ibs instytut badań strukturalnych

IBS WORKING PAPER 04/2026
JUNE 2026

DISTRIBUTIONAL PROPERTIES OF COST ALLOCATION RULES IN MULTI-FAMILY ENERGY RETROFITS: EVIDENCE FROM 4 MILLION APARTMENTS

Jakub Sokołowski
Stefan Bouzarovski

ibs working paper 04/2026

June 2026

DISTRIBUTIONAL PROPERTIES OF COST ALLOCATION RULES IN MULTI-FAMILY ENERGY RETROFITS: EVIDENCE FROM 4 MILLION APARTMENTS[•]

Jakub Sokołowski[♦]

Stefan Bouzarovski[♣]

Abstract

Collective decisions to retrofit multi-family residential buildings require co-owners to agree on how the total cost is divided among dwellings, yet the distributional properties of alternative allocation rules have been insufficiently investigated at scale. Using harmonised Energy Performance Certificate (EPC) microdata covering over 4 million apartments in almost 450,000 buildings across Poland, England and Wales, and the Netherlands, we simulate five allocation rules: area-proportional, progressive-area, emissions-proportional, inefficiency-proportional, and Shapley-value allocation. For each building, we evaluate the resulting cost-share distributions using within-building inequality indices, size-progressivity measures, and cooperative-game-theoretic stability criteria. We find that performance-based rules produce within-building Gini coefficients 2 to 11 times higher than area-proportional allocation, with systematic variation across national building stocks. These rules are also less proportional in terms of dwelling size, assigning larger cost shares to smaller dwellings than their floor-area shares warrant. For retrofit governance in multi-owner buildings, allocation design should therefore be treated as a central component of policy implementation rather than a technical-administrative choice.

Keywords: multi-family housing; energy-efficiency retrofits; cost allocation; Shapley value; cooperative game theory; Energy Performance Certificates.

JEL: D63; Q48; R31; C71; H23; Q52

[•] Jakub Sokołowski's contribution was supported by the National Science Centre, Poland (NCN), OPUS 27, grant no. 2024/53/B/HS4/03039. Stefan Bouzarovski's work on this paper was supported by: 1) the Centre for Joined-Up Sustainability Transformations (JUST), funded by the United Kingdom Research and Innovation Building a Green Future strategic theme, and the Economic and Social Research Council [grant number ES/Z504130/1]; 2) the Energy Demand Research Centre (EDRC) funded by the Engineering and Physical Sciences Research Council and the Economic and Social Research Council [grant number EP/Y010078/1]. Additional support was provided by the EU LIFE Programme project LOCATEE (LIFE23-CET-LOCATEE), grant agreement no. 101167621, via the Institute for European Energy and Climate Policy Foundation. The usual disclaimers apply. The views expressed are those of the authors only and do not necessarily reflect those of the funders. All errors and omissions are our own.

[♦] Institute for Structural Research (IBS); University of Warsaw. E-mail: jakub.sokolowski@ibs.org.pl.

[♣] The University of Manchester; The Institute for European Energy and Climate Policy (IEECP). E-mail: stefan.bouzarovski@manchester.ac.uk.

ISSN 2451-4373 | IBS ®

1. Introduction

Achieving climate and energy efficiency goals requires massive investments in retrofitting residential buildings. In Europe, programs aimed at improving insulation, heating systems, and overall energy performance in multi-family housing are expanding (Bouzarovski et al., 2018). Multi-family buildings present a distinct policy challenge relative to single-family housing. Retrofit decisions require collective agreement among co-owners, and how the costs are divided among apartment owners is a less-studied but consequential lever. Perceptions of fairness in cost allocation can significantly influence residents' willingness to undertake a retrofit project (Maestre-Andrés et al., 2019). An allocation perceived as unfair by a subset of residents may lead to disputes or vetoes in homeowner associations, delaying or blocking renovations (Grossmann, 2019). The choice of allocation rule, therefore, matters not only for distributional outcomes within buildings but for the pace of the residential energy transition itself.

Traditional cost-sharing practices in multi-owner residential buildings typically rely on dwelling size or equal shares. Maintenance and improvement costs are commonly allocated equally per unit or proportionally to floor area (Lujanen, 2010). These rules are easy to implement and broadly accepted as transparent. However, the growing use of Energy Performance Certificates (EPCs) has made dwelling-level energy performance increasingly visible. EPCs provide standardised information on the modelled energy performance of buildings or dwellings and are intended to inform owners, buyers, and tenants, but they are also increasingly linked to renovation targeting, minimum performance standards, rental restrictions, and transaction requirements (Li et al., 2019).

This creates tension for within-building cost sharing. Standard allocation rules do not account for differences in energy performance between dwellings in the same building. An apartment with poor asset-rated efficiency and therefore higher modelled energy demand pays the same (under equal shares) or only slightly more (under floor area) than an efficient one, a pattern inconsistent with performance-based allocation rules, like the "polluter-pays principle" (Luppi et al., 2012). Performance-based allocation rules, which assign costs to dwellings with worse asset-rated performance, would change this pattern but raise concerns of their own. When the worst-performing units are occupied by lower-income households, performance-based rules can conflict with ability-to-pay considerations (Anderson et al., 2003). More fundamentally, dwelling-level performance is largely a structural inheritance from the building (envelope, position, orientation, construction era) rather than something current occupants choose, which complicates any straightforward application of polluter-pays logic at this scale.

Rather than being an administrative triviality, the choice of cost-allocation rule is, therefore, fundamentally a distributional and normative decision that requires systematic attention. This leads us to ask three research questions. First, how do alternative cost allocation rules differ in the within-building distributions of cost shares they produce? Second, do the rules differ systematically in their structural distributional properties across building stocks, or are the patterns specific to the national context? Third, do allocation rules motivated by cooperative-game-theoretic fairness satisfy the cooperative-stability axioms, which would justify their use in voluntary collective-investment decisions?

We address these questions through the first large-scale empirical analysis characterising the structural distributional properties of five cost allocation rules at the within-building level. We leverage detailed data from Energy Performance Certificate databases in Poland, England and Wales in the United Kingdom (UK), and the Netherlands, covering over 4 million apartments in almost 450,000 buildings. For each building, we simulate floor-area-proportional, progressive-area, emissions-proportional, inefficiency-proportional, and Shapley values, and evaluate the resulting cost-share distributions using Gini coefficients, Atkinson indices, Theil decomposition, Kakwani size-progressivity indices, and cooperative game theory criteria (core stability and individual rationality). This approach addresses several knowledge and policy gaps. It moves beyond small samples and simulations to real building populations at scale; it introduces within-building inequality decomposition as a novel methodological contribution; and it connects the cooperative game theory literature on Shapley fairness with the public economics literature on distributional impacts.

Our contributions are threefold. First, we provide the first empirical characterisation at scale of how cost-share distributions differ across allocation rules. Performance-based rules produce within-building Gini coefficients higher than those under area-proportional allocation. Second, we adopt a within-building decomposition framework that aligns the unit of analysis with the unit of decision-making, i.e. the homeowner association, condominium, or housing community, rather than the household or country. Third, we document that the Shapley rule, often advocated for shared-infrastructure cost-sharing on cooperative-game-theoretic grounds, violates the individual-rationality condition regularly, allocating more than the cost of retrofitting that dwelling alone to at least one dwelling.

These findings have direct implications for retrofit governance in multi-owner buildings. Simple, transparent rules produce both lower dispersion and satisfy cooperative stability criteria in nearly all buildings, facilitating agreement. Performance-based methods, while theoretically appealing for incentive alignment, generate systematically larger within-building dispersion and routinely violate cooperative stability. Policymakers seeking to accelerate retrofits may benefit from default guidelines that favour transparent, size-based allocation rules, supplemented by external support mechanisms to address affordability concerns at the household level rather than within the allocation rule itself.

The remainder of the paper is organised as follows: Section 2 reviews the relevant literature, while Section 3 describes the data sources and institutional context, defines the allocation methods, and presents the distributional and stability metrics. In Section 4, we present the analysis results, then move on to country-specific implications, conceptual policy lessons, and theoretical and practical implications (in Section 5). Section 6 concludes.

2. Literature review

The existing empirical literature on retrofit cost allocation remains rather limited. Weber & Wolff (2018) conducted a field study of 10 buildings in Germany. They found that, despite achieving reductions in energy consumption, more than half of households faced increased total costs due to rent increases, demonstrating that current allocation practices can shift cost burdens in ways that disadvantage some occupants. Ahlrichs & Rockstuhl (2022) developed a max-min fairness approach based on Rawlsian principles to determine fair rent increases following retrofits; however, their work remains theoretical without large-scale empirical validation. Galvin (2023), analysing over 2 million German property advertisements, documented that sales premiums for energy efficiency are approximately four times higher than rental premiums, helping explain underinvestment in tenant-occupied buildings. However, these studies focus on landlord-tenant dynamics in rental markets rather than cost allocation among owner-occupiers in condominiums, and none systematically compares alternative allocation methods using formal inequality metrics.

A separate strand of literature applies cooperative game theory to the allocation of energy costs. The Shapley value, which allocates costs based on each participant's marginal contribution to the collective, has theoretical appeal due to its axiomatic properties of efficiency, symmetry, and the null-player condition (Shapley, 2016). Cremers et al. (2023) developed computational approximations that enable Shapley calculations for larger groups, validated on 200 UK prosumers. Beyond energy, cooperative game theory has been applied to cost-sharing in shared infrastructure, including transport (Frisk et al., 2010), with similar findings: theoretical stability properties can fail in real applications. However, this literature relies predominantly on simulations or small samples, focusing on stability rather than the dispersion of allocated shares. Whether Shapley's theoretical properties translate into structurally stable and well-dispersed allocations at the scale of real building populations remains untested.

Research on the distributional impacts of energy and climate policies provides essential context. Ohlendorf et al. (2021) conducted a meta-analysis of 53 studies across 39 countries, finding that clean energy policies tend to be regressive in higher-income countries. Borenstein & Davis (2016) documented that the top income quintile receives approximately 60% of US clean energy tax credits. Dorband et al. (2019) demonstrated an inverse-U relationship between the progressivity of carbon pricing and national income levels. A growing literature on decision-support tools for retrofit choice (Jafari and Valentin, 2017; Si and Marjanovic-Halburd, 2018) and on behavioural and procedural barriers in homeowner-association decision-making (Carr and Boyd Kramer, 2022; Ebrahimigharehbaghi et al., 2019) shows that the process of arriving at an allocation matters as much as the substantive choice of rule. Yet this literature operates at the household or national level, examining aggregate policy impacts rather than allocation rules within buildings. Finally, the polluter-pays principle, extensively debated in environmental philosophy (García-Portela, 2024), has been theorised primarily for industrial pollution and international climate burden-sharing, rather than for cost allocation among neighbours in a condominium, a setting where the attribution of "pollution" to individual occupants is conceptually fraught.

Energy Performance Certificate (EPC) databases offer rich potential for analysing building-level outcomes. Fuerst et al. (2015) and Brounen & Kok (2011) pioneered hedonic analyses using UK and Dutch EPCs, respectively, establishing price premiums for energy-efficient homes. Ali et al. (2024) linked EPCs to neighbourhood characteristics to examine efficiency gaps. However, EPCs report standardised, asset-rated values rather than measured energy use, and a substantial literature documents systematic gaps between modelled and metered consumption (van den Brom et al., 2018). No study has used EPC data to systematically characterise the distributional properties of alternative cost allocation methods, and, critically, none has examined inequality within buildings rather than across households or regions. This within-building perspective matters because it aligns with how decisions are made. In homeowner associations and cooperatives, what matters is each resident's share relative to their neighbours, not relative to households elsewhere.

3. Data, Methods, and Metrics

3.1. Data Sources and Sample

This study draws on Energy Performance Certificate (EPC) databases from three European countries: Poland, the United Kingdom, and the Netherlands. These countries were selected to represent a range of building stock characteristics and policy environments within the EU's Energy Performance of Buildings Directive framework. Each country's EPC database provides information on residential properties, including floor area and energy modelled energy and emissions values, which are the inputs to the allocation rules we evaluate (Table 1). EPC databases record dwelling characteristics; they do not record household composition, income, occupancy, or metered consumption. Our analysis is therefore at the dwelling level throughout, and we discuss the interpretive consequences in 3.1.1.

Table 1. The data overview

Country	Description	Details
Poland	We obtained data from the Central Register of Energy Performance of Buildings. Since 2009, Poland has mandated the issuance of energy certificates for buildings at the time of sale, rental, or construction, resulting in a large dataset of energy efficiency information for residential units. Our Polish sample is dominated by multi-family blocks (often from the socialist era). A key limitation is that Polish EPCs record building-level floor area rather than individual apartment areas, which precludes the computation of area-proportional allocation methods for Poland.	175,247 buildings; 932,461 apartments; median building size = 4 apartments (mean 5.3).
United Kingdom	We use the Domestic EPC Register for England and Wales, accessed via the Open Data Communities portal. This database contains over 25 million certificates, with details including each dwelling's floor area, current and potential energy efficiency (e.g., fuel consumption, CO_2 emissions), and a unique building identifier. The UK sample features many buildings with only 2–3 flats. It also contains information on tenure (owner-occupied vs rental), allowing additional stratified analysis.	182,672 buildings; 2,713,007 apartments; median building size = 2 apartments (mean 14.9). The small median reflects many two-flat conversions. The UK sample is skewed toward very small buildings, which has implications for cost-sharing dynamics (essentially two-party negotiations in many cases).
Netherlands	We utilise the EP-Online database maintained by the Netherlands Enterprise Agency. Dutch EPCs (Energy Performance Certificates) are based on the NTA 8800 calculation method and provide data on energy usage and apartment characteristics for a broad cross-section of the housing stock. The Dutch sample comprises many mid-sized apartment buildings (often 3–5 apartments) and reflects the country's historically strong commitment to energy-efficient buildings.	84,329 buildings; 370,284 apartments; median building size = 3 apartments (mean 4.4). Dutch multi-family buildings in this sample are often 2–4 unit buildings or small complexes, with a mix of uniform and varied apartment sizes.

Source: Own elaboration based on the Polish, British and Dutch EPC data.

We focus on multi-family residential buildings where cost-sharing would be applicable. We apply the following sample filters: (1) only apartments/flats (excluding single-family homes). (2) Each building must have at least two apartments with valid EPC records. (3) Apartments with missing or implausible floor area data are excluded. After cleaning, our analytical sample comprises 442,248 buildings and over 4 million apartments across the three countries.

3.1.1. On the use of EPC-modelled energy values[1]

Three of our five allocation rules, i.e., emissions-proportional, inefficiency-proportional, and the Shapley value, depend on energy and emissions values reported in EPC databases.

EPCs do not record measured energy consumption or measured emissions. They report modelled estimates derived from standardised assessment procedures in the three countries. Each methodology takes as an input the physical characteristics of the dwelling, i.e. floor area, envelope U-values, glazing, heating system, ventilation, and applies standardised occupancy patterns, internal temperatures, and use profiles to compute

---

[1] A description of the national EPC systems, including data sources, country-specific indicators, coverage, and harmonisation procedures, is provided in Appendix A1.

an asset-rated annual energy demand and associated CO_2 emissions. The values are properties of the dwelling, not of the household occupying it.

A substantial empirical literature documents systematic discrepancies between EPC-modelled and metered energy consumption, i.e. the "performance gap" or "prebound effect" (Galvin, 2023). Modelled values typically overstate consumption in inefficient dwellings and understate it in efficient ones, because actual occupants of cold dwellings under-heat relative to standardised assumptions, while occupants of efficient dwellings often heat more (rebound).

Two implications follow. First, because EPC values do not measure realised consumption or emissions, we use "performance-based allocation" throughout to describe rules that load costs onto dwellings with worse asset-rated performance, and reserve "polluter-pays principle" for the underlying philosophical concept (5.2). Whether asset-rated performance is even the appropriate basis for residential-sector polluter-pays implementation is a normative question we do not adjudicate. Second, we argue the analysis remains policy-relevant precisely because EPC-modelled values are the values institutional actors would use if performance-based allocation were adopted. Homeowner associations, freeholders, housing associations, and policymakers do not have access to per-dwelling metered data in multi-family buildings, and even where smart-meter coverage exists, the legal and informational frameworks for using such data in cost-sharing decisions are largely absent. EPC values, by contrast, are routinely available, legally recognised, and already used in service-charge apportionment in some jurisdictions. The structural distributional properties we document are therefore properties of feasibly implementable rules, not theoretically ideal ones.

3.1.2. Institutional Context

The three countries represent distinct multi-family housing systems with different ownership structures, governance frameworks, and existing cost-allocation practices. Understanding these institutional contexts is essential for interpreting what our simulated allocations would mean if implemented. These institutional differences have direct implications for the policy relevance of each allocation rule.

The UK's multi-family buildings in our sample are typically small, with a median of only two apartments per building. This reflects the prevalence of converted Victorian and Edwardian houses and period properties subdivided into flats, which comprise approximately 14% of the private rented sector (Poon and Garratt, 2012). These conversions differ fundamentally from purpose-built apartment blocks: they often have unclear apportionment of repair responsibilities, varying lease terms across units, and a building fabric not initially designed for subdivision. English flats operate under a leasehold system, in which approximately 5 million dwellings (around 20% of the housing stock) are held as time-limited interests, of which 69% are flats (Hamnett and Randolph, 2021). This creates an ownership structure distinct from continental European condominium models. Leaseholders pay service charges, but freeholders typically control major works decisions, generating split incentives absent from Dutch or Polish frameworks (Merwe, 2015). The Right-to-Buy programme, which has transferred 2.2 million former council dwellings to private ownership since 1980, has created widespread mixed-tenure buildings where owner-occupiers and social tenants share the same blocks, and must negotiate collective retrofit investments together (Pearce and Vine, 2014).

A complication in interpreting UK results is that tenure recorded in EPCs does not directly map onto ownership structure. Owner-occupied flats are predominantly leaseholds held by occupants, and multi-owner buildings where the allocation question is directly relevant. Private rental properties present a mixed picture: some buildings are entirely owned by a single landlord (making the multi-owner allocation question inapplicable), while others, particularly converted properties, contain multiple buy-to-let investors who must collectively agree on major works. Social housing is typically owned by a single entity (a housing association or a local authority), so the multi-owner negotiation we model does not arise directly. Our analysis remains relevant in two social housing contexts: (a) mixed-tenure buildings where some flats are owner-occupied leaseholds alongside remaining social tenancies; and (b) decisions by landlords about how to allocate retrofit costs across dwellings when setting service charges.

Poland's buildings tend to be larger, with a median of four apartments per building. The Polish housing stock is shaped by its socialist legacy: approximately 60,000 large-panel buildings were constructed between the 1960s and 1989 using standardised prefabrication systems (Muczyński, 2023), featuring highly uniform apartments. Mass privatisation in the 1990s, typically at 80–90% discounts, transformed Poland into a 'homeownership' society with 87% owner-occupation (Lux et al., 2011). Buildings with four or more separately owned units operate as housing communities under the 1994 Ownership of Units Act (Frankowski et al., 2025), with costs allocated based on floor-area share and decisions made by majority vote weighted by ownership shares. This governance structure, although formally similar to Western cooperatives, emerged from circumstances in which tenants became owners overnight, without prior experience in collective property management, a phenomenon comparative housing scholars term 'fragmented ownership without governance capacity' (Lux and Sunega, 2014). A data limitation in our Polish analysis is that the EPC database records building-level floor area rather than individual apartment areas. This precludes the computation of area-proportional and progressive area allocations, as well as the size-progressivity (Kakwani, Suits) indices, for Poland.

The Netherlands occupies an intermediate position, with many 3-unit buildings reflecting the traditional Dutch upstairs-downstairs conversion form. Dutch apartment ownership is managed through the Vereniging van Eigenaren (VvE), a homeowners' association automatically created when a building is divided into separate units. Unlike English leaseholders or Polish cooperative members, Dutch apartment owners benefit from legal requirements introduced in 2018 that mandate either a Multi-Year Maintenance Plan or minimum annual reserve contributions of 0.5% of rebuilding value (Kendall, 2021), the strongest mandatory reserve framework among our three countries. Cost allocation follows a fractional-share system, typically calculated by floor area and recorded in the original deed of division. The Netherlands also has Europe's largest social housing sector proportionally (29% of stock), managed by non-profit housing corporations; when these associations sell units, mixed collaborations result (Hoekstra, 2017).

3.2. Cost Allocation Methods[2]

We evaluate five rules for allocating a fixed building-level retrofit cost across apartments. For each building, the allocation rule produces a vector of non-negative cost shares that sums to one. Apartment-level monetary burdens are obtained by multiplying these shares by the total retrofit cost. Because EPC databases do not capture occupancy, household income, or wealth, all rules operate only on dwelling-level attributes. The resulting allocations should therefore be interpreted as structural cost-sharing rules rather than household-level welfare measures.

The first two rules are area-based. The standard area-proportional rule allocates costs based on each apartment's share of the total floor area. The progressive area rule applies a convex transformation to floor area, assigning larger dwellings a share of their proportional area and smaller dwellings a share that is less.

---

[2] Formal definitions of the five cost-allocation rules, including notation, implementation choices, and robustness specifications, are provided in Appendix A2.

We include this rule as a stylised alternative that strengthens the capacity-to-pay logic often associated with floor-area allocation, while recognising that dwelling size is an imperfect proxy for household resources.

The remaining rules are performance-based. The emissions-proportional rule allocates costs according to each dwelling's EPC-modelled emissions intensity, while the inefficiency-proportional rule allocates costs according to modelled improvement potential relative to a target performance level. These rules capture different versions of a responsibility, or performance-based allocation logic, but they rely on asset-rated EPC values rather than realised household behaviour.

Finally, we evaluate the Shapley value allocation from cooperative game theory. In the baseline specification, where the characteristic function is additive in modelled improvement potential, the Shapley allocation coincides with the inefficiency-proportional rule. We nevertheless report it separately because the Shapley rule carries a distinct theoretical interpretation and motivates the cooperative-stability tests developed later, including individual rationality and core-related criteria. Robustness checks consider alternative characteristic functions based on emissions and floor area.

Throughout the analysis, the total retrofit cost of each building is held fixed. The purpose is not to model retrofit take-up, behavioural responses, or alternative investment packages, but to compare how different allocation principles distribute the same collective cost among dwellings in the same building.

3.3. Distributional Metrics[3]

To characterise the distributional properties of the cost-share vectors these rules produce, we employ a set of inequality and redistribution metrics, drawing from economic theory on income distribution, tax progressivity, and cooperative game theory. A key aspect of our methodology is that we compute these metrics within each building and then aggregate or summarise across buildings. This recognises that collective retrofit decisions are made at the building level.

We evaluate each allocation rule along three dimensions: within-building inequality, size-progressivity, and cooperative stability. First, for every building, we treat the vector of dwelling-level cost shares as a distribution of contributions and compute standard inequality measures, including the Gini coefficient, Atkinson index, Theil indices, coefficient of variation, and maximum-to-minimum ratios. These measures characterise the dispersion of cost burdens generated by each rule within the relevant decision-making unit: the building. They are not interpreted as measures of household income or welfare inequality. Building-level results are aggregated using means, medians, threshold shares, and bootstrapped confidence intervals, treating each building as a single collective decision context.

Second, we assess whether cost shares are proportional, superproportional, or subproportional across dwelling sizes. Using floor area as a structural proxy for capacity, we compute size-progressivity measures, primarily the Kakwani index, with the Suits index used as a robustness check. These indices indicate whether larger dwellings pay more or less than their proportionate share of total floor area. Because Polish EPC data do not include apartment-level floor area, this analysis is restricted to the UK and the Netherlands. We interpret these results as measures of size-progressivity, not welfare progressivity, since we do not observe household income, wealth, or liquidity constraints.

---

[3] Definitions, formulas, interpretation, and implementation details for the inequality, size-progressivity, and cooperative-stability indicators are provided in Appendix A3.

Third, we test whether allocation rules satisfy cooperative-stability criteria. In particular, we examine whether any dwelling is allocated a cost share exceeding a stand-alone benchmark, proxied by its area-proportional share. This provides a necessary individual-rationality condition for voluntary collective investment: if a dwelling is charged more than its stand-alone reference cost, it has a structural incentive to reject the joint allocation. We report both a strict version of this test and a 1.5-times-tolerant version, along with the maximum excess ratio. These tests do not rely on interpreting EPC-modelled values as measures of responsibility or ability to pay; they assess whether the allocation rule itself generates cost burdens that may undermine cooperation within the building.

3.4. Empirical Strategy

Our analysis proceeds in four stages. First, for each building and each method, we generate the cost-share vector $(\pi_1, \dots, \pi_n)$. For the area-based rules, this uses floor area; for emissions-proportional, harmonised EPC-modelled emissions intensity (kg CO_2/m^2/year); for inefficiency-proportional and Shapley, modelled improvement potential computed against the country-specific target documented in 3.1.1. and Appendix A1. Where characteristic data are missing for an apartment, we default to equal shares $(1/n)$ for that building.

Second, for each cost-share vector, we compute within-building inequality indices (Gini, Atkinson at $\varepsilon = 0.5, 1, 2$, Theil T, Theil L, CV, coefficient of variation, max/min ratio) and, where applicable, the size-progressivity indices (Kakwani, Suits) using floor area as the size attribute. We then compute cooperative-stability metrics (individual-rationality violations under both the strict and 1.5×-tolerant definitions, and the maximum excess ratio).

Third, we summarise across buildings. We report the mean and median for each metric, with nonparametric bootstrap confidence intervals (1,000 replications, building-level resampling). We report the percentage of buildings falling below or above policy-relevant thresholds (e.g., Gini < 0.1 as "low dispersion"). Each building is treated as a single observation, since it constitutes its own decision context. We also compare rules pairwise, asking how often one rule produces higher within-building dispersion than another in the same building. We replicate all analyses separately by country, and for the UK, we additionally stratify by tenure type and construction cohort.

Fourth, we conduct a battery of robustness checks (Appendix B). These include: alternative Shapley characteristic functions (current emissions, floor area); three alternative specifications of the post-retrofit reference X^{target} (within-country median energy intensity, fixed 50 kWh/m^2 benchmark, within-building minimum); total versus per-m^2 normalised emissions; varying the progressive area exponent α between 1.0 and 2.5; trimming buildings with extreme heterogeneity; bootstrap confidence intervals on key statistics; building-size, tenure, and construction-cohort stratifications; and random sample splits. Headline conclusions hold under all alternatives.

4. Results

4.1. Within-Building Inequality Across Allocation Methods

Our core findings are summarised in the within-building Gini coefficients for each method and country (Figure 1). This provides a first measure of how unequal the cost burdens are for neighbours under different allocation schemes.

Figure 1. Mean within-building Gini coefficient by allocation method and country.

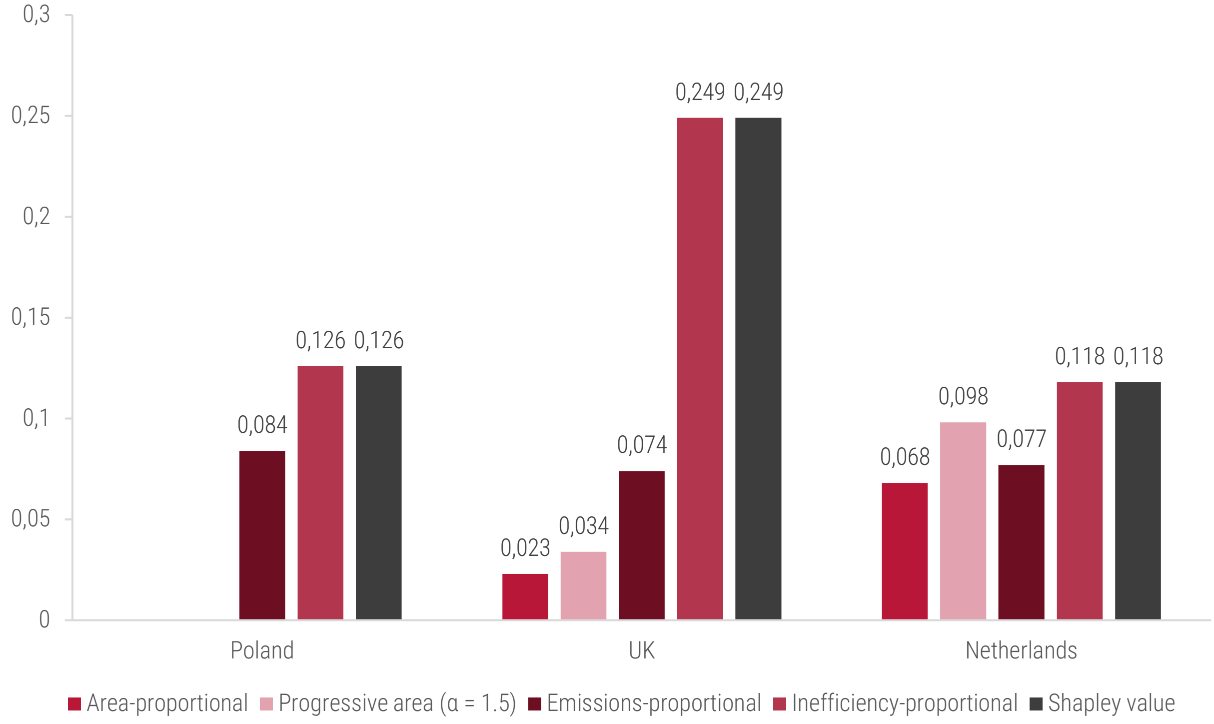

Source: Own elaboration based on the Polish, British and Dutch EPC data. Each cell is the mean of the within-building Gini coefficient across all multi-apartment buildings in that country, computed under the harmonised specifications described in 3.1.1 (per-m² emissions intensity; per-dwelling SAP target for UK; fixed 50 kWh/m²/year NZEB target for NL and PL). Inefficiency-proportional and Shapley coincide by construction under our additive characteristic function (3.2). Higher values indicate more dispersed cost shares within buildings; bootstrap 95% confidence intervals on the mean are tight (±0.001) and reported in Appendix B.

Area-proportional allocation produces low within-building dispersion. The mean Gini is 0.023 in the UK and 0.068 in the Netherlands, with 96% of UK buildings and 75% of Dutch buildings falling below the 0.1 dispersion threshold (Figure 2).

Progressive area allocation (with $\alpha = 1.5$) produces modestly higher dispersion than simple area, with a mean Gini of 0.034 in the UK and 0.098 in the Netherlands. Sensitivity to the convexity exponent α (tested from 1 to 2.5) is discussed in 4.5; even at $\alpha = 2.5$, the mean Gini remains well below the levels produced by performance-based rules. The choice of allocation principle (size-based versus performance-based) is far more consequential for dispersion than tuning the convexity within the size-based family.

Emissions-proportional allocation, harmonised to per-m² intensity, produces moderate within-building dispersion. The mean Gini is 0.074 in the UK, 0.077 in the Netherlands, and 0.084 in Poland. Once size-confounding is removed by per-m² normalisation, within-building variation in modelled emissions intensity is structurally similar across the three building stocks. The implications of moving from per-m² to total emissions are explored in 4.5, although the headline conclusions are unchanged.

Inefficiency-proportional and Shapley value allocations produce the highest within-building dispersion among the rules we examine. The mean within-building Gini is 0.249 in the UK, 0.118 in the Netherlands, and 0.126 in Poland. Inefficiency and Shapley coincide exactly by construction, since under our additive characteristic function, the Shapley value of an apartment is its modelled improvement potential as a fraction of the building's total.

The cross-country pattern is informative. The UK Gini is more than ten times the area-proportional value, reflecting the heterogeneity of the UK building stock, which comprises more than 1 million two-flat conversions, where the two units typically differ substantially in modelled energy intensity due to floor position, orientation, and individual heating-system upgrades (Madeddu and Clifford, 2023). Polish and Dutch values are roughly half the UK level, reflecting the more uniform construction of post-war Polish prefabricated blocks and Dutch purpose-built apartment buildings (Tsenkova, 2017). Even at this lower level, however, the dispersion under performance-based rules is several times higher than under area-based rules in the same building stock.

Across all three countries, the ratio of performance-based to area-based dispersion is large. UK Shapley dispersion is 10.7 times the area-proportional dispersion; the Netherlands ratio is 1.7. The mechanism is the same in both, as performance-based rules amplify modest differences in modelled energy intensity into substantial differences in cost shares, while size-based rules track an attribute that is more uniformly distributed within buildings. It is also worth noting the consistency of the ranking across countries. Where comparable rules are computed, we observe the ordering (from lowest to highest dispersion):

$$\text{Area} < \text{Progressive} < \text{Emissions} < \text{Inefficiency} \approx \text{Shapley}.$$

This holds across building stocks of substantially different ages, ownership structures, and architectural types. The structural distributional properties of these rules are robust and not artefacts of a single national context.

Figure 2. Percentage of Buildings with Low Inequality (Gini < 0.1)

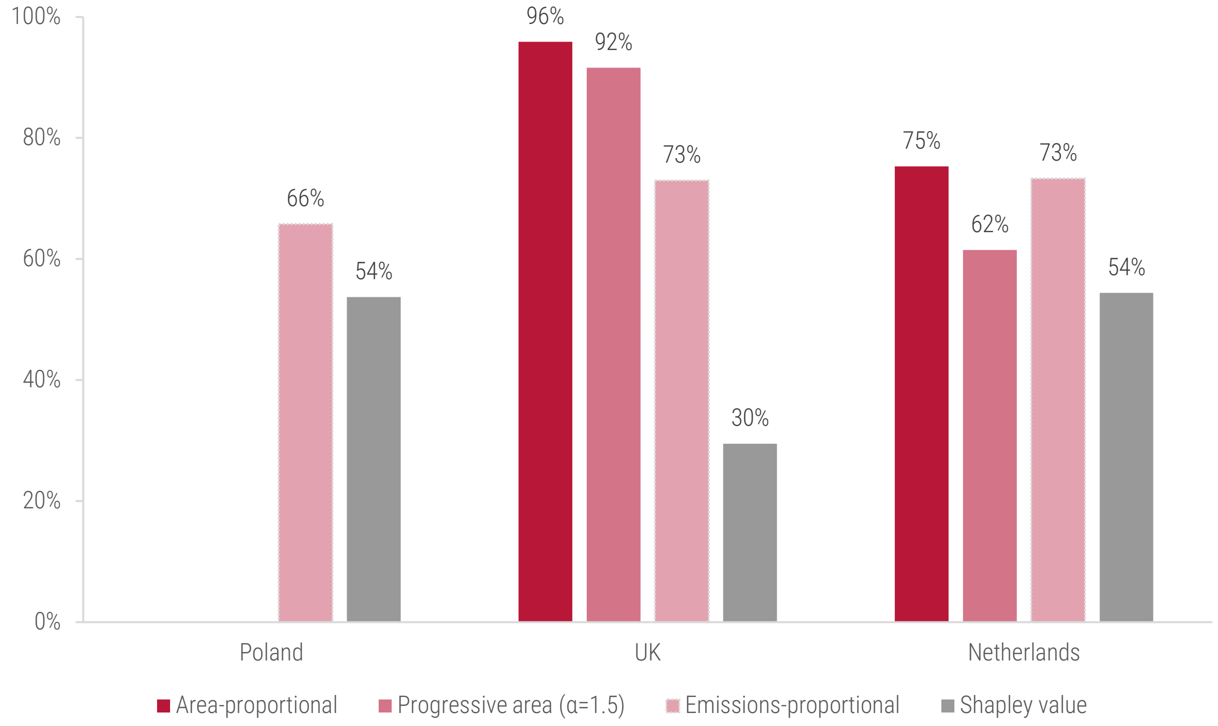

Source: Own elaboration based on the Polish, British and Dutch EPC data.

Performance-based rules systematically reduce the share of buildings in the low-dispersion range. Under Shapley, the share of UK buildings with Gini < 0.1 falls from 96% (area) to 30%, a 66-percentage-point drop. The Dutch and Polish drops are smaller in absolute terms (75% → 54%) but reach similar relative magnitudes.

In other words, switching from a size-based to a performance-based rule moves the modal building from “low dispersion” to "moderate dispersion or higher" in all three countries.

Median max/min ratios within buildings (Appendix B.3) tell a similar story. Under area allocation, the median UK building has a max/min cost ratio of 1.04, meaning the highest-paying flat contributes only 4% more than the lowest-paying flat. Under Shapley, the median UK ratio is 1.93. The Polish and Dutch Shapley ratios are 1.40 and 1.53, respectively. These ratios reflect the harmonised specifications described in 2.1.1; the ranking across rules remains unchanged.

4.2. Extended Inequality Metrics: Welfare and Tail Sensitivity

We computed Atkinson indices at multiple levels of inequality aversion and Theil indices to characterise the shape of the within-building cost-share distribution beyond the Gini summary statistic. Full per-country results are reported in Tables B.4 and B.5 in Appendix B.

The Atkinson indices preserve the ordering established by the Gini and amplify the contrast at higher aversion parameters. In the UK, the mean Atkinson index under Shapley rises from 0.053 at $\varepsilon = 0.5$ to 0.174 at $\varepsilon = 2$, while under area-proportional it rises only from 0.002 to 0.008. The Atkinson family is informative as a structural measure of how concentrated the cost-share distribution is when inequality is more or less heavily penalised.

The Theil decomposition is more informative. For size-based rules (area, progressive area, emissions), we observe Theil L > Theil T in all three countries, indicating that within-building dispersion under these rules is concentrated at the lower end of the cost-share distribution. Under the inefficiency-proportional and Shapley measures, however, the relationship reverses: Theil L < T (in the UK, 0.166 vs 0.241; in the Netherlands, 0.041 vs 0.066; in Poland, 0.068 vs 0.077). This indicates top-end concentration. Under performance-based rules, the within-building cost-share distribution is shaped by a small number of apartments paying very large shares.

This is a substantive finding about the type of dispersion that performance-based rules generate. Size-based rules produce many small differences across all apartments in a building. Performance-based rules produce a few large differences; we hypothesise that they are typically loaded onto the apartment(s) with the worst modelled energy performance. This top-heavy concentration is precisely the structural feature that the cooperative-stability tests in 4.4 identify as creating individual-rationality violations.

4.3. Size-Progressivity

We now turn to the Kakwani and Suits indices, which characterise each allocation rule’s distributional position relative to dwelling floor area. We measure size-progressivity rather than welfare progressivity: a negative index indicates that the rule is sub-proportional in dwelling size, that is, smaller dwellings receive larger cost shares relative to their share of total floor area. This is a structural property of the rule given the building stock. Polish EPCs lack apartment-level area data, so Kakwani and Suits are computed only for the UK and the Netherlands.

All five rules are sub-proportional in dwelling size in both countries, but the magnitudes differ. UK Kakwani values cluster tightly between -0.17 and -0.19; Dutch values are around -0.67 to -0.75 (Table B.6). The difference reflects the underlying distribution of floor area: Dutch buildings exhibit greater within-building variation in apartment size than UK buildings, which makes any given allocation farther from size-proportionality in the Netherlands.

Within each country, the Kakwani values are remarkably stable across rules. UK values range from -0.172 (progressive area) to -0.193 (emissions), a span of just 0.02. Even strong convexity in the size-based rule (α = 1.5) is insufficient to make the rule super-proportional in size, because the underlying within-building floor-area distribution is too compressed for any tractable convexity parameter to dominate. Performance-based rules are no more or less sub-proportional in size than size-based ones, indicating that modelled energy intensity is, on average, weakly correlated with dwelling size within buildings.

This last point is relevant to the polluter-pays interpretation. If small dwellings were systematically more efficient than large ones, performance-based rules would be more size-proportional than area-based rules. They are not. Apartments with the worst modelled performance can be of any size, and this structural feature is what gives performance-based rules their distinctive shape: they target individuals on the basis of dwelling characteristics that are uncorrelated with size.

4.4. Cooperative Stability

The cooperative-stability tests are the analysis least dependent on contested interpretations of EPC-modelled values. The individual-rationality and core conditions are structural properties of the allocation rule applied to the building stock, requiring no assumptions about occupant preferences, capacity, or responsibility. We report results under both the strict and the 1.5×-tolerant test, with bootstrap 95% confidence intervals based on building-level resampling.

Table 2. Percentage of buildings satisfying individual rationality, by allocation rule and country.

	UK		NL		PL	
Method	Strict	1.5× tolerant	Strict	1.5× tolerant	Strict	1.5× tolerant
Progressive area	13.3%	100%	7.4%	99.8%	N/A	
Equal	13.3%	98.7%	7.4%	90.1%	99.7%	99.8%
Emissions	2.3%	93.2%	0.3%	75.9%	29.2%	85.9%
Inefficiency / Shapley	3.6%	52.5%	0.6%	67.5%	27.6%	74.7%

Source: Own elaboration. Bootstrap 95% confidence intervals are tight (UK: ± 0.04 percentage points; NL: ± 0.20 pp; PL: ± 0.25 pp; full intervals reported in Appendix B.7). The strict test treats the area-proportional share as the additive stand-alone-cost benchmark; the 1.5×-tolerant test allows up to 50% overpayment relative to that benchmark.

Two patterns dominate (Table 2). First, performance-based rules (emissions, inefficiency, Shapley) fail the strict test almost universally. Inefficiency / Shapley satisfies strict cooperative stability in 3.6% of UK buildings, 0.6% of Dutch buildings, and 27.6% of Polish buildings. The corresponding tolerant figures are higher, but the gap between strict and tolerant is itself revealing. In approximately half of UK buildings, the Shapley rule allocates to at least one apartment a share more than 50% larger than that apartment's stand-alone retrofit cost. In nearly all UK buildings, the Shapley rule allocates to some apartment more than its stand-alone cost outright.

Second, Polish buildings are notably more stable under performance-based rules than UK or Dutch buildings. The Polish prefabricated building stock, characterised by more standardised apartment layouts and more uniform construction, produces relatively homogeneous modelled improvement potentials within buildings. The Shapley value can therefore approximate area-proportional sharing in Polish buildings, whereas in heterogeneous UK conversion buildings it produces shares that diverge sharply from any size-proportional benchmark. This is consistent with the finding that UK Shapley dispersion (Gini 0.249) is roughly twice the Polish or Dutch values (0.126, 0.118).

The distance between the strict and tolerant figures matters for interpretation. Under the strict test, a building fails if any apartment's share strictly exceeds its area-proportional benchmark, even by a fraction of a percentage point. Under the 1.5×-tolerant test, a building fails only if some apartment is allocated more than 1.5 times its area-proportional benchmark. Both definitions point in the same direction: performance-based rules systematically allocate more to some apartments than the additive, stand-alone-cost benchmark. Whether this translates into actual refusal depends on residents' preferences, information, and outside options, which we do not observe. The structural property is informative regardless.

4.5. Robustness

Headline conclusions hold across a battery of robustness checks (Appendix B). We highlight three findings here that bear directly on the methodological concerns most relevant to a careful reader.

The choice of post-retrofit reference X^{target} for the inefficiency rule is consequential and varies depending on what is recorded in each EPC system. Figure 3 plots the Shapley/inefficiency Gini under four alternative target specifications: per-dwelling regulatory targets, country-fixed national targets, country-median apartment-level energy intensity, and within-building minimum. The four specifications produce different magnitudes: using country-median targets produces larger Ginis than fixed targets in the Netherlands and Poland, while within-building-minimum targets yield the largest Ginis everywhere, and the qualitative ranking (performance-based dispersion > area-proportional dispersion) is preserved under all specifications. The Dutch Eis specification produces near-zero Gini because the NTA 8800 regulatory threshold is fixed by building type and construction year and is therefore approximately constant within most buildings; we discuss the implications for regulatory targeting in 5.2.

Absolute-emissions allocation may conflate dwelling size with intensity. Appendix B.2 confirms this concern is small in magnitude under the harmonised specifications applied. The UK emissions Gini falls from 0.075 (total CO_2) to 0.074 (per-m^2 intensity); Dutch and Polish numbers are unchanged because their EPC fields are already in per-m^2 units. The within-building dispersion under per-m^2 normalisation is bounded above 0.001 in our data; size confounding within UK buildings is therefore quantitatively small once the harmonisation is applied.

Cross-tabulating the within-building Gini by building size (Appendix B.8), UK tenure type (B.9), and construction period (B.10) does not change the cross-country ranking of rules. Two patterns are worth noting. First, UK social rental buildings exhibit higher within-building Shapley dispersion (mean Gini 0.262) than UK owner-occupied buildings (0.218). This is consistent with greater heterogeneity in within-building energy performance in social housing, possibly reflecting uneven historical maintenance, and has equity implications that we develop in 4. Second, larger UK buildings (≥11 apartments) exhibit substantially lower within-building Shapley dispersion (mean Gini 0.048) than 2-flat buildings (0.236), because performance heterogeneity averages out across many units.

Figure 3. Sensitivity of performance-based dispersion to target specification

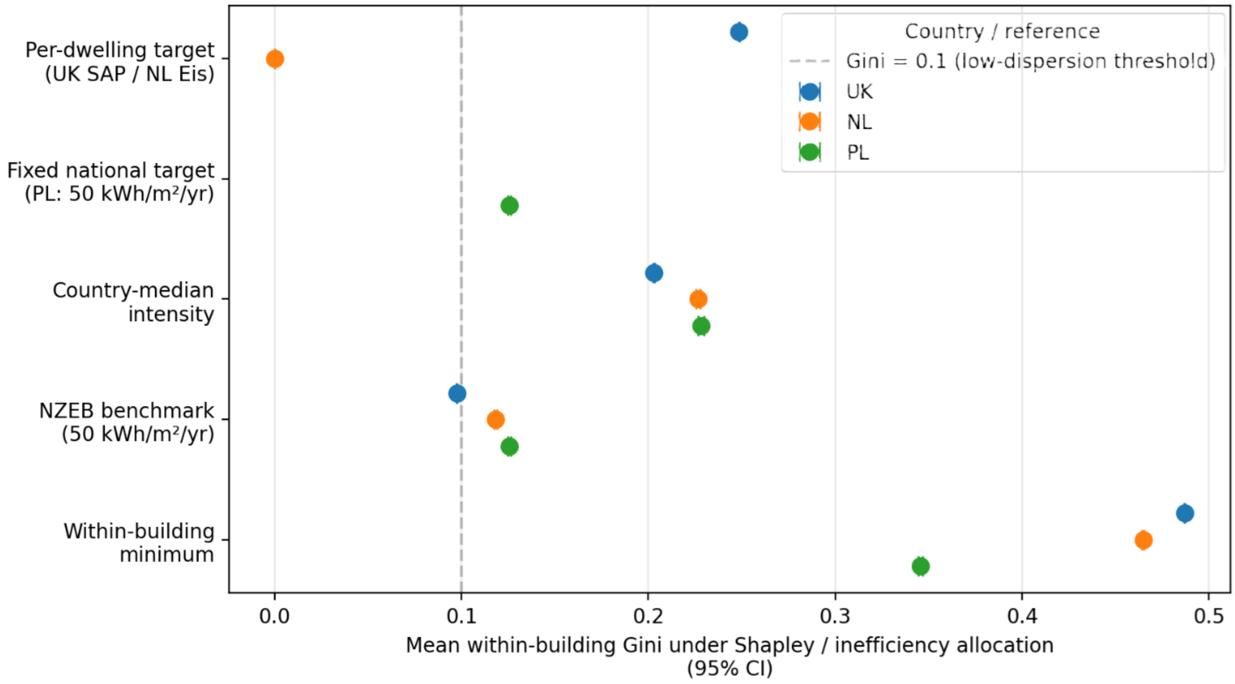

Note: Each row reports a different definition of the post-retrofit reference level, X^{target}*. Markers show the mean within-building Gini coefficient produced by the Shapley/inefficiency rule under each target definition; horizontal bars indicate 95% bootstrap confidence intervals. The dashed vertical line marks Gini=0.1, used as the threshold between low and moderate within-building dispersion. Target–country combinations are omitted where the relevant target cannot be computed, for example, where no per-dwelling target is available. The within-building-minimum target should be interpreted as a hard upper-bound specification, since at least one apartment in each building is assigned a zero share by construction.*

5. Discussion

This section interprets the findings and develops their implications. We first read the results across the three institutional contexts, reassess the normative status of the polluter-pays principle at the within-building scale, and draw general lessons for fair and effective community investment. We then develop three theoretical implications that bear on adjacent literatures, three recommendations that translate the structural findings into retrofit programme design, and an explicit statement of the scope and limitations of the analysis.

5.1. Country-Specific Implications for Cost Allocation Practice

Our results have country-specific implications that depend on the structure of the building stock, the legal framework governing cost allocation, and the size of the residual social-housing sector.

For Poland, our results indicate that performance-based allocation would be more cooperatively stable than in the other two countries, but would still introduce substantial within-building dispersion compared with the equal-shares benchmark. The Polish building stock is unusual in that within-building energy heterogeneity is moderate rather than extreme, reflecting standardised construction (Attia et al., 2022). This makes performance-based rules less destabilising in absolute terms. The implication is that Poland's existing practice of building-level area-proportional allocation is well-suited to its building stock and institutional capacity, and that introducing performance-based elements would deliver smaller gains than in the UK while introducing larger procedural risks.

The UK case is the most distributionally stark of the three. The combination of small buildings, high within-building energy heterogeneity, and the leasehold ownership structure produces the highest within-building Shapley dispersion in our sample and the lowest strict cooperative-stability rates. In leasehold buildings, the freeholder controls major works decisions and recovers costs through service charges apportioned in accordance with the lease terms. Despite being charged for the service, leaseholders have limited influence over either the decision to retrofit or the allocation of costs, and cannot block works the freeholder deems necessary (Crosby et al., 2003). Our results reveal the distributional consequences of different apportionment approaches the freeholder might adopt. A Shapley apportionment in a two-flat conversion can produce cost-share ratios of approximately 1.93:1 (median) or higher; in 96% of UK buildings, the strict individual-rationality condition is violated under Shapley. This creates a potential for disputes, legal challenges, and resistance to paying service charges, which can delay retrofits even when the freeholder has the authority to proceed. Lease terms, leaseholder access to legal advice, and the freeholder's appetite for litigation could also play a role here (Cole and Robinson, 2000). In mixed-tenure buildings created by Right-to-Buy, where some units remain social housing, and others are privately owned leaseholds, performance-based apportionment could prove especially contentious if it concentrates costs on lower-income owner-occupiers in less efficient units.[4]

Social housing presents a more fundamental disconnect. Council tenants do not pay retrofit costs directly, as local authorities fund improvements from capital budgets and housing revenue accounts, with tenants paying only rent (Bright et al., 2019). Housing association tenants may face some indirect pass-through via service charges or rent adjustments, but this operates quite differently from leasehold cost recovery. Our finding that UK social rental buildings exhibit higher within-building dispersion under Shapley than owner-occupied buildings shows how costs would be distributed across dwellings if allocated according to modelled inefficiency. Whether this translates into tenant burden depends on the social landlord’s pass-through mechanism, a question outside the scope of our dwelling-level analysis but central to the policy implications.

The social housing results remain relevant in two narrower contexts. First, in mixed-tenure buildings, where some flats are owner-occupied leaseholds alongside remaining social tenancies, our allocation analysis applies to the leaseholder portion, and performance-based apportionment could concentrate costs on owner-occupiers in less efficient units. Second, if social landlords sought to recover costs through differentiated service charges (allocating more to less efficient units), our results indicate that this would burden tenants in poorly positioned flats, raising concerns about affordability and equity that would likely face resistance (Sovacool, 2015). For these reasons, size-based rules based on floor area or equal shares would better protect residents in contexts where costs are passed through to occupants.

The Netherlands sits at an intermediate institutional position. Our results indicate Dutch within-building dispersion under Shapley is moderate, but the strict cooperative-stability rate is the lowest of the three countries. The combination of moderate dispersion, very low strict stability, and strong institutional absorption capacity is mechanically explained by within-building uniformity in apartment characteristics, as most Dutch apartments in a building are similar enough that Shapley produces only moderate variation, but never exactly matches area-proportional, so any positive variation registers as a strict-test violation. The 1.5×-tolerant rate is much higher, suggesting the dispersion under Shapley is real but bounded. Whether Dutch

---

[4] A methodological caveat is necessary regarding our analysis of UK tenure. While we report results stratified by tenure type recorded in EPCs, the relevance of cost allocation methods varies significantly across categories. For owner-occupied leaseholds, our analysis directly applies. These owners face service charges apportioned according to lease terms and would bear the distributional consequences of different allocation principles. For private rental properties, landlords absorb retrofit costs (potentially recovering them through rent adjustments over time), making the dwelling-level allocation we model one step removed from the direct impacts on tenants.

VvEs would tolerate the implied cost-share differences depends on the existing share-of-cost arrangements recorded in each building's deed of division, which would have to be unanimously amended to accommodate any rule other than the recorded fractional-share system (Hoppe and Lulofs, 2008; Ploeger and Groetelaers, 2014).[5]

A common pattern emerges across these three institutional contexts. Existing legal frameworks and housing stock characteristics strongly favour size-based cost-sharing for retrofits. Performance-based methods, while appealing for aligning cost with the "cause" of energy waste, systematically allocate cost shares that violate cooperative-stability axioms and that would be expected to face disproportionate resistance from individual co-owners. This applies most strongly in the UK and least strongly in Poland, but holds in all three countries.

5.2. The Polluter-Pays Principle in the Within-Building Setting

We dwell on the polluter-pays principle because it is increasingly invoked, explicitly in EU energy-justice policy framings (Kaschny, 2023; van der Kooij, 2025), and implicitly in the energy-community literature's enthusiasm for performance-based cost-sharing in shared retrofit investments (Cremers et al., 2022), as a normative justification for performance-based cost allocation in the residential sector. We argue that the principle's normative force depends on conditions that hold for industrial pollution but fail systematically within multi-family residential buildings. The polluter-pays principle was developed for industrial pollution, where the polluter is identifiable, the harm is attributable to specific actions, and the actor has the capacity to alter behaviour. None of these conditions holds straightforwardly within building energy retrofit cost allocation.

First, what is being measured is not pollution. EPC-modelled energy and emissions values are properties of the dwelling, computed under standardised occupancy and use assumptions. Charging an apartment based on its asset-rated emissions therefore charges it based on a property of the dwelling rather than the environmental harm caused. Second, dwelling-level performance is largely structurally inherited. Envelope quality, position within the building, orientation, and construction materials, the principal determinants of asset-rated performance, are properties of the building or of the apartment-as-located, not actions taken by current occupants. A leaseholder who buys a north-facing top-floor flat in a Victorian conversion inherits its modelled energy profile. The performance-based allocation that follows charges the leaseholder for architectural decisions made by previous owners, builders, and architects, possibly decades ago.

Third, the "payer" cannot unilaterally alter the structural attributes that drive performance differences. Even an occupant motivated to reduce their building's modelled energy intensity faces strong constraints. Most fabric improvements require freeholder consent (UK leasehold), unanimous deed amendment (Dutch VvE), or majority approval (Polish housing community), and none of these can be unilaterally undertaken by a single apartment owner. The polluter-pays principle would charge individuals without the capacity to change what they are being charged for.

---

[5] An additional Dutch-specific finding deserves discussion. The Dutch EPC system records the per-dwelling NTA 8800 regulatory primary-energy requirement, which is the legally-prescribed maximum for the dwelling type and construction year. Using this regulatory threshold as the post-retrofit reference target for the inefficiency rule produces near-zero within-building dispersion. Because the threshold is fixed by building type and construction year, both of which are constant within a building, apartments in the same building face essentially the same nominal reference, and the resulting within-building improvement-potential vector is approximately uniform. This is a substantive finding about regulatory targeting rather than a methodological artefact. The NTA 8800 framework was not designed to discriminate among apartments in the same building, and using it for within-building cost allocation results in nearly equal shares. The implication for Dutch retrofit policy is that performance-based allocation cannot be implemented through reference to existing regulatory thresholds; it would require an externally constructed, non-regulatory benchmark.

The principle, therefore, loses much of its normative force when applied to the within-building setting. Our use of "performance-based allocation" to describe rules that impose costs on dwellings with worse asset-rated performance reflects this distinction: such rules invoke the polluter-pays logic but do not deliver the principle's normative content.

The implication for policy is that performance-based logic may be more appropriately applied at the building level than at the within-building level. A grant programme that directs more public subsidy to lower-performing buildings (and less to already-efficient ones) implements polluter-pays logic against an actor, the building or its collective owners, that does have the capacity to undertake the corrective action being financed. A rule that imposes additional costs on lower-performing apartments within a building applies the same logic to an actor whose capacity to act is structurally constrained. The two scales are different policy levers despite invoking the same principle.

5.3. General Implications for Fair and Effective Community Investment

Fairness in practice encompasses perceptions of both the outcome and the process (Bal et al., 2023). A method might achieve an economically "fair" outcome by one criterion. Still, if it violates another principle, such as the notion of horizontal equity, or if it feels punitive, participants may resist (Heffron and Sokołowski, 2024). For example, the Shapley allocation is grounded in a rigorous fairness concept from cooperative game theory, in which each player pays according to their marginal contribution to the coalition's costs. Yet, this framework does not transfer perfectly to residential retrofits. As highlighted in section 5.2 above, the dwelling-level "contribution" measured by EPC values reflects building characteristics beyond occupant control. Applying the cooperative-game-theoretic apparatus to such characteristics imputes responsibility to actors who structurally cannot change what they are being charged for. Thus, decision-makers should consider multiple fairness criteria: Is the allocation equitable? Does it respect that similar people are treated similarly? Is anyone paying so much that they effectively subsidise others? Our approach can serve as a template for others to follow. Practically, before finalising a cost allocation scheme, communities or planners might simulate its impact to check for extreme cases.

The simplest schemes (e.g., area-proportional) were not only among the least dispersive in terms of outcomes, but are also easy for people to understand and, crucially, legally implementable within existing frameworks. In the UK, lease terms specify apportionment methods (typically floor area, rateable value, or fixed percentages) that cannot be unilaterally changed. Implementing the Shapley allocation would require renegotiating every lease in a building, which is practically impossible. In the Netherlands, shares are recorded in the original deed of division and require the unanimous consent of all owners to be modified. Polish law mandates floor-area-based allocation. Shapley or emissions-based methods are thus not merely complex. They are largely unimplementable within current legal structures without wholesale legislative reform. Even if such reform occurred, complexity breeds mistrust. If residents cannot easily verify how their share was computed, they may suspect bias or manipulation. Area-based allocation, by contrast, is transparent (as anyone can measure floor space), stable (it does not change with building conditions), and aligns with established legal practice across all three jurisdictions.

Community investments often require trust that everyone is paying their fair share; a convoluted formula could undermine that trust even if well-intentioned (Goedkoop and Devine-Wright, 2016). In our context, floor area is a readily observable and verifiable basis for sharing costs. It is observable (any party can independently measure or verify it), legally entrenched (it is the recorded fractional-share basis in two of three jurisdictions), and stable across time. These institutional properties are what make it a workable focal point for agreement. Performance-based, by contrast, requires accepting that an abstract measure (CO_2 or inefficiency points) serves as the basis; some might question the accuracy of those EPC measures or argue that they should not pay for past neglect by a previous owner who left their unit inefficient.

As countries push to decarbonise buildings, governments and industry bodies may issue guidelines on how retrofit costs should be divided. Our results suggest such guidelines should err on the side of transparency and stability rather than pure environmental accounting. A reasonable default for publicly subsidised retrofits would be size-based apportionment, with deviations possible by unanimous agreement of co-owners. This preserves the option for buildings to adopt performance-based allocation where all parties accept the trade-offs, while ensuring that the default does not inherit the structural problems documented in 4.4.

One insight is that extremes are problematic. Equal shares ignore differences (which violates notions of fairness tied to size), whereas Shapley exaggerates them (which violates notions of solidarity). A middle ground could be hybrid approaches. A compromise could reduce dispersion compared to a 100% performance-based approach, while still incorporating the marginal-contribution intuition that motivates performance-based rules in the first place. Or they might put caps, e.g., no one pays more than 50% of what anyone else pays, regardless of the metrics. This cap would correspond to the 1.5×-tolerant cooperative-stability test in 4.4 and would substantially increase the share of buildings with stable allocations. Another idea would be to apply the polluter-pays principle to a portion of costs specifically targeted at unit-level improvements (such as additional insulation around a specific flat), while using equal/area for the truly collective parts (roof, foundation, etc.). Future research and practical experimentation can explore these hybrid models to balance multiple distributional goals.

If a particular building's agreed allocation is still felt to be burdensome for some members, external policy measures can help. For example, means-tested grants or loans to low-income apartment owners can ensure their share of the retrofit cost does not force them out. This separation of concerns, by using a transparent, size-based rule within the building, and addressing affordability concerns through external means-tested support, is more legally tractable, politically defensible, and protective of low-income occupants than embedding affordability proxies into the cost-allocation formula itself.

The polluter-pays principle is widely endorsed in environmental economics, as seen in the use of carbon taxes and pollution permits (Ambec & Ehlers, 2016). However, in a small community such as a building, our study shows that performance-based allocation can have contrary social outcomes. The structural conditions that make polluter-pays defensible at the industrial scale are absent within residential buildings, as developed in 5.2. Implementation at the building or programme level, where these conditions partially apply, may be more appropriate than implementation at the within-building level.

The cooperative game analysis highlights that an allocation that leaves any subgroup feeling they are better off opting out is a risky strategy. Our strict cooperative-stability tests show that performance-based allocation rules violate this property in the overwhelming majority of buildings in all three countries. Thus, strategies to maintain cooperation, either by choosing inherently stable allocations (such as the area method) or by adjusting unstable ones (with compensation payments, negotiation, etc.), are crucial. Communities might consider formal mechanisms, such as side agreements, in which those paying more receive extra credit or future consideration, though this can become complex.

While our focus is on energy retrofits, similar issues arise in any shared investment, such as installing an elevator in a condo, rooftop solar, or flood defences in a neighbourhood. The tools we used can be applied broadly. One can simulate different cost splits and measure how unequal or sub-proportional they are, as well as whether any participant would prefer a subset solution. Our findings suggest that if the benefit or risk is shared by all, equal or proportionate contributions are perceived as fair. If the benefit is uneven, one might try to adjust, but over-adjusting can cause new distributional issues. The Shapley value is often advocated for infrastructure cost sharing (e.g., among municipalities), but our results caution that if Shapley produces outcomes in which one party pays a lion's share, it may not be politically or socially feasible without compensation.

5.4. Theoretical Implications

First, cooperative game-theoretic axioms do not deliver fairness when the characteristic function is exogenous. The Shapley value satisfies efficiency, symmetry, the null-player condition, and additivity. In its original formulation (Shapley, 1953), these axioms apply to settings in which the characteristic function, i.e., the value generated by each coalition, reflects players' chosen actions or capabilities. The Shapley value's normative appeal in those settings rests on charging each player according to their marginal contribution to the outcomes they helped produce. Our results show that the cooperative-stability conditions are systematically violated when the framework is applied at scale to real residential building stocks.

The implication is that cooperative game-theoretic axiomatic appeal should not be confused with cooperative stability in application. The Shapley rule's theoretical attractiveness does not survive empirical implementation at scale when characteristic-function values are exogenous to the players whose contributions they purport to measure. This finding is relevant beyond residential retrofits: similar concerns may apply to Shapley-based allocations in shared-renewable-energy systems (Alonso Pedrero et al., 2024), peer-to-peer energy sharing (Han et al., 2021), and any setting where "marginal contribution" is measured against fixed physical attributes rather than chosen actions.

Second, the polluter-pays principle does not transfer to the within-building scale. In 5.2, we argued that the polluter-pays principle loses its normative content when applied within residential buildings, where dwelling-level performance is structurally inherited, and individual occupants cannot unilaterally alter the attributes for which they would be charged. This conceptual failure is reinforced by an empirical observation. Performance-based rules invoking polluter-pays logic produce within-building cost-share dispersion that is statistically dominated by structural attributes (envelope quality, position, orientation) rather than behavioural choices. The implication is that polluter-pays-style allocation is more appropriately implemented at scales where the principle's conditions can be met, at the building level, where the building's owners can collectively undertake corrective action; or at the programme level, where public subsidies can be calibrated by building performance, than at the within-building level, where individual leaseholders cannot.

This implication has theoretical bearing on the broader environmental-justice and energy-justice literatures, which have increasingly engaged with intra-household and intra-building distributional questions (Bouzarovski and Simcock, 2017; Heffron and McCauley, 2017; Walker and Day, 2012). Our results suggest that fine-grained distributional analyses at sub-building scales should be cautious about importing principles developed for industrial-scale environmental governance, where the actor-action-harm-capacity chain is qualitatively different.

Third, the unit of analysis matters for distributional findings. Existing distributional analyses of climate and energy policy operate primarily at the household or national scale (Antosiewicz et al., 2022; Montenegro et al., 2021), examining aggregate effects across populations. We have argued that the binding distributional question for retrofit decisions plays out within buildings, where co-owners must agree. Our Theil decomposition (Appendix B.1) confirms that more than 97% of total inequality across all apartments in our pooled dataset arises from between-building differences rather than from within-building differences. A household-level distributional analysis would identify the between-building component and miss the within-building component entirely.

This points to a methodological implication for distributional analysis of climate policy more broadly. Where decisions are made at sub-national, sub-household, or sub-population scales, e.g. homeowner associations, neighbourhood retrofit cooperatives, district-heating consumer groups, the unit-of-analysis choice is consequential. Aggregate analyses can identify whether a policy is regressive in expectation, but they cannot determine whether the implementation rule fails at the level where it actually occurs. Both scales matter; neither subsumes the other.

Beyond individual buildings, our results have implications for the design of public retrofit subsidies. Several large EU and national programmes, e.g. the UK Warm Homes scheme, the Horizon Europe Built Environment cluster, and forthcoming Social Climate Fund implementations, must specify how subsidy is calibrated to building characteristics and how the residual non-subsidised cost is distributed within multi-family buildings.

5.5. Policy recommendations

Three concrete design recommendations follow from our analysis. First, it is important to calibrate subsidies to building-level performance, not within-building performance. A grant programme that gives more public money to lower-performing buildings (and less to already-efficient ones) implements polluter-pays logic at a scale where the principle’s conditions partially apply: the building is the unit being charged, the building's collective owners can undertake corrective action, and enforcement runs against an institutional actor (the homeowner association or freeholder). A grant programme that requires within-building cost-sharing to follow performance-based allocation imports the polluter-pays logic into a setting where those conditions do not hold. The first design implements the principle responsibly; the second does not.

Second, specify the within-building cost-sharing default. Where a programme conditions a subsidy on retrofit completion, it implicitly delegates the within-building cost-sharing decision to the building's homeowner association or freeholder. Our results suggest that the default should be size-based (area-proportional in the UK and the Netherlands; equal-shares-or-as-existing in Poland), with deviations possible by unanimous agreement of co-owners. This default minimises within-building dispersion, satisfies cooperative stability, and is compatible with existing legal frameworks across all three countries.

Third, address affordability concerns through external means-tested support rather than through the allocation formula. As argued in 5.3, attempting to embed ability-to-pay considerations into a within-building allocation rule encounters insurmountable data and procedural obstacles: occupant incomes are not observable, affordability proxies (such as floor area) are not robust, and any deviation from the recorded fractional shares requires unanimous consent. A means-tested grant or loan programme that supports low-income apartment owners separately from the within-building allocation is more legally tractable, more politically defensible, and more protective of vulnerable occupants than embedding affordability proxies into the cost-allocation rule itself.

These three design choices together suggest a coherent framework: performance-based logic operates between buildings (subsidy targeting), size-based logic operates within buildings (default cost-sharing), and ability-to-pay considerations operate at the household level (means-tested support). Each principle is implemented at the scale where its underlying conditions are met.

5.6. Limitations and Scope

Our analysis has well-defined limits, which we summarise here for the reader's benefit and to circumscribe the paper's claims. First, EPC databases describe dwellings, not households. We do not observe income, household composition, occupancy, or metered consumption. Welfare interpretations of any of our distributional findings would require additional data, through tax-record linkage, census matching, or survey overlays, which is not currently feasible at scale in any of the three countries we study. Our progressivity indices (Kakwani, Suits) measure whether allocation rules are super-proportional or sub-proportional in dwelling floor area; they do not measure progressivity in the welfare-economics sense.

Second, EPC values are computed under standardised assumptions about occupancy, internal temperatures, and use profiles. The performance-gap literature documents systematic discrepancies between modelled and metered consumption. Our analysis remains policy-relevant because EPC values are the values institutional actors would use if they implemented performance-based allocation, but rules that operate on EPC values do not implement the polluter-pays principle in any meaningful sense (5.2).

Third, the three EPC systems report energy and emissions in different units and with different reference benchmarks. We harmonise to a per-m² emissions intensity for the emissions-proportional rule across all three countries, and we use country-appropriate reference targets for the inefficiency rule. The implications of alternative target specifications are explored in Appendix B.1 and discussed in 4.5; headline conclusions are robust, but the level of within-building dispersion is target-specific.

Third, we allocate the total retrofit cost; we do not allocate the subsequent benefits (energy savings, comfort improvements, property value increases). A net-cost analysis incorporating benefit allocation might lead to different distributional outcomes, particularly if higher-cost-bearing apartments also achieve larger absolute savings. Whether this would offset the dispersion documented in 4.1 depends on actual realised energy savings, which the performance-gap literature suggests are systematically lower than EPC-modelled savings. We flag this as a direction for future research in 6.

Fourth, we assume that the total cost is fixed and is divided across apartments. We do not model how the choice of allocation rule might itself affect the decision to retrofit, the scope of the retrofit, or the willingness of co-owners to participate. The cooperative-stability tests in 4.4 identify structural incentives for individual apartments to refuse, but whether and how those incentives translate into observed behaviour requires linked data on actual decisions, which would be a substantial separate research undertaking.

Fifth, the UK, Dutch, and Polish housing stocks span a meaningful range, heterogeneous building conversions, planned post-war apartment blocks, and standardised socialist-era prefabricated buildings, but they do not span the full diversity of European multi-family housing. We expect the qualitative ranking of rules to hold in other contexts, but the magnitudes of within-building dispersion and the cooperative-stability rates would vary with local building-stock characteristics.

6. Conclusions

This paper has provided the first large-scale empirical characterisation of the within-building distributional properties of cost allocation rules for residential retrofits, drawing on Energy Performance Certificate microdata covering 4 million apartments in 442,248 buildings across Poland, the UK, and the Netherlands. We applied five allocation rules: area-proportional, progressive area, emissions-proportional, inefficiency-proportional, and the Shapley value, to every multi-apartment building in our sample and characterised the resulting cost-share distributions using inequality, size-progressivity, and cooperative-stability metrics.

Three findings dominate the analysis. First, performance-based allocation rules produce within-building cost-share dispersion that is 2 to 11 times higher than that of size-based rules, depending on the country. The cross-country ordering: area < progressive area < emissions < inefficiency ≈ Shapley, is identical across building stocks of substantially different ages, ownership structures, and architectural types. The structural distributional properties of these rules are therefore robust, rather than artefacts of national contexts.

Second, every rule we examine charges smaller dwellings a larger share of cost than their share of total floor area would imply, and to a similar degree across rules within each country. The choice of allocation principle (size-based versus performance-based) matters far more for distributional outcomes than the specific performance metric or the convexity parameter chosen within either family.

Third, the Shapley rule violates the strict cooperative-stability axiom of individual rationality in 96% of UK buildings, 99% of Dutch buildings, and 72% of Polish buildings, and allocating to at least one apartment more than its stand-alone retrofit cost. Even under a 1.5×-tolerant test allowing 50% overpayment relative to the additive stand-alone benchmark, violations occur in approximately half of UK buildings. Polish prefabricated buildings, with their more homogeneous within-building energy profiles, are the most cooperatively stable under performance-based rules; the UK's are the least. The cooperative-game-theoretic axiomatic appeal of the Shapley value does not survive empirical implementation at scale when characteristic-function values are exogenous to the players whose contributions they purport to measure.

The implications for retrofit policy are practical and immediate. The case for performance-based within-building allocation rests on the assumption that asset-rated EPC values reliably capture the dwelling-level "responsibility" the rule is meant to charge for. Implementing performance-based allocation responsibly would therefore require investment in infrastructure for measured-consumption data, linked smart-meter records, and appropriate privacy and legal frameworks, none of which currently exists at scale in the three countries we study.

Until such infrastructure is in place, the responsible default is to keep performance-based logic at scales where its conditions are met. Public subsidy targeting can implement performance-based logic at the building level, where the building's collective owners have the capacity to undertake the corrective action being financed. Within-building cost allocation can default to transparency, legal entrenchment, cooperative stability, and consistency with existing decision-making frameworks. Affordability concerns can be addressed through external, means-tested support for low-income occupants, separate from the cost-allocation formula. Each of the three principles – performance-based, size-based, and ability-to-pay – operates at the scale where its underlying conditions hold.

Our analysis leads to three directions for future research. First: the need to integrate cost-side and benefit-side allocation. We allocate retrofit costs but do not allocate the subsequent benefits, i.e. energy savings, comfort improvements, and property value increases. A net-cost analysis incorporating benefit allocation might yield different distributional outcomes, particularly if higher-cost-bearing apartments also realise larger absolute savings. Whether this would offset the dispersion documented here depends on actual realised energy savings, which the performance-gap literature suggests are systematically lower than EPC-modelled savings. Field data on realised post-retrofit consumption would resolve this. Second, behavioural responses to allocation rules would be beneficial. We assume the total cost is fixed rather than modelling how the choice of allocation rule might itself affect the scope of retrofit, the willingness of co-owners to participate, or the possibility of side-payments. The cooperative-stability tests in 4.4 identify structural incentives for individual apartments to refuse, but observed responses require linked data on actual collective decisions and outcomes, a substantial separate research undertaking. Third, the link between dwelling-level allocation and household-level welfare is missing. EPC data describes dwellings, not households; a fully welfare-oriented distributional analysis would require linkage to tax records, census, or survey data with apartment-level identifiers. Such linkages are not currently feasible at scale in any of the three countries. Whether the structural sub-proportionality in dwelling size we document translates into welfare regressivity depends on the joint distribution of dwelling size, household income, and location, which varies across countries and across building types within countries.

Beyond these direct extensions, two broader research agendas follow from our methodological contribution. The within-building decomposition framework we develop is applicable to any setting where collective investment decisions are made by small numbers of individuals with heterogeneous characteristics, e.g. district-heating consumer groups, neighbourhood retrofit cooperatives, shared renewable-energy installations, and other forms of collective infrastructure. Whether the structural patterns we document generalise to those settings is an open empirical question. Procedural fairness, i.e. how the decision is made, by whom, and through what process, is likewise beyond our scope but central to whether any allocation rule will be accepted in practice. Our distributional analysis identifies which rules produce structurally tractable cost shares; whether residents perceive them as legitimate and consent to them is a separate question that quantitative survey work and qualitative case studies could address.

Our findings are an empirical foundation for matching cost-allocation principles to the scales at which their underlying conditions hold. Performance-based logic operates between buildings, where collective owners can undertake corrective action; size-based logic operates within buildings, where shares are observable, legally entrenched, and cooperatively stable; ability-to-pay logic operates at the household level, where it can be addressed through external means-tested support. Each principle implements a coherent normative claim at the scale where the claim is actually meaningful.

References

Ahlrichs, J., Rockstuhl, S., 2022. Estimating fair rent increases after building retrofits: A max-min fairness approach. Energy Policy 164, 112923.

Ali, U., Bano, S., Shamsi, M.H., Sood, D., Hoare, C., Zuo, W., Hewitt, N., O'Donnell, J., 2024. Urban residential building stock synthetic datasets for building energy performance analysis. Data in Brief 53, 110241.

Alonso Pedrero, R., Pisciella, P., Crespo del Granado, P., 2024. Fair investment strategies in large energy communities: A scalable Shapley value approach. Energy 295, 131033.

Ambec, S., Ehlers, L., 2016. Regulation via the Polluter-pays Principle. Econ J 126, 884–906.

Anderson, L.M., St. Charles, J., Fullilove, M.T., Scrimshaw, S.C., Fielding, J.E., Normand, J., 2003. Providing affordable family housing and reducing residential segregation by income: A systematic review. American Journal of Preventive Medicine 24, 47–67.

Antosiewicz, M., Fuentes, J.R., Lewandowski, P., Witajewski-Baltvilks, J., 2022. Distributional effects of emission pricing in a carbon-intensive economy: The case of Poland. Energy Policy 160, 112678.

Aronson, J.R., Lambert, P.J., 1994. Decomposing the gini coefficient to reveal the vertical, horizontal, and reranking effects of income taxation. National Tax Journal 47, 273–294.

Atkinson, A.B., 1970. On the measurement of inequality. Journal of Economic Theory 2, 244–263.

Attia, S., Kosiński, P., Wójcik, R., Węglarz, A., Koc, D., Laurent, O., 2022. Energy efficiency in the polish residential building stock: A literature review. Journal of Building Engineering 45, 103461.

Bal, M., Stok, M., Bombaerts, G., Huijts, N., Schneider, P., Spahn, A., Buskens, V., 2023. A fairway to fairness: Toward a richer conceptualization of fairness perceptions for just energy transitions. Energy Research & Social Science 103, 103213.

Borenstein, S., Davis, L.W., 2016. The Distributional Effects of US Clean Energy Tax Credits. Tax Policy and the Economy 30, 191–234.

Bouzarovski, S., Frankowski, J., Tirado Herrero, S., 2018. Low-Carbon Gentrification: When Climate Change Encounters Residential Displacement. International Journal of Urban and Regional Research 42, 845–863.

Bouzarovski, S., Simcock, N., 2017. Spatializing energy justice. Energy Policy.

Bright, S., Weatherall, D., Willis, R., 2019. Exploring the complexities of energy retrofit in mixed tenure social housing: a case study from England, UK. Energy Efficiency 12, 157–174.

Brounen, D., Kok, N., 2011. On the economics of energy labels in the housing market. Journal of Environmental Economics and Management 62, 166–179.

Carr, M.F., Boyd Kramer, D., 2022. Homeowners' associations: Barriers or bridges to more sustainable residential development? Landscape and Urban Planning 224, 104419.

Cole, I., Robinson, D., 2000. Owners yet Tenants: The Position of Leaseholders in Flats in England and Wales. Housing Studies 15, 595–612.

Cremers, S., Robu, V., Hofman, D., Naber, T., Zheng, K., Norbu, S., 2022. Efficient methods for approximating the shapley value for asset sharing in energy communities. In: Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, E-Energy '22. Association for Computing Machinery, New York, NY, USA, pp. 320–324.

Cremers, S., Robu, V., Zhang, P., Andoni, M., Norbu, S., Flynn, D., 2023. Efficient methods for approximating the Shapley value for asset sharing in energy communities. Applied Energy 331, 120328.

Crosby, N., Gibson, V., Murdoch, S., 2003. UK Commercial Property Lease Structures: Landlord and Tenant Mismatch. Urban Studies 40, 1487–1516.

Dorband, I.I., Jakob, M., Kalkuhl, M., Steckel, J.C., 2019. Poverty and distributional effects of carbon pricing in low- and middle-income countries – A global comparative analysis. World Development 115, 246–257.

Ebrahimigharehbaghi, S., Qian, Q.K., Meijer, F.M., Visscher, H.J., 2019. Unravelling Dutch homeowners' behaviour towards energy efficiency renovations: What drives and hinders their decision-making? Energy Policy 129, 546–561.

Frankowski, J., Prusak, A., Sokołowski, J., Mazurkiewicz, J., Świetlik, T., 2025. Energy Commons During Crisis: Understanding Housing Cooperative Strategies Through Dramaturgical Sociology. Int J Commons 19.

Fuerst, F., McAllister, P., Nanda, A., Wyatt, P., 2015. Does energy efficiency matter to home-buyers? An investigation of EPC ratings and transaction prices in England. Energy Economics 48, 145–156.

Galvin, R., 2023. How prebound effects compromise the market premium for energy efficiency in German house sales. Building Research & Information 51, 501–517.

García-Portela, L., 2024. A Minimal Capabilities-Based Account of Loss and Damage. Journal of Human Development and Capabilities 25, 170–193.

Goedkoop, F., Devine-Wright, P., 2016. Partnership or placation? The role of trust and justice in the shared ownership of renewable energy projects. Energy Research & Social Science 17, 135–146.

Grossmann, K., 2019. Using conflicts to uncover injustices in energy transitions: The case of social impacts of energy efficiency policies in the housing sector in Germany. Global Transitions 1, 148–156.

Hamnett, C., Randolph, B., 2021. Cities, Housing and Profits: Flat Break-Up and the Decline of Private Renting. Routledge, London.

Han, L., Morstyn, T., McCulloch, M., 2021. Estimation of the Shapley Value of a Peer-to-peer Energy Sharing Game Using Multi-Step Coalitional Stratified Sampling. Int. J. Control Autom. Syst. 19, 1863–1872.

Heffron, R.J., McCauley, D., 2017. The concept of energy justice across the disciplines. Energy Policy 105, 658–667.

Heffron, R.J., Sokołowski, M.M., 2024. Resolving energy policy failure: Introducing energy justice as the solution to achieve a just transition. Energy Policy 187, 114042.

Hoekstra, J., 2017. Reregulation and Residualization in Dutch Social Housing: A Critical Evaluation of New Policies. Critical Housing Analysis 4, 31–39.

Hoppe, T., Lulofs, K., 2008. The Impact of Multi-Level Governance on Energy Performance in the Current Dutch Housing Stock. Energy & Environment 19, 819–830.

Jafari, A., Valentin, V., 2017. An optimization framework for building energy retrofits decision-making. Building and Environment 115, 118–129.

Kakwani, N.C., 1977. Measurement of Tax Progressivity: An International Comparison. Econ J 87, 71–80.

Kaschny, L., 2023. Energy Justice and the Principles of Article 194(1) TFEU Governing EU Energy Policy. Transnational Environmental Law 12, 270–294.

Kendall, S.H., 2021. Developments toward a residential fit-out industry: the key to a sustainable housing stock. In: Residential Architecture as Infrastructure. Routledge.

Li, Y., Kubicki, S., Guerriero, A., Rezgui, Y., 2019. Review of building energy performance certification schemes towards future improvement. Renewable and Sustainable Energy Reviews 113, 109244.

Lujanen, M., 2010. Legal challenges in ensuring regular maintenance and repairs of owner-occupied apartment blocks. International Journal of Law in the Built Environment 2, 178–197.

Luppi, B., Parisi, F., Rajagopalan, S., 2012. The rise and fall of the polluter-pays principle in developing countries. International Review of Law and Economics, The Economics of Efficiency and the Judicial System 32, 135–144.

Lux, M., Sunega, P., 2014. Public Housing in the Post-Socialist States of Central and Eastern Europe: Decline and an Open Future. Housing Studies 29, 501–519.

Lux, M., Sunega, P., Katrňák, T., 2011. Classes and Castles: Impact of Social Stratification on Housing Inequality in Post-Socialist States. European Sociological Review 29, 274–288.

Madeddu, M., Clifford, B., 2023. The conversion of buildings to housing use: England’s permitted development rights in comparative perspective. Progress in Planning, The conversion of buildings to housing use: England’s permitted development rights in comparative perspective 171, 100730.

Maestre-Andrés, S., Drews, S., van den Bergh, J., 2019. Perceived fairness and public acceptability of carbon pricing: a review of the literature. Climate Policy 19, 1186–1204.

Montenegro, R.C., Fragkos, P., Dobbins, A.H., Schmid, D., Pye, S., Fahl, U., 2021. Beyond the Energy System: Modeling Frameworks Depicting Distributional Impacts for Interdisciplinary Policy Analysis. Energy Technology 9, 2000668.

Muczyński, A., 2023. Collective renovation decisions in multi-owned housing management: the case of public–private homeowners associations in Poland. J Hous and the Built Environ 38, 2105–2127.

Ohlendorf, N., Jakob, M., Minx, J.C., Schröder, C., Steckel, J.C., 2021. Distributional Impacts of Carbon Pricing: A Meta-Analysis. Environ Resource Econ 78, 1–42.

Pearce, J., Vine, J., 2014. Quantifying residualisation: the changing nature of social housing in the UK. J Hous and the Built Environ 29, 657–675.

Ploeger, H., Groetelaers, D., 2014. Managing multi-owned, mixed-use developments in The Netherlands: The opinion of professional managing agents. International Journal of Law in the Built Environment 6, 271–285.

Poon, J., Garratt, D., 2012. Evaluating UK housing policies to tackle housing affordability. International Journal of Housing Markets and Analysis 5, 253–271.

Shapley, L.S., 1953. Stochastic Games*. Proceedings of the National Academy of Sciences 39, 1095–1100.

Shapley, L.S., 2016. 17. A Value for n-Person Games. In: Kuhn, H.W., Tucker, A.W. (Eds.), Contributions to the Theory of Games, Volume II. Princeton University Press, pp. 307–318.

Si, J., Marjanovic-Halburd, L., 2018. Criteria weighting for green technology selection as part of retrofit decision making process for existing non-domestic buildings. Sustainable Cities and Society 41, 625–638.

Sovacool, B.K., 2015. Fuel poverty, affordability, and energy justice in England: Policy insights from the Warm Front Program. Energy 93, 361–371.

Suits, D.B., 1977. Measurement of Tax Progressivity. The American Economic Review 67, 747–752.

Theil, H., 1968. On the geometry and the numerical approximation of cost of living and real income indices. De Economist 116, 677–689.

Tsenkova, S., 2017. Housing Change in East and Central Europe: Integration or Fragmentation? Routledge.

van den Brom, P., Meijer, A., Visscher, H., 2018. Performance gaps in energy consumption: household groups and building characteristics. Building Research & Information 46, 54–70.

van der Kooij, P., 2025. Putting the three-tenets energy justice framework into policy practice: A case study in the Netherlands. Energy Research & Social Science 127, 104300.

Walker, G., Day, R., 2012. Fuel poverty as injustice: Integrating distribution, recognition and procedure in the struggle for affordable warmth. Energy policy.

Weber, I., Wolff, A., 2018. Energy efficiency retrofits in the residential sector – analysing tenants' cost burden in a German field study. Energy Policy 122, 680–688.

Appendix A. Detailed Methods

A1. National EPC systems, data sources, and harmonisation

The three national EPC systems report similar concepts but with different units and reference benchmarks. For inefficiency-based rules we use, in each country, the asset-rated primary energy demand in kWh/m²/year: current energy consumption from the UK Domestic EPC Register (a final-energy measure); primary fossil energy (*Primaire Fossiele Energie*) from the Dutch EP; and primary energy indicator (*wskaźnik energii pierwotnej*) from the Polish CRChE. We use UK final energy rather than constructing a primary-energy equivalent because primary-energy weights are not consistently reported across the SAP register; we discuss the implications and present a robustness check using energy consumption potential ratios in Appendix B.

The post-improvement reference required for the inefficiency rule varies in availability by country:

- England and Wales: We use the per-dwelling energy consumption potential, which the SAP/RdSAP assessor calculates by applying a standardised package of cost-effective recommended improvements (loft insulation, cavity-wall insulation, condensing boiler, etc.) to the assessed dwelling.
- Netherlands: We use the per-dwelling NTA 8800 regulatory primary-energy requirement, which varies by building type and construction year. Where this field is missing, we substitute a fixed national reference of 70 kWh/m²/year, approximating the Dutch low-energy housing standard.
- Poland: The Polish EPC system reports a primary energy indicator for current dwellings but does not consistently report a per-dwelling post-improvement value. We use a fixed national reference of 50 kWh/m²/year, corresponding to the EU near-zero energy building (NZEB) primary-energy benchmark.

These three reference choices are not equivalent, and the choice carries weight. We present results under three alternative target specifications as robustness checks (Appendix B): country-median apartment-level energy intensity; the fixed 50 kWh/m²/year NZEB benchmark applied uniformly across countries; and within-building minimum energy intensity (a hard upper bound that always assigns one apartment a zero share). Headline conclusions hold across these alternatives, but the level of within-building dispersion depends on the target, as we document.

Finally, the three national systems report CO_2 in different units. UK SAP reports absolute annual emissions in tonnes per year. Dutch EP-Online reports per-m² intensity in kg CO_2/m²/year. Polish EPC has per-m² intensity in tonnes CO_2/m²/year. Naïve cross-country comparison of an emissions-proportional rule operating directly on these fields would conflate different operationalisations. In the Netherlands and Poland, the rule is already size-neutral by construction, whereas in the UK it imposes costs on larger dwellings through the size component of total emissions. We therefore harmonise to per-m² intensity throughout, multiplying or dividing by floor area as needed, and present results under both the harmonised intensity specification and the unharmonised total specification in Appendix B.

A2. Cost-allocation rules: formal definitions and implementation

We evaluate five methods for allocating the total cost of a building retrofit among the apartments in that building. For a building B with total retrofit cost C and n apartments B, an allocation is defined by a vector of cost shares $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ such that $\pi_i \geq 0$ and $\sum_i \pi_i = 1$. Apartment i's monetary cost burden is $\pi_i \times C$. Each method corresponds to a different rule for determining π. As discussed in 2.1.1, EOC databases do not record occupancy, household composition, or income; the allocation rules we evaluate operate only on dwelling-level attributes. Below, we describe each method and its underlying principle.[6]

Area-Proportional Allocation. Costs are divided in proportion to each apartment's floor area. If A_i is the floor area of apartment i, then $\pi_i^{area} = \frac{A_i}{\sum_i A_j}$. This rule reflects the principle that each owner's stake in the building (and thus in the improvement) is proportional to their unit size.

Progressive Area Allocation. This rule applies a convex weighting to floor area. It makes larger apartments pay more than their proportional share, and smaller units pay less than their proportional share of the total cost. We operationalise this by raising the floor area to an exponent $\alpha > 1$. In our baseline, we use $\alpha = 1.5$

---

[6] All five rules are computed for the UK and the Netherlands. For Poland, area-proportional and progressive area cannot be computed (apartment-level area is unavailable); equal, emissions-proportional, inefficiency-proportional, and Shapley are computed. The full availability map is given in Table 2.

(moderately convex). Thus, $\pi_i^{prog} = \frac{A_i^\alpha}{\sum_j A_j^\alpha}$. If $\alpha = 1$, this reduces to the standard area rule; higher α increases the convexity. This method is sometimes proposed as a way to shift more costs onto larger dwellings, on the assumption that larger dwellings correlate with higher household income. Whether that correlation holds in any given building stock is empirical and beyond the scope of our analysis (see 2.3.2). We include the rule because it is occasionally discussed as a middle-ground alternative to strict area-proportional allocation.

Emissions-Proportional Allocation. Costs are allocated in proportion to each apartment's asset-rated CO_2 emissions intensity in kg/m^2/year. Let E_i be apartment i's EPC-modelled emissions intensity (per 2.1.1, harmonised to a common per-m^2 basis across countries). Then $\pi_i^{emiss} = \frac{E_i}{\sum_j E_j}$. This rule operationalises a performance-based allocation logic by charging according to asset-rated emissions intensity. We discuss the relationship to the polluter-pays principle in 4.2; for the empirical analysis, we treat it as one of several rules whose distributional properties we characterise. We additionally present results using total (size-confounded) emissions in Appendix B.

Inefficiency-Proportional Allocation. This method, closely related to emissions, allocates costs according to each unit's modelled improvement potential. As detailed in 2.1.1, $I_i = \max(0, X_i^{current} - X_i^{target})$ measures the modelled energy savings potential. Then $\pi_i^{ineff} = \frac{I_i}{\sum_j I_j}$. This is another formulation of the polluter-pays principle, focusing on who has the most to gain from retrofit. In our setup, the Shapley value allocation (below) under an additive model is theoretically equivalent to this inefficiency-proportional rule. The inefficiency rules charge based on modelled improvement potential under standardised assessor assumptions; it is not a measure of realised waste or harm. The choice of X^{target} is consequential and varies across countries; we document our per-country choices in 2.1.1 and present three alternative target specifications as robustness checks in Appendix B.

Shapley Value Allocation. The Shapley value comes from cooperative game theory (Shapley, 1953) and is often proposed as a fair division rule for joint costs or benefits. We treat the retrofit cost allocation as a cooperative game in which each apartment is a player, and the value of any coalition is the cost savings from retrofitting that coalition. The Shapley value of a player is their average marginal contribution to all possible coalitions. In practical terms, it allocates the cost based on each unit's marginal contribution to the total modelled improvement potential. The Shapley value satisfies axioms of efficiency (all costs allocated), symmetry (treat equals equally), the null player property (if an apartment adds no extra cost or inefficiency, it pays zero), and additivity. In our context, with an additive characteristic function, the Shapley formula reduces to the inefficiently proportional allocation described above. We retain the Shapley label specifically because it allows us to invoke cooperative-stability concepts: the core, individual rationality, that motivate much of the energy-community literature's interest in this allocation, and which we evaluate in 3.4. In robustness checks (Appendix B), we test alternative characteristic functions, including current emissions and floor area.

It is essential to note that we assume the total cost, C, is fixed for the building (i.e., determined by a chosen retrofit package) and does not vary with the allocation method.[7] We are interested in how the same total bill

---

[7] Our analysis focuses on the allocation of cost shares rather than absolute monetary amounts. We assume that the total retrofit cost C for each building is fixed and exogenously determined by the chosen retrofit package; our interest lies in how this total is divided among apartments, not in what determines C. Each allocation method generates a vector of proportions $(\pi_1, \pi_2, \ldots, \pi_n)$ summing to one, where apartment i would pay $\pi_i \times C$. This approach has two advantages: it eliminates the need to estimate building-specific retrofit costs (which vary widely due to building

can be split differently. We do not consider behavioural responses or different retrofit choices under each scheme; our analysis is positive, assuming the cost is given and examining the distributional outcomes.

A3. Distributional, progressivity, and cooperative-stability metrics

For each building, we have an allocation of shares $\{\pi_i\}$, and compute its dispersion, redistributive properties, and stability. We then look at the distribution of these metric values across the population of buildings.

A3.1. Within-Building Inequality

For each building, we treat the vector of cost shares π as a distribution of contributions and compute standard inequality indices. This is a structural application of inequality indices to allocation outputs; we make no inference about household-level economic status.

The Gini index is a measure of inequality ranging from 0 (perfect equality: all π_i equal) to 1 (maximum inequality: one person pays everything). We use the formula $G = \frac{1}{2n} \sum_i \sum_j |\pi_i - \pi_j|$, which captures the average absolute difference between two randomly chosen shares.[8] We will often classify within-building Gini < 0.1 as "low inequality," 0.1–0.2 as moderate, and >0.2 as high, for descriptive purposes. We compute the Gini including zero shares (apartments paying nothing), which can arise under the inefficiency-proportional and Shapley rules when an apartment's modelled energy intensity is at or below the reference target. Excluding zero shares would understate dispersion; we report the inclusive measure as our headline and the exclusive measure as a robustness check.

The Atkinson index is a parametric family of inequality measures that incorporates different levels of inequality aversion. Atkinson, (1970) introduced this index to answer how much total welfare is lost due to inequality, given a social aversion parameter ε. Higher ε means society is more averse to inequality (giving more weight to disparities at the lower end). We compute Atkinson indices at multiple ε values to assess the sensitivity of our comparisons to the emphasis on the bottom of the distribution. For $\varepsilon = 1$, the Atkinson can be interpreted via the geometric mean: $A(\varepsilon = 1) = 1 - \frac{\mathrm{GeometricMean}(\pi)}{\mathrm{ArithmeticMean}(\pi)}$. In our context, since the π's sum to 1, the arithmetic mean is $\frac{1}{n}$; a higher geometric mean relative to $\frac{1}{n}$ indicates lower inequality. The Atkinson index is used to illustrate extreme cases.[9]

We use the **Theil T** (which is sensitive to disparities at the top end) and Theil L (sensitive to the bottom end). Theil T is calculated as $T = \frac{1}{n} \sum_i \frac{\pi_i}{\pi} ln\left(\frac{\pi_i}{\pi}\right)$. Theil L is $L = \frac{1}{n} \sum_i ln\left(\frac{\overline{\pi}}{\pi_i}\right)$. Both are 0 when all π equal, and higher values indicate more inequality (Theil, 1968). A benefit of Theil indices is that they can be decomposed into within-group and between-group components. We exploit this by treating each building as a group and

---

characteristics, retrofit scope, and local prices), and it separates the distributional question from the cost estimation question. For the inequality metrics, we compute Gini coefficients, Atkinson indices, and Theil indices; only the relative shares matter, not their absolute values. The proportions are computed directly from EPC data: floor area for area-based methods, CO_2 emissions for emissions-proportional allocation, and improvement potential (defined as the difference between current and potential energy consumption, clipped at zero) for inefficiency-based and Shapley allocations. Where characteristic data is missing for an apartment, we default to equal shares $(1/n)$ for that building.

[8] For example, in a two-apartment building, a Gini of 0.33 would result if one pays 67% and the other 33%.

[9] For instance, at $\varepsilon = 2$ (a very high aversion), some allocations, such as Shapley, produce very high Atkinson values. This is a structural welfare-loss interpretation of the inequality index given a hypothetical social welfare function; it is not a claim about residents' actual welfare, which we cannot observe.

decomposing the total inequality across all individuals (i.e., all apartments in the dataset) into within-building and between-building components. This tells us whether a given allocation method systematically causes high inequality between buildings (e.g., if some buildings as a whole pay more than others, which in our setup might happen if total costs vary by building size) versus within.

We also compute simpler measures, such as the **coefficient of variation (CV)** of shares (standard deviation divided by the mean) and the **maximum/minimum ratio** within each building. The max/min ratio tells us how many times larger the biggest payer's share is compared to the smallest payer's share. For example, a maximum/minimum of 3 means one apartment pays three times what another apartment pays. We report median values of such ratios for different methods.

After computing these metrics per building, we typically report the mean or median across all buildings, or the percentage of buildings that fall below certain thresholds (like Gini < 0.1). We treat each building equally in these aggregations (i.e., not weighting by the number of apartments, since each building is a decision context in its own right). Confidence intervals for building-level means are obtained using a nonparametric bootstrap with 1,000 replications, sampling buildings with replacement.

A3.2. Size-Progressivity

We use floor area as a structural attribute of dwellings against which to measure the redistributive properties of each allocation rule. The Kakwani and Suits indices, as we apply them, measure whether an allocation rule is super-proportional or sub-proportional in dwelling size, which is a structural property of the rule given the building stock. Whether sub-proportionality in dwelling size translates into regressive welfare outcomes depends on the joint distribution of dwelling size and household income, which we do not observe and which is heterogeneous in practice.

We use established indices from public economics, applied here to measure size-progressivity rather than welfare progressivity. Polish EPCs lack apartment-level area data, so the size-progressivity indices are computed for the UK and the Netherlands only.[10]

The Kakwani Index (Kakwani, 1977) measures the difference between the distribution of the cost burden and the distribution of the size attribute. It is defined as $\Pi_K = C_t - G_x$, where C_t is the concentration coefficient of the cost shares (which is like a Gini but ranking individuals by size instead of their cost share) and G_x is the Gini of the size variable (e.g., the Gini of floor area across apartments). A positive Kakwani indicates a super-proportional in size (larger units pay more than their proportionate share of floor area), zero indicates a size-proportional, and negative indicates a sub-proportional in size (smaller units pay more than their share of total floor area).

---

[10] Additionally, we consider the Reynolds-Smolensky Index, which measures the redistributive effect of moving from a proportional benchmark to the given allocation. It is the difference in the Gini of capacity before and after the cost allocation. In formula, $\Pi_{RS} = G_x - G_{x-T}$, where G_{x-T} is the Gini of post-cost incomes. This index can also be related to Kakwani by $\Pi_{RS} = \frac{t}{1-t} \times \Pi_K$ for a flat tax rate t (share of total income taken as tax). In our case, since the "tax" (retrofit cost) is a one-time charge and we do not have incomes, we mainly use Kakwani as a simpler indicator of progressivity. However, conceptually the RS index reminds us that even a highly progressive rate structure might have minimal effect on inequality if the overall amount is small relative to income (in a building context, if retrofit costs are low relative to property values or other wealth, the redistributive effect on residents' overall economic status could be minor – something we note but do not directly quantify due to lack of income data).

The Suits Index is closely related to Kakwani and is often used in public finance to assess tax progressivity (Suits, 1977). It is defined as $S = 1 - 2L$, where L is the area under the concentration curve plotting cumulative cost burden vs cumulative size. In practice, the Suits index will give the same qualitative assessment as Kakwani (the two move in tandem in our results, so we primarily report Kakwani). Both range from -1 (extremely sub-proportional) to 1 (extremely super-proportional), with 0 indicating size-proportionality.

The fairness ideal is that similar units should be treated similarly (horizontal equity) and the ordering of who pays more should align with the ordering of who can pay (no reranking). We address this through the Atkinson-Plotnick Index for Reranking, which evaluates the extent to which imposing cost allocation alters the rank ordering of individuals by economic position (Aronson and Lambert, 1994). In our context, if capacity is equal to area, reranking would occur if, for example, a smaller apartment ends up paying more than a larger apartment. That would be a violation of the expected order. The index can be defined as $R = G_y - C_y$, the difference between the Gini of capacity and the concentration index of cost after sorting by capacity. We desire $R = 0$ (no reranking) ideally. Our methods generally do not cause reranking if capacity equals area for the area-based methods.[11]

A.3.3. Cooperative Stability

The cooperative-stability tests in this subsection do not require any interpretation of EPC-modelled values as proxies for ability or responsibility. They are structural tests of whether the allocation rule produces shares that exceed stand-alone benchmarks.

Since one of our methods (Shapley) is derived from cooperative game theory, we also evaluate each allocation against stability criteria:

The core of a cost allocation game is the set of allocations where no subset of players (apartments) would be better off splitting off and doing the retrofit by themselves (or in a smaller coalition). For an allocation to be in the core, every coalition's allocated cost must be no more than the cost that coalition would incur on its own. For computational tractability, we test the necessary individual-rationality condition: no single apartment pays more than its stand-alone retrofit cost. If any single apartment's allocated cost $\pi_i^* C$, exceeds the cost of improving that apartment alone, that apartment has an incentive to leave the coalition. For each apartment, we proxy the stand-alone cost by that apartment's fraction of the total building floor area, which serves as a natural reference point under the assumption that retrofit costs scale with dwelling size. We report core stability under two definitions:

- Strict. The building is in the core if no apartment's allocated share exceeds its area-proportional stand-alone share. This is the textbook cooperative-game-theoretic definition.
- 1.5×-tolerant. The building is in the core if no apartment's allocated share exceeds 1.5 times its area-proportional stand-alone share. This is a "meaningful violation" threshold that permits modest overpayments that might be tolerated in practice through informal side-payments or cooperative norms.

---

[11] By design, the area method keeps rank by area monotonic in burden; equal can cause reranking since all pay equal despite different areas, and polluter-pays can cause significant re-ranking if inefficiency is not perfectly correlated with size. The literature sometimes decomposes the total redistributive effect into a vertical part (assuming no re-rank, how much inequality reduction (or increase) happens) and a loss due to horizontal inequity and reranking (Atkinson-Plotnick's approach). In our study, since we are not evaluating an initial income distribution but rather cost shares, we do not calculate this aspect. But we qualitatively note when a method might violate horizontal equity.

Related to the core, an allocation is subsidy-free if no group of apartments pays less than the incremental cost they impose on the system. In other words, each coalition pays at least the cost of retrofitting itself, minus the cost of retrofitting the rest of the building. We check a simplified version: no single apartment pays less than the incremental cost it adds. For additive costs, this condition is typically met automatically if it is in the core.

We also calculate the largest ratio by which any apartment's allocated cost exceeds its stand-alone cost ("maximum excess ratio"). A value above 100% means some apartments are paying more than twice the cost of their isolated retrofit.

An allocation that violates individual rationality, i.e. that charges some apartment more than its stand-alone retrofit cost, creates a structural incentive for that apartment to refuse to participate in the joint project. Whether this incentive translates into actual refusal depends on residents' preferences, information, and outside options, which we do not observe. The structural property itself is a property of the rule and the building stock, and is informative regardless of how that translation is ultimately resolved.

Appendix B. Additional results

B1. Theil Decomposition: Within vs. Between Buildings

Total inequality in cost shares across all apartments in our pooled sample can be decomposed into within-building and between-building components. The between-building component dominates: it accounts for the great majority of total cost-share inequality in absolute terms across our sample, regardless of allocation method. This pattern holds because the total cost of retrofitting a building is itself heterogeneous across our sample, and a single apartment's absolute cost burden is therefore determined first by which building it is in and only secondarily by how that building's total cost is divided.

This decomposition supports our framing of the analysis around within-building inequality rather than population-level apartment inequality. From a household-level perspective, the dominant determinant of cost-share inequality is between-building variation in retrofit cost, which is exogenous to the choice of allocation rule and which public subsidies could, in principle, equalise. From the perspective of a co-owner negotiating with their neighbours, however, the binding distributional question is how the (fixed) total cost of their building is divided among the apartments in it. The choice of allocation rule has no first-order effect on between-building inequality (since each rule allocates the same total cost), but it has a substantial effect on within-building dispersion, as documented in 4.1.

B2. Subgroup Analyses

We perform tenure-based stratification for the UK to identify systematic differences in within-building dispersion across ownership categories, after consolidating the raw EPC tenure values into four categories: Owner-occupied, Rented (social), Rented (private), and Unknown/other. The patterns are as follows.

UK tenure	N buildings	Area Gini	Shapley Gini	Difference	Ratio
Rented (social)	301,912	0.018	0.262	+0.244	14.2
Rented (private)	343,594	0.025	0.246	+0.221	9.9
Owner-occupied	160,736	0.023	0.218	+0.195	9.6
Unknown/other	41,190	0.014	0.065	+0.051	4.7

Source: Own elaboration based on UK EPC data.

UK social rental buildings exhibit higher within-building Shapley dispersion than owner-occupied buildings, with a 14× ratio between Shapley and area-proportional dispersion in social housing, the highest in our analysis. The difference reflects greater heterogeneity in within-building modelled energy performance in social housing, possibly due to uneven historical maintenance or to individual apartments receiving different amounts of capital investment over time. Whatever the cause, the structural pattern is that performance-based allocation, applied to social housing, would concentrate cost shares disproportionately on individual dwellings within the sector that already house lower-income occupants.

The Unknown/other category, populated by EPCs without reliable tenure information, exhibits much lower dispersion than the named categories. This pattern is consistent with the consolidation rule absorbing rows with otherwise-incomplete data, which may also have lower-quality energy fields. Results in this category should not be taken at face value.

As discussed in 5.1, the policy relevance of these tenure-stratified results varies. For owner-occupied leaseholds, the within-building cost-allocation question is directly faced by the leaseholder. For private rentals, landlords absorb costs and may pass them through to tenants only partially. For social housing, landlords typically fund retrofits centrally, and the pass-through to tenants is governed by separate frameworks. The structural pattern documented here describes how costs would be distributed across dwellings if performance-based allocation were applied; the welfare implications depend on the governance and pass-through mechanisms specific to each tenure (Sovacool, 2015).

B3. Median Max/Min Ratio Within Buildings

Allocation Method	Poland	UK	Netherlands
Area-proportional	N/A	1.04	1.25
Progressive area ($\alpha = 1.5$)	N/A	1.07	1.40
Emissions-proportional	1.27	1.25	1.38
Shapley value	1.40	1.93	1.53

Source: Own elaboration. Each cell is the median across all multi-apartment buildings of the within-building ratio of the highest cost share to the lowest cost share. Values close to 1 indicate roughly equal cost shares; values further from 1 indicate larger gaps between the highest- and lowest-paying apartments.

The maximum-to-minimum ratio is a more intuitive summary than the Gini for non-technical readers: a value of 1.93 in the median UK building under Shapley means that the highest-paying flat contributes nearly twice as much as the lowest-paying flat. The Polish and Dutch ratios under Shapley are more compressed (1.40 and 1.53, respectively), reflecting more uniform within-building energy profiles in those housing stocks. The cross-rule ranking (Area < Progressive < Emissions < Shapley) is preserved across countries.

B4. Atkinson Indices (Mean Per Building)

Country	Method	$\varepsilon = 0.5$	$\varepsilon = 1.0$	$\varepsilon = 2.0$
UK	Area-proportional	0.002	0.004	0.008
UK	Progressive area	0.004	0.008	0.015
UK	Emissions	0.011	0.021	0.040
UK	Inefficiency / Shapley	0.053	0.106	0.174

NL	Area-proportional	0.011	0.022	0.038
NL	Progressive area	0.021	0.041	0.070
NL	Emissions	0.009	0.018	0.034
NL	Inefficiency / Shapley	0.018	0.036	0.067
PL	Emissions	0.018	0.035	0.064
PL	Inefficiency / Shapley	0.029	0.058	0.104

Source: Own elaboration. Each value is the mean across multi-apartment buildings of the within-building Atkinson index at the corresponding inequality-aversion parameter ε.

Higher ε weights inequality among apartments paying smaller shares more heavily; the contrast across allocation rules is preserved at all ε values, with performance-based rules exhibiting two-to-five-fold higher Atkinson indices than size-based rules in every country.

B5. Theil Indices (Mean Per Building)

Country	Method	Theil T	Theil L	T > L?
UK	Area-proportional	0.004	0.004	–
UK	Progressive area	0.008	0.009	–
UK	Emissions	0.021	0.023	–
UK	Inefficiency / Shapley	0.241	0.166	Yes
NL	Area-proportional	0.022	0.027	–
NL	Progressive area	0.042	0.056	–
NL	Emissions	0.018	0.019	–
NL	Inefficiency / Shapley	0.066	0.041	Yes
PL	Emissions	0.035	0.042	–
PL	Inefficiency / Shapley	0.077	0.068	Yes

Source: Own elaboration. Theil T is sensitive to inequality at the top of the cost-share distribution; Theil L is sensitive to inequality at the bottom. T > L indicates that within-building dispersion is concentrated among the highest-paying apartments rather than the lowest-paying ones.

The Theil decomposition reveals a substantive pattern not visible in the Gini summary. For size-based rules (area, progressive area, emissions), Theil L slightly exceeds Theil T in all three countries; the within-building dispersion these rules produce is shaped by many small differences across all apartments. For the inefficiency-proportional and Shapley measures, the relationship reverses: Theil T markedly exceeds Theil L. This indicates that performance-based rules generate within-building dispersion through a small number of apartments paying very large shares, rather than through broad-based variation across all apartments. The top-end concentration is the structural feature that gives rise to the cooperative-stability violations documented in 4.4.

B6. Size-Progressivity Indices (Kakwani and Suits)

Country	Method	Kakwani	Suits
UK	Area-proportional	−0.177	−0.177
UK	Progressive area	−0.172	−0.172
UK	Emissions	−0.193	−0.192
UK	Inefficiency / Shapley	−0.187	−0.187
NL	Area-proportional	−0.687	−0.749

NL	Progressive area	−0.666	−0.743
NL	Emissions	−0.743	−0.767
NL	Inefficiency / Shapley	−0.748	−0.768

Source: Own elaboration. Negative values indicate that the rule is sub-proportional in dwelling floor area: smaller dwellings receive larger cost shares relative to their share of total floor area. The interpretation is structural (a property of the rule applied to the building stock) and not welfare-related (3.3.2). Polish EPCs lack apartment-level area data and are excluded from this analysis (3.1.2).

Within each country, Kakwani and Suits values cluster tightly across rules, indicating that the principle (size-based vs performance-based) matters more than the specific performance metric. UK values span a 0.02-point range across all four rules; Dutch values span 0.08 points. The much larger absolute magnitudes in the Netherlands reflect greater within-building variation in Dutch apartment sizes than in the UK, where two-flat conversions tend to contain similarly-sized units.

B7. Cooperative Stability by Allocation Method

This is the appendix counterpart to Table 6 in the body. Each entry reports the percentage of buildings satisfying the individual-rationality condition, alongside its bootstrap 95% confidence interval (1,000 replications, building-level resampling).

Country	Method	Strict (95% CI)	1.5×-tolerant (95% CI)
UK	Progressive area	13.3 (13.2–13.3)	100.0 (100.0–100.0)
UK	Equal	13.3 (13.2–13.3)	98.7 (98.7–98.8)
UK	Emissions	2.3 (2.3–2.3)	93.2 (93.1–93.2)
UK	Inefficiency / Shapley	3.6 (3.5–3.6)	52.5 (52.4–52.6)
NL	Progressive area	7.4 (7.2–7.5)	99.8 (99.8–99.8)
NL	Equal	7.4 (7.2–7.5)	90.1 (90.0–90.3)
NL	Emissions	0.3 (0.3–0.4)	75.9 (75.7–76.1)
NL	Inefficiency / Shapley	0.6 (0.6–0.7)	67.5 (67.3–67.8)
PL	Equal	99.7 (99.6–99.7)	99.8 (99.8–99.8)
PL	Emissions	29.2 (28.9–29.4)	85.9 (85.7–86.2)
PL	Inefficiency / Shapley	27.6 (27.4–27.9)	74.7 (74.4–74.9)

Source: Own elaboration. Confidence intervals are tight enough to support the headline cross-method ranking with high precision. The strict test treats the area-proportional share as the additive stand-alone-cost benchmark; the 1.5×-tolerant test allows up to 50% overpayment relative to that benchmark. Polish entries for area and progressive area are not computed (3.1.2).

B8. Results by Building Size

Country	Building Size	N buildings	Area Gini	Shapley Gini	Difference	Ratio
UK	2 apartments	1,053,770	0.022	0.236	+0.214	10.8×
UK	3–5	180,638	0.032	0.323	+0.291	10.1×
UK	6–10	1,712	0.042	0.364	+0.322	8.7×
UK	11–50	322	0.013	0.048	+0.035	3.7×
NL	2	49,583	0.059	0.080	+0.021	1.4×
NL	3–5	43,715	0.068	0.122	+0.054	1.8×
NL	6–10	21,139	0.067	0.145	+0.077	2.2×

NL	11–50	16,033	0.082	0.170	+0.088	2.1×
NL	51+	3,441	0.131	0.226	+0.096	1.7×
PL	2	46,613	N/A	0.073	–	–
PL	3–5	35,241	N/A	0.145	–	–
PL	6–10	19,642	N/A	0.167	–	–
PL	11–50	13,025	N/A	0.187	–	–
PL	51+	2,144	N/A	0.211	–	–

Source: Own elaboration. The Shapley-to-area ratio is not computable for Poland (2.1.2).

Two patterns are worth noting. First, in the UK, the Shapley/area ratio falls dramatically at building sizes above 10 apartments, from 10.8× at two-flat conversions to 3.7× at 11–50 apartment buildings, because larger UK buildings tend to be purpose-built and contain more uniform apartments, whereas the dominant 2-flat sample comprises Victorian and Edwardian conversions with substantial within-building heterogeneity. Second, the Polish and Dutch within-building Shapley dispersion increases with building size, while the UK pattern reverses. This reflects the different age and composition of the building stocks: Polish prefabricated and Dutch purpose-built blocks are larger and more uniform internally, but larger buildings in those stocks contain proportionally more heterogeneous mixes of efficient and inefficient apartments than smaller blocks do.

instytut
badań
strukturalnych

www.ibs.org.pl