

VU Research Portal

Agglomeration and Human Capital Spillovers

Verstraten, P.W.A.H.

2019

document version

Publisher's PDF, also known as Version of record

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

citation for published version (APA)

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

General rights

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

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

Take down policy

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

E-mail address:

vuresearchportal.ub@vu.nl

Chapter 4

The scope of the external return to higher education

4.1 Introduction⁴⁸

Economists have long stressed the importance of education to welfare (Becker, 1994; Schultz, 1961, Schultz, 1988). Reviews of the empirical literature estimate the private return to an additional year of schooling, in terms of individual earnings, to be around 5–10 percent (Ashenfelter, 1999; Card, 1999; Harmon et al., 2003). It is speculated, however, that the social return to education exceeds the private return. Moretti (2004b) distinguishes between education spillovers that occur because schooling lowers the probability that individuals engage in criminal activities (e.g., Lochner, 2011), education spillovers that foster political behavior (e.g., Friedman, 1962; Hanushek, 2002) and those that lead to productivity spillovers (e.g., Lucas, 1988; Acemoglu, 1996; 1998). Furthermore, Krueger and Lindahl (2001) argue that the former category most likely stems from improvements at the lower end of the education distribution, whereas the latter is expected from expansions in higher education.

This chapter focuses on the productivity spillovers from higher education that capitalize into individual wages. In particular, I aim to analyze whether the productivity spillovers from a large share of highly educated workers occur within

⁴⁸ Apart from minor changes, this chapter was published as: Verstraten, P. (2018). *The scope of the external return to higher education*. CPB Discussion Paper (No. 381).

regions, sectors, firms or a combination of these three work environments. This is an interesting perspective because productivity spillovers may percolate through various mechanisms, and these mechanisms can take place within different work environments. Marshall (1890), Lucas (1988) and Glaeser (1999), for instance, argue that human capital spillovers occur through sharing of knowledge and skills between people. This mechanism is expected to be largely limited within the immediate surroundings of the individual, such as a firm because face-to-face contact is crucial to the exchange of tacit knowledge (Storper and Venables, 2004). Dosi (1988) also claims that tacit knowledge is essential to innovative processes, while Teece (1988) argues that the transmission of tacit knowledge is costly and, therefore, more prevalent within firms. Other potential mechanisms through which the human capitals spillovers operate are less restricted in scope and may occur at the level of labor markets, i.e. within regions and economic sectors. For instance, it is theorized that education spillovers arise due to search frictions in the labor market when individuals and firms must make large ex-ante investments in, respectively, education and physical capital (Acemoglu, 1996) or because new technologies are complementary to high-skilled workers (Acemoglu, 1998). Developing a better understanding of the scope of the external return to higher education is helpful to get an idea which mechanisms are most important.

This is the first attempt to estimate the external return to higher education across three separate work environments: regions, sectors and firms.⁴⁹ Considering all these work environments simultaneously is crucial because the share of highly educated workers within these environments is mutually correlated, and misguided conclusions can be reached when analyzing only a subset. For instance, suppose for the sake of the argument that spillovers only occur within sectors. If we ignore the

⁴⁹ A fourth dimension could be the temporal scope. The temporal scope makes a distinction between spillovers that lead to static productivity advantages, as reflected by a wage-level premium, and dynamic productivity advantages, as reflected by a wage-growth premium. In order to keep the empirical model computationally tractable, I focus on the spatial, sectoral and organizational scope of the static education spillover only. This is in line with the existing micro-based empirical literature (see Section 4.2). Micro-evidence for the existence of dynamic education spillovers is, to my knowledge, still lacking, although there is substantial evidence on a macroeconomic level (e.g., Barro, 1991; 1997; Barro and Sala-i-Martin, 1992). Several micro-based empirical studies have already analyzed the temporal scope of productivity spillovers stemming from agglomeration size (e.g., Glaeser and Maré, 2001; De la Roca and Puga, 2017; and Chapter 3 of this thesis). It would be an interesting angle for future research to examine the temporal scope of education spillovers on a microeconomic level.

sectoral scope of these spillovers, then we would still detect a significant productivity effect from the average schooling level within regions. This is because the sectoral composition is not homogeneous across space, leading to correlations between sectoral and regional schooling levels. Similarly, when ignoring the organizational scope, it may appear that spillovers occur within sectors or regions while they in reality only take place within firms.

In order to analyze the scope of higher education spillovers, I use matched employer-employee panel data on individual earnings, educational attainment and a host of other individual-, region-, sector-, and firm-specific variables. The use of worker fixed effects accounts for non-random spatial sorting of individuals. A solution for dependent disturbances within work locations is provided by following a two-stage procedure to estimate a Mincerian wage equation. The results indicate that the scope of the external return to higher education is very limited. Most of the identified spillovers occur within firms, being a factor of 2–3 larger than those operating outside the firm. The spillovers that take place outside the firm are restricted within the own sector and only occur on short distances from the working place. I find no strong evidence of spillovers coming from outside the own sector or on distances beyond 10 kilometer. The relatively narrow scope of the external return to higher education is in line with the idea that productivity spillovers mainly percolate through information sharing.

The remainder of this chapter is structured as follows. In the next section I review the related literature and pay special attention to previous attempts to identify the scope of education spillovers. Section 4.3 describes the panel data, which contain detailed information on individual earnings, educational attainment and other characteristics of workers, regions, sectors and firms. Then, in Section 4.4, I describe the empirical identification strategy and present the corresponding results. An analysis of heterogeneous external returns to higher education across firms and education groups, as well as robustness checks with alternative specifications, is provided in Section 4.5. Section 4.6 concludes.

4.2 Related literature

This section reviews the empirical literature on the external productivity benefits of education. The discussion begins with the earliest attempts to identify education spillovers, which were mostly focused on the spatial level. These studies also point out some important empirical problems, for which no perfect solution currently exists. Then, I explore the literature that analyzes the relationship between individual wages and schooling levels within sectors and organizations.

On a spatial level, Rauch (1993) was the first to identify the external return to education. His findings indicate a positive correlation between individual wages and average schooling levels at the regional level, even when controlling for the private return to education. Acemoglu and Angrist (2000), however, stress that a causal interpretation of Rauch's estimates may be confounded by problems related to reverse causality and omitted variable bias. In an attempt to address this endogeneity problem, Acemoglu and Angrist have used historic differences in state compulsory attendance laws and child labor laws as instruments for present-day spatial variation in average schooling levels. These instrumental variables estimates offer little evidence for sizeable external returns to education. Rudd (2000) addresses the endogeneity issue by introducing additional region-specific controls, which also leads to negative results.

Krueger and Lindahl (2001) argue that productivity spillovers from education mostly stem from individuals at the higher end of the education distribution. This would explain the small spillover effect found by Acemoglu and Angrist (2000) because their instruments primarily affect the number of people achieving low to medium education levels. For this reason, Moretti (2004a) exploits spatial variation in the share of college graduates, and uses the lagged demographic structure and the presence of a land-grant college as instruments. He also emphasizes that a positive relationship between individual earnings and the share of college graduates does not necessarily indicate the existence of a spillover. Instead, the positive wage effect of college graduates may exist because high- and low-skilled workers are imperfect substitutes in production (Katz and Murphy, 1992). Therefore, Moretti demonstrates the existence of spillovers by showing that the positive spillover effect overcomes the negative productivity effect of increased supply for a sample of high-skilled workers.

Yet, other studies that addressed the issue of imperfect substitutability across workers have reported mixed results. For instance, Groot and De Groot (2014) and Heuermann (2011), who follow an approach similar to Moretti (2004a), find negative and positive evidence for the existence of higher education spillovers, respectively. Muravyev (2008) solves the endogeneity problem by exploiting the abrupt end of communism in Russia. His results do indicate a positive external return to higher education for a sample of high-skilled workers. Ciccone and Peri (2006) and Iranzo and Peri (2009), who employ a constant composition approach to deal with imperfect substitutability, report negative and positive evidence for the existence of spillovers, respectively. Finally, there are relatively few attempts to pin down the exact spatial scope of the higher education spillover. Fu (2007) and Rosenthal and Strange (2008) use a concentric ring-based specification to estimate the spatial attenuation rate, and find that the spillover attenuates rapidly beyond three and eight kilometer, respectively.

Literature on the sectoral scope of education spillovers is relatively scarce. The main empirical challenge of these studies is in finding a proper instrument for the average schooling level within economic sectors, and I am not aware of any study that has succeeded in finding one. Most research on this topic reports positive correlations between wages and schooling levels within sectors (Sakellariou and Maysami, 2004; Winter-Ebmer, 1994). Sakellariou (2001), on the other hand, finds negative evidence. Kirby and Riley (2008), who do find evidence in favor of education spillovers at the sectoral level, also pay some attention to the issue of imperfect substitutability across workers. Their results, however, are not in line with the predictions of imperfect substitutability. That is to say, Kirby and Riley find no inverse relationship between individual schooling levels and the external return to average schooling levels.

Finally, there exists strong evidence in favor of education spillovers at the organizational level. Various studies have reported positive correlations between individual wages and the average education level within firms or establishments (Barth, 2002; Battu et al., 2003; Munch and Skaksen, 2008). Yet, these studies have paid no attention to endogeneity bias, probably because valid instruments at a firm level are hard to find. Martins and Jin (2010) have proposed lagged education levels and the lagged share of workers that are of retiring age as instruments for current

education levels within firms. It is, however, unlikely that these variables have no direct effect on productivity. Mas and Moretti (2009) find that the introduction of highly productive co-workers into a work group has positive and persistent effects on individual productivity. Braakmann (2009) and Canton (2009) have analyzed both the spatial and organizational scope of education spillovers. They conclude that, conditional on schooling levels within firms, there is no significant wage effect from the schooling level within a region.

On balance, it can be argued that the nature of education spillovers is not yet fully understood. In particular, difficulties in finding proper instruments for average schooling levels and imperfect substitutability between high- and low-skilled labor are first-order problems for empirical research. It is for this reason that I aim to approach the issue from a different angle by unraveling the scope at which the higher education spillover operates. This approach can generate valuable insights about the mechanisms that drive the external return to higher education, even when it is not possible to fully address the issues related to endogeneity and imperfect substitution.

4.3 Data description

In this section I describe the matched employer-employee panel data on individual earnings (subsection 4.3.1) and the process of computing the share of highly educated workers within firms, within sectors and at various distances from the work location (subsection 4.3.2). The crucial feature of this chapter is the use of a dataset on individual educational attainment that covers about 73 percent of all Dutch employees. The large coverage allows the calculation of accurate figures for the share of highly educated workers, even at the level of individual firms. This level of detail is necessary to distinguish between the various work environments at which the external return to higher education may operate.

4.3.1 Panel data on individual earnings

Non-public microdata from Statistics Netherlands (CBS) are used to construct a panel dataset on individual earnings over the years 2006 to 2015. The yearly nominal wage of each individual is calculated by summing the regular pre-tax wages, overtime payments, bonuses, thirteenth month salaries and company cars. Dividing the sum of

these annual financial rewards by the reported number of hours worked, which consists of both regular and overtime hours, provides the nominal hourly wage for every individual in a particular year. When an individual has been employed in more than one job during a year, I only preserve the job with the highest number of hours worked during that particular year.

The geographic work location of an individual is exactly known when the job is provided by a firm with only one establishment. When the individual is employed at a firm with multiple establishments, the individual could in principle work in either one of these establishments. In order to geographically assign the work location of these individuals, I use an imputed dataset provided by Statistics Netherlands. This dataset is constructed using an algorithm that assigns employees to the various establishments by minimizing the distance between the individual's place of work and place of residence, while ensuring that the reported number of employees per establishment is matched. Of course, this procedure is not without error, although I believe that the benefit of retaining a large part of the sample outweighs the downside of measurement error.

Information on the highest educational attainment of individuals is retrieved from a dataset that combines various data sources, such as education registrations and surveys, in order to reach the highest possible coverage ratio. The temporal variation in individual schooling levels is, however, less reliable because the exact date on which an individual achieves a higher education level may not coincide with the moment on which the individual self-reports. Therefore, each individual is assigned with the highest known education level and this is assumed to be constant over time. Six different educational attainment levels are distinguished: one for primary education, two for secondary education, and three for higher education.⁵⁰ I remove individuals from the constructed wage panel if the highest educational attainment is unknown.

The data are further restricted as follows. Jobs with less than 12 hours of work per week do not meet the official definition of Statistics Netherlands for being employed, and are therefore removed from the sample. I also decided to exclude the

⁵⁰ These education levels correspond, respectively, to the following SOI 2006 classifications: 10–20, 31–33, 41–43, 51–53, 60 and 70. For a methodological clarification of the SOI 2006 classifications, see Statistics Netherlands (2017a).

following non-regular job types: owner-director, intern, outsourced worker, on-call worker, and WSW-worker.⁵¹ Individuals that are younger than 18 or older than 65 years old are deleted as well. Economic sectors in agriculture, forestry and fishing are excluded because the productivity of these sectors is strongly linked to the availability of natural resources. The public sectors are also removed because earnings in these sectors are heavily regulated.⁵² Since natural barriers may substantially hinder social interaction, I decided to remove five municipalities that are islands without a fixed connection to the mainland. Nominal hourly wages below the legal minimum wage or above 20 times this minimum wage are considered to be outliers, and are therefore removed. Individuals with a non-consecutive employment history are removed because the employment gap may be related to life-changing events. I also exclude jobs for which the part-time status is unknown. Individuals with only one observation over the period 2006–2015 do not contribute to the estimation results of worker fixed effect models, and are therefore deleted. After having cleaned the data, approximately 2.4 million individuals remain. Table 4.1 contains descriptive statistics for the years 2006, 2015 and the complete 2006–2015 period.

4.3.2 Computation of the share of highly educated workers

The share of highly educated workers is calculated as follows:

$$SHE_{r,k,f} = \frac{h_{r,k,f}^{high}}{h_{r,k,f}^{high} + h_{r,k,f}^{medium-low}},$$

where $h_{r,k,f}^{high}$ and $h_{r,k,f}^{medium-low}$ denote the total number of hours worked by high and medium- to low-educated workers, respectively, in municipality r , sector k and firm f .⁵³ This measure assumes that part-time employees contribute less to the overall human capital stock compared to full-time workers. In Section 4.5, I analyze the robustness of the results by calculating the share of highly educated workers on the basis of a headcount.

⁵¹ The WSW is a Dutch law aimed to foster the employment of persons with disabilities.

⁵² This leaves a total of 70 economic sectors, based on the two-digit NACE classification.

⁵³ The distinction between high and medium- to low-educated workers corresponds to the following SOI 2006 classifications: 10–43 and 51–70.

The exact spatial scope of the higher education spillover may not correspond to the size and shape of most municipalities. For this reason, I use GIS tools to compute a concentric ring variable that measures the share of highly educated workers within 10 kilometer straight line distance from the average job in a municipality ($SHE_r^{0-10\ km}$). This measure incorporates the share of highly educated workers in neighboring municipalities, as well as the fact that the geographic centroid of the municipality may not correspond to the economic center of gravity. To examine whether the higher education spillover operates on distances beyond 10 kilometer, I also calculate ring variables measuring the share of highly educated workers on 10–40 and 40–80 kilometer distance.⁵⁴

When calculating the concentric ring variables, I assume a homogeneous mix of high and medium-low educated workers within a municipality. This assumption is necessary because I have no information on how these workers are spatially mingled within the municipalities. Nevertheless, I do take into account the spatial distribution of total employment within municipalities by using the LISA employment register. This dataset contains information on the spatial distribution of jobs at the level of four-digit postal codes. These postal codes are on average 10 times smaller than a municipality. The bottom of Table 4.1 provides descriptive statistics for the share of highly educated workers within 10 kilometer distance, within the sector and within the firm.

⁵⁴ A concentric ring-based approach ignores natural barriers, such as water bodies. Alternatively, I could calculate the ring variables on the basis of travel times rather than straight-line distances. However, the use of travel times is not free of complications either (see Chapter 2).

Table 4.1. Descriptive statistics

	2006	2015	2006-2015
Number of workers	1,161,845	1,444,209	2,413,497
Number of municipalities	388	388	388
Number of sectors	70	70	70
Number of firms	150,014	150,035	330,296
Hourly wage in euro's	19.44 (11.07)	24.28 (15.66)	21.89 (13.62)
<i>Individual characteristics</i>			
Age	36.00 (10.40)	38.18 (11.49)	36.87 (11.21)
Female	0.34	0.35	0.35
Part-time	0.28	0.33	0.31
Foreign-born	0.10	0.10	0.11
Foreign-born parent(s)	0.16	0.18	0.18
Primary education	0.06	0.04	0.05
Secondary education, first phase	0.15	0.12	0.14
Secondary education, second phase	0.46	0.48	0.47
Higher education, first phase	0.22	0.23	0.22
Higher education, second phase	0.11	0.12	0.11
Higher education, third phase	0.01	0.01	0.01
<i>Sectoral composition</i>			
Manufacturing	0.18	0.17	0.17
Construction	0.08	0.07	0.08
Logistics	0.07	0.07	0.07
Wholesale	0.13	0.14	0.13
Retail	0.10	0.11	0.11
Consumer services	0.04	0.04	0.04
Hospitality industry	0.03	0.05	0.04
ICT	0.07	0.07	0.07
Financial services	0.07	0.07	0.08
Business services	0.21	0.21	0.21
<i>Share of highly educated workers</i>			
Within 10 km from work location	0.39 (0.08)	0.39 (0.08)	0.39 (0.08)
Within own sector	0.32 (0.21)	0.31 (0.21)	0.32 (0.21)
Within own firm	0.32 (0.27)	0.33 (0.27)	0.32 (0.27)

Note: Standard deviations of continuous variables are in parentheses.

The share of highly educated workers is not calculated without error. In fact, educational attainment is unknown for about 27 percent of the Dutch workforce. This is a source of measurement error, which is likely to be more severe for small firms. Also, the measurement error may be systematic because individuals with an unknown educational attainment are on average older and, therefore, they are more abundant at the lower end of the education distribution. This issue is addressed in subsection 4.4.4, using the following formula for the share of workers with an unknown education level:

$$SUE_f = \frac{h_f^{unknown}}{h_f^{high} + h_f^{medium-low} + h_f^{unknown}}.$$

4.4 Estimating the external return to higher education across three work environments

This section describes the empirical identification strategy and presents the corresponding results. In order to obtain an accurate estimate of the external return to higher education, it is important to distinguish between the private and external return to education, and to account for non-random sorting of high-skilled workers into regions, sectors and firms. For this reason, I follow a two-stage procedure to estimate a Mincerian wage equation with worker fixed effects. Subsection 4.4.1 provides a more thorough description of this two-stage estimation procedure and it contains a discussion on how to correctly specify the first-stage equation.

After having decided about the preferred first-stage specification, I analyze the three work environments at which the higher education spillover may operate. This is done by the estimation of various versions of the second-stage equation. Analyzing three work environments, which may also interact, can easily lead to overly complicated model specifications. It is for this reason that I aim to unravel the scope of the external return to higher education like peeling off an onion. First, I estimate the spatial scope of the spillover in subsection 4.4.2. After having established the maximum spatial extent at which the higher education spillover operates, I include the sectoral scope into the picture in subsection 4.4.3. This subsection examines whether the positive correlation between individual wages and the share of highly educated

workers is caused by spatial proximity, by sectoral proximity (i.e. working within the same sector) or by a combination of both. Then, in subsection 4.4.4, I take into account the most detailed environment: the organizational scope.

4.4.1 Two-stage estimation procedure of the Mincerian wage equation

A two-stage estimation procedure is preferred over estimating the external return to higher education in one single regression equation because it offers a solution to the dependent disturbances within regions, sectors and firms. Dependent disturbances may arise because individuals that share some characteristic, such as a work location or sector, may affect each other and/or could be subject to the same economic shocks. Moulton (1990) demonstrates that this issue can lead to downward biased standard errors and, therefore, overestimated p -values. The standard solution to calculate cluster robust standard errors is not compatible with the estimation of a worker fixed effects model because individuals may change their work location and sector over time. Hence, in order to obtain correct standard errors, I follow the two-stage procedure proposed by Combes et al. (2008).

As already mentioned, it is important to account for individual characteristics in order to get an accurate estimate of the external return to higher education. In particular, it is necessary to distinguish between the private and external return to education, and to account for non-random sorting of high-skilled workers. Keeping this in mind, an appealing specification of the Mincerian wage equation is:

$$\log w_{i,t} = \beta x_{i,t} + \sigma_i + \tau_t + \sigma_{r(i,t),k(i,t),f(i,t)} + \varepsilon_{i,t}, \quad (4.1)$$

where $w_{i,t}$ is the hourly wage of individual i in year t . $x_{i,t}$ is a vector of observed (time-varying) worker characteristics with parameter β , and σ_i is a worker-specific fixed effect, which captures all time-invariant worker characteristics. τ_t is a vector of year-specific fixed effects. $\sigma_{r,k,f}$ denotes a vector of fixed effects that are specific to region r , sector k and firm f . Finally, $\varepsilon_{i,t}$ is an error term.

The estimation of Equation (4.1) allows $\sigma_{r,k,f}$ to capture all wage effects that are related to the region/sector/firm but that are unrelated to the characteristics of the individual itself. Hence, if we expect productivity effects from the share of highly educated workers, then $\sigma_{r,k,f}$ should depend on the share of highly educated workers

($SHE_{r,k,f}$). This hypothesis can be tested by estimating the following second-stage equation:

$$\sigma_{r,k,f} = \alpha + \gamma SHE_{r,k,f} + \varepsilon_{r,k,f}. \quad (4.2)$$

The results of the second-stage equation depend heavily on the exact specification of the first-stage equation. In order to get a more precise idea of how the exact first-stage specification shapes the second-stage outcomes, I have reported four different specifications in Table 4.2. For the sake of simplicity, this table only examines how individual wages are related to the share of highly educated workers within 10 kilometer straight line distance from region r . Column (1), which is the most basic first-stage specification containing only region- and year-specific fixed effects, shows a strong association between individual wages and the share of highly educated workers within 10 kilometer: a one percentage point increase in the share of highly educated workers is on average associated with a 0.8 percent increase in individual wages. Of course, this estimate is purely descriptive and has no causal interpretation because the specification does not yet account for non-random sorting of high-skilled individuals and does not yet distinguish between the private and external return to higher education.

I introduce controls for observed worker characteristics that are unrelated to the individual's education level in column (2). The sign of the control variables is in line with literature on the gender wage gap (Weichselbaumer and Winter-Ebmer, 2005), the part-time wage penalty (Manning and Petrongolo, 2008) and the immigrants wage disadvantage (Barrett et al., 2012). The second-stage estimate of the external return to higher education, though, appears to be insensitive to the inclusion of these controls. Column (3) includes controls for individual educational attainment. The wage effect of education follows a distinct pattern: wages are strictly increasing in individual education levels. It is also clear that a substantially lower estimate of the higher education spillover is obtained once we account for the private return to education: the point estimate falls from 0.771 to 0.407.

Table 4.2. First-stage results and the external return to higher education

Column: First-stage equation:	(1) (4.1)	(2) (4.1)	(3) (4.1)	(4) (4.1)
Year fixed effects	YES	YES	YES	YES
Region fixed effects	YES	YES	YES	YES
Worker characteristics				
Age		0.117*** (0.000)	0.096*** (0.000)	–
Age squared / 100		–0.129*** (0.000)	–0.010*** (0.000)	–0.099*** (0.000)
Female		–0.121*** (0.001)	–0.125*** (0.001)	–
Part-time		–0.188*** (0.001)	–0.154*** (0.001)	–0.016*** (0.000)
Foreign-born		–0.171*** (0.001)	–0.118*** (0.001)	–
Foreign-born parent(s)		–0.067*** (0.001)	–0.034*** (0.001)	–
Secondary education, first phase			0.079*** (0.001)	–
Secondary education, second phase			0.231*** (0.001)	–
Higher education, first phase			0.525*** (0.001)	–
Higher education, second phase			0.708*** (0.001)	–
Higher education, third phase			0.860*** (0.004)	–
Worker fixed effects	NO	NO	NO	YES
R^2	0.066	0.421	0.570	0.941
Second-stage equation:	(4.2)	(4.2)	(4.2)	(4.2)
Region characteristics				
Share of highly educated workers within 10 km	0.783*** (0.074)	0.771*** (0.058)	0.407*** (0.039)	0.083*** (0.008)
R^2	0.273	0.376	0.237	0.204

Notes: First- and second-stage estimates are based on 14,845,265 worker-year observations and 388 region fixed effects, respectively. Robust standard errors, which are clustered by worker in the first-stage estimates, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Finally, in column (4), I add worker fixed effects to the model. The inclusion of worker fixed effects prohibits the estimation of wage effects related to observed time-invariant worker characteristics, such as gender and immigrant background. Although

individual education levels may in principle vary over time, they are assumed to be time-invariant due to data limitations (see Section 4.3). The linear age variable could not be estimated because the linear effect of aging cannot be distinguished from the year-specific wage effects when having worker fixed effects included in the model. The second-stage estimate of the external return to higher education is much smaller when worker fixed effects are included to the model. A one percentage point increase in the share of highly educated workers is associated with 0.08 percent higher individual wages, compared to 0.41 percent in a model without worker fixed effects. This estimate is, however, still statistically significant at the one percent level.

This subsection underlines the importance of properly accounting for the private return to education and other unobserved time-invariant worker characteristics. In particular, a misspecified first-stage equation leads to overestimation of the external return to higher education in the second stage. Hence, in the remainder of this chapter, I rely on first-stage equations that include worker fixed effects, such as in column (4) of Table 4.2. Since the first-stage results are not directly relevant for this study, and because they hardly differ between the various models, I will only report the second-stage estimation results throughout the rest of the chapter.

4.4.2 Spatial scope

In this subsection, I analyze the spatial scope of the external return to higher education. That is to say, the empirical strategy is focused on obtaining the maximum spatial extent at which the spillover operates. To this end, I first estimate a fully specified first-stage equation that contains region-specific fixed effects (σ_r), and then estimate the following second-stage equation:

$$\begin{aligned} \sigma_r = & \alpha + \gamma_1 SHE_r^{0-10 \text{ km}} + \gamma_2 SHE_r^{10-40 \text{ km}} + \gamma_3 SHE_r^{40-80 \text{ km}} + \\ & \rho_1 Z_r^{0-10 \text{ km}} + \rho_2 Z_r^{10-40 \text{ km}} + \rho_3 Z_r^{40-80 \text{ km}} + \varepsilon_r. \end{aligned} \quad (4.3)$$

The right-hand side of this equation contains concentric ring variables measuring the share of highly educated workers within a set of distance intervals: 0–10 kilometer, 10–40 kilometer, and 40–80 kilometer. Also, the specification includes controls for

third spatial characteristics (z_r) that may confound a causal interpretation of the parameters γ .

Table 4.3. The spatial scope of the higher education spillover

Column:	(1)	(2)	(3)	(4)	(5)	(6)
Second-stage equation:	(4.3)	(4.3)	(4.3)	(4.3)	(4.3)	(4.3)
Region characteristics						
Share of highly educated workers within 10 km	0.061*** (0.009)	0.070*** (0.010)	0.037* (0.020)	0.047** (0.019)	0.040** (0.016)	0.034** (0.016)
Share of highly educated workers between 10 and 40 km	0.069*** (0.013)	0.002 (0.018)	-0.006 (0.040)	0.014 (0.038)	0.030 (0.036)	
Share of highly educated workers between 40 and 80 km	0.055*** (0.015)	-0.074*** (0.026)	-0.013 (0.043)	0.009 (0.038)		
Share of part-time workers	NO	YES	YES	NO	NO	NO
Share of female workers	NO	YES	YES	YES	YES	YES
Unemployment rate	NO	NO	YES	NO	NO	NO
Public university	NO	NO	YES	YES	YES	YES
Log employment	NO	NO	YES	YES	YES	YES
R^2	0.293	0.405	0.439	0.400	0.381	0.320

Notes: Second-stage estimates are based on 388 region fixed effects obtained from the first-stage equation. All first-stage results are the same for every column in this table, and can be found in column (4) of Table 4.2. Robust standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The estimation results of Equation (4.3) are presented in Table 4.3.⁵⁵ Column (1) shows the results without any second-stage control variables. All three concentric ring variables are statistically significant at the one percent level. The results change substantially once I control for the share of part-time workers and female workers on the various distance intervals. The ring variable that measures the share of highly educated workers on 10–40 kilometer distance becomes statistically insignificant, and the 40–80 kilometer ring variable becomes negative. The estimated parameter of the share of highly educated workers within 10 kilometer is insensitive to the inclusion of these controls. Then, in column (3), I introduce controls for the unemployment rate,

⁵⁵ To facilitate readability of the tables, I have not reported the estimated coefficients of the control variables. The coefficient of log employment is positive, which is in line with the literature on the urban wage premium (e.g., Melo et al., 2009; Puga, 2010), and the effect of unemployment is negative, which is in accordance with the wage curve literature (e.g., Blanchflower and Oswald, 1994).

the presence of a public university and employment levels at the various distance intervals. With this specification, only the share of highly educated workers within 10 kilometer is statistically significant, although only at the 10 percent significance level.

Angrist and Pischke (2009, p. 47) argue that “more control is not always better”. In particular, some control variables can be considered to be an outcome of the mechanism that we aim to identify. For instance, highly educated workers might affect productivity through their decision to work part-time more often than low-educated workers do. Similarly, highly educated workers have a lower probability of being unemployed, which may have a positive effect on regional productivity levels. In column (4) I exclude these so-called ‘bad’ controls from the model. The influence of this omission on the results is limited. If anything, the coefficients of the share of highly educated workers increases, which is what we would expect if the share of part-time workers and the unemployment rate are mechanisms through which the share of highly educated workers affects productivity. Finally, in columns (5) and (6) I exclude the 10–40 and 40–80 kilometer concentric ring variables from the model to see how this affects the coefficient of the 0–10 kilometer ring variable.⁵⁶

The advantage of a concentric ring-based specification is that it allows for flexibility in the spatial decay pattern, which is not necessarily monotonic in distance (see Chapter 2). However, if the spatial decay effect can be described by a monotonic function of distance, then the concentric ring-based specification is inefficient because it does not allow for spatial decay within the distance intervals. Efficiency gains can

⁵⁶ I cannot dismiss the possibility that the results are (partially) driven by reverse causality. For instance, high regional productivity levels may influence the educational attainment of its residents. I have tried to address this issue by using instrumental variables (IV). A variable qualifies as an appropriate IV if it correlates with the share of highly educated workers, but it should have no independent relationship to wages. The latter condition makes it difficult to find a good IV. I have experimented with a measure of cultural heritage (the number of squared kilometers that has an official protected status) because this type of amenity is known to attract highly educated households in the Netherlands (Van Duijn and Rouwendal, 2013). Also, I have used the share of the housing stock that was rent-controlled in 1981 because this is expected to decrease the share of highly educated workers (Kattenberg, 2014). Admittedly, these instruments are most likely not really exogenous, and may therefore not qualify as valid instruments. For instance, cultural heritage is considered to be a consumption amenity, which leads to higher land prices and, therefore, negatively influences labor productivity because firms use relatively less land (Combes et al., 2008). Outstanding cultural heritage is also associated with an attractive hotel and catering industry, which may facilitate informal gatherings. Furthermore, rent-controlled dwellings are more likely to be liberalized during economic booms (Hochstenbach and Musterd, 2018). The results from the IV-estimates confirm these doubts. The use of instrumental variables increases the coefficient of the share of highly educated workers, which is counterintuitive. Also, the Hansen *J* over-identification test rejects the null-hypothesis that the instruments are uncorrelated to the error term.

then be achieved by employing a pre-defined spatial decay function. This approach, however, requires assumptions about the exact functional form of the decay effect and about the maximum spatial extent of the spillover.

In order to explore which spatial decay function best fits the data, I follow an optimization procedure that has been used by Koster (2013) in a very similar context. First, I calculate a spatially weighted share of highly educated workers based on two different kernel functions (π). One tricube kernel function:

$$SHE_r^{\omega, tricube} = \frac{\sum_c^C \left(h_c^{high} \left(1 - \left(\frac{d_{r,c}}{\omega} \right)^3 \right)^3 I(d_{r,c} \leq \omega) \right)}{\sum_c^C \left((h_c^{high} + h_c^{low}) \left(1 - \left(\frac{d_{r,c}}{\omega} \right)^3 \right)^3 I(d_{r,c} \leq \omega) \right)}, \quad (4.4)$$

and one exponential kernel function:

$$SHE_r^{\omega, exponential} = \frac{\sum_c^C \left(h_c^{high} \left(1 - \sqrt{\frac{d_{r,c}}{\omega}} \right) I(d_{r,c} \leq \omega) \right)}{\sum_c^C \left((h_c^{high} + h_c^{low}) \left(1 - \sqrt{\frac{d_{r,c}}{\omega}} \right) I(d_{r,c} \leq \omega) \right)}. \quad (4.5)$$

In both equations, $d_{r,c}$ denotes the distance in kilometers between region r and another region c . The bandwidth of the distance decay function is denoted by ω and $I(\cdot)$ is an indicator function that equals one when the statement is true, and zero otherwise. Figure 4.1 gives a graphical illustration of the tricube and exponential kernel functions. The tricube function is characterized by a slow spatial decay on short distances that gradually accelerates. In contrast, the exponential function has a steep spatial decay on short distances that slows down as the distance from region r gets larger. Both equations attach zero weight to regions that are located on distances larger than ω . I calculate both Equations (4.4) and (4.5) for 50 integer values of ω : $\{1, 2, 3, \dots, 50\}$.

Then, I execute the optimization procedure by estimating Equation (4.6) for every spatially weighted share of highly educated workers ($SHE_r^{\omega, \pi}$):

$$\sigma_r = \alpha + \gamma_1 SHE_r^{\omega, \pi} + \rho_1 z_r^{0-10 \text{ km}} + \rho_2 z_r^{10-40 \text{ km}} + \rho_3 z_r^{40-80 \text{ km}} + \varepsilon_r \quad (4.6)$$

The concentric ring variables measuring the share of female workers, the presence of a public university and log employment on three different distance intervals are used

as control variables.⁵⁷ Equation (4.7) calculates the mean squared error (MSE) of Equation (4.6):

$$MSE(\omega, \pi) = \frac{1}{R} \sum_{r=1}^R (\sigma_r - \hat{\sigma}_r(\omega, \pi))^2 \quad (4.7)$$

The values for π and ω that minimize the MSE represent the optimal spatial decay function and the maximum spatial extent of the higher education spillover, respectively.

Figure 4.1. A tricube and exponential kernel function





Figure 4.2 provides a graphical representation of the calculated MSEs of the two kernel functions. I conclude that the tricube and exponential kernel function yield very similar results. This is in line with Koster (2013), who also finds that the results are insensitive to the exact specification of the decay effect. Furthermore, the figure shows that the MSE is minimized for values of ω around seven. A bootstrap procedure of 1,000 iterations is used to investigate the confidence intervals around this value. This yields a minimum value of ω that is smaller than 10 kilometer in 95 percent of the instances. The results in this subsection are in line with the findings of Fu (2007)

⁵⁷ It is necessary to specify the control variables in Equation (4.6) as concentric ring variables rather than distance decay functions because I have no ex-ante knowledge about the spatial scope of these controls. An optimization procedure to determine the spatial scope of all four variables simultaneously is computationally too demanding, as it would require me to estimate more than five million regression equations.

and Rosenthal and Strange (2008), who conclude that the external return to college graduates attenuates rapidly beyond three and eight kilometer distance from the working place, respectively. Hence, in the remainder of this chapter, I will ignore the share of highly educated workers on distances beyond 10 kilometer.

Figure 4.2. Graphical representation of the mean squared errors



4.4.3 Sectoral scope

The previous subsection concluded that the external return to higher education does not stretch across distances beyond 10 kilometer. These results, however, do not necessarily prove that the spillover percolates through geographic space. Instead, the higher education spillover might operate within economic sectors only. The outcomes from the previous section may then be the result of spurious correlation because the sectoral composition is not homogeneous across space. In order to examine this hypothesis, I estimate region-sector fixed effects ($\sigma_{r,k}$) in the first-stage equation, which are then used as a dependent variable in the second-stage equation:

$$\begin{aligned} \sigma_{r,k} = & \alpha + \gamma_1 SHE_r^{0-10 \text{ km}} + \gamma_2 SHE_k + \\ & \rho_1 z_r^{0-10 \text{ km}} + \rho_2 z_k + \varepsilon_{r,k}, \end{aligned} \quad (4.8)$$

where the right-hand side contains the share of highly educated workers within 10 kilometer ($SHE_r^{0-10 \text{ km}}$) and within the own sector (SHE_k), as well as two vectors of corresponding control variables ($z_r^{0-10 \text{ km}}$ and z_k). $\varepsilon_{r,k}$ is an error term.

The results of second-stage Equation (4.8) are reported in columns (1) and (2) of Table 4.2. The first column, which does not contain sector-specific controls, indicates that the spillover mostly occurs within the own sector. The coefficient of the share of highly educated workers within the own sector is about twice as large as the regional counterpart. Moreover, the share of highly educated workers within 10 kilometer is statistically insignificant at conventional levels. When sector-specific controls are introduced to the model, both coefficients slightly increase and the higher education spillover within 10 kilometer becomes statistically significant at the 10 percent level.

The positive relationship between wages and the share of highly educated workers within 10 kilometer is either the result of spillovers operating on short distances within the own sector, on short distances outside the own sector, or both. To explore these possibilities, I split the share of highly educated workers within 10 kilometer into two more detailed variables: within the own sector ($SHE_{r,k}^{0-10\text{ km, within }k}$) and outside the own sector ($SHE_{r,k}^{0-10\text{ km, outside }k}$). This gives me second-stage Equation (4.9):

$$\begin{aligned} \sigma_{r,k} = & \alpha + \gamma_2 SHE_k + \\ & \rho_1 Z_r^{0-10\text{ km}} + \rho_2 Z_k + \\ & \gamma_{1a} SHE_{r,k}^{0-10\text{ km, within }k} + \gamma_{1b} SHE_{r,k}^{0-10\text{ km, outside }k} + \varepsilon_{r,k}. \end{aligned} \quad (4.9)$$

Both columns (3) and (4) of Table 4.4 indicate that no external return to higher education occurs outside the own sector. The spillover within the own sector, on the other hand, is positive and statistically significant. Column (3) also shows that the share of highly educated workers within 10 kilometer within the own sector is statistically significant, even conditional on the nationwide share of highly educated workers within the own sector. This result implies that the external return to higher education predominantly operates within sectors and that it is even larger when geographic distances to highly educated workers within the own sector are short. Column (4) includes region- and sector-specific fixed effects to the second-stage equation. With this specification, the wage effect of the share of highly educated

workers is identified on variation within regions and sectors only. The estimates confirm that the spillover takes place only within the own sector.

Table 4.4. The spatial and sectoral scope of the higher education spillover

Column: Second-stage equation:	(1) (4.8)	(2) (4.8)	(3) (4.9)	(4) (4.9)
Region characteristics				
Share of highly educated workers within 10 km	0.028 (0.020)	0.034* (0.019)		–
Share of female workers within 10 km	YES	YES	YES	–
Public university within 10 km	YES	YES	YES	–
Log employment within 10 km	YES	YES	YES	–
Region fixed effects	NO	NO	NO	YES
Sector characteristics				
Share of highly educated workers within own sector	0.059*** (0.004)	0.075*** (0.004)	0.040*** (0.007)	–
Share of female workers within own sector	NO	YES	YES	–
Log employment within own sector	NO	YES	YES	–
Sector fixed effects	NO	NO	NO	YES
Region-sector characteristics				
Share of highly educated workers within 10 km within own sector			0.039*** (0.007)	0.017** (0.007)
Share of highly educated workers within 10 km outside own sector			0.003 (0.020)	–0.506 (0.352)
R^2	0.022	0.072	0.074	0.202

Notes: Second-stage estimates are based on 19,707 region-sector wage effects obtained from the first-stage equation. All first-stage results are the same for every column, and are available upon request.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.4.4 Organizational scope

In this final subsection, I include the organizational scope into the picture. Each firm is considered to be a separate organization, which is nested within sectors and not nested within regions. Therefore, the first-stage equation contains region-firm fixed effects ($\sigma_{r,f}$), which are used as a dependent variable in the following second-stage equation:

$$\begin{aligned} \sigma_{r,f} = & \alpha + \gamma_2 SHE_{k(f)} + \gamma_3 SHE_f + \\ & \rho_1 Z_r^{0-10 \text{ km}} + \rho_2 Z_{k(f)} + \rho_3 Z_f + \\ & \gamma_{1a} SHE_{r,k(f)}^{0-10 \text{ km, outside } k} + \gamma_{1b} SHE_{r,k(f)}^{0-10 \text{ km, within } k} + \varepsilon_{r,f}. \end{aligned} \quad (4.10)$$

Equation (4.10) is very similar to Equation (4.9), with the exception that the model includes the share of highly educated workers within firms (SHE_f), as well as other firm-specific characteristics (Z_f).

The results are presented in Table 4.5. Column (1) shows that, conditional on the share of highly educated workers within the firm, there is no statistically significant relationship between individual wages and the share of highly educated workers within the own sector. This result indicates that the previously detected higher education spillover within the own sector is in fact the result of spillovers that occur within the firm. In accordance with the previous subsection, I do find a positive coefficient for the share of highly educated workers within the own sector on short distances from the working place. The higher education spillover on short distances outside the own sector is negative, although this result is not robust to other specifications. In column (2) I include controls for the share of female workers and total employment within firms, and in column (3) I incorporate region- and sector-specific fixed effects. This leaves the estimates largely unchanged.

Table 4.5. Spatial, sectoral and organizational scope of the higher education spillover

Column: Second-stage equation:	(1) (4.10)	(2) (4.10)	(3) (4.10)	(4) (4.10)	(5) (4.10)	(6) (4.10)
Region characteristics						
Share of female workers within 10 km	YES	YES	–	–	–	–
Public university within 10 km	YES	YES	–	–	–	–
Log employment within 10 km	YES	YES	–	–	–	–
Region fixed effects	NO	NO	YES	YES	YES	YES
Sector characteristics						
Share of highly educated workers within own sector	0.004 (0.010)	0.006 (0.010)	–	–	–	–
Share of female workers within own sector	YES	YES	–	–	–	–
Log employment within own sector	YES	YES	–	–	–	–
Sector fixed effects	NO	NO	YES	YES	YES	YES
Region-sector characteristics						
Share of highly educated workers within 10 km within own sector	0.076*** (0.010)	0.079*** (0.010)	0.077*** (0.006)	0.062*** (0.007)	0.062*** (0.008)	0.056*** (0.009)
Share of highly educated workers within 10 km outside own sector	–0.044** (0.021)	–0.048** (0.020)	0.173 (0.133)	0.385** (0.151)	0.054 (0.164)	0.164 (0.191)
Firm characteristics						
Share of highly educated workers within firm	0.058*** (0.002)	0.053*** (0.002)	0.057*** (0.002)	0.146*** (0.004)	0.164*** (0.004)	0.174*** (0.005)
Share of female workers within firm	NO	YES	YES	YES	YES	YES
Log employment within firm	NO	YES	YES	YES	YES	YES
Exclude micro firms (< 10 fte)	NO	NO	NO	YES	YES	YES
Exclude firm when share of workers with unknown education > 0.5	NO	NO	NO	NO	YES	YES
Include share of workers with unknown education as a control	NO	NO	NO	NO	YES	YES
Exclude firm when share of highly educated workers <0.1 or >0.9	NO	NO	NO	NO	NO	YES
Observations	402,548	402,548	402,548	151,807	84,309	54,807
R ²	0.060	0.065	0.093	0.296	0.351	0.344

Notes: The dependent variable is the region-firm fixed effect obtained from the first-stage equation. All first-stage results are the same for every column, and are available upon request. Robust standard errors, which are clustered by region-sector, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Columns (4) to (6) in Table 4.5 analyze whether the results are driven by some peculiar firm characteristics. First, I exclude all micro-firms from the sample because the region-firm fixed effects are less precisely estimated for firms with few employees compared to those with many employees. This might lead to measurement error and, as a consequence, downward biased results in the second-stage equation. And indeed, column (4) shows that the exclusion of micro-firms leads to a considerably larger coefficient for the share of highly educated workers within the firm.⁵⁸ Second, I address potential measurement error of the share of highly educated workers within the firm (see also subsection 4.3.2). I aim to limit random measurement error by excluding those firms of which more than 50 percent of the staff has an unknown educational attainment. The measurement error may also contain a systematic component because the average worker with an unknown education level differs from the typical worker. For instance, workers with an unknown education level are on average older because the coverage of the dataset on educational attainment increases over time. Systematic measurement error is targeted by including an additional control that reflects the share of workers with an unknown educational attainment. Column (5) shows that these actions have a limited influence on the results. The point estimate of the spillover within the firm increases slightly, which is what we would expect if random measurement error biased the results. Third, I examine the influence of firms that have extreme values for the share of highly educated workers. To this end, I exclude firms with a share of highly educated workers that is smaller than 10 or larger than 90 percent. These firms make up 35 percent of the remaining dataset and are thus relatively abundant. The results, as shown in column (6), are quite insensitive to this exclusion.

The identified external return to higher education is of sizable economic significance. For instance, based on column (4) of Table 4.5, it can be claimed that an increase of one standard deviation in the share of highly educated workers within the own firm leads on average to 4.0 percent higher wages. A one standard deviation increase in the share of highly educated workers within 10 kilometer within the own

⁵⁸ A more elegant solution to the problem of inaccurately estimated firm fixed effects would be to use a feasible generalized least squares (FGLS) estimator (Gobillon, 2004). This approach is, however, computationally very demanding when having more than 400,000 firm fixed effects. Moreover, other studies have shown that OLS and FGLS generally lead to very similar point estimates (Combes et al., 2008; De la Roca and Puga, 2017, and Chapter 3 of this thesis).

sector adds 1.3 percent to individual wages. These effect sizes are moderate when compared to the international literature. On a regional level, for example, Moretti (2004a) reports considerably larger external returns to higher education, whereas Groot and De Groot (2014) find smaller or even negative effects. The higher education spillover found by Heuermann (2011) is similar in size to the estimates in this chapter.

4.5 Heterogeneous effects and alternative specifications

This section analyzes whether the results in Table 4.5, which apply to the average worker/firm in the Netherlands, conceal heterogeneous effects. To this end, I provide separate estimates by firm size and knowledge intensity of the sector in subsection 4.5.1, and estimates by education group in subsection 4.5.2. The latter is particularly interesting because it accounts for imperfect substitution across workers (see Section 4.2). Subsection 4.5.3 examines the sensitivity of the results to alternative specifications, such as running the first-stage equations without worker fixed effects.

4.5.1 Estimates by firm size and knowledge intensity of the sector

The scope of the external return to higher education possibly varies across firms with different characteristics. A natural candidate to cause heterogeneous effects would be firm size. Small firms, after all, dispose of less opportunities to interact within the own organization, which may lead them to look for valuable interactions outside the firm. To investigate this hypothesis, I have split the sample in three subsamples: small (<50 fte), medium-sized (50–500 fte) and large firms (> 500 fte). Columns (1) to (3) in Table 4.6 seem to confirm that the higher education spillover within the own firm is more pronounced for larger firms than for smaller firms. Also, smaller firms appear to benefit more strongly from higher education spillovers outside the own firm.

Table 4.6. The external return to higher education for firms with differences in size and knowledge intensity

Column: Second-stage equation:	(1) (4.10)	(2) (4.10)	(3) (4.10)	(4) (4.10)	(5) (4.10)
Subsample:	Small firms (<50 fte)	Medium-sized firms (50–500 fte)	Large firms (> 500 fte)	Knowledge extensive sectors	Knowledge intensive sectors
Region characteristics					
Region fixed effects	YES	YES	YES	YES	YES
Sector characteristics					
Sector fixed effects	YES	YES	YES	YES	YES
Region-sector characteristics					
Share of highly educated workers within 10 km within own sector	0.069*** (0.009)	0.051*** (0.010)	0.015 (0.011)	0.100*** (0.011)	0.034*** (0.009)
Share of highly educated workers within 10 km outside own sector	0.469** (0.190)	0.496*** (0.211)	–0.167 (0.235)	0.511*** (0.198)	–0.295 (0.249)
Firm characteristics					
Share of highly educated workers within firm	0.120*** (0.005)	0.166*** (0.006)	0.269*** (0.009)	0.204*** (0.005)	0.116*** (0.005)
Share of female workers within firm	YES	YES	YES	YES	YES
Log employment within firm	YES	YES	YES	YES	YES
Exclude micro firms (< 10 fte)	YES	–	–	YES	YES
Observations	78,450	46,307	27,043	93,029	42,848
R^2	0.203	0.367	0.616	0.236	0.212

Notes: The dependent variable is the region-firm fixed effect obtained from the first-stage equation without worker fixed effects. First-stage results are available upon request. Robust standard errors, which are clustered by region-sector, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Next, I examine whether firms operating in knowledge intensive economic sectors gain more from higher education spillovers compared to firms in knowledge extensive sectors. To this end, I use a classification provided by Eurostat (n.d.) to divide the firms into a knowledge intensive and extensive subsample.⁵⁹ It is only

⁵⁹ The classification provided by Eurostat indicates whether a two-digit NACE sector can be regarded as a technological and knowledge intensive or extensive economic sector. Manufacturing industries are categorized according to R&D intensity, while service industries are classified on the basis of the share of tertiary educated persons. A few NACE sectors could not be assigned to one of the two categories: 6, 8, 9, 35–43.

natural to expect that knowledge spillovers are more important for economic sectors operating in technological and knowledge intensive activities. Yet, this is not what I observe from columns (4) and (5) in Table 4.6. In fact, the opposite is true: the higher education spillover is more pronounced for firms in knowledge extensive sectors. A possible explanation for this counter intuitive result is that a marginal increase in the share of highly educated workers adds less to productivity when this share is already at a high level. After all, the share of highly educated workers is much larger within knowledge intensive sectors than in knowledge extensive sectors: 50 versus 19 percent.

4.5.2 Estimates by education group

Moretti (2004a) argues that a positive relationship between individual wages and the share of highly educated workers does not necessarily indicate the existence of a spillover. Instead, the relationship may arise because high- and low-skilled individuals are imperfect substitutes in production. This implies that an increase in the share of highly educated workers increases the productivity of low-educated workers, even in the absence of a spillover effect. Highly educated workers, on the other hand, only benefit from an increase in the share of highly educated workers if the positive spillover effect dominates over the negative effect of increased supply.

In order to examine the possibility that imperfect substitution drives the results, I have split the sample into three separate education groups: low, medium and high.⁶⁰ Then, I re-estimate the equations for these three education groups. The theory of imperfect substitution predicts that low-educated workers, compared to highly educated workers, benefit more from an increase in the share of highly educated workers. This is what I observe from Table 4.7. It can also be seen that highly educated workers experience a significantly positive wage effect from an increase in the share of highly educated workers. This implies that the positive spillover effect is strong enough to overcome the negative supply effect.

⁶⁰ This corresponds to the following SOI 2006 classification: 10–41, 42–43 and 51–70.

Table 4.7. The external return to higher education across education groups

Column: Second-stage equation:	(1) (4.10)	(2) (4.10)	(3) (4.10)
Subsample:	Low-educated workers	Medium-educated workers	Highly educated workers
Region characteristics			
Region fixed effects	YES	YES	YES
Sector characteristics			
Sector fixed effects	YES	YES	YES
Region-sector characteristics			
Share of highly educated workers within 10 km within own sector	0.071*** (0.010)	0.052*** (0.008)	0.056*** (0.009)
Share of highly educated workers within 10 km outside own sector	0.232 (0.218)	0.303 (0.187)	0.316 (0.194)
Firm characteristics			
Share of highly educated workers within firm	0.178*** (0.005)	0.157*** (0.004)	0.118*** (0.005)
Share of female workers within firm	YES	YES	YES
Log employment within firm	YES	YES	YES
Exclude micro firms (< 10 fte)	YES	YES	YES
Observations	106,512	129,380	96,982
R^2	0.206	0.234	0.148

Notes: The dependent variable is the region-firm fixed effect obtained from the first-stage equation without worker fixed effects. First-stage results are available upon request. Robust standard errors, which are clustered by region-sector, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.5.3 Alternative specifications

In the remainder of this section, I examine whether the results are robust to alternative specifications. First, I report the results when calculating the share of highly educated workers on the basis of a headcount, rather than using a working hours weighted measure. After all, the potential for knowledge spillovers is possibly not so strongly related to the amount of time that individuals work together. The estimates, which are shown in columns (1) to (3) in Table 4.8, are quite insensitive to this alternative calculation.

Table 4.8. Alternative specifications (a)

Column: Second-stage equation:	(1) (4.10)	(2) (4.10)	(3) (4.10)	(4) (4.10)	(5) (4.10)	(6) (4.10)
Subsample:	Low- educated workers	Medium- educated workers	Highly educated workers	Low- educated workers	Medium- educated workers	Highly educated workers
Alternative specification:	Share of highly educated workers on the basis of a headcount			First-stage equation contains no worker fixed effects		
Region characteristics						
Region fixed effects	YES	YES	YES	YES	YES	YES
Sector characteristics						
Sector fixed effects	YES	YES	YES	YES	YES	YES
Region-sector characteristics						
Share of highly educated workers within 10 km within own sector	0.082*** (0.010)	0.063*** (0.008)	0.066*** (0.009)	0.106*** (0.013)	0.103*** (0.012)	0.126*** (0.014)
Share of highly educated workers within 10 km outside own sector	0.111 (0.166)	-0.091 (0.142)	0.070 (0.173)	-0.016 (0.291)	-0.203 (0.268)	0.525 (0.325)
Firm characteristics						
Share of highly educated workers within firm	0.189*** (0.005)	0.166*** (0.004)	0.125*** (0.005)	0.377*** (0.008)	0.432*** (0.007)	0.396*** (0.007)
Share of female workers within firm	YES	YES	YES	YES	YES	YES
Log employment within firm	YES	YES	YES	YES	YES	YES
Exclude micro firms (< 10 fte)	YES	YES	YES	YES	YES	YES
Observations	106,512	129,380	96,982	106,664	129,493	97,123
R ²	0.157	0.174	0.122	0.363	0.409	0.303

Notes: The dependent variable is the region-firm fixed effect obtained from the first-stage equation. First-stage results are available upon request. Robust standard errors, which are clustered by region-sector, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 4.9. Alternative specifications (b)

Column:	(1)	(2)	(3)
Alternative specification:	Absolute number of highly educated workers		
Subsample:	Low-educated workers	Medium-educated workers	Highly educated workers
Region characteristics			
Region fixed effects	YES	YES	YES
Sector characteristics			
Sector fixed effects	YES	YES	YES
Region-sector characteristics			
Log number of highly educated workers within 10 km within own sector	0.008*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Log number of highly educated workers within 10 km outside own sector	-0.112** (0.051)	-0.224*** (0.052)	-0.093** (0.045)
Firm characteristics			
Log number of highly educated workers within firm	0.017*** (0.001)	0.020*** (0.001)	0.025*** (0.001)
Share of female workers within firm	YES	YES	YES
Log employment within firm	YES	YES	YES
Exclude micro firms (< 10 fte)	YES	YES	YES
Observations	99,618	122,808	96,438
R^2	.203	0.236	0.148

Notes: The dependent variable is the region-firm fixed effect obtained from the first-stage equation. First-stage results are available upon request. Robust standard errors, which are clustered by region-sector, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Second, I exclude the worker fixed effects from the first-stage equations. This implies that the region-firm fixed effects are cleaned for observed worker characteristics only. This robustness check is informative because a model with worker fixed effects relies solely on individuals who move between municipalities, sectors and firms, and these movers may not be representative for the whole population. Columns (4) to (6) in Table 4.8 report the results of this exercise. Similar to the main specification, these estimates show no evidence in favor of spillovers outside the own sector. Also, the higher education spillover within the own firm is considerably larger than the spillover on short distances within the own sector. However, with this alternative specification I see no inverse relationship between individual schooling levels and the external return to higher education. Finally, in

Table 4.9, I examine the wage effect of the absolute number of highly educated workers, rather than the share. Again, I find that the positive education spillovers are stronger within the firm compared to those that occur outside the firm.⁶¹

4.6 Conclusion

This chapter examines the spatial, sectoral and organizational scope of the external return to higher education. In order to do so, I combine a large dataset on the educational attainment of Dutch workers and matched employer-employee panel data on individual earnings. This unique dataset enables me to distinguish between higher education spillovers that operate at the level of regions and sectors from those that occur within firms. The empirical identification strategy involves a two-stage estimation procedure of a Mincerian wage equation with worker fixed effects. Various forms of the second-stage equation have been estimated to unravel the scope at which the higher education spillover operates. To pin down the spatial scope of the spillover, I follow a concentric ring-based approach, as well as an optimization procedure to minimize the mean squared error of various distance decay functions.

I find that the spatial scope of the higher education spillover is very limited. Most of the identified spillovers take place within firms, being a factor 2–3 larger than those occurring outside the firm. The external return to higher education that operates outside the firm is restricted within about 10 kilometer distance from the work location and only occurs within the own sector. This finding suggests that the higher education spillover is predominantly driven by the exchange of tacit knowledge, which heavily depends on face-to-face contact. The findings are less favorable towards skill-technology complementarities because they are expected to foster aggregate productivity at the level of labor markets, i.e. regions and/or economic sectors.

⁶¹ Table 4.9 also shows that individual wages are negatively related to the number of highly educated workers on short distances outside the own sector. A potential explanation for this negative wage effect can be found in the literature on job market signaling (e.g., Spence, 1973; Weiss, 1995). Education may generate negative spillovers when educational attainment simply serves as a signal for innate skill and does not directly foster productivity (Moretti, 2004b).