Agglomeration and Human Capital Spillovers

Verstraten, P.W.A.H.

2019

document version
Publisher's PDF, also known as Version of record

Chapter 2

The spatial scope of agglomeration economies

2.1 Introduction[3]

Although the agglomeration literature has reached no consensus on the maximum spatial extent of agglomeration economies[4], most studies do agree on how these agglomeration spillovers decay across geographic space. In general, the literature concludes that the decay effect of agglomeration spillovers is a monotonic function of distance (Rice et al., 2006; Arzaghi and Henderson, 2008; Di Addario and Patacchini, 2008; Rosenthal and Strange, 2008; Koster, 2013; Ahlfeldt et al., 2012). For some studies, this consistent finding in the literature has even been a reason to assume a priori that agglomeration spillovers decay monotonically across space (e.g., Rice et al., 2006; Koster, 2013). Market-potential functions based on Harris (1954), which are built into many economic geography models, also rely on this particular assumption (e.g., Davis and Weinstein, 2003; Head and Mayer, 2004; Hanson, 2005).

This chapter argues that the spatial scope of agglomeration economies is much more complex than is often assumed. This is for three main reasons. First, it is a

[3] Apart from minor changes, this chapter was published as: Verstraten, P., Verweij, G. and Zwaneveld, P.J. (2019). Complexities in the spatial scope of agglomeration economies. *Journal of Regional Science*, 59(1), 29–55. Reprinted with permission of the co-authors and Wiley-Blackwell. Available at https://doi.org/10.1111/jors.12391.

[4] Estimates range from 40–80 kilometer (Rosenthal and Strange, 2008) to only a few kilometers (Arzaghi and Henderson, 2008; Ahlfeldt et al., 2012), and everything in between (e.g., Di Addario and Patacchini, 2008; Koster, 2013).

misconception to think that the net effect of all different agglomeration spillovers must be described by a monotonically declining distance decay function, even though there are sound arguments to think that individual spillovers do decay monotonically.[5] The spatial decay of the net agglomeration spillover may not be a monotonic function of distance because individual spillovers operate over different distances and can vary in terms of intensity and direction of the effect, i.e. positive or negative (Harvey, 1973; Li and Brown, 1980). For instance, positive agglomeration spillovers related to knowledge diffusion and labor market pooling are expected to have a short spatial extent, whereas the benefits of input sharing are expected to operate over larger distances (Ellison et al., 2010). Urban congestion, such as pollution and high traffic volumes, on the other hand, represents a negative spillover effect with a relatively short spatial scope (Zhou and Levy, 2007). The net effect of all these spillovers may be characterized by a wide variety of spatial decay functions, which are not necessarily monotonic in distance.

Second, empirical research on the spatial scope of agglomeration economies relies on rents or wages. However, from a theoretical perspective, there is no reason to assume that the spatial scope of spillovers capitalizing into rents is similar to those that capitalize into wages. Hence, it is difficult to draw conclusions about the spatial scope of agglomeration economies from studies that analyze only either of these two prices. Third, there is not any particular reason to assume that the spatial scope of agglomeration economies applies to all regions equally. This assumption, however, is explicitly built into many empirical models.

This chapter provides empirical evidence for this complexity by showing that agglomeration on short distances (<5 kilometer) does not significantly affect wages in the Netherlands, whereas it has a significant and positive effect on medium distances (5–10 kilometer). This effect attenuates across geographic space and becomes insignificant after 40–80 kilometer. These results are in sharp contrast with the findings of Koster (2013), who concludes that agglomeration within five kilometer is strongly related to rents of commercial property in the Netherlands, whereas they are unrelated on longer distances. Hence, this apparent contradiction suggests that

---

[5] Marshall (1890) was the first to distinguish between different sources of agglomeration economies. Duranton and Puga (2004) provide an extensive theoretical overview of these individual mechanisms, and Puga (2010) reviews empirical evidence.

spillovers capitalizing into rents and wages differ in terms of their spatial scope. Furthermore, our results indicate that agglomeration within five kilometer is not irrelevant to the wage formation. In fact, the data show that only highly urbanized areas benefit from agglomeration on longer distances. This implies that not all regions benefit equally from economies of agglomeration.[6]

The finding that agglomeration economies stretch across a relatively large distance also raises questions about the role of foreign agglomerations in the domestic wage formation. After all, the Netherlands is a small country, part of the European Single Market, and shares a common language with the Northern part of Belgium. Therefore, in order to assess the influence of foreign agglomerations, we have constructed a unique dataset containing information on the current spatial distribution of employment and historical population censuses for both Belgium and Germany. Despite the openness of the Dutch economy, our analysis shows that foreign economic mass does not affect wages in the Netherlands. This result is consistent with the bulk of the literature, which finds substantial border barriers (e.g., Brakman et al., 2002).

In order to reveal the complexities underlying the spatial decay function, this chapter employs panel data on individual wages with a high level of geographic detail: Dutch postal codes with a mean area of only nine km^2. The use of this dataset has two advantages compared to earlier work. First, the spatial richness of the dataset enables us to construct narrow concentric ring variables, which is an important prerequisite to disentangle the effects of agglomeration on very short distances from those on longer distances. Similar studies on the spatial scope of agglomeration economies have relied on spatial units that are much larger than the Dutch postal code: e.g., 6,522 km^2 (Rosenthal and Strange, 2008), 1,394 km^2 (Rice et al., 2006) and 889 km^2 (Di Addario and Patacchini, 2008). It is evident that this lack of spatial detail in most other studies makes it difficult to identify the effects of agglomeration on various short distances.

---

[6] This chapter fits into a large body of research examining agglomeration economies in the Netherlands on wages (Groot and De Groot, 2014; Groot et al., 2014), commercial rents (Koster, 2013; Koster et al., 2014), innovation intensity (Van Oort, 2002; Van der Panne, 2004), employment growth (Van Soest et al., 2006; Van Oort, 2007), GDP per hour worked (Broersma and Oosterhaven, 2009) and firm formation (Van Oort and Atzema, 2004).

The second key advantage relates to the longitudinal nature of the wage data. By following workers over time, we are able to control for both observed and unobserved worker characteristics. This is crucial for the identification of agglomeration economies since it is well established that a considerable part of the urban wage premium is driven by sorting of high-skilled workers into urban areas (Combes et al., 2008). In contrast, other studies on the spatial scope of agglomeration economies (e.g., Rosenthal and Strange, 2008; Di Addario and Patacchini, 2008; Rice et al., 2006), have controlled for observed worker characteristics only. Therefore, these studies run the risk of omitting important unobserved differences in labor quality.

The remainder of this chapter is structured as follows. Section 2.2 discusses the micro-econometric model, whereas Section 2.3 describes the wage data and the process of constructing concentric ring variables. In Section 2.4 we report the results regarding the spatial scope of agglomeration economies, and in Section 2.5 we provide checks for robustness. Section 2.6 examines the magnitude of the wage-agglomeration relationship. Section 2.7 concludes.

2.2 Methodology

In order to analyze the relationship between wages and agglomeration, this chapter employs a two-stage estimation approach as proposed by Combes et al. (2008). In the first stage of this approach, we estimate area fixed effects using a Mincerian wage equation. Then, in the second stage, we explain these area fixed effects using concentric ring variables that measure the employment levels at various distances. This concentric ring-based approach was first used by Rosenthal and Strange (2003).

An important benefit of this two-stage estimation approach is the elegant solution of the dependent disturbances within the spatial units. Non-independent disturbances may arise because observations sharing the same geographic space might influence each other and/or might be subject to the same local shocks. Neglecting this dependence often leads to downward biased standard errors (Moulton, 1990).The standard solution of calculating cluster robust standard errors assumes nesting of the workers within the same spatial cluster. However, our study relies on workers who change their work location, making the default use of cluster robust standard errors not applicable. We will further elaborate on this two-stage approach in the remainder of this section.

2.2.1 Two-stage estimation approach

A profit-maximizing and perfectly competitive firm in area r, industry k and year t pays wages equal to the marginal product of labor. Hence, following Combes et al. (2008), the hourly wage of worker i in year t can be described as:

$$ln\, w_{i,t} = \beta X_{i,t} + \sigma_r R_{r(i,t)} + \sigma_k K_{k(i,t)} + \sigma_t T_{t(i,t)} + \varepsilon_{i,t}, \tag{2.1}$$

where the log-transformed hourly wage $w_{i,t}$ is explained by a vector of worker characteristics $X_{i,t}$ and productivity effects unrelated to worker characteristics. The latter consists of a vector of area-dummies $R_{r(i,t)}$ indicating the individual's place of work, a vector of industry-dummies $K_{k(i,t)}$, and a vector of year-dummies $T_{t(i,t)}$. The vectors β, σ_r, σ_k and σ_t contain the parameters to be estimated, and $\varepsilon_{i,t}$ is a random error term.[7]

It is commonly acknowledged in the agglomeration literature that the urban wage premium might be driven by the sorting of high-skilled workers into urban areas. This implies that $\mathrm{cov}\left(X_{i,t}, R_{r(i,t)}\right) \neq 0$. In order to identify the area-specific productivity effects under the presence of sorting, it is necessary to include variables that capture all relevant worker characteristics $X_{i,t}$. To this end, all studies to date that examine the geographic scope of agglomeration economies using wage data (e.g., Rosenthal and Strange, 2008; Di Addario and Patacchini, 2008; Rice et al., 2006), have employed observed characteristics to control for worker heterogeneity. However, these studies run the risk of having omitted some worker characteristics that correlate with the area-specific productivity effects.

[7] In this specification, we ignore potential interactions between the area-, industry- and time-specific productivity effects. This is for practical reasons, as estimating the full interaction set $(\sigma_{r,k,t} R_{r(i,t)} K_{k(i,t)} T_{t(i,t)})$ would require the inclusion of roughly 2.3 million fixed effects, besides the 2.3 million worker fixed effects. Equation (2.1), on the other hand, would require only 3,800 additional fixed effects. Since our main interest is ultimately in the effect of agglomeration on wages, it is not strictly necessary to include an area-year interaction because the spatial distribution of economic mass hardly varies over time (on the postal code-level, the correlation coefficient between the number of jobs in 2006 and 2014 equals 0.982).

In contrast to these earlier works, our study relies on worker fixed effects to control for all time-invariant worker characteristics. The age of workers and its square are used as a proxy for worker experience. The regression equation then becomes:

$$\ln w_{i,t} = \sigma_i + \beta_1 \widetilde{age}_{i,t} + \beta_2 \widetilde{age}_{i,t}^2 + \sigma_r R_{r(i,t)} + \sigma_k K_{k(i,t)} + \sigma_t T_{t(i,t)} + \varepsilon_{i,t}, \qquad (2.2)$$

where σ_i is a worker fixed effect and $\widetilde{age}_{i,t}$ denotes the age of a worker. The worker's age is centered around its industry-average to account for the fact that some industries tend to hire older/younger workers (Combes et al., 2008). The squared term captures any concave effects of experience on wages.

It should be noted that, with this specification, the area-specific effects on wages are assumed to be static. This means that the model ignores potential area-specific wage-growth effects, such as dynamic agglomeration economies. Although these dynamic effects can potentially bias the estimates of the static effect, we know from De la Roca and Puga (2017) that standard worker fixed effects estimates of the static gains from agglomeration are, under reasonable circumstances, insensitive to the existence of dynamic effects. Given these considerations, we will rely on the standard fixed effects model of Equation (2.2).

The area fixed effect estimates that are obtained from Equation (2.2) reflect spatial wage differences. The equation below describes how these wage differences are the result of both positive and negative agglomeration spillovers at various distances:

$$\sigma_r = \phi \sum_e E_e \mathrm{B}(D_{r,e}) - \omega \sum_e E_e \mathrm{C}(D_{r,e}) + \upsilon_r, \qquad (2.3)$$

where σ_r is the area fixed effect parameter from Equation (2.2), and E_e denotes total employment at establishment e, which we use as a measure of agglomeration.[8] $\mathrm{B}(D_{r,e})$ and $\mathrm{C}(D_{r,e})$ represent the distance decay functions of the positive and negative agglomeration spillovers, respectively. These distance decay functions provide

[8] Much debate exists in the literature about whether agglomeration economies arise from the concentration of industries (localization) or from the overall size of the market (urbanization). In this chapter, our main interest lies in the effect of urbanization, measured in terms of total employment. This is in line with most studies that examine the spatial extent of agglomeration economies (e.g., Rice et al., 2006; Rosenthal and Strange, 2008; Di Addario and Patacchini, 2008; Koster, 2013; Rice et al., 2006).

weights to employment at various distances and, without loss of generality, can take any value between zero and one, depending on the straight line distance between area r and establishment e ($D_{r,e}$). The parameters ϕ and ω represent the wage effect of the spatially weighted agglomeration measures, and υ_r is a random error term.

Estimating Equation (2.3), however, is not possible because both the positive and negative agglomeration spillover stem from the same source (E_e), which makes it virtually impossible to disentangle these two effects. It is for this reason that this chapter, as the other studies in this field of research, estimates the net effect of agglomeration on wages:

$$\sigma_r = \gamma \sum_e E_e \mathrm{N}\big(D_{r,e}\big) + \upsilon_r. \tag{2.4}$$

Even when the individual agglomeration spillovers, $\mathrm{B}\big(D_{r,e}\big)$ and $\mathrm{C}\big(D_{r,e}\big)$, decay monotonically across space, there is no ex-ante reason to expect that the net effect of these spillovers decays monotonically as well. In fact, the decay function of the net effect can take a wide variety of (non-monotonic) functional forms. Figure 2.1 shows one set of the possible spatial decay functions.

Figure 2.1. The net effect of two monotone distance decay functions

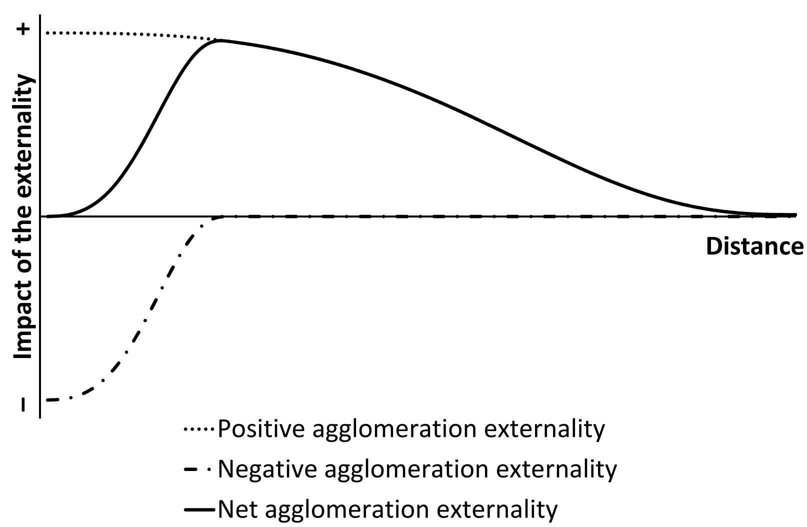

Notes: This figure shows how a non-monotonically declining distance decay effect might be the net outcome of two monotone distance decay functions with a positive and a negative effect. Of course many other functional forms are possible. This figure is similar to the one of Li and Brown (1980).

In order to estimate Equation (2.4), we construct a set of concentric ring variables measuring total employment at various distance intervals, e.g., within five kilometer, between five and 10 kilometer, etc. This flexible strategy is preferred over strategies that employ a pre-defined monotonically declining decay function (e.g., Rice et al., 2006; Koster, 2013) because of the aforementioned reason that the net effect of agglomeration might change non-monotonically across space. In line with the study of Rosenthal and Strange (2003), the regression equation of the second stage then becomes:

$$\sigma_r = \sum_{D_d} \gamma_d \sum_{D_{r,e} \in D_d} E_e + \upsilon_r. \tag{2.5}$$

The first summation is over all concentric rings at various distance intervals D_d. The second summation term aggregates all employment that falls within that particular distance interval. The estimated parameters of the ring variables (γ_d) give the percentage wage effect of an additional unit of employment within a particular distance interval.[9] A numerical comparison of these parameters provides information on how the net benefit of agglomeration differs across geographic space.[10]

When determining the width of the distance intervals, we encounter practical limitations. Although the construction of very narrow distance intervals will, in theory, render a detailed distance decay pattern, it will also lead to serious multicollinearity problems. In particular, this problem tends to become more severe as the distance from area r gets larger. Therefore, in order to avoid problems of multicollinearity, the distance intervals must be somewhat wider on longer distances

[9] Note that our empirical model equally weighs all employment at a particular distance, regardless of the spatial concentration of employment. In order to assess whether spatially concentrated employment delivers the same productivity benefits as spatially diffuse employment, we should incorporate a measure of spatial concentration into the model. It is, however, unclear how this measure can be combined with a concentric ring approach because it should not only take into account spatial concentration within rings but also between rings.

[10] Throughout this chapter we use standard OLS and IV regressions for the second-stage estimation, although this will generally lead to biased and inefficient estimates (Combes et al., 2008). The size of this bias and inefficiency depends on the standard error of the estimated area fixed effects in the first stage. We have re-estimated the model with a feasible generalized least squares (FGLS) estimator (Gobillon, 2004). The estimations provided results comparable to the standard estimation strategy. The difference in estimated parameters and standard errors was generally below 10 percent. We therefore conclude that the influence of estimation errors of the area fixed effects from the first stage can be neglected during the second stage.

than on shorter distances. We use the following set of cutoff values for our distance intervals: 5, 10, 20, 40, 80, and 120 kilometer.[11]

2.2.2 Instrumental variable approach

Finally, a word on one of the classical problems in the agglomeration literature: endogeneity of the agglomeration measure. This issue of endogeneity means that the estimated relationship between agglomeration and wages might be driven by omitted variables, such as (non-)human local endowments, and/or reverse causality. To tackle this endogeneity problem, the literature has suggested several approaches; see Rosenthal and Strange (2004) and Combes et al. (2010a) for an extensive discussion.

In order to address endogeneity issues, this chapter applies the instrumental variable (IV) approach with two sets of instruments. First, we compute concentric ring variables that measure historical (year 1840) population counts. This set of variables will be used as an instrument for the concentric ring variables that measure current employment. The assumption underlying this IV is that (non-)human local endowments that have influenced the spatial distribution of population until the mid-19th century, are no longer important for productivity in a modern, 21st century economy, except through their influence on current employment. Historical population censuses are a relevant IV because the spatial distribution of population is strongly autocorrelated over time, possibly due to path-dependency caused by self-reinforcing spillovers from agglomeration (Bleakley and Lin, 2012).

The second set of instrumental variables consists of concentric ring variables that measure the total number of railway stations in 1870, which is similar to the instrument used by Koster (2013). This instrument is correlated to current employment levels because the opening of railway stations during the 19th century drastically increased the area's accessibility and therefore triggered the formation of urban areas. Nowadays, however, these railway stations are only one of the many links in the infrastructure network. Hence, railway stations that have opened before

---

[11] Compared to earlier studies that employ concentric ring variables to explain the urban wage premium, we use a rather narrow and comprehensive set of concentric ring variables. For example, Rosenthal and Strange (2008) use cutoff values 8, 40, 80 and 160 kilometer (they use terrestrial miles as their unit of length, which corresponds to cutoff values of 5, 25, 50, and 100 miles), whereas Di Addario and Patacchini (2008) choose 4, 8, 12 and 16 kilometer. In Section 2.5.1 of this chapter, we experiment with an even more narrow set of concentric ring variables.

1870 are not expected to influence labor productivity today. In fact, almost half of these stations are no longer operational.

2.3 Data description

2.3.1 Microdata and summary statistics

Our empirical model requires three key datasets. First, we use wage data containing individual information for all employees in the Netherlands[12] on pretax wages and other financial rewards, hours worked, date of birth, sectoral classification of the employer (two-digit NACE), place of work at the postal code level, and job-type. This dataset is based on own calculations using non-public microdata from Statistics Netherlands (CBS): fiscal data (Polisadministratie), census data (Sociaal Statistisch Bestand), and firm data (Algemeen Bedrijven Register). Based on this information we construct an unbalanced panel (2006–2014) with yearly observations for each individual.

The wage data do not only contain regular pre-tax wages, but also overtime payments, paid holidays, bonuses, thirteenth salaries and company cars. 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. Due to limitations of the dataset, this calculation of total hourly labor costs is prone to measurement errors when a worker has not been employed for the full year at the same employer. For this reason we drop observations that are not based on a complete year of work at the same employer. We present a robustness analysis in Section 2.5.1 using only regular hourly pre-tax wages. This alternative wage definition permits the inclusion of these dropped observations.

[12] The dataset does not contain self-employed workers, who comprise 10–15 percent of the Dutch labor force.

Table 2.1. Summary statistics of the longitudinal wage data

	2006	2010	2014
Number of workers	1,456,067	1,192,499	1,115,038
Hourly wages in euro's (price level 2006)			
Mean (standard deviation)	20.6 (11.5)	21.6 (12.4)	21.4 (12.9)
1st percentile	8.3	8.6	8.4
Median	17.8	18.4	18.0
99th percentile	67.2	71.8	74.4
Age			
Mean (standard deviation)	40.0 (10.7)	40.9 (10.9)	41.8 (11.1)
1st percentile	20.5	20.6	21.0
Median	39.3	40.8	42.2
99th percentile	61.4	62.7	63.3
Industrial composition (in percentages)			
Manufacturing	23.2	22.1	22.7
Construction	11.3	10.4	8.1
Logistics	7.2	7.6	7.6
Wholesale	15.1	15.4	16.2
Retail	7.0	7.4	6.8
Consumer services	3.2	3.5	3.5
Hospitality industry	4.2	4.7	5.0
ICT	4.7	5.4	6.3
Financial services	3.6	3.0	3.0
Business services	20.4	20.7	21.0

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.[13] Jobs in agriculture and the fishing industry are excluded from the sample because these sectors are strongly linked to the location of natural resources.

[13] The WSW is a Dutch law aimed to foster the employment of persons with disabilities.

Also the public sectors are excluded because it is improbable that these sectors meet our underlying assumption that employers are perfectly competitive and profit-maximizing. Jobs provided by a firm with more than one establishment, could not be assigned geographically and had to be removed from the sample. 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. 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, over one million observations per year remain. Table 2.1 summarizes the data remaining for estimation in the years 2006, 2010 and 2014.

2.3.2 Spatial variables

The second key dataset contains information on the spatial distribution of both current employment and historical population in the Netherlands and neighboring countries.[14] We constructed this dataset by combining several data sources, which are listed in Appendix A.1. As can be seen from Table A.1, our spatial unit of analysis, the four-digit postal code in the Netherlands, is rather small with an average area of only 8.86 km^2. This high level of spatial detail allows us to examine the decay pattern of agglomeration economies on short as well as long distances. The third key dataset contains coordinates of all historic railway stations that have been operational during the year 1870. This amounts to a total of 235 railway stations, of which 106 stations were no longer operational by the year 2006 (the first year of our wage data).

Using GIS tools, we construct concentric ring variables that measure the current employment levels and historical population counts within particular distance intervals. First, we draw concentric rings around the geographic centroid of the postal codes and then calculate for each geographic unit in our sample which percentage of the area falls within the concentric ring. As previously discussed, we choose a total of seven concentric rings with a respective radius of 5, 10, 20, 40, 80, and 120 kilometer. Then we assume that, within geographic units, employment and population are homogeneously distributed across space, which enables us to approximate current employment and historical population levels within each concentric ring. Finally, we

[14] Because we estimate area fixed effects for the period 2006–2014, we take the middle year (2010) as our measure of current employment.

first-difference the concentric ring variables to obtain total employment and population within distance intervals. In a similar way, we have calculated the total number of historical railway stations within distance intervals. The domestic and foreign concentric ring variables measuring current employment are graphically presented in Appendix A.4.

Table 2.2 contains a correlation matrix of the ring variables that measure employment, historical population (in brackets), and the total number of historical railway stations (in parentheses). This table shows that, although the ring variables are mutually correlated, the correlation is limited due to the increasing distance intervals. Hence, it appears that concerns regarding multicollinearity of the exogenous regressors will be limited.

Table 2.2. Correlation matrix of the concentric ring variables

	0–5 km	5–10 km	10–20 km	20–40 km	40–80 km	80–120 km
0–5 km	1.000 [1.000] (1.000)					
5–10 km	0.660 [0.264] (0.399)	1.000 [1.000] (1.000)				
10–20 km	0.376 [0.184] (0.232)	0.633 [0.357] (0.294)	1.000 [1.000] (1.000)			
20–40 km	0.247 [0.148] (−0.243)	0.362 [0.244] (−0.228)	0.591 [0.446] (0.019)	1.000 [1.000] (1.000)		
40–80 km	0.277 [0.137] (−0.022)	0.390 [0.218] (0.017)	0.520 [0.374] (0.125)	0.702 [0.615] (0.412)	1.000 [1.000] (1.000)	
80–120 km	−0.070 [−0.056] (−0.086)	−0.116 [−0.084] (−0.052)	−0.136 [−0.104] (0.014)	−0.058 [−0.005] (0.371)	0.194 [0.258] (0.566)	1.000 [1.000] (1.000)

Note: The table shows the correlation coefficients of the concentric ring variables measuring domestic employment, historical population (in brackets) and the number of historical railway stations (in parentheses).

2.4 Results

2.4.1 The spatial decay effect of agglomeration on wages

Since the outcomes of the first-stage regressions are not directly relevant for this chapter, this section presents the outcomes of the second-stage regressions only.[15] Column (2) of Table 2.3 shows the results of the second-stage IV estimates with the full set of concentric ring variables. We conclude that employment within five kilometer distance does not significantly affect wages. Between five and 10 kilometer we observe a relatively strong effect of employment on wages. More specifically, wages increase by 0.77 percent when employment between five and 10 kilometer distance increases by 100,000.[16] The net benefits of agglomeration attenuate rapidly after 10 kilometer, although the effect remains significant until at least 40–80 kilometer. We find no significant effect of employment on wages after 80 kilometer. A graphical representation of these results can be seen in Appendix A.5. Similar results can be obtained when the concentric rings are based on population rather than employment levels, see Appendix A.2.

When comparing column (1) and (2) of Table 2.3, we see that both the OLS and IV estimates are very similar, apart from the first concentric ring variable. This suggests that endogeneity is not a big concern at longer distances, whereas it does play a role at short distances. And indeed, according to the endogeneity test, the data reject the null hypothesis that the first concentric ring variable can be treated as an exogenous regressor.[17] Furthermore, the Hansen *J* overidentification test cannot reject the null hypothesis that the instruments are uncorrelated with the error term, and the Kleibergen-Paap rk Wald *F*-statistic on the weak instruments identification test exceeds all thresholds proposed by Stock and Yogo (2005). See Appendix A.3 for the

[15] Results of the first-stage regression are available upon request. The parameters of the age variables are as one would expect: significant positive for the linear variable and significant negative for the quadratic variable.

[16] Note that employment within concentric rings is expressed as the total number of jobs in millions. A change of 100,000 jobs between five and 10 kilometer is equal to a 0.8 standard deviation.

[17] This result is likely to be driven by a classical omitted variable bias. For instance, (unobserved) local endowments have a positive effect on productivity and, as a consequence, also a positive effect on nearby agglomeration size. This explains why the coefficient of the first ring decreases when we apply instruments. It makes intuitive sense that local endowments have no strong effect on agglomeration further away, which explains why the coefficients of the other rings hardly change when using instruments.

first-stage IV regression results, including a Sanderson-Windmeijer F-statistic for each endogenous variable. Since the endogeneity test has rejected the use of OLS, we will primarily focus on IV regressions in the remainder of this chapter.

Table 2.3. The spatial scope of agglomeration economies

Column: Specification:	(1) All rings	(2) All rings	(3) Five rings	(4) Four rings	(5) Three rings	(6) Two rings	(7) One ring
Employment 0 to 5 km	0.048*** (0.017)	0.022 (0.019)	0.021 (0.019)	0.022 (0.019)	0.020 (0.020)	0.020 (0.019)	0.141*** (0.014)
Employment 5 to 10 km	0.061*** (0.015)	0.077*** (0.020)	0.078*** (0.020)	0.080*** (0.019)	0.078*** (0.020)	0.144*** (0.016)	
Employment 10 to 20 km	0.023*** (0.007)	0.024*** (0.010)	0.023** (0.009)	0.025*** (0.009)	0.048*** (0.008)		
Employment 20 to 40 km	0.012*** (0.003)	0.011*** (0.003)	0.011*** (0.003)	0.016*** (0.003)			
Employment 40 to 80 km	0.004*** (0.001)	0.003** (0.002)	0.004** (0.002)				
Employment 80 to 120 km	0.000 (0.002)	0.001 (0.002)					
IV	NO	YES	YES	YES	YES	YES	YES
F-statistic weak identification test		97.400	118.338	146.381	193.032	266.197	900.577
p-value Hansen J statistic		0.287	0.403	0.596	0.154	0.606	0.537
Max VIF [Mean VIF]	2.59 [2.11]	3.94 [2.75]	3.91 [2.90]	3.83 [2.76]	3.81 [2.76]	2.36 [2.36]	
R^2	0.054	–	–	–	–	–	–

Notes: 3,722 observations (area fixed effects obtained from the first-stage equation). Robust standard errors are in parentheses. Employment is expressed as the total number of jobs in millions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

It is revealing to see what happens when the model contains only a limited set of concentric ring variables and thus ignores employment at longer distances. Therefore, Table 2.3 reports a total of six IV regressions; each containing one additional concentric ring variable. By looking only at column (7), we may conclude that employment within five kilometer affects wages positively and significantly. However, this estimate suffers from an omitted variable bias, even despite the fact that we have included instrumental variables. To see how this works, we must take a look at Table 2.2. This table shows that each concentric ring variable correlates with adjacent ring variables, which will be a source of omitted variable bias if we do not include the full set of ring variables. The instrumental variables are inappropriate to tackle this kind of endogeneity since they are themselves also correlated with adjacent concentric rings and therefore with the error term. The most efficient way to deal with this omitted variable bias is to add more concentric ring variables to the model. Column (6) shows that the coefficient of the first concentric ring becomes insignificant when an adjacent concentric ring is included. When we continue this process of adding additional ring variables to the model, the results remain quite stable.

The results indicate a relatively wide spatial scope of agglomeration economies, stretching across 40–80 kilometer straight line distance. This is in line with the findings of Rice et al. (2006) and Rosenthal and Strange (2008). Di Addario and Patacchini (2008) find a smaller spatial scope. We can only speculate which individual agglomeration spillover operates over such large distances, although there is some evidence that benefits of input sharing exhibit a large spatial scope (Ellison et al., 2010). Since mechanisms related to matching and learning are expected to operate over shorter distances, they may explain the relatively large wage effect of employment between five and 10 kilometer.

We offer three explanations, which are not mutually exclusive, for our finding that wages and agglomeration are not significantly related on short distances. First, it is possible that straight line geographic distance fails to be a good predictor of proximity on short distances. After all, the bivariate correlation between travel time and straight line distance decreases as the straight line distances become smaller (Phibbs and Luft, 1995). This potential measurement error could bias the estimates downwards. Testing this hypothesis is complicated because of two main reasons. First of all, an adequate measure of generalized travel costs is not readily available because such a measure needs to incorporate both travel times and costs for all means of transportation (e.g., car, public transport, bicycle, walking). A second reason is that the use of generalized travel costs will inevitably lead to problems of reverse causality because the more productive and dense areas tend to have more and better transport connections than less productive areas.

Second, as we have argued in Section 2.2.1, this non-monotone spatial decay effect can emerge when negative agglomeration spillovers (e.g., traffic congestion and pollution) are substantial on short distances, and decay more rapidly than positive spillovers. Benefits of agglomeration will then be offset by negative spillovers on short distances, whereas the benefits dominate on longer distances. This explanation is in line with Li and Brown (1980), who find that the house price-effect of proximity to commercial establishments is not a monotonic function of distance.

Third, agglomeration economies within five kilometer might capitalize into rents rather than wages. Interestingly, Koster (2013) finds that agglomeration in the Netherlands is strongly related to rents of commercial property on short distances (<5 kilometer) and unrelated on longer distances. Although our results appear to be the exact opposite of the findings of Koster, it should be noted that, from a theoretical perspective, it is unclear whether benefits of agglomeration will capitalize into rents or wages. Hence, the results of both studies can be reconciled if the spatial scope of spillovers that capitalize into rents differs from the spillovers that capitalize into wages. The outcome may also be the result of a bargaining game over the gains from agglomeration: landowners may have more bargaining power when land is scarce, i.e. on short distances to urban areas, whereas workers may have strong bargaining power when land is abundant, i.e. on longer distances to urban areas.

Although the results of this chapter appear, at first sight, to contradict earlier studies, reconciliation with their results is straightforward. For instance, the apparent contradiction with earlier studies that find a strong and positive relationship on short distances (e.g., Arzaghi and Henderson, 2008; Rosenthal and Strange, 2008; Di Addario and Patacchini, 2008; Ahlfeldt et al., 2012), can be attributed to differences in spatial detail of the datasets and the area under scope. More specific, this chapter employs Dutch postal codes with a mean area of nine km^2 as the geographic unit of analysis, whereas Rosenthal and Strange (2008) and Di Addario and Patacchini (2008) have used respectively US place-of-work PUMA's with a mean area of 6,522 km^2 and Italian local labor markets with a mean area of 889 km^2. Evidently, a high level of spatial detail is necessary to disentangle agglomeration-wage effects on short distances. If such a dataset is not available, then problems concerning collinearity between the concentric ring variables and measurement error will bias the results.

Arzaghi and Henderson (2008) and Ahlfeldt et al. (2012) did analyze a spatially detailed dataset. However, the spatial scope of their dataset is limited. The areas under study were Manhattan and Berlin, respectively, which precludes the detection of agglomeration economies with a large spatial extent. Hence, the crucial feature of this type of study is to analyze a nationwide wage panel with a high level of spatial detail.

Our results are consistent with the few studies that did have access to a nationwide spatially detailed dataset. In particular, Rosenthal and Strange (2003), who also use data at a postal code level, find that agglomeration is not always positively related to the birth-rate of new establishments, especially at short distances. The authors also attribute this finding to the interplay of positive and negative agglomeration spillovers. Duranton and Overman (2005), although they focus on localization rather than agglomeration, also provide evidence that the location pattern of industries does not always decline monotonically.

2.4.2 Regional heterogeneities in the spatial scope of agglomeration economies

All previous estimates apply to the average area in the Netherlands. There are, however, good arguments to expect that some areas benefit differently from agglomeration economies, or do not benefit from them at all. For instance, the less urbanized areas might not meet a critical city size that is required to benefit from agglomeration. In this case, we expect to see no effect of employment on wages for the less urbanized areas, but a positive effect for the most urbanized areas. Another possibility is that, above a certain point that reflects the optimal level of employment, every additional unit of employment raises total congestion more than total gains. In this case, we expect a positive effect of employment for the less urbanized regions but not for the most urbanized regions.

To examine these possible heterogeneities across regions, we split the total sample of 3,722 postal codes in subgroups based on their level of agglomeration. Evidently, how to define agglomeration size is not obvious and, therefore, we use three different measures: total employment within five kilometer distance, within 10 kilometer distance, and employment density within the postal code's administrative borders. Because agglomeration levels are highly skewed to the right, we choose to split our sample into three unequally sized groups, using the 50th and 75th percentile as cutoff points. For convenience, we label these subgroups 'highly urbanized', 'moderately urbanized' and 'little urbanized'. See Appendix A.6 for a map of the Netherlands indicating these three subgroups for the three agglomeration measures. Then, we estimate the second-stage equation (2.5) on each of these subsamples.

Table 2.4. Regional heterogeneity in the spatial scope of agglomeration economies (a)

Column:	(1)	(2)	(3)	(4)	(5)	(6)
Subsample:	Highly urbanized within 5 km	Moderately urbanized within 5 km	Little urbanized within 5 km	Highly urbanized within 10 km	Moderately urbanized within 10 km	Little urbanized within 10 km
Employment 0 to 5 km	−0.005 (0.020)	−0.245 (0.636)	−0.588 (1.302)	0.013 (0.019)	−0.187 (0.444)	0.067 (0.349)
Employment 5 to 10 km	0.102*** (0.019)	−0.011 (0.051)	0.015 (0.121)	0.083*** (0.021)	−0.220 (0.409)	0.082 (0.322)
Employment 10 to 20 km	0.024*** (0.009)	0.013 (0.017)	0.058* (0.032)	0.023** (0.009)	0.025 (0.022)	0.027 (0.033)
Employment 20 to 40 km	0.019*** (0.004)	0.009** (0.005)	0.009 (0.008)	0.011*** (0.004)	0.010* (0.006)	0.010 (0.009)
Employment 40 to 80 km	0.001 (0.002)	0.004 (0.003)	0.004 (0.003)	0.005*** (0.002)	0.005 (0.004)	0.004 (0.003)
Employment 80 to 120 km	0.007*** (0.002)	−0.005 (0.003)	0.001 (0.003)	0.006** (0.002)	0.002 (0.004)	−0.002 (0.003)
IV	YES	YES	YES	YES	YES	YES
Observations	930	931	1,861	930	931	1,861

Notes: Robust standard errors are in parentheses. Employment is expressed as the total number of jobs in millions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Tables 2.4 and 2.5 report the results, from which we derive two conclusions. First, all coefficients of employment within five kilometer distance are insignificant, which suggests that there are no nonlinearities in the wage-agglomeration relationship on short distances. Second, as reflected by the significance levels of the ring variables, employment on more than five kilometer distance is much more important for the highly urbanized areas compared to the moderately and little urbanized areas. These are interesting results, as they indicate that areas should have sufficient agglomeration on short distances to benefit from agglomeration on longer distances.

Table 2.5. Regional heterogeneity in the spatial scope of agglomeration economies (b)

Column:	(1)	(2)	(3)
Subsample:	Highly urbanized within postal code	Moderately urbanized within postal code	Little urbanized within postal code
Employment 0 to 5 km	–0.031 (0.021)	–0.114 (0.091)	–0.134 (0.321)
Employment 5 to 10 km	0.116*** (0.020)	0.088*** (0.033)	0.022 (0.078)
Employment 10 to 20 km	0.023*** (0.008)	0.021* (0.012)	0.020 (0.027)
Employment 20 to 40 km	0.010*** (0.003)	0.008* (0.004)	0.016** (0.007)
Employment 40 to 80 km	0.001 (0.002)	0.003 (0.002)	0.004 (0.003)
Employment 80 to 120 km	0.008*** (0.002)	–0.004 (0.002)	0.000 (0.003)
IV	YES	YES	YES
Observations	930	931	1,861

Notes: Robust standard errors are in parentheses. Employment is expressed as the total number of jobs in millions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.4.3 The influence of foreign agglomeration

So far, the results indicate that agglomeration economies have a wide spatial scope. For a small open economy like the Netherlands, this raises questions about the role of foreign economic mass in shaping domestic wages. Furthermore, if foreign economic mass does in fact affect wages, then the previous results may be confounded by an omitted variable bias. To examine the influence of foreign agglomeration, we add concentric ring variables measuring foreign agglomeration to the model.

Table 2.6. The influence of foreign agglomeration on domestic wages

Column: Specification:	(1) Domestic and foreign employment (separately)	(2) Domestic and foreign employment (separately)	(3) Total employment (domestic + foreign)	(4) Total employment (domestic + foreign)
Domestic employment 0 to 5 km	0.052*** (0.017)	0.025 (0.020)		
Domestic employment 5 to 10 km	0.060*** (0.015)	0.080*** (0.020)		
Domestic employment 10 to 20 km	0.026*** (0.007)	0.028*** (0.010)		
Domestic employment 20 to 40 km	0.014*** (0.003)	0.013*** (0.003)		
Domestic employment 40 to 80 km	0.005*** (0.001)	0.004** (0.002)		
Domestic employment 80 to 120 km	0.001 (0.002)	0.002 (0.002)		
Foreign employment 0 to 10 km	0.172 (0.202)	0.252 (0.273)		
Foreign employment 10 to 20 km	0.061 (0.085)	–0.012 (0.111)		
Foreign employment 20 to 40 km	–0.006 (0.019)	–0.006 (0.021)		
Foreign employment 40 to 80 km	0.005* (0.003)	0.008*** (0.003)		
Foreign employment 80 to 120 km	–0.001 (0.001)	–0.002 (0.001)		
Total employment 0 to 5 km			0.053*** (0.017)	0.023 (0.020)
Total employment 5 to 10 km			0.062*** (0.015)	0.082*** (0.020)
Total employment 10 to 20 km			0.027*** (0.007)	0.025*** (0.009)
Total employment 20 to 40 km			0.014*** (0.003)	0.012*** (0.003)
Total employment 40 to 80 km			0.004*** (0.001)	0.005*** (0.002)
Total employment 80 to 120 km			–0.001 (0.001)	–0.001 (0.001)
IV	NO	YES	NO	YES
R^2	0.057	–	0.057	–

Notes: 3,722 observations (area fixed effects obtained from the first-stage equation). Robust standard errors are in parentheses. Employment is expressed in millions of jobs (domestic) and employed people (foreign). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.6 shows the result of this analysis. We do not find compelling evidence that foreign economic mass influences domestic wages. All concentric ring variables measuring foreign employment are insignificant, except the one measuring foreign employment between 40 and 80 kilometer.[18] Furthermore, when estimating the effects of total employment, which is calculated by summing domestic and foreign employment, we find that the coefficients do not change substantially from the ones measuring domestic employment only.

We conclude that foreign economic mass has, at best, only a limited influence on domestic wages. This finding indicates the existence of substantial border barriers, and it fits within a large strand of the literature dealing with border effects. Brakman et al. (2002), for instance, also find that market potential stemming from abroad does not affect wages in Germany. The good news, however, is that the estimates of earlier studies, which have ignored foreign economic mass, are most likely not biased.

2.5 Sensitivity analyses

In this section, we analyze whether the results in Table 2.3 are robust to alternative specifications (subsection 2.5.1) and whether they are subject to industrial heterogeneities (subsection 2.5.2).

2.5.1 Alternative specifications

To examine the robustness of the main results, we present the estimates of some alternative specifications in Table 2.7. We start with splitting the first concentric ring variable into two variables that measure employment within 2½ kilometer distance and between 2½ and five kilometer distance. The rationale behind this is that the insignificant estimate from Table 2.3 might conceal two significant effects with opposite directions. Column (1) in Table 2.7 shows that splitting the first ring variable into two smaller rings does not yield any significant effects of agglomeration on short distances. The most notable differences between these results and those in Table 2.3 are the VIF-values and the size of the standard errors on short distances, which are

[18] The first two foreign concentric rings are merged to form one variable measuring employment within 10 kilometer because foreign employment is not accurately measured on short distances. Nevertheless we did run the regressions with the full set of concentric ring variables, which lead to similar (insignificant) estimates.

now much larger. Hence, we conclude that multicollinearity issues make it difficult to disentangle the effect of agglomeration within 2½ kilometer distance from agglomeration between 2½ and five kilometer.

Second, we examine the role of measurement error. The assumption that individuals work at the geographic centroid of the postal code and that employment within a postal code is homogeneously spread across space (see Section 2.3.2) is a source of measurement error. This measurement error is likely to be random and, therefore, may bias the estimates downwards. Since this problem is most severe for the large postal codes, we have checked for robustness by excluding the postal codes with an average radius larger than two kilometer. The results, which are reported in column (2) of Table 2.7, show that the coefficient of the 0–5 kilometer ring decreases slightly, which leads us to conclude that the role of measurement error is limited.

Column (3) in Table 2.7 recalculates hourly wages by excluding financial rewards other than the worker's regular pre-tax wage. With this recalculation, the dependent variable does no longer reflect total labor costs as it excludes thirteenth salaries, holiday entitlements, cash bonuses, etc. An advantage of this recalculation, however, is that we can retain those years in which a worker has been employed for less than the full year at the same employer, which increases the number of observations for the first-stage regression.[19] A comparison between the original and newly estimated area fixed effects shows that the recalculation has substantial implications for the area fixed effect estimates. First, both sets of area fixed effects are not as strongly correlated as one may expect (the correlation is 0.66). Second, the recalculation reduces the dispersion of the area fixed effect estimates substantially. More specifically, the variance of these area fixed effects falls from 0.005 to 0.003.

Despite the fact that this alternative definition of hourly wages has a considerable impact on the estimates of the area fixed effects in the first stage, we find that the second-stage estimates are still consistent with the original estimates. When comparing column (3) in Table 2.7 with column (2) in Table 2.3, we find that the most notable change occurs at the 80–120 kilometer ring variable, which turns significant at the five percent level. The significance levels of the other parameters remain similar

[19] This increase in the number of first-stage observations is also the reason why the number of second-stage observations (number of estimated area fixed effects) increases from 3,722 to 3,753.

to the original estimates, although the point estimates are somewhat lower. The lower point estimates indicate that non-regular financial rewards are an important mechanism for agglomeration economies to capitalize into wages.

Table 2.7. Alternative specifications

Column: Specification:	(1) Split first ring	(2) Exclude PC4 with a radius larger than two km	(3) Basic wages	(4) Access to railway stations	(5) Access to highway ramps	(6) Access to transport infra-structure
Employment 0 to 2½ km	0.040 (0.079)					
Employment 2½ to 5 km	0.010 (0.060)					
Employment 0 to 5 km		0.017 (0.021)	0.007 (0.014)	0.022 (0.021)	0.028 (0.020)	0.024 (0.021)
Employment 5 to 10 km	0.079*** (0.023)	0.094*** (0.020)	0.038** (0.016)	0.075*** (0.020)	0.081*** (0.020)	0.079*** (0.020)
Employment 10 to 20 km	0.024*** (0.010)	0.024** (0.011)	0.017** (0.007)	0.023** (0.009)	0.026*** (0.010)	0.024*** (0.009)
Employment 20 to 40 km	0.011*** (0.003)	0.014*** (0.004)	0.006** (0.002)	0.011*** (0.003)	0.011*** (0.003)	0.011*** (0.003)
Employment 40 to 80 km	0.003** (0.002)	0.002 (0.002)	0.004*** (0.001)	0.003** (0.002)	0.004** (0.002)	0.004** (0.002)
Employment 80 to 120 km	0.001 (0.002)	0.002 (0.002)	0.003** (0.001)	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)
Distance to nearest present-day railway station (km)				–0.000 (0.000)		–0.000 (0.000)
Distance to nearest highway ramp (km)					0.001 (0.000)	0.001 (0.000)
IV	YES	YES	YES	YES	YES	YES
Max VIF [Mean VIF]	9.22 [4.26]	3.80 [2.69]	3.96 [2.76]	3.93 [2.56]	3.97 [2.55]	3.97 [2.44]
Observations	3,722	2,851	3,753	3,722	3,722	3,722

Notes: Robust standard errors are in parentheses. Employment is expressed as the total number of jobs in millions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Columns (4) to (6) of Table 2.7 analyze whether the results are sensitive to the accessibility to infrastructure, in the form of straight line distance to the nearest present-day railway station and highway ramp. Access to infrastructure is an input of the production function of firms and is therefore expected to be positively related to wages. This might lead to biased estimates since the outskirts of cities have, in general, good access to highway ramps and have relatively much employment located between five and 10 kilometer distance (see Appendix A.4). Indeed, the bivariate correlation coefficient between employment on 5–10 kilometer distance and the straight line distance to the nearest highway ramp equals –0.347 and is statistically significant at the one percent level. Also, if distance to the nearest present-day railway station is relevant to the wage formation, then omitting this variable from the model will render our historical instrumental variables invalid. The results, however, show that both infrastructure variables have, conditional on the concentric ring variables, no significant effect on wages. Furthermore, we find that the coefficients of the concentric ring variables are fairly insensitive to the inclusion of these additional variables.

2.5.2 Industrial heterogeneities

Next, we investigate whether the spatial scope of agglomeration economies is subject to industrial heterogeneities. This analysis is appropriate because reviews of the empirical literature show that agglomeration economies are stronger for service industries than for manufacturing (e.g., Melo et al., 2009). Ideally, we would analyze industrial heterogeneities at an alphabetical, or even two-digit, NACE-code. However, the first-stage regression requires sufficient observations within each spatial unit to avoid problems related to significance and identification. Hence, the small size of the postal codes places some restrictions on the level of industrial detail we can achieve. Given these considerations, we only distinguish manufacturing (NACE 11–33) and services (NACE 44–99).

Table 2.8 reports the results for two different regressions on a sample of workers employed in the manufacturing and service industries, respectively. We conclude that the concentric ring variables have much more explanatory power for the service industries. The effect of urban agglomeration on wages in the service industries is mostly significant until 40–80 kilometer, which closely resembles our previous estimates. For manufacturing, on the other hand, the estimates of the ring variables are mostly insignificant. We conclude from this analysis that service industries benefit more strongly from agglomeration economies compared to manufacturing industries.[20] More interestingly, we find that service industries do benefit from agglomeration within five kilometer, although the effect is only significant at the 10 percent level and substantially smaller than the effect from agglomeration on 5–10 kilometer.

Table 2.8. Industrial heterogeneities

Column: Subsample:	(1) Manufacturing	(2) Services
Employment 0 to 5 km	–0.085 (0.067)	0.039* (0.022)
Employment 5 to 10 km	0.125** (0.062)	0.095*** (0.025)
Employment 10 to 20 km	0.019 (0.025)	0.013 (0.017)
Employment 20 to 40 km	0.002 (0.009)	0.011** (0.005)
Employment 40 to 80 km	–0.002 (0.004)	0.006*** (0.002)
Employment 80 to 120 km	0.007* (0.004)	–0.004 (0.002)
IV	YES	YES
Observations	2,362	3,626

Notes: Robust standard errors are in parentheses. Employment is expressed as the total number of jobs in millions. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2.6 The magnitude of the wage-agglomeration relationship

The elasticity of wages with respect to agglomeration size has been frequently estimated in the agglomeration literature. As we know from Briant et al. (2010), these estimated coefficients depend, besides the specification, also on the size of the spatial units. It is, therefore, often argued that the disaggregated spatial units should match the boundaries of the economic phenomenon under scope. It is for this reason that many studies aim to identify agglomeration economies at the level of employment

[20] We have tried to disentangle the effects of own industry employment, i.e. localization, from the effects of other industry employment, i.e. urbanization. Unfortunately, this exercise turned out to be unfruitful due to multicollinearity issues: for some variables the calculated VIF-value was larger than 50. A finer grained industrial subdivision could potentially temper these multicollinearity problems, but would also increase identification issues in the first-stage regressions.

areas (Combes et al., 2008), urban areas (De la Roca and Puga, 2017), or NUTS3 areas (Groot et al., 2014).

The decision on which spatial level best fits the mechanisms under scope remains somewhat arbitrary and may turn out to be wrong. Also, as researchers decide to use larger areas to capture large-scale mechanisms, much of the spatial detail will inevitably be lost. The resulting measurement errors could generate biased estimates if they are systematic (Briant et al, 2010). This study is able to overcome most of these issues in estimating the wage-agglomeration elasticity. First, we do not make any presumptions about the scope at which the spillovers operate, but instead determine this scope empirically. Second, since we employ relatively small spatial units, the risk of systematic measurement error is limited.

We obtain an overall elasticity of wages with respect to agglomeration size by re-estimating the second-stage equation. Instead of including the individual concentric rings as explanatory variables, we now include only one log-transformed agglomeration variable (E_r):

$$\sigma_r = \gamma \ln E_r + v_r$$

For this analysis, we calculate four different agglomeration variables. The first one ignores employment outside the postal code's own administrative borders. The second variable takes into account that agglomeration economies can reach 40–80 kilometer, and therefore sums all employment between zero and 80 kilometer. The third one is identical to the second, but excludes employment within five kilometer. Finally, we calculate a weighted employment variable, by multiplying each concentric ring variable by the corresponding point estimates (Table 2.3, column 2) and then sum these results over the rings.

The results are presented in Table 2.9. The elasticity of wages with respect to the postal code's own agglomeration size is 0.007, which is relatively low compared to international standards. For instance, Combes et al. (2008) and De la Roca and Puga (2017), who also estimate worker fixed effects models, find elasticities of around 0.030 and 0.020, respectively. The main difference with these studies is the size of the spatial unit: Dutch postal codes versus the larger French employment areas and Spanish urban areas. Since agglomeration economies can stretch across relatively large distances, postal codes are too small to capture the full scope at which these spillovers operate.

The wage-agglomeration elasticity increases substantially once we include employment at longer distances into our agglomeration measure. Summing all employment within an 80 kilometer radius yields an elasticity of 0.021. The weighted employment elasticity is even larger: 0.025. These elasticities are also much more similar to those found by Combes et al. (2008) and De la Roca and Puga (2017). We conclude from this analysis that research on the magnitude of the wage-agglomeration relationship, should use the correct spatial scale at which agglomeration economies operate. Underestimating the spatial scope of agglomeration economies will result in downward bias of the wage-agglomeration elasticities.

Table 2.9. The magnitude of the wage-agglomeration relationship

Column:	(1)	(2)	(3)	(4)
Specification:	Local postal code	0–80 km summation	5–80 km summation	0–120 km weighted
Log employment	0.007***	0.021***	0.020***	0.025***
	(0.002)	(0.002)	(0.002)	(0.002)
IV	YES	YES	YES	YES

Notes: 3,722 observations (area fixed effects obtained from the first-stage equation). Robust standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Another way to examine the magnitude of the wage-agglomeration relationship is to calculate the expected value of the area fixed effects given the employment levels on various distances: $E\left(\sigma_r | \sum_{D_{r,e} \in D_d} E_e\right)$. Taking the exponential of these values provides us with the expected effect of agglomeration on wages in percentages compared to a hypothetical area with zero employment within 120 kilometer distance. We carry out this exercise using the coefficients from columns (1) to (3) in Table 2.4. Furthermore, we only use coefficients that are statistically significant at conventional levels.

Figure 2.2 plots the results of this exercise. The benefits of agglomeration are concentrated in the highly urbanized areas. The extent to which they enjoy a wage benefit depends on the amount of agglomeration within 120 kilometer distance. The region of Amsterdam benefits most, enjoying a wage advantage of 7.5 to 10 percent. The Hague, Rotterdam and Utrecht have a wage benefit of five to 7.5 percent, whereas this wage premium is between one and five percent for other cities.

Figure 2.2. The expected wage effect of agglomeration

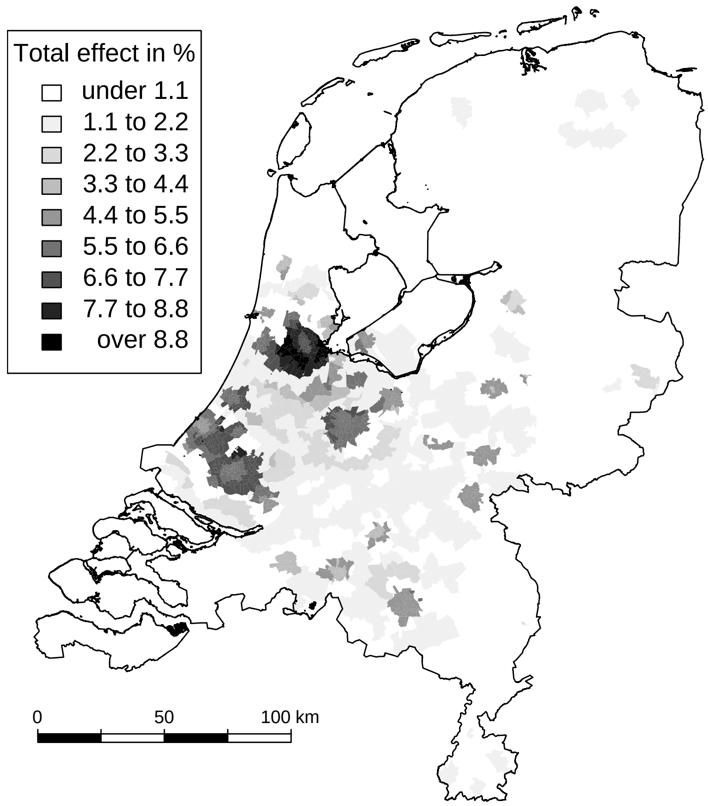

Notes: The predicted effect of agglomeration on wages is calculated using the significant coefficients from columns (1) to (3) in Table 2.4. In percentages, this effect is compared to a hypothetical region with zero employment within 120 kilometer.

2.7 Conclusion

The main contribution of this chapter is to empirically reveal complexities in the spatial scope of agglomeration economies. To this end, we analyze panel data on individual wages with a high level of spatial detail: Dutch postal codes with a mean area of only nine km^2. This high level of spatial detail, which is absent in similar studies, enables us to analyze the effect of agglomeration on long as well as on short distances. The panel structure of the wage data is used to control for sorting of high-skilled labor into urban areas.

The results show that the spatial decay effect of agglomeration on wages is not a monotone function of distance. Wages and agglomeration are not significantly related on short distances (<5 kilometer), whereas they are strongly related on medium distances (5–10 kilometer). The effects of agglomeration on wages attenuate rapidly across geographic space after 10 kilometer, becoming insignificant after 40–80 kilometer. This non-monotone spatial decay effect can arise when negative agglomeration spillovers are substantial on short distances, whereas the positive spillovers dominate on longer distances. Another explanation for the weak relationship between wages and nearby agglomeration is that agglomeration economies capitalize into rents rather than wages on short distances.

The conclusion that wages and agglomeration within five kilometer are unrelated on short distances, however, does not imply that nearby agglomeration is irrelevant to the wage formation. In fact, the data show that only highly urbanized areas benefit from agglomeration on longer distances. Furthermore, this chapter contributes to strands of the literature dealing with border effects, by showing that foreign economic mass has no influence on domestic wages. Although the Netherlands is a small and open economy, this lack of cross-border diffusion of agglomeration economies suggests that national borders still hinder economic interaction.

Other results in this chapter show that the magnitude of the wage-agglomeration elasticity heavily depends on spatial scale. The elasticity of wages with respect to a postal code's own employment level turns out to be relatively low with a point estimate of 0.007. Once we take into account the correct spatial scale at which agglomeration spillovers operate, this elasticity increases to 0.025.