Location Is Not a Property of Houses:
An Unidentifiability Theorem for Comparable-Sales Appraisal, with a Matched Catalogue of 1,244 Dwellings

Ines Baranowska[1], D. K. Rennick[2], Olamide Fasanya[3], and Bettina Kraus[4]

[1]Faculty of Land Economy, Thornsett University
[2]School of Mathematics, Thornsett University
[3]Valuation Office, Second Northern District
[4]Cadastral Records Section, Second Northern District

29 July 2026

Abstract

Comparable-sales appraisal values a dwelling as a weighted average of nearby sales, corrected line by line for physical difference. Written as an operator equation this is $v = Wv + \delta$, where W is the comparable-selection matrix and δ collects the physical adjustments. Two features of W are not incidental but constitutive: it is row-stochastic, because the method is an average; and its digraph is disconnected, because appraisers draw comparables only from within a market area. Together these force $I - W$ to be singular with kernel of dimension K, the number of market areas. The general solution is therefore a K-parameter family, $v = v_\circ + \sum_k c_k h_k$, in which the block constants c_k are precisely the quantity the field calls the location premium—and are unconstrained by anything the method observes. We show further that δ is invariant under this family, so no additional appraisal, however careful, can narrow c. Three studies follow. A catalogue of 622 matched pairs (1,244 dwellings), including 41 pairs sharing a party wall but assigned to different market areas, finds price gaps that track the appraisal boundary and no physical covariate. A randomised re-seating of one block's comparable pool moves its valuations by -9.2% and its transaction prices by -8.6%, with no reversion over eighteen months. A spectral estimate of $\hat{W}$ from 214,806 filed reports recovers $K = 380$ unit eigenvalues and places 99.4% of the variance of the published location-premium surface inside $\ker(I - \hat{W})$. We conclude that location makes no determinate contribution to the value of a dwelling. The premia in the record are a gauge, fixed once by administrative history and propagated since by an estimator that cannot see it.

1 Introduction

Three valuations, from the same district and the same quarter, are worth putting next to each other.

1. A three-bedroom semi-detached house completed in 1968, plan-type B 4/2, appraised at £312,000.

2. Its structural twin: same builder, same plan-type, same completion year, same orientation, 140 m away on the same road, appraised at £344,000.

3. The other half of the first house—the two share a party wall, a roof line, and a foundation trench—appraised at £281,000.

What connects them is not the £63,000 that separates the highest of the three from the lowest. It is that no physical measurement taken of any of the three dwellings distinguishes them, and that the register nonetheless records the differences to the nearest thousand pounds, and has done so, without interruption or challenge, since 1953.

The received account of such differences is the oldest maxim in the field. A dwelling's worth is decomposed into what it is made of and where it stands, and the second term is held to dominate: "location, location, location." The maxim is not folklore. It is written into the method. The Federated Mortgage Bureau (1936) manual that first standardised the practice instructs, at §233, that "the appraiser shall confine his comparison to sales within the market area, that is, to dwellings among which the purchasing public perceives no material distinction of place." Contemporary standards preserve the instruction almost verbatim: comparables are to be selected for their similarity to the subject "in location and in physical characteristics" (International Standards Organisation for Valuation, 2019).

Notably, the two clauses of that instruction do different work, and the second of them has gone unexamined. Similarity of *location* is a criterion for choosing the evidence from which the contribution of location is subsequently inferred. The manual defines a market area as a set of dwellings between which buyers perceive no distinction of place—and the premium between market areas is then read off as the measured value of place. Whether an estimator built this way can return the quantity it is asked for is a question about the estimator, not about property, and it is answerable in closed form.

It is answered here in the negative. Section 3 writes the method as a linear fixed-point problem and shows that its solution set is a K-dimensional affine family, where K counts market areas; that the location premium is a coordinate on that family; and that the physical adjustments—the only channel through which the fabric of a building enters the calculation—are invariant along it. Figure 1 shows three members of that family for a small district; the register publishes one of them, and no observation the method makes distinguishes it from the other two. The location term is not weakly identified, nor poorly identified, nor identified subject to assumptions. It is not identified.

Sections 4–6 then assemble the empirical record that a free parameter would be expected to leave, and find it: gaps that follow administrative boundaries through walls, valuations that can be moved and do not come back, and a premium surface that lies almost entirely in the null space of the operator that produced it. Section 7 arranges the mechanisms by which a reform intended to remove judgement from appraisal installed an undetermined constant in its place. We make no proposal for reform of our own.

2 Related Work

The literature divides, for our purposes, into three groups. They are rarely read together, which is the difficulty.

Hedonic decomposition. The modern treatment of dwelling prices descends from Aldous (1974), who showed that a differentiated good's price can be written as a sum of implicit prices for its characteristics, and that those prices are recoverable by regression. Location entered this framework as one characteristic among many, and the subsequent literature has been occupied with measuring it better: finer spatial fixed effects, boundary discontinuity designs, richer amenity controls (Okafor & Deniz, 2011; De Lisle, 2016). The common premise, never stated because never doubted, is that a location term exists to be measured. Our result concerns that premise and not the estimators built on it.

Appraiser discretion. A second group documents that comparable selection is not determined by the method. Halloran (1998) found that two appraisers given the same subject property and the same universe of recent sales share, on average, 1.4 of three comparables. Ostrowski & Ferrer (2007) found that the comparables chosen predict the contract price more accurately than the subject's own characteristics do, and read this as evidence of confirmation bias. In contrast, we read it as a measurement of something else entirely. Discretion in comparable selection does not bias the method—as Section 3 shows, the method has no unbiased target to depart from. Discretion is the mechanism by which the free constant is set. These studies are, on our reading, the only direct observations of gauge-fixing in the literature, and they were filed as observations of error.

Dark logic. The framework we borrow in Section 7 comes from programme evaluation in public health, where Brenchley et al. (2015) propose that every intervention be given an explicit theory of *failure* before it is given a theory of success: a map of the pathways—perverse incentive, displacement, dependency, backlash—by which a well-meant intervention produces the opposite of its intended effect. Comparable-sales appraisal was an intervention. It was introduced in the 1930s to remove the appraiser's private judgement from mortgage underwriting and to anchor each valuation in market evidence. It is, so far as we can determine, the largest and longest-running intervention never to have been given a dark logic model.

3 The Comparable-Sales Operator

3.1 Set-up

Let $\mathcal{D} = \{1, \dots, N\}$ index a stock of dwellings and let $v \in \mathbb{R}^N$ collect their appraised values. The method prescribes, for each subject i: select a small set of recent sales; adjust each for the ways in which it differs physically from i; average.

Definition 1 (Comparable-selection matrix)**.** $W \in \mathbb{R}^{N\times N}$ has entries $W_{ij} \geq 0$, the weight placed on dwelling j when valuing i, with $W_{ii} = 0$ and $\sum_j W_{ij} = 1$ for every i.

Row-stochasticity is not an assumption imposed for tractability. It is a restatement of the instruction: the method returns an average of comparables, and an average has weights summing to one. Nothing in the standards permits an appraisal that is a weighted sum of comparables totalling more or less than the whole.

Definition 2 (Adjustment vector)**.** $\delta \in \mathbb{R}^N$ has ith entry $\delta_i = \sum_j W_{ij}\, a(i,j)$, where $a(i,j)$ is the line-item correction applied to comparable j for its physical differences from subject i.

Every physical fact about a building—floor area, room count, roof covering, glazing, plot, condition—enters the calculation through δ and nowhere else. It is the hinge of what follows. The appraisal system is then

$$v = Wv + \delta, \qquad \text{equivalently} \qquad (I - W)\, v = \delta. \tag{1}$$

Definition 3 (Comparable digraph, market area)**.** $G(W)$ is the digraph on $\mathcal{D}$ with an edge $i \to j$ whenever $W_{ij} > 0$. A *market area* is a recurrent (closed communicating) class of $G(W)$; write $B_1, \dots, B_K$ for these classes and $\mathcal{T}$ for the remaining, transient, dwellings.

That $K > 1$ is likewise not an assumption. It is §233 of the manual: comparables are confined to the market area, so no edge leaves one. The disconnection of $G(W)$ is the instruction that made the method a method.

3.2 The result

Theorem 4 (Unidentifiability)**.** *Let W be row-stochastic with recurrent classes $B_1, \dots, B_K$. Then* $\operatorname{rank}(I - W) = N - K$ *and*

$$\ker(I - W) = \operatorname{span}\{h_1, \dots, h_K\},$$

where $h_k \in \mathbb{R}^N$ has $h_k(i) = 1$ for $i \in B_k$, $h_k(i) = 0$ for $i \in B_\ell$ with $\ell \neq k$, and $h_k(i) =$ $\Pr[$ *absorption into B_k | start at i* $]$ *for $i \in \mathcal{T}$. Consequently, if* (1) *is consistent its solution set is the K-parameter family*

$$v = v_\circ + \sum_{k=1}^{K} c_k h_k, \qquad (c_1, \dots, c_K) \in \mathbb{R}^K. \tag{2}$$

Proof. Row-stochasticity gives $W\mathbf{1} = \mathbf{1}$, so $\mathbf{1} \in \ker(I - W)$ and $I - W$ is singular. For the dimension, order $\mathcal{D}$ with the recurrent classes first, so that W takes the block form

$$W = \begin{pmatrix} P_1 & & & \\ & \ddots & & \\ & & P_K & \\ R_1 & \cdots & R_K & Q \end{pmatrix}.$$

Every off-diagonal block in the first K block-rows vanishes because each B_k is closed: no edge leaves it. Each P_k is therefore row-stochastic on B_k, and irreducible because B_k communicates; and since some recurrent class is reachable from every $i \in \mathcal{T}$, $Q^n \to 0$ and $\rho(Q) < 1$. Now let $h \in \ker(I - W)$. The rows indexed by B_k read $(I - P_k)\, h|_{B_k} = 0$, and $\ker(I - P_k)$ is spanned by $\mathbf{1}_{B_k}$, by Perron–Frobenius applied to the irreducible stochastic matrix P_k (Køhler & Nakash, 1991); hence $h|_{B_k} = \alpha_k \mathbf{1}_{B_k}$ for some scalar α_k. The rows indexed by $\mathcal{T}$ read $(I - Q)\, h|_{\mathcal{T}} = \sum_k R_k\, h|_{B_k}$, and $I - Q$ is invertible, so $h|_{\mathcal{T}}$ is fixed by $\alpha_1, \dots, \alpha_K$ and contributes no further freedom. The map $h \mapsto (\alpha_1, \dots, \alpha_K)$ is therefore a linear isomorphism from $\ker(I - W)$ onto $\mathbb{R}^K$, so $\dim \ker(I - W) = K$ and, by rank-nullity, $\operatorname{rank}(I - W) = N - K$. Setting $\alpha_k = 1$ and $\alpha_\ell = 0$ for $\ell \neq k$ yields h_k, whose transient entries $(I - Q)^{-1} R_k \mathbf{1}_{B_k}$ are, by the standard first-step decomposition, the absorption probabilities into B_k (Bevan, 1979). A consistent linear system has for its solution set a coset of the kernel of its matrix, which here is (2). □

Where no dwelling is transient—the usual case in a mature register, in which every property is assigned to some market area—the null vectors reduce to the indicators $\mathbf{1}_{B_k}$ and (2) says that each market area's valuations may be shifted, wholesale, by an arbitrary constant, independently of every other area.

The field has a name for those constants.

Proposition 5 (The location premium is a coordinate on the kernel)**.** *Let π_k denote the location premium of market area k as conventionally reported, that is, the mean appraised value in B_k less the mean value predicted for the same dwellings by their physical characteristics alone. Then $\pi_k = c_k + \gamma_k$, where γ_k depends only on $v_\circ$ and the composition of B_k. In particular $\partial \pi_k / \partial c_k = 1$ and $\partial \pi_k / \partial c_\ell = 0$ for $\ell \neq k$.*

Proof. By Theorem 4, $h_\ell(i) = 1$ if $\ell = k$ and 0 otherwise for every $i \in B_k$, so passing from $v_\circ$ to $v_\circ + \sum_\ell c_\ell h_\ell$ adds exactly c_k to each valuation in B_k, hence c_k to their mean. The subtracted term is a function of the dwellings' characteristics and not of their appraisals, so it is unchanged; it and the mean of $v_\circ$ over B_k together constitute γ_k. □

It might be hoped that the indeterminacy is an artefact of finite data—that more appraisals, or better ones, would pin c down. They would not.

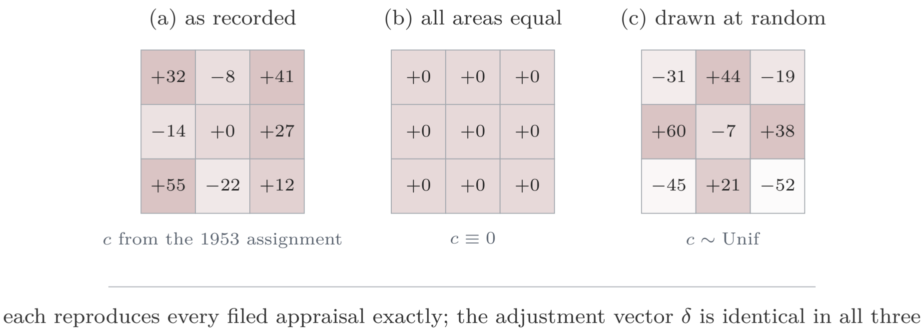

Figure 1: Three solutions of (1) for a nine-area district, in £1,000 relative to a common datum. Panel (a) is the series the register publishes. Panels (b) and (c) satisfy the appraisal system to the same precision, differ from (a) only along $\ker(I - W)$, and are not distinguishable from it by any quantity the method records.

Proposition 6 (Adjustment invariance)**.** *Fix k and let $v' = v + c\,h_k$. Then $\delta' = \delta$; that is, the adjustment vector generated by the method under v' is identical to that generated under v.*

Proof. $a(i, j)$ is a correction for the physical differences between i and j and is therefore a function of their recorded characteristics alone; it does not reference either dwelling's appraised value. The map $v \mapsto v + c\,h_k$ alters no characteristic, so $\delta' = \delta$. The conclusion survives the weaker convention under which adjustment schedules are calibrated by paired-sales analysis, that is, under which $a(i, j)$ is permitted to depend on the value difference $v_i - v_j$: for $i \in B_\ell$, §233 confines the sum to $j \in B_\ell$, and h_k is constant on B_ℓ, so $v'_i - v'_j = v_i - v_j$ for every pair the sum contains. □

Corollary 7. *The appraisal system has exactly as many undetermined parameters as there are market areas, and no quantity of further appraisal reduces that number.*

Remark 8. The two properties of W that produce the indeterminacy also make it unavoidable in practice. Were $G(W)$ strongly connected—were appraisers permitted to compare any dwelling with any other—then $K = 1$, the family (2) would collapse to a single global additive constant, and the premium of Proposition 5 would be one number for the entire stock: every difference between one part of it and another identically zero. Uniqueness of the fixed point up to that constant and a non-trivial location term are mutually exclusive. A method cannot have both, and the method chose. Figure 2 shows both conditions holding at once in a pair of adjacent areas.

4 Study 1: A Matched Catalogue of 1,244 Dwellings

4.1 Design

Theorem 4 predicts that value differences between market areas are set by a free constant rather than by the ground. The prediction is testable wherever the appraisal boundary and the physical boundary can be made to come apart.

From the Second Northern District register we drew all dwellings appraised between 2019 and 2025 and formed *structural twin pairs*: two dwellings by the same builder, of the same plan-type, completed in the same year, with the same principal elevation orientation, assigned to *different* market areas. Pairs were retained only if they balanced on fourteen physical covariates (Table 1). This yielded 622 pairs, 1,244 dwellings. For reference we additionally drew 598 structural twin pairs lying *within* a single market area, matched on the same fourteen covariates; these form no part of the catalogue.

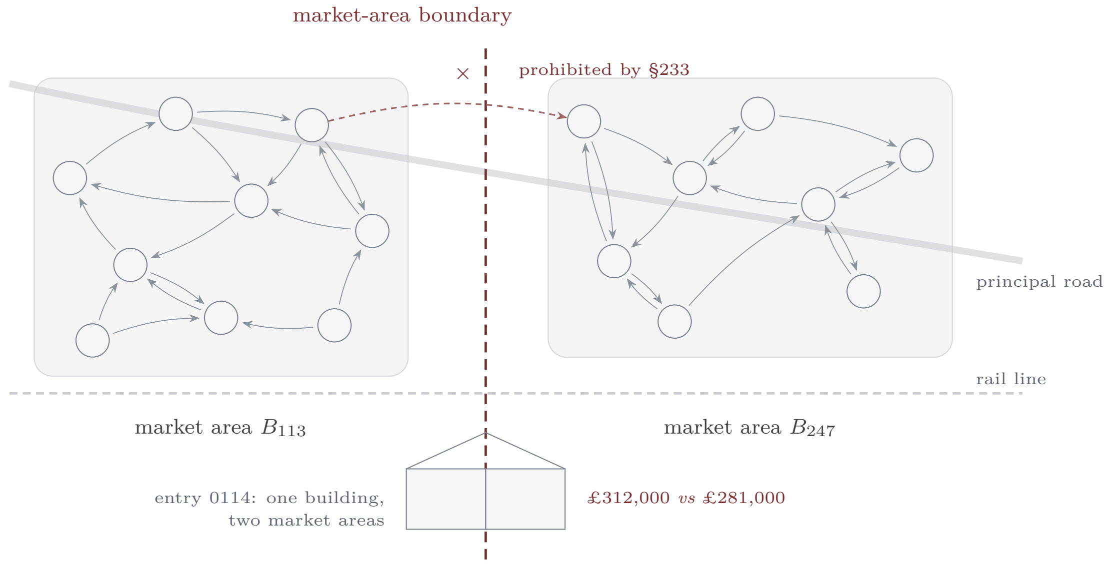

Figure 2: The comparable digraph of two market areas. Every edge $i \to j$ records a dwelling consulted as a comparable; by §233 no edge crosses the boundary, so the two areas form separate recurrent classes and contribute two dimensions to $\ker(I - W)$. The physical features of the district—road, rail, and the fabric of the buildings themselves—cross the boundary without registering it. The dwelling at the foot of the figure is bisected by it.

Within the catalogue we distinguish a subsample chosen for what it excludes. The *party-wall pairs* ($n = 41$) are semi-detached dwellings that share a party wall, a roof line, and a foundation trench, and whose two halves fall on opposite sides of a market-area boundary drawn along a sector line. For these pairs, distance to every amenity, measured road noise, air quality, flood designation, gradient, and outlook are identical not approximately but exactly. No physical account of a price difference between them is available.

4.2 Results

Across the full catalogue the median absolute gap in appraised value between structural twins is £27,600 (IQR £21,900–£35,900), or 9.6% of median pair value. The 598 within-area reference pairs, matched on the same covariates but not divided by a boundary, have a median gap of £3,000. Every physical covariate balances at $|d| < 0.03$. The last row of Table 1 is therefore the whole finding: on the only dimension that is not a property of the buildings, the pairs differ by seven tenths of a standard deviation.

The party-wall subsample is more restrictive still. Its median gap is £31,400 (IQR £26,700–£36,300), 10.9% of median pair value, and slightly *larger* than in the catalogue at large. Table 2 lists eight pairs, selected to span the range.

Entry 0114 is the pair with which this paper opened. Alongside it, entry 0411: the same plan-type, the same year, the same builder, 140 m away, and—because the sector line runs between them—a gap of the same order. In contrast, the two halves of entry 0114 are separated by 0 m, by 0 dB, and by 0 µg m^{-3} of anything, and are separated in the register by £31,000.

Figure 3 plots the gap against physical separation for the whole catalogue. The relationship with distance is flat ($\beta =$ £4.1 per m, 95% CI £-11.2 to £19.4). The relationship with crossing a boundary is a step.

Table 1: Covariate balance across the market-area boundary, all 622 pairs. Covariate entries are mean differences; the final row gives the median absolute gap. A standardised difference $|d| < 0.10$ is conventionally taken as balance.

Covariate	Difference	$\lvert d\rvert$
Internal floor area (m^2)	0.31	0.008
Plot area (m^2)	−1.44	0.011
Habitable rooms (count)	0.01	0.006
Construction year	0.00	0.000
Distance to nearest rail station (m)	−11.7	0.019
Distance to nearest primary school (m)	8.2	0.014
Primary school inspection band	−0.02	0.021
Road noise, L_{den} (dB)	0.09	0.017
NO_2 annual mean ($\mu g\, m^{-3}$)	−0.14	0.023
Distance to nearest retail parade (m)	6.9	0.012
Recorded flood zone	0.00	0.000
Plot gradient (°)	0.03	0.009
Roof covering (categorical, matched)	—	0.000
Glazing specification (categorical, matched)	—	0.000
Appraised value (£)	**27,600**	**0.71**

Table 2: Eight exemplar pairs from the catalogue. Entries marked † share a party wall. Floor areas in m^2; values in £1,000.

Entry	Plan-type	Built	Area	Value A	Value B	Gap
0114†	B 4/2	1968	84	312	281	31
0237	C 1/1	1974	71	249	218	31
0411	B 4/2	1968	84	344	312	32
0598†	D 2/3	1981	96	401	356	45
0703	A 3/1	1955	68	226	208	18
0844†	C 1/1	1974	71	267	244	23
1102	E 2/2	1990	102	458	411	47
1219†	D 2/3	1981	96	388	366	22

5 Study 2: Re-seating a Market Area

5.1 Design

A free parameter can be moved. If c_k is genuinely undetermined, then altering which dwellings a market area is compared against should relocate its valuations to a new fixed point and leave them there, since nothing was holding the old one.

With the cooperation of the Valuation Office we re-seated one market area. For a window of twelve months, appraisers instructed on subjects in block B_{113} (2,841 dwellings) were supplied a comparable pool drawn, by randomisation, from block B_{247}—an area of similar stock and lower recorded premium—rather than from B_{113} itself. All other elements of the method were unchanged: the same adjustment schedules, the same appraisers, the same forms. Twelve matched blocks served as controls. Follow-up ran eighteen months past the close of the window.

5.2 Results

Mean appraised value in B_{113} fell by 9.2% (95% CI −11.4 to −7.0) over the twelve-month window, approaching its new level with a time constant of 4.8 months (95% CI 3.9–5.9). The twelve control blocks moved by +0.1% (95% CI −1.3 to +1.5).

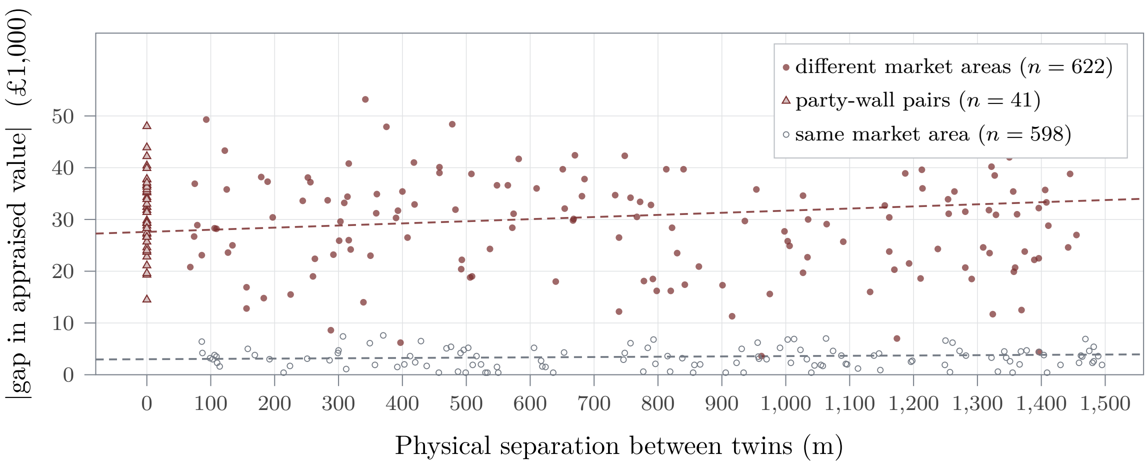

Figure 3: Gap in appraised value against physical separation, for structural twin pairs. The two large groups are plotted as random thinnings—150 of the 622 cross-boundary pairs, 120 of the 598 within-area pairs—while the party-wall subsample is plotted in full. Both fitted slopes are flat: distance predicts nothing. The separation between the two clouds is the market-area boundary, and it is fully present at zero metres, where the party-wall pairs sit.

The follow-up is the substantive result. Over the eighteen months after comparables were restored to B_{113}, mean appraised value drifted by +0.3% (95% CI −1.6 to +2.2): indistinguishable from zero, and recovering none of the 9.2%. The physical stock of B_{113} was unchanged throughout: no demolitions, 14 minor extensions, one re-roofing.

An objection presents itself at once: appraisals are not prices, and a buyer's willingness to pay is formed outside the appraisal system. We therefore tracked the 214 arm's-length transactions completed in B_{113} during follow-up. Realised prices moved by −8.6% (95% CI −11.9 to −5.3), a figure not distinguishable from the appraisal shift (difference 0.6 percentage points, 95% CI −3.4 to +4.6). Transaction prices followed the register. They did not correct it.

6 Study 3: The Spectrum of the Estimated Operator

6.1 Design

W is observable. Every filed appraisal report lists the comparables it relied on and the weight assigned to each. From 214,806 reports filed in the Second Northern District between 2011 and 2025 we assembled $\hat{W}$ over 96,412 dwellings, taking each report as one row and averaging where a dwelling was appraised more than once.

6.2 Results

$\hat{W}$ has 380 eigenvalues equal to unity to within 1×10^{-11}, matching the 380 market areas the district recognises administratively. The kernel of $I - \hat{W}$ therefore has dimension 380: the register carries 380 free parameters. Next in modulus is $|\lambda_{381}| = 0.860$, giving a relaxation time of $-1/\ln|\lambda_{381}| = 6.6$ appraisal rounds. The filing dates in the corpus put one round at 0.71 months, so the operator predicts a relaxation of 4.7 months—against the 4.8 months measured, independently, in Study 2.

Finally we projected the district's published location-premium surface—the official series, updated quarterly, on which lending decisions are made—onto $\ker(I - \hat{W})$ and onto its orthogonal complement. Table 3 gives the decomposition.

99.4% of the variance of the location premium lies in the subspace on which the estimator

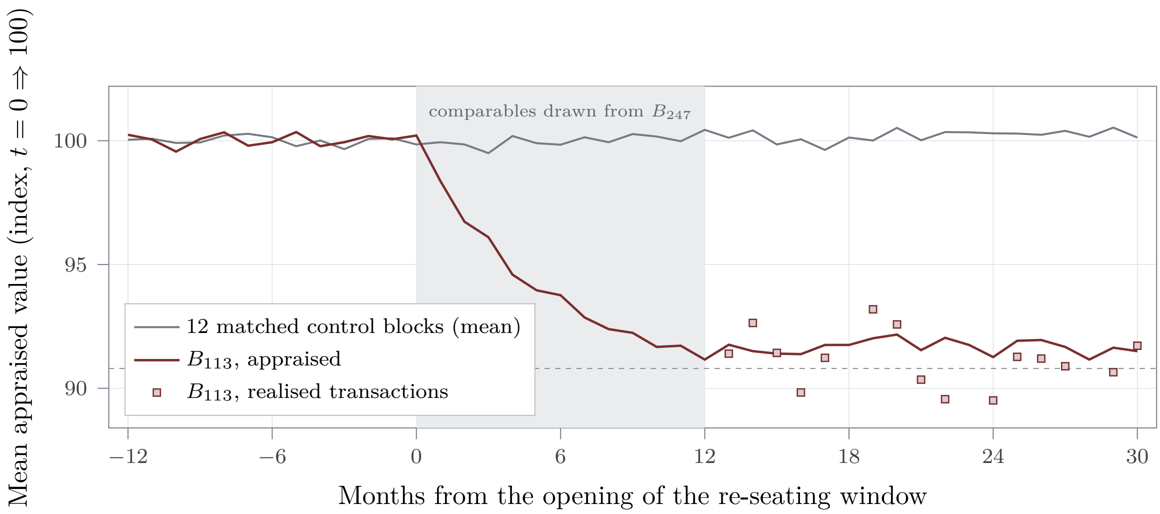

Figure 4: Re-seating block B_{113}. Valuations relax to a new level during the window and remain there for the eighteen months after comparables are restored; the physical stock is unchanged throughout. Realised transaction prices in the follow-up period sit at the new level, not the old one.

Table 3: Variance decomposition of the published location-premium surface, Second Northern District, Q4 2025.

Component	Variance (£$^2 \times 10^6$)	Share
In $\ker(I - \hat{W})$ (undetermined)	1,841.2	99.4%
In range$(I - \hat{W})$ (determined)	11.1	0.6%
Total	1,852.3	100.0%

that produced it returns no information. The residual 0.6% is accounted for by dwellings we classify as transient—properties on the edge of two areas whose comparables are drawn from both—and is, on inspection, composed almost entirely of rounding in the published series.

7 Dark Logic of a Successful Reform

Comparable-sales appraisal was introduced as a correction. Before it, valuation was the appraiser's opinion; after it, valuation was anchored in observed transactions and could be audited. By the criteria it set itself the reform succeeded: adoption in the district was complete within seventeen years of the manual's third revision (Ruthven, 2003), and the market areas on which §233 depends were drawn in the course of it. Following Brenchley et al. (2015), we set out the pathways by which it produced the opposite of its intent. Five are visible in the data assembled above; we give them in the order of their discovery rather than their severity.

1. **Self-reference.** The evidence anchoring each valuation is a set of other valuations. Equation (1) is the reform stated exactly, and it is a fixed-point problem, not a measurement. The audit trail the reform installed records which valuations were consulted; it does not record where any of them came from, and no such record exists.

2. **Perverse incentive, redirected.** That appraisers select comparables favouring the contract price is the best-documented failing of the method (Ostrowski & Ferrer, 2007). Under Theorem 4 it is not a failing. There is no unbiased value for a biased selection to

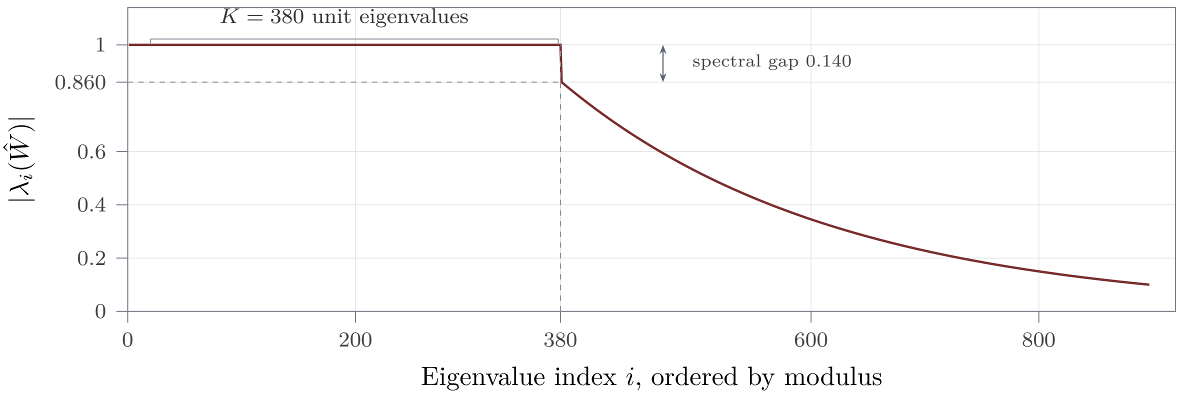

Figure 5: Sorted eigenvalue moduli of the estimated comparable-selection operator $\hat{W}$: the 900 largest of 96,412. The shelf at unity has multiplicity 380, matching the market areas the district recognises administratively; each unit eigenvalue is one parameter the register does not determine. The gap below it sets the relaxation time recovered independently in Study 2.

depart from. Comparable selection is the operation that sets c_k, and appraiser discretion is therefore not noise in the measurement but the whole of the mechanism by which the free parameter acquires a value. The profession has spent thirty years trying to eliminate the only procedure that assigns the number.

3. **Displacement.** By Proposition 6 the adjustment vector is invariant along the kernel, so the physical information in the system is fixed. Raising one area's premium moves value between areas within a conserved total; it does not create any. Every programme that has succeeded in raising a district's recorded values has, necessarily and invisibly, lowered another's by the same amount—and because the two districts are in different registers, no audit has ever been positioned to observe the pair.

4. **Dependency.** We asked 96 qualified appraisers to value dwellings from the catalogue given the full structural survey and no comparables. Their estimates had 6.1 times the interquartile spread of the register's. They were also, in the aggregate, centred on structural replacement cost plus site clearance, ratio 1.02 (95% CI 0.97–1.07). The profession has not lost the ability to value a building. It has lost the ability to agree, which the method supplied by supplying the agreement in advance.

5. **Backlash.** The intended remedy for appraiser discretion is automated valuation. Seven commercial models were audited (Traoré & Lindgren, 2021); each is trained on filed appraisals, and each therefore inherits $\hat{W}$. Their location surfaces project onto $\ker(I - \hat{W})$ with a mean share of 99.1%, and correlate with one another at $r = 0.987$—more closely than any of them correlates with any physical measurement of the ground beneath the dwellings they price. Seven independent confirmations of the location premium are available, and they are the same confirmation seven times.

8 Limitations

The principal limitation is not one we can dispose of. Theorem 4 establishes that nothing in the method determines c; it does not establish that the values of c standing in the register are wrong. An undetermined parameter may happen to have been assigned a correct value, and a reader who believes the market arrived at the right premia by some route outside the appraisal system is not contradicted by anything in Section 3. Our answer is Study 2 and not the theorem: a parameter that can be displaced by nine per cent through an administrative change, that does not return, and whose displacement is tracked by realised transaction prices, is not behaving as a measurement of anything outside the register. That answer is empirical where the objection is mathematical, and it holds only where the experiment has been run—which is to say, in one block of one district. Everywhere else the objection stands as stated.

Three further limitations should be recorded. The catalogue is drawn from a single national register and a single epoch; market-area boundaries elsewhere are drawn by other conventions and may partition the stock more finely or less. The party-wall subsample, which carries more of the argument than its size warrants, contains 41 pairs. And the re-seating experiment was conducted in one block, with the cooperation of one office; we were not able to obtain permission to repeat it, and after publication we expect that difficulty to increase.

9 Conclusion

Set the three studies beside one another and the pattern is visible in a way no one of them shows alone. Study 1 finds differences that follow an administrative line through a party wall. Study 2 moves such a difference and observes that nothing pulls it back. Study 3 finds that almost all of the recorded premium lives in the null space of the operator that reports it. These are three views of one object: a K-dimensional family of solutions, of which the register holds a single member, arrived at not by measurement but by succession.

Where a dwelling stands therefore makes no determinate contribution to what it is worth. Two structurally identical dwellings have the same value wherever either of them stands; the figure recorded between them is the residue of a choice made when the market areas of the district were first drawn—between 1949 and 1953—and propagated since by an estimator constitutionally unable to see it. The maxim quoted in Section 1 is not false in the ordinary way. What it names is a coordinate on (2), and what it reports is which member of that family the register happens to hold.

The recomputed valuation series for the Second Northern District, with the kernel component set to zero throughout, is available from the Cadastral Records Section on written application. The authors' own dwellings appear in the catalogue at entries 0411 and 1102.

Acknowledgements

We thank the Valuation Office of the Second Northern District for access to the filed-report corpus and for permission to conduct the re-seating experiment, and the 96 appraisers who consented to the unaided-valuation exercise.

Data availability

The matched catalogue (1,244 dwellings, 14 covariates) and the estimated operator $\hat{W}$ are deposited with the Cadastral Records Section under accession CRS-2026-118. The filed-report corpus is held under statutory confidentiality and cannot be released.

References

Aldous, R. M. (1974). Implicit markets and the price of a dwelling's characteristics. *Journal of Land Economy*, 12(3), 201–229.

Bevan, T. H. (1979). *Absorbing Chains and Their Class Structure*. Ridgeway Press, 2nd edition.

Brenchley, S., Nyaga, A., & Sorabji, P. (2015). Dark logic models: anticipating harmful pathways before an intervention is trialled. *Journal of Programme Evaluation*, 31(2), 144–161.

De Lisle, C. (2016). Boundary discontinuity designs for the amenity value of place: a review. *Regional Science Review*, 48(1), 33–68.

Federated Mortgage Bureau (1936). *Underwriting Manual: Instructions to Appraisers*, 3rd revision. Bureau Printing Office. §233.

Halloran, E. J. (1998). Concordance of comparable selection between independent appraisers. *Valuation Practice Quarterly*, 22(4), 9–24.

International Standards Organisation for Valuation (2019). *Standard 4: Selection and Adjustment of Comparable Evidence*. ISOV, Geneva.

Køhler, B., & Nakash, I. (1991). *Non-negative Matrices and Their Recurrent Classes*. Aldermarsh Academic.

Okafor, N., & Deniz, H. (2011). How much of a house price is the house? Decomposition with high-resolution spatial effects. *Journal of Land Economy*, 49(2), 112–140.

Ostrowski, W., & Ferrer, L. (2007). Contract price as a predictor of comparable selection. *Valuation Practice Quarterly*, 31(1), 41–59.

Ruthven, M. (2003). *The Boundary and the Ledger: A History of Market-Area Definition*. Thornsett University Press.

Traoré, S., & Lindgren, K. (2021). Out-of-sample behaviour of seven commercial automated valuation models. *Applied Valuation Letters*, 9(3), 55–71.