nl](mailto:vuresearchportal.ub@vu.nl)

# Chapter 3

## Why do wages grow faster in urban areas? Sorting of high potentials matters

### 3.1 Introduction<sup>21</sup>

Urban economists have devoted considerable effort to analyzing the strong and positive relationship between wages and city size. Descriptions of the urban-rural wage differential date back to the 19<sup>th</sup> century (Weber, 1899), although attempts to empirically identify the source are more recent (e.g., Glaeser and Maré, 2001; Combes et al., 2008; Rosenthal and Strange, 2008; Groot et al., 2014; De la Roca and Puga, 2017). The urban wage premium continues to attract a great deal of interest because it implies that labor is more productive in big cities than in more sparsely populated areas. After all, firms require compensation for locating and staying in these high wage areas (Moretti, 2011). Having proper knowledge of the sources underlying the urban wage premium is therefore indispensable for developing a full understanding of the urban economy.

Although substantial progress has been made in explaining the urban wage premium, much is left to be resolved. In particular, reviews of the empirical literature (e.g., Melo et al., 2009; Puga, 2010) point out that non-random sorting of high-skilled

---

<sup>21</sup> Apart from minor changes, this chapter was published as: Verstraten, P., Verweij, G. and Zwaneveld, P.J. (2018). *Why do wages grow faster in urban areas? Sorting of high potentials matters*. CPB Discussion Paper (No. 377).

individuals and more productive firms into urban areas remains a first-order problem when quantifying the benefits of urban areas. Another key concern is related to the temporal scope of the urban production advantages. As extensively discussed by Duranton and Puga (2004), there exists a plethora of mechanisms that foster the productivity of urban workers and some of them take time to become effective, such as human capital spillovers. Hence, some mechanisms are expected to influence the urban wage premium via a wage-level effect, whereas others materialize through wage growth, or a mixture of both.<sup>22</sup> This chapter sheds light on these two identification issues and, in particular, on how spatial sorting explains the observed urban wage-growth premium.

Traditionally, most of the empirical literature considers the urban wage premium to be a wage-level phenomenon driven by agglomeration economies. That is to say, urban workers are assumed to enjoy a wage-level premium that is instantaneously obtained by working in an urban area and lost upon relocating to a more rural area. This view is graphically illustrated by the bottom two lines in Figure 3.1. Estimates of the wage-level elasticity of agglomeration are, however, likely to be confounded by spatial sorting on productivity by both workers and firms. For instance, workers with more valuable innate abilities may self-select into bigger cities because they are attracted to the higher returns to education (Costa and Kahn, 2000) or the wide variety of urban amenities (Lee, 2010; Van Duijn and Rouwendal, 2013). Also, larger firms tend to be more productive (Melitz, 2003) and have stronger incentives to relocate their production facilities to big markets (Baldwin and Okubo, 2005). Faberman and Freedman (2016) find empirical evidence that high-wage establishments are most likely to relocate towards bigger cities. For these reasons, empirical studies include controls for (un)observed heterogeneity in wage levels among workers (e.g., Combes et al., 2008) and firms (e.g., Mion and Naticchioni, 2009).


Ever since Glaeser and Maré's (2001) seminal contribution and the follow-up studies by Wheeler (2006) and Yankow (2006), it is known that the urban wage

---

<sup>22</sup> In this chapter, we prefer to speak of 'wage-level' and 'wage-growth' effects, as opposed to the terminology used by De la Roca and Puga (2017), who refer to these phenomena as 'static' and 'dynamic' effects, respectively. The reason is to avoid confusion between the different types of mechanisms that drive agglomeration spillovers and the way in which they capitalize into individual wages. For instance, learning mechanisms, which are considered to be a dynamic process, do not necessarily capitalize into wages via growth effects alone. Instead, they can also result in higher wage levels as it increases the region's capacity to adapt new technology.

premium does not entirely consist of wage-level effects. Instead, urban workers appear to become more productive over time compared to rural workers, giving rise to an urban wage-growth premium that is portable to other areas (see the top two lines in Figure 3.1). This empirical finding supports the long-held view that cities are important in facilitating human capital accumulation (Marshall, 1890; Glaeser, 1999). More recently, De la Roca and Puga (2017) have generated renewed interest in this topic. Using detailed data for Spain, the authors demonstrate that urban wage-growth effects can fully account for the wage effect that was generally understood to be caused by sorting on time-invariant worker characteristics. This result suggests that spatial sorting is less important for explaining the urban wage premium than traditionally deduced from standard wage-level equations.

Figure 3.1. Wage-level and wage-growth effects of urban areas



*Note:* The idea for this figure is inspired by De la Roca and Puga (2017), see, e.g., Figure 3 in their article.

This chapter's main contribution is to extend the wage equation proposed by De la Roca and Puga (2017) by accounting for individual-, industry- and firm size-specific differences in wage growth. The inclusion of these extra variables is appropriate because earlier empirical work indicates that substantial heterogeneity in the growth rate of individual earnings exists (e.g., Baker, 1997). Moreover, earlier

evidence suggests that rapid urban wage growth may (partly) be the result of high individual-specific returns to experience (D’Costa and Overman, 2014) and knowledge spillovers within large firms (Lehmer and Möller, 2010). Our empirical analysis, which accounts for all these factors, demonstrates that the wage-growth controls account for approximately two-thirds of the urban wage-growth premium. Although the remaining one-third is still large enough to be economically significant, we conclude that spatial sorting is the most important determinant of the urban wage-growth premium.

Furthermore, this chapter shows that the wage-agglomeration elasticities, as applicable to the complete sample, conceal large heterogeneities among different types of workers. In particular, the wage-level benefits of urban areas increase with the individual’s educational attainment level, and are also higher for younger workers and those employed in knowledge intensive sectors. Heterogeneous effects among workers are less important when quantifying the wage-growth effects of urban areas. As well as for the complete sample, most estimates become statistically insignificant when controlling for spatial sorting. However, we do find evidence that younger workers are more receptive to the wage-growth benefits of urban areas, even after controlling for all types of spatial sorting.

The remainder of this chapter is structured as follows. Section 3.2 gives a description of the panel data, which contain detailed information on individual earnings and characteristics of both workers and firms in the Netherlands. In Section 3.3, we estimate a wage-level equation, similar to Combes et al. (2008), which confirms the conventional view that spatial sorting is an important source underlying the urban wage premium. Section 3.4 replicates the results of De la Roca and Puga (2017) for the Netherlands. Individual wages are shown to grow more rapidly in bigger cities, and these wage-growth effects can fully account for the wage gap that was thought to be the result of sorting on time-invariant skills of workers. Then, in Section 3.5, we introduce controls for wage-growth determining characteristics of workers and firms. This natural extension of De la Roca and Puga’s (2017) model shows that the urban wage-growth premium is to a large extent driven by spatial sorting. Section 3.6 addresses four key estimation issues, and Section 3.7 examines heterogeneities among different types of workers. Section 3.8 concludes.

## 3.2 Data description

### 3.2.1 Treatment of the microdata and summary statistics

Our empirical model requires three sets of non-public microdata from Statistics Netherlands (CBS): fiscal data (*Polisadministratie*), census data (*Sociaal Statistisch Bestand*), and firm data (*Algemeen Bedrijven Register*). Together, these datasets contain individual information for all non-self-employed employees in the Netherlands on pre-tax wages and other financial rewards, hours worked, date of birth, gender, educational attainment, sectoral classification of the employer (two-digit NACE), place of work at the municipality level, firm size, and job type. Based on this information we construct a panel (2006–2014) with yearly observations for each individual.

The wage data not only contain regular pre-tax wages, but also overtime payments, bonuses, thirteenth month salaries and company cars. Paid holidays are not included in the wages because they could not be assigned to a specific year. As this component comprises in general a fixed wage premium of eight percent, omitting it will not influence the estimation results. The reported number of hours worked consists of both regular and overtime hours. Dividing the sum of these annual financial rewards by the number of hours worked and deflating them with the consumer price index, provides an adequate approximation of the total hourly labor costs of each employee in a particular year.

The data are further restricted as follows. We excluded all workers under 18 and above 65 years old. Also, jobs with less than 12 hours of work per week, the official definition by Statistics Netherlands for being employed, are excluded from the sample. In order to limit the influence of non-regular workers, we decided to drop the following job types: owner-director, intern, outsourced worker, on-call worker, and WSW-worker.<sup>23</sup> Jobs in agriculture, forestry, and the fishing industry are excluded from the sample because these sectors are strongly linked to the location of natural resources. Also the public sectors are excluded because wages are heavily regulated in these sectors.<sup>24</sup> Jobs provided by a firm with establishments in more than one

---

<sup>23</sup> The WSW is a Dutch law aimed to foster the employment of persons with disabilities.

<sup>24</sup> This leaves us with a total of 70 economic sectors, based on the two-digit NACE classification.

municipality, could not be assigned geographically and had to be removed from the sample.<sup>25</sup> Furthermore, for those people with more than one job during a year, we restrict the analysis to the job with the highest number of hours worked during that particular year. Workers with only one observation over the period 2006–2014 do not contribute to the estimation results and are excluded, as are workers with a non-consecutive employment history. We excluded the data from seven municipalities because they are either islands or very small (less than 30 job changes over the period 2006–2014).<sup>26</sup> Outliers are defined as hourly wages below the legal minimum wage and above 20 times this minimum wage, and they are removed. After cleaning the data, nearly 700,000 observations per year remain.<sup>27</sup> Table 3.1 summarizes the longitudinal data remaining for estimation in the years 2006, 2010 and 2014.<sup>28</sup>

---

<sup>25</sup> This data limitation mainly concerns large firms. As a consequence, workers employed at large firms are underrepresented in our sample when compared to the national average. According to Statistics Netherlands, approximately 60 percent of the Dutch workforce is employed at firms with over 100 employees, whereas this is only 24 percent in our sample, see Table 3.1.

<sup>26</sup> This leaves us with 396 municipalities. Municipalities that have been merged between the years 2006 and 2014 are made time-consistent using information about work locations at the postal code level.

<sup>27</sup> The number of observations per year is larger in the middle of the sample period compared to the tails. There are two main reasons for this. First, workers who enter/quit the labor market in the last/first year of our sample period represent only one worker-year observation and are therefore removed. Second, we have excluded workers with nonconsecutive employment histories (interruptions of at least one year) between 2006 and 2014. Employment gaps during the first or last few years of our sample period are not considered to be an interruption, which preserves observations in the middle of the sample period.

<sup>28</sup> To verify the representativeness of the sample with respect to industrial composition, we have compared our sample to the National Accounts of Statistics Netherlands. Both figures are quite similar, although retail and financial services seem to be somewhat underrepresented in our sample, whereas construction and business services are overrepresented. Differences are largely driven by the exclusion of jobs with less than 12 hours of work per week. Also, our dataset does not contain information on the self-employed workers. The self-employed comprise between 10 to 15 percent of the total Dutch working population.

Table 3.1. Summary statistics of the longitudinal wage data

	2006	2010	2014
<i>Number of workers</i>	576,480	708,461	538,437
<i>Hourly wages in euro's (price level 2006)</i>			
Mean (standard deviation)	18.1 (9.9)	18.8 (10.9)	18.9 (11.6)
1 <sup>st</sup> percentile	7.8	8.1	8.2
Median	15.7	15.9	15.5
99 <sup>th</sup> percentile	56.8	61.8	65.4
<i>Age</i>			
Mean (standard deviation)	36.5 (10.5)	37.0 (11.2)	37.6 (11.3)
1 <sup>st</sup> percentile	19.8	19.8	20.8
Median	35.3	35.9	36.0
99 <sup>th</sup> percentile	59.8	61.8	62.6
<i>Firm size, number of employees (in percentages)</i>			
1–9	29.6	30.5	28.3
10–99	48.4	45.9	47.1
100–999	19.6	20.8	21.0
≥ 1000	2.5	2.9	3.6
<i>Industrial composition (in percentages)</i>			
Manufacturing	21.1	19.3	20.1
Construction	10.7	10.5	8.2
Logistics	5.9	6.3	6.5
Wholesale	13.8	13.8	14.6
Retail	7.4	7.9	7.9
Consumer services	3.1	3.5	3.8
Hospitality industry	5.6	6.3	7.0
ICT	6.0	6.4	7.0
Financial services	3.5	2.8	2.7
Business services	22.9	23.3	22.3



As will become clear in the next sections, our identification strategy relies on workers who migrate across municipalities between 2006 and 2014.<sup>29</sup> Table 3.2 shows that approximately 20 percent of the workers in our sample moves between municipalities at least once. The other 80 percent does not contribute to the identification of area-specific wage effects, although they are valuable for a proper estimation of the control variables. The strong dependency on movers to identify the area-specific wage effects might be troublesome if movers differ systematically from the non-movers. Table 3.2 reveals that both groups are quite similar in terms of observables. Only workers older than 40 years are relatively more numerous among the non-movers compared to the movers.

*Table 3.2. Movers and non-movers between 2006 and 2014*

	<b>Movers</b>	<b>Non-movers</b>
<i>Total workers</i>	256,480	7 ,23,334
<i>By gender</i>		
Male (%)	68	63
Female (%)	32	37
<i>By age (in years, 2010)</i>		
[18;30] (%)	42	38
[30;40] (%)	30	26
[40;65] (%)	28	36
<i>By education level</i>		
Low (%)	29	31
Medium (%)	42	41
High (%)	29	29

<sup>29</sup> In this research we use place of work rather than place of residence to assign a geographic location. Hence, in this setting, migration does not necessarily imply change of residence.

### 3.2.2 Computation of the spatial variables

The size of municipalities is computed very similar to De la Roca and Puga (2017). Using GIS tools, we first draw concentric rings with a radius of 10 kilometer around each postal codes' geographic centroid. Then, using the LISA employment register, we sum the total number of jobs within each concentric ring.<sup>30</sup> These figures are aggregated to the level of municipalities, using weights corresponding to the number of jobs within a postal code. This agglomeration measure gives us the employment level within 10 kilometer from the municipality's average job. The same procedure is used to calculate employment levels within 40, 80, and 120 kilometer distance. These additional ring variables will be used in Section 3.6.


As explained clearly by De la Roca and Puga (2017), the use of this agglomeration measure has some advantages compared to employment density, which is more commonly used as a measure of size (e.g., Combes et al., 2008; Mion and Naticchioni, 2009; Groot et al., 2014). Most importantly, density figures are highly sensitive to the degree of tightness to which the administrative borders are drawn around the built-up area. Also, density measures fail to account for economic activity in neighboring municipalities. This might introduce substantial noise, especially when dealing with polycentric city structures. Figure 3.2 shows the implications of our newly computed agglomeration measure by plotting it against employment density. The variables are strongly correlated, although the relationship is far from perfect. For instance, on the basis of employment density, we would consider the municipality of Urk to be a relatively large urban area. However, upon closer examination it becomes clear that Urk's administrative borders are tightly drawn around a rather small city, even to Dutch standards, while there is almost no economic activity beyond that border. The opposite holds for the municipality of Landsmeer, which is adjacent to Amsterdam, but contains large amounts of unbuilt land. The size of Rotterdam would also be underestimated when considering density alone. This is because a large share of Rotterdam's area consists of sparsely populated harbor districts, while most of Rotterdam's economic activity is located in the center of a large metropolitan area.

---

<sup>30</sup> To measure city size, we prefer employment levels over population counts because the place of work is more directly linked to economic activity than place of residence. Obviously, spatially aggregated employment and population figures are highly correlated (0.98 at the municipality level). Hence, this decision is unlikely to affect the results.

Finally, we make use of historical spatial data to compute instrumental variables (see also Section 3.6). First, we have collected population censuses of 1,217 Dutch municipalities in the year 1840, which were linked to a historical map of the Netherlands made available by Boonstra (2007). Then, similar to the construction of the agglomeration measure, we use GIS tools to calculate the historical population level within a 10 kilometer radius. Second, we made an intersection between the Netherlands' current surface area and the historical map to calculate the percentage of the municipality's area that was covered by water in 1840. Third, using spatial coordinates of ancient Roman forts obtained from Talbert (2000), we have estimated the minimum straight line distance between postal codes and forts. These straight line distances were aggregated to the level of municipalities using the number of jobs within a postal code as a weight.

Figure 3.2. Two measures of urban scale



### 3.3 The urban wage premium and the wage-level equation

In this section we estimate a traditional wage-level equation, which is frequently used in the literature to quantify the benefits of urban areas (e.g., Combes et al., 2008; Di Addario and Patacchini, 2008; Rosenthal and Strange, 2008; Mion and Naticchioni, 2009; Groot et al., 2014). This type of wage equation assumes that only current

worker characteristics, such as age and work location, are important to the wage setting. Detailed information on the worker's employment history is not included and, therefore, assumed irrelevant. With such a specification, the urban wage premium is treated as a pure level effect, which will be lost to the worker when he or she relocates to a more rural area (see Figure 3.1).

As already discussed in the introduction of this chapter, the wage-agglomeration relationship is likely confounded by spatial sorting on productivity by both workers and firms. Hence, to obtain an unbiased estimate of the urban wage effect, the wage equation should include controls for both worker and firm heterogeneity. Bearing this in mind, an appropriate identification strategy would be to estimate the following wage-level equation:

$$\log w_{i,t} = \sigma_{r(i,t)} + \sigma_{f(i,t)} + \sigma_{k(i,t)} + \sigma_i + \beta x_{i,t} + \tau_t + \varepsilon_{i,t}. \quad (3.1)$$

$w_{i,t}$  is the hourly wage of worker  $i$  in year  $t$ .  $\sigma_{r(i,t)}$  is an area fixed effect, which captures the wage effect associated with the individual's place of work, such as benefits of agglomeration. Wage effects stemming from firm-specific characteristics are controlled for by including firm-size  $\sigma_{f(i,t)}$  and industry  $\sigma_{k(i,t)}$  fixed effects.  $\sigma_i$  is a worker fixed effect, which captures all time-independent worker characteristics, and  $x_{i,t}$  is a vector of observed, time-dependent worker characteristics with scale parameter  $\beta$ . Finally,  $\tau_t$  is a set of year dummies and  $\varepsilon_{i,t}$  is an error term.<sup>31</sup>

---

<sup>31</sup> Labor and capital are considered to be imperfect substitutes under a Cobb-Douglas production function. Hence, the firm's capital-labor ratio is another aspect influencing the wage formation. We aim to control for this by including industry and firm-size fixed effects. Although this approach is admittedly indirect, it is likely to be sufficient for our purpose, as Abowd et al. (1999) find that worker heterogeneity is substantially more important in explaining wage differentials than firm heterogeneity. Alternative strategies to account for capital-labor ratios, such as the inclusion of firm fixed effects (e.g., Abowd et al., 1999; Mion and Naticchioni, 2009) or a TFP-based approach (e.g., Combes et al., 2012), face limitations as well. For instance, including firm fixed effects would destroy much of the variation needed to identify the area fixed effect. The variation available for identification would only stem from workers who relocate within the same firm or between firms with establishments in more than one area. This strong dependence of the results on multi-establishment firms raises additional concerns with respect to the influence of the location of the other establishments on the worker's wage. TFP-based approaches, on the other hand, have difficulties in accounting for spatial differences in labor quality.

After having estimated Equation (3.1), we obtain a wage-agglomeration elasticity  $\xi_S$  by regressing the area-specific wage effect  $\sigma_r$  on our log-transformed agglomeration measure  $E_r^{0-10km}$ :

$$\sigma_r = \alpha + \xi_S \log E_r^{0-10km} + \varepsilon_r, \quad (3.2)$$

where  $\alpha$  is a constant. This two-stage procedure is preferred over estimating the elasticity in one single stage because it offers an elegant solution to the dependent disturbances within areas.<sup>32</sup>

Before turning to the estimation results, two remarks should be made. First, when urban wage-growth effects do exist but are ignored as in Equation (3.1), the wage-agglomeration elasticity  $\xi_S$  will reflect a mixture of both level and growth effects. De la Roca and Puga (2017), however, show that  $\xi_S$  can provide an accurate estimate of the urban wage-level elasticity of agglomeration when worker fixed effects are included to the model and two reasonable conditions are satisfied: balanced migration and full portability of the wage-growth effects.<sup>33</sup> Second, Equation (3.1) strongly relies on migrants moving from one area to another because the wage effect stemming from the area cannot be distinguished from the individual fixed effects of non-movers. This aspect of the model is also one of its main drawbacks: migrants may differ systematically from non-movers, yielding unrepresentative wage-agglomeration elasticities. However, we did not find strong evidence of this bias in our data (see Section 3.2).

Table 3.3 shows the results of the wage-level equation in various forms. Column (1) is the most basic version, having only year dummies, and it gives us a first impression of the magnitude of the urban wage premium. The crude wage-agglomeration elasticity amounts to 0.056, which corresponds to a wage increase of approximately four percent ( $2^{0.056}$ ) every time the area doubles in size. Our estimate falls within the range of values typically found in the literature and is particularly close to the elasticity of 0.049 found by Combes et al. (2008).

---

<sup>32</sup> See Moulton (1990) for a more detailed discussion on the econometric problems that arise from observations sharing the same geographic space. The use of worker fixed effects in this study makes the standard solution of calculating clustered robust standard errors not applicable (Combes et al., 2008).

<sup>33</sup> We will show in Section 3.4 that the estimates of our fully specified Equation (3.1), i.e. including worker fixed effects, are indeed quite insensitive to the existence of urban wage-growth effects.

In columns (2) to (4) we gradually introduce controls for observed worker and firm characteristics. All coefficients have the expected sign and they explain a large share of the variation in individual earnings.<sup>34</sup> Also, the positive wage effect of firm size is of the same size as the gender wage gap, which is in line with the empirical literature (Oi and Idson, 1999). The wage-agglomeration elasticity falls from 0.056 to 0.026, implying that more than half of the crude urban wage premium is explained by spatial sorting on observables. Finally, column (5) shows the results of the fully specified wage-level equation.<sup>35</sup> The introduction of worker fixed effects lowers the estimated elasticity even further from 0.026 to 0.012, amounting to a drop of 54 percent. This worker fixed effect-induced decline is close to the 47 percent found by De la Roca and Puga (2017).

Having controlled for unobserved worker characteristics, the wage-agglomeration elasticity of 0.012 is relatively low compared to international standards. For instance, Combes et al. (2008) and De la Roca and Puga (2017) find substantially larger elasticities of 0.032 and 0.024 for France and Spain, respectively. Nevertheless, our results are by no means unique. Mion and Naticchioni (2009) and Andersson et al. (2014) find even lower elasticities of 0.007 and 0.005 for Italy and Sweden. In Section 3.6 we demonstrate that significantly larger elasticities can be obtained when taking into account the Netherlands' polycentric urban structure. However, this involves a more complicated concentric ring-based specification. For the sake of comparison with other empirical studies, we derive our main results using the standard approach.

---

<sup>34</sup> The  $R^2$  of 0.52 in column (4) is comparable to the 0.50 obtained by both Groot et al. (2014) and De la Roca and Puga (2017), while it is substantially larger than the 0.41 and 0.22 found by Di Addario and Patacchini (2008) and Rosenthal and Strange (2008), respectively.

<sup>35</sup> Note that the linear age variable is omitted from the worker fixed effects model. This is because the year dummies can fully account for the linear age effects when having only within-variation.

Table 3.3. Wage benefits of urban areas

Column:	(1)	(2)	(3)	(4)	(5)
First-stage equation:	(3.1)	(3.1)	(3.1)	(3.1)	(3.1)
Age		0.078*** (0.000)	1 fe5*** (0.000)	0.0695*** (0.000)	–
Age squared		–0.001*** (2.5e-06)	–0.001*** (2.4e-06)	–0.001*** (2.3e-06)	–0.001*** (2.2e-06)
Female		–0.208** (0.001)	–0.195*** (0.001)	–0.1625*** (0.001)	–
Medium-educated		0.166*** (0.001)	0.166*** (0.001)	0.1468*** (0.001)	–
Highly educated		0.472*** (0.001)	0.456*** (0.001)	0.4008*** (0.001)	–
Firm size: 10–99 employees			0.087*** (0.001)	0.0749*** (0.001)	0.0179*** (0.000)
Firm size: 100–999 employees			0.171*** (0.001)	0.1479*** (0.001)	0.0348*** (0.001)
Firm size: ≥1000 employees			0.234*** (0.002)	0.2386*** (0.002)	0.076*** (0.002)
Year effects	YES	YES	YES	YES	YES
Area indicators	YES	YES	YES	YES	YES
Industry effects	NO	NO	NO	YES	YES
Worker fixed effects	NO	NO	NO	NO	YES
$R^2$	0.057	0.448	0.467	0.518	0.190
Second-stage equation:	(3.2)	(3.2)	(3.2)	(3.2)	(3.2)
<b>Wage benefit</b>					
Log employment within 10 km	0.056*** (0.005)	0.034*** (0.003)	0.030*** (0.003)	0.026*** (0.003)	0.012*** (0.001)
$R^2$	0.260	0.214	0.207	0.215	0.226

Notes: First- and second-stage estimates are based on 6,130,091 worker-year observations and 396 area-specific wage effects, respectively. Industry indicators are based on two-digit NACE. Robust standard errors, which are clustered by worker in the first-stage estimates, are in parentheses. The first-stage  $R^2$  in column (5) is within workers. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

### 3.4 The urban wage-growth premium

Although the wage-level equation is helpful for unraveling spatial wage disparities, it does not provide a complete description of the sources underlying the urban wage premium. As we know from Glaeser and Maré (2001), a part of the urban benefits accrues to workers over time and remains with them after relocating to more rural areas. The wage equation proposed by De la Roca and Puga (2017) addresses this

issue by including an additional term for the area-specific experience of workers. A slightly modified version of their wage equation has the following form:

$$\log w_{i,t} = \sigma_{r(i,t)} + \sigma_{f(i,t)} + \sigma_{k(i,t)} + \sigma_i + \beta x_{i,t} + \tau_t + \sum_{r=1}^R \delta_r e_{r(i,t)} + \varepsilon_{i,t}, \quad (3.3)$$

where the new summation term reflects the total value of experience acquired in all areas.  $e_{r(i,t)}$  is total experience accumulated by worker  $i$  up and until year  $t$  in area  $r$ , and  $\delta_r$  is a parameter indicating the value of this experience acquired in area  $r$ . Then, we obtain a wage-growth-agglomeration elasticity  $\xi_D$  by regressing the estimated parameter  $\delta_r$  on our log-transformed agglomeration measure  $E_r^{0-10km}$ :

$$\delta_r = \alpha + \xi_D \log E_r^{0-10km} + \varepsilon_r. \quad (3.4)$$

This wage-growth elasticity of agglomeration reflects how the value of experience depends on the size of the area where it was acquired. A positive wage-growth elasticity thus implies that working in urban areas speeds up individual wage growth.

Column (1) in Table 3.4 presents the results of Equation (3.3) and corresponding second-stage Equations (3.2) and (3.4). Our results indicate that both wage-agglomeration elasticities are statistically significant at the one percent level. Furthermore, the magnitude of the wage-growth elasticity of agglomeration is relatively large. It takes less than five years of local experience to have an equal share of both wage-level and wage-growth benefits from urban areas. If the worker does not relocate after five years, the share of urban wage-growth benefits will become increasingly larger. These results are also graphically presented by the top-panel of Figure 3.3, which depicts the temporal wage evolution of a worker located in Amsterdam (or in a median-sized Dutch city) relative to a worker from a rural area. In particular, this figure reveals that rural workers with five years of experience in Amsterdam, will earn higher wages than workers from medium-sized cities who have worked in these cities for 10 years.



Although we employ a slightly modified version of De la Roca and Puga's (2017) wage equation, the results are remarkably similar.<sup>36</sup> For instance, accounting for wage-growth effects of urban areas has only a minor effect on the wage-level elasticity of agglomeration. Compared to column (5) in Table 3.3, the wage-level elasticity of agglomeration declines roughly 8 percent, which closely resembles the 9 percent found by De la Roca and Puga (2017). In addition, when applying their approach of calculating a medium-term wage-agglomeration elasticity, which consists of a wage-level effect plus a medium-term wage-growth effect, we find that this elasticity is only 23 percent larger compared to the wage-level elasticity of agglomeration obtained without worker fixed effects (column (4) in Table 3.3).<sup>37</sup> De la Roca and Puga (2017) find this difference to be approximately 12 percent.

This non-significant difference between the medium-term elasticity and the wage-agglomeration elasticity without worker fixed effects is considered to be an important result. De la Roca and Puga (2017) argue that this similarity suggests that urban workers do not have higher time-invariant ability. Instead, urban workers become more productive over time compared to rural workers, which can fully account for the wage gap that was thought to be the result of sorting on time-invariant

---

<sup>36</sup> There are three main differences between Equation (3.3) and De la Roca and Puga's (2017) model. First, their model allows the value of experience to depend on the worker's current location, which is used to test whether urban wage-growth effects are portable to more rural areas. We choose to omit this feature from the model because the empirical literature strongly supports the portability of wage-growth effects (Glaeser and Maré, 2001; Matano and Naticchioni, 2016; De la Roca and Puga, 2017). Second, our model does not include squared experience terms to allow for concave returns to experience. This is due to practical limitations, as we have incomplete information on the workers' employment history before the start of our sample period (see also Section 3.2). The lack of squared terms, however, is only a minor restriction because our primary interest is in the overall impact of experience. Moreover, this data limitation is not problematic for estimating the linear experience terms since all experience accumulated before the start of the sample period is absorbed by the worker fixed effects. Third, De la Roca and Puga (2017) divide the urban areas into three larger subgroups, based on their population size. This avoids problems related to identification and statistical significance, especially when having many interaction terms, but it does not facilitate comparison with other empirical studies. In this chapter, we prefer to derive a more general result by estimating the wage-growth effects for each area, and use these parameters in a second-stage regression to obtain a wage-growth elasticity of agglomeration.

<sup>37</sup> We use parameters  $\sigma_r$  and  $\delta_r$  to calculate a medium-term wage-agglomeration elasticity  $\xi_M$ , based on the average duration of employment within a municipality (8.9 years in the Netherlands). To this end, we estimate the following second-stage equation:  $(\sigma_r + 8.9 \times \delta_r) = \xi_M \log E_r^{0-10km} + \varepsilon_r$ . The estimated elasticity equals 0.032 with a standard error of 0.003.

ability.<sup>38</sup> In Section 3.5, we will explore whether this conclusion is robust to controlling for wage-growth determining characteristics of workers and firms.

### 3.5 Spatial sorting on wage-growth determining characteristics

Now we extend Equation (3.3) in a natural way by introducing a wage-growth control for every wage-level counterpart, which gives us the following equation:

$$\log w_{i,t} = \sigma_{r(i,t)} + \sigma_{f(i,t)} + \sigma_{k(i,t)} + \sigma_i + \beta x_{i,t} + \tau_t + \sum_{r=1}^R \delta_r e_{r(i,t)} + \sum_{f=1}^F \delta_f e_{f(i,t)} + \sum_{k=1}^K \delta_k e_{k(i,t)} + \delta_i e_{i,t} + \varepsilon_{i,t}, \quad (3.5)$$

where  $e_{f(i,t)}$  and  $e_{k(i,t)}$  reflect total experience acquired by worker  $i$  up and until year  $t$  within, respectively, a firm of size  $f$  and industry  $k$ .  $\delta_f$  and  $\delta_k$  are corresponding parameters indicating the value of these different types of experience. Finally,  $e_{i,t}$  is the worker's own total experience with individual-specific return to experience  $\delta_i$ .

The main objective of this extension is to evaluate whether Equation (3.3) is capable of inferring an unbiased estimate of the relationship between wage growth and the city size. In particular, we will argue that the commonly applied wage-level controls, although necessary for estimating the wage-level effects of urban areas, are not sufficient for obtaining an unbiased estimate of the wage-growth effects of urban areas, especially when firms and workers with inherently higher wage growth sort into urban areas. For instance, if Equation (3.5) holds and we use Equation (3.3) for estimating the wage-growth effects of urban areas, then  $\delta_i e_{i,t}$  will be absorbed into the error term. For an individual worker  $\text{cov}(e_{r(i,t)}, e_{i,t}) > 0$  and  $\delta_i > 0$ , creating an upward bias in the estimated  $\delta_r$  (Wooldridge, 2002, pp. 61–63). As long as workers are distributed randomly across areas, every  $\delta_r$  will be biased by the same amount. However, if urban areas attract relatively more workers with a higher return to experience  $\delta_i$  then the estimated  $\delta_r$  for these areas will be upward biased compared to more rural areas. This is an important concern because empirical studies suggest that sorting into big cities of both workers (D'Costa and Overman, 2014) and firms

---

<sup>38</sup> Baum-Snow and Pavan's (2012) structural model estimation also supports the idea that wage-level and wage-growth effects are the most important sources behind the urban wage premium, whereas sorting on time-invariant ability contributes relatively little.

(Lehmer and Möller, 2010) is partly based on wage-growth determining characteristics.<sup>39</sup> Evidently, including wage-level variables, such as individual fixed effects, will not capture these omitted wage-growth variables.

Columns (2) to (4) in Table 3.4 gradually introduce the wage-growth controls from Equation (3.5) and report the results of the corresponding second-stage estimates. Starting with column (2), we see that wage growth is positively related to the size of firms and that these wage-growth effects of large firms are substantial. Also, the inclusion of firm size-specific wage-growth controls reduces the wage-growth elasticity of agglomeration from 0.0023 to 0.0020. This decline is, however, relatively small in magnitude and insignificant at conventional levels. A larger reduction in the wage-growth elasticity occurs when industry-specific wage-growth controls are included as well, see column (3). Together, the wage-growth controls for firm characteristics reduce the wage-growth elasticity by 43 percent, which is a significant reduction at the one percent level.

We now turn to our fully specified Equation (3.5), which includes controls for worker-specific returns to experience. Adding these controls to the model implies that the wage-growth effects of areas  $\delta_r$  are identified on the basis of migrants only. After all, for non-movers it is not possible to disentangle the wage-growth effects stemming from the individual characteristics and those from the work location. Estimating such a large set of individual-specific wage-growth effects, in addition to the usual worker fixed effects, is computationally challenging. In order to estimate this fully specified equation, we apply the Frisch-Waugh-Lovell theorem to eliminate all worker-specific controls in a preliminary regression step, and estimate the remaining parameters of Equation (3.5) on the transformed variables.<sup>40</sup> The results in column (4) show that this fully specified equation yields an even lower, insignificant wage-growth elasticity

---

<sup>39</sup> De la Roca and Puga (2017) do explore whether there are complementarities between the individual's time-invariant ability  $\sigma_i$  and the wage-growth effects of urban areas  $\delta_r$ . They incorporate this possible interaction into the framework by imposing a proportionality assumption between  $\sigma_i$  and the returns to experience from working in bigger cities. Although this approach provides insight into the heterogeneities in earnings profiles, it is not fit to control for spatial sorting on wage-growth determining characteristics.

<sup>40</sup> The reported  $R^2$  of 0.005 in column (4) is relatively low because it relates to the regression on the transformed variables. That is to say, the worker fixed effects and worker-specific trends were already absorbed in the preliminary regression step. The relatively low value of the  $R^2$  implies that characteristics of firms and areas add relatively little explanatory power to the model once we have controlled for worker-specific characteristics.

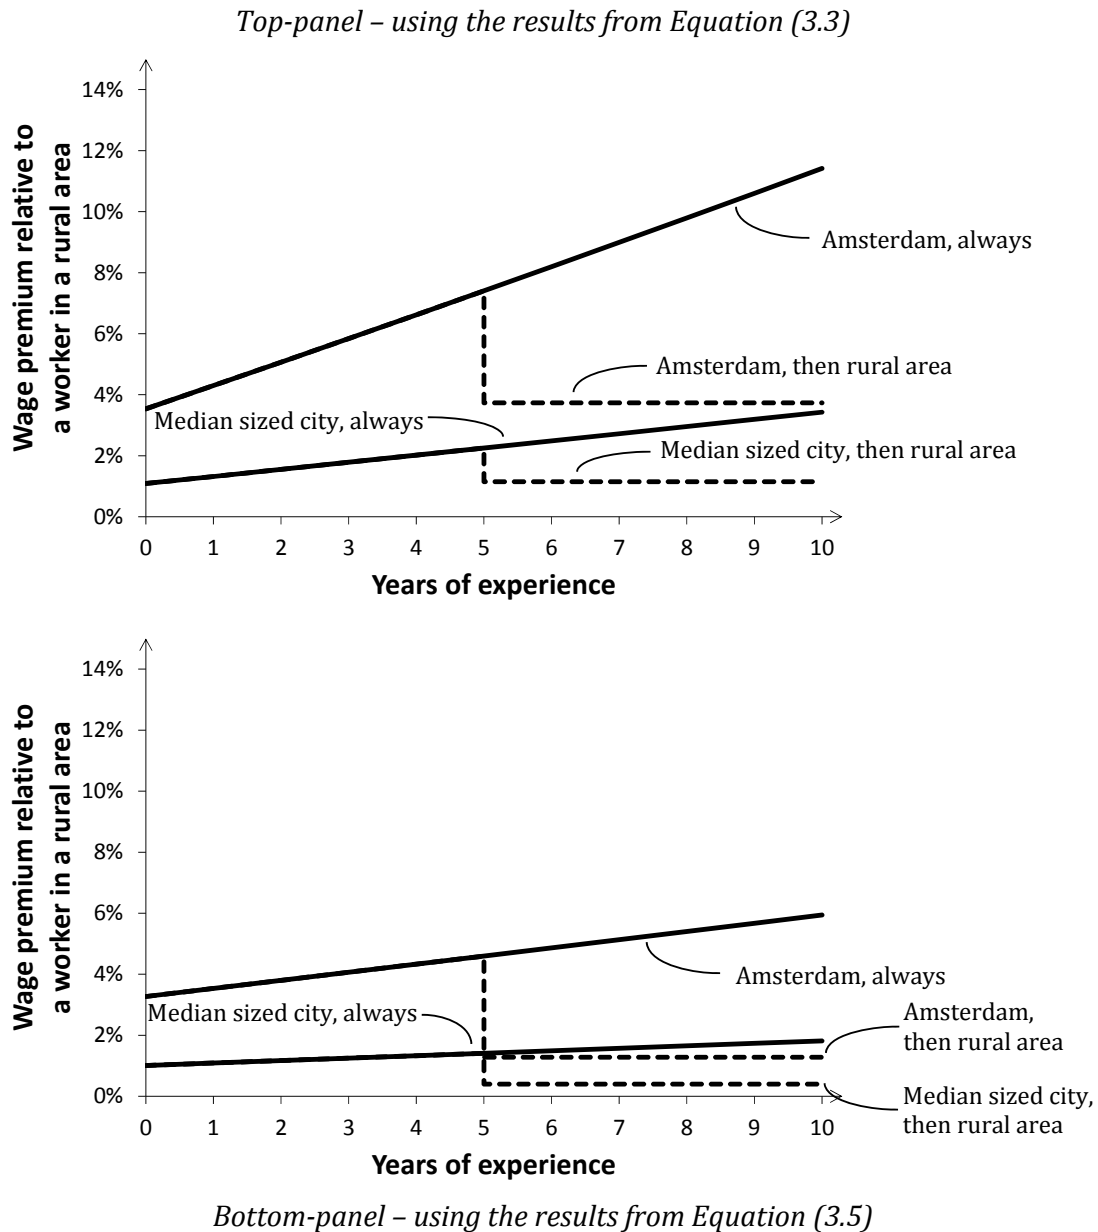
of 0.0008. This estimate also differs significantly from the elasticity obtained in column (1) using Equation (3.3). Hence, we conclude from this analysis that spatial sorting is the most important source underlying the urban wage-growth premium.

*Table 3.4. Wage-level and wage-growth benefits of urban areas*

Column: First-stage equation:	(1) (3.3)	(2) (3.5)	(3) (3.5)	(4) (3.5)
Age squared	-0.001*** (2.2e-06)	-0.001*** (2.2e-06)	-0.001*** (2.2e-06)	-0.001*** (1.4e-07)
Wage-level effect of firm size: 10–99 employees	0.018*** (0.000)	0.016*** (0.000)	0.015*** (0.000)	0.009*** (0.000)
Wage-level effect of firm size: 100–999 employees	0.034*** (0.001)	0.028*** (0.001)	0.028*** (0.001)	0.020*** (0.001)
Wage-level effect of firm size: ≥1000 employees	0.070*** (0.002)	0.058*** (0.002)	0.055*** (0.002)	0.037*** (0.002)
Wage-growth effect of firm size: 10–99 employees		0.005*** (0.000)	0.004*** (0.000)	0.001** (0.000)
Wage-growth effect of firm size: 100–999 employees		0.011*** (0.000)	0.009*** (0.000)	0.001 (0.001)
Wage-growth effect of firm size: ≥1000 employees		0.016*** (0.000)	0.017*** (0.000)	0.004*** (0.002)
Year effects	YES	YES	YES	YES
Wage-level/growth area indicators	YES	YES	YES	YES
Wage-level industry effects	YES	YES	YES	YES
Wage-level worker effects (FE)	YES	YES	YES	YES
Wage-growth industry effects		NO	YES	YES
Wage-growth worker effects		NO	NO	YES
$R^2$	0.198	0.201	0.208	0.059
Second-stage equations:	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)
<b>Wage-level benefit</b>				
Log employment within 10 km	0.011*** (0.001)	0.011*** (0.001)	0.011*** (0.001)	0.010*** (0.001)
$R^2$	0.192	0.194	0.209	0.165
<b>Wage-growth benefit</b>				
Log employment within 10 km	0.0023*** (0.0003)	0.0020*** (0.0002)	0.0013*** (0.0002)	0.0008 (0.0006)
$R^2$	0.159	0.158	0.091	0.005

*Notes:* First- and second-stage estimates are based on 6,130,091 worker-year observations and 396 area-specific wage effects, respectively. Industry indicators are based on two-digit NACE. Robust standard errors, which are clustered by worker in the first-stage estimates, are in parentheses. The first-stage  $R^2$  in columns (1) to (3) is within workers, whereas the first-stage  $R^2$  in column (4) is within workers including the absorption of the worker-specific trends. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Figure 3.3. Graphical representation of the main results



Finally, it is worth noting that most of the wage-growth controls for firm size remain statistically significant in column (4), although their magnitude is much lower than in column (3). This implies that, conditional on the worker's intrinsic wage-growth trend, wages increase more rapidly in larger firms. In contrast, the wage-growth elasticity of agglomeration is no longer statistically significant at conventional levels. Obviously, this statistically insignificant result does not allow for the conclusion

that wage-growth effects of urban areas are nonexistent or irrelevant, especially because the point estimate is still sufficiently large to be economically significant. The bottom-panel of Figure 3.3 shows that the benefits of working in Amsterdam, compared to a rural area, increase from an initial wage premium of 3.3 percent to a premium of 5.9 percent in 10 years.

### 3.6 Second-stage estimation issues

The estimation of the second-stage Equations (3.2) and (3.4) entails several identification problems. In this section we address four key issues in turn order: (i) possible inaccuracies in the estimates of the area-specific wage-level  $\sigma_r$  and wage-growth  $\delta_r$  effects, (ii) endogeneity concerns related to our agglomeration measure, (iii) nonlinearities in the wage-agglomeration relationship, and (iv) assumptions about the spatial scope of agglomeration economies.

In the previous sections we have used robust OLS to estimate the second-stage equations. This estimator, however, might produce biased and inconsistent estimates since the dependent variables  $\sigma_r$  and  $\delta_r$  do not contain observed but rather estimated values. This problem will be most eminent for areas with relatively little migrants, as this will likely yield larger standard errors in the first stage. In order to address this potential bias in the standard errors and inefficiencies, we have re-estimated the second-stage equations using a feasible generalized least squares (FGLS) estimator (Gobillon, 2004). Results are reported in column (1) of Table 3.5. Both the wage-level and wage-growth elasticities of agglomeration are only slightly smaller than those obtained with OLS, which is in line with the findings of Combes et al. (2008). We conclude from this exercise that the influence of estimation errors in the first stage can be neglected when estimating the second-stage equations.

Another source of concern is that the agglomeration measure can be endogenous. This widely recognized problem in urban economics states that the positive association between wages and city size may not reflect any causal relationship. Instead, there may be third factors influencing both variables, such as local endowments or variation in capital intensity (Moomaw, 1981; Combes et al, 2010a). The standard approach to tackle this issue is to find instrumental variables (IV) that correlate with the agglomeration measure, but that have no independent

relationship to wages. In this chapter we employ this IV approach using three different sets of instruments.

First, we follow Ciccone and Hall's (1996) pioneering work and use historical population censuses as an instrumental variable. The justification is that all factors influencing the spatial distribution of people before the industrial revolution, are no longer important for productivity in a modern economy. Still, the historical variable does strongly correlate with the current spatial distribution of economic activity – possibly due to path dependency in the size of cities (Bleakley and Lin, 2012) – which satisfies the relevance condition. Our second IV is a geological variable: percentage of the area that has been drained since 1840. The intuition behind this IV is that the decision to drain historical water bodies has no direct impact on today's productivity, except that it made land available for construction.<sup>41</sup> Third, we compute for each area the minimal distance to a Roman fort. Although the presence of an ancient fort is not expected to affect productivity, its influence is still noticeable through the early construction of infrastructure.

At the bottom of column (2) in Table 3.5 we present the results of the 2SLS first-stage regression. All three instruments have the expected sign and are highly statistically significant in explaining city size. Furthermore, the Kleibergen-Paap under-identification test and the Hansen  $J$  over-identification test confirm that our instruments are relevant and uncorrelated to the error term for both Equations (3.2) and (3.4). The results of column (2) reveal that the wage-growth elasticity of agglomeration is largely unaffected by the use of instruments. The point estimate declines from 0.0008 to 0.0003 and remains statistically insignificant. Also, the Hausman test for endogeneity does not reject the use of OLS when estimating the urban wage-growth premium. This is different for the wage-level elasticity of agglomeration, which increases significantly from 0.010 to 0.014.<sup>42</sup> Hence, in the remainder of this chapter we choose to derive all results using the 2SLS estimator. For

---

<sup>41</sup> Rosenthal and Strange (2008) were among the first to use geological variables to instrument current city size. The exact instruments used in this study – seismic/landslide hazard and sedimentary rock – are less appropriate for our setting because the Netherlands is rather homogeneous in geographical dimensions.

<sup>42</sup> The literature finds mixed results from IV-analyses. Some studies report reductions in the elasticity (e.g., Combes et al., 2008; Mion and Naticchioni, 2009), whereas others find an increase (e.g., Di Addario and Patacchini, 2008; Rosenthal and Strange, 2008; Groot et al., 2014). Overall, most of the empirical literature agrees that endogeneity of the agglomeration measure is only a second-order issue.

the sake of consistency and to facilitate comparison, this applies to both the wage-level and wage-growth elasticities of agglomeration.

*Table 3.5. FGLS and 2SLS estimates of the second-stage equations*

Column: Estimator: Second-stage equations:	(1) FGLS (3.2 & 3.4)	(2) 2SLS (3.2 & 3.4)
<b>Wage-level benefit</b>		
Log employment within 10 km	0.010*** (0.001)	0.014*** (0.002)
<i>p</i> -value Kleibergen-Paap rk LM statistic		0.000
<i>p</i> -value Hansen <i>J</i> statistic		0.457
<i>p</i> -value Hausman test for endogeneity		0.000
<b>Wage-growth benefit</b>		
Log employment within 10 km	0.0007 (0.0005)	0.0003 (0.0007)
<i>p</i> -value Kleibergen-Paap rk LM statistic		0.000
<i>p</i> -value Hansen <i>J</i> statistic		0.777
<i>p</i> -value Hausman test for endogeneity		0.205
<b>First-stage results of the 2SLS estimation</b>		
Log population in 1840		0.904*** (0.047)
Log minimal distance to Roman fort		-0.222*** (0.031)
Percent of area covered by water in 1840		0.990*** (0.266)
$R^2$		0.616

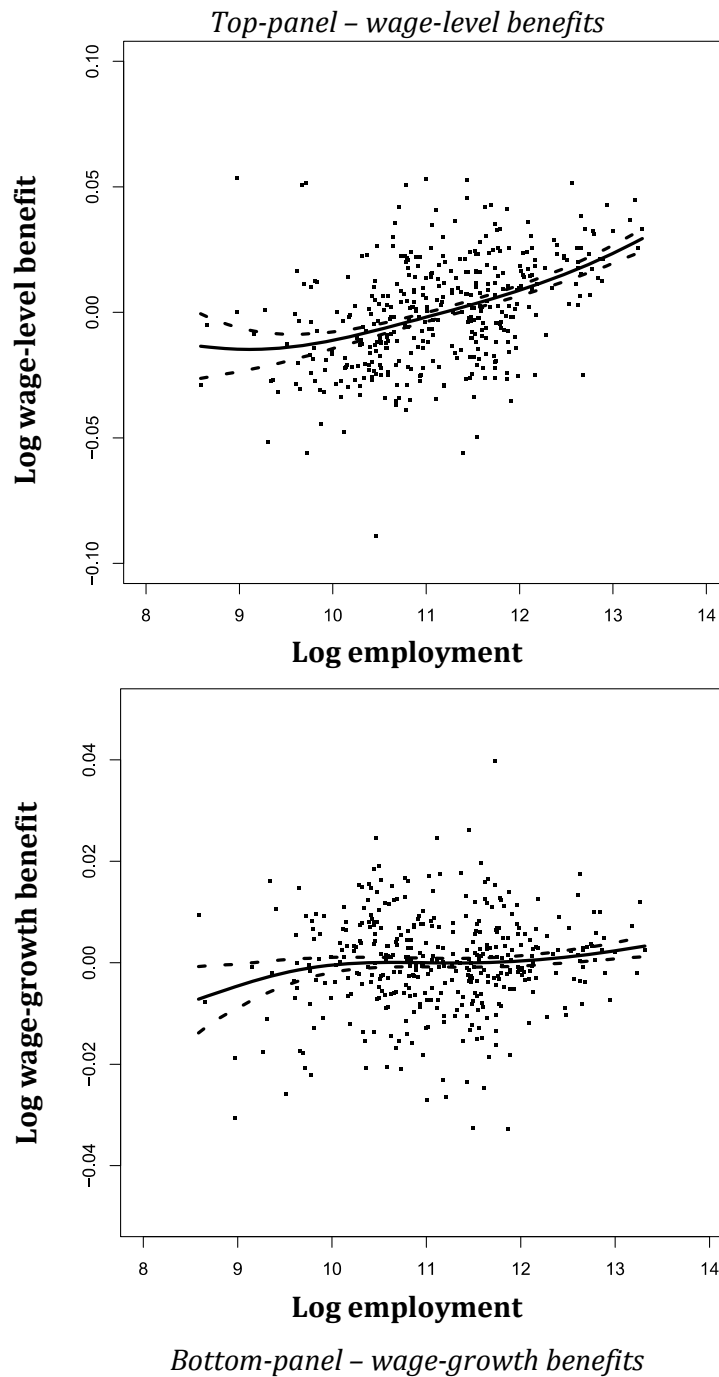
*Notes:* First-stage estimates for  $\sigma_r$  and  $\delta_r$  are obtained from Equation (3.5), see also column (4) in Table 3.4. Second-stage estimates are based on 396 area-specific wage effects. Standard errors, which are non-robust in column (1) and robust in column (2), are in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

We now turn to the possible nonlinearities in the wage-agglomeration relationship. Cities are generally understood to be the outcome of a trade-off between urban benefits and congestion costs (Duranton and Puga, 2004). This trade-off can result in a nonlinear relationship between wages and city size if, within some range of the city size distribution, urban benefits/costs become relatively more important for every additional unit of economic mass. For instance, it is possible that wage-growth benefits of areas exhibit an inverted U-curve as the area grows in size. This would



imply a positive/negative relationship between agglomeration and wage growth at the lower/higher end of the city size distribution.

Figure 3.4. Non-parametric kernel estimation of the wage-agglomeration relationship



*Notes:* Smoothing bandwidth equals 1. Kernel is Gaussian. The dashed lines represent 90 percent confidence intervals obtained by 1,000 bootstrap replications.

We examine potential nonlinearities non-parametrically in Figure 3.4. The non-parametric kernel estimations do not yield any indication for the existence of nonlinearities in the wage-agglomeration relationship. Consequently, we consider the traditional double-log regression equations to be an appropriate tool to analyze the wage-agglomeration relationship.

A fourth, and final, concern is related to the spatial scale at which agglomeration economies are expected to operate. In order to allow for proper comparison with other empirical work, we have until now employed a quite local agglomeration measure, which sums all employment within a 10 kilometer radius. This geographic scope, however, might be too restrictive for this setting, especially because the Netherlands consists of multiple middle-sized though well-connected cities. The previous chapter already confirmed that agglomeration economies can operate at a distance of 40–80 kilometer in the Netherlands.<sup>43</sup> Hence, by ignoring employment at distances beyond 10 kilometer, the estimated elasticities are likely to be downward biased. This might also explain why the elasticities found in this chapter are relatively low compared to international standards.

In order to determine the relevant spatial scope of agglomeration economies, we employ a concentric ring-based strategy. This approach was first proposed by Rosenthal and Strange (2003) and involves the computation of employment levels, and corresponding long-lagged population instruments, within a set of distance intervals: 0–10 kilometer, 10–40 kilometer, 40–80 kilometer, and 80–120 kilometer. These concentric ring variables can be used in the second-stage equations:

$$\sigma_r = \gamma_{S,1}E_r^{0-10km} + \gamma_{S,2}E_r^{10-40km} + \gamma_{S,3}E_r^{40-80km} + \gamma_{S,4}E_r^{80-120km} + \varepsilon_r, \quad (3.6)$$

$$\delta_r = \gamma_{D,1}E_r^{0-10km} + \gamma_{D,2}E_r^{10-40km} + \gamma_{D,3}E_r^{40-80km} + \gamma_{D,4}E_r^{80-120km} + \varepsilon_r, \quad (3.7)$$

where the  $E_r$  is the employment level within a particular distance interval from area  $r$ .  $\gamma_S$  and  $\gamma_D$  reflect the effect of the concentric ring on the wage-level and wage-growth effect, respectively. Also note that both equations take a log-linear form. This enables us to evaluate the effect of one unit of economic mass in a particular distance interval

---

<sup>43</sup> Similar results are obtained by Rosenthal and Strange (2008) and Rice et al. (2006) for the US and Great Britain, respectively.

relative to one unit in another distance interval. A numerical comparison between the four parameters provides information about the rate of distance decay.

Table 3.6 presents the results of this concentric ring-based approach. Column (1) conveys that the wage-level effects of agglomeration are relatively large within 10 kilometer distance and decay rapidly across space. Yet, the effect remains significantly positive until 40–80 kilometer, which closely resembles the findings of the previous chapter. A more remarkable result is that employment on 80 to 120 kilometer distance negatively affects the wage-level benefit. Taking this result at face value, may lead us to conclude that congestion costs dominate agglomeration economies on long distances. An alternative, and arguably more plausible, explanation is that the variable is picking up a periphery effect. This would be in line with Brakman et al. (2002), who find that, conditional on market potential, German regions close to the national border pay lower wages.

*Table 3.6. The spatial scope of agglomeration economies*

Column: Second-stage equation:	(1) (3.6)	(2) (3.2)	(3) (3.7)
Employment <10 km	0.048*** (0.012)		–0.001 (0.005)
Employment 10–40 km	0.010*** (0.002)		0.001 (0.001)
Employment 40–80 km	0.004*** (0.001)		0.002** (0.001)
Employment 80–120 km	–0.003** (0.001)		0.000 (0.001)
Log employment 0–80 km		0.017*** (0.002)	
IV	YES	YES	YES
<i>p</i> -value Kleibergen-Paap rk LM statistic	0.000	0.000	0.000
<i>p</i> -value Hansen J statistic	0.584	0.439	0.318
<i>p</i> -value Hausman test for endogeneity	0.614	0.132	0.164

*Notes:* First-stage estimates for  $\sigma_r$  and  $\delta_r$  are obtained from Equation (4.5), see also column (4) in Table 3.4. Second-stage estimates are based on 396 area-specific wage effects. Robust standard errors are in parentheses. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

The result from column (1) in Table 3.6 can be used to obtain an alternative wage-level elasticity of agglomeration based on the spatial scale at which the agglomeration economies actually operate. For this alternative estimation, all

employment within 80 kilometer is summed, log-transformed, and then used to estimate Equation (3.2). Column (2) shows that this alternative wage-level elasticity of agglomeration is significantly larger than the one found in Table 3.5 (0.017 compared to 0.014) and also closer to the 0.020 found by De la Roca and Puga (2017). Notably, the Hausman test for endogeneity does not reject the use of OLS in columns (1) and (2). This indicates that most of the previously detected endogeneity problems stem from employment on further distances.

Finally, column (3) gives the concentric ring estimates for the wage-growth effects. The coefficients, which are mostly insignificant and small in magnitude, confirm our previous conclusions. The only exception relates to the 40–80 kilometer ring variable, which turns out to be significant at the five percent level. It is difficult to make sense of this result, as it does by no means accord with standard economic intuition. All in all, the concentric ring-based results give us little reason to believe that false assumptions about geographic scope are undermining estimates of the wage-growth elasticity of agglomeration.

In this section we addressed four potentially confounding estimation issues. Two of them – nonlinearities and inaccurate estimates of  $\sigma_r$  and  $\delta_r$  – appear to be of minor concern. Endogeneity of the agglomeration measure and assumptions about the spatial scope of agglomeration economies, on the other hand, have proved to be more pressing problems. This is particularly true for the wage-level elasticity of agglomeration. Based on our analyses, we are confident that both issues share the same roots: employment at more than 10 kilometer distance has a significant effect on wage levels. Hence, to tackle this issue, we should either control for these variables directly, by following a concentric ring-based approach, or address the issue indirectly by using instruments. To allow for comparison with other empirical work, we prefer to derive all results in Section 3.7 by using the more traditional IV-method, rather than the less conventional concentric ring-based strategy.

### 3.7 Heterogeneous effects

All previous estimates in this chapter are based on a representative sample of workers from the Netherlands. Although these analyses are appropriate to examine the importance of spatial sorting in explaining the urban wage premium, we must keep in

mind that these estimates apply to the average worker. As emphasized by Combes et al. (2010b), this average value may conceal large heterogeneities among different types of workers. Uncovering these different effects will foster our understanding of the urban economy and may provide a glance at the mechanisms that drive the urban increasing returns. In this section, we will explore possible heterogeneities across four dimensions: education level, knowledge intensity of the economic sector, age and gender. The results, as reported in Table 3.7 and 3.8, are obtained by estimating Equation (3.5) and corresponding second-stage equations. For reference purposes, we have included the estimates from Equation (3.3) in Appendix B.

Columns (1) to (3) in Table 3.7 present the results for low, medium and highly educated people.<sup>44</sup> The results display a distinct pattern: the wage-level effects of urban areas are strongly increasing in the individual's education level. Compared to a low-educated worker, the wage-agglomeration elasticity is 72 percent larger for medium-educated workers and 163 percent larger for the highly educated. These findings concur with other empirical studies, which find a larger urban wage differential for highly educated workers (Wheeler, 2001; Groot and De Groot, 2014), white-collar workers (Gould, 2007), and people with strong cognitive skills (Bacolod et al., 2009). This education gap in the returns to agglomeration could also explain why high-skilled workers tend to sort themselves into urban areas. The wage-growth effects of urban areas, on the other hand, do not reveal any heterogeneity across education levels: all elasticities are insignificant.

---

<sup>44</sup> Three levels of highest attained education level were distinguished. The low-educated category comprises primary education and the more practical oriented secondary educations (VMBO). Middle educated consists of the more theoretical oriented secondary educations (HAVO and VWO) plus the more practical oriented tertiary educations (MBO). Finally, the highly educated category contains the more theoretical oriented tertiary educations (HBO and University including PhDs).

Table 3.7. Heterogeneous effects by educational attainment and knowledge intensity

Column:	(1)	(2)	(3)	(4)	(5)
First-stage equation:	(3.5)	(3.5)	(3.5)	(3.5)	(3.5)
Subsample:	Low-educated	Medium-educated	Highly educated	Knowledge extensive sectors	Knowledge intensive sectors
Age squared	-0.001*** (2.3e-07)	-0.001*** (2.2e-07)	-0.001*** (3.2e-07)	-0.001*** (2.1e-07)	-0.001*** (3.2e-07)
Wage-level effect of firm size: 10–99 employees	0.007*** (0.001)	0.008*** (0.001)	0.015*** (0.001)	0.007*** (0.001)	0.009*** (0.001)
Wage-level effect of firm size: 100–999 employees	0.016*** (0.001)	0.017*** (0.001)	0.026*** (0.001)	0.016*** (0.001)	0.018*** (0.001)
Wage-level effect of firm size: ≥1000 employees	0.042*** (0.004)	0.028*** (0.003)	0.045*** (0.003)	0.036*** (0.004)	0.028*** (0.003)
Wage-growth effect of firm size: 10–99 employees	0.003*** (0.001)	0.001 (0.001)	-0.000 (0.001)	0.001 (0.001)	-0.003*** (0.001)
Wage-growth effect of firm size: 100–999 employees	0.005*** (0.001)	0.001 (0.001)	-0.002 (0.001)	0.001 (0.001)	-0.003** (0.001)
Wage-growth effect of firm size: ≥1000 employees	-0.011*** (0.003)	0.003 (0.003)	0.011*** (0.002)	-0.006** (0.003)	0.002 (0.002)
Year effects	YES	YES	YES	YES	YES
Wage-level/growth area indicators	YES	YES	YES	YES	YES
Wage-level industry effects	YES	YES	YES	YES	YES
Wage-level worker effects (FE)	YES	YES	YES	YES	YES
Wage-growth industry effects	YES	YES	YES	YES	YES
Wage-growth worker effects	YES	YES	YES	YES	YES
Observations	1,866,829	2,528,101	1,735,161	2,878,025	1,789,866
R <sup>2</sup>	0.050	0.067	0.070	0.051	0.058
Second-stage equations:	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)
<b>Wage-level benefit</b>					
Log employment within 10 km	0.008*** (0.002)	0.013*** (0.002)	0.021*** (0.003)	0.012*** (0.002)	0.016*** (0.003)
IV	YES	YES	YES	YES	YES
p-value Kleibergen-Paap rk LM statistic	0.000	0.000	0.000	0.000	0.000
p-value Hansen J statistic	0.103	0.154	0.080	0.234	0.071
p-value Hausman test for endogeneity	0.139	0.000	0.001	0.010	0.009
<b>Wage-growth benefit</b>					
Log employment within 10 km	-0.0007 (0.0012)	0.0002 (0.0009)	-0.0011 (0.0018)	0.0005 (0.0008)	0.0025 (0.0028)
IV	YES	YES	YES	YES	YES
p-value Kleibergen-Paap rk LM statistic	0.000	0.000	0.000	0.000	0.000
p-value Hansen J statistic	0.086	0.332	0.749	0.637	0.431
p-value Hausman test for endogeneity	0.401	0.082	0.904	0.902	0.346

Notes: Second-stage estimates are based on 396 area-specific wage effects. Industry indicators are based on two-digit NACE. Robust standard errors, which are clustered by worker in the first-stage estimates, are in parentheses. First-stage R<sup>2</sup> is within workers including the absorption of the worker-specific trends. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 3.8. Heterogeneous effects by age groups and gender

Column:	(1)	(2)	(3)	(4)	(5)
First-stage equation:	(3.5)	(3.5)	(3.5)	(3.5)	(3.5)
Subsample:	Aged [18;30]	Aged [30;40]	Aged [40;65]	Male	Female
Age squared	0.001*** (2.7e-07)	-0.001*** (2.8e-07)	-0.002*** (2.0e-07)	-0.001*** (1.7e-07)	-0.001*** (2.5e-07)
Wage-level effect of firm size: 10–99 employees	0.010*** (0.001)	0.010*** (0.001)	0.005*** (0.001)	0.010*** (0.001)	0.008*** (0.001)
Wage-level effect of firm size: 100–999 employees	0.027*** (0.001)	0.018*** (0.001)	0.012*** (0.001)	0.020*** (0.001)	0.019*** (0.001)
Wage-level effect of firm size: ≥1000 employees	0.045*** (0.004)	0.041*** (0.003)	0.023*** (0.003)	0.039*** (0.002)	0.032*** (0.003)
Wage-growth effect of firm size: 10–99 employees	0.002*** (0.001)	0.002*** (0.001)	0.003*** (0.001)	0.001* (0.001)	0.001 (0.001)
Wage-growth effect of firm size: 100–999 employees	0.003** (0.001)	0.001 (0.001)	0.005*** (0.001)	0.001 (0.001)	0.001 (0.001)
Wage-growth effect of firm size: ≥1000 employees	0.011*** (0.004)	0.008*** (0.003)	0.009*** (0.002)	0.003* (0.002)	0.005* (0.003)
Year effects	YES	YES	YES	YES	YES
Wage-level/growth area indicators	YES	YES	YES	YES	YES
Wage-level industry effects	YES	YES	YES	YES	YES
Wage-level worker effects (FE)	YES	YES	YES	YES	YES
Wage-growth industry effects	YES	YES	YES	YES	YES
Wage-growth worker effects	YES	YES	YES	YES	YES
Observations	2,025,533	1,768,036	2,336,522	4,069,173	2,060,918
$R^2$	0.117	0.065	0.044	0.063	0.053
Second-stage equations:	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)	(3.2 & 3.4)
<b>Wage-level benefit</b>					
Log employment within 10 km	0.015*** (0.002)	0.016*** (0.002)	0.008*** (0.002)	0.014*** (0.002)	0.014*** (0.002)
IV	YES	YES	YES	YES	YES
$p$ -value Kleibergen-Paap rk LM statistic	0.000	0.000	0.000	0.000	0.000
$p$ -value Hansen J statistic	0.026	0.213	0.347	0.527	0.276
$p$ -value Hausman test for endogeneity	0.000	0.002	0.781	0.001	0.011
<b>Wage-growth benefit</b>					
Log employment within 10 km	0.0025** (0.0010)	0.0003 (0.0011)	0.0009 (0.0010)	0.0008 (0.0008)	-0.0008 (0.0011)
IV	YES	YES	YES	YES	YES
$p$ -value Kleibergen-Paap rk LM statistic	0.000	0.000	0.000	0.000	0.000
$p$ -value Hansen J statistic	0.092	0.547	0.506	0.989	0.122
$p$ -value Hausman test for endogeneity	0.988	0.829	0.845	0.695	0.149

Notes: Second-stage estimates are based on 396 area-specific wage effects. Industry indicators are based on two-digit NACE. Robust standard errors, which are clustered by worker in the first-stage estimates, are in parentheses. First-stage  $R^2$  is within workers including the absorption of the worker-specific trends. \*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

The second explored heterogeneity dimension is the knowledge intensity of the economic sectors. After all, if wage-growth benefits of urban areas are driven by the diffusion of knowledge and ideas, we would expect these effects to be more pronounced in knowledge intensive sectors. Using the Eurostat (n.d.) indicators on technological and knowledge intensity of economic sectors, we have divided nearly all two-digit NACE industries into two broader categories, i.e. knowledge intensive and extensive sectors.<sup>45</sup> Then, we split the sample along each of these two categories and remove workers who have moved between them. The results from this analysis, presented in columns (4) and (5) in Table 3.7, indicate that the wage-level effects of urban areas are larger for workers in knowledge intensive sectors. However, the wage-growth elasticities of agglomeration are insignificant for both subsamples, despite the relatively large point estimate for the knowledge intensive sectors.

In Table 3.8 workers are divided into subsamples on the basis of age and gender.<sup>46</sup> The results show no significant difference in the benefits of urban areas between men and women. This is in line with the findings of D'Costa and Overman (2014), although they differ from the estimates of De la Roca and Puga (2017). The latter study finds an almost two times higher medium-term elasticity for males compared to females. It should be noted here, however, that the observed gender gap in the study of De la Roca and Puga (2017) might be the result of discrepancies in individual-specific returns to experience rather than differences in wage-growth effects of urban areas. This is corroborated by the results in Appendix B, which find significantly larger wage-growth benefits for males compared to females when we do not account for spatial sorting on wage-growth determining characteristics.

The heterogeneity analysis for age groups conveys significantly smaller wage-level effects for older people ( $\geq 40$  years) compared to younger people. This may be due to the lower labor mobility of older workers (see also Table 3.2), which constrains their ability to gain from agglomeration economies via mechanisms related to matching. Even more interesting, we find that wage-growth effects of urban areas are

---

<sup>45</sup> Manufacturing industries are categorized according to R&D intensity, while service industries are classified on the basis of the share of tertiary educated persons. A few two-digit NACE industries could not be assigned to one of the knowledge intensity categories: 6, 8, 9, 35–43.

<sup>46</sup> Since the age of workers increases over time, it is not obvious to which age group a worker should be assigned. In this chapter, we choose to assign workers to age groups on the basis of their age in the year 2010, the middle of our sample period. This implies that workers are, at a particular point in time, at most four years older or younger than the cutoff values of the age group to which they were assigned.



significantly positive for workers below the age of 30 years. Additionally, the point estimate is also relatively large in magnitude.<sup>47</sup> These results, which are in line with D'Costa and Overman (2014), indicate that younger workers are more receptive to wage-growth benefits of urban areas, even after controlling for individual-specific returns to experience. These outcomes support the view that urban areas stimulate the accumulation of human capital.

### 3.8 Conclusion

This research has focused on two important issues when estimating the benefits of working in urban areas. First, big cities are generally more successful than smaller cities in attracting the most talented people and the most productive firms. Hence, the urban wage premium may not reflect a causal link between wages and city size, but rather represent an omitted variable bias in the form of spatial sorting. Second, agglomeration economies can materialize through both wage-level and wage-growth effects.

Using panel data on individual earnings, we show that the urban wage premium consists of both wage-level and wage-growth effects. However, after introducing wage-level and wage-growth controls for worker and firm characteristics, the magnitude of the wage-agglomeration elasticities declines substantially. In fact, the wage-growth effect of urban areas becomes statistically insignificant at conventional levels. This result challenges the conventional view that human capital spillovers are responsible for the existence of an urban wage-growth premium. Instead, we conclude that the positive association between city size and wage growth is to a large extent driven by spatial sorting of workers and firms.

We have also examined heterogeneities in the wage-agglomeration relationship among different types of workers. The analyses indicate that wage-level effects are important to all workers and particularly pronounced for highly educated workers and those employed in knowledge intensive economic sectors. Older

---

<sup>47</sup> We have tried to estimate the wage-growth elasticity of agglomeration for even finer-grained age groups, but the low number of workers within each age group raises problems related to identification and statistical significance. Nevertheless, the pattern that emerges from this exercise supports the view that younger people are more receptive to wage-growth benefits of urban areas. For workers between the age of 18 and 24 we obtain a point estimate of 0.0034 with a standard error of 0.0020, whereas the point estimate is 0.0023 with a standard error of 0.0011 for workers between 24 and 30.

workers, on the other hand, benefit significantly less from wage-level effects of agglomeration. Estimates of the urban wage-growth effect, on the other hand, are mostly insignificant. Young workers are a notable exception.

As with any empirical research, the external validity of our results to other contexts is not without limitations. Most importantly, the Netherlands is somewhat special due to its socioeconomic equality and polycentric urban structure, having multiple middle-sized though well-connected cities. These aspects may partly explain why the wage-agglomeration elasticities reported in this chapter are small compared to the international literature. Although we acknowledge that these limitations are real, we are nevertheless confident that the main message from this chapter is relevant to other countries: sorting of high potentials matters for identifying the urban wage-growth premium.