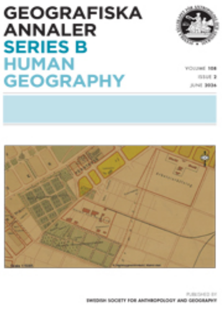

Routledge
Taylor & Francis Group

Geografiska Annaler: Series B, Human Geography

ISSN: 0435-3684 (Print) 1468-0467 (Online) Journal homepage: www.tandfonline.com/journals/rgab20

The Problem of Uncertain Contextual Characteristic (PUCC): does it matter how contextual poverty is measured for the neighbourhood effect estimation?

Jérôme Francisco Conceicao, Ana Petrović, Maarten van Ham & David Manley

To cite this article: Jérôme Francisco Conceicao, Ana Petrović, Maarten van Ham & David Manley (22 May 2026): The Problem of Uncertain Contextual Characteristic (PUCC): does it matter how contextual poverty is measured for the neighbourhood effect estimation?, Geografiska Annaler: Series B, Human Geography, DOI: 10.1080/04353684.2026.2675923

To link to this article: https://doi.org/10.1080/04353684.2026.2675923

GEOGRAFISKA ANNALER: SERIES B, HUMAN GEOGRAPHY
https://doi.org/10.1080/04353684.2026.2675923

RESEARCH ARTICLE

OPEN ACCESS

Check for updates

The Problem of Uncertain Contextual Characteristic (PUCC): does it matter how contextual poverty is measured for the neighbourhood effect estimation?

Jérôme Francisco Conceicao [a], Ana Petrović [a], Maarten van Ham [a] and David Manley [b]

[a]Department of Urbanism, Delft University of Technology, Delft, The Netherlands; [b]School of Geographical Sciences, University of Bristol, Bristol, United Kingdom

ABSTRACT
This paper investigates the sensitivity of neighbourhood effect estimates to the operationalization of contextual poverty. It introduces the Problem of Uncertain Contextual Characteristic (PUCC), which refers to uncertainty surrounding what is measured and represented when constructing contextual variables, potentially resulting in estimation bias. Using longitudinal micro-data from Dutch population registers (2011–2020), we assess four key parameters when operationalizing poverty: poverty dimensions, reference groups, poverty-line thresholds, and aggregation statistics. We undertake a systematic analysis modelling the effect of each poverty indicator while keeping all other factors constant. We also generate models including different residential context scales and geographies to compare the effects of PUCC with other sources of estimation variation. Results show that the operationalization of contextual poverty substantially influences the estimated neighbourhood effects on individual income. In our analyses, the operationalization of contextual poverty introduced greater variation than the residential context's scale or the geographical extent of the study. Findings further suggest that PUCC and the Modifiable Areal Unit Problem (MAUP) are closely related, as the impact of contextual poverty measures varies significantly across spatial scales.

ARTICLE HISTORY
Received 11 February 2025
Revised 11 February 2026
Accepted 12 March 2026

KEYWORDS
Neighbourhood effect; poverty measurements; operationalization biases; spatial inequalities; income

1. Introduction

A growing body of research highlights the important role that neighbourhood environments play in shaping individual life outcomes, ranging from education and employment to income, and health (Andersson and Malmberg 2015; Chetty and Hendren 2018; Jivraj et al. 2020; Nieuwenhuis, Kleinepier, and van Ham 2021; Urban 2009). Central to this literature is the common assumption that living in socioeconomically disadvantaged areas negatively impacts individual trajectories. Among the various outcome variables examined, individual income stands out as one of the most frequently studied indicators of neighbourhood effects (e.g. Galster, Andersson, and Musterd 2016; Hedman et al. 2015; Miltenburg and Van Der Meer 2018).

CONTACT Jérôme Francisco Conceicao j.franciscoconceicao@tudelft.nl Department of Urbanism, Delft University of Technology, 2628 BZ Delft, The Netherlands
This article has been corrected with minor changes. These changes do not impact the academic content of the article.
Supplemental data for this article can be accessed online at https://doi.org/10.1080/04353684.2026.2675923

Yet, despite decades of research, comparing findings across studies is challenging due to the wide range of methodological approaches employed. For instance, neighbourhood effects are often assumed to be weaker in Europe than in the US due to higher residential segregation and a weaker welfare state in the US (Friedrichs, Galster, and Musterd 2003; Galster 2012). However, this assumption is only weakly supported by empirical evidence, and robust cross-national comparisons are missing. Previous research has also identified several potential sources of estimation bias stemming from model misspecification, and among these issues, the Modifiable Areal Unit Problem (MAUP) – which arises from the arbitrary spatial definitions of neighbourhoods – has received considerable attention (Bolster et al. 2007; Petrović, van Ham, and Manley 2022; Stein 2014). Other studies emphasize the influence of factors such as the duration of exposure (Chetty and Hendren 2018; Musterd, Galster, and Andersson 2012; Troost, Janssen, and Ham 2023), and the choice of the econometric model (see review in Galster et al. 2008; Galster, Andersson, and Musterd 2016).

An issue which has received much less attention in the literature is the impact of the selection of contextual measures of poverty on the estimated neighbourhood effects. The neighbourhood effects literature frequently uses the neighbourhood poverty level as a contextual measure of interest. However, even this seemingly straightforward variable is operationalized in various ways (see review in Section 2.1.), and we know very little about how these differences impact the neighbourhood effect estimates. A few studies have compared the impact of different contextual variables, offering some insight. For example, Andersson and Musterd (2010) found that in Sweden's metropolitan areas, the proportion of low-income residents had a stronger effect on individual income than did the shares of non-Western migrants, the unemployed, or high-income residents. Similarly, Andrews, Green, and Mangan (2004) found that average neighbourhood income and vocational qualification levels influenced youth employment outcomes in Australia, while the proportion of highly educated residents had no significant effect. Despite these valuable insights, no study has systematically examined how different operationalizations of the same variable, such as poverty, might impact modelling outcomes, nor have they explored the conceptual implications of these choices. Such a systematic approach is necessary to be able to assess and compare the results of different studies analysing the effect of contextual poverty on individual income.

This paper addresses this gap by examining the variability in neighbourhood effect estimates that arises from the choice of contextual poverty measure. Rather than focusing solely on poverty, this study adopts a broad framework to examine how multiple sources of uncertainty jointly shape neighbourhood effect estimates. Our goal is not to quantify a contextual effect per se, as much of the literature on which we are building does, but to assess how varying definitions of poverty contribute to variation in neighbourhood effect estimates.

Drawing on the MAUP and Kwan's (2012) Uncertain Geographic Context Problem (UGCoP), we also conceptualize the problem at the core of this issue: the uncertainty surrounding what is really being measured behind neighbourhood effects and how to measure it, which in turn can lead to potentially significant misspecifications. We term this phenomenon the Problem of Uncertain Contextual Characteristic (PUCC). The PUCC operates alongside the MAUP and UGCoP, as all originate from challenges in identifying and capturing the multiplicity of neighbourhood effect mechanisms. Moreover, these dimensions are likely to interact, since both the operationalization of poverty and the spatial definition of context shape how effectively specific neighbourhood mechanisms are measured.

Through systematic analysis, we model the effects of various poverty indicators independently, while keeping all other factors constant across models. By using this approach, we contribute to understanding the extent to which neighbourhood effect studies are sensitive to the operationalization of contextual variables. Specifically, we examine the influence of four key parameters: dimensions of poverty (income, unemployment, education, and access to social benefits), reference groups (national and municipal poverty lines), poverty line thresholds (40% and 60% of median income), and the summary statistics used in data aggregation (poverty rate and Watts index). As a second contribution, we also build models that incorporate varying spatial scales for residential contexts and multiple geographic areas, allowing us to compare the impact of PUCC with other sources of variability in neighbourhood effect estimates and, more importantly, investigate whether they interact with each other.

Our results indicate that the definition of contextual poverty is not merely a technical detail as it significantly affects the estimation of neighbourhood effects. Moreover, the influence of residential scale varies across different poverty measures, showing that PUCC interacts with the MAUP and UGCoP rather than operating independently. These sources of uncertainty are thus interrelated: both the operationalization of poverty and the spatial delineation of neighbourhoods shape how effectively specific mechanisms are captured. Consequently, depending on how context is defined, neighbourhood effect studies may be measuring distinct social dynamics.

2. Theory

2.1. Uncertainty in the spatial definition of context: MAUP and UGCoP

The socio-economic composition of a residential context can vary considerably depending on its spatial definition. For example, a household may be located in a small pocket of concentrated poverty adjacent to less deprived areas. As the spatial scale used to define the neighbourhood increases, the measured level of poverty may change considerably due to aggregation of smaller areas.

However, changes in spatial scale of context do not merely alter its composition but also the relevance of the mechanisms through which contextual conditions influence individual outcomes. Galster (2012) identifies fifteen potential mechanisms through which poverty may affect individuals, grouped into four broad categories: social-interactive, environmental, geographical, and institutional effects. These mechanisms are not expected to operate uniformly across space. Some, such as peer-group socialization and network effects, are likely to function at relatively small spatial scales, whereas others, including labour-market dynamics or stigmatization, may operate at larger scales or simultaneously across multiple scales (Petrović, Manley, and Van Ham 2020).

The multiplicity of mechanisms and the various geographical scales at which they can occur introduce considerable complexity into neighbourhood-effects modelling, increasing the risk of misspecification. In this regard, empirical studies explicitly addressing the challenge of defining spatial context have already highlighted the lack of theoretical guidance on the appropriate geographical scale at which neighbourhood effects operate (see Plum and Knies 2015). As Petrović, Manley, and Van Ham (2020, 1115) note, selecting a single spatial scale risks 'cutting through various mechanisms, capturing relevant scales for some and less relevant scales for others'.

Since the late 1980s, researchers have examined this problem of spatial-context-related model misspecification, most notably through the MAUP (Openshaw 1984; Wong 2004). The MAUP refers to inconsistencies in empirical results arising from the aggregation of data using different areal unit configurations and consists of two interrelated components: scale effects and zonation effects. Scale effects occur when results differ across levels of spatial aggregation, while zonation effects concerns biases originating from the delimitation of an area boundary.

In neighbourhood-effects research, both components of the MAUP have been shown to influence estimated contextual effects. Several studies have shown that the spatial scale at which contextual variables are measured has an effect on modelled neighbourhood effects (Bolster et al. 2007; Petrović, van Ham, and Manley 2022; Stein 2014). Others like Andersson and Musterd (2010) demonstrate that neighbourhood-effect estimates vary depending on whether residential contexts are delineated using grid cells, administrative units such as Small Market Statistics (SAMS), or municipalities. In the absence of clearly defined 'natural boundaries' for neighbourhoods, uncertainty surrounding spatial delineation, combined with the instability of results under the MAUP, introduces substantial variation in the estimated effect of neighbourhood context.

Building on this literature, Kwan (2012) introduced the Uncertain Geographic Context Problem (UGCoP), highlighting the limited understanding of the true causally relevant context(s) in terms of space and time and the statistical biases generated in the aggregation of data. This lack of knowledge of the specific spatial and temporal context(s) within which a phenomenon occurs can result in, at worst, arbitrary or at best, partial, selection of context. As a result, misspecification of the geographical context can generate inferential errors, leading to the apparent presence of an effect when there is none (type I) or the omission of an effect when there is an important relationship present (type II).

The issue of contextual uncertainty is particularly a central theoretical and empirical challenge in studies examining the impact of context due to the complex nature of the underlying processes (i.e. Galster's (2012) fifteen potential mechanisms) and the likelihood that different mechanisms operate across multiple, overlapping spatial scales.

2.2. The problem of uncertain contextual characteristic (PUCC)

The uncertainty about what is being measured when assessing neighbourhood effects also extends to the operationalization of poverty itself, which can lead studies examining the same question to measure different phenomena and produce estimation biases. This issue is the focus of this paper and is referred to here as PUCC. Researchers examining contextual poverty face a multitude of decisions regarding its operationalization. Measuring (contextual) poverty is inherently complex, involving decisions around the various socio-economic dimensions, the selection of reference groups for assessing income levels and the choice of representative aggregate measures at the contextual level. As with the MAUP, the wide range of possible configurations in this regard results in potentially arbitrary but significant decisions. Even though different poverty indicators are correlated, we can expect their spatial distributions not to fully overlap, especially across socio-economic dimensions. For example, a neighbourhood may have a high share of low-income residents while appearing less deprived on other dimensions, such as unemployment rates.

Connecting these alternative poverty measures to the diversity of mechanisms through which contextual poverty affects lifetime economic outcomes further amplifies the problem of operationalization. A given poverty indicator may be more causally relevant for certain mechanisms than for others. This diversity of mechanisms makes the search for a single 'true' contextual condition virtually impossible. Instead, we may need to find the most appropriate indicator(s) that capture the majority of the mechanisms at play. Moreover, the PUCC can also interact with other factors of variations, adding further layers of complexity. As discussed above, specific causal pathways may operate at different geographical scales, and their relative importance therefore depends on the spatial delineation of the residential context. As the mechanisms captured may differ across contextual scales, the relevance of any given operationalization of poverty is also expected to vary with scale.

Though it goes beyond the scope of the paper, other factors could also have an important impact on the importance of each mechanism, such as the duration of exposure (as it takes time for social ties to influence the individual via socialization or network effect, while other effects, such as stigmatization and spatial mismatch, are expected to be instantaneous) and the timing of exposure (exposure to poverty might be more impactful on children via socialization and social contagion).

As a consequence of this uncertainty as well as data availability, the literature on contextual effects has used a wide range of contextual poverty measures to investigate the same phenomenon. In terms of socio-economic dimensions, studies have evaluated contextual poverty through income and housing (Van Ham and Manley 2010), education (Andersson et al. 2007; Andersson and Malmberg 2015; Andrews, Green, and Mangan 2004), as well as labour market participation (Andersson et al. 2007; Weinberg, Reagan, and Yankow 2004). Some studies focus on a single dimension (Galster et al. 2008; Galster, Andersson, and Musterd 2016; Petrović, van Ham, and Manley 2022), whereas others employ a multidimensional framework by incorporating multiple contextual variables in the same regression or constructing a composite index (Andersson et al. 2007; Andrews, Green, and Mangan 2004; Bolster et al. 2007; Miltenburg and Van Der Meer 2018).

Among the income-based contextual poverty indicators, the poverty rate is frequently employed as a straightforward means to operationalize contextual poverty. However, even within this measure, variations exist in terms of the income threshold used to define the poverty line. For instance, some studies have placed the poverty line at 40% of the national median income (Miltenburg and Van Der Meer 2018; Petrović, Manley, and Van Ham 2022; Petrović, van Ham, and Manley 2022), while others have employed 50% of the national median income (Musterd, Galster, and Andersson 2012), the 30th percentile of the national income distribution (Andersson et al. 2007; Hedman et al. 2015; Musterd, Galster, and Andersson 2012), or official government-recognized poverty lines (Chetty, Hendren, and Katz 2016; Quillian 2003). Like the size and shape of an area for the MAUP, different parameters in the operationalization of poverty can impact the evaluation of poverty, which in turn can lead to variation in the estimation of the effect of contextual poverty.

While there is relative consensus in the literature that the socio-economic composition of a neighbourhood can influence individual income (negatively when focusing on contextual poverty), further comparing the empirical evidence remains challenging, if not impossible, as each study has its own approach to measuring neighbourhood effects depending on its own research goal, data accessibility, and specific context. In addition to the operationalization of poverty, we find a large variety of methodological differences even when we compare studies from the same country and use the same data source (e.g. population registers data from Netherlands Statistics).

For instance, Miltenburg and Van Der Meer (2018) defined residential context using administrative neighbourhoods (buurten, with an average population of about 1,500) and measured contextual poverty through a composite deprivation index. This index included the average personal income, the proportion of low-income residents (bottom 40% nationally), the share of high-income residents (top 80%), and the percentage of welfare benefit recipients. To assess the cumulative and lingering effects of residential histories, they employed a cross-classified multilevel model that assessed the influence of multiple past neighbourhoods simultaneously. Their dependent variable, i.e. individual income, includes yearly income from work as well as social security benefits, private funds, and alimony, minus premiums for income insurance.

In comparison, the study of Petrović, Manley, and Van Ham (2022) examined how the estimated neighbourhood effect varies with the spatial scale at which context is measured. Instead of relying on a single, predefined administrative unit, they constructed 101 bespoke areas around each individual. More importantly, they decided to measure contextual poverty via a poverty line based on 40% of the national income distribution. They applied a within-individual fixed-effect model to control for the time-invariant variables. As for the dependent variable, we also find some differences: although they also consider the yearly income from work, they did not include other sources of income like social security benefits. Moreover, their study focused solely on working-age men living in the largest urban regions of the Netherlands, whereas Miltenburg and Van der Meer included the broader working-age population.

This brief comparison illustrates the wide range of methodological decisions, concerning spatial units, poverty definitions, populations studied, and modelling strategies, that complicate meaningful comparison across studies. Understanding how each of these variations may introduce bias or reflect different underlying causal mechanisms is essential to building a more cumulative and coherent body of evidence on neighbourhood effects.

2.3. Four potentially important parameters

To explore the effect of PUCC, we consider four potentially significant parameters: the dimensions of poverty, the reference groups, the level of the poverty line and the summary statistics employed in the aggregation of data.

2.3.1. Socio-economic dimensions

Studies investigating neighbourhood effects have considered poverty through diverse socioeconomic dimensions (e.g. income, housing, education, and labour market participation). This plurality of indicators can be, at least in part, explained by the numerous outcome variables that have been investigated, the wide array of ways in which contexts can impact individuals (recall Galster (2012) and the 15 mechanisms), as well as the multidimensionality of poverty (Cantillon 2011; Kakwani and Silber 2008). The latter is based on the idea that the different dimensions of inequalities can complement and reinforce each other. However, even though they are all positively correlated – increases in one dimension of poverty are often associated with increases elsewhere – they can reflect different concepts and measures.

Some socio-economic dimensions may better reflect some mechanisms of the neighbourhood effect than others. As an example, income-based contextual poverty indicators may be more appropriate to estimate network effects compared to education-based indicators, given that they measure economic resources specifically and information (e.g. on well-paid job opportunities) accessible via ties with neighbours (e.g. recommendation; Burt 2004; Granovetter 1973; Lin 2008). In contrast, we can expect the significance of the education level of neighbours to be more suited for measuring social contagion and collective socialization, processes through which behaviours and attitudes spread through social networks and interactions or via role models and community norms (e.g. recommendation; Galster 2012). However, those two mechanisms are more relevant before adulthood and appear after a longer exposure.

2.3.2. Conception of poverty and reference groups

A key consideration when we construct an income-based contextual poverty indicator is whether it should measure an absolute or a relative conception of poverty. In other words, should the poverty line be set based on an income distribution or a consumer price index. However, the absolute approach is rarely employed in the literature analysing developed countries. Poverty is generally considered a relative notion as it is contingent on the general level of prosperity in a country or population group at a given time (Atkinson et al. 2002; Cantillon 2011). This view is explicit, for instance, in Townsend's definition of poverty: the resources of economically deprived individuals are 'so seriously below those commanded by the average individual or family that they are, in effect, excluded from ordinary patterns, customs and activities' (Townsend 1979, 31).

Assuming that poverty is relative, a crucial question remains: what is the most appropriate reference group to evaluate an individual's position? One common method is to set the poverty line relative to the national income distribution. However, using local poverty lines may more appropriately estimate poverty than using a national one (see for example Jesuit, Rainwater, and Smeeding (2002) and Ayala, Jurado, and Pérez-Mayo (2014)), as it would better approximate the local living standards and the conditions faced during daily lives. This is far from a trivial issue, as these two approaches measure fundamentally different forms of inequality. Studies employing national poverty lines better reflect inequalities between localities rather than the inequalities within them. Depending on the degree of inequalities between municipalities, we might not detect poverty in wealthier cities with a national poverty line, even though some people might be relatively poor locally. If we therefore consider that regional labour market effects are not conceptually part of neighbourhood effects but merely a confounding factor, using an income-poverty indicator based on the municipal income distribution might be favoured over one based on a national one.

2.3.3. Poverty thresholds

After selecting a reference group, we need to set the poverty line within the income distribution. Studies that explore neighbourhood effects using an income poverty indicator have employed a wide variety of poverty lines (e.g. set at x% of the median income, an income percentile, or using official nationally recognized poverty lines). This decision is critical and can have significant impacts on any subsequent study because it determines the intensity of poverty considered in the analysis. Despite its importance, little is known about the effects of this choice. In theory, a poverty line that is too low can underestimate the existence of poor areas. At the same time, a poverty line that is too high would not reflect poverty, as it could include more affluent middle-class members. As Cantillon (2011) stresses, a poverty indicator should always relate to 'severe economic hardship' and 'reflect what people need, given prevailing local patterns of living, in order to be able to participate minimally in social and economic life' (21). However, in the case of neighbourhood effects, we still do not know what level of contextual poverty is most impactful for the individual (e.g. should we focus on the severe cases of poverty or use a more inclusive poverty indicator?).

To address this, we compare two commonly used poverty thresholds: 40% and 60% of the median income. The 60% threshold is widely used in European poverty research and is considered the standard for identifying individuals at risk of poverty (Goedemé and Rottiers 2011). The 40% threshold, however, is frequently used in neighbourhood effects studies in the Netherlands (Miltenburg and Van Der Meer 2018; Petrović, Manley, and Van Ham 2022; Petrović, van Ham, and Manley 2022). The rationale behind this choice is that this lower threshold better captures severe material deprivation due to the prosperity of the Netherlands. This choice can have significant consequences on the evaluation of poverty as the share of the Dutch population classified as poor drops from 13.3% under the 60% threshold to just 3.3% under the 40% threshold (Eurostat). Thus, the lower threshold identifies only a small, highly deprived segment of the population. It remains an open question whether the choice of threshold generates stronger or weaker neighbourhood effects. It may be that different poverty thresholds correspond to qualitatively distinct neighbourhood phenomena.

2.3.4. Summary statistics

The choices above relate to the imposition of thresholds and comparisons. However, the contextual poverty effect estimation may also be impacted by the choice of summary statistics used to aggregate data. Although the share of low-income individuals is commonly used in the neighbourhood effect literature, other poverty indicators with distinct properties have also been considered outside of this academic debate and might be worth applying to neighbourhood effect analyses. For instance, there are distribution-sensitive indicators, such as the Watts index[1] (Watts 1969), which aims at providing a more comprehensive view of poverty by considering the severity of poverty and the inequality within the low-income group. For instance, any rise in income among the low-income population would result in a drop in the Watts index (Monotonicity axiom). The Watts index also penalizes inequality within the low-income group by giving greater weight to the poorest individuals. A transfer of income from a deprived individual to an even poorer one, should lead to a decrease in the Watts index (transfer axiom). Additionally, this distribution sensitivity indicator has been recognized to meet several other criteria for a good poverty indicator, including its scale invariance, translation invariance, and normalization (Ravallion, 2015; Zheng 1993). Nonetheless, it is essential to note that the improved properties of the Watts index come at the cost of its interpretability. In contrast, the poverty rate is a relatively 'crude' indicator of poverty, as it is insensitive to any variation in income below the poverty line (Sen 1976). When it comes to properties mentioned above, the poverty rate only satisfies the scale invariance and normalization axioms (Ravallion, 2015; Zheng 1993). We can, therefore, assume that the selection of a summary statistic will result in different regression outcomes, as these two aggregation methods (indices of relative deprivation such as the Watts index, and poverty rates) have very different properties and reflect different conceptions of contextual poverty.

3. Data

To assess the potential impact of the PUCC on modelling neighbourhood effects, we employ individual-level longitudinal register data covering all individuals aged 25 to 65 who resided in the Netherlands from 2011 to 2020. These data allow us to follow individuals across nine annual observations, with a one-year time lag between neighbourhood exposure and individual income outcome. In addition to these criteria, we exclude observations for individuals who are unemployed or economically inactive (e.g. students, early retirees, or long-term sick or disabled), even if they may receive some income from work as a secondary source, as our focus is on income derived from active employment. This data, provided by Statistics Netherlands (CBS; see Bakker, Van Rooijen, and Van Toor (2014)), includes a wide range of socio-economic information, such as personal and household income, education, employment status, and demographic characteristics. It also provides detailed residential location data, with resolutions of 100 m by 100 m and 500 m by 500 m grids. This comprehensive dataset allows us to compare a broad set of indicators, which are typically limited by data availability, and analyse their effects on the entire population. However, it is important to acknowledge the potential selection bias, particularly due to missing data on education levels, which particularly affects older individuals and first-generation migrants (see descriptive statistics in Annex 3.1, available in the Online Supplementary Material). Regarding the residential context definition, we used the national spatial grid system (100 m by 100 m and 500 m by 500 m grids) instead of administrative boundaries as it offers standardized areas which can be scaled consistently.

Before introducing the contextual variables used in our analysis, we outline the broader economic and spatial characteristics of the Netherlands. The Netherlands is an affluent and relatively egalitarian country. In 2020, its real GDP per capita was €46,810, notably higher than the EU-27 average of €30,510 (Eurostat 2026). The Dutch welfare system is characterized by generous social assistance, public pensions, and income redistribution mechanisms. As a result, the Netherlands has one of the lowest levels of income inequality among OECD countries, with a Gini coefficient of 0.26 in 2020, compared to 0.3 for the EU-27 and 0.4 for the United States (World Bank 2026). The country's social model, including its large de-commodified social housing sector and policy interventions in disadvantaged neighbourhoods, has also been linked to relatively low levels of residential segregation (Musterd and Ostendorf 2013). For instance, empirical studies indicate that cities such as Amsterdam experience lower levels of economic segregation than many other European cities (Haandrikman et al. 2023; Musterd et al. 2017). Nonetheless, there is notable variation across urban regions in the Netherlands. San Millán, Cottineau-Mugadza, and Van Ham (2025) report that Gini coefficients of Dutch urban regions can range from 0.2 to 0.4, and that ROITI segregation values vary from 0.03 in Gouda to 0.11 in Groningen. It should also be noted that Dutch cities have experienced a liberalization of the housing sector since the 1990s. This has led, among other things, to a reduction in the social housing stock in favour of owner-occupied and market-rate housing, a growing concentration of low-income households in social housing, and gentrification processes, especially in cities like Rotterdam (Custers and Willems 2024; Van Gent and Hochstenbach 2020).

A significant share of the population resides in the Randstad, a densely populated conurbation encompassing the four largest cities (Amsterdam, Rotterdam, The Hague, and Utrecht) as well as other urban centres such as Leiden, Haarlem, Dordrecht, and Amersfoort. Although the country's small geographical size and high population density reduce the physical distance between urban and rural areas, regional income disparities persist. For example, North Holland, which includes Amsterdam and Haarlem, reported a GDP per capita of €70,285 in 2020, nearly double that of the northern province of Drenthe (€37,146). Despite these differences, regional economic disparities in the Netherlands remain comparatively moderate within the European context (Muštra and Škrabić 2014).

In this study, we constructed twelve contextual variables and compared their effects on the individuals' annual income from work (see Figure 1, and Annexes 2 and 4 for more details available in the Online Supplementary Material). These independent variables are measured with a one-year time lag (at t-1) to account for the lagged effect of spatial context on individual income. We created these contextual variables by modifying three parameters: the summary statistic; the poverty line; and the reference group. For the first parameter, we selected the poverty rate, as it is the most

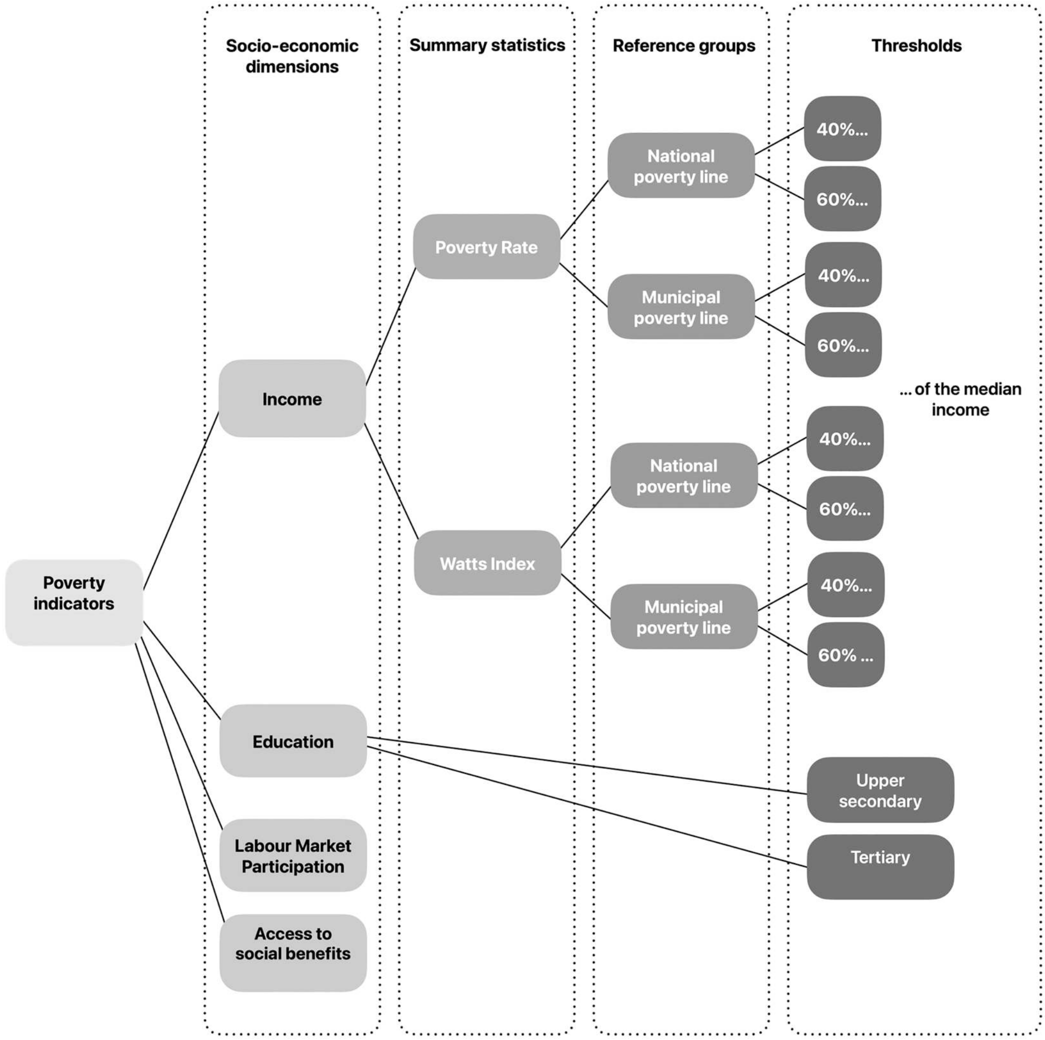

Figure 1. Construction of contextual poverty indicators.

common measure of poverty in academia and public institutions. We compared this to the Watts index to explore whether employing a distribution-sensitive indicator would significantly impact the neighbourhood effect estimation. Finally, we tested different poverty lines by modifying the reference group and the poverty level. We evaluated the significance of the geographical scope of the reference group by setting the poverty lines relative to the national and local median incomes. We also assessed whether setting poverty lines at 40% or 60% of the median income impacts the regression outcomes. To compute these indicators, we use the individuals' annual disposable household income after taxes and transfers instead of their personal income from work, as it better reflects the purchasing power of the individuals. Additionally, the household income is standardized (based on the number of adults and children) to facilitate comparing households of different sizes and compositions.

To assess the potential impact of the PUCC on modelling neighbourhood effects, we use individuals' annual income from work (logged and adjusted for inflation) as the main outcome variable. This includes all sources of employment-related income.[2] We retain individuals with zero work related income in the dataset, considering them valid observations, as excluding them would introduce substantial selection bias, since unemployment is a relevant dimension of neighbourhood labour market effects.[3] As such, the study results can be interpreted as an upper-bound estimate of neighbourhood effects on income. Following the same reasoning, we exclude individuals who are economically inactive – such as students, retirees, or those on long-term sick leave or disability – by treating their work-related income as missing. These groups differ structurally in their labour market participation, and including them could obscure the estimation of neighbourhood effects. Negative income values are also treated as missing since they are primarily reported by self-employed individuals experiencing business losses. According to the Centraal Planbureau (CPB), even among the top 1% of earners, some individuals report negative incomes due to temporary losses in specific years (Planbureau 2024).

In contrast, we use equivalized annual disposable household income for the contextual income-based poverty indicators (i.e. poverty rates), as it aligns with standard practices in poverty research and public institutions such as Eurostat (see At-risk-of-poverty rate) (Decancq et al. 2014; Goedemé and Rottiers 2011). This measure adjusts for household size and composition, capturing the shared economic resources within a household and providing a more accurate reflection of an individual's purchasing power and the risk of social exclusion and material deprivation.

Whilst a clear focus of much of the literature, income is not the only measure through which contextual poverty can be measured. We also used contextual data on other dimensions of socio-economic characteristics: labour market participation, access to social benefits and the level of education. We measured the first two aspects by calculating the unemployment rate and the share of social benefit recipients in the residential context. The unemployment rate reflects the proportion of unemployment benefit recipients in the active population, while the second indicator considers all individuals receiving social benefits (see Annexes 2 and 4, available in the Online Supplementary Material). To assess the education level of neighbours, we use two education thresholds as we do for the income-based indicators. We calculate the proportion of individuals without an upper secondary or a tertiary education degree. This study also controls for various socio-economic factors, including an individual's age, gender, country of origin, education level, and the urban region in which they reside (see Annex 1 and 4, available in the Online Supplementary Material).

We included controls for the individual's age (and age squared), education level, household type, and urban region (more explanation in Annex 1, available in the Online Supplementary Material). These variables are well-known confounding factors in neighbourhood effect studies, as they can influence both individual mobility patterns and income levels. Dummy variables for the urban region of residence are added to control for regional economic effects. Additionally, as further explained below, this study controls for the effect of time-invariant factors by employing within-fixed effect specifications.

4. Method

To assess the PUCC, we undertake a systematic analysis assessing the relationship of each poverty indicator with income separately through regression models while keeping all other factors constant. Given that the main objective of this study is to explore potential biases originating from the operationalization of contextual poverty in neighbourhood effect research, we employ the widely applied within-individual fixed-effects specification. This method isolates the impact of changes in neighbourhood context on changes in an individual's income over time. In this sense, the estimated coefficients reflect to what extent a shift in exposure to contextual poverty – either through residential mobility or changes in neighbourhood socio-economic composition – is associated with an income change for the same individual. This method minimizes potential unobserved heterogeneity and omitted variables biases. Consequently, the time-invariant variables, i.e. the individual's gender, and country of origin, are omitted from these models. A drawback of this method is that it restricts our analysis to within-individual variations, causing us to lose part of the relationship being studied. However, this limitation does not prevent the study from exploring the variations in the neighbourhood effect estimation originating from the operationalization of contextual poverty.

The structure of these fixed effect models is presented in the following equation:

$$
\begin{aligned}
(Y_{i,t0} - \bar{Y}_i) = {} & \beta 0 + \beta 1(Contextual\ Poverty_{i,\ t-1} - \overline{contextual\ Poverty_i}) + \beta 2(Age_{i,t0} - \overline{Age_i}) \\
& + \beta 3(Age^2_{i,t0} - \overline{Age^2_i}) + \beta 4(Education{:}Medium_{i,t0} - \overline{Education{:}Medium_i}) \\
& + \beta 5(Education{:}\ High_{i,t0} - \overline{Education{:}\ High_i}) + \beta 6(Household\ Type{:}\ Single\ Parent_{i,t0} \\
& - \overline{Household\ Type{:}\ Single\ Parent_i}) \\
& + \beta 7(Household\ Type{:}\ Couple\ without\ child_{i,t0} - \overline{Household\ Type{:}\ Couple\ without\ child_i}) \\
& + \beta 8(Household\ Type{:}\ Couple\ with\ child_{i,t0} - \overline{Household\ Type{:}\ Couple\ with\ child_i}) \\
& + \beta 9(Household\ Type{:}\ Other_{i,t0} - \overline{Household\ Type{:}\ Other_i}) + \beta 10(Urban\ Region{:}Rotterdam_{\ i,t-1} \\
& - \overline{Urban\ Region{:}Rotterdam_i}) \\
& + \beta 11(Urban\ Region{:}The\ Hague_{i,t-1} - \overline{Urban\ Region{:}The\ Hague_i}) \\
& + \beta 12(Urban\ Region{:}Utrecht_{i,\ t-1} - \overline{Urban\ Region{:}Utrecht_i}) \\
& + (e_{i,t0} - \bar{e}_i)
\end{aligned}
$$

where Y is the individual income from work, the βs (from $\beta 1$ to $\beta 12$) are the beta coefficients and $(e_{i,t0} - \bar{e}_i)$ is the error term. Depending on the model, the beta coefficient for the contextual poverty ($\beta 1$) is standardized to facilitate the comparison between the different indicators. The contextual poverty indicator and the number of urban regions selected vary across models.

We applied the same regression modelling approach to three different geographies. First, we assessed the neighbourhood effect in the four largest urban regions[4] of the Netherlands (Amsterdam, Rotterdam, The Hague, and Utrecht). Second, we enlarged the scope of the analysis to the eight urban regions composing the Randstad conurbation, which also include Leiden, Haarlem, Dordrecht, and Amersfoort. Finally, we examined this contextual effect for the whole of the Netherlands. This enabled us to assess the impact of another source of bias on PUCC. The geographic selection of a sample may lead to an underestimation of the neighbourhood effect by including the rural areas, as it might only be an urban phenomenon. Moreover, the impact of contextual poverty could work at a different spatial scale and different extent in rural and urban areas. This could relate to different sizes of neighbourhoods and population density, as well as different transportation means through which individuals move through urban and rural spaces. Yet, the literature often limits the geographical scope of the study to a few core agglomerations. This can also generate important selection biases by over-representing the less mobile individuals (the analysis omits any individual who moved out of the area during the timeframe studied). Finally, to account for possible biases originating from the selection of a geographical scale (MAUP), the study examined two spatial contexts, one measuring 100m-by-100 m and the other 500m-by-500 m. Previous research has shown that the strongest neighbourhood effect is found within this range (Petrović, van Ham, and Manley 2022).

5. Results

Figure 2 shows the regression outcomes regarding the relationships between the twelve contextual poverty indicators and the individual's income one year later (summary table in Annex 5, available in the Online Supplementary Material).[5] It summarizes the results of seventy-two fixed effect models within which we vary the following: the operationalization of poverty, the spatial scale, and the area studied. Only the estimates relating to the relationships between the contextual poverty

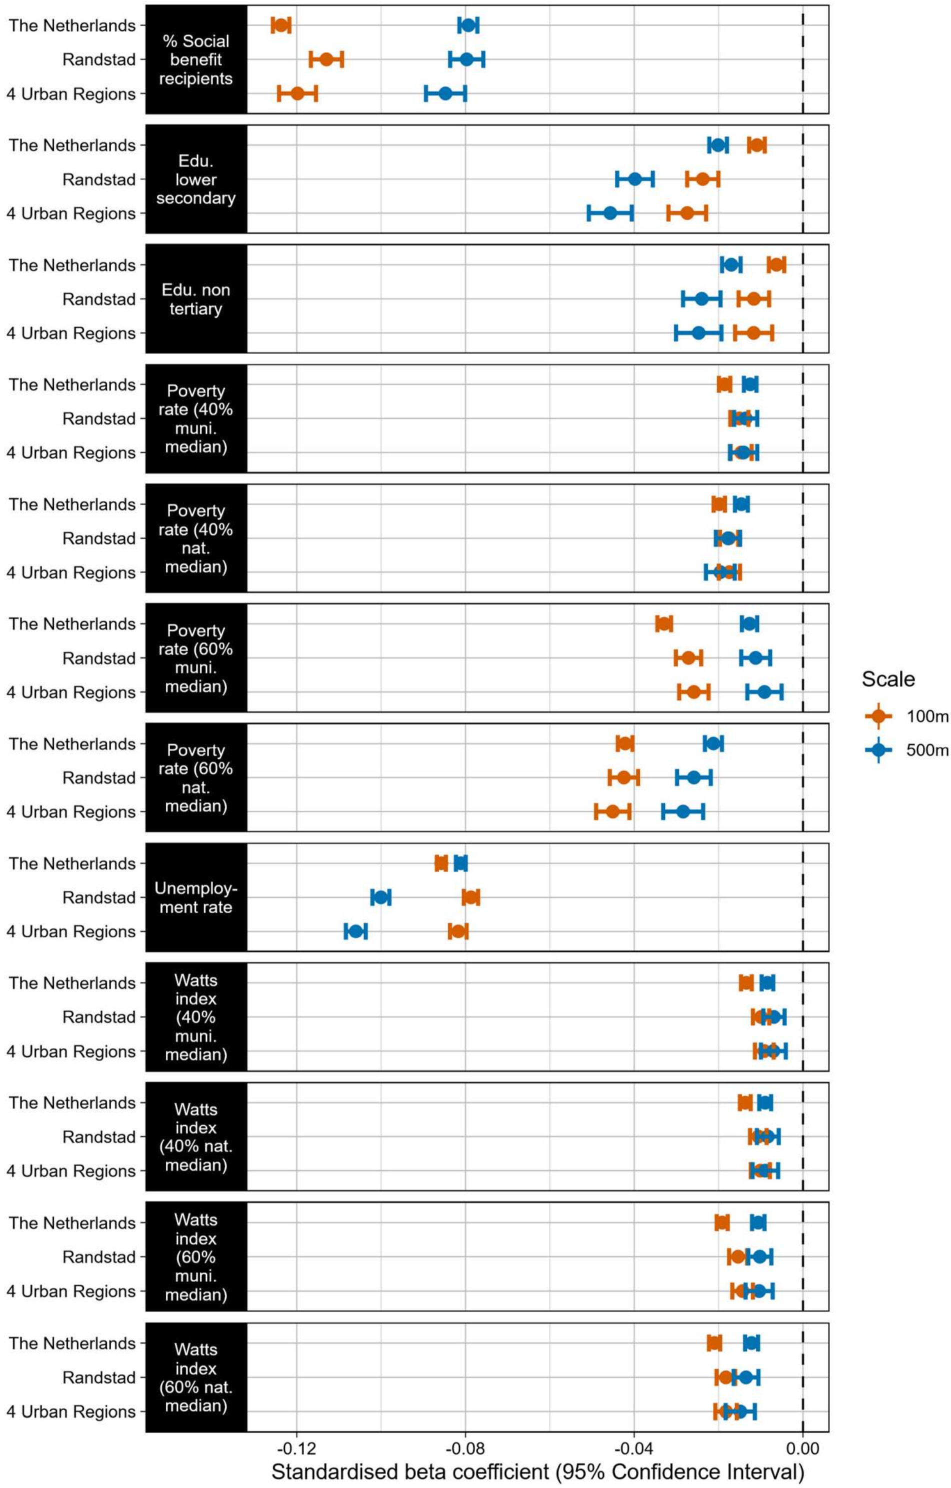

Figure 2. Regression outcomes regarding the effect of contextual poverty indicators on the individual's income from work.

indicator and income are reported.[6] To make the comparison of the different models possible, we standardized the beta coefficients, because each contextual variable has different units and distribution. This standardization is achieved by rescaling the contextual variables based on their standard deviation (one unit is equal to one standard deviation). The previous intermediate step of the analysis employed unstandardized beta coefficients, which are shown in Annexes 6 and 7 (available in the Online Supplementary Material). In the forest plot (Figure 2), each grey row corresponds to one poverty indicator and displays the results for the three areas studied (the Netherlands, Randstad, four largest urban regions). In addition, a colour code is applied to distinguish the regression results using different residential context scales i.e. 100 m (dark orange) and 500 m grid cells (blue).

As shown in Annex 5 (available in the Online Supplementary Material), although the beta coefficients vary depending on the study design (i.e. how poverty is operationalized, the neighbourhood scale, and the geographic scope of the analysis), the overall R^2 remains unchanged at 0.839, and the within-unit R^2 varies only slightly, from 0.008 to 0.010. This suggests that the explanatory power of the model as a whole including the variation explained within units (the individual level) remains largely stable across different model specifications.

5.1. Variation across contextual poverty indicators

When we focus on the results of the regression models for the four largest urban regions, the immediate residential context (100m-by-100 m grid cells), we notice that the percentage of social benefit recipients in the residential context has the highest standardized beta coefficient among the contextual variables ($\beta = -0.120$ and 95% CI: −0.124, −0.115). This coefficient indicates that an increase of one standard deviation in the percentage of social benefit recipients (SD = 13.24%) is associated with an 11.3%[7] decrease in the individual's income from work one year later, holding all other variables constant. In contrast, the Watts indexes with a poverty line set at 40% of national or municipal median income show the weakest relationship with the dependent variable. The regression outcomes suggest that a one-standard-deviation rise in one of these three indicators is associated with a drop in annual income of around 1%, as their βs of −0.010 and −0.009. Z-tests[8] reveal that the gap between these three beta coefficients is statistically insignificant ($-1.96 < Z < 1.96$) (see Clogg, Petkova, and Haritou 1995; Paternoster et al. 1998). Hence, for the set of models focusing on the four largest cities and analysing the smallest residential scale (100m-by-100 m), we observe that the effect of the percentage of social benefit recipients is around 12 times greater than that of the Watts indexes using 40% of the (national or municipal) median income as a poverty threshold.

However, this range varies with the geographical scope and scale of the residential context selected. For instance, when we consider the whole of The Netherlands, the variation in results across contextual variables is even wider. The relationship between the percentage of social benefit recipients and individual income from work ($\beta = -0.124$) is 20 times stronger than the one with the percentage of individuals without an upper secondary education degree ($\beta =$ −0.006, and 95% CI: −0.008, −0.004) ($Z = -113.844$). We also notice a slightly stronger gap in models analysing the larger residential context. More importantly, we observe notable variation in regression outcome regarding the unemployment rate and the share of social benefit. For instance, when we employ the 500m-by-500 m grid cells as context and the residents of the four urban regions as sample, the association between the percentage of social benefit recipients and income weakens to a $\beta = -0.085$ (95% CI: −0.089, −0.080), while the relationship between the unemployment rate and income increases and becomes the strongest association at this contextual scale, with a β of −0.106 (95% CI: −0.124, −0.115). The relationship is 15 times stronger than the weakest association, i.e. the Watts index using 40% of municipal median income as a poverty line with a β of −0.007 ($Z = -58.574$). These variations across the geographical scope and residential context are further discussed below.

The other contextual variables have varying degrees of association with income, falling within the range of these extremes. For instance, coming back to regression models regarding the four largest urban regions and 100m-by-100 m referential context, the unemployment rate has the second strongest negative relationship with the individual's income. This is evidenced by a standardized beta coefficient of −0.082 (95% CI: −0.084, 0.080). It means that an increase in this poverty rate of one SD is correlated with a decrease of 8.2% in annual income for an average individual. It is followed by the poverty rate using 60% of the national median income displays a beta coefficient of −0.045 (95% CI: 0.049, −0.041). The share of neighbours without upper-secondary education degree and the poverty rate using 60% municipal median income comes are next, with a respective β of −0.026 and −0.027 (the difference between these values is statistically insignificant: Z = 0.539). In other words, a one SD rise in these two indicators is associated with around a 2.6% decrease in income. Then, the poverty rate using 60% municipal median income and two poverty rates based on 40% of the national and municipal median incomes show the weakest associations with income, as they have a respective β of −0.017 (95% CI: −0.020, −0.015) and −0.015 (95% CI: −0.017, −0.012). It means that a one-SD increase of each poverty rate in the residential context is associated with a respective loss of 1.7% and 1.5% of individual income. However, the difference between these βs is statistically insignificant (Z = −1.537).

The analysis of the Watts indexes reveals a comparable trend to the poverty rates. Yet, the individual's income has a less pronounced negative association with the Watts indexes than with the poverty rates based on identical poverty thresholds. The Watts index with a poverty line set at 60% of the national median income has the strongest relationship with income, with a value of −0.018 (95% CI: −0.021, −0.016). It indicates that a one-SD (=0.08) rise in the Watts index is associated with −1.8% loss in an average annual income. Meanwhile, the Watts index based on 60% of municipal median income has a slightly weaker association with a β of −0.014 (95% CI: −0.017, −0.012) – which is statistically different from the β of the first indicator (Z = −2.178). For a one-SD (=0.08) increase in this index, we would, therefore, expect the individual's annual income to reduce by 1.4%. Lastly, as mentioned above, the two indexes that use poverty thresholds set at 40% of the national or municipal median incomes have beta coefficients of −0.010 (with a 95% CI of −0.012 to −0.008) and −0.009 (with a 95% CI of −0.011 to −0.007) – these are significantly weaker than the β of the second Watts index (with respective Z-values of 2.467, and 3.046).

The relationship between the education-based poverty indicators and the individual income from work depends significantly on the education threshold selected. While the share of individuals without a tertiary education degree is among the indicators with relatively weak estimates (β of −0.012 and 95% CI: −0.016, −0.007), the relationship between the share of individuals without an upper secondary education level and the individual's income is more than twice as strong (β = −0.027, and 95% CI: −0.032, −0.023). Based on these results, we expect the annual income of an average individual to decrease by 1.2% and 2.7% for every SD increase in these contextual indicators (with a respective SD of 23.83 and 17.83).

5.2. Variations across spatial contexts

The relationship between an individual's income from work and the contextual poverty indicators can vary with the scale of the residential context. This is particularly true when it comes to the percentage of social benefit recipients. As we increase the residential scale to a 500m-by-500 m cell, the effect of the negative relationship reduces by around a third, i.e. from a β of −0.120 (95% CI: −0.124, −0.115) to a β of −0.085 (95% CI: −0.089, −0.080) (Z = −10.786). We can notice a similar pattern for the poverty rate using 60% of the national and municipal median incomes as poverty lines. The standardized beta coefficient for the poverty rate based on the national bounds drops from −0.045 (95% CI: −0.049, −0.041) to −0.028 (95% CI: −0.033, −0.024) when we increase the size of the residential context (Z = −5.304). This represents a 40% decrease in the strength of the association. Regarding the poverty rate based on the municipal medians, its negative relationship with income weakens by 65.4% in the model using the 500m-by-500 m cell grid cells, i.e. from −0.026 (95% CI: −0.029, −0.022) to β = −0.009 (95% CI: −0.013, −0.005) (Z = −6.13). In contrast, the unemployment rate is the only indicator with an estimate that is stronger when we study the larger residential context (500m-by-500 m). With a standardized beta coefficient of 0.106 (95% CI: −0.108, −0.104), its association becomes even stronger than for the percentage of social benefit recipients (Z = 15.535). In this specific setting (using the 500m-by-500 m grid cells and focusing on the four urban regions), its beta coefficient is even lower than the one corresponding to the share of social benefit recipients (Z = −8.013). Apart from those cases, the other variables remain constant or only slightly vary when we modify the size of the residential context (when we focus on the four urban regions). For example, the variation in β for the Watts index using 40% national income is insignificant, as indicated by a Z-value of −1.135.

5.3. Variations across geographical scopes

There are variations in the contextual effect estimation across the three areas studied (i.e. the four urban regions, the Randstad conurbation, and the Netherlands) when analysing the regression results for the smaller residential contexts (100 m by 100 m). However, in comparison to the differences explored so far, these are less prominent than those generated by selecting a contextual poverty indicator or the scale of the residential context. For instance, the negative association between the percentage of social benefit recipients and individual income strengthens varies slightly as the geographic scope of the study expands. The beta coefficient is −0.120 (95% CI: −0.124, −0.115) in regression models focusing on the four largest urban regions, and slightly improves to −0.113 (95% CI: −0.117, −0.109) when considering the entire Randstad conurbation. However, the estimate drops again to −0.124 (95% CI: −0.126, −0.122) at the national level, which is statistically indistinguishable from the result for the four largest cities (Z = 1.593).

In addition, the beta coefficients for the percentage of unemployed individuals and most income-based indicators vary by around the same degree (by around 5% to 20%) when we expand the geographical scope of the analysis from the four urban regions to the whole Dutch territory. For example, the negative relationship between the unemployment rate and income strengthens by 4.9%; its β decreases from −0.082 to −0.086, with a Z-value of 3.521. Similarly, all Watts indexes – except the one using 60% of the national median income as the poverty line – as well as the poverty rate based on the municipal median income, also show a minor but statistically significant strengthening of their negative association with income. For instance, the beta coefficient for the Watts index with a poverty line set at 60% of the municipal median income declines from −0.014 to −0.019 (Z = 2.828). The estimate for the poverty rate calculated using municipal poverty lines (set at 60% of the median income) shows the strongest variation. We observe a 26.9% increase in its β, from −0.026 (95% CI: −0.029, −0.022) to −0.033 (95% CI: −0.035, −0.031) (Z = 3.428). The other income-based poverty indicators remain unchanged when varying the geographical scope of the analysis ($-1.96 < Z < 1.96$).

The regression outcomes also reveal that the negative association between the individual's income and the poverty indicators based on education gradually weakens as the geographical scope of the analysis expands. To illustrate this, the beta coefficient measuring the relationship between the percentage of individuals without an upper secondary education degree and income increase from −0.027 (95% CI: −0.032, −0.023) to −0.011 (95% CI: 0.013, −0.009), which represents a 59% drop of this negative relationship (Z = −6.704). Likewise, the negative association between income and the proportion of individuals without a tertiary education level decreases from a β of −0.012 (95% CI: −0.016, −0.015) to a β of −0.006 (95% CI: −0.008, −0.004), with a Z-value of −2.237. The strength of the negative association is reduced by 50%.

5.4. Variations across geographical scopes and spatial contexts

When we compare the results of the regression models analysing the 500m-by-500 m contexts to those utilizing the smallest neighbourhoods (100m-by-100 m), we notice slight differences in the variations of the estimated contextual effect across the different study areas. The most notable change concerns the unemployment rate. The modelling outcomes concerning the smaller residential context show that the negative relationship between the unemployment rate and individual income becomes 4.9% stronger when we study the Dutch territory instead of the four urban regions (its β drops from −0.082 to −0.086, with a Z-value of 3.521). However, in models investigating the larger residential context, we observe a gradual weakening of this relationship as the study's geographical scope increases. While it has a beta coefficient of −0.106 (95% CI: −0.108, −0.104) in the models focusing on the four largest cities, it increases progressively to −0.081 (95% CI: −0.082, −0.080) at the national level. This equates to a 23.6% weakening of this negative association (Z = −18.609). We notice a similar but fainter trend for the Watts indexes. Among them, the index using 60% of the national median income as the poverty line shows the most substantial variation in its association with income. When considering the 100m-by-100 m residential context, the negative relationship between this indicator and income remains stable across models using different geographical scope (Z = −0.129). In the models analysing the larger contexts, it weakens by 20% when we broaden the area studied to the entire Dutch territory (from $\beta = -0.015$ to $\beta = -0.012$, with a Z = −2.733). In comparison, the regression results regarding the poverty rates do not show any clear divergence in outcomes when comparing models using different residential contexts.

Regarding the share of individuals without an upper secondary education degree, the models analysing the 500-m-by-500-m residential contexts show similar results to those considering the 100m-by-100-m grid cells. Its negative relationship with income weakens when the study's geographical scope increases from the four largest cities to the whole country. The beta coefficient rises from −0.046 (95% CI: −0.04, −0.029) to −0.020 (95% CI: −0.024, −0.019), which represents a decrease in the relationship's strength of 56% (Z = −9.101). We also observe slight variation for the share of individuals without a tertiary education degree. In models analysing the smaller residential contexts, the negative relationship weakens by 56.5%, from a beta coefficient of −0.012 (95% CI: −0.016, −0.007) to −0.006 (95% CI: −0.008, −0.004) (Z = −2.237). This difference drops to −32%, from −0.025 (95% CI: −0.030, −0.0) to −0.017 (95% CI: −0.019, −0.015), in the models considering the 500m-by-500 m residential contexts (Z = −2.605).

6. Discussion

6.1. Effect of selecting a contextual poverty measure

The results of this study highlight the influence of researcher choices in terms of the wide range of potential contextual poverty measurements on the neighbourhood effect estimation – an issue we have termed the Problem of Uncertain Contextual Characteristic (PUCC). PUCC captures a conceptual and methodological challenge that goes beyond mere technical variation: it reflects fundamental uncertainty about which mechanisms are driving neighbourhood effects and how to measure their aggregated effect. To investigate this problem, we focused on the effect of neighbourhood poverty on individuals' income. Put simply, the choice around the operationalization of contextual poverty introduces more variation in the results than either the spatial scale or geographical extent of the analysis – often the focus of research in this field. Of the measures used, the percentage of social benefit recipients and the unemployment rate in the residential context have the strongest negative relationship with the individual's income from work, with associations up to 20 times stronger than for other measures. This underlines the importance of the operationalization of poverty in estimating neighbourhood effect.

The fact that the strongest estimates are found for the share of social benefit recipients and the unemployed could be attributed to the stronger conceptual alignment between these indicators and labour market outcomes. These indicators reflect more sustained forms of labour market exclusion,[9] making them more directly relevant for grasping cumulative disadvantages in neighbourhoods. This is especially true for the percentage of social benefit recipients, which includes not only the unemployed but also individuals on long-term assistance. This group typically experiences multiple and compounding deprivations, such as long-term unemployment, minimal income from social assistance, and the absence of savings or assets. Social assistance benefits are available only to those without sufficient income, capital or other provisions (e.g. unemployment benefit), making them a clear marker of severe economic hardship. These deeper forms of deprivation likely amplify neighbourhood effect mechanisms, such as socialization, network effect, and stigmatization.

We found similar patterns for the education-based contextual variables. The share of neighbours without an upper secondary education degree has a stronger negative association with individual income than the share of those without a tertiary education degree. This reflects the more important role of secondary education for accessing the labour market. In the Netherlands, like in most Western economies, completing upper secondary education has become a minimum threshold for stable employment. A high concentration of individuals lacking this qualification may therefore signal deeply rooted disadvantages. By contrast, the absence of tertiary education may be less consequential for the individuals. That said, the analysis of education-based indicators has limitations, which reduce their comparability to other context indicators. The higher share of missing values, especially for older individuals or first-generation migrants, may lead to an underestimation of the proportion of lower-educated residents in disadvantaged areas. As a result, it could reduce the strength of the relationship between education-based contextual variables and individual income.

The weaker relationships observed for income-based contextual poverty indicators may reflect the limitations of relying solely on income to define poverty. While such measures are commonly used in neighbourhood effects research, this approach might oversimplify poverty as a one-dimensional phenomenon. Focusing exclusively on income can misclassify households that have low income but also other resources shielding them from poverty (e.g. homeownership). It could also be that neighbourhood effect is more a matter of labour market integration than purchasing power. If so, calculating poverty based on the household unit and correcting for the household composition might introduce confounding factors in the modelling of neighbourhood effect. This could potentially reduce the sensitivity of income-based indicators to the broader mechanisms driving neighbourhood effects.

Yet we must be cautious in interpreting these results as evidence that neighbourhood effects are exclusively linked to severe material deprivation. One might expect a lower poverty threshold to yield stronger effects, given it reveals more extreme poverty. However, our results show the opposite. The standardized regression outcomes indicate that poverty indicators using the higher threshold (60% of median income) have stronger association with individual income than those using the lower threshold (40%). We conclude, therefore, that using a 40% threshold is too restrictive, excluding individuals who experience relative disadvantage, even if not in extreme poverty. By definition this intermediary group represents a significantly larger share of the Dutch population than those falling below the 40% line, as evidenced by the average poverty rates (see Annex 3, available in the Online Supplementary Material).

The results also support the assumption that the summary statistics selection can influence the neighbourhood effect estimation in terms of intensity, as the Watts indexes show lower coefficients than those using poverty rates. However, this does not mean that the poverty rate has a higher impact on income than the Watts index, as these results could be due to the binary nature of the poverty rate leading to an overestimation of the neighbourhood effect. Contrary to the poverty rate, the Watts index provides a clearer picture of contextual poverty by considering the severity of deprivation and inequality within the low-income group (Ravallion, 2015; Zheng 1993, 2002).

This analysis also explored whether setting the poverty line at national or municipal levels influences the estimation of the contextual poverty effect. Here, we found minor variations in beta coefficients. For instance, the standardized beta coefficients for the poverty rates and the Watts indexes using 60% of the median income slightly decrease when we set the poverty line at the municipal level. This could be because part of the estimated neighbourhood effect is attributed to regional economic disparities instead of the local inequality patterns, as argued by Jesuit, Rainwater, and Smeeding (2002) and Ayala, Jurado, and Pérez-Mayo (2014). However, we do not find any statistically significant variations in the beta coefficient of the Watts indexes using 40% of the median income when varying the reference group.

6.2. Relationships between contextual measures, neighbourhood scale, and studied geographies

These findings also indicate that the strength of each relationship between the contextual poverty indicators and individual income is sensitive to study design. This is a central insight captured by the PUCC: the significance of each mechanism behind the aggregated neighbourhood effect may vary depending on the study design and the relevance of the contextual measures to measure this aggregated effect will vary with them. As a result, comparing findings across studies (each using different indicators, spatial scales, or populations) is far from straightforward and can lead to diverging conclusions regarding the size of neighbourhood effects.

This study shows that the estimation of the contextual poverty effect highly varies with the scale of the residential context for some poverty indicators. These results confirm the importance of the scale effect from the MAUP. More importantly, they also reveal an interaction between the MAUP and PUCC, suggesting that spatial scale and the operationalization of the poverty measure jointly influence the estimation of neighbourhood effects. For instance, the regression models using the 100m-by-100 m grid cells as the residential context show a stronger relationship between the percentage of social benefit recipients and the individual's income from work than those using the 500m-by-500 m grid cells. We see the opposite pattern for the unemployment rate when we focus on the largest urban regions and the Randstad conurbation. As a result, we observe that the share of social benefit always has a stronger relationship with income than the unemployment rate at the smallest residential context scale (100m-by-100 m), and this is reversed at the largest scale (500m-by-500 m) for the largest urban regions and the Randstad conurbation. These outcomes support our prior assumption that each poverty indicator can correspond to different causal mechanisms, which can occur at different contextual scales as highlighted by Petrović, Manley, and Van Ham (2020). Our results also draw attention to the need to further investigate the nature of the relationship between different forms of poverty and neighbourhood effect mechanisms. Even though we observe evidence of this heterogeneity, our understanding of the specific causal pathways and their interplay with different forms of poverty and contextual scales remains limited.

Surprisingly, we observed little variation across models using different geographical extents. This indicates that the neighbourhood effect exists beyond the urban reaches and is also relevant for less densely inhabited areas in the Netherlands. We also initially expected that the neighbourhood effect would operate at a larger spatial scale and manifest differently in rural areas compared to urban ones, due to factors such as lower population density and differences in transportation means. Additionally, we considered that the sample with narrower geographic extent, focusing on highly urbanized areas, may overestimate the neighbourhood effect by over-representing poorer, less mobile individuals. However, the results across the contextual variables and residential scales do not support these hypotheses. One notable exception is the unemployment rate at the 100m-by-100 m scale, where the strength of the association decreases when the analysis is expanded to the national level. This may point to regional differences in how unemployment is spatially concentrated, though this pattern was not observed across other indicators. While these assumptions cannot be definitively dismissed, further research is necessary to explore them in greater detail. Moreover, caution is warranted when generalizing these findings, as countries with lower population densities may experience a more pronounced rural-urban divide than the Netherlands.

7. Conclusion

This study examined the sensitivity of neighbourhood effect estimates to the operationalization of contextual poverty, an overlooked yet potentially critical source of estimation variation. We argue that greater attention to this issue is essential for comparing empirical findings between studies, understanding what is truly being measured by neighbourhood effects estimates, and improving methodological approaches in the field. Our systematic analysis showed that the choice of contextual poverty measure significantly influences estimated effects, with the operationalization of poverty introducing more variation than either the scale of the residential context or the geographical extent of the study.

To conceptualize this challenge, we introduced the PUCC, which highlights the uncertainty that surrounds what is really being measured behind the neighbourhood effect and how it is represented through the construction of contextual variables such as poverty. This issue, inherent to all neighbourhood effects studies, can lead researchers to make arbitrary decisions when operationalizing contextual poverty, which in turn can result in estimation biases.

This uncertainty originates from the fact that different poverty measures can correspond to distinct social and economic realities and, since neighbourhood effects are produced by multiple underlying mechanisms, each may be more accurately represented by a different indicator. The challenge lies in identifying which indicators best reflect specific mechanisms (or a meaningful combination of them) and in determining how significant those mechanisms are within the aggregated neighbourhood effect. Adding to the complexity, the significance of each causal pathway will vary depending on the study design.

In this regard, our findings highlight that the PUCC does not operate in isolation. It interacts with other well-known sources of variation, such as the MAUP. Since the relevance of particular mechanisms can shift with the spatial scale considered, the effectiveness of any given poverty measure may also vary accordingly. In this way, the PUCC complements the MAUP in explaining how the measurement of contextual characteristics affects both estimation and interpretation of neighbourhood effects.

While this study brings attention to a largely underexplored dimension of neighbourhood effects research, it remains exploratory. Further work is needed to clarify the relationships between specific neighbourhood effect mechanisms and the various contextual indicators assessed in this analysis, and to examine how the PUCC interacts with other design choices, including population characteristics and the duration of exposure. Although we focused on income from work – one of the most commonly studied outcomes in the neighbourhood effects literature, and well suited to illustrating the implications of the PUCC – the uncertainty surrounding the operationalization of context is likely to extend to other outcomes, such as educational attainment or health. While patterns of variation may differ, the underlying issue is essentially the same: the PUCC. Future research should continue to illuminate this problem in order to better understand how this source of estimation variation influences the modelling of different individual outcomes.

Notes

1. The Watts index is computed as follows: $W = \sum \frac{ln(z) - ln(Yi)}{N} = \frac{1}{N} \sum \ln\left(\frac{z}{Yi}\right), \forall Yi < z$. Here, z is the poverty line, Yi is the individual's disposable household income, and N is the number of individuals in the area.
2. It includes the wage of an employee, the salary of a civil servant, the salary and other for directors, the wage in kind, the income from other works, reimbursed Zorgverzekeringswet Premium, and the income from an individual's own business.
3. To keep these null income values when applying the log transformation, we converse them into 1.
4. According to CBS, an urban region (in Dutch, 'Stadsgewest') consists of a metropolitan agglomeration and its surrounding area with smaller centres (towns, villages, hamlets). Two criteria have been used to assign the surrounding areas to each agglomeration, namely: the commuter flows and moves within and between municipalities. We deliberately included city peripheries because they are part of the city's sphere of influence and to maintain consistency with approaches commonly found in the literature.

5. We conducted robustness checks concerning the inclusion of education level as a control variable, as potential selection bias may arise from the high proportion of missing values among migrant populations and older cohorts in the CBS database. We also assess the influence of outliers in sparsely populated areas: 3.4% of observations reside in 100m-by-100m grid cells with only five individuals in the Netherlands dataset, a share that drops to 0.4% when using 500m-by-500m grid cells. Re-estimating the models after excluding observations in grid cells with five or fewer residents yields very similar results. In both cases, estimated coefficients remain stable, suggesting a negligible impact on neighbourhood effect estimates.
6. We present an example of the full regression model in Annex 4 (available in the Online Supplementary Material). This model focuses on the four largest urban regions in The Netherlands and using the income-based poverty rate (with a poverty line set at 60% of national median) in a 100m-by-100m grid cell as contextual poverty indicator. The other full regression models are available upon request.
7. Income change (in %): $100 \times (exp(\beta) - 1)$.
8. Z-test: $Z = \frac{\beta 1 - \beta 2}{\sqrt{SE\beta 1^2 + SE\beta 2^2}}$. Where SEβ is the standard error of β. The null hypothesis is rejected if the Z-value is greater than 1.96 or less than −1.96.
9. The unemployment rate and the share social benefit recipients are conservative indicators, as they are based on individuals' primary income source over the entire year rather than a single month.

Acknowledgements

Due to disclosure restrictions set by Centraal Bureau voor de Statistiek (CBS) to protect privacy, the data set is not available. This article is part of a larger project that has received approval from the Human Research Ethics Committee of the main author's institution.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Jérôme Francisco Conceicao http://orcid.org/0000-0002-7023-5636
Ana Petrović http://orcid.org/0000-0002-6337-964X
Maarten Van Ham http://orcid.org/0000-0002-2106-0702
David Manley http://orcid.org/0000-0003-1864-5141

References

Andersson, E. K., and B. Malmberg. 2015. "Contextual Effects on Educational Attainment in Individualised, Scalable Neighbourhoods: Differences across Gender and Social Class." *Urban Studies* 52 (12): 2117–2133. https://doi.org/10.1177/0042098014542487

Andersson, R., and S. Musterd. 2010. "What Scale Matters? Exploring the Relationships between Individuals' Social Position, Neighbourhood Context and the Scale of Neighbourhood." *Geografiska Annaler: Series B, Human Geography* 92 (1): 23–43. https://doi.org/10.1111/j.1468-0467.2010.00331.x

Andersson, R., S. Musterd, G. Galster, and T. M. Kauppinen. 2007. "What Mix Matters? Exploring the Relationships between Individuals' Incomes and Different Measures of Their Neighbourhood Context." *Housing Studies* 22 (5): 637–660. https://doi.org/10.1080/02673030701474602.

Andrews, D., C. Green, and J. Mangan. 2004. "Spatial Inequality in the Australian Youth Labour Market: The Role of Neighbourhood Composition." *Regional Studies* 38 (1): 15–25. https://doi.org/10.1080/0034340031000163228O.

Atkinson, T., B. Cantillon, E. Marlier, and B. Nolan. 2002. *Social Indicators: The EU and Social Inclusion*. Oxford: Oxford University Press. doi:10.1093/0199253498.001.0001.

Ayala, L., A. Jurado, and J. Pérez-Mayo. 2014. "Drawing the Poverty Line: Do Regional Thresholds and Prices Make a Difference?" *Applied Economic Perspectives and Policy* 36 (2): 309–332. https://doi.org/10.1093/aepp/ppt053.

Bakker, B. F., J. Van Rooijen, and L. Van Toor. 2014. "The System of Social Statistical Datasets of Statistics Netherlands: An Integral Approach to the Production of Register-Based Social Statistics." *Statistical Journal of the IAOS* 30 (4): 411–424. https://doi.org/10.3233/SJI-140803.

Bolster, A., S. Burgess, R. Johnston, K. Jones, C. Propper, and R. Sarker. 2007. "Neighbourhoods, Households and Income Dynamics: A Semi-parametric Investigation of Neighbourhood Effects." *Journal of Economic Geography* 7 (1): 1–38. https://doi.org/10.1093/jeg/lbl013.

Burt, R. S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110 (2): 349–399. https://doi.org/10.1086/421787

Cantillon, B. 2011. "The Paradox of the Social Investment State: Growth, Employment and Poverty in the Lisbon Era." *Journal of European Social Policy* 21 (5): 432–449. https://doi.org/10.1177/0958928711418856.

Chetty, R., and N. Hendren. 2018. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *The Quarterly Journal of Economics* 133 (3): 1107–1162. https://doi.org/10.1093/qje/qjy007.

Chetty, R., N. Hendren, and L. F. Katz. 2016. "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment." *American Economic Review* 106 (4): 855–902. https://doi.org/10.1257/aer.20150572.

Clogg, C. C., E. Petkova, and A. Haritou. 1995. "Statistical Methods for Comparing Regression Coefficients between Models." *American Journal of Sociology* 100 (5): 1261–1293. https://doi.org/10.1086/230638.

Custers, G., and J. J. Willems. 2024. "Rotterdam in the 21st Century: From 'Sick Man' to 'Capital of Cool'." *Cities* 150:105009. https://doi.org/10.1016/j.cities.2024.105009

Decancq, K., T. Goedemé, K. Van den Bosch, and J. Vanhille. 2014. "The Evolution of Poverty in the European Union: Concepts, Measurement and Data." In *Reconciling Work and Poverty Reduction in Europe: How Successful Are European Welfare States*, edited by Bea Cantillon and Frank Vandenbroucke, 60–93. Oxford: Oxford University Press.

Eurostat. 2026. *Gross Domestic Product (GDP) and Main Components Per Capita (nama_10_pc)*. https://ec.europa.eu/eurostat/databrowser/bookmark/16617f93-b35f-400d-9776-65b10fe6750b?lang=en.

Friedrichs, J. R., G. Galster, and S. Musterd. 2003. "Neighbourhood Effects on Social Opportunities: The European and American Research and Policy Context." *Housing Studies* 18 (6): 797–806. https://doi.org/10.1080/0267303032000156291

Galster, G., R. Andersson, and S. Musterd. 2016. "Neighborhood Social Mix and Adults' Income Trajectories: Longitudinal Evidence from Stockholm." *Geografiska Annaler: Series B, Human Geography* 98 (2): 145–170. https://doi.org/10.1111/geob.12096.

Galster, G., R. Andersson, S. Musterd, and T. M. Kauppinen. 2008. "Does Neighborhood Income mix Affect Earnings of Adults? New Evidence from Sweden." *Journal of Urban Economics* 63 (3): 858–870. https://doi.org/10.1016/j.jue.2007.07.002.

Galster, G. C. 2012. "The Mechanism (s) of Neighbourhood Effects: Theory, Evidence, and Policy Implications." In *Neighbourhood Effects Research: New Perspectives*, edited by Maarten van Ham, David Manley, Nick Bailey, Ludi Simpson, and Duncan Maclennan, 23–56. Dordrecht: Springer. https://doi.org/10.1007/978-94-007-2309-2_2.

Goedemé, T., and S. Rottiers. 2011. "Poverty in the Enlarged European Union. A Discussion about Definitions and Reference Groups." *Sociology Compass* 5 (1): 77–91. https://doi.org/10.1111/j.1751-9020.2010.00350.x.

Granovetter, M. S. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78 (6): 1360–1380. https://doi.org/10.1086/225469

Haandrikman, K., R. Costa, B. Malmberg, A. F. Rogne, and B. Sleutjes. 2023. "Socio-economic Segregation in European Cities. A Comparative Study of Brussels, Copenhagen, Amsterdam, Oslo and Stockholm." *Urban Geography* 44 (1): 1–36. https://doi.org/10.1080/02723638.2021.1959778

Hedman, L., D. Manley, M. Van Ham, and J. Östh. 2015. "Cumulative Exposure to Disadvantage and the Intergenerational Transmission of Neighbourhood Effects." *Journal of Economic Geography* 15 (1): 195–215. https://doi.org/10.1093/jeg/lbt042.

Jesuit, D., L. Rainwater, and T. Smeeding. 2002. *Regional Poverty Within the Rich Countries*.

Jivraj, S., E. T. Murray, P. Norman, and O. Nicholas. 2020. "The Impact of Life Course Exposures to Neighbourhood Deprivation on Health and Well-Being: A Review of the Long-Term Neighbourhood Effects Literature." *European Journal of Public Health* 30 (5): 922–928. https://doi.org/10.1093/eurpub/ckz153

Kakwani, N., and J. Silber. 2008. *Many Dimensions of Poverty*. London: Palgrave Macmillan. https://doi.org/10.1057/9780230592407.

Kwan, M. P. 2012. "The Uncertain Geographic Context Problem." *Annals of the Association of American Geographers* 102 (5): 958–968. https://doi.org/10.1080/00045608.2012.687349.

Lin, N. 2008. "A Network Theory of Social Capital." In *The Handbook of Social Capital*, edited by Dario Castiglione, Jan van Deth, and Guglielmo Wolleb, 50–69. New York: Oxford University Press.

Miltenburg, E. M., and T. W. Van Der Meer. 2018. "Lingering Neighbourhood Effects: A Framework to Account for Residential Histories and Temporal Dynamics." *Urban Studies* 55 (1): 151–174. https://doi.org/10.1177/0042098016639012.

Musterd, S., G. Galster, and R. Andersson. 2012. "Temporal Dimensions and Measurement of Neighbourhood Effects." *Environment and Planning A: Economy and Space* 44 (3): 605–627. https://doi.org/10.1068/a44298.

Musterd, S., S. Marcińczak, M. Van Ham, and T. Tammaru. 2017. "Socioeconomic Segregation in European Capital Cities. Increasing Separation between Poor and Rich." *Urban Geography* 38 (7): 1062–1083. https://doi.org/10.1080/02723638.2016.1228371

Musterd, S., and W. Ostendorf. 2013. *Urban Segregation and the Welfare State: Inequality and Exclusion in Western Cities*. London: Routledge.

Muštra, V., and B. Škrabić. 2014. "Regional Inequalities in the European Union and the Role of Institutions." *Review of Urban & Regional Development Studies* 26 (1): 20–39. https://doi.org/10.1111/rurd.12017.
Nieuwenhuis, J., T. Kleinepier, and M. van Ham. 2021. "The Role of Exposure to Neighborhood and School Poverty in Understanding Educational Attainment." *Journal of Youth and Adolescence* 50 (5): 872–892. https://doi.org/10.1007/s10964-021-01427-x
Openshaw, S. 1984. *The Modifiable Areal Unit Problem*. Norwich: GeoBooks.
Paternoster, R., R. Brame, P. Mazerolle, and A. Piquero. 1998. "Using the Correct Statistical Test for the Equality of Regression Coefficients." *Criminology* 36 (4): 859–866. https://doi.org/10.1111/j.1745-9125.1998.tb01268.x.
Petrović, A., D. Manley, and M. Van Ham. 2020. "Freedom from the Tyranny of Neighbourhood: Rethinking Sociospatial Context Effects." *Progress in Human Geography* 44 (6): 1103–1123. https://doi.org/10.1177/0309132519868767.
Petrović, A., D. Manley, and M. Van Ham. 2022. "Multiscale Contextual Poverty in the Netherlands: Within and Between-Municipality Inequality." *Applied Spatial Analysis and Policy* 15 (1): 95–116. https://doi.org/10.1007/s12061-021-09394-3
Petrović, A., M. van Ham, and D. Manley. 2022. "Where Do Neighborhood Effects End? Moving to Multiscale Spatial Contextual Effects." *Annals of the American Association of Geographers* 112 (2): 581–601. https://doi.org/10.1080/24694452.2021.1923455
Planbureau, C. 2024. *Inkomens en belastingen aan de top*. https://www.cpb.nl/sites/default/files/omnidownload/CPB-Publicatie-Inkomens-en-belastingen-aan-de-top.pdf.
Plum, A., and G. Knies. 2015. *Does Neighbourhood Unemployment Affect the Springboard Effect of Low Pay?*
Quillian, L. 2003. "How Long Are Exposures to Poor Neighborhoods? The Long-Term Dynamics of Entry and Exit from Poor Neighborhoods." *Population Research and Policy Review* 22 (3): 221–249. https://doi.org/10.1023/a:1026077008571.
Ravallion, M. 2015. *The Economics of Poverty: History, Measurement, and Policy*. Oxford: Oxford University Press.
San Millán, J., C. Cottineau-Mugadza, and M. van Ham. 2025. "The Economic Urban Divide: A Detailed Study of Income Inequality and Segregation in Dutch Urban Areas (2011–2022)." *Tijdschrift Voor Economische en Sociale Geografie* 116 (4): 508–530. https://doi.org/10.1111/tesg.70011.
Sen, A. 1976. "Poverty: An Ordinal Approach to Measurement." *Econometrica: Journal of the Econometric Society* 44 (2): 219–231. https://doi.org/10.2307/1912718.
Stein, R. E. 2014. "Neighborhood Scale and Collective Efficacy: Does Size Matter?" *Sociology Compass* 8 (2): 119–128. https://doi.org/10.1111/soc4.12127.
Townsend, P. 1979. *Poverty in the United Kingdom: A Survey of Household Resources and Standards of Living*. Berkeley: Univ of California Press.
Troost, A. A., H. J. Janssen, and M. V. Ham. 2023. "Neighbourhood Histories and Educational Attainment: The Role of Accumulation, Duration, Timing and Sequencing of Exposure to Poverty." *Urban Studies* 60 (4): 655–672. https://doi.org/10.1177/00420980221112351.
Urban, S. 2009. "Is the Neighbourhood Effect an Economic or an Immigrant Issue? A Study of the Importance of the Childhood Neighbourhood for Future Integration into the Labour Market." *Urban Studies* 46 (3): 583–603. https://doi.org/10.1177/0042098008100996.
Van Gent, W., and C. Hochstenbach. 2020. "The Neo-Liberal Politics and Socio-Spatial Implications of Dutch Post-crisis Social Housing Policies." *International Journal of Housing Policy* 20 (1): 156–172. https://doi.org/10.1080/19491247.2019.1682234
Van Ham, M., and D. Manley. 2010. "The Effect of Neighbourhood Housing Tenure mix on Labour Market Outcomes: A Longitudinal Investigation of Neighbourhood Effects." *Journal of Economic Geography* 10 (2): 257–282. https://doi.org/10.1093/jeg/lbp017.
Watts, H. W. 1969. An Economic Definition of Poverty.
Weinberg, Bruce A., Patricia B. Reagan, and Jeffrey J Yankow. 2004. "Do Neighborhoods Affect Hours Worked? Evidence from Longitudinal Data." *Journal of Labor Economics* 22 (4): 891–924. https://doi.org/10.1086/423158.
Wong, D. W. 2004. "The Modifiable Areal Unit Problem (MAUP)." In *WorldMinds: Geographical Perspectives on 100 Problems: Commemorating the 100th Anniversary of the Association of American Geographers 1904–2004*, edited by Donald G. Janelle, Barney Warf, and Kathy Hansen, 571–575. Dordrecht: Springer. https://doi.org/10.1007/978-1-4020-2352-1_93.
World Bank. 2026. *Gini index (SI.POV.GINI)*. (https://databank.worldbank.org/reports.aspx?source=2&series=SI.POV.GINI.
Zheng, B. 1993. "An Axiomatic Characterization of the Watts Poverty Index." *Economics Letters* 42 (1): 81–86. https://doi.org/10.1016/0165-1765(93)90177-E
Zheng, B. 2002. "Aggregate Poverty Measures." *Journal of Economic Surveys* 11 (2): 123–162. https://doi.org/10.1111/1467-6419.00028.