



GEAR

GRADUATE
PROGRAM
IN APPLIED
ECONOMIC
RESEARCH

working paper

Social Housing and Local Economic Segregation: Evidence from Local Tenure Change in the Netherlands

Diego Buitrago-Mora
Clémentine Cottineau-Mugadza

Working paper:
2026-01

This collection belongs to:

**Departament
d'Economia Aplicada**

UAB

Avinguda de l'Eix Central
Edifici B2
Campus de la UAB
08193 Bellaterra
(Cerdanyola del Vallès)
Barcelona - Spain
Tel. +34 93 581 16 80
Fax +34 93 581 22 92
d.econ.aplicada@uab.cat
www.uab.cat/departament/
economia-aplicada/

Coordinator: Rosella Nicolini (rosella.nicolini@uab.cat)

This collection is licensed under a Creative Commons Attribution-NonCommercial NoDerivatives 4.0 International License.



This collection includes a selection of research by students of the PhD Program in Applied Economics (UAB) and the Master of Applied Research in Economics and Business (MAREB) - specialization in Applied Economics. Research contributions can be published in English or Spanish.

Social Housing and Local Economic Segregation: Evidence from Local Tenure Change in the Netherlands

Diego Buitrago-Mora¹
TU Delft

Clémentine Cottineau-Mugadza²
TU Delft

Abstract

This paper examines how changes in local social-housing tenure composition affect neighbourhood income exposure in the Netherlands under the post-2010 income-targeted allocation regime. Using administrative panel data for 2009–2023 aggregated to 100 m grid cells, we instrument the local social-housing tenure share with baseline ownership interacted with leave-one-municipality-out national changes in owner-occupation. A 10-percentage-point increase in the social-housing share reduces exposure to households between the 5th and 10th percentiles of the national income distribution by approximately 0.87 percentage points, while the estimate for households in the bottom 5% is small and imprecise. A tenure decomposition shows that the P5–P10 effect operates primarily through social renters and, to a lesser extent, private renters, with little evidence of a corresponding effect among owner-occupiers. The central result is robust across alternative samples, exposure measures, trend specifications, inference procedures, and major-municipality exclusions. The findings show that changes in social-housing tenure composition can reshape local income exposure through broader rental-sector reallocation, with heterogeneous effects even within the bottom decile.

¹Department of Urbanism, TU Delft, Faculty of Architecture and the Built Environment, 2629, the Netherlands

²Department of Urbanism, TU Delft, Faculty of Architecture and the Built Environment, 2629, the Netherlands

1. Introduction

Social-housing allocation systems do more than determine access to subsidised dwellings. Through repeated allocations, they can reshape the spatial distribution of households and therefore the income composition of neighbourhoods. This raises a central question for urban and housing policy: when access to social housing becomes more strongly targeted by income, does local economic segregation increase, or do broader adjustments across housing tenures produce more differentiated spatial effects?

The Netherlands provides an informative setting to examine this question. The country has long been recognised as an international reference for the effectiveness of its social rental housing system. More than one-third of the housing stock is dedicated to social housing, and the Netherlands combines relatively low income segregation with high sociospatial mobility by European standards (Nieuwenhuis et al., 2020).

Housing associations (HAs) are entities tasked with providing affordable housing under public-service obligations and have played a central role in the expansion of the Dutch social-housing stock. However, structural reforms since the late 1980s have transformed their operational framework, shifting decision-making responsibilities from national authorities towards municipalities. At the same time, the mandate assigned to housing associations has moved from broad housing provision towards serving the most vulnerable segments of the population. The introduction of the so-called 90% rule in 2010 represents an important stage in this transition (Elsinga and Lind, 2013; Flynn and Montalbano, 2024; Nelson et al., 2024).

The 90% rule resulted from an agreement between the Dutch government and the European Commission requiring that at least 90% of newly available social-housing units be allocated to households below the *DAEB income limit*¹. The rationale was to reduce competition considered unfair between subsidised housing associations and regular landlords by restricting social-housing provision more closely to low-income renters (Gruis and Priemus, 2008; Elsinga and Lind, 2013).

The rule altered the institutional framework governing new social-housing allocations, and its effects accumulated over time in the local composition of housing tenure. We therefore examine whether variation in the local social-housing tenure share under the post-2010 regime affects neighbourhood income exposure.

These possible spatial consequences are important because the residualisation of social housing may reinforce social-renter stigmatisation and the concentration of poverty (Van Gent and Hochstenbach, 2020). Recent evidence shows that rent control combined with centralised waiting lists can generate misallocation among Dutch public-housing tenants (Buurma-Olsen et al., 2025). More generally, neighbourhood context can shape employment opportunities, social networks, health, and other socioeconomic outcomes

¹The DAEB income limit refers to the annual household-income ceiling used to determine whether a household belongs to the target group for social-housing allocation. Under the 90% rule, at least 90% of newly available social-housing units must be allocated to households below this income limit. DAEB stands for *Diensten van Algemeen Economisch Belang*, or Services of General Economic Interest (SGEI). In the Dutch housing context, DAEB housing refers to the regulated social housing segment provided by housing associations as part of their public-service task. This classification is rooted in EU state-aid rules, which allow public support for activities defined as services of general economic interest.

through a range of social, institutional, and spatial mechanisms (Pinkster, 2014; Galster, 2012; van Ham et al., 2012; Sampson, 2012). Exposure to place during childhood also has lasting effects on intergenerational mobility, with concentrated poverty and segregation associated with poorer outcomes for children from low-income families (Chetty and Hendren, 2018).

Against this background, we examine whether changes in local social-housing tenure composition reshape neighbourhood income exposure, with particular attention to households at the bottom of the income distribution. Causal evidence on this relationship remains limited. Existing research on housing policy and segregation in the Netherlands has primarily documented persistent sociospatial inequalities and the increasing differentiation of neighbourhoods across income groups (Uitermark et al., 2023; van Duijne and Ronald, 2018).

We use the evolution of the local social-housing tenure share as the observable empirical margin through which repeated allocations and broader tenure adjustments become visible. To address the endogeneity of local tenure composition, we instrument this share using baseline ownership interacted with leave-one-municipality-out national changes in owner-occupation. The resulting estimates capture how variation in local social-housing tenure composition within the broader post-2010 allocation regime affects neighbourhood income exposure.

The analysis uses administrative panel data for 2009–2023 aggregated to 100 m grid cells. We measure neighbourhood income exposure as the average income-group share across neighbouring grid cells within 400 metres, excluding the grid itself. This fine spatial scale allows us to examine changes in the immediate residential environment without relying on administrative neighbourhood boundaries.

The paper makes three contributions. First, it provides evidence on the relationship between local social-housing tenure composition and economic segregation using a national, highly granular spatial panel and an instrumental-variable strategy to address the endogeneity of local tenure composition. Second, it shows that broad definitions of the low-income population can conceal substantial heterogeneity. We distinguish households in P0–P5, P5–P10, and the cumulative bottom decile P0–P10, rather than treating all households below broad DAEB or *Passend Toewijzen* eligibility thresholds as a single group. Third, we connect the estimated exposure effects to broader housing-tenure reallocation by examining whether they operate through social renting, private renting, or owner-occupation. In doing so, the paper links the literature on social-housing allocation to research on neighbourhood exposure and local economic segregation.

The results reveal substantial heterogeneity within the bottom decile. A 10-percentage-point increase in the local social-housing tenure share reduces exposure to households between the 5th and 10th percentiles by approximately 0.87 percentage points, whereas the corresponding estimate for households in the bottom 5% is small and imprecise. A tenure decomposition shows that the P5–P10 effect operates primarily through social renters and, to a lesser extent, private renters, with little evidence of a corresponding effect among owner-occupiers. The central result is robust across alternative samples, exposure measures, trend specifications, inference procedures, and major-municipality exclusions.

The remainder of the paper is structured as follows. Section 2 describes the institu-

tional evolution of social-housing allocation in the Netherlands. Section 3 presents the data and variable construction. Section 4 documents changes in housing tenure and the residential contexts of different income groups. Section 5 introduces the local-exposure measure and presents descriptive spatial patterns. Section 6 outlines the empirical strategy and identification approach. Section 7 presents the main estimates, distributional and tenure mechanisms, heterogeneity across urban contexts, and robustness analyses. Section 8 concludes.

2. Social housing allocation in the Netherlands

This section outlines the institutional evolution of social housing allocation in the Netherlands. We focus on the sequence of regulatory changes that gradually narrowed access to social housing, beginning with the 90% rule and later reinforced by subsequent reforms. This historical context is essential for understanding how allocation practices shifted from broad access toward income-based targeting, and from relatively independent allocation systems, closely tied to the geographical presence of housing associations, to a more centralized regime with municipal implementation. Importantly, we do not treat the 90% rule as a one-time treatment, but as the starting point of a broader allocation system whose effects accumulate through repeated housing allocations over time. These changes provide a relevant setting for studying local economic segregation.

2.1. Historical Context

The development of the Dutch social rental housing system is rooted in a long tradition of strong state involvement and regulation, particularly between the 1950s and the 1990s. During these decades, extensive public support enabled a significant expansion of social housing, with its share of the total housing stock reaching 41% by 1990. According to [van Deursen \(2023\)](#), Housing Associations (HAs)—then known as *woningcorporaties*—played a pivotal role in this growth, becoming the primary vehicle for delivering the national government’s large-scale postwar housing agenda. Their expansion and operational scale were largely supported by central government policies, which treated access to affordable housing for low-income groups as a basic social right. This institutional framework positioned the Netherlands among the European countries with the highest levels of social housing provision.

In the late 1980s and early 1990s, HAs and the housing market underwent a process of deregulation and increased integration with the European market. This process granted HAs greater administrative and financial independence from the central government following the Balancing Act of 1995². By 1998, housing associations had been fully severed from direct state financial support and emerged as independent social enterprises. They owned approximately 70% of the entire rental stock, amounting to around 2.4 million dwellings at the time ([van Deursen, 2023](#)). The liberalization process aimed to create a more competitive market, encourage private sector participation, promote

²This financial independence was achieved by offsetting the net present value of future subsidies against the governmental loans owed by housing associations (amounting to approximately €32 billion in today’s money), providing HAs with a clean financial slate.

homeownership, and attract international investment, while maintaining the coverage and quality of the Dutch public rental system. As independent entities, HAs were no longer subject to direct state regulation. However, a regulatory framework was established, including the Guarantee Fund for Social Housing (Waarborgfonds Sociale Woningbouw, WSW), which ensured continued government involvement. This framework produced a low-risk structure that allowed HAs to secure loans at discounted rates. Combined with high housing demand, HAs were placed in a strong financial position. In 2016, HA-owned housing stock was valued at €250 billion, with €87 billion (34%) financed through loans and €173 billion (65%) as equity.³ Furthermore, HAs often set rental prices below the rent control limit⁴ (Priemus and Gruis, 2011), making some units eligible for support even if their features might have placed them above regulated thresholds. HAs could also purchase land from municipalities at below-market prices (Priemus and Gruis, 2011; van Deursen, 2023). Between 2007 and 2009, the global financial crisis impacted the Dutch economy and disrupted the housing market. Residential property sales halved between 2006 and 2013, according to Statistics Netherlands (Statistics Netherlands (CBS), 2018). The average price of existing owner-occupied homes decreased by about 6.5% in 2009, with recovery to pre-crisis levels only occurring in 2017.

The Post-2010 Income-Targeted Allocation Regime

The financial crisis exposed vulnerabilities in both the housing market and the regulatory framework established after the liberalization. Both private developers and HAs faced major challenges. The dominant position of HAs, supported by state aid, was increasingly seen as unfair to private developers. According to van Deursen (2023), the scale and commercialization of HAs appeared incompatible with their social mission. Scandals involving financial mismanagement⁵ further weakened public trust. In response, the Dutch Association of Institutional Investors (IVBN) filed a complaint with the European Commission (EC), citing the distortion of competition due to state support (Priemus and Gruis, 2011). The EC raised concerns in 2009, identifying three key distortive channels:

- **State Guarantees:** Free public guarantees for HAs constituted indirect transfers.
- **Subsidies and Loans:** Favorable loans and grants gave HAs selective advantage.
- **Land Sales:** Land sold below market value was seen as state aid.

Allocation Changes and Income Limit Adjustments

Two fundamental changes reshaped the Dutch social housing allocation system after 2010. First, housing associations, which had previously enjoyed considerable autonomy

³For context, Chile's GDP in the same year was approximately €250 billion, while Finland's GDP was €238 billion.

⁴In the Netherlands, a points system defines thresholds based on a dwelling's characteristics and the owner's tax obligations. If a dwelling exceeds this threshold, it enters the "free market," where rents are not regulated. Below the threshold, rents are government-regulated and eligible for housing allowances.

⁵Notable examples include the director of Rochdale using a company Maserati, Woonbron's failed cruise ship investment, and Vestia's disastrous derivatives deals (van Deursen, 2023, p. 23).

in their allocation decisions, became subject to stricter legal requirements. While income verification had existed in practice for assessing rent affordability and eligibility for *huurtoeslag* (government rent subsidy), there had been no binding national income thresholds, particularly not for social-housing allocation. Many associations applied their own internal guidelines or followed local agreements that prioritized specific low-income or vulnerable groups. However, middle-income households were still eligible for social housing under this flexible system.

Second, although the 90% rule was conceived as a national policy, its implementation became effectively decentralized. Local municipalities gained a key role in ensuring compliance, while housing associations adapted their practices to meet the DAEB thresholds. Furthermore, the 2015 Housing Act deepened this shift by embedding the *Passend Toewijzen* mechanism, which linked rent ceilings more tightly to household income. Of the 90% of units required to meet DAEB eligibility, at least 95% had to follow this appropriate allocation principle.

It is important to note that eligibility rules do not mechanically determine who lives in social housing. Income thresholds immediately define the pool of eligible households, but final allocation also depends on local allocation systems, waiting lists, point accumulation, priority rules, applications, offers, and household choices. The 90% rule therefore changes the institutional constraint under which new allocations take place, rather than assigning dwellings directly to specific households. Observing the period 2009–2023 allows us to capture how repeated allocations under this new constraint gradually reshape the composition of social housing and its local neighbourhood effects.

In practice, these changes narrowed the eligible pool of households and formalized a system in which income-based allocation was centrally regulated but locally enforced. This solidified the re-regularization trajectory of the Dutch social housing sector, replacing decades of flexible and broad access with a tightly monitored, rule-based system focused on the most vulnerable.

In Figure 1, we present a schematic view of the allocation limits. The figure summarizes how access to social housing is determined by a combination of income thresholds (DAEB limit) and rent-based allocation rules under *Passend Toewijzen*. The dark blue area represents the group of households eligible for social housing, within which at least 90% of new allocations must occur.

From the 1990s onward, rhetoric shifted from “social housing for everyone who wanted it” to “housing for vulnerable groups” (van Deursen, 2023), implying prioritization of low-income families but without formal income restrictions. During this earlier period, allocation practices relied largely on local discretion, and although affordability checks existed, there were no binding national income ceilings restricting access to social housing. After the 90% rule, eligibility became subject to DAEB income limits, and was later further constrained through rent-based allocation rules under *Passend Toewijzen*. A concise comparison of the main allocation rules before and after the 90% rule is summarized in Table 1.

Compliance with the 90% rule has consistently met its yearly allocation targets since its establishment. According to the report from [Ministerie van Binnenlandse Zaken en Koninkrijksrelaties \(2019\)](#), housing corporations have regularly met or slightly exceeded the minimum allocation requirement. For instance, in 2018, the allocation rate reached




Figure 1: Social housing allocation thresholds in the Netherlands. The figure illustrates the institutional structure of income and rent thresholds governing access to social housing. The shaded (dark blue) area represents the segment of households eligible under the DAEB income limit, within which at least 90% of newly allocated social housing units must be assigned. Additional rent-based restrictions (*Passend Toewijzen*) further constrain allocation within this eligible group. The figure is illustrative and intended to clarify the allocation framework rather than to provide direct empirical evidence.

Table 1: Summary of Social Housing Allocation Rules Before and After the 90% Rule

Dimension	Before 2010 (Pre-90% Rule)	After 2010 (Post-90% Rule)
Income eligibility	No legal income ceiling; local discretion prioritized low-income or priority groups.	Mandatory income threshold (DAEB limit); 90% of units reserved for eligible households.
Allocation autonomy	High autonomy for HAs; variation across municipalities.	Centralized regulation and coordination with municipalities; income testing required.
Market access	Regulated market broadly accessible; middle-income households often included.	Access restricted to DAEB-eligible households under income and rent caps.
Oversight	Limited state oversight; sector largely self-regulated.	Stronger state control through Autoriteit Woningcorporaties (AW).
Rent-income link	Rent controlled, but not strictly linked to income eligibility.	Income-rent coordination enforced under the 90% rule and the <i>Passend Toewijzen</i> mechanism.

93%.⁶

Together, these reforms shifted housing allocation from broad access toward targeted eligibility and centralized compliance, effectively ending the period of liberalized, market-

⁶From 2024 onwards, the limit increases to 92.5%.

oriented housing allocation that had characterized the previous two decades. This marked the beginning of a new phase of re-regularization, in which the Dutch state reasserted stronger regulatory control over housing associations in both allocation and oversight.

3. Data and Variable Construction

Our analysis uses access-restricted, georeferenced administrative data from Statistics Netherlands (CBS) for the period 2009–2023. The System of Social Statistical Datasets (SSD) provides information on household characteristics, housing tenure, and income at a highly granular spatial scale.⁷ Although the post-2010 social-housing allocation regime became effective from 2011 onward, we begin the panel in 2009 (the earliest year in which all variables used in the analysis were available) to observe pre-implementation spatial patterns and construct baseline variables, including the initial ownership share used in the empirical strategy.

We aggregate household records into 100 m × 100 m grid cells. In accordance with CBS confidentiality requirements, we exclude grid-year observations containing fewer than 10 households. The resulting dataset contains approximately 210,000 populated grids in 2009 and 3.27 million grid-year observations over the 15-year period.

We then restrict the sample to grids with a positive social-housing share in 2009. This allows us to focus on changes in local tenure composition within areas where social housing was already present before the allocation reforms. The restriction reduces the possibility of entry into the sample through newly created social-housing locations, while retaining changes in the existing social-housing stock and in other tenure categories. This restriction reduces the sample to 1,674,393 grid-year observations.

Our outcome variables measure exposure to the income composition of neighbouring grid cells within a 400-metre radius. To use a common sample across income groups, we exclude observations for which any of the spatial-lag measures cannot be constructed. The preferred IV analytical sample contains 60,216 unique grid cells and 903,240 grid-year observations. Table 2 reports descriptive statistics for this sample in 2009 and 2023.

We also use household net income to classify households into national income-percentile groups, calculated separately for each year. P0–P5 denotes households in the bottom 5% of the national income distribution, while P5–P10 denotes those between the 5th and 10th percentiles; together, these groups form the bottom decile, P0–P10. We additionally distinguish P10–P30, P30–P50, P50–P70, P70–P90, P90–P95, and P95–P100 to examine changes across the full income distribution. Our main focus is on the lower tail, where the allocation reforms are most likely to matter. Defining the groups annually provides a consistent measure of households' relative position in the national income distribution and avoids comparisons driven by changes in nominal income thresholds over time.

⁷*Data access and source statement.* Results are based on calculations by TU Delft in project number 9428 (SEGUE; <https://erc-segue.nl/>) using non-public microdata from Statistics Netherlands. Under certain conditions, these microdata are accessible for statistical and scientific research. For further information, see *Microdata: Conducting your own research* | CBS.

Table 2: Summary statistics for the final sample, 2009 and 2023

Variable	2009		2023		Δ 23–09
	Mean	SD	Mean	SD	Mean
Population per grid	86.273	51.189	87.659	53.150	+1.386
Households per grid	42.302	29.129	44.391	31.726	+2.089
Population density	8627.264	5118.948	8765.923	5315.000	+138.659
Average household size	2.151	0.488	2.093	0.465	−0.058
Social-housing share	0.532	0.317	0.482	0.318	−0.050
Single-person household share	0.368	0.182	0.400	0.181	+0.032
Couples without children share	0.271	0.121	0.255	0.115	−0.016
Income share: P0–P5	0.046	0.055	0.041	0.060	−0.005
Income share: P5–P10	0.051	0.056	0.058	0.061	+0.007
Income share: P10–P30	0.234	0.126	0.268	0.155	+0.034
Income share: P30–P50	0.223	0.094	0.216	0.087	−0.007
Income share: P50–P70	0.202	0.088	0.186	0.095	−0.016
Income share: P70–P90	0.178	0.103	0.166	0.108	−0.012
Income share: P90–P95	0.037	0.045	0.036	0.045	−0.001
Income share: P95–P100	0.029	0.049	0.029	0.049	0.000

Notes: Statistics are calculated across 100 m × 100 m grid cells in the final sample. Population density is measured as inhabitants per square kilometre. Housing, household-composition, and income-group variables are expressed as shares between zero and one. Income percentile groups are defined annually using the national household-income distribution. The final column reports the difference between the 2023 and 2009 means.

4. Tenure Reallocation and Spatial Context

While this paper focuses on social housing as the tenure directly affected by allocation rules, changes in social housing cannot be understood in isolation from broader tenure dynamics. Allocation reforms operate within a housing system where social rent, private rent, and owner-occupation are closely interconnected through household mobility and tenure transitions. For this reason, we situate social housing within the evolution of the full tenure structure across income groups. Doing so helps clarify how changes in the local social-housing tenure share interact with parallel shifts in private renting and ownership, and why observed patterns of local income exposure may reflect broader tenure rebalancing rather than changes within social housing alone.

Table 3 presents changes in the total housing stock between 2009 and 2023, disaggregated by tenure and income group. Over this period, the total number of dwellings increased from 7,158,080 in 2009 to 8,173,870 in 2023. This net expansion was primarily driven by growth in ownership, with more than 530,000 additional units, and by a substantial increase in private rental housing, which expanded by approximately 410,000 units. By contrast, the social rental sector remained broadly stable in aggregate terms. The changes presented in Table 3 are consistent with broader accounts of Dutch housing

Table 3: Change in Housing Tenure by Income Group, 2009–2023

Income	09–23 Δ (Units)			09–23 Δ (%)		
	Social Rent	Private Rent	Own.	Social Rent	Private Rent	Own.
P0–P10	95,474	156,085	-145,894	32.8	107.2	-53.1
P10–P30	275,012	94,864	-165,644	34.3	56.5	-36.0
P30–P50	-49,971	122,385	129,684	-8.9	79.6	18.1
P50–P70	-132,617	49,053	285,115	-37.6	32.2	30.7
P70–P90	-102,471	3,848	299,904	-51.1	2.6	27.6
P90–P100	-19,536	-15,743	136,220	-54.1	-21.2	22.5
Total	65,891	410,492	539,385	2.9	48.8	13.2

Notes: Changes in household units and percentage changes (%) relative to 2009. Total housing stock.

policy that emphasise the long-term promotion of owner-occupation, the expansion of market-oriented rental provision, and the gradual narrowing of social housing toward lower-income households (Hochstenbach, 2022).

Focusing on social rental housing, the total stock increased by 2.9% over the fifteen-year period, corresponding to a net gain of around 65,900 dwellings once demolitions, replacements, and new developments are taken into account. Behind this aggregate stability, however, the distribution of social rental housing across income groups changed significantly. Households in the lower part of the income distribution (up to P30) experienced substantial increases in both absolute and relative terms, with social rental units rising by 32.8% for P0–P10 and by 34.3% for P10–P30. In contrast, middle- and higher-income groups saw pronounced declines in social rental housing, both in units and as a share of their respective housing stock. This redistribution is consistent with a process of increased income targeting within the social housing sector and aligns with the notion of residualization discussed in the literature, whereby access to social housing becomes increasingly concentrated among lower-income households (Van Gent and Hochstenbach, 2020; Priemus, 2002; Scanlon et al., 2015).

Private rental housing exhibits the strongest proportional expansion across all tenures. Between 2009 and 2023, the private rental stock grew by nearly 49% relative to its 2009 level. This expansion is observed across the income distribution but is particularly pronounced at the lower end: for households in P0–P10, the number of private rental units more than doubled (an increase of 107.2%), while substantial growth is also observed for P10–P30 and P30–P50. This pattern suggests that private renting increasingly accommodates households that previously would have relied on social rental housing or entry into owner-occupation.⁸

In absolute terms, the largest tenure expansion occurred in owner-occupation, with an increase of approximately 539,000 dwellings between 2009 and 2023. Owner-occupation

⁸Whether the strong expansion of private renting among lower-income households reflects affordability constraints, reduced access to social housing, or broader changes in housing market institutions is beyond the scope of this study and is left for future research.

remains the dominant tenure in the Netherlands and expanded for most income groups over the period. However, this growth is unevenly distributed: while ownership increased among middle- and higher-income households, it declined sharply for the lowest income groups. For P0–P10 and P10–P30, ownership fell by 53.1% and 36.0%, respectively.⁹

These aggregate tenure shifts also conceal important differences in the residential contexts of households within the bottom decile. In 2009, P5–P10 households were located in grids with an average social-housing share of 0.634, compared with 0.537 for P0–P5 households. By 2023, the corresponding figures were 0.638 and 0.507. Conversely, the average private-rental share of the grids in which P0–P5 households were located increased from 0.153 to 0.228, compared with an increase from 0.116 to 0.130 for P5–P10 households. These figures suggest that P5–P10 households remain more closely associated with social-housing environments, whereas P0–P5 households are increasingly located in grids with lower social-housing shares and higher private-rental shares.

All in all, these patterns point to a housing market characterized by pronounced tenure reallocation across income groups. Despite aggregate stability in the size of the social rental sector, social housing becomes increasingly concentrated among lower-income households, while middle- and higher-income groups exit toward ownership or private renting. At the same time, the expansion of private rental housing and continued growth of owner-occupation coexist with a contraction of ownership among the lowest income groups, who increasingly concentrate in rental tenures.

Taken together, these patterns help place social housing within the broader context of tenure changes observed over the period. While the empirical analysis that follows focuses on local social-housing tenure-share variation, this descriptive evidence provides the background for interpreting changes in local income exposure and spatial segregation.

Furthermore, it is important to mention that social housing is not evenly distributed across urban space. Its location reflects historical construction patterns, municipal land policies, and the spatial development of housing associations. In Dutch municipalities, the spatial autocorrelation of social housing was on average positive in 2024 (Moran's $I=0.43$), but varied between 0.03 (no spatial clustering of social housing) in Vlieland to more than 0.6 (high spatial concentration) in Wassenaar [San Millán et al. \(2026\)](#). This spatial concentration matters for our empirical setting because allocation rules affect who can access social housing, but their neighbourhood effects also depend on where the existing stock is located.

5. Local Economic Segregation: Exposure Patterns in the Netherlands

Measuring local exposure

In this paper, we conceptualise local economic segregation as exposure to neighbours from specific income groups. We therefore focus on the exposure dimension of segregation, rather than broader dimensions such as evenness or centrality ([Massey and Denton, 1988](#); [Reardon and O'Sullivan, 2004](#)). To capture this local exposure, we measure the

⁹The expansion of owner-occupation during this period coincides with a broader policy and market environment that encouraged homeownership, including mortgage market developments and fiscal incentives (see, among others, [Elsinga et al., 2014](#); [Boelhouwer, 2018](#)).

income composition of the neighbourhood surrounding each 100m grid cell using a spatial-lag approach.

For each grid i and year t , we measure local exposure to income group k as:

$$\text{spatial-lag}_{it}^k = \frac{\sum_j w_{ij} s_{jt}^k}{\sum_j w_{ij}}, \quad (1)$$

where s_{jt}^k is the share of households in income group k in grid j , and w_{ij} is a binary spatial weight equal to 1 if grid j lies within 400 metres of grid i , and 0 otherwise. We exclude own-cell information by setting $w_{ii} = 0$ for all i . In other words, the index j refers to neighbouring grids in the spatial weighting matrix, while the index k denotes income-percentile groups (e.g. P0–P10, P90–P100).

The baseline spatial lag is constructed as a cell-weighted measure: for each grid cell, it averages the income-group shares of neighbouring 100m cells within the 400m buffer. This provides a place-based measure of the surrounding local environment and keeps the spatial unit of exposure consistent across the country.

Because neighbouring 100m grids can share part of the same 400m buffer, the resulting exposure values are spatially related. We therefore interpret the measure as a continuous local neighbourhood-exposure indicator, rather than as independent neighbourhood observations.

This measure captures the average share of each income group in the immediate neighbourhood surrounding each grid cell. Rather than relying on global segregation indexes or local autocorrelation statistics such as Moran’s I , the spatial-lag provides a direct and interpretable measure of local economic segregation. It is closely aligned with exposure-based approaches in the segregation literature, which focus on the composition of the surrounding population as the relevant dimension of neighbourhood context (Massey and Denton, 1988; Reardon and O’Sullivan, 2004; Bischoff and Reardon, 2012). In this sense, the spatial-lag provides a transparent representation of the local economic environment, which is the relevant outcome for our policy question.

The choice of the 400-metre radius is guided by the empirical literature on neighbourhood boundary definitions. Studies commonly use distances around 400 metres to proxy local walkable environments, corresponding to roughly a five-minute walking distance (e.g. McCormack and Shiell (2010)). In urban studies, this distance is often interpreted as the immediate social and physical surroundings residents engage with on a daily basis. Jenks and Dempsey (2007) emphasise that even when neighbourhoods do not have clear physical boundaries, a 400-metre buffer is useful to capture the area that residents implicitly consider part of their neighbourhood. Their review highlights that neither administrative units nor visually drawn boundaries reliably approximate residents’ own definitions, whereas a buffer approach offers a more consistent solution in dense urban settings.

Using this buffer-based spatial-lag has several advantages for our specific context. First, local economic segregation is inherently spatial: it describes the income context around each location, not only the characteristics of a single 100m cell. Relying only on the income composition of the grid itself would miss this relational dimension and produce a highly volatile measure due to the small population counts in many cells.

Second, the spatial-lag avoids the mechanical overlap that arises when both the treatment (share of social housing within the grid) and the outcome rely on the same cell value. By construction, the spatial-lag reflects the surrounding neighbourhood rather than the interior of the grid, creating a cleaner distinction between exposure and treatment. Third, it allows us to fully exploit the granularity of the Dutch micro-data: with 100m grids, the 400-metre buffer provides a comprehensive set of neighbours for the vast majority of grids, avoiding a substantial loss of observations.

Our approach also aligns with established practices in the neighbourhood-effects literature, particularly the use of egohoods to define local environments around individuals or locations (Hipp, 2007). Local socioeconomic conditions are commonly measured using spatially aggregated indicators within short buffers (Bayer et al., 2007; Ellen and O'Regan, 2011; Galster, 2012). Compared with alternative measures such as global Moran's I, Theil's H, D-indexes, or entropy-based segregation metrics, our approach is deliberately local and non-parametric. It does not impose assumptions about administrative boundaries or functional neighbourhoods, and it can be computed consistently for all grids in the Netherlands, from dense urban areas to less urbanised ones. In densely built environments, the 400-metre buffer closely matches a pedestrian-scale neighbourhood; in more sparsely populated areas, it still captures the closest available residential context (recall that we omit grids with fewer than 10 households).

Exposure patterns across urban contexts

Having defined the measure of local exposure, we first use it descriptively to examine how neighbourhood income exposure changed between 2009 and 2023 across urban contexts.

To examine differences in local exposure across income groups and urban contexts, we rely on CBS's 2009 degree-of-urbanisation thresholds: extremely urbanised ($\geq 2,500$), strongly urbanised (1,500–2,499), moderately urbanised (1,000–1,499), and rural/sparsely populated ($< 1,000$).¹⁰ The use of these categories allows us to describe how exposure patterns vary across different urban contexts.

The spatial lag ranges from 0 to 1 and measures local exposure: for each grid cell i (100 m $\times$ 100 m), it captures the average income-group share across neighbouring 100m grid cells within a 400 m buffer. A value of 0 indicates that the income group is absent from the neighbouring grids, while higher values indicate greater local exposure to that group. Table 4 reports the average change in the spatial-lag of income groups between 2009 and 2023 across degrees of urbanisation.

Table 4 provides a first descriptive summary of changes in local exposure. The clearest pattern is found at the bottom of the income distribution. Exposure to Po–P5 households falls across all degrees of urbanisation, suggesting a modest but systematic decline in local exposure to the very bottom of the income distribution. By contrast, exposure to P5–P10 households remains closer to zero, with small increases in strongly and moderately urbanised areas and almost no change in rural areas. This indicates that the decline in local exposure is clearer for Po–P5 than for the next-lowest group.

¹⁰These categories are based on address density, measured as the number of addresses within a 1 km radius.

Table 4: Average change in spatial-lag of income shares by percentile group and degree of urbanisation (2009–2023)

	Ext.	Str.	Mod.	Rur.
Lag-P0–P5	–0.006	–0.013	–0.011	–0.008
Lag-P5–P10	–0.001	+0.005	+0.002	–0.0004
Lag-P10–P30	+0.014	+0.012	+0.003	–0.003
Lag-P30–P50	+0.002	+0.006	+0.009	+0.013
Lag-P50–P70	–0.003	+0.005	+0.007	+0.002
Lag-P70–P90	–0.004	–0.004	–0.003	–0.007
Lag-P90–P95	–0.002	–0.003	–0.003	+0.001
Lag-P95–P100	–0.002	–0.004	–0.006	–0.002

Notes: Values report the average change in the spatial-lag (400 m buffer) of the share of households in each national income-percentile group between 2009 and 2023, averaged across all grids with more than 10 households. “Ext.”, “Str.”, “Mod.” and “Rur.” correspond to CBS degree-of-urbanisation categories.

For lower-middle income groups, exposure increases particularly for P10–P30 and P30–P50. These increases are most visible in extremely and strongly urbanised municipalities. Changes at the upper tail of the distribution remain small, suggesting that local exposure among the highest-income households changes less over the period.

However, the values in Table 4 are averages across many grids and should be interpreted as descriptive summaries rather than aggregate segregation indexes. They do not show whether the average changes are widespread across grids or driven by specific local patterns. For this reason, we complement the table with distributional evidence at the grid and municipality levels.

Figures 2 and 3 show the distribution of grid-level changes by degree of urbanisation. For P0–P5, the median change is negative across all categories, indicating that the decline in exposure to the bottom 5% is not only an average pattern but is also visible in the central tendency of the grid-level distribution. The distribution is more dispersed in extremely urbanised areas, where some grids experience increases, but the median remains below zero.

For P5–P10, the distribution is centred closer to zero. This is consistent with Table 4, where average changes for this group are smaller and less uniform. The figure also shows that changes in this group are more spatially heterogeneous: some grids experience increases in exposure, while others experience declines. Overall, the grid-level distributions suggest that changes in local exposure are clearer and more systematic for P0–P5, while P5–P10 shows a more mixed pattern.

Municipal and within-city heterogeneity

The previous figures show how changes vary across degrees of urbanisation. We now examine whether these patterns are also visible across municipalities. Figures 4 and 5 report the distribution of grid-level changes for municipalities with at least 50,000 inhabitants. These figures are not intended to estimate municipality-specific effects, but to illustrate how local exposure changes are distributed within and across cities.

Figure 4 shows the distribution of grid-level changes in exposure to P0–P5 households




Figure 2: Distribution of changes in the spatial-lag of Po-P5 households by degree of urbanisation, 2009–2023




Figure 3: Distribution of changes in the spatial-lag of P5–P10 households by degree of urbanisation, 2009–2023

across municipalities with at least 50,000 inhabitants. The line inside each box represents the median grid-level change, while the dot represents the municipal mean. In most municipalities, both the median and the mean are below zero, which is consistent with the average decline reported in Table 4. However, the boxes also show substantial within-municipality dispersion, indicating that even municipalities with a negative central

tendency contain grids where exposure to Po–P5 increases. Delft and Leiden stand out as municipalities with positive average changes, while several large cities, including Amsterdam, The Hague, Utrecht, Tilburg, Haarlem, and Eindhoven, show negative average changes. This suggests that the decline in exposure to Po–P5 is not driven by a single city, but also that local variation within municipalities remains important.

Figure 5 presents the analogous distribution for P5–P10 households. The pattern is less uniform than for Po–P5. In several municipalities, the median remains close to zero, and the mean often differs from the median, indicating a more uneven distribution of local changes. The boxes frequently cross both sides of the zero line, suggesting that increases and decreases in exposure coexist within the same municipality. This confirms that changes in exposure to P5–P10 households are more heterogeneous and less clearly directional than changes for Po–P5. In this sense, the municipality-level figures help clarify why Table 4 should be interpreted as a descriptive summary rather than as a complete account of local exposure dynamics.

To illustrate the local heterogeneity captured by the spatial lag, we map changes in local exposure in the city of Amsterdam. Figure 6 shows the change in the spatial-lag for Po–P5 at the 100 m grid level. In several parts of the municipality, with a notable concentration in New West and other more peripheral areas, local exposure to Po–P5 households increases (darker green areas), while other areas, particularly around the centre, show declines. This pattern is consistent with broader accounts of gentrification and tenure restructuring in Amsterdam, which document growing pressure on lower-income households in central areas and changing access to the social rental stock (Hochstenbach, 2017). However, because the map describes changes in neighbourhood exposure rather than individual residential moves, it should not be interpreted as direct evidence of displacement. Instead, it shows that the average reduction in Po–P5 exposure coexists with localised increases in some more peripheral parts of the city.

Figure 7 presents the analogous map for the top group (P95–P100). Here, the pattern is different: around the central districts, local exposure to high-income households increases over the period, consistent with growing exposure to richer households in and around the core.

Beyond serving as an illustrative example, these maps help clarify the interpretation of our measure: each grid receives a neighbourhood-specific exposure value, and these local values reveal spatial variation that is not visible in aggregate averages. These patterns motivate the next step of the analysis, where we examine whether changes in the local share of social housing help explain changes in neighbourhood income exposure.

6. Empirical Strategy

Our empirical strategy estimates how changes in the local social-housing tenure share affect neighbourhood income exposure at the 100-metre grid level. The treatment is a relative measure of local tenure composition and should not be interpreted as the number of new social-housing allocations or as the absolute social-housing stock. We instrument the local social-housing share using baseline ownership structure interacted with subsequent national ownership trends. The analysis therefore identifies the effect of




Figure 4: Distribution of changes in the spatial-lag of Po–P5 households by municipality, 2009–2023

instrument-induced variation in the local social-housing tenure share within the broader post-2010 allocation regime.

Accordingly, the treatment variable should be interpreted as a relative measure of changes in local tenure composition within the broader post-2010 allocation regime, rather than as a direct indicator of a single institutional reform.

To test this idea, we estimate whether the evolution of the share of social housing




Figure 5: Distribution of changes in the spatial-lag of P5–P10 households by municipality, 2009–2023

in each grid shapes the level of local income exposure, measured as the spatial-lag of income within a 400-metre buffer. Because the treatment is a stock share, its evolution may reflect changes in occupied social housing, other tenure categories, and the total number of classified households. We therefore interpret the coefficient as the effect of net local social-housing tenure-share variation, rather than as the effect of pure allocation flows.

Amsterdam: Spatial Lag, Q05(smoothed)




Figure 6: Evolution of grid-level exposure to Po-P5 households

Amsterdam: Spatial Lag, Q95(smoothed)




Figure 7: Evolution of grid-level exposure to P95-P100 households

We exploit the panel structure of the data to identify how changes in the local share of social housing within a grid are associated with changes in the income composition of its surrounding area over time. The empirical strategy relies on within-grid variation, while controlling for common national trends and municipality-specific shocks. By combining fine spatial resolution with grid fixed effects, year fixed effects, and municipality-by-year fixed effects, the specification absorbs persistent grid characteristics, common national trends, and municipality-specific shocks that may jointly affect housing markets and local income composition.

Baseline model

The starting point is the following equation:

$$\text{spatial-lag}_{it}^k = \alpha + \beta \text{SocH}_{it} + \mathbf{X}_{it}'\gamma + \mu_i + \tau_t + \lambda_{m(i),t} + \varepsilon_{it}^k, \quad (2)$$

SocH_{it} denotes the share of dwellings that belong to the social housing sector in grid i at time t . $\text{spatial-lag}_{it}^k$ is the spatial-lag income measure for income group k , within 400 metres of grid i for a year t between 2009 and 2023. $\mathbf{X}_{it}'$ includes grid-level time variant controls that absorb structural differences between grids:

- Population density, capturing long-term urban form;
- Household size and composition ¹¹ which affect local income structures;

We also include a set of fixed effects that absorb the main time-invariant characteristics at the grid level:

- Year fixed effects (τ_t)
- Grid fixed effects (μ_i)
- Municipality-year fixed effects ($\lambda_{m(i),t}$), absorbing municipal policies and local market conditions.

This is important because municipalities and neighbourhoods differ in the historical concentration and clustering of social housing, which may be correlated with neighbourhood income composition [San Millán et al. \(2026\)](#). The inclusion of grid fixed effects means that identification does not rely on cross-sectional differences between places with high and low baseline social-housing concentration, but on changes within grids over time. Municipality-by-year fixed effects further absorb local shocks, policies, and housing-market conditions common to grids within the same municipality and year.

An OLS relationship between the local share of social housing and local income exposure may suffer from simultaneity and policy targeting. Grids experiencing rising vulnerability may experience an increase in the local social-housing tenure share, while

¹¹Here we include two variables: one captures single-person households and the other captures households with children. Although we have access to a broader set of household-composition variables, they are collinear. We therefore keep the variables that best capture long-term effects.

changes in that share may in turn reshape the surrounding income composition. This two-way relationship would bias the OLS estimate of β . To address this endogeneity concern, we adopt a 2SLS strategy using an instrument based on a Bartik (shift-share) structure.

6.1. Instrument: Expected Ownership

The instrument follows a shift-share logic: a fixed local tenure structure interacts with an external national tenure shift. The local exposure component is the initial share of owner-occupied dwellings in 2009, which captures how strongly each grid is exposed to subsequent national changes in owner-occupation. The instrument therefore exploits the fact that grids with different baseline tenure structures are differentially exposed to common national ownership dynamics, after excluding the municipality from the national shift. We interpret this source of variation as predicting changes in the local share of social housing within otherwise comparable grids in the same municipality and year.

The instrument takes the following form:

$$Z_{it} = \underbrace{\text{Own}_{i,2009}}_{\text{Local exposure}} \times \underbrace{\left(\bar{\text{Own}}_t^{-m(i)} - \bar{\text{Own}}_{2009}^{-m(i)} \right)}_{\text{National shift (leave-one-out)}}, \quad (3)$$

where:

- $\text{Own}_{i,2009}$ is the ownership share in grid i in 2009;
- $\bar{\text{Own}}_t^{-m(i)}$ is the national ownership average at time t , excluding municipality $m(i)$;
- the difference relative to 2009 captures the national trend.

The identifying assumption is that, conditional on grid fixed effects, year fixed effects, municipality-year fixed effects, and time-varying controls, baseline ownership interacted with national ownership shifts affects local income exposure only through changes in the local social-housing tenure share. The leave-one-out construction prevents the national shift from being mechanically driven by housing-market developments in municipality $m(i)$, while municipality-year fixed effects absorb local shocks such as price cycles, construction activity, and policy changes. The instrument therefore identifies variation in relative local tenure composition, rather than changes in the absolute number of social-housing households or in pure allocation flows.

6.2. First stage

The first stage predicts the social housing share using the Bartik instrument:

$$\text{SocH}_{it} = \pi_0 + \pi_1 Z_{it} + \mathbf{X}'_{it} \rho + \mu_i + \tau_t + \lambda_{m(i),t} + u_{it}. \quad (4)$$

The coefficient π_1 captures how much the instrument shifts the local social-housing tenure share, after conditioning on fixed effects and controls. The fitted values, $\widehat{\text{SocH}}_{it}$, isolate the exogenous component of social housing variation ¹².

¹²Table 6 reports the first-stage results.

With municipality–year fixed effects, identification comes from comparing grids within the same municipality and year that differ in baseline ownership exposure to the common national shift.

6.3. Second stage

In the second stage, we replace the endogenous regressor with its predicted value:

$$\text{spatial-lag}_{it}^k = \alpha + \beta \widehat{SocH}_{it} + \mathbf{X}'_{it}\gamma + \mu_i + \tau_t + \lambda_{m(i),t} + \varepsilon_{it}^k. \quad (5)$$

The coefficient β provides the effect of instrument-induced variation in the local social-housing tenure share on local income exposure within a 400-meter radius.

Our empirical strategy identifies the effect of the variation in the local social-housing tenure share on local income exposure. While the 90% rule motivates the analysis and plausibly contributed to shifts in allocation and sorting, our estimand is not a reduced-form “policy effect” of the reform itself, but rather the neighbourhood-composition impact of tenure-share variation within the broader post-2010 allocation regime.

7. Empirical Results

We present the empirical results in four steps. We first report the preferred estimates for the lower tail of the income distribution, together with the first-stage, reduced-form, and dynamic timing results. We then examine the broader distributional pattern and decompose the central P5–P10 result by housing tenure. Next, we assess heterogeneity across urbanisation categories. Finally, we evaluate the stability of the results across alternative samples, controls, exposure measures, trend specifications, inference procedures, and major-municipality exclusions.

7.1. Baseline Estimates

We begin by comparing pooled OLS and fixed-effects estimates following Equation 2, and instrumental-variable estimates following Equation 5, for the lower tail of the income distribution. Table 5 reports results for three outcomes: exposure to households in the bottom 5% of the income distribution (Po–P5), exposure to households between the 5th and 10th percentiles (P5–P10), and exposure to the cumulative bottom decile (Po–P10). This distinction is important because the bottom decile combines two groups that may respond differently to local social-housing tenure-share variation.

Column (5) reports the preferred IV estimate, using the ownership-based shift–share instrument. The IV estimate for P5–P10 is negative and statistically significant, indicating that instrument-induced increases in the local share of social housing reduce local exposure to households between the 5th and 10th percentiles. The difference between the fixed-effects and IV estimates is consistent with upward bias in the non-instrumented estimates, potentially reflecting endogenous targeting, residential sorting, or local housing-market dynamics that correlate changes in social housing with changes in low-income exposure.

The result for Po–P5 is different. The IV estimate is positive but statistically insignificant. As shown below, this coefficient is also less stable across sample restrictions and heterogeneous specifications. We therefore treat the Po–P5 estimate as imprecise rather

than as the central result. In contrast, the negative estimate for P5–P10 is more stable and provides the main evidence that local social-housing tenure-share variation reduces exposure for part of the lower-income distribution. The cumulative Po–P10 coefficient reflects these two offsetting patterns.

Table 5: OLS, Fixed-Effects, and IV Estimates for Lower-Income Exposure

	(1) OLS	(2) OLS	(3) FE	(4) FE	(5) IV
Panel A: Po–P5 exposure					
Share of social housing	0.0057 ^a (0.0005)	-0.0114 ^a (0.0005)	0.0012 ^a (0.0003)	0.0013 ^a (0.0003)	0.0071 (0.0061)
Panel B: P5–P10 exposure					
Share of social housing	0.0232 ^a (0.0004)	0.0141 ^a (0.0003)	0.0009 ^a (0.0003)	0.0002 (0.0003)	-0.0870 ^a (0.0077)
Panel C: Po–P10 exposure					
Share of social housing	0.0289 ^a (0.0007)	0.0027 ^a (0.0007)	0.0021 ^a (0.0004)	0.0021 ^a (0.0007)	-0.0798 ^a (0.0107)
Controls	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Grid FE	No	No	Yes	Yes	Yes
Municipality × Year FE	No	No	No	Yes	Yes
Ownership: IV	No	No	No	No	Yes
Kleibergen–Paap Wald F-stat	–	–	–	–	343.12
Observations	903,240	903,240	903,240	903,240	903,240
Groups (Grids)	60,224	60,224	60,224	60,224	60,216

Notes: Robust standard errors clustered at the grid level. Significance levels: ^a p<0.01, ^b p<0.05, ^c p<0.1. Controls include population density, household size, single-person households, and couples without children. Column (5) instruments the share of social housing using the ownership-based shift-share instrument.

The first stage follows Equation 4 and its results are reported in Table 6. The ownership-based shift-share instrument has the expected positive first-stage relationship with the local share of social housing, indicating that the instrument predicts variation in social-housing share in the intended direction. The corresponding reduced-form regressions estimate the effect of the instrument directly on local exposure to Po–P5 and P5–P10 households, using the same controls and fixed effects as in the main specification. The reduced-form estimates have the same signs as the IV estimates: positive for Po–P5 and negative for P5–P10.

Table 6: First-Stage and Reduced-Form Regressions

	First stage	Reduced form	Reduced form
Dependent variable	Social housing share	P0–P5 exposure	P5–P10 exposure
Instrument: expected ownership	1.508 ^a (0.0814)	0.0107 (0.0092)	-0.131 ^a (0.0092)
Population density	Yes	Yes	Yes
Household size	Yes	Yes	Yes
Household composition	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Grid FE	Yes	Yes	Yes
Municipality $\times$ Year FE	Yes	Yes	Yes
F-test excluded instrument	343.12	–	–
Kleibergen–Paap rk LM statistic	341.136	–	–
Kleibergen–Paap Wald F-stat	343.125	–	–
Observations	903,240	903,240	903,240
Groups (Grids)	60,216	60,216	60,216

Notes: Robust standard errors clustered at the grid level.

Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

All regressions include grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls. The reduced-form columns regress the outcome directly on the ownership-based shift-share instrument.

Figure 8 reports a dynamic reduced-form timing check for P5–P10 based on interactions between baseline ownership and year indicators. The relationship is slightly positive in 2010 and becomes negative from 2011 onward, consistent with the timing and sign of the main result.

7.2. Distributional Effects and Tenure Mechanisms

Although the main focus of the study is the lower tail of the income distribution, we estimate effects across all income groups to assess whether changes in low-income exposure are part of a broader redistribution of neighbourhood income composition. Table 7 presents the preferred IV estimates across the income distribution, obtained by varying income group k in Equation 5. All specifications include the full set of time-varying controls, grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the ownership-based shift-share instrument. The period of analysis covers 2009–2023, corresponding to the pre-implementation, implementation, and consolidation phases of the allocation reforms.

The coefficients should be interpreted as changes in exposure shares. Since both the dependent variable and the social-housing variable are measured as shares, a coefficient β implies that a one-percentage-point increase in the local share of social housing changes the spatial-lag of the corresponding income group by $\beta/100$ units, or equivalently by β percentage points. The estimates therefore describe changes in the income composition of the surrounding 400-metre neighbourhood environment, rather than changes in




Figure 8: Dynamic reduced-form timing check for P5–P10 exposure. The figure reports coefficients and 95% confidence intervals from interactions between baseline 2009 ownership share and year indicators. The omitted reference year is 2009. The vertical dashed line marks the beginning of the post-2010 social-housing allocation regime.

individual residential mobility.

The upper panel reports results for the lower and middle parts of the income distribution. For P0–P5, the coefficient is positive but statistically insignificant. A one-percentage-point increase in the local share of social housing is associated with a 0.007 percentage-point increase in local exposure to P0–P5 households and a 0.087 percentage-point decrease in local exposure to P5–P10 households. The P0–P5 estimate is small and imprecise.

By contrast, the coefficient for P5–P10 is negative and statistically significant. A one-percentage-point increase in the local social-housing tenure share reduces the average P5–P10 income share across neighbouring 100m grid cells within 400 metres by approximately 0.087 percentage points. Equivalently, a 10-percentage-point increase reduces P5–P10 exposure by approximately 0.87 percentage points. This is the clearest result within the lower tail of the income distribution. It shows that the effect is not uniform among low-income households: while the estimate for the very bottom group is small and statistically insignificant, exposure to the second half of the bottom decile declines substantially.

This contrast is consistent with the descriptive patterns documented in Section 4: P5–P10 households are more strongly concentrated in social-housing environments, whereas P0–P5 households are increasingly located in grids with lower social-housing shares and higher private-rental shares.

Table 7: Effects of local social-housing tenure-share variation on local income exposure by income-percentile group

Lower and Middle Income Groups				
	P0–P5	P5–P10	P10–P30	P30–P50
Share of social housing	0.007 (0.0061)	-0.087 ^a (0.0077)	0.175 ^a (0.0147)	0.118 ^a (0.0133)
Upper Income Groups				
	P50–P70	P70–P90	P90–P95	P95–P100
Share of social housing	-0.044 ^a (0.0119)	-0.121 ^a (0.0136)	-0.027 ^a (0.0066)	-0.021 ^a (0.0063)
Population density			Yes	
HH size			Yes	
HH composition			Yes	
Year FE			Yes	
Grid FE			Yes	
Municipality × Year FE			Yes	
Ownership: IV			Yes	
Kleibergen–Paap Wald F-stat			343.125	
Observations			903,240	
Groups (Grids)			60,216	

Notes: Robust standard errors clustered at the grid level.

Significance levels: ^a p<0.01, ^b p<0.05, ^c p<0.1.

IV uses the ownership-based shift-share instrument. All specifications include grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls.

The results for the remaining income groups reveal a broader redistribution of local exposure across the income distribution. Exposure to households in P10–P30 and P30–P50 increases, while exposure to households above the median decreases. Because income-group shares are compositional, increases in local exposure for some groups necessarily correspond to reductions for others. Figure 9 summarises this pattern by plotting the estimated coefficients and their confidence intervals. The figure highlights the non-monotonic structure of the results: the effects differ within the lower tail, become positive for lower-middle and middle-income groups, and turn negative for higher-income groups.

Taken together, the distributional estimates are consistent with a rental-sector reallocation mechanism rather than a uniform concentration of all low-income households. P5–P10 households are more strongly concentrated in social-housing environments.

To examine the tenure segments underlying the main P5–P10 result, we re-estimate Equation 5 after replacing the dependent variable with tenure-specific measures of local exposure. For each tenure category $r \in \{\text{social rent, private rent, ownership}\}$, the specification is:

$$\text{spatial-lag}_{it}^{P5-P10,r} = \alpha + \beta_r \widehat{\text{SocH}}_{it} + X'_{it}\gamma + \mu_i + \tau_t + \lambda_{m(i),t} + \varepsilon_{it}^r \quad (6)$$

where $\text{spatial-lag}_{it}^{P5-P10,r}$ measures cell-weighted local exposure to P5–P10 households belonging to tenure category r . The treatment, controls, fixed effect, instrument, and

estimation sample follow the preferred IV specification.

Table 8: Tenure Decomposition of P5–P10 Exposure

	Social renters	Private renters	Owner-occupiers
Share of social housing	-0.0551 ^a (0.0061)	-0.0211 ^a (0.0029)	-0.0009 (0.0019)
Population density	Yes	Yes	Yes
HH size	Yes	Yes	Yes
HH composition	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Grid FE	Yes	Yes	Yes
Municipality $\times$ Year FE	Yes	Yes	Yes
Ownership: IV	Yes	Yes	Yes
Kleibergen–Paap Wald F-stat	343.125	343.125	343.125
Observations	903,240	903,240	903,240
Groups (Grids)	60,216	60,216	60,216

Notes: Robust standard errors clustered at the grid level are reported in parentheses. The dependent variables measure cell-weighted local exposure to P5–P10 households separately for social renters, private renters, and owner-occupiers. All regressions include grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls. The local social-housing share is instrumented using the ownership-based shift-share instrument. The three tenure-specific coefficients sum to approximately -0.0771 , compared with -0.0870 for total P5–P10 exposure in the common estimation sample. The components do not sum exactly to the total because approximately 1% of households have missing or unclassified tenure information. Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

The tenure-specific estimates show (Table 8), that the negative P5–P10 effect is concentrated among social renters and, to a lesser extent, private renters, while the owner-occupier component is small and statistically insignificant. The three tenure-specific coefficients account for most of the total P5–P10 coefficient; the remaining difference, around 1% of the sample, reflects households with missing or unclassified tenure information. This pattern indicates that changes in local social-housing tenure composition primarily reshape the spatial distribution of P5–P10 households through rental tenures.

The contrast with P0–P5 further indicates that the spatial consequences of tenure reallocation differ within the bottom decile. The results therefore do not indicate a common response across all low-income households, but differentiated effects within the lower tail of the income distribution.

Overall, Table 7 points to heterogeneous effects within the lower tail and across the income distribution. The evidence is strongest for P5–P10, where increased social-housing share reduces local exposure. The positive coefficients for P10–P50 and the negative coefficients for higher-income groups indicate that instrument-induced local social-housing tenure-share variation is associated with a broader reshaping of neighbourhood income composition at the 400-metre scale.

7.3. Heterogeneity across Urban Contexts

The baseline results presented above document average effects of local social-housing tenure-share variation on local income exposure across the Netherlands. However, these




Figure 9: Plot coefficients and confidence intervals. Effects of local social-housing tenure-share variation on local income exposure by income-percentile group.

effects are unlikely to be uniform across space. Urban structure shapes housing demand, tenure composition, and competitive pressure for social housing, all of which may condition how allocation rules translate into local income exposure outcomes. Dense urban areas face stronger demand pressures and tighter housing markets, while less urbanised and rural areas operate under different spatial constraints. Examining heterogeneity by degree of urbanisation therefore helps clarify where changes in local tenure composition matter most.

To assess this heterogeneity, we interact the local share of social housing with indicators for the degree of urbanisation of the municipality in which each grid is located. We rely on CBS's 2009 classification, distinguishing extremely urbanised, strongly urbanised, moderately urbanised, and less urbanised/rural areas. These categories capture meaningful differences in density and housing-market structure across the Netherlands, ranging from dense metropolitan cores such as Amsterdam, Rotterdam, and The Hague to medium-sized cities and more peripheral municipalities.

The estimating equation is given by:

$$\text{spatial-lag}_{it}^k = \alpha + \sum_{u \in \{Ext, Str, Mod, Rur\}} \beta_u (SocH_{it} \cdot D_{m(i)}^u) + \gamma' \mathbf{X}_{it} + \mu_i + \tau_t + \lambda_{m(i),t} + \varepsilon_{it}^k, \quad (7)$$

where u denotes the degree of urbanisation of the municipality where grid i is located. The interaction term $SocH_{it} \cdot D_{m(i)}^u$ allows the effect of social housing on local exposure, $\text{spatial-lag}_{it}^k$, to vary across urban contexts. As in the baseline IV specification, the social-housing interactions are treated as endogenous and instrumented using the corresponding interactions between the ownership-based shift-share instrument and the urbanisation

indicators. The superscript k indicates that the dependent variable is estimated separately for each income group. The coefficients β_{it} therefore measure the effect of changes in the local share of social housing within each urbanisation category, while the fixed effects ensure that identification comes from within-grid changes over time rather than from cross-sectional differences between areas.

Table 9 reports the results for selected income groups, focusing on the lower tail (Po–P5 and P5–P10) and the upper tail (P90–P95 and P95–P100) of the income distribution.

Table 9: Effects of Local Social-Housing Tenure-Share Variation by Income-Percentile Group and Degree of Urbanisation

	Po–P5	P5–P10	P90–P95	P95–P100
Social H. $\times$ Extremely urbanised	-0.017 (0.0104)	-0.038 ^a (0.0079)	-0.043 ^a (0.0075)	-0.018 ^a (0.0067)
Social H. $\times$ Strongly urbanised	-0.013 (0.0158)	-0.069 ^a (0.0241)	0.006 (0.0240)	-0.039 (0.0266)
Social H. $\times$ Moderately urbanised	-0.029 (0.0194)	-0.064 ^b (0.0283)	0.021 (0.0339)	0.025 (0.0330)
Social H. $\times$ Less urbanised / Rural	0.001 (0.0432)	-0.060 (0.0456)	0.015 (0.0615)	0.107 (0.0821)
Population density			Yes	
Household size			Yes	
Household composition			Yes	
Ownership-based IV interactions			Yes	
Year FE			Yes	
Grid FE			Yes	
Municipality $\times$ Year FE			Yes	
Kleibergen–Paap Wald F-stat			6.595	
Wald test: equal urbanisation effects	0.21	0.73	2.39	1.18
Wald test p-value	0.8874	0.5366	0.0671	0.3139
Observations			296,205	
Unique grids			19,747	

Robust SEs in parentheses, clustered at the grid level. Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. The Wald test reports the joint test that the social-housing coefficients are equal across urbanisation categories. Sample composition by urbanisation category is as follows: extremely urbanised, 135,495 observations and 9,033 grids; strongly urbanised, 92,925 observations and 6,195 grids; moderately urbanised, 65,535 observations and 4,369 grids; less urbanised/rural, 2,250 observations and 150 grids. The social-housing interactions are instrumented using the corresponding interactions between the ownership-based shift-share instrument and the urbanisation indicators.

The heterogeneous estimates show that the lower-tail results are not uniform across income groups or urban contexts. For Po–P5, all coefficients are statistically insignificant. The point estimates are negative in extremely, strongly, and moderately urbanised municipalities and close to zero in less urbanised/rural areas.

For P5–P10, the pattern is clearer. The coefficients are negative and statistically significant in extremely, strongly, and moderately urbanised municipalities. A one-percentage-point increase in the local share of social housing reduces exposure to P5–P10

households by 0.038 percentage points in extremely urbanised municipalities, 0.069 percentage points in strongly urbanised municipalities, and 0.064 percentage points in moderately urbanised municipalities. The coefficient remains negative in less urbanised and rural areas, but it is not statistically significant. The Wald test does not reject equality across urbanisation categories ($p = 0.5366$).

These results reinforce the main finding from the baseline specification: within the lower tail, the strongest evidence of reduced exposure appears for households between the 5th and 10th percentiles. The P5–P10 estimates are negative in all four urbanisation categories and statistically significant in the three more urbanised categories, although the equality test does not indicate statistically significant differences across categories.

The upper-tail results are broadly consistent with the baseline estimates. For P90–P95 and P95–P100, the coefficients are negative and statistically significant only in extremely urbanised municipalities. The estimates for the other urbanisation categories are mixed and statistically imprecise.

Taken together, the heterogeneity analysis qualifies the baseline result rather than changing it. The evidence for P0–P5 remains weak, while the P5–P10 coefficients are negative in every urbanisation category and statistically significant in the three more urbanised categories. However, the interaction specification has a weaker joint first stage than the baseline model, with a Kleibergen–Paap Wald F-statistic of 6.595, and the less urbanised/rural category contains only 150 grids. The category-specific estimates should therefore be interpreted cautiously. Overall, the results support a negative P5–P10 pattern across urban contexts, but the Wald test does not provide evidence that the effects differ significantly across categories.

7.4. *Robustness*

We assess the stability of the preferred estimates along six dimensions: alternative sample restrictions, the inclusion of private rental share as a sensitivity control, differential trends associated with baseline neighbourhood characteristics, household-weighted exposure, the exclusion of major municipalities one at a time, and municipality-clustered inference.

Table 10 first examines alternative sample restrictions. Column (1) reproduces the preferred specification, which restricts the sample to grids with at least 10 households and some social housing in 2009. Column (2) relaxes the baseline social-housing condition and includes grids without social housing in 2009, increasing the sample from 903,240 to 1,698,330 observations. For P5–P10, the coefficient remains negative and statistically significant, although it becomes smaller in absolute value. This indicates that the negative effect for P5–P10 does not depend on the 2009 social-housing restriction. At the same time, the smaller coefficient in column (2) is consistent with the sample logic discussed above: including grids without pre-existing social housing may also capture variation related to new social-housing supply. For this reason, we keep column (1) as the preferred specification.

Column (3) imposes a stricter household-count restriction by requiring at least 20 households per grid, while keeping the baseline social-housing condition. The sample falls to 770,629 observations. The coefficient for P5–P10 remains negative, statistically

significant, and close to the preferred estimate. By contrast, the coefficient for Po–P5 is more sensitive to these changes in the sample definition.

Table 10: Alternative Sample Restrictions

	(1)	(2)	(3)
Po–P5 exposure	0.0071 (0.0061)	0.0481 ^a (0.0032)	-0.0035 (0.0071)
P5–P10 exposure	-0.0870 ^a (0.0077)	-0.0434 ^a (0.0031)	-0.1012 ^a (0.0093)
Household-count per grid	≥10	≥10	≥20
Social housing in 2009 > 0	Yes	No	Yes
Population density	Yes	Yes	Yes
HH size	Yes	Yes	Yes
HH composition	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Grid FE	Yes	Yes	Yes
Municipality × Year FE	Yes	Yes	Yes
Ownership: IV	Yes	Yes	Yes
Kleibergen–Paap Wald F-stat	343.125	1416.647	273.943
Observations	903,240	1,698,330	770,629
Groups (Grids)	60,216	113,222	53,299

Notes: Robust standard errors clustered at the grid level. Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$. Column (1) uses the baseline sample: grids with at least 10 households and some social housing in 2009. Column (2) removes the baseline sample restrictions. Column (3) imposes a stricter household-count restriction while retaining the baseline social-housing condition.

Table 11 adds the share of private rental housing as a sensitivity control, rather than treating it as a separate treatment variable. We do not include this variable in the preferred specification because private rental housing forms part of the broader tenure adjustment documented above and may itself respond to local housing-market changes. Controlling for it may therefore absorb part of the adjustment process rather than isolate an additional exogenous source of variation.

The estimates remain close to the preferred specification. For Po–P5, the coefficient changes from 0.0071 to 0.0069 and remains statistically insignificant. For P5–P10, the coefficient changes from -0.0870 to -0.0837 and remains negative and statistically significant. This indicates that the main P5–P10 result is not driven by changes in the private rental share.

Table 11: Private Rental Share Control

	Po–P5 exposure	P5–P10 exposure
Share of social housing	0.0069 (0.0059)	-0.0837 ^a (0.0069)
Private rental share	0.0024 (0.0028)	-0.0426 ^a (0.0033)
Population density	Yes	Yes
HH size	Yes	Yes
HH composition	Yes	Yes
Year FE	Yes	Yes
Grid FE	Yes	Yes
Municipality $\times$ Year FE	Yes	Yes
Ownership: IV	Yes	Yes
Kleibergen–Paap Wald F-stat	482.933	482.933
Observations	903,240	903,240
Groups (Grids)	60,216	60,216

Notes: Robust standard errors clustered at the grid level.

Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

This specification adds the private rental share as an additional control. All regressions include grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls. The share of social housing is instrumented using the ownership-based shift-share instrument.

Table 12 re-estimates the preferred IV specification while allowing 2009 neighbourhood characteristics to follow differential linear trends. The P5–P10 estimate remains negative and statistically significant and becomes somewhat larger in magnitude. The Po–P5 estimate is sensitive to specification: its magnitude and statistical significance change under alternative sample definitions and trend controls, and its sign is not consistent across all robustness checks. The persistence of the P5–P10 result indicates that it is not driven by differential trends associated with baseline population density, private-rental share, household size, or the initial income composition of the grid.

Table 12: Baseline Neighbourhood Characteristics Interacted with Linear Time Trends

	Po–P5 exposure	P5–P10 exposure
Share of social housing	0.0383 ^a (0.0059)	-0.1085 ^a (0.0080)
Time-varying controls	Yes	Yes
2009 population density × trend	Yes	Yes
2009 private-rental share × trend	Yes	Yes
2009 household size × trend	Yes	Yes
2009 income-group share × trend	Yes	Yes
Year FE	Yes	Yes
Grid FE	Yes	Yes
Municipality × Year FE	Yes	Yes
Ownership: IV	Yes	Yes
Kleibergen–Paap Wald F-stat	402.405	382.637
Observations	903,240	903,240
Groups (grids)	60,216	60,216

Notes: Robust standard errors clustered at the grid level are reported in parentheses. The dependent variables are the cell-weighted local-exposure measures for Po–P5 and P5–P10 households. Both specifications include the full set of time-varying controls, grid fixed effect, year fixed effect, and municipality-by-year fixed effect. The specifications additionally interact 2009 population density, private-rental share, household size, and the corresponding baseline income-group share with a linear time trend. The local social-housing share is instrumented using the ownership-based shift-share instrument. Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

Table 13 replaces the baseline cell-weighted spatial lag with a household-weighted exposure measure. This alternative gives more weight to neighbouring grids containing more households. The P5–P10 coefficient remains negative and statistically significant and is close to the preferred estimate, while the Po–P5 coefficient is effectively zero. The main result therefore does not depend on assigning equal weight to neighbouring grid cells.

Table 13: Robustness: Household-Weighted Local Income Exposure

	Po–P5 exposure	P5–P10 exposure
Share of social housing	-0.0003 (0.0069)	-0.0939 ^a (0.0085)
Population density	Yes	Yes
HH size	Yes	Yes
HH composition	Yes	Yes
Year FE	Yes	Yes
Grid FE	Yes	Yes
Municipality $\times$ Year FE	Yes	Yes
Ownership: IV	Yes	Yes
Kleibergen–Paap Wald F-stat	340.400	340.400
Observations	902,385	902,385
Groups (Grids)	60,159	60,159

Notes: Robust standard errors clustered at the grid level are reported in parentheses. The dependent variables are household-weighted measures of local exposure. For each focal grid, they measure the share of all households in neighbouring 100m grid cells within 400 metres that belong to the corresponding income group. The focal grid is excluded. All regressions include grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls. The local social-housing share is instrumented using the ownership-based shift-share instrument. The household-weighted outcome is unavailable for a small number of grid-year observations, resulting in a slightly smaller estimation sample. Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

Table 14 addresses whether the main results are driven by one of the largest municipalities. We re-estimate the preferred specification excluding Amsterdam, Rotterdam, The Hague, and Utrecht one at a time.

The P5–P10 coefficient remains negative and statistically significant in all cases. Excluding Amsterdam and Rotterdam makes the coefficient more negative than in the preferred specification. Excluding The Hague leaves the estimate unchanged, while excluding Utrecht makes it slightly smaller in absolute value. Overall, these changes do not alter the main result: increased social-housing share reduces local exposure to P5–P10 households, and this pattern is not driven by a single major municipality.

Panel A confirms the weaker pattern for Po–P5. The coefficient changes sign across exclusions and is only statistically significant when Utrecht is excluded. This is consistent with the previous robustness checks and supports treating Po–P5 as suggestive rather than central.

Table 14: Excluding Major Municipalities One by One

	Excl. Amsterdam	Excl. Rotterdam	Excl. The Hague	Excl. Utrecht
Panel A: P0–P5 exposure				
Share of social housing	-0.0048 (0.0071)	0.0026 (0.0060)	0.0071 (0.0061)	0.0175 ^a (0.0060)
Panel B: P5–P10 exposure				
Share of social housing	-0.1133 ^a (0.0102)	-0.0924 ^a (0.0077)	-0.0870 ^a (0.0077)	-0.0852 ^a (0.0078)
Population density	Yes	Yes	Yes	Yes
HH size	Yes	Yes	Yes	Yes
HH composition	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Grid FE	Yes	Yes	Yes	Yes
Municipality × Year FE	Yes	Yes	Yes	Yes
Ownership: IV	Yes	Yes	Yes	Yes
Kleibergen–Paap Wald F-stat	242.600	353.624	343.125	332.030
Observations	845,610	854,610	903,240	878,820
Groups (Grids)	56,374	56,974	60,216	58,588

Notes: Robust standard errors clustered at the grid level.

Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

Each column excludes one major municipality from the estimation sample. All regressions include grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls. The share of social housing is instrumented using the ownership-based shift-share instrument.

Finally, because the local-exposure measure is constructed using overlapping 400 m buffers, residuals may be spatially correlated across nearby grids. To address this concern, we complement the baseline grid-clustered standard errors with municipality-clustered standard errors, keeping the specification, instrument, controls, and fixed effects unchanged. Table 15 reports the results. The coefficient for P5–P10 remains negative and statistically significant, while the P0–P5 coefficient is not statistically distinguishable from zero.

Table 15: Alternative Inference: Municipality-Clustered Standard Errors

	Po–P5 exposure	P5–P10 exposure
Share of social housing	0.0071 (0.0235)	-0.0870 ^a (0.0301)
Population density	Yes	Yes
Household size	Yes	Yes
Household composition	Yes	Yes
Year FE	Yes	Yes
Grid FE	Yes	Yes
Municipality × Year FE	Yes	Yes
Ownership: IV	Yes	Yes
Municipality-clustered SE	Yes	Yes
Kleibergen–Paap Wald F-stat	38.120	38.120
Observations	903,240	903,240
Municipality clusters	253	253

Notes: Robust standard errors clustered at the municipality level in parentheses. The specification is identical to the preferred IV model, with grid fixed effect, year fixed effect, municipality-by-year fixed effect, and the full set of controls. The share of social housing is instrumented using the ownership-based shift-share instrument. Significance levels: ^a $p < 0.01$, ^b $p < 0.05$, ^c $p < 0.1$.

8. Conclusions

This paper examines how variation in the local social-housing tenure share within the post-2010 allocation regime affects local economic segregation in the Netherlands between 2009 and 2023. Using administrative panel data aggregated to 100 m grid cells, we measure neighbourhood income exposure as the average income-group share across neighbouring cells within 400 metres. To address the endogeneity of local tenure composition, we instrument the social-housing tenure share using baseline ownership interacted with leave-one-municipality-out national changes in owner-occupation.

The central finding is concentrated among households between the 5th and 10th percentiles of the income distribution. A 10-percentage-point increase in the local social-housing tenure share reduces exposure to this group by approximately 0.87 percentage points. By contrast, the estimate for households in the bottom 5% is small and imprecise, showing that the effects differ even within the bottom decile. The descriptive analysis provides context for this contrast: between 2009 and 2023, grids containing households in the bottom 5% had a lower average social-housing share and a markedly higher private-rental share, whereas households between the 5th and 10th percentiles remained more closely associated with social-housing environments. This descriptive divergence helps interpret the weaker estimated response for the bottom-5% group. The tenure decomposition further indicates that the reduction in P5–P10 exposure is concentrated among social renters and, to a lesser extent, private renters, with little evidence of a corresponding effect among owner-occupiers. The central P5–P10 estimate remains stable across alternative sample restrictions, the inclusion of the private-rental share as a sensitivity control, differential trends associated with baseline neighbourhood characteristics,

household-weighted exposure, municipality-clustered inference, and major-municipality exclusions.

The results also indicate that local social-housing tenure-share variation reshapes neighbourhood income composition beyond the lower tail. Exposure to lower-middle and middle-income households increases, while exposure to higher-income households decreases. Because income-group shares are compositional, these results should be interpreted as a redistribution of local exposure across income groups, rather than as independent effects for each segment of the distribution. The P5–P10 coefficient is also negative in all four urbanisation categories and statistically significant in the three more urbanised categories. However, the category-specific estimates are less precisely identified than the baseline result and should therefore be interpreted cautiously.

These patterns are consistent with a broader process of rental-sector reallocation. Although the aggregate social-housing stock remains relatively stable over the period, social housing becomes increasingly concentrated among lower-income households. At the same time, private renting expands substantially, while owner-occupation grows primarily among middle- and higher-income households. Together with the tenure decomposition, these developments suggest that income targeting within social housing does not mechanically produce higher local exposure to all low-income households. Its spatial consequences depend on how households are distributed across social renting, private renting, and owner-occupation.

Importantly, the estimates should not be interpreted as the isolated policy effect of the 90% rule or as the effect of individual allocation flows. They capture the neighbourhood-composition effects of variation in the local social-housing tenure share within the broader post-2010 allocation regime. The preferred estimates also pertain to grids in which social housing was already present in 2009. They refer to very local exposure, measured within approximately 400 metres, and should not be interpreted as evidence about residential mobility, displacement, or segregation at larger spatial scales.

Overall, the findings underline that the spatial implications of income-targeted social housing cannot be assessed in isolation from the wider housing system. They depend on the existing location of the social-housing stock, the structure of local housing markets, and the availability of alternative tenures. Future research should examine additional measures and spatial scales of economic segregation and provide more direct evidence on the household and tenure transitions underlying the estimated exposure effects. The central contribution of this paper is to show that broad definitions of the low-income population can conceal materially different local exposure responses: changes in local social-housing tenure composition affect households in P5–P10 differently from those at the very bottom of the income distribution.

References

- Bayer, P., Ferreira, F., and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy*, 115(4):588–638.
- Bischoff, K. and Reardon, S. F. (2012). Residential segregation and the preservation of concentrated poverty. *Social Science Research*, 41(5):1114–1127.
- Boelhouwer, P. (2018). The housing market in the netherlands as a social system. *Journal of Housing and the Built Environment*, 33(2):229–246.
- Buurma-Olsen, J., Koster, H. R., van Ommeren, J., and Sinninghe Damsté, J. (2025). Quantifying misallocation of public housing. *Journal of Public Economics*, 242:105272.
- Chetty, R. and Hendren, N. (2018). The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates*. *The Quarterly Journal of Economics*, 133(3):1163–1228.
- Ellen, I. G. and O'Regan, K. (2011). How low-income neighborhoods change: Entry, exit, and enhancement. *Regional Science and Urban Economics*, 41(2):89–97.
- Elsinga, M., Haffner, M., and van der Heijden, H. (2014). The changing role of social housing in the netherlands. *Journal of Housing and the Built Environment*, 29(4):569–590.
- Elsinga, M. and Lind, H. (2013). The Effect of EU-Legislation on Rental Systems in Sweden and the Netherlands. *Housing Studies*, 28(7):960–970.
- Flynn, L. B. and Montalbano, G. (2024). A turn towards post-neoliberalism? Housing policy paradigm crisis in Europe. *West European Politics*, pages 1–30.
- Galster, G. C. (2012). The mechanism(s) of neighbourhood effects: Theory, evidence, and policy implications. In Ham, M. V., Manley, D., Bailey, N., Simpson, L., and Maclennan, D., editors, *Neighbourhood Effects Research: New Perspectives*, pages 23–56. Springer Netherlands, Dordrecht.
- Gruis, V. and Priemus, H. (2008). European Competition Policy and National Housing Policies: International Implications of the Dutch Case. *Housing Studies*, 23(3):485–505.
- Hipp, J. R. (2007). Block, tract, and levels of aggregation: Neighborhood structure and crime and disorder as a case in point. *American Sociological Review*, 72(5):659–680.
- Hochstenbach, C. (2017). State-led gentrification and the changing geography of market-oriented housing policies. *Housing, Theory and Society*, 34(4):399–419.
- Hochstenbach, C. (2022). *Uitgewoond: Waarom het hoog tijd is voor een nieuwe woonpolitiek*. Das Mag Uitgeverij, Amsterdam.
- Jenks, M. and Dempsey, N. (2007). Defining the neighbourhood: Challenges for empirical research. *Town Planning Review*, 78(2):153–177.

- Massey, D. S. and Denton, N. A. (1988). The dimensions of residential segregation. *Social Forces*, 67(2):281–315.
- McCormack, G. R. and Shiell, A. (2010). The relationship between neighbourhood environment and physical activity: A review. *International Journal of Behavioral Nutrition and Physical Activity*, 7(1):1–11.
- Ministerie van Binnenlandse Zaken en Koninkrijksrelaties (2019). Cijfers over wonen en bouwen 2019. Accessed: 2025-04-30.
- Nelson, R., Warnier, M., and Verma, T. (2024). Housing inequalities: The space-time geography of housing policies. *Cities*, 145:104727.
- Nieuwenhuis, J., Tammaru, T., Van Ham, M., Hedman, L., and Manley, D. (2020). Does segregation reduce socio-spatial mobility? Evidence from four European countries with different inequality and segregation contexts. *Urban Studies*, 57(1):176–197.
- Pinkster, F. M. (2014). Social networks, social resources and social mobility in a low-income neighbourhood in the netherlands. *Urban Studies*, 51(7):1404–1420.
- Priemus, H. (2002). Social housing policy in the european union: Past, present and perspectives. *Urban Studies*, 39(2):191–200.
- Priemus, H. and Gruis, V. (2011). Social housing and illegal state aid: The agreement between european commission and dutch government. *International Journal of Housing Policy*, 11(1):89–104.
- Reardon, S. F. and O’Sullivan, D. (2004). Measures of spatial segregation. *Sociological Methodology*, 34(1):121–162.
- Sampson, R. J. (2012). *Great American city: Chicago and the enduring neighborhood effect*. Univ. of Chicago Press, Chicago, Ill.
- San Millán, J., Cottineau-Mugadza, C., and van Ham, M. (2026). Why inequality does not always lead to segregation: The role of social housing in the netherlands. *Urban Housing*. In press.
- Scanlon, K., Whitehead, C., and Fernández Arrigoitia, M., editors (2015). *Social Housing in Europe*. Wiley Blackwell, Oxford.
- Statistics Netherlands (CBS) (2018). The netherlands 10 years after lehman brothers. Accessed: 2025-04-30.
- Uitermark, J., Hochstenbach, C., and Groot, J. (2023). The impact of public policy on multiple dimensions of spatial inequality. *Urban Geography*.
- van Deursen, H. (2023). The people’s housing: Woningcorporaties and the dutch social housing system, part 1: The history. Available online.

- van Duijne, R. J. and Ronald, R. (2018). The unraveling of amsterdam's unitary rental system. *Journal of Housing and the Built Environment*, 33(3):633–651.
- Van Gent, W. and Hochstenbach, C. (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.
- van Ham, M., Manley, D., Bailey, N., Simpson, L., and Maclennan, D., editors (2012). *Neighbourhood Effects Research: New Perspectives*. Springer, Dordrecht.

Appendix A. Additional Municipality-Level Exposure Tables

Table A.16: Mean of Local Exposure to Po–P5 and P5–P10 Households by Municipality and Degree of Urbanisation: Extremely urbanised

Municipality	Grids Pop. 2023		Po–P5			P5–P10		
			Lag 2009	Lag 2023	Δ	Lag 2009	Lag 2023	Δ
Amsterdam	5,233	736,522	0.081	0.076	-0.005	0.075	0.075	0.000
Rotterdam	4,803	568,907	0.065	0.064	0.000	0.068	0.066	-0.002
The Hague	3,900	484,595	0.072	0.067	-0.005	0.059	0.059	0.000
Utrecht	2,674	289,373	0.092	0.080	-0.011	0.059	0.053	-0.005
Eindhoven	2,765	200,339	0.067	0.065	-0.002	0.052	0.049	-0.003
Tilburg	2,432	187,128	0.073	0.065	-0.008	0.056	0.050	-0.006
Haarlem	1,411	148,905	0.047	0.037	-0.010	0.044	0.040	-0.004
Zaandam	1,547	133,550	0.035	0.030	-0.004	0.038	0.039	0.001
Amersfoort	1,685	130,313	0.039	0.027	-0.012	0.036	0.035	-0.001
Zoetermeer	1,240	111,403	0.033	0.023	-0.010	0.028	0.031	0.003
Leiden	1,130	108,555	0.083	0.095	0.012	0.050	0.045	-0.005
Dordrecht	1,237	106,283	0.044	0.029	-0.015	0.042	0.042	0.000
Delft	839	88,733	0.104	0.125	0.021	0.052	0.047	-0.005
Alkmaar	1,133	86,638	0.041	0.028	-0.013	0.037	0.039	0.002
Hilversum	997	80,581	0.047	0.035	-0.012	0.034	0.037	0.002
Purmerend	943	72,150	0.034	0.025	-0.009	0.030	0.034	0.004
Schiedam	606	71,745	0.051	0.043	-0.008	0.052	0.051	-0.001
Leidschendam-Voorburg	810	70,298	0.042	0.035	-0.007	0.030	0.033	0.003
Spijkensisse	835	66,477	0.027	0.024	-0.003	0.030	0.035	0.005
Vlaardingen	647	65,979	0.037	0.032	-0.005	0.044	0.050	0.007
Gouda	742	65,264	0.038	0.027	-0.011	0.038	0.042	0.004
Velsen	748	61,936	0.034	0.024	-0.011	0.031	0.038	0.008
Capelle aan den IJssel	682	60,476	0.038	0.032	-0.006	0.033	0.036	0.003
Nieuwegein	740	56,926	0.034	0.023	-0.011	0.028	0.028	0.000
Katwijk	589	54,731	0.028	0.017	-0.010	0.029	0.032	0.003
Rijswijk	444	45,601	0.048	0.038	-0.010	0.040	0.047	0.007
Zwijndrecht	523	40,687	0.031	0.023	-0.008	0.029	0.035	0.006
Beverwijk	442	37,550	0.030	0.026	-0.004	0.035	0.038	0.003
Bussum	422	30,368	0.048	0.030	-0.017	0.033	0.030	-0.003
Gorinchem	373	30,151	0.035	0.027	-0.009	0.032	0.040	0.008
Papendrecht	386	28,687	0.024	0.019	-0.005	0.026	0.033	0.007
Maassluis	294	28,313	0.027	0.025	-0.002	0.031	0.039	0.008
Leiderdorp	302	24,742	0.036	0.025	-0.011	0.024	0.026	0.002
Diemen	233	22,425	0.060	0.053	-0.007	0.039	0.040	0.001
Sliedrecht	259	20,884	0.024	0.019	-0.005	0.025	0.038	0.014
Zandvoort	209	15,334	0.057	0.034	-0.023	0.040	0.042	0.002
Rozenburg	163	11,054	0.025	0.024	-0.002	0.026	0.038	0.012

Note: Table reports municipality-level averages of local exposure to Po–P5 and P5–P10 households within a 400 m buffer. Grids refers to the number of 100m grid cells used in the calculation. Δ indicates the change between 2009 and 2023. Population refers to the 2023 population contained in the 100m grids included in the filtered sample, not total municipal population. Positive changes in the two Δ columns are highlighted in red.

Table A.17: Mean of Local Exposure to Po–P5 and P5–P10 Households by Municipality and Degree of Urbanisation: Strongly urbanised

Municipality	Grids	Pop. 2023	Po–P5			P5–P10		
			Lag 2009	Lag 2023	Δ	Lag 2009	Lag 2023	Δ
Almere	2,197	168,803	0.048	0.032	-0.016	0.040	0.036	-0.004
Breda	2,163	152,983	0.060	0.047	-0.013	0.044	0.040	-0.004
Den Bosch	1,755	122,821	0.046	0.031	-0.014	0.041	0.039	-0.002
Haarlemmermeer	1,725	122,082	0.033	0.023	-0.010	0.025	0.025	0.000
Amstelveen	903	74,868	0.045	0.038	-0.007	0.027	0.031	0.004
Alphen aan den Rijn	764	63,537	0.033	0.018	-0.015	0.025	0.027	0.002
Lelystad	1,049	63,163	0.051	0.033	-0.018	0.044	0.042	-0.002
Hoorn	911	62,886	0.038	0.023	-0.015	0.034	0.038	0.004
Bergen op Zoom	813	53,472	0.044	0.024	-0.019	0.037	0.040	0.003
Zeist	695	49,005	0.053	0.036	-0.017	0.031	0.035	0.005
Den Helder	737	48,196	0.033	0.023	-0.010	0.047	0.044	-0.003
Oosterhout	740	45,124	0.033	0.022	-0.011	0.028	0.033	0.005
Heerhugowaard	666	43,515	0.031	0.020	-0.010	0.033	0.031	-0.002
Barendrecht	541	40,821	0.031	0.018	-0.014	0.024	0.025	0.001
Middelburg	648	40,483	0.045	0.034	-0.011	0.044	0.040	-0.004
Houten	588	40,192	0.033	0.017	-0.016	0.021	0.020	-0.001
Ridderkerk	497	39,393	0.026	0.022	-0.004	0.027	0.038	0.011
Vlissingen	609	39,058	0.047	0.034	-0.012	0.047	0.047	0.000
Pijnacker-Nootdorp	505	38,843	0.036	0.025	-0.011	0.025	0.022	-0.002
Soest	609	38,431	0.038	0.024	-0.014	0.030	0.031	0.001
Huizen	566	36,661	0.039	0.023	-0.015	0.034	0.034	0.000
Veldhoven	643	35,737	0.033	0.018	-0.015	0.027	0.025	-0.002
Tiel	512	34,098	0.038	0.027	-0.010	0.043	0.049	0.006
Hellevoetsluis	474	33,525	0.031	0.022	-0.009	0.026	0.031	0.005
Maarssen	437	33,384	0.034	0.018	-0.017	0.024	0.024	0.000
Heemskerk	443	32,654	0.026	0.019	-0.006	0.026	0.028	0.003
Teylingen	409	30,147	0.040	0.021	-0.019	0.026	0.029	0.003
IJsselstein	385	29,764	0.030	0.016	-0.014	0.026	0.026	-0.001
Krimpen aan den IJssel	358	25,709	0.029	0.020	-0.009	0.023	0.031	0.008
Edam-Volendam	372	24,993	0.022	0.013	-0.009	0.021	0.024	0.003
Hendrik-Ido-Ambacht	310	24,412	0.026	0.015	-0.011	0.021	0.023	0.002
Uithoorn	330	23,501	0.028	0.022	-0.006	0.022	0.030	0.008
Culemborg	335	22,904	0.039	0.020	-0.019	0.033	0.032	-0.001
Heemstede	351	22,744	0.046	0.023	-0.023	0.024	0.018	-0.006
Noordwijk	278	21,834	0.049	0.027	-0.022	0.033	0.031	-0.002
Waddinxveen	279	21,500	0.029	0.020	-0.010	0.024	0.038	0.014
Veenendaal	292	21,222	0.035	0.013	-0.021	0.030	0.028	-0.002
Oud-Beijerland	307	20,861	0.029	0.016	-0.014	0.023	0.026	0.003
Voorschoten	306	20,767	0.041	0.025	-0.017	0.024	0.023	-0.001
Baarn	305	20,384	0.042	0.022	-0.020	0.034	0.027	-0.007
Oegstgeest	274	19,612	0.049	0.033	-0.016	0.029	0.027	-0.002
Lisse	263	19,410	0.032	0.019	-0.013	0.024	0.020	-0.003
Albrandswaard	278	19,344	0.026	0.021	-0.005	0.022	0.028	0.006
Stede Broec	323	17,883	0.030	0.019	-0.011	0.029	0.027	-0.002
Nieuwerkerk aan den IJssel	232	17,873	0.028	0.016	-0.012	0.026	0.026	0.001
Hillegom	250	17,651	0.033	0.021	-0.011	0.029	0.031	0.003
Geertruidenberg	293	17,341	0.024	0.017	-0.007	0.026	0.036	0.010
Alblasserdam	217	17,246	0.025	0.021	-0.004	0.029	0.038	0.009
Bunschoten	242	17,237	0.024	0.019	-0.005	0.023	0.030	0.007
Enkhuizen	249	16,664	0.039	0.025	-0.014	0.039	0.049	0.010
Schagen	272	15,626	0.032	0.017	-0.015	0.035	0.037	0.002
Weesp	184	15,446	0.036	0.029	-0.007	0.031	0.041	0.010
Hardinxveld-Giessendam	207	14,438	0.022	0.019	-0.003	0.021	0.031	0.010
Naarden	223	13,929	0.044	0.022	-0.022	0.024	0.023	-0.001
Ouder-Amstel	162	11,477	0.037	0.031	-0.006	0.021	0.024	0.003
Uitgeest	158	10,813	0.028	0.015	-0.013	0.020	0.020	0.000
Schoonhoven	141	10,299	0.051	0.038	-0.013	0.033	0.040	0.007
Bolsward	159	8,551	0.039	0.022	-0.016	0.049	0.047	-0.002
Oostzaan	120	7,720	0.030	0.024	-0.006	0.024	0.029	0.005
Geldrop-Mierlo	74	4,165	0.030	0.023	-0.007	0.026	0.019	-0.007

Note: Table reports municipality-level averages of local exposure to Po–P5 and P5–P10 households within a 400 m buffer. Grids refers to the number of 100m grid cells used in the calculation. Δ indicates the change between 2009 and 2023. Population refers to the 2023 population contained in the 100m grids included in the filtered sample, not total municipal population. Positive changes in the two Δ columns are highlighted in red.

