

The spatial scope of agglomeration economies: Revealing complexities with spatially rich wage data^{*}

Paul Verstraten, Gerard Verweij, and Peter Zwaneveld[†]

August 2017

Abstract

This article argues that the spatial decay function of agglomeration economies is much more complex than is often assumed in the agglomeration literature. We provide insight into this issue by analyzing a nationwide and spatially rich wage panel. The key finding is that wages and agglomeration are not significantly related on short distances (<5 km), while strongly and positively related on medium distances (5-10 km). This positive effect attenuates across geographic space and becomes insignificant after 40-80 km. This result, however, does not imply that agglomeration on short distances is irrelevant for productivity. The data show that areas must meet a critical threshold of agglomeration on short distances in order to benefit from agglomeration on further distances. Also, in contrast to manufacturing, we find that service industries do benefit from agglomeration within 5 km distance. Finally, this article finds no evidence that foreign economic mass affects wages in the Netherlands, which suggests that national borders are still a substantial barrier for economic interaction.

JEL Codes: R12; J24; J31

Keywords: Agglomeration economies; Panel data analysis; Spatial wage disparities; Wages.

^{*} We are grateful to Steven Brakman, Joost van Gemen, Stefan Groot, Henri de Groot, Mark Kattenberg, Clemens Kool, Hans Koster, and participants at the CPB seminar for their comments on an earlier version of this article. Furthermore, we thank Thomas Tichelbäcker for his assistance in gathering data on German historical population censuses. The usual disclaimer applies.

[†] All authors are affiliated with the CPB Netherlands Bureau for Economic Policy Analysis. Views expressed in this article are our own and do not necessarily reflect those of the CPB. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sector. Corresponding author: Paul Verstraten (p.verstraten@cpb.nl).

1. Introduction

Although the agglomeration literature has reached no consensus on the maximum spatial range of agglomeration economies¹, most studies do agree on how these agglomeration externalities decay across geographic space (Rice et al., 2006; Arzaghi and Henderson, 2008; Di Addario and Patacchini, 2008; Rosenthal and Strange, 2008; Koster, 2013; Ahlfeldt et al., 2015). In general, the literature concludes that the relationship between the spatial concentration of employment and productivity is strongest on short distances and decays rapidly across geographic space. For some studies, this consistent finding in the literature has even been a reason to assume a priori that agglomeration externalities decay monotonically across space (e.g. Rice et al., 2006; Koster, 2013).²

This article argues that the spatial decay function of agglomeration economies is much more complex than is often assumed. Although there are sound arguments to think that individual mechanisms underlying agglomeration externalities³ can be described by a monotonically declining distance decay function, it is a misconception to assume the same for the net effect of all these externalities, especially when individual mechanisms work in opposite directions. This article provides empirical evidence for this complexity by showing that agglomeration on short distances (<5 km) does not significantly affect wages, whereas it has a strong and significant positive effect on medium distances (5-10 km). This effect attenuates across geographic space and becomes insignificant after 40-80 km.

In addition, we show that the spatial scope of agglomeration economies is very heterogeneous across regions and industries. Only the highly urbanized areas are able to benefit from agglomeration on 5 to 80 km distance, whereas the less urbanized areas largely fail to benefit from nearby urban areas. This result indicates that agglomeration on short distances, though not directly related to wages, must meet some critical threshold in order for areas to benefit from agglomeration on larger distances. Hence, we conclude that agglomeration economies only arise within large urban systems. Similarly, we find evidence that benefits from agglomeration are much stronger for service industries than for

¹ Estimates range from 40-80 km (Rosenthal and Strange, 2008) to only a few kilometers (Arzaghi and Henderson, 2008; Ahlfeldt et al., 2015), and everything in between (e.g. Di Addario and Patacchini, 2008). This lack of consensus can however be attributed to differences regarding the area under scope. For instance, it is self-evident that studies that analyze data within one particular city (e.g. Arzaghi and Henderson, 2008; Ahlfeldt et al., 2015) are well-suited for identifying agglomeration economies with a narrow spatial scope, although they are, by construction, unable to identify agglomeration economies that stretch beyond city borders. The opposite goes for studies that use nationwide data at a highly aggregated spatial level (e.g. Rosenthal and Strange, 2008; Rice et al., 2006; Di Addario and Patacchini, 2008).

² Studies using a market-potential function based on Harris (1954) also assume there is a monotone relationship between measures of distance and the weight attached to another region's economic mass, e.g. Hanson (2005) and Combes et al. (2008).

³ For instance, it is hypothesized that dense urban surroundings foster the diffusion of ideas (learning), promote efficient labor market coordination (matching), and increase the availability of differentiated intermediate inputs (sharing). See Duranton and Puga (2004) for an extensive theoretical overview of these individual mechanisms. As opposed to the gains from agglomeration economies, there also exist mechanisms that lead to congestion costs, e.g. traffic congestion, pollution, and small lot sizes.

manufacturing. This is particularly true for the effect of agglomeration within 5 km distance, which is significantly positive for services.

Our finding that agglomeration economies stretch across a relatively large distance also raises questions about the role of foreign agglomerations in 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 (19th century) population censuses for both Belgium and Germany. Despite the openness of the Dutch economy, our analysis provides no compelling evidence that foreign economic mass affects wages in the Netherlands. This result is nevertheless 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 article employs a nationwide wage panel with a high level of geographic detail: Dutch postal codes with a mean area of only 9 km². The use of this dataset has two advantages compared to earlier work. First, the spatial richness of the dataset enables us to construct concentric ring variables which are very narrow. This high level of spatial detail is an important prerequisite to disentangle the effects of agglomeration on very short distances (e.g. within 2.5 km) 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 km² (Rosenthal and Strange, 2008), 1,394 km² (Rice et al., 2006) and 889 km² (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.⁴

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 known that a considerable part of the urban wage premium is driven by the sorting of high-skilled workers into urban areas (Combes et al., 2008). Other studies on the spatial scope of agglomeration economies using wage data (e.g. Rosenthal and Strange, 2008; Di Addario and Patacchini, 2008; Rice et al., 2006), however, have controlled for observed worker characteristics only. Therefore, these studies run the risk of omitting important unobserved differences in labor quality. This spatial sorting of workers by skills is also the main reason why we prefer the use of wages over a TFP-based approach, which has difficulties in accounting for regional differences in labor quality.

⁴ In fact, some of the few studies that have employed a dataset with a high level of spatial detail to examine the spatial scope of agglomeration economies, do find evidence against the idea that agglomeration economies decline monotonically across space (Rosenthal and Strange, 2003; Duranton and Overman, 2005). These studies analyze the birth-rate of new establishments and location patterns of industries. Arzaghi and Henderson (2008) and Ahlfeldt et al. (2015) do not find evidence against a monotonically declining spatial decay function, despite the high level of spatial detail in these studies. It should be noted, however, that these studies examine relatively small areas (Manhattan and Berlin, respectively), which makes it practically impossible to detect agglomeration economies with a larger spatial extent, see also footnote 1.

The remainder of this article is structured as follows. Section 2 discusses the micro-econometric specification, which enables us to estimate the spatial scope of agglomeration economies. In Section 3 we describe the wage data and the process of constructing concentric ring variables. Section 4 reports the main results regarding the spatial scope of agglomeration economies and analyzes the influence of foreign agglomeration. Section 5 investigates regional and industrial heterogeneities in the spatial scope of agglomeration economies. In Section 6 we examine the magnitude of the wage-agglomeration relationship. Section 7 concludes.

2. Methodology

In order to analyze the relationship between wages and agglomeration, this article 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. These area fixed-effects can be interpreted as area-specific productivity indices. Then, in the second-stage, we explain these area fixed-effects using concentric ring variables that measure the employment levels at various distances, as proposed by Rosenthal and Strange (2003). An important benefit of this two-stage estimation approach is the elegant solution of the dependent disturbances within the regional units.⁵ We will further elaborate on this two-stage approach in the remainder of this section.

2.1 Two-stage estimation approach

In spatial and competitive equilibrium, a profit-maximizing 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

$$\log w_{i,t} = \beta X_{i,t} + \sigma_r R_{r(i,t)} + \sigma_k K_{k(i,t)} + \sigma_t T_{t(i,t)} + \varepsilon_{i,t}, \quad (1)$$

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.⁶

⁵ 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 regional cluster. However, our study relies on workers who change their working location, making the default use of cluster robust standard errors not applicable.

⁶ 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 1, on the other hand, would require only 3,800 additional fixed-effects. Since our main interest is

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 $\text{cov}(X_{i,t}, R_{r(i,t)}) \neq 0$. 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 region-specific productivity effects. 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

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

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 age 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, like 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).

The area fixed-effect estimates that are obtained from Equation (2) reflect regional differences in productivity. The equation below describes how these productivity differences are the result of a complex interplay between agglomeration economies and urban congestion costs at various distances:

$$\sigma_r = \phi \sum_e E_e B(D_{r,e}) - \omega \sum_e E_e C(D_{r,e}) + \varepsilon_r \quad (3)$$

where σ_r is the area fixed-effect parameter from Equation (2), and E_e denotes total employment at establishment e , which we use as a measure of agglomeration.⁷ $B(D_{r,e})$ and $C(D_{r,e})$ represent the

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 does not vary much over time (on the postal code-level, the correlation coefficient between the number of jobs in 2006 and 2014 equals 0.982). Nevertheless we will show in Section 4.2 that our results are robust to the inclusion of an area-year interaction.

⁷ 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 article, 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

distance decay functions of the benefits of agglomeration and urban congestion costs, respectively. These distance decay functions provide weights to employment at various distances and, without loss of generality, can take any value between 0 and 1, depending on the straight line distance between area r and establishment e ($D_{r,e}$).⁸ The parameters ϕ and ω represent the wage effect of the spatially weighted agglomeration measures, and ε_r is a random error term.

Estimating Equation (3), however, is not possible because both the benefits of agglomeration and congestion costs stem from the same source (E_e), which makes it virtually impossible to disentangle these two effects. It is for this reason that this article, as the other studies in this field of research, estimates the net effect of agglomeration:

$$\sigma_r = \gamma \sum_e E_e N(D_{r,e}) + \varepsilon_r. \quad (4)$$

Estimating the net effect of agglomeration has one crucial implication for our setting. Even when the partial effects of agglomeration, $B(D_{r,e})$ and $C(D_{r,e})$, decay monotonically across space, there is no a priori reason to expect that the net effect decays monotonically as well. In fact, the decay function of the net effect is characterized by a wide variety of functional forms. Figure A.1 in the Appendix shows one set of possible functional forms for the decay function.

In order to estimate Equation (4) we construct a set of concentric ring variables that measure total employment at various distance intervals, e.g. within 2.5 km, between 2.5 and 5 km, etc.⁹ 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 decline non-monotonically across space. In line with the work of Rosenthal and Strange (2003), the regression equation of the second-stage then becomes

$$\sigma_r = \sum_{D_d} \gamma_d \sum_{D_{r,e} \in D_d} E_e + \varepsilon_r. \quad (5)$$

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

Addario and Patacchini, 2008; Koster, 2013; Rice et al., 2006). It should however be noted that the patterns of spatial decay can differ between localization and urbanization economies (Rosenthal and Strange, 2003).

⁸ Ideally, we would use measures of effective distance, taking into account natural barriers and transport infrastructure, rather than straight line distance. Yet, the use of effective distance would inevitably lead to problems of reverse causality, because the more productive and dense regions tend to have more and better transport connections than less productive areas. For this reason, we prefer the use of straight line distance, which is arguably a more exogenous measure than effective distance. The results of this article are nevertheless unlikely to be very different when using travel time on roads as a measure of effective distance because this is generally highly correlated to straight line distances (Phibbs and Luft, 1995). Based on our own calculations, we find that straight line distances between postal codes in the Netherlands can explain around 94% of the variation in travel time on roads as predicted by the LMS National Model System for Traffic and Transport.

⁹ Note that this approach implies that we use a Riemann summation as an approximation of the true distance decay function.

estimated parameters of the ring variables (γ_d) give the percentage wage effect of an additional unit of employment within a particular distance interval. A numerical comparison of these parameters provides information on how the net agglomeration economies differ 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.¹¹ Therefore, in order to avoid problems of multicollinearity, the distance intervals must be somewhat wider on longer distances than on shorter distances. We use the following set of cutoff values for our distance intervals: 2½, 5, 10, 20, 40, 80, and 120 kilometer.¹²

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, like (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. (2010) for an extensive discussion.

This article 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

¹⁰ Throughout this article we use standard OLS and IV regressions for the second-stage estimation, although this will generally lead to biased and inefficient estimates as Combes et al. (2008) point out. 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 the correct equations provided by Gobillon (2004) for an unbiased estimate of the standard errors in the second-stage and with an efficient FGLS estimator. Both estimations provided results comparable to the standard estimation strategy. The difference in estimated parameters and standard errors was generally below 10% (results are available upon request). 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.

¹¹ The circumference of a concentric ring increases proportionally to its radius, which makes that more spatial detail is lost when employment is aggregated at longer distances from a particular area. For that reason, the bivariate correlations tend to be larger between rings on long distances than between those on short distances.

¹² We want the first concentric ring to cover most of the postal code's own area. Given the size of most postal code areas, this implies that the width of the first concentric ring cannot be smaller than approximately 2½ kilometer. Then we keep doubling the cutoff values, in order to avoid multicollinearity problems, until 120 km. As we will see in Section 4 of this article, 120 km is far enough to get an insignificant estimate on long distances. 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 km (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 km.

population is strongly autocorrelated over time, for example due to path-dependency caused by self-reinforcing spillovers from agglomeration (Bleakley and Lin, 2012).

The second instrumental variable is the distance to the nearest railway station in 1870 (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 1870 are not expected to influence labor productivity today. In fact, almost half of these stations are no longer operational.

3. Data description

3.1 Microdata and summary statistics

Our empirical model requires four key datasets. First, we use wage data containing individual information for all employees in the Netherlands¹³ 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 a 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 (see next section) using only regular hourly pre-tax wages. This alternative wage definition permits the inclusion of these dropped observations.

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, temporary worker, and WSW-worker.¹⁴ Jobs in agriculture and the fishing industry are excluded from the

¹³ The dataset does not contain self-employed workers, who comprise 10-15% of the Dutch labor force.

¹⁴ The WSW is a Dutch law aimed to foster the employment of persons with disabilities.

sample, because we do not expect substantial agglomeration benefits to occur in these sectors. Also the public sectors are excluded because it is improbable that these sectors meet our underlying assumption that employers are 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 1 million observations per year remain. Table 1 summarizes the data remaining for estimation in the years 2006, 2010 and 2014. The left panel of Figure A.2 shows the mean hourly wages in euros per four-digit postal code.

Table 1
Summary statistics of the longitudinal wage data

	2006	2010	2014
<i>Number of workers</i>	1,456,067	1,192,499	1,115,038
<i>Hourly wages in euro's (price level 2006)</i>			
Mean (standard deviation)	20.6 (11.5)	21.6 (12.4)	21.4 (12.9)
Median	17.8	18.4	18.0
1 st percentile	8.3	8.6	8.4
99 th percentile	67.2	71.8	74.4
<i>Age</i>			
Mean (standard deviation)	40.0 (10.7)	40.9 (10.9)	41.8 (11.1)
Median	39.3	40.8	42.2
1 st percentile	20.5	20.6	21
99 th percentile	61.4	62.7	63.25
<i>Industrial composition (in percentages)</i>			
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

3.2 Spatial variables

The second key dataset contains information on the spatial distribution of both current¹⁵ employment and historical population in the Netherlands and neighboring countries. We constructed this dataset by combining several data sources, which are listed in Table A.1. As can be seen from this table, our spatial unit of analysis, the four-digit postal code in the Netherlands, is rather small with an average area of only 8.86 km². This high level of spatial detail allows us to examine the decay pattern of agglomeration economies on short as well as long distances.

The right panel of Figure A.2 shows the spatial distribution of employment in the Netherlands. Comparing the left and right panel reveals that areas with relatively high employment densities also tend to have relatively high mean wages. This observation is of course merely descriptive and does not imply a causal relationship between wages and agglomeration. Figure A.3 gives an overview of the domestic and foreign spatial distribution of current employment (left panel) and historical population counts (right panel).

Using GIS-software, 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 its area falls within the concentric ring. As previously discussed, we choose a total of seven concentric rings with a respective radius of 2½, 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 the level of current employment and historical population within each concentric ring. Finally, we first-difference the concentric ring variables in order to get total employment and population within particular distance intervals. The domestic and foreign concentric ring variables are graphically presented in Figures A.4 and A.5, respectively.

Table A.2 contains a correlation matrix of the ring variables that measure employment, and those that measure historical population in brackets. This table shows that, although the ring variables are mutually correlated, this correlation is limited due to the increasing distance intervals. Hence, it appears that concerns regarding multicollinearity of the exogenous regressors will be limited.

The third key dataset contains coordinates of all 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). This dataset is used to calculate for each postal code the straight line distance to the nearest railway station in 1870, see Figure A.6. Similarly, we have used the fourth dataset, which contains the coordinates of all highway ramps in 2012, to calculate the straight line distance to the nearest highway ramp, see Figure A.7.

¹⁵ Because we estimate area fixed-effects for the period 2006-2014, we take the middle year (2010) as our measure of current employment.

4. The spatial scope of agglomeration economies

4.1 Main results

Since the outcomes of the first-stage regressions are not directly relevant for our article, this section presents the outcomes of the second-stage regressions only.¹⁶ Column (2) of Table 2 shows the results of the second-stage IV estimates with the full set of concentric ring variables (Equation 5). We conclude that employment within 5 km distance does not significantly affect wages. Between 5 and 10 km we observe a relatively strong effect of employment on wages. More specifically, wages increase by 0.84% when employment between 5 to 10 km distance increases by 100,000.¹⁷ After 10 km the net benefits of agglomeration decrease, although the effect still remains significant until at least 40 km. After 80 km we find no significant effect of employment on wages. A graphical representation of these results can be seen in Figure A.8.¹⁸

When comparing column (1) and (2) of Table 2, 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 (within 2½ km). 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. Furthermore, the Kleibergen-Paap underidentification test and the Hansen J overidentification test confirm that the instruments are both relevant and valid. Hence, in what follows, we will primarily focus on the IV regressions.

It is revealing to see what happens when the model contains only a limited set of concentric ring variables and thus ignores employment at further distances. Therefore, Table 2 reports a total of seven IV regressions; each containing one additional concentric ring variable. By looking only at column (8), we may conclude that employment within 2½ km 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 A.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 IV's to tackle this kind of endogeneity, since they are themselves also correlated with adjacent concentric rings and therefore also with the error term.

¹⁶ 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.

¹⁷ Note that employment within concentric rings is expressed as the total number of jobs in millions. A change of 100,000 jobs within 5 to 10 km is equal to 0.8 standard deviation.

¹⁸ Some of the reported Variance Inflation Factors (VIF) turns out to be rather high, especially for the 2SLS estimates, and some would argue that this is too high given the frequently used rule-of-thumb that a VIF larger than 10 indicates serious multicollinearity. O'Brien (2007), however, rejects the common practice of relying blindly on these general rules, and suggests that researchers should also motivate whether the inclusion or exclusion of variables makes sense from a theoretical perspective. As there are good reasons to expect that agglomeration economies behave differently on various distances, especially on short distances, we prefer this extensive set of concentric ring variables. Nevertheless, we will show in Table 3 that the VIF values can be greatly reduced by merging the first and second concentric ring (0 to 2½ and 2½ to 5 km).

Indeed, we can see from column (8) that the Hansen J test of overidentification rejects the null-hypothesis that the instruments are valid (i.e. uncorrelated with the error term).

Table 2
The spatial scope of agglomeration economies

Dependent variable: first-stage area fixed-effects								
	(1) All rings	(2) All rings	(3) Six rings	(4) Five rings	(5) Four rings	(6) Three rings	(7) Two rings	(8) One ring
Employment 0 to 2.5 km	0.1278** (0.0530)	0.0389 (0.0723)	0.0388 (0.0724)	0.0388 (0.0722)	0.0236 (0.0743)	0.0914 (0.0728)	-0.0132 (0.0697)	0.3339*** (0.0391)
Employment 2.5 to 5 km	-0.0043 (0.0430)	0.0079 (0.0538)	0.0080 (0.0539)	0.0077 (0.0536)	0.0212 (0.0553)	-0.0280 (0.0537)	0.2368*** (0.0404)	
Employment 5 to 10 km	0.0701*** (0.0171)	0.0837*** (0.0220)	0.0829*** (0.0219)	0.0856*** (0.0218)	0.0816*** (0.0220)	0.1532*** (0.0181)		
Employment 10 to 20 km	0.0219*** (0.0071)	0.0208** (0.0097)	0.0199** (0.0095)	0.0219** (0.0095)	0.0458*** (0.0077)			
Employment 20 to 40 km	0.0116*** (0.0026)	0.0123*** (0.0032)	0.0118*** (0.0031)	0.0163*** (0.0027)				
Employment 40 to 80 km	0.0041*** (0.0013)	0.0031** (0.0016)	0.0035** (0.0015)					
Employment 80 to 120 km	0.0003 (0.0016)	0.0010 (0.0018)						
IV	NO	YES	YES	YES	YES	YES	YES	YES
P-value Kleibergen- Paap rk LM statistic		0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
P-value Hansen J statistic		0.6258	0.6533	0.7504	0.4170	0.9613	0.1547	0.0349
Max VIF [Mean VIF]	4.76 [2.74]	10.15 [4.52]	10.14 [4.94]	10.02 [5.19]	9.99 [5.78]	9.58 [6.11]	5.01 [5.01]	
R ²	0.0542	0.0537	0.0537	0.0519	0.0430	0.0332	0.0195	0.0135

Notes: 3,722 observations. 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

The most efficient way to deal with this omitted variable bias is to add other concentric ring variables to the model. Column (7) shows that the coefficient of the first concentric ring becomes insignificant when an adjacent concentric ring is included. Hence, the previous model, containing only one concentric ring, was indeed confounded by an omitted variable bias. The Hansen J test now marginally fails to reject the null-hypothesis at a 10%-level. When adding a third concentric ring variable to the model, the coefficient of the second concentric ring variable becomes insignificant as well. When we continue this process of adding additional ring variables to the model, the results remain quite stable.

4.2 Alternative specifications

To examine the robustness of the main results, we present the estimates of four alternative specifications in Table 3. We start with combining the first and second concentric ring variable into one variable that measures employment within 5 km distance. We do this because the relatively high VIF-values reported in Table 2 raise questions about the accuracy of the estimates, especially the estimates of the first and second ring variable. Column (1) in Table 3 shows that the action of combining these ring variables greatly reduces the VIF, although the results remain roughly similar compared to those in Table 2. The most notable difference between these results and those in Table 2 is the size of the standard errors on short distances. The percentage wage effect of employment within 5 km is now more precisely estimated (see Figure A.9 for a graphical representation), which gives a significant difference at the 10% level between the coefficients for the 0-5 km and 5-10 km ring variables.

Table 3
Alternative specifications

Dependent variable: first-stage area fixed-effects				
	(1) Combine 2 rings	(2) Highway ramps	(3) Basic wages	(4) Area-year interaction
Employment 0 to 5 km	0.0205 (0.0183)	0.0255 (0.0184)	-0.0008 (0.0130)	0.0298 (0.0183)
Employment 5 to 10 km	0.0815*** (0.0195)	0.0855*** (0.0198)	0.0434*** (0.0154)	0.0741*** (0.0192)
Employment 10 to 20 km	0.0210** (0.0097)	0.0222** (0.0097)	0.0167** (0.0075)	0.0240*** (0.0091)
Employment 20 to 40 km	0.0123*** (0.0032)	0.0122*** (0.0032)	0.0061*** (0.0024)	0.0115*** (0.0029)
Employment 40 to 80 km	0.0031** (0.0016)	0.0036** (0.0016)	0.0036*** (0.0011)	0.0035** (0.0014)
Employment 80 to 120 km	0.0010 (0.0018)	0.0011 (0.0018)	0.0027** (0.0013)	0.0002 (0.0015)
Distance to nearest highway ramp (km)		0.0006 (0.0004)		
IV	YES	YES	YES	YES
P-value Kleibergen- Paap rk LM statistic	0.0000	0.0000	0.0000	0.0000
P-value Hansen J statistic	0.6381	0.9528	0.0406	0.4234
Max VIF [Mean VIF]	4.10 [2.81]	4.14 [2.60]	4.11 [2.82]	4.07 [2.78]
N	3,722	3,722	3,752	32,179
R ²	0.0536	0.0546	0.0464	0.0530

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

The second robustness analysis examines whether the results are sensitive to the accessibility to highways. Access to infrastructure can be considered to be part 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 5 and 10 km distance (see Figure A.4). Indeed, the bivariate correlation coefficient between employment on 5-10 km distance and the straight line distance to the nearest highway ramp equals -0.347 and is significant at the 1% level. Column (2) in Table 3 accounts for the straight line distance to the nearest highway ramp but finds no significant effect of this variable on wages. Furthermore, we find that the coefficients of the concentric ring variables are fairly insensitive to the inclusion of this additional variable.

Column (3) in Table 3 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-effects 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-effects estimates substantially. More specifically, the variance of these area fixed-effects falls from 0.0054 to 0.0027.

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 3 with column (2) in Table 2, we find that the most notable change occurs at the 80-120 km ring variable, which turns significant at the 5%-level. The significance levels of the other parameters remain similar to the original estimates, though the point estimates are somewhat lower. These results support the original finding that agglomeration at short distances does not affect wages.

The fourth and last robustness check involves a recalculation of the first-stage equation using an area-year interaction. The inclusion of an area-year interaction captures any variation over time of the area fixed-effects, which provides a more flexible estimator. The downside of including an area-year interaction, is that the number of fixed-effects to be estimated increases by around 30,000, making the calculations very time-consuming with standard regression algorithms.²⁰ Column (4) shows that this alternative specification yields estimates very similar to the base specification.

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 3,752.

²⁰ For this robustness analysis, we used the STATA-module REGHDFE for estimating models with many fixed-effects (Correia, 2014).

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 results in Tables 2 and 3 may be confounded by an omitted variable bias. Table 4 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 until 40 km.²¹ The other two ring variables, 40-80 km and 80-120 km, are significant, but have opposite signs, for which we cannot provide an obvious explanation.

Table 4
The influence of foreign agglomeration on domestic wages

Dependent variable: first-stage area fixed-effects		
	(1) Domestic rings	(2) Foreign rings
Employment	0.0222	}
0 to 5 km	(0.0188)	
Employment	0.0810***	0.2107
5 to 10 km	(0.0197)	(0.2696)
Employment	0.0230**	0.0338
10 to 20 km	(0.0099)	(0.1741)
Employment	0.0132***	-0.0411
20 to 40 km	(0.0035)	(0.0578)
Employment	0.0049***	0.0151**
40 to 80 km	(0.0018)	(0.0061)
Employment	0.0017	-0.0039*
80 to 120 km	(0.0023)	(0.0022)
IV	YES	
P-value Kleibergen-Paap rk LM statistic	0.0000	
P-value Hansen J statistic	0.6269	
Max VIF	37.01	
[Mean VIF]	[9.13]	
R ²	0.0543	

Notes: 3,722 observations. Robust standard errors are in parentheses. Employment is expressed in millions of jobs (domestic) and employed people (foreign). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The domestic and foreign concentric ring variables are estimated simultaneously using one single regression equation. However, in order to save space, the regression results of the domestic and foreign concentric ring variables are presented next to each other in two distinct columns.

²¹ The first three foreign concentric rings are merged to form one variable measuring employment within 10 km. We did this 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.

4.4 Discussion

In summary, Section 4 has shown that agglomeration and wages are not significantly related on short distances (<5 km), but strongly and positively related on longer distances (between 5 and 40-80 km). Furthermore, the estimates are quite robust to a variety of alternative specifications. An intuitive explanation for this spatial pattern would be that the costs of agglomeration (e.g. traffic congestion, pollution, and small lot sizes) decay more rapidly than the gains. In this case, the congestion externalities will offset the productivity gains on short distances, whereas the productivity gains will dominate the congestion costs on longer distances.

Another conclusion that can be drawn from this analysis is that models that include only employment at short distances, are confounded by an omitted variable bias. Therefore, we argue that studies on the relationship between wages and agglomeration should take into account employment at various distances. Instrumental variables cannot adequately tackle these endogeneity issues, because employment at various distances is mutually correlated. Note that ‘classical’ estimates of the wage-density elasticity, which focus on the relationship between urban density and wages at a local level (e.g. at the level of postal codes or municipalities), are likely to suffer from this omitted variable bias. We will address this issue in Section 6.

Although the results of this article 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., 2015), can be attributed to differences in spatial detail of the datasets and the area under scope. To be specific, this article has used Dutch postal codes with a mean area of 9 km² 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 km² and Italian local labor markets with a mean area of about 889 km². 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. In contrast, Arzaghi and Henderson (2008) and Ahlfeldt et al. (2015) did analyze a spatially detailed dataset, although their spatial scope is limited. The areas under study were, respectively, Manhattan and Berlin, which makes by definition the detection of agglomeration economies with a large spatial extent impossible. 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 agglomeration

economies and congestion costs. 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 and can be bumpy. Also, we want to touch upon the study of Koster (2013), who used a spatially detailed dataset to analyze the relationship between agglomeration and rents of commercial property in the Netherlands. Interestingly, Koster finds that commercial rents and agglomeration are strongly related on short distances (<5 km) and unrelated on longer distances. Although our results appear to be the exact opposite of the findings of Koster, it should be noted that the two studies estimate two distinct phenomena: the effect of agglomeration on the marginal productivity of respectively land and labor.²²

Finally, we conclude from the estimations in Table 4 that foreign economic mass has, at best, only a limited influence on domestic wages, which indicates the existence of substantial border barriers. This finding 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. Since we find a relatively large spatial extent of agglomeration economies, we conclude that productivity levels in Dutch border regions can be fostered by reducing border barriers.

5. Regional and industrial heterogeneities

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

To examine 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 the level of agglomeration is not obvious and, hence, we use three different measures of agglomeration: total employment within 5 and 10 km distance, and employment density within the postal code's administrative borders. Because the agglomeration level is 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.

²² The results of both studies can be reconciled if there are mechanisms at work that split the gains from agglomeration unequally among landowners and workers. For instance, 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.

For convenience, we label these subgroups as ‘highly urbanized’, ‘moderately urbanized’ and ‘little urbanized’. See Figure A.10 for a map of the Netherlands indicating the three subgroups using the three agglomeration measures. Then, we run second-stage equation (5) on each of these 9 subsamples.

Tables 5 and 6 report the results, from which we draw two main conclusions. First, all coefficients of employment on 0-5 km distance are insignificant, which suggests that there are no nonlinearities in the wage-agglomeration relationship on short distances. Second, for employment on more than 5 km distance we detect many significant coefficients for the highly urbanized areas, but much less for the moderately and little urbanized areas. Also, as reflected by the R^2 , agglomeration on various distances explains a much larger share of the variation in the estimated postal code fixed-effects for the highly urbanized areas compared less urbanized areas. These are interesting results, as it indicates that areas must meet a critical threshold of agglomeration on short distances in order to benefit from agglomeration on further distances. Another interpretation of these results is that agglomeration economies operate within large urban systems only.

Table 5
Regional heterogeneities in the spatial scope of agglomeration economies (a)

Dependent variable: first-stage area fixed-effects						
	(1) Highly urbanized within 5 km	(2) Moderately urbanized within 5 km	(3) Little urbanized within 5 km	(4) Highly urbanized within 10 km	(5) Moderately urbanized within 10 km	(6) Little urbanized within 10 km
Employment 0 to 5 km	-0.0092 (0.0199)	0.3783 (0.8526)	-0.3449 (1.4515)	0.0128 (0.0186)	-0.9936 (0.7348)	0.2342 (0.3740)
Employment 5 to 10 km	0.1083*** (0.0181)	-0.0097 (0.0506)	-0.0163 (0.1226)	0.0893*** (0.0215)	-0.9556 (0.7013)	0.1484 (0.3520)
Employment 10 to 20 km	0.0241*** (0.0093)	0.0046 (0.0177)	0.0522 (0.0326)	0.0241** (0.0099)	0.0203 (0.0233)	0.0120 (0.0337)
Employment 20 to 40 km	0.0161*** (0.0040)	0.0118** (0.0046)	0.0103 (0.0078)	0.0116*** (0.0036)	0.0095 (0.0061)	0.0116 (0.0090)
Employment 40 to 80 km	0.0012 (0.0019)	0.0033 (0.0026)	0.0041 (0.0031)	0.0046** (0.0020)	0.0086 (0.0060)	0.0038 (0.0032)
Employment 80 to 120 km	0.0073*** (0.0023)	-0.0048* (0.0028)	0.0009 (0.0028)	0.0072*** (0.0024)	-0.0011 (0.0049)	-0.0027 (0.0028)
IV	YES	YES	YES	YES	YES	YES
P-value Kleibergen- Paap rk LM statistic	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
P-value Hansen J statistic	0.2233	0.1539	0.3459	0.2816	0.6764	0.8684
Max VIF [Mean VIF]	2.70 [1.99]	2.79 [1.98]	3.74 [2.85]	1.82 [1.46]	34.46 [13.27]	4.47 [2.67]
N	931	931	1,860	931	931	1,860
R^2	0.1702	0.0282	0.0223	0.0949	0.0373	0.0174

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

Table 6
Regional heterogeneities in the spatial scope of agglomeration economies (b)

Dependent variable: first-stage area fixed-effects			
	(1) Highly urbanized within postal code	(2) Moderately urbanized within postal code	(3) Little urbanized within postal code
Employment 0 to 5 km	-0.0359 (0.0210)	-0.0901 (0.0951)	-0.0931 (0.3333)
Employment 5 to 10 km	0.1181*** (0.0208)	0.0888** (0.0345)	0.0090 (0.0776)
Employment 10 to 20 km	0.0237*** (0.0085)	0.0148 (0.0132)	0.0167 (0.0269)
Employment 20 to 40 km	0.0093*** (0.0029)	0.0103** (0.0044)	0.0173** (0.0070)
Employment 40 to 80 km	0.0011 (0.0016)	0.0024 (0.0023)	0.0039 (0.0030)
Employment 80 to 120 km	0.0069*** (0.0019)	-0.0037 (0.0022)	0.0003 (0.0029)
IV	YES	YES	YES
P-value Kleibergen-Paap rk LM statistic	0.0000	0.0000	0.0000
P-value Hansen J statistic	0.0186	0.3869	0.9350
Max VIF	3.62	4.47	3.59
[Mean VIF]	[2.24]	[2.86]	[2.72]
N	931	931	1,860
R ²	0.1549	0.0889	0.0222

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

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 that we can achieve. Given these considerations, we choose to make a distinction between manufacturing (NACE 11-33) and services (NACE 44-99).

Table 7 reports the results for two different regressions on a sample of workers employed in the manufacturing and service industries, respectively. By comparing columns (1) and (2), 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 km, which closely resembles our previous estimates. Even more interestingly, we do find a significant positive effect of agglomeration within 5 km for service industries. For manufacturing, on the other hand, the R^2 is very close to zero and 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. On short distances (<5 km) these benefits are large enough to offset the congestion costs.²³

Table 7
Industrial heterogeneities in the spatial scope of agglomeration economies
Dependent variable: first-stage area fixed-effects

	(1) Manufacturing	(2) Services
Employment 0 to 5 km	-0.0949 (0.0656)	0.0480** (0.0203)
Employment 5 to 10 km	0.1166* (0.0625)	0.0901*** (0.0242)
Employment 10 to 20 km	0.0294 (0.0257)	0.0127 (0.0177)
Employment 20 to 40 km	-0.0037 (0.0090)	0.0103** (0.0051)
Employment 40 to 80 km	-0.0007 (0.0040)	0.0060*** (0.0020)
Employment 80 to 120 km	0.0076* (0.0039)	-0.0037 (0.0023)
IV	YES	YES
P-value Kleibergen-Paap rk LM statistic	0.0000	0.0000
P-value Hansen J statistic	0.8586	0.8145
Max VIF	4.11	4.08
[Mean VIF]	[2.69]	[2.78]
N	2,362	3,626
R ²	0.0017	0.0458

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

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

²³ 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 enormous 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 it would also increase identification issues in the first-stage regressions.

The decision on which spatial level best fits the mechanisms under scope remains, however, 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 be inevitably lost. This will lead to measurement error, which results in a biased estimate if it is 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 these externalities operate, but instead determine this empirically. Second, since we employ relatively small spatial units, the potential risk of systematic measurement error is limited.

We obtain the elasticity of wages with respect to agglomeration size by re-estimating the second-stage equation. Instead of including the concentric rings as explanatory variables, we now include only one log-transformed agglomeration variable (E_r), see Equation (6) below. 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 km, and therefore sums all employment between 0 and 80 km. The third one is identical to the second, but excludes employment within 5 km. Finally, we calculate a weighted employment variable, by multiplying each concentric ring variable by the corresponding point estimates (Table 3, column 1) and then sum these results over the rings.

$$\sigma_r = \gamma \log E_r + \varepsilon_r \quad (6)$$

The results are presented in Table 8. The elasticity of wages with respect to the postal code's own agglomeration size is about 0.01 (column 1), which is relatively low compared to international standards. For instance, Combes et al. (2008) and De la Roca and Puga (2017), who also estimate a worker fixed-effects model, find elasticities of around 0.03 and 0.025 respectively.²⁴ The main difference with these studies is the size of the spatial unit: Dutch postal codes versus 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 externalities operate.

The wage-agglomeration elasticity increases drastically once we include employment at further distances into our density measure. Summing all employment within an 80 km radius yields an elasticity of approximately 0.022 and the weighted employment elasticity is even larger. 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 the agglomeration economies operate. Underestimating the spatial scope of agglomeration economies will result in downward bias of the wage-agglomeration elasticities.

²⁴ A note of caution when comparing the results: the elasticities in these studies are estimated with respect to employment density and population levels respectively.

Table 8
The magnitude of the wage-agglomeration relationship

Dependent variable: first-stage area fixed-effects				
	(1) Local postal code	(2) 0-80 km summation	(3) 5-80 km summation	(4) 0-120 km weighted
Log employment	0.0098*** (0.0025)	0.0216*** (0.0018)	0.0213*** (0.0018)	0.0252*** (0.0018)
IV	YES	YES	YES	YES
P-value Kleibergen-Paap rk LM statistic	0.0000	0.0000	0.0000	0.0000
P-value Hansen J statistic	0.0688	0.1239	0.0922	0.5234
R ²	0.0354	0.0415	0.0408	0.0509

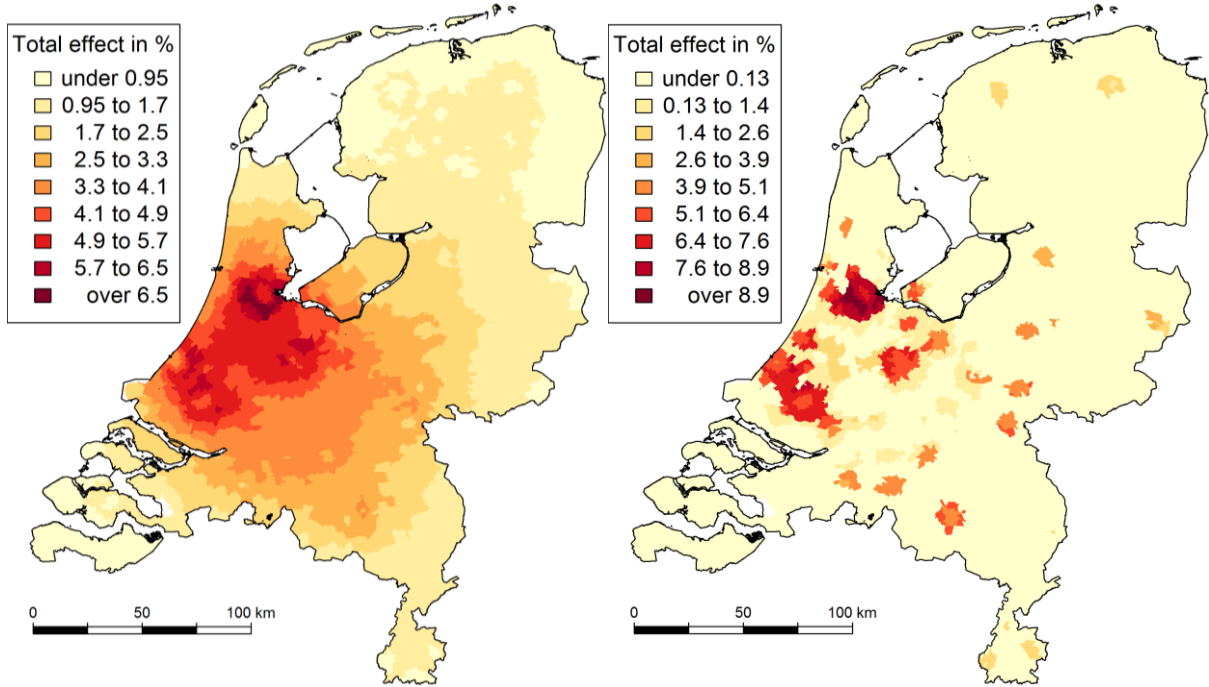
Notes: 3,722 observations. Robust standard errors are in parentheses. * p < 0.1, ** p < 0.05, *** p < 0.01

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: $E(\sigma_r | \sum_{D_{r,e} \in D_d} E_e)$. 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 km distance. To account for regional heterogeneities, we carry out this exercise using the coefficients from both column (1) in Table 3 and those from columns (1) to (3) in Table 5. Furthermore, we only use coefficients that are significant at conventional levels.

Figure 1 plots the results of this exercise. In the left panel, which is based on the overall spatial scope without regional heterogeneities, the benefits of agglomeration are relatively evenly spread across geographic space. Areas that benefit most from agglomeration economies are situated in mid-western part of the country and enjoy a wage advantage of more than 4%. Evidently, this region benefits most due to the close proximity of metropolitan areas Amsterdam, Rotterdam, The Hague and Utrecht. This spatially even distribution of agglomeration economies contrasts the distribution of economic activity, which exhibits a spikier pattern (see the right panel of Figure A.2).

The discrepancy between the distribution of economic activity and the benefits of agglomeration economies raises the following question: why do people and firms flock to dense urban surroundings if the potential benefits from agglomeration are almost just as large outside the urban areas? The answer, as we know from Section 5, is that areas must meet a critical agglomeration threshold to benefit from agglomeration economies. The right panel of Figure 1 plots the results when accounting for regional heterogeneities (columns 1 to 3 in Table 5). The pattern that emerges from these calculations is much more in line with the observed spiky urbanization pattern. Only the highly urbanized areas benefit from agglomeration economies, and the extent to which they enjoy a wage benefit depends on the amount of agglomeration within 120 km distance.

Figure 1
Graphical representation of the expected effect of agglomeration on wages in percentages
(a) using significant coefficients from column (1) Table 3
(b) using significant coefficients from columns (1)-(3) in Table 5



Notes: The predicted effect of agglomeration on wages is calculated using the significant coefficients from column (1) in Table 3 (left panel) and those from columns (1) to (3) in Table 5 (right panel). In percentages, this effect is compared to a hypothetical region with zero employment within 120 km.

7. Conclusions

The main methodological novelty of this article is to estimate the spatial scope of agglomeration economies by analyzing a nationwide and spatially rich dataset on individual wages. This nationwide high level of spatial detail, which is absent in similar studies, enables us to analyze the behavior of agglomeration economies on long as well as on short distances. For the complete sample, we find an insignificant relationship between wages and agglomeration on short distances (<5 km) and a strong positive relationship on medium distances (5-10 km). This positive effect attenuates across geographic space and becomes insignificant after 40-80 km. This spatial pattern can be explained as the net effect of two externalities with opposite effects: agglomeration economies and congestion costs. However, we also conclude that agglomeration on short distances is not irrelevant: areas must meet a critical threshold of agglomeration on short distances in order to benefit from agglomeration on further distances. Furthermore, we find that service industries do directly benefit from urban agglomeration on short distances. This contrasts manufacturing, for which we find only very modest agglomeration economies.

Another finding in this article is that the magnitude of the wage-agglomeration relationship depends heavily 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 0.01. Once we take into

account the spatial scale at which agglomeration externalities operate, this elasticity increases to roughly 0.025, which is close to the estimates of earlier studies using a worker fixed-effects model. Finally, this article contributes to strands of the literature dealing with border effects, by showing that foreign economic mass does not affect wages in the Netherlands. 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. Hence, policymakers can foster productivity levels in border regions by reducing barriers to cross-border economic interaction.

References

- Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., & Wolf, N. (2015). The economics of density: Evidence from the Berlin Wall. *Econometrica*, 83(6), 2127-2189.
- Arzaghi, M., & Henderson, J. V. (2008). Networking off madison avenue. *The Review of Economic Studies*, 75(4), 1011-1038.
- Bleakley, H., & Lin, J. (2012). Portage and path dependence. *The quarterly journal of economics*, 127(2), 587.
- Brakman, S., Garretsen, H., & Schramm, M. (2002). The final frontier? Border effects and German regional wages (No. 197). Hamburg Institute of International Economics (HWWA).
- Briant, A., Combes, P. P., & Lafourcade, M. (2010). Dots to boxes: Do the size and shape of spatial units jeopardize economic geography estimations?. *Journal of Urban Economics*, 67(3), 287-302.
- Combes, P. P., Duranton, G., & Gobillon, L. (2008). Spatial wage disparities: Sorting matters!. *Journal of Urban Economics*, 63(2), 723-742.
- Combes, P. P., Duranton, G., Gobillon, L., & Roux, S. (2010). Estimating agglomeration economies with history, geology, and worker effects. In *Agglomeration Economics* (pp. 15-66). University of Chicago Press.
- Correia, S. (2015). REGHDFE: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects. *Statistical Software Components*.
- De la Roca, J., & Puga, D. (2017). Learning byWorking in Big Cities. *The Review of Economic Studies*, 84(1), 106-142.
- Di Addario, S., & Patacchini, E. (2008). Wages and the city. Evidence from Italy. *Labour Economics*, 15(5), 1040-1061.
- Duranton, G., & Overman, H. G. (2005). Testing for localization using micro-geographic data. *The Review of Economic Studies*, 72(4), 1077-1106.
- Duranton, G., & Puga, D. (2004). Micro-foundations of urban agglomeration economies. *Handbook of regional and urban economics*, 4, 2063-2117.
- Gobillon, L. (2004). The estimation of cluster effects in linear panel models. Processed, ined.
- Groot, S., Groot, H. L., & Smit, M. J. (2014). Regional wage differences in the Netherlands: Micro evidence on agglomeration externalities. *Journal of Regional Science*, 54(3), 503-523.
- Hanson, G. H. (2005). Market potential, increasing returns and geographic concentration. *Journal of international economics*, 67(1), 1-24.

- Harris, C. D. (1954). The Market as a Factor in the Localization of Industry in the United States. *Annals of the association of American geographers*, 44(4), 315-348.
- Koster, H. R. (2013). Rocketing rents The magnitude and attenuation of agglomeration economies in the commercial property market.
- Melo, P. C., Graham, D. J., & Noland, R. B. (2009). A meta-analysis of estimates of urban agglomeration economies. *Regional science and urban Economics*, 39(3), 332-342.
- Moulton, B. R. (1990). An illustration of a pitfall in estimating the effects of aggregate variables on micro units. *The review of Economics and Statistics*, 334-338.
- O'Brien, R. M. (2007). A caution regarding rules of thumb for variance inflation factors. *Quality & Quantity*, 41(5), 673-690.
- Phibbs, C. S., & Luft, H. S. (1995). Correlation of travel time on roads versus straight line distance. *Medical Care Research and Review*, 52(4), 532-542.
- Rice, P., Venables, A. J., & Patacchini, E. (2006). Spatial determinants of productivity: analysis for the regions of Great Britain. *Regional Science and Urban Economics*, 36(6), 727-752.
- Rosenthal, S. S., & Strange, W. C. (2003). Geography, industrial organization, and agglomeration. *review of Economics and Statistics*, 85(2), 377-393.
- Rosenthal, S. S., & Strange, W. C. (2004). Evidence on the nature and sources of agglomeration economies. *Handbook of regional and urban economics*, 4, 2119-2171.
- Rosenthal, S. S., & Strange, W. C. (2008). The attenuation of human capital spillovers. *Journal of Urban Economics*, 64(2), 373-389.

Appendix

Table A.1
Data sources for constructing the concentric ring variables

	Geographic unit	Number of geographic units	Year	Mean area in km ²	Data source
<i>Current employment</i>					
Netherlands	Four-digit postal code	3950	2010	8.86	LISA
Germany*	Municipality	1445	2010	57.00	Statistik der Bundesagentur für Arbeit
Belgium	Municipality	589	2010	52.13	Vlaamse Arbeidsrekening
<i>Historical population</i>					
Netherlands	Municipality	1232	1840	26.38	CBS (Volkstellingen)
Germany*	Municipality	1445	1867	57.00	See on this page below
Belgium	Municipality	589	1846	52.13	Statistics Belgium

* Our dataset does not contain all German municipalities, but only those that belong to the Bundesländer Lower Saxony, Bremen and North Rhine-Westphalia. This is sufficient for our analysis.

Data sources of the German historical population counts

Statistisches Bureau Preussen (1874). Die Gemeinden und Gutsbezirke des Preussischen Staates und ihre Bevölkerung: Nach den Urmaterialien der allgemeinen Volkszählung vom 1. December 1871 (11): Die Gemeinden und Gutsbezirke der Rheinprovinz und ihrer Bevölkerung: nebst einem Anhang, betreffend die Hohenzollerschen Lande. Berlin, Verlag des königlichen Statistischen Bureaus.

Statistisches Bureau Preussen (1873). Die Gemeinden und Gutsbezirke des Preussischen Staates und ihre Bevölkerung: Nach den Urmaterialien der allgemeinen Volkszählung vom 1. December 1871 (8): Die Gemeinden und Gutsbezirke der Provinz Hannover. Berlin, Verlag des königlichen Statistischen Bureaus.

Statistisches Bureau Preussen (1874). Die Gemeinden und Gutsbezirke des Preussischen Staates und ihre Bevölkerung: Nach den Urmaterialien der allgemeinen Volkszählung vom 1. December 1871 (9): Die Gemeinden und Gutsbezirke der Provinz Westfalen und ihrer Bevölkerung: nebst einem Anhang, betreffend die Fürstenthümer Waldeck und Pyrmont. Berlin, Verlag des königlichen Statistischen Bureaus.

Statistisches Bureau (1871). Statistische Nachrichten über das Grossherzogtum Oldenburg (12): Ergebnisse der Volkszählung vom 3. December 1867. Oldenburg, Gerhard Stalling.

Kraus, Antje (1980). In Köllmann, Wolfgang (Ed.), Quellen zur Bevölkerungs-, Sozial- und Wirtschaftsstatistik Deutschlands: 1815-1875, Bd. 1. Quellen zur Bevölkerungsstatistik Deutschlands 1815-1875, p. 329-335. Boppard am Rhein, Boldt.

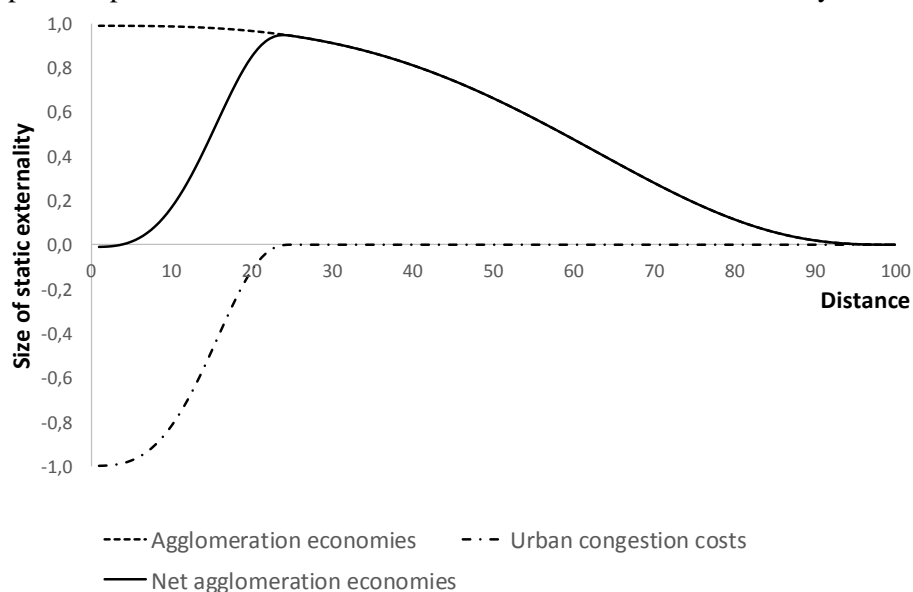
Table A.2
Correlation matrix of domestic concentric ring variables

	0-2½ km	2½-5 km	5-10 km	10-20 km	20-40 km	40-80 km	80-120 km
0-2½ km	1.0000 [1.0000]						
2½-5 km	0.8083 [0.6085]	1.0000 [1.0000]					
5-10 km	0.4837 [0.0778]	0.7096 [0.3450]	1.0000 [1.0000]				
10-20 km	0.2918 [0.1366]	0.3928 [0.1838]	0.6320 [0.3567]	1.0000 [1.0000]			
20-40 km	0.1969 [0.1069]	0.2501 [0.1484]	0.3582 [0.2424]	0.5890 [0.4461]	1.0000 [1.0000]		
40-80 km	0.2198 [0.0919]	0.2809 [0.1400]	0.3853 [0.2156]	0.5173 [0.3718]	0.7014 [0.6143]	1.0000 [1.0000]	
80-120 km	-0.0571 [-0.0472]	-0.0762 [-0.0567]	-0.1202 [-0.0877]	-0.1420 [-0.1100]	-0.0634 [-0.0096]	0.1867 [0.2524]	1.0000 [1.0000]

Notes: The correlation coefficients of concentric ring variables measuring historical population are in brackets.

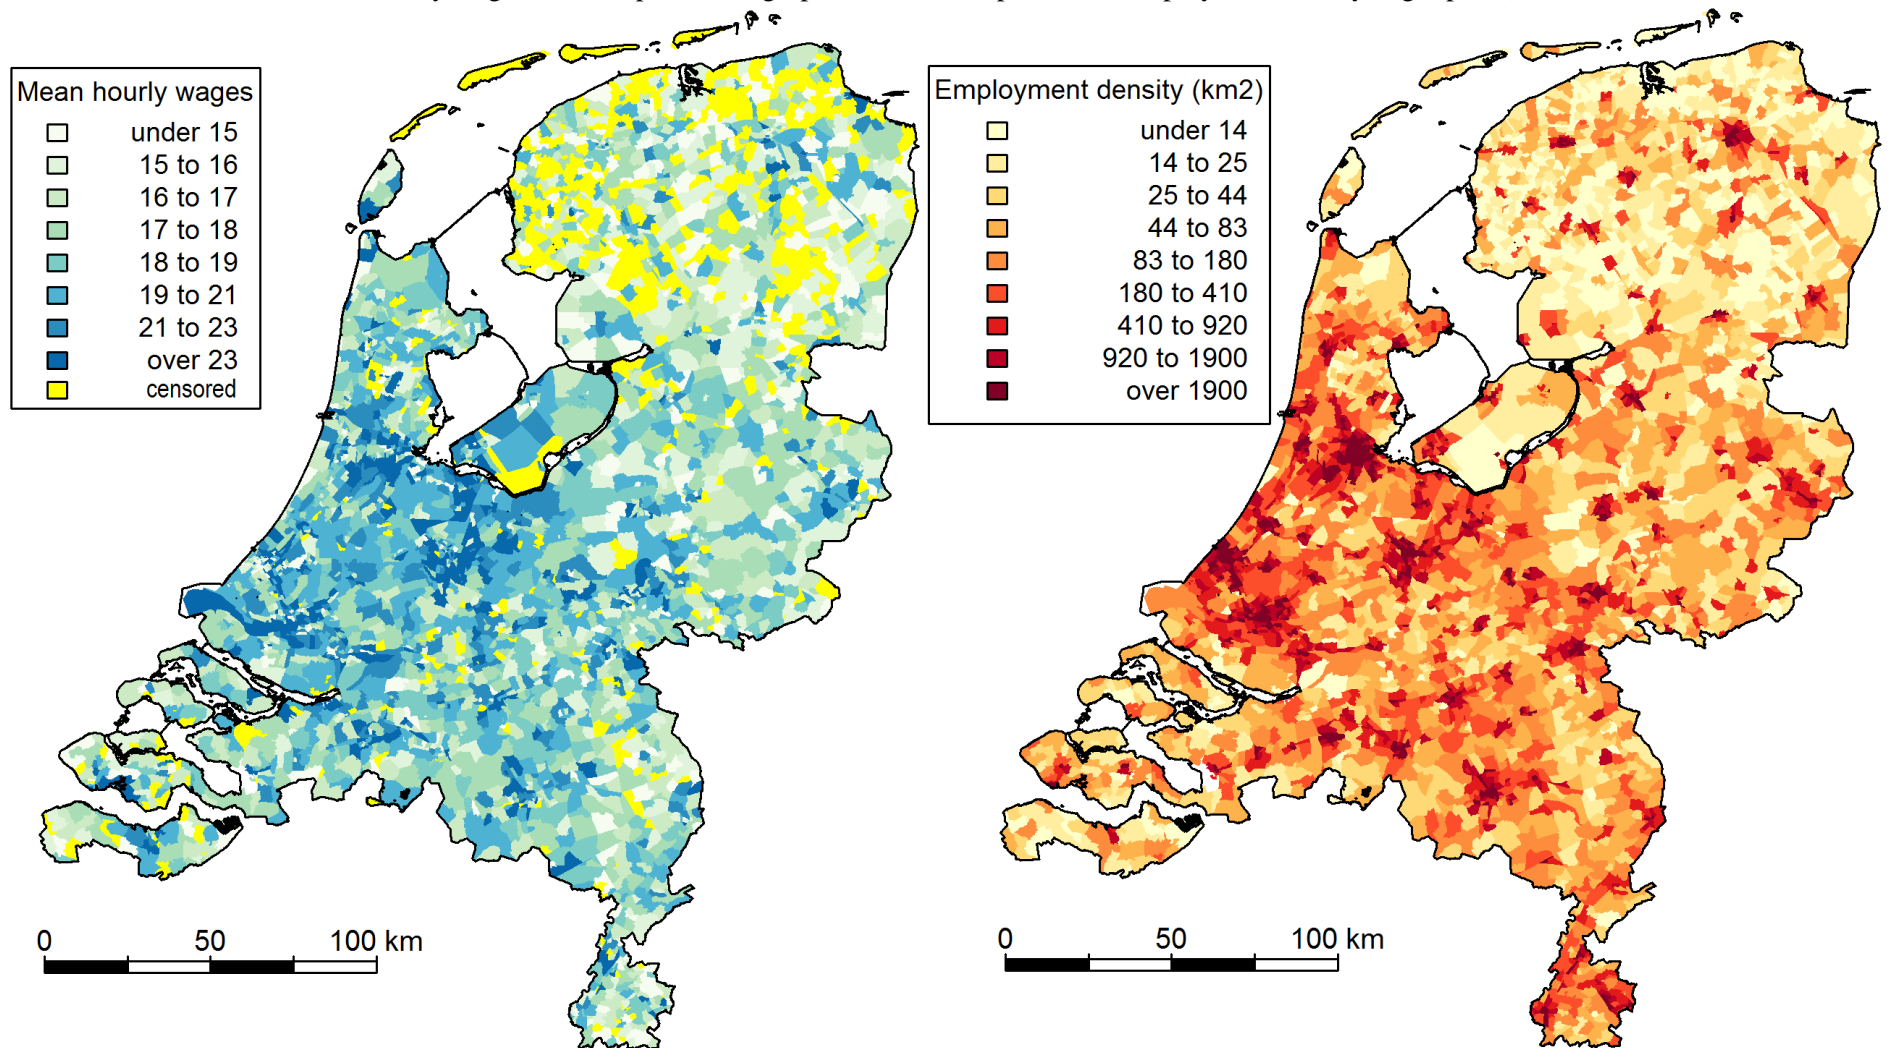
Figure A.1

Graphical representation of the net effect of two monotone distance decay functions



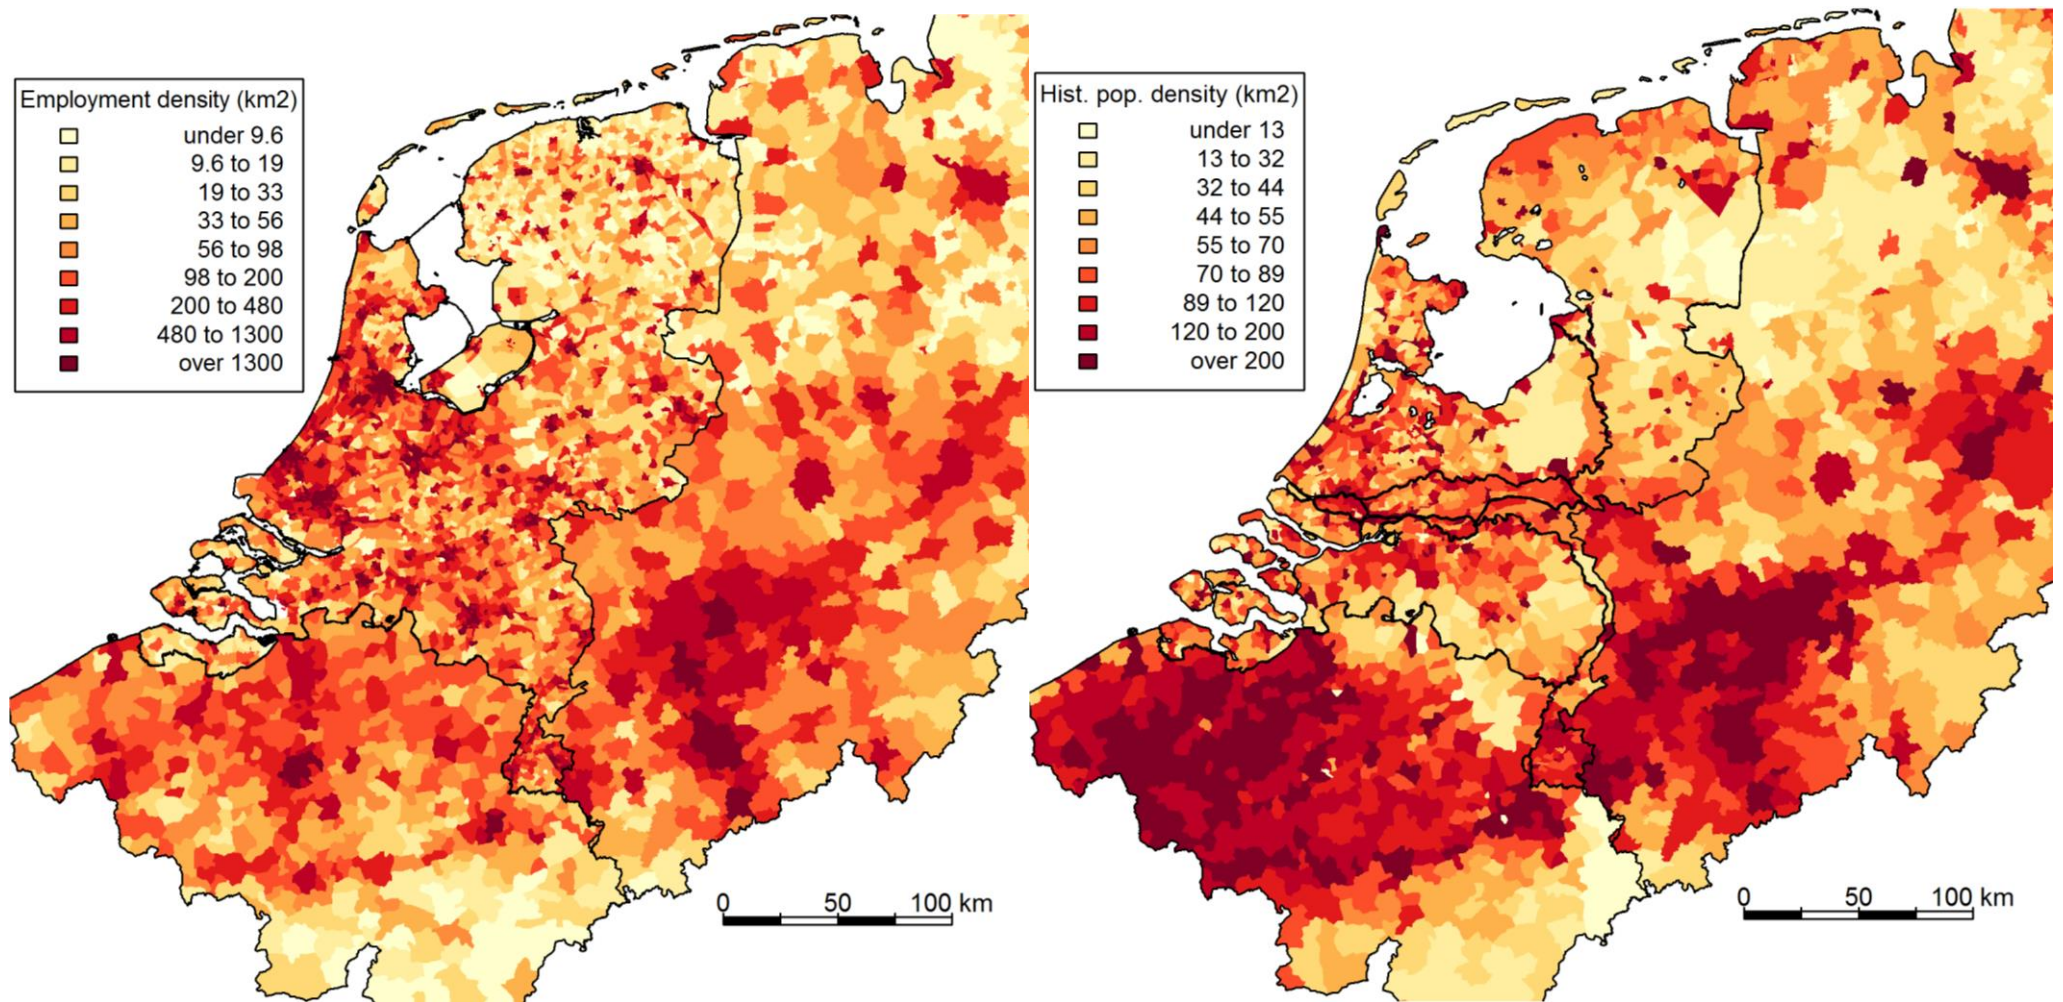
Notes: the purpose of this figure is to show how a non-monotonically declining distance decay pattern might be the net outcome of two monotonically declining distance decay functions with opposite effects. Of course many other functional forms are possible.

Figure A.2
Mean hourly wages in euros per four-digit postal code (left panel) and employment density (right panel)



Notes: In the left figure, approximately 17% of the four-digit postal codes is censored because these postal codes contain less than 10 workers in our dataset.

Figure A.3
Current employment density in the Netherlands, Germany and Belgium (left panel) and historical population density (right panel)



Notes: The figures above are purely descriptive and should be interpreted with caution. For example, the legend displays 9 classes based on quantiles, which may not accurately reflect the fact that the densities are actually highly skewed to the right. Also, the map of the Netherlands in the right panel is based on the 1840 division of municipalities (source: Dr. O.W.A. Boonstra (2007): NLGis shapefiles. DANS. <http://dx.doi.org/10.17026/dans-xb9-t677>), whereas the German and Belgian maps are based on the current municipal division. As a consequence, the Dutch municipalities are on average only half the size of the Belgian and German municipalities, which affects the graphical representation. It is however true that Belgium was more densely populated than the Netherlands during the 19th century. According to population censuses, the Netherlands' total population transcended Belgium's total population around 1930.

Figure A.4
Graphical representation of the concentric ring variables measuring domestic employment

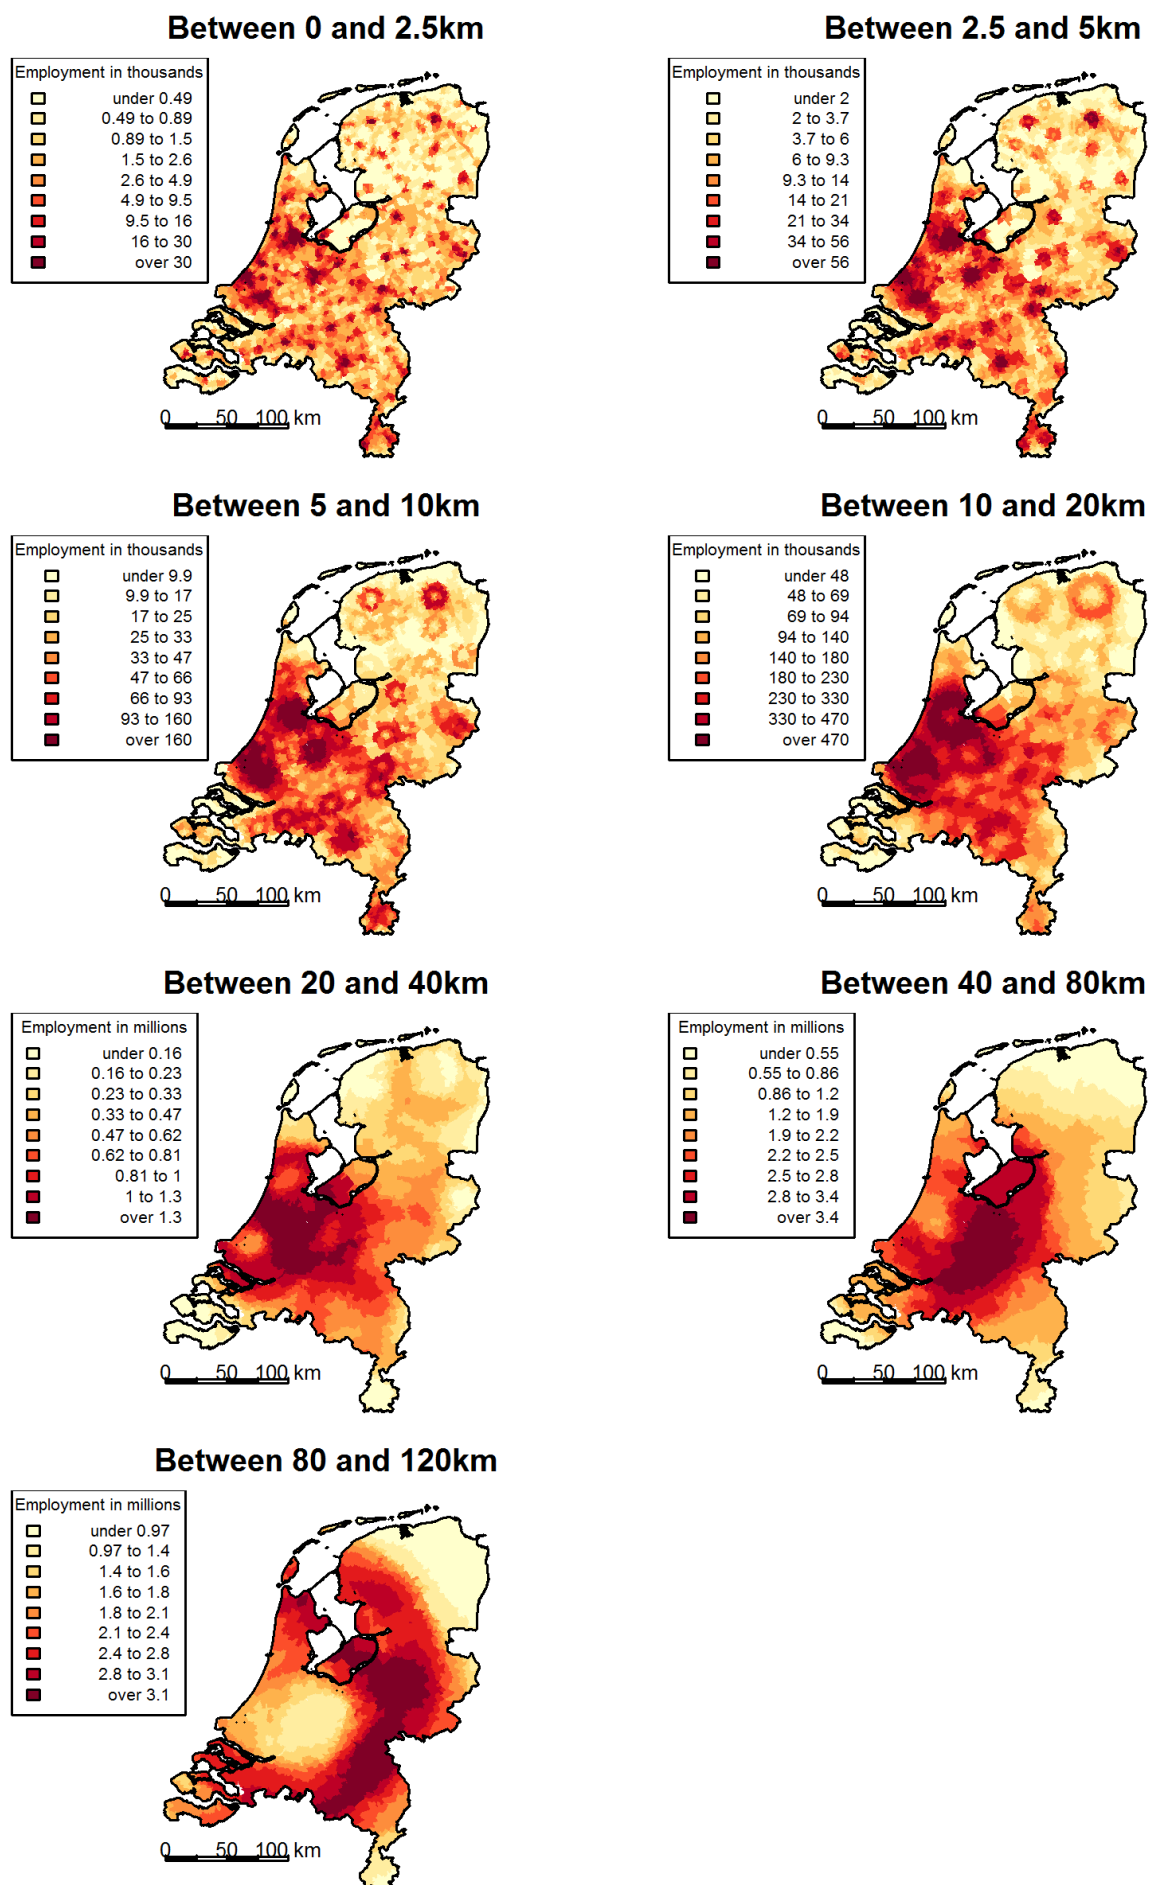


Figure A.5
Graphical representation of the concentric ring variables measuring foreign employment

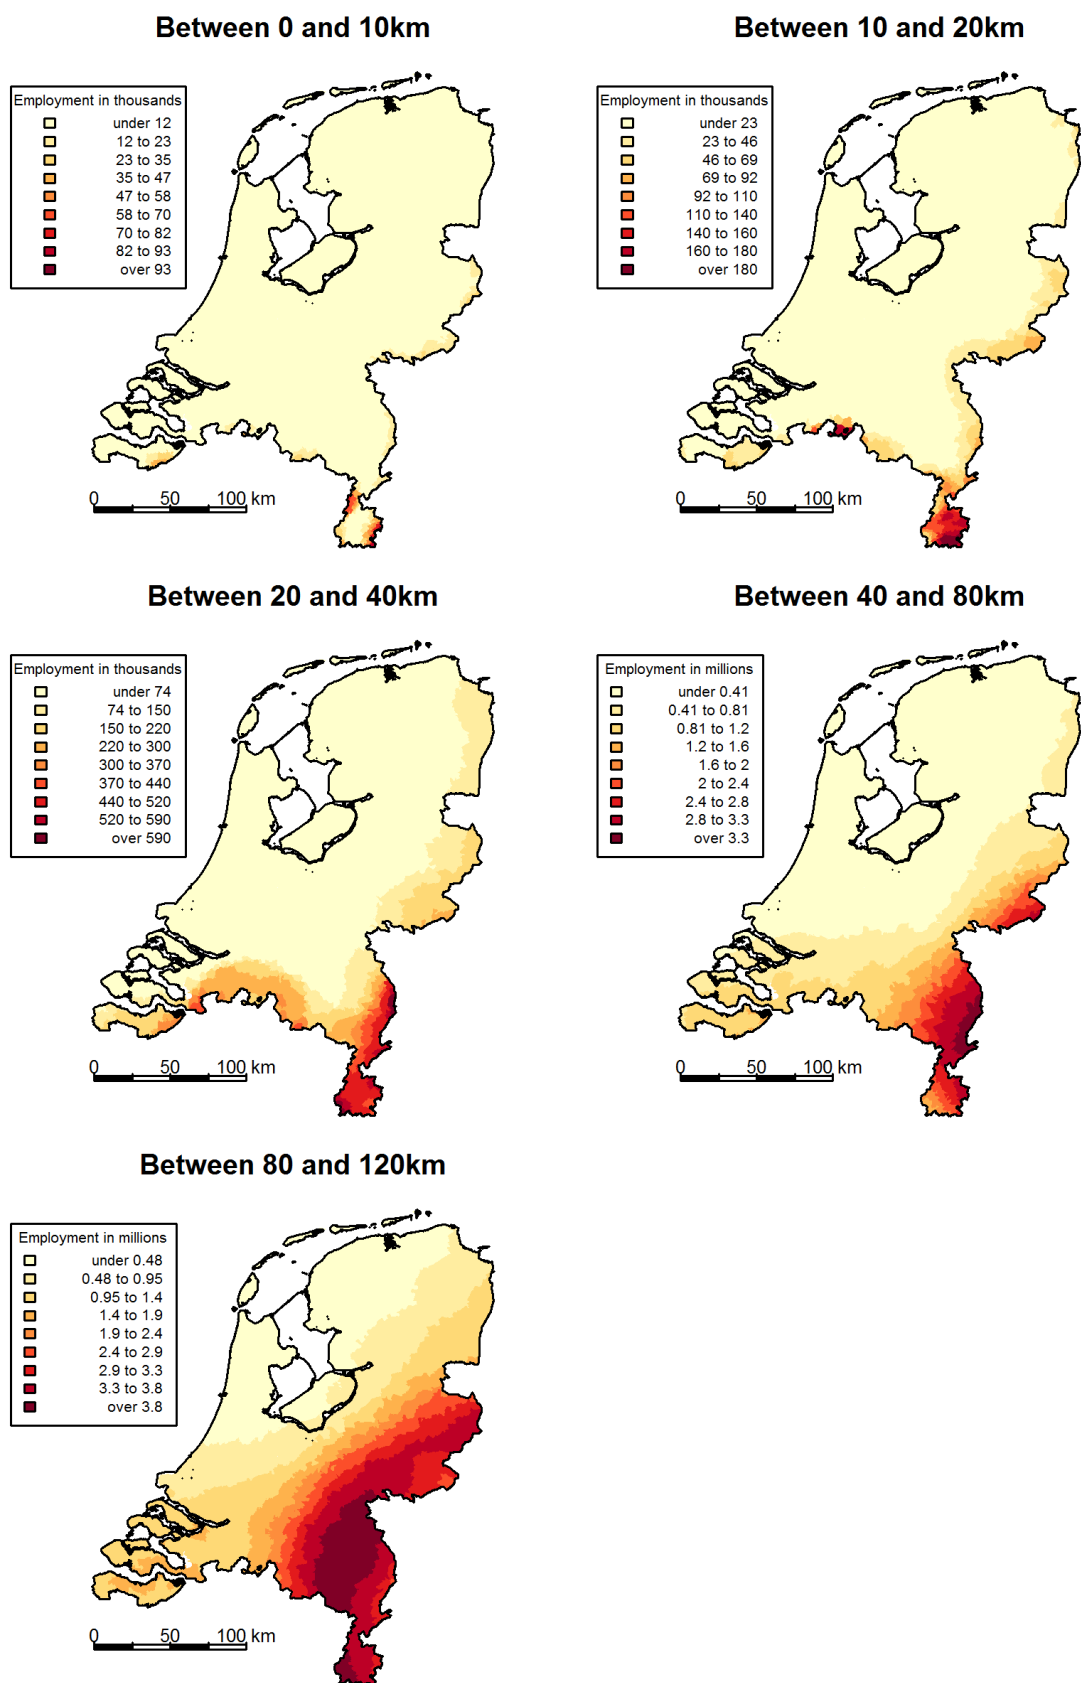


Figure A.6
 Straight line distance to nearest railway station in 1870 and highway ramp in 2012 (right panel)

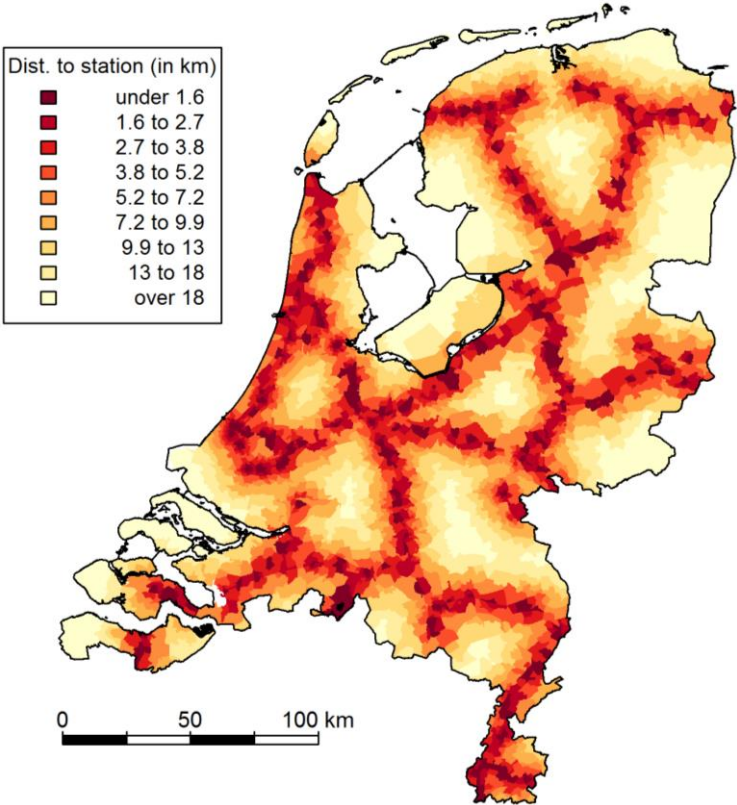


Figure A.7
 Straight line distance to nearest highway ramp in 2012

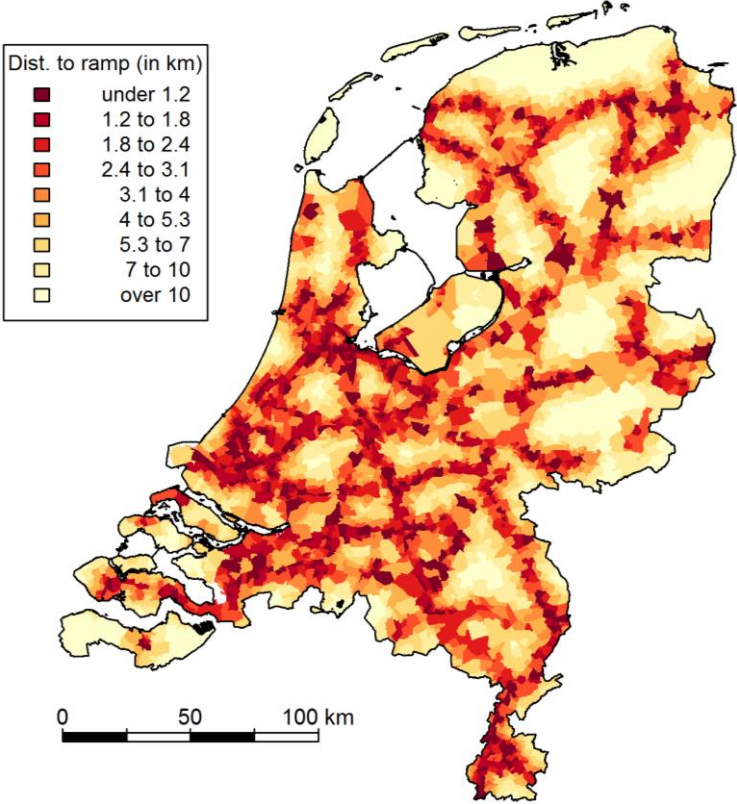
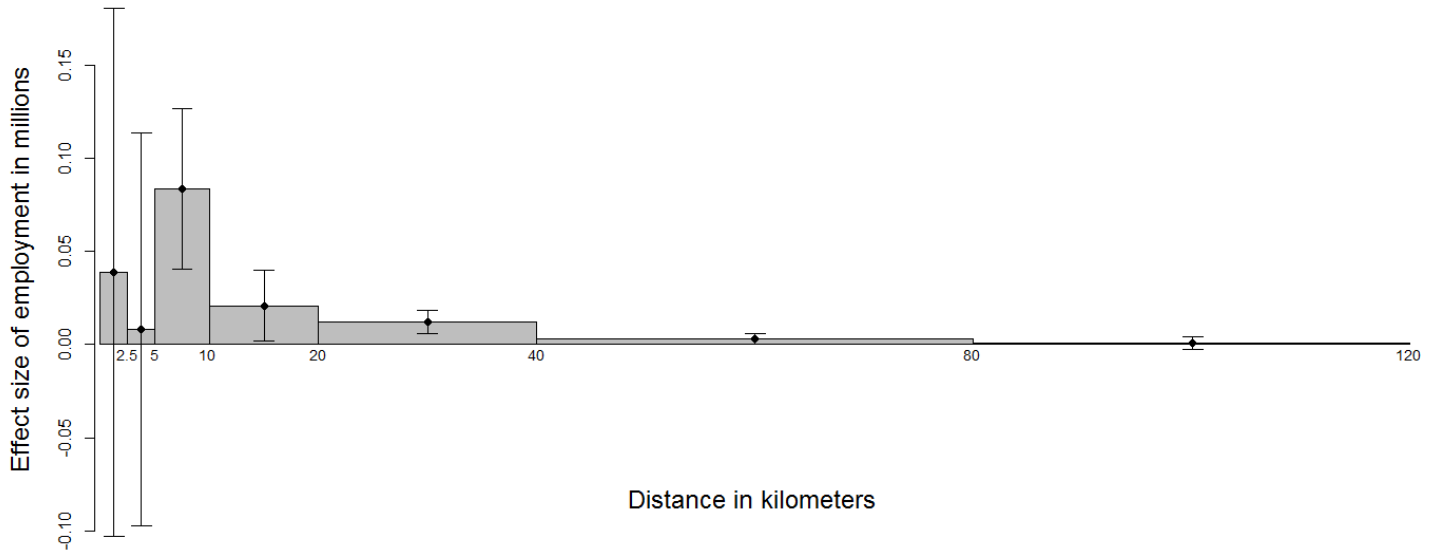
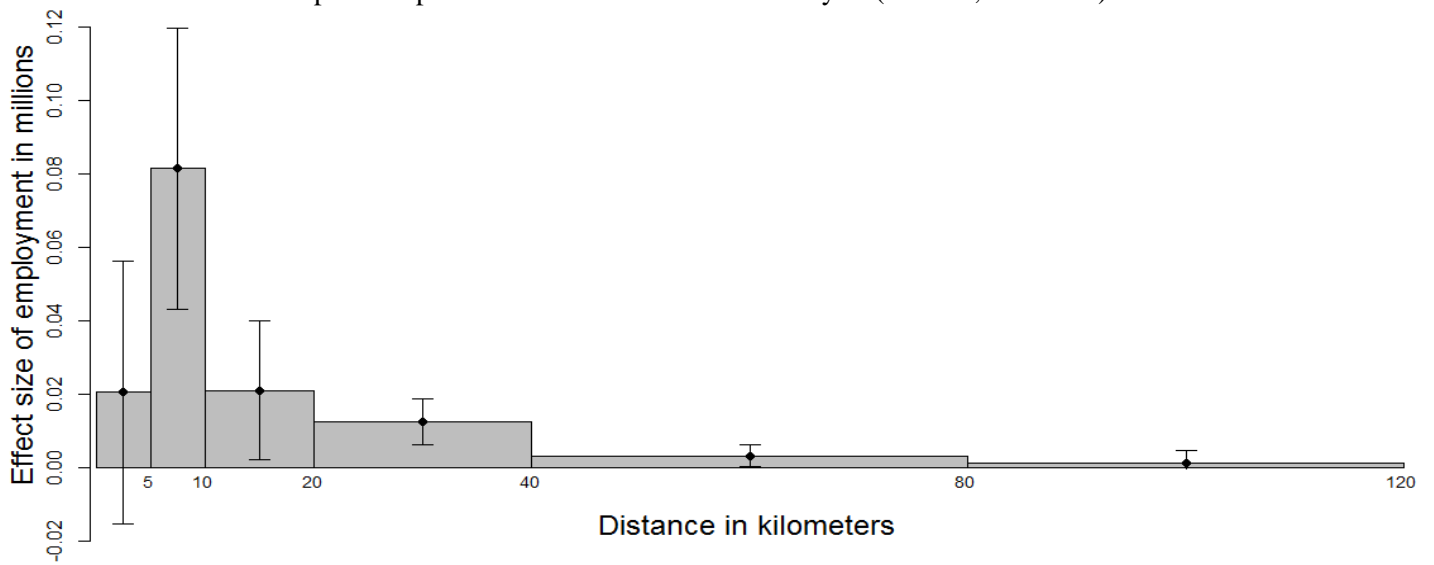


Figure A.8
Graphical representation of the main results (Table 2, column 2)



Notes: The height of the bars represents the point estimates of the concentric ring's parameter (see Table 2, column 2). The bar's width represents the corresponding distance intervals. The error bars show the 95% confidence interval.

Figure A.9
Graphical representation of the robustness analysis (Table 3, column 1)



Notes: The height of the bars represents the point estimates of the concentric ring's parameter (see Table 3, column 1). The bar's width represents the corresponding distance intervals. The error bars show the 95% confidence interval.

Figure A.10
Three measures of agglomeration

