Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Can TDR Policy Reshape Property-Price Covariance? Development-Rights Allocation and Network Reconfiguration through Development-Saturation Similarity in Seoul

Dahoon Kim

zealot8100@naver.com

Korea Institute of Local Finance

Research Article

Keywords: transferable development rights, development saturation, spatial weight matrix, property-price covariance, null benchmark, Seoul

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11698100/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Can TDR Policy Reshape Property-Price Covariance?

Development-Rights Allocation and Network Reconfiguration through Development-Saturation Similarity in Seoul

Abstract

Can transferable development rights (TDR) reshape property price covariance? This pape r shows how alternative allocations of development capacity reorganize modeled price de pendence through development saturation. A capacity-conserving transfer changes saturati on, which defines a similarity matrix among buildings of comparable size. A spatial erro r model links this matrix to price covariance. Using 6,413 Seoul commercial and mixed-use property transactions, the analysis compares equal-volume allocations that vary sende r and receiver selection. The results show that although the extent of covariance reconfig uration and the counterpart relationships affected vary with sender and receiver selection, covariance changes outside receiving areas exceed those within them. Despite this reorg anization of relationships among properties, absolute changes in aggregate uncertainty re main small. The main geographical and counterpart contrasts across allocation scenarios persist when saturation is measured using above-ground floor area, supporting the robust ness of the findings.

Keywords: transferable development rights; development saturation; spatial weight matri x; property-price covariance; null benchmark; Seoul

1. Introduction

Seoul has been discussing transferable development rights (TDR) to move unused development capacity from conservation-constrained sites to locations with development potential (Seoul Metropolitan Government, 2025). Can equal-volume transfers reshape property-price covariance differently depending on the choice of sending and receiving properties? This paper extends TDR appraisal to the relationships among properties that capacity allocation may reorganize. The framework links capacity allocation to development saturation, property similarity and modeled covariance. With existing buildings unchanged, receiving rights lowers the ratio of registered floor area to nominal capacity, while relinquishing rights raises it. A spatial error model combines saturation-based similarity among comparably sized buildings with geographic weights. Covariance refers to within-year residual log-price levels after controlling for property characteristics and district-by-year effects. Using 6,413 commercial and mixed-use transactions in Seoul during 2018–2025, the analysis compares six equal-volume allocations, supported by randomized benchmarks, historical validation and sensitivity checks. Concentrating receivers in Gangnam increases its share of absolute covariance change from 20.1% to 30.7%, although most change remains elsewhere. Selecting receivers nearest the district center adds no detectable geographical concentration relative to random eligible receivers within Gangnam, but produces greater receiver–undesignated reconfiguration. Sender selection also affects the overall extent of change. These relational shifts coexist with small aggregate movements: across all four receiving rules, median annual changes in the standard deviation of fixed-group average residuals remain within ±1%. Equal-volume allocations thus differ in where modeled covariance changes occur and which property relationships they involve.

2. Related research and contributions

2.1 TDR appraisal: adding a relational outcome to program evaluation

To TDR appraisal this paper adds an outcome that existing evaluation frameworks do not report, the reorganization of modeled price dependence among selected sites and their peers, with a comparison design that holds transferred volume and participant counts fixed across allocation rules. The appraisal literature establishes that program performance depends on the development opportunities created at receiving sites and the restrictions imposed at sending sites. Machemer and Kaplowitz (2002) organize program characteristics into an evaluation framework; Walls and McConnell (2007) examine design, implementation and outcomes across US communities; McConnell and Walls (2009) review the use of transferable rights to reconcile preservation with development. Receiving demand is central, because a nominal right has value only where a developer can use it: Pruetz and Standridge (2009) identify demand for additional development and the design of receiving areas among the recurring conditions of successful programs, and McConnell, Walls and Kopits (2006) connect zoning and TDR to development density. Markets also depart from

an idealized exchange of homogeneous rights: Linkous (2016, 2017) emphasizes restructu ring for development incentives, timing, information, land heterogeneity and limited parti cipation; Shahab, Clinch and O'Neill (2018) estimate the transaction costs borne by parti cipants in four Maryland programs; Bruno et al. (2023) extend the agenda to ecosystem services and receiving-area consequences. These studies evaluate programs by their devel opment outcomes and costs and agree that sending and receiving choices must be assess ed jointly. Among the studies reviewed here, none asks which other properties an allocat ion ties to the selected sites in a price model, and none compares allocation rules at a c ommon transfer volume; this paper does both.

2.2 Valuation: from the price level of entitlements to the dependence they create

To valuation research it adds a use of entitlement information in the covariance structure rather than in the mean equation, where that research has placed it. Rosen's (1974) hed onic framework relates differentiated attributes to market values; Cheshire and Hilber (20 08), Hilber and Vermeulen (2016) and Glaeser and Gyourko (2018) show how supply re strictions capitalize into prices and supply responses. That literature asks how developme nt conditions relate to an individual property's expected price. Saturation uses the same e ntitlement information to describe a relative development position, distinguishing parcels with substantial headroom from those at or beyond the benchmark, and conditioning on building size recognizes that an equally saturated small shop and a large commercial buil ding need not face comparable conditions. The object here is different: which residual de viations tend to move together once observed characteristics and local annual price levels are removed, and how that pattern changes when a policy changes the entitlement comp onent of saturation.

2.3 Non-geographic proximity: a policy-driven weight matrix at parcel level

To the spatial-weights literature it adds a matrix one of whose ingredients is the quantity a planning instrument reallocates, so that a policy scenario implies a rule-based change to the matrix rather than an analyst's choice. Economic relationships need not coincide with physical distance: Bathelt and Glückler (2003) develop a relational view of econom ic geography and Boschma (2005) separates geographical from organizational, cognitive and other proximities. Spatial econometrics supplies the tools: Dubin (1988) treats spatial ly autocorrelated errors, Pinkse, Slade and Brett (2002) use flexible representations of m arket relationships, Seya, Yamagata and Tsutsumi (2013) show that selecting the weight matrix is a substantive modeling choice in hedonic work, Elhorst, Lacombe and Piras (2 012) address specification with several matrices, and Kelejian and Piras (2014) treat end ogenous weighting matrices. Property research has begun to use non-geographic relations hips directly. Bailey et al. (2018) trace housing-market effects through social networks. Y ang, Gao and Li (2025) is the closest antecedent: they construct a socioeconomic distanc e between cities, show that house-price covariance declines with it, and use the covarian

ce to explain homeownership decisions. Their relationship is between cities and is taken as given, so the link between non-geographic similarity and price covariance is itself established. What the present paper adds is narrower: the relationship is built at parcel level from nominal development capacity, the quantity a planning instrument reallocates, and the question is how alternative uses of that instrument at equal volume reorganize it. No new estimator is proposed; the contribution is the mapping of a capacity-conserving allocation into the matrix and the comparison of equal-volume allocations on the covariance they imply.

2.4 Benchmarks: separating placement from selection

A change to any node attribute in a nearest-neighbor graph rewires links, so reporting that an allocation "reconfigures covariance" is uninformative unless one can say how much of the reconfiguration, and how much of its geography, follows from the mere placement of participants. The benchmark design is therefore the second contribution. A mechanical exposure rule says what the district shares would be under a simple accounting of capacity changes; randomized saturation perturbations say what any change of the same magnitude would do with and without participant locations preserved; and a matched draw of random eligible receivers inside the receiving district, under the same senders, volume and conservation constraint, isolates what the within-district selection rule adds. Together these separate two levels of allocation choice, where to place receivers and how to select them within that area, which the results of Section 4 show to have different consequences. Historical basket validation plays a supporting role: it checks whether the baseline saturation representation carries predictive information under spatially blocked cross-validation (Roberts et al., 2017) and whether cross-property covariance matters when prices are aggregated, scored by proper rules (Gneiting and Raftery, 2007).

3. Pathway, data and design

3.1 Capacity transfers change saturation and similarity

For parcel i, let A_i denote lot area, F_i registered gross floor area, E_i the standardized entitlement ratio and q_i the allocated capacity in square metres, positive for receiving and negative for sending. Saturation under allocation q is

$$s_i(q) = \frac{F_i}{A_i E_i + q_i} \tag{1}$$

with $\Sigma_i\, q_i = 0$ across parcel representatives. Holding F_i fixed, receiving lowers saturation and sending raises it, so a receiver moves away from highly saturated peers and toward parcels with more headroom even though nothing has been built. This index then defines the saturation matrix $W_s(q)$, which compares properties within the same year and among buildings of comparable size. Each observation's candidate set is the nearest tenth of that year's observations in log floor area, rounded down and no smaller than 40; the eight

candidates closest in saturation receive equal weights of 1/8, with self-links excluded. A separate geographic matrix W_e connects the eight nearest same-year observations in projected metres. W_e and the size candidate sets are held fixed across allocations. Saturation entering W_s is clipped at the baseline 99th percentile (4.739), with the same cutoff before and after allocation, so changes that stay above the cutoff do not alter weights.

3.2 The spatial error model translates W into covariance

For log transaction price per unit of lot area, $y = \log(P/A)$, the mean contains Lasso-selected property controls and district-by-year effects. Residuals u follow a two-weight spatial error process,

$$u = \lambda_e W_e u + \lambda_s W_s(q)\, u + \varepsilon, \qquad \varepsilon \sim N\left(0, \sigma^2 I\right) \tag{2}$$

$$\Sigma(q) = \sigma^2 \left[\left(I - \lambda_e W_e - \lambda_s W_s(q)\right)^{\mathsf{T}} \left(I - \lambda_e W_e - \lambda_s W_s(q)\right) \right]^{-1} \tag{3}$$

The policy comparison is $\Sigma(q) - \Sigma(0)$. Because Equation 3 inverts, a change at designated parcels propagates through indirect paths to the modeled covariance of undesignated ones. If $\lambda_s = 0$ the allocation has no effect. The calculation is a conditional scenario: buildings, mean coefficients, dependence parameters and innovation variance are held fixed while W_s is fully recomputed. Figure 1 summarizes the pathway and the quantities held fixed.

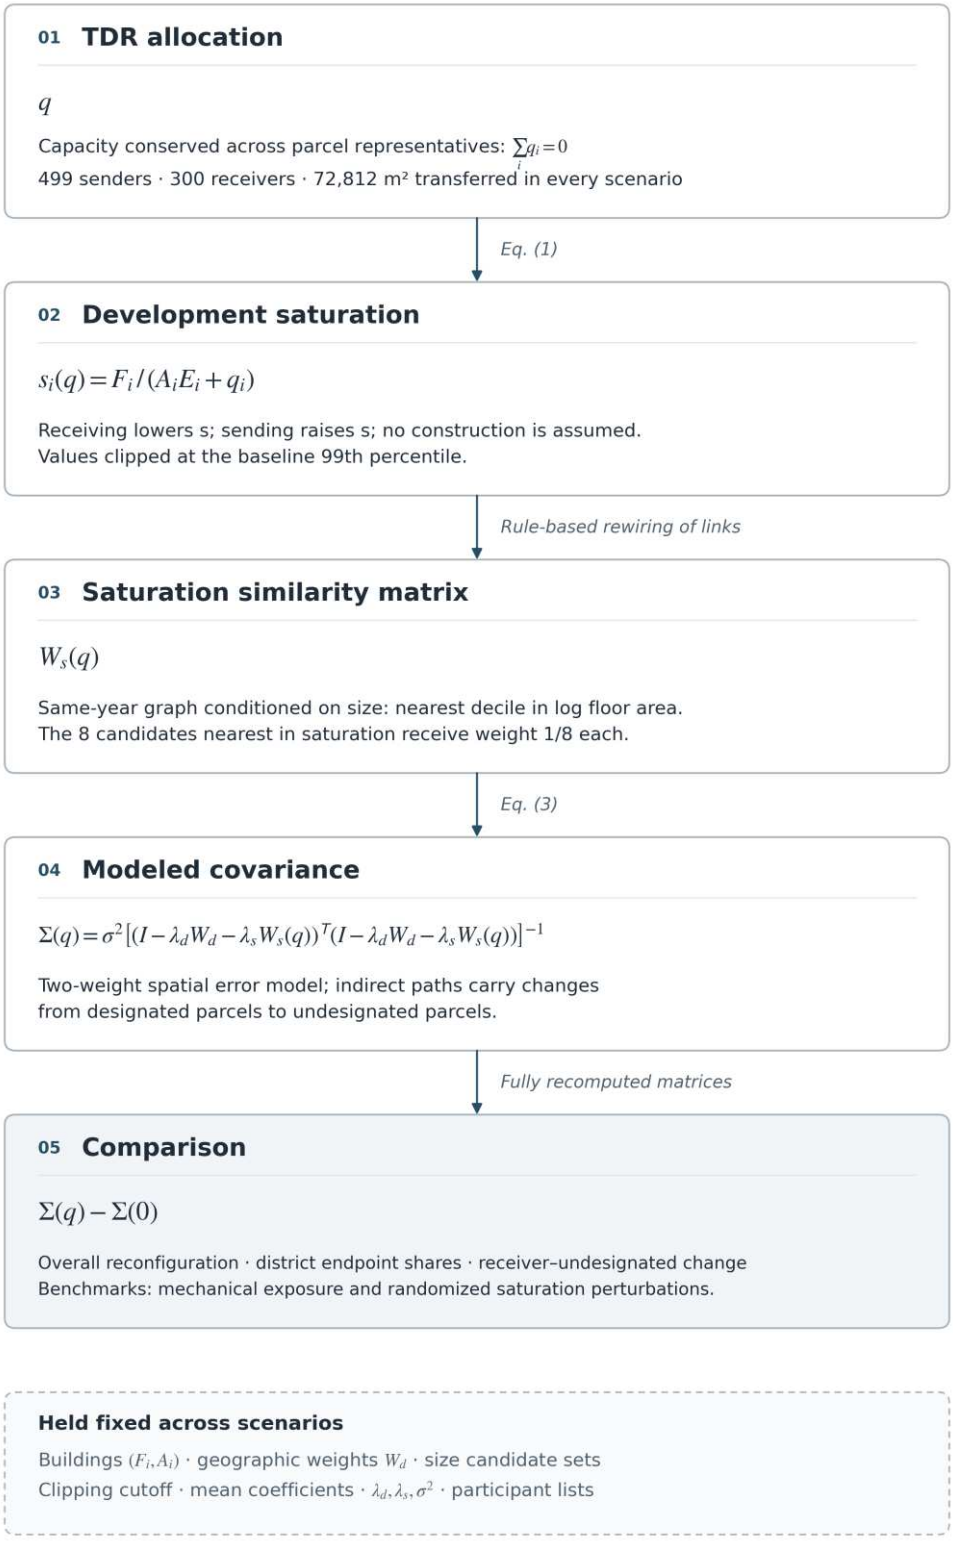


Figure 1 Policy pathway from TDR allocation to modeled covariance. Arrows are the three step s of Sections 3.1–3.2; the dashed box lists the inputs held fixed while W_s is recomputed. Each s cenario is compared with the baseline through $\Sigma(q) - \Sigma(0)$ using the measures of Section 3.5 and the benchmarks of Section 3.6.

3.3 Data and the saturation distribution

The sample contains 6,413 Seoul commercial and mixed-use transactions during 2018–2025 after excluding three records without coordinates, representing 5,627 parcels. Prices include land and structures. Registered gross floor area defines scale and enters saturation; lot area normalizes price and nominal capacity. Registered gross floor area includes below-ground space because the scale of an existing building is not fully represented by its above-ground area. Basements form part of the constructed asset and entail investment in excavation, structural works and building services. Excluding them would omit part of the building's physical scale and capital investment when identifying comparable properties. The similarity matrix therefore first restricts comparisons to buildings of similar total floor area and then identifies neighbors by saturation. Using total floor area in the saturation numerator describes the extent of existing construction relative to standardized nominal capacity; it is not intended to measure the proportion of legally permitted floor area already used. The above-ground specification provides an alternative measure and tests whether the allocation comparisons depend on including below-ground space. Nominal capacity is lot area multiplied by a standardized ratio by zone from the Seoul Ordinance on Urban Planning (Article 55): 10 for central commercial, 8 for general commercial, 6 for neighborhood commercial, 4 for semi-residential and semi-industrial, 2.5, 2.0 and 1.5 for general residential III, II and I, 1.0 and 1.2 for exclusive residential I and II, and 0.5 for green and restricted categories. Online Resource A gives the sample construction.

Within this sample, Table 1 describes the resulting index. The median transaction is at 1.03 times nominal capacity and 52.1% of transactions exceed it. The excess is concentrated in residential zones, where 63.2% of transactions are above one, and is rare in commercial zones (18.7%), whose higher ratios leave most parcels with headroom. The pattern reflects the inclusion of below-ground floors in registered gross floor area and the omission of site-specific incentives, district unit plans and legacy permissions from the standardized ratio (Section 6). An alternative that counts only above-ground floor area, which the data record as a proportion of gross area, has a median of 0.86 and places 26.2% of transactions above one; it ranks parcels almost identically (Spearman 0.95, log-scale correlation 0.92), retains every sending-pool parcel and 83% of receiving-pool parcels, but reproduces only 44% of the baseline eight-nearest-neighbor saturation links (Online Resource A5). The graph is therefore sensitive to the measure even where the ranking is not; Section 4.4 re-estimates the model and the policy comparison with the alternative measure. The sending pool, parcels whose maximum observed raw saturation is below 0.8, contains 1,962 parcels; the disjoint receiving pool, parcels whose representative saturation is at least 0.8, contains 3,665.

Table 1 Development saturation in the Seoul sample, by zone group

Zone group	Transactions	P25	Median	P75	Share > 1 (%)	Share < 0.8 (%)
Commercial (central, general, neighborhood)	1,055	0.26	0.45	0.80	18.7	75.0
Semi-residential and semi-industrial	665	0.42	0.71	1.03	27.2	59.9
Residential and other	4,693	0.86	1.17	1.47	63.2	22.2
All	6,413	0.60	1.03	1.39	52.1	34.8

Raw saturation, registered gross floor area / (lot area × standardized ratio), for the 6,413 transactions with coordinates. The baseline 99th percentile is 4.739 and the clipped mean is 1.101. Gangnam-gu accounts for 1,357 transactions (21.2%).

Turning to estimation, the primary specification, H1, includes district-by-year effects in both Lasso selection and the spatial refit, so it examines residual relationships beyond local annual price levels. Lasso selects 22 of 118 property and land controls; dependence parameters are then estimated jointly with the selected mean by concentrated maximum likelihood. H2, with year effects only, is retained as a comparison. Historical validation uses five outer spatial folds of 1 km cells that keep each parcel's repeated transactions together; held-out prices never enter estimation. The primary evaluation unit is a same-year basket of 20 held-out parcels similar in floor area and saturation, with at most two parcels per district; there are 1,200 overlapping baskets. Estimation and validation settings are in Online Resource B.

The baseline saturation graph can link properties across Seoul (Figure 2). Two 2024 transactions were selected to illustrate long-distance links. The eight saturation neighbors are 7.4–25.7 km from the focal property in Songpa-gu and 6.4–22.0 km from the one in Gangseo-gu. The eight geographic neighbors lie within 2.34 km and 1.36 km of the same properties, respectively.

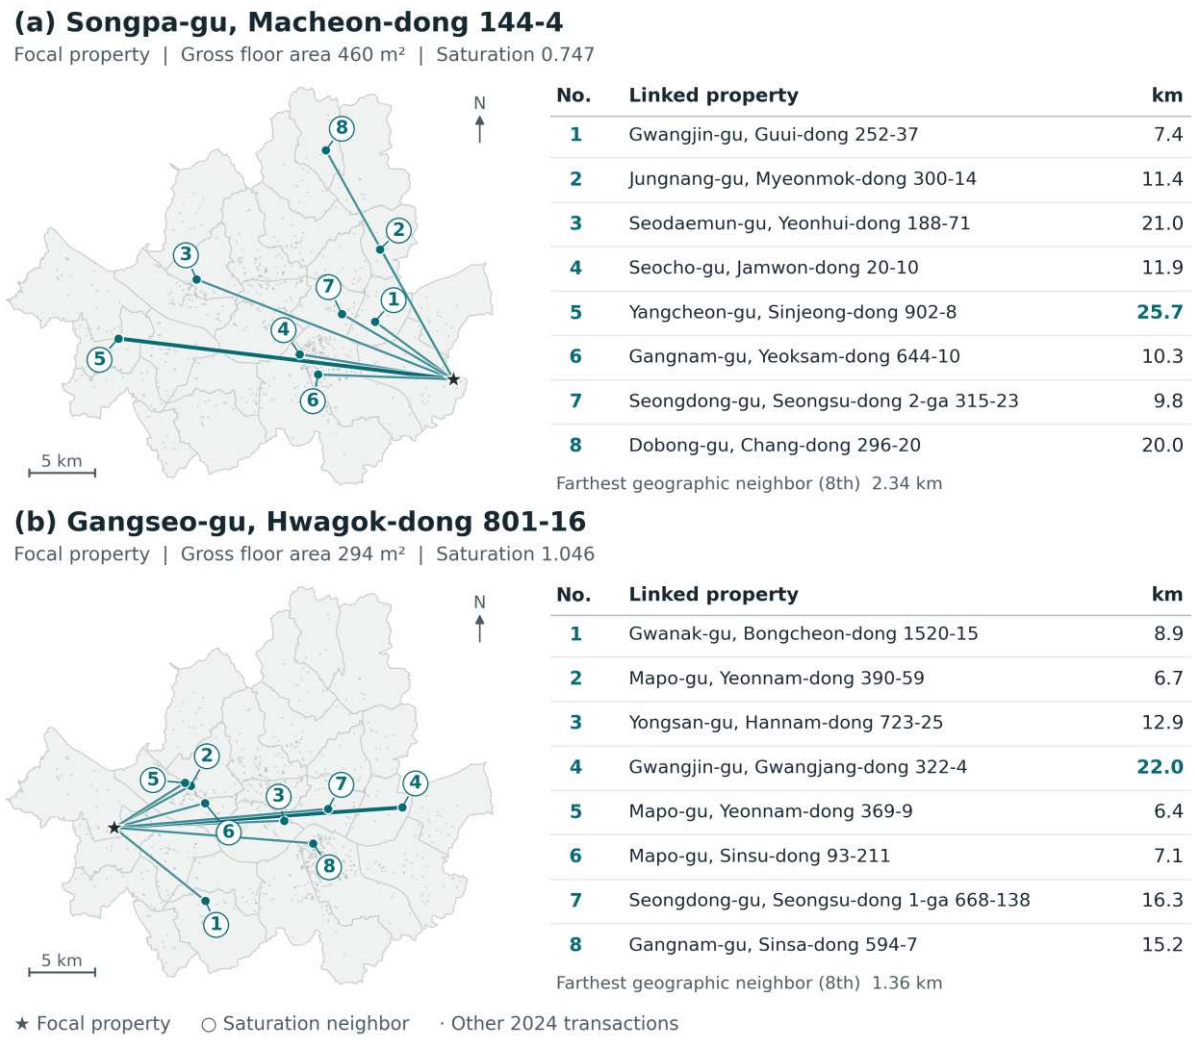


Figure 2 Baseline development-saturation links for two Seoul properties in 2024. Each numbered peer receives a weight of 1/8 from the focal property in the baseline saturation matrix. Numbers rank the eight peers by saturation similarity within the 77 same-year candidates nearest in log g ross floor area. Distances are straight-line centroid distances in EPSG:5186. District boundaries a re from Statistics Korea (2013).

3.4 Six equal-volume allocations

Each of six rules moves 72,812 m² from 499 senders to 300 receivers, about 1.0% of th e sample parcels' 7.10 million m² of nominal capacity. The volume is an experimental sc ale chosen to keep allocations comparable, not a policy target: it is the amount that rand om senders release under the 20% cap and the floor-area constraint, and every other rule is held to the same quantity so that differences among rules reflect design rather than v olume. Parcel representatives are the latest observed transactions, so repeated sales do no t multiply the transferred quantity. The sending comparison varies low-saturation, high-eli gible-saturation and random senders with citywide receivers fixed; the receiving comparis on holds random senders fixed and varies citywide random, district-balanced (12 per dist rict), three-center (100 eligible sites nearest each of the Jongno–Jung, Gangnam and Yeo ngdeungpo centers) and Gangnam-centered (the 300 eligible sites nearest the Gangnam c

enter) receivers. The shared random-sender, citywide-receiver case belongs to both comp arisons. All 300 Gangnam-centered receivers lie within Gangnam-gu (Figure 3).

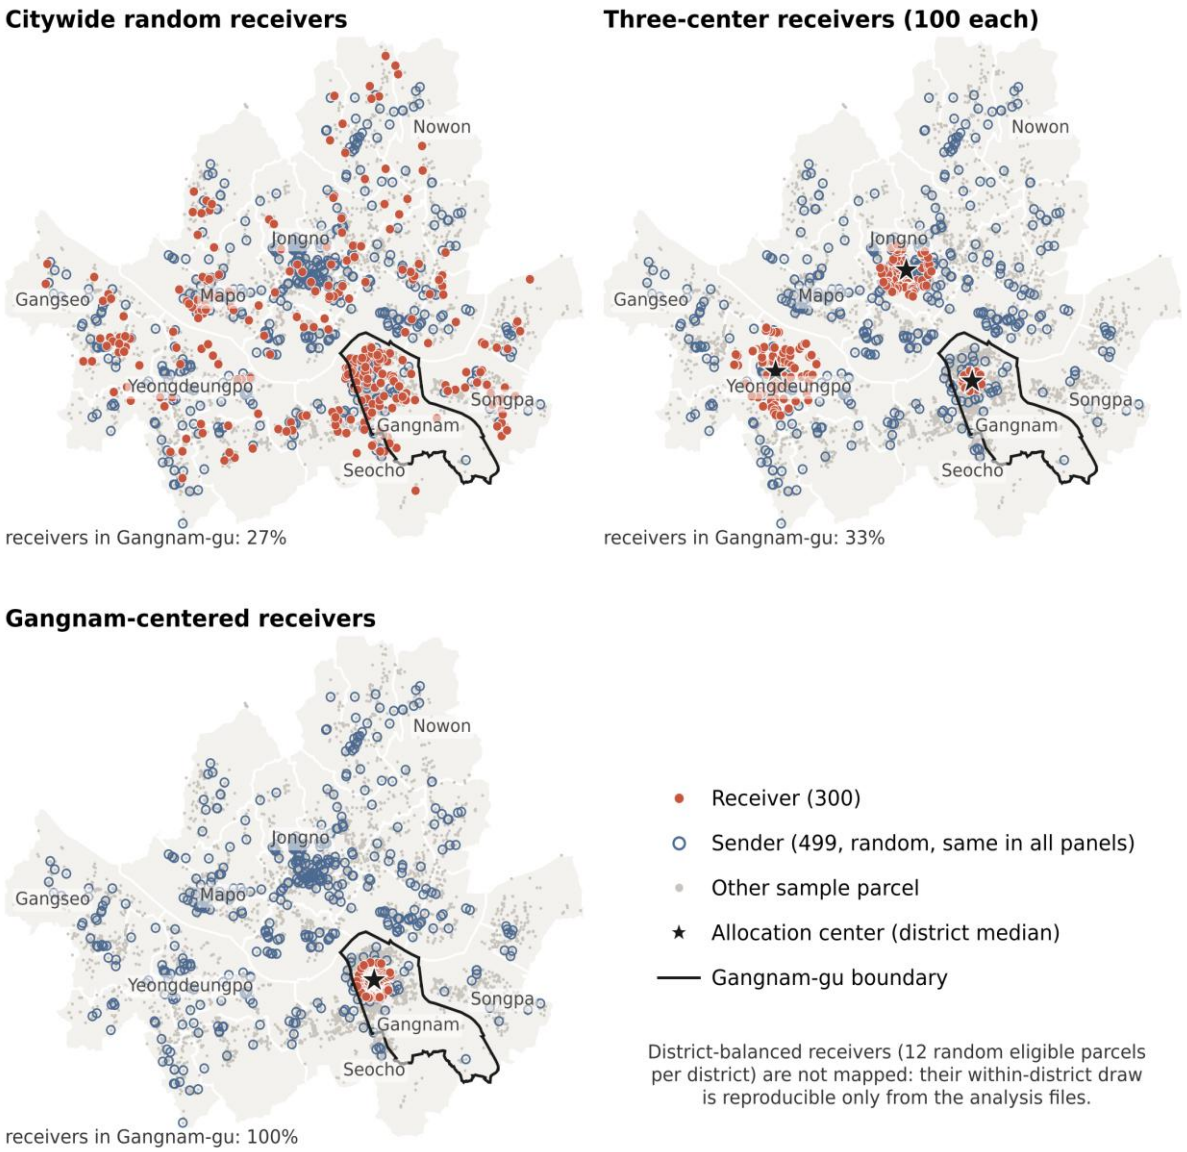


Figure 3 Participant geography under three receiving rules. Filled markers are the 300 receivers, hollow markers the 499 random senders shared by all panels; grey points are the remaining sam ple parcels. Stars mark the district-median centers used by the three-center and Gangnam-centere d rules; Gangnam-gu is outlined. District boundaries: Statistics Korea (2013); coordinates EPSG: 5186. Sending cuts and receiving expansions are given in Online Resource Table S2. The district -balanced rule is not mapped because its receivers are a random draw within each district.

Under every rule, sending reductions are proportional, capped at 20% and cannot leave c apacity below registered floor area; a proportional expansion at receivers absorbs the sa me volume. Because selected sites differ in baseline capacity, equal volumes imply differ ent percentage cuts and expansions (Online Resource C, Table S2). The allocation is defi ned at parcel level and mapped to every transaction on the parcel by a common capacity ratio, holding transaction-specific attributes fixed.

3.5 Measures

Four measures describe the amount, location, counterpart composition and aggregate cons equence of change. Overall reconfiguration is $100 \times \Sigma|\Delta\Sigma_{ij}| / \Sigma|\Sigma_{ij}(0)|$ over unordered with in-year off-diagonal pairs; a value of 10% means absolute pairwise changes sum to one t enth of baseline covariance mass, not that every covariance rises by 10%. District shares assign half of each pair's absolute change to each endpoint district and sum to 100%; Gangnam's share and its complement, the outside-Gangnam share, are reported. Receiver –undesignated reconfiguration applies the absolute calculation to pairs with exactly one r eceiver and one undesignated endpoint, normalized by that pair group's own baseline; a signed counterpart lets increases and decreases offset. Fixed-group uncertainty is the stan dard deviation of a group's average residual, $\sqrt{(w'\Sigma w)}$, for senders, receivers and their fi xed undesignated peer groups, with membership and equal weights held fixed within eac h scenario; its percentage change, summarized by the median over the eight years, distin guishes reorganization of pairwise relationships from a change in group-level uncertainty.

3.6 Benchmarks and uncertainty

Two benchmarks say what the district shares would look like without the covariance mo del. The *mechanical exposure* benchmark replaces $|\Delta\Sigma_{ij}|$ with $a_i + a_j$ for each pair and ap plies the same endpoint accounting; the receiver-only version uses $a_i = \max(q_i, 0)$ and th e sender-inclusive version uses $a_i = |q_i|$. These are declared accounting rules, not models of diffusion. The *random perturbation* benchmark permutes the full vector of parcel capa city multipliers, recomputes W_s and Σ with dependence parameters fixed, and records the district shares. A size-conditioned null permutes within five floor-area quintiles and ther eby scatters capacity changes across districts; a district-and-size-conditioned null permute s within district × quintile cells and so preserves the geography of participants. Each nul l uses 199 replicates. Neither permutation null imposes eligibility, the common volume o r capacity conservation, so a third, matched benchmark does: 99 draws of 300 receivers taken at random from the 1,002 eligible parcels in Gangnam-gu, with the same random s enders, the same 72,812 m^2 and the same proportional expansion rule, each run through the full pipeline. The permutation nulls measure how much participant locations matter; t he matched benchmark isolates what the nearest-to-center rule adds beyond location and the numerical constraints. Both tails are reported for every reference distribution.

Beyond these benchmarks, conditional standard errors for λ_e and λ_s come from the obser ved Hessian of the concentrated negative log likelihood with the selected support fixed. Policy metrics are ratios in which σ^2 cancels, so their standard errors follow by the delta method from the joint covariance of the two dependence estimates; paired differences b etween scenarios use the same joint covariance. These intervals quantify fitted-parameter

uncertainty.

4. Results

4.1 Concentrated receiving concentrates, but most change is assigned elsewhere

Receiving rules differ in where modeled change falls far more than in how much of it th ere is (Table 2). Citywide reconfiguration is 8.1–9.9% across the four receiving rules; the marginal intervals share the same fitted parameters, so differences between rules are jud ged by paired contrasts rather than overlap. Under citywide random receiving Gangnam's endpoint share is 20.1%, close to its 21.2% share of transactions; under district-balance d receiving it falls to 15.5%; under Gangnam-centered receiving it rises to 30.7%. The p aired citywide-to-Gangnam-centered difference is 10.62 percentage points with conditiona l 95% interval [10.25, 10.98].

Table 2 Equal-volume allocations: amount and location of covariance change

Allocation rule	Overall change % [95% interval]	Gangnam share %	Outside Gangnam % [95% interval]
Low-saturation senders	7.27 [6.66, 7.89]	23.98	76.02 [75.85, 76.18]
High-eligible senders	8.53 [7.77, 9.29]	23.09	76.91 [76.81, 77.00]
Citywide receivers	9.87 [8.96, 10.77]	20.09	79.91 [79.86, 79.96]
District-balanced receivers	9.58 [8.73, 10.43]	15.47	84.53 [84.35, 84.71]
Three-center receivers	8.07 [7.37, 8.78]	19.23	80.77 [80.62, 80.92]
Gangnam-centered receivers	9.78 [8.89, 10.67]	30.71	69.29 [68.88, 69.70]

All scenarios transfer 72,812 m^2 between 499 senders and 300 receivers. Rows 1–3 shar e receivers; rows 3–6 share senders. Overall change is relative to baseline absolute off-di agonal covariance. Shares assign half of each pair's absolute change to each endpoint dis trict; the Gangnam share is 100 minus the outside share. Brackets are conditional 95% d elta-method intervals. H1 controls, fully updated weights.

Notably, even when every receiver is in Gangnam, 69.3% of absolute covariance change is assigned to parcels in other districts. Pairs with both endpoints outside Gangnam acco unt for about half of total change and Gangnam–other pairs for about two fifths, half of which enters the outside share. Part of this reflects the senders, which are drawn citywid e in every receiving scenario, so the full-allocation result does not isolate transmission o utward from receivers; Section 4.3 separates the two.

By contrast, with receivers fixed, sender selection changes the amount of reconfiguration. Low-saturation, high-eligible-saturation and random senders produce 7.27%, 8.53% and 9.87%; the paired differences from random sending are -2.59 [-2.88, -2.30] and -1.34 [-1 .49, -1.19] points. The same transferred volume therefore produces different amounts of

network change depending on which parcels give up capacity, because a proportional cut moves a parcel's relative saturation by different amounts in different parts of the distrib ution.

4.2 Relationships with undesignated properties

Gangnam-centered receiving reorganizes relationships between receivers and undesignated properties far more than the other rules: 36.8% of that pair group's baseline covariance in absolute terms, against 20.5–21.8% for the three dispersed rules (Table 3). The margin al intervals are wide and share the same fitted parameters, so the comparison rests on th e paired contrast: the Gangnam-centered minus citywide difference is 16.3 points with co nditional 95% interval [11.0, 21.5]. Because the own-group denominator differs across sc enarios, Table 3 also expresses the same receiver–undesignated change on two common bases: relative to the baseline covariance of all pairs it is 1.90% under citywide and 2.85 % under Gangnam-centered receiving (paired difference 0.95 [0.85, 1.05]), and as a shar e of total absolute change 19.3% against 29.2% (9.9 [9.6, 10.2]). Signed changes are sm all and negative under every rule (−0.3% to −1.7%), so the gross reorganization consists of offsetting increases and decreases rather than a net rise in dependence.

Table 3 Receiver–undesignated covariance change by receiving rule

Receiving rule	Absolute change, own-group denominator % [95% interval]	Signed change, own-group denominator % [95% interval]	Absolute change, all-pairs denominator % [95% interval]	Share of total absolute change % [95% interval]
Citywide	20.51 [17.77, 23.25]	-0.33 [-0.40, -0.26]	1.90 [1.64, 2.16]	19.26 [18.36, 20.17]
District-balanced	21.81 [18.83, 24.79]	-1.02 [-1.15, -0.88]	1.88 [1.63, 2.13]	19.61 [18.68, 20.53]
Three-center	20.55 [17.18, 23.91]	-0.65 [-0.71, -0.59]	1.72 [1.51, 1.93]	21.30 [20.52, 22.09]
Gangnam-centered	36.77 [29.32, 44.21]	-1.67 [-1.80, -1.54]	2.85 [2.50, 3.20]	29.16 [28.21, 30.11]

Pairs with exactly one receiver and one undesignated endpoint. The own-group denomina tor is each scenario's receiver–undesignated baseline absolute covariance, which differs a cross scenarios; the all-pairs denominator is the baseline absolute covariance of all withi n-year pairs, common to all scenarios; the last column is the receiver–undesignated sum as a share of total absolute change. Brackets are conditional 95% delta-method intervals. Paired contrasts are in Online Resource F, Table S10.

At the same time, the fixed-group check shows that this reorganization is not an increase in aggregate uncertainty, with one qualification. Figure 4 plots the annual percentage ch ange in the standard deviation of average residuals for senders, receivers and their two fi xed peer groups under each receiving rule. Across the four rules and four groups the me

dians over eight years all lie within ±1%, but the year-level values spread more widely: 119 of the 128 rule-group-year values (93.0%; Online Resource Table S9) lie within ±1 %, the extremes are +2.1% (Gangnam-centered receivers, 2019) and −1.6% (senders, 201 9), and under high-eligible sending the sender group falls by 1.4% at the median and 2.2 % in the worst year. Peer groups, whose covariance with designated parcels is the chann el of interest, never move by more than 0.72% in any year. Substantial pairwise reconfig uration therefore coexists with small, year-dependent changes in group-level uncertainty, concentrated in the designated groups themselves.

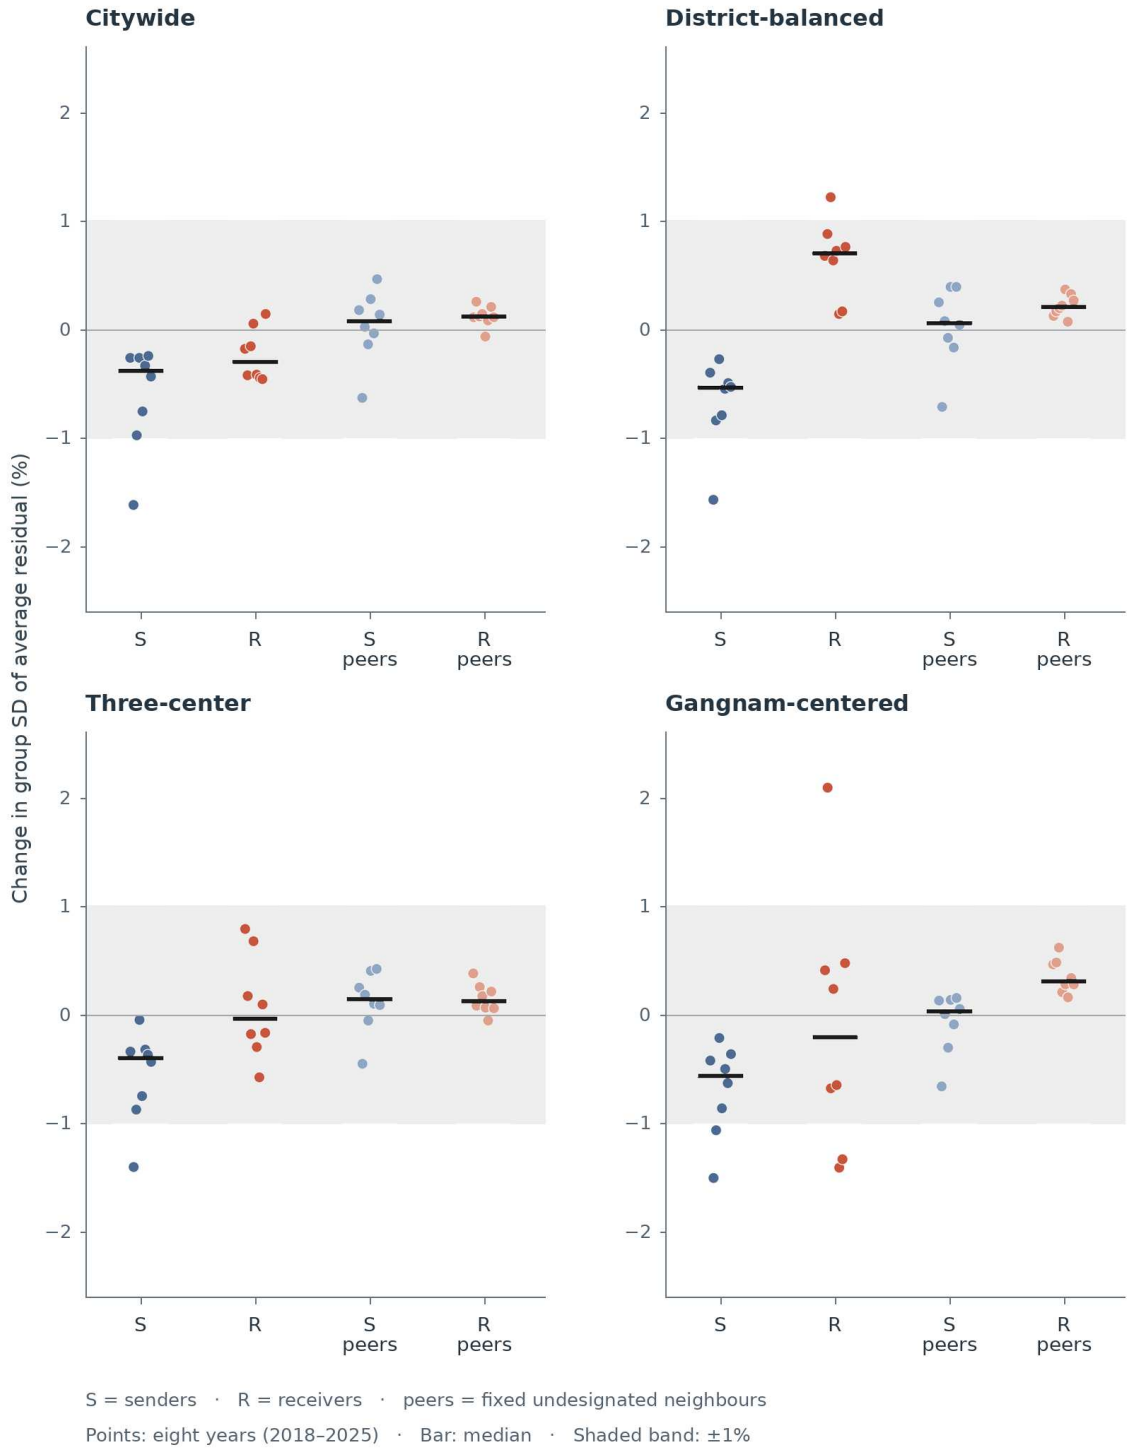


Figure 4 Annual change in the standard deviation of fixed-group average residuals, by receiving rule. S and R are senders and receivers; peers are the undesignated parcels linked to them in the baseline saturation graph. Each point is one year (2018–2025), the bar is the median across years and the shaded band marks ±1%. Values are $100 \times [\sqrt{(w'\Sigma(q)w)}/\sqrt{(w'\Sigma(0)w)} - 1]$ with one observation per parcel-year and equal weights.

4.3 What the benchmarks attribute to participant locations

The geography of the Gangnam-centered allocation is accounted for by the placement of participants; the within-district selection rule adds no detectable geographical concentrati on. Table 4 establishes this by placing the observed outside-Gangnam share of 69.3% ag ainst five benchmarks. A receiver-only exposure rule, which spreads each receiver's capa city gain over its model pairs, expects only 39.5% outside Gangnam, so relative to wher e the capacity goes the model assigns 29.8 points more change outside the receiving dist rict. Including the citywide senders raises the mechanical expectation to 62.5%, leaving a n excess of 6.8 points; Figure 5 maps the district geography of that benchmark. Participa nt locations under a simple accounting rule therefore explain most of the nonlocal footpr int; the remaining question is whether the model adds anything beyond location.

Table 4 Benchmarks for the outside-Gangnam share under Gangnam-centered recei ving

Benchmark	Outside share % [interval or range]	Observed − benchmark mean (pp)	Position of observed value
Observed model	69.29 [68.88, 69.70]	—	—
Receiver-only exposure	39.50	+29.79	fixed rule
Sender + receiver exposure	62.49	+6.80	fixed rule
Size-conditioned null	78.27 [77.32, 79.15]	-8.98	below range by 8.0 pp (lower-tail p = 0.005)
District + size null	70.59 [69.90, 71.28]	-1.30	below range by 0.6 pp (lower-tail p = 0.005)
Matched random Gangnam receivers	69.15 [68.67, 69.74]	+0.14	inside range (two-sided p = 0.60)

The observed row carries a conditional 95% interval. Mechanical expectations have no s ampling interval because locations and amounts are fixed. Permutation nulls reshuffle par cel capacity multipliers within strata (199 draws each; brackets are 2.5th–97.5th percentil es, reference ranges rather than confidence intervals; Monte Carlo p = (1 + count)/200, minimum 0.005). The matched benchmark draws 300 random eligible receivers inside G angnam-gu with the same senders, volume and expansion rule (99 draws; two-sided p = 2 × min(1 + #draws ≤ observed, 1 + #draws ≥ observed)/100, minimum 0.02). Differenc es are from benchmark means; the last column gives the distance from the nearer limit o f the range.

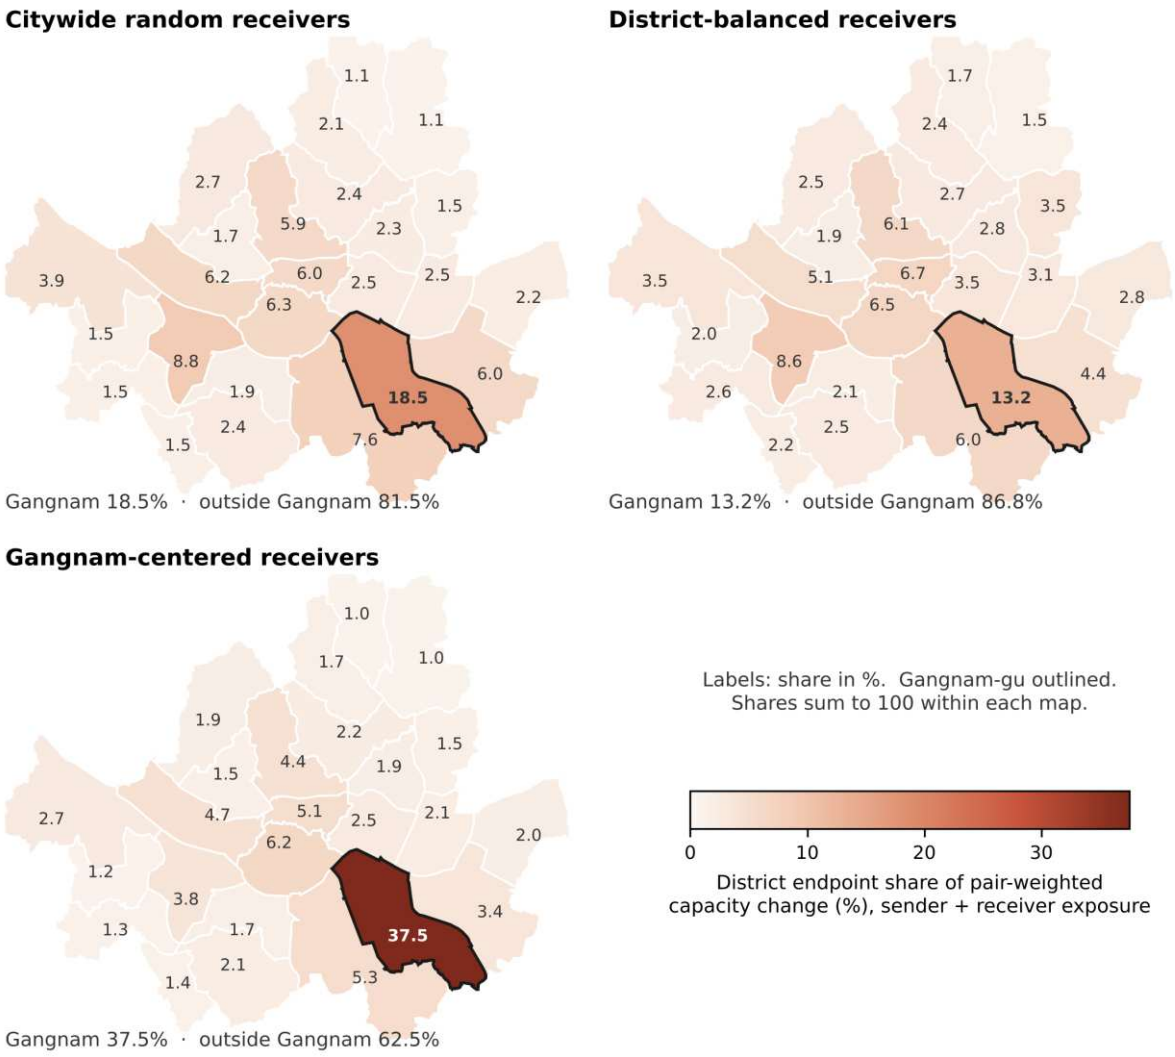


Figure 5 District endpoint shares under the sender-and-receiver mechanical exposure benchmark, three receiving rules. Each district is shaded by the share of pair-weighted absolute capacity change assigned to its endpoints (Section 3.6); labels are percentages and Gangnam-gu is outlined. The Gangnam-centered panel corresponds to the sender-and-receiver benchmark in Table 4. The benchmark places 37.5% in Gangnam even though every receiver is there, because the 499 senders are dispersed citywide and every within-year pair contributes to both endpoints.

Turning to the permutation nulls, when the same capacity multipliers are reshuffled without regard to district, so that capacity changes scatter across the city, the null mean outside share is 78.3% with central range [77.3, 79.1]; every one of 199 draws lies above the observed 69.3% (lower-tail p = 0.005). The Gangnam-centered allocation is thus more concentrated than any location-blind perturbation of equal magnitude. Reshuffling within district-and-size cells, which preserves the geography of participants, lowers the null mean to 70.6% and the range to [69.9, 71.3], and every draw still exceeds the observed value. Participant locations therefore account for most of the gap, but the permutation null cannot say whether the remaining 1.3 points reflect the nearest-to-center rule or merely the eligibility, volume and conservation constraints that permutation ignores.

The matched benchmark answers that question. When 300 receivers are drawn at random

from the eligible parcels of Gangnam-gu under the same senders, the same 72,812 m^2 and the same conservation and expansion rules, the outside share is 69.2% with range [68.7, 69.7], and the observed 69.3% lies inside it (70 of 99 draws below, 29 above; two-sided p = 0.60). Gangnam's own share is likewise inside its matched range (30.7% against [30.3, 31.3]). No additional geographical concentration attributable to the nearest-to-center selection is therefore detected once eligible receivers are placed in Gangnam under the numerical constraints. This does not make the broad placement choice immaterial: placing receivers in Gangnam rather than citywide is what moves Gangnam's share from 20.1% to 30.7%, and that choice is a policy decision in its own right. The matched benchmark isolates the within-district selection rule from it. What the rule does change is the counterpart composition: receiver–undesignated reconfiguration is 36.8% under the observed allocation against 23.5–25.8% across all 99 matched draws, and overall reconfiguration is slightly lower (9.78% against a matched range of [9.84, 10.35], 1 draw below). Clustering receivers around a single center concentrates their similarity links on a compact set of undesignated peers and reorganizes those relationships far more than dispersed receivers within the same district would.

4.4 Baseline validation and robustness

The saturation representation is supported by held-out spatial prediction, but only under district-by-year controls, and the policy comparisons are robust to the saturation measure, the participant draw and the specification alternatives. Adding W_s to a geography-only model reduces similar-basket RMSE from 0.08446 to 0.08177, a 3.19% reduction with improvement in all five folds (Table 5). The fitted dependence parameters are λ_e = 0.592 (SE 0.016; 95% interval [0.560, 0.625]) and λ_s = 0.171 (SE 0.024; [0.125, 0.218]); the saturation coefficient is modest but precisely estimated. The scenario calculations use the full-sample estimates λ_e = 0.591 and λ_s = 0.168 (Online Resource F). Retaining the full predictive covariance rather than its diagonal raises nominal 95% basket coverage from 88.2% to 90.8% and improves the mean negative log score by 0.025, which supports the relevance of cross-property covariance when prices are aggregated.

Table 5 Historical spatial prediction: adding saturation links

Specification and unit	Geography-only RMSE	With saturation	Reduction %
H1 similar baskets	0.08446	0.08177	3.19
H1 individual prices	0.35206	0.35127	0.23
H2 similar baskets	0.08515	0.08520	−0.06

H1 includes district-by-year effects; H2 year effects only. Each specification scores 6,413 transactions and 1,200 overlapping similar-property baskets. Positive reduction favors adding saturation.

That said, the validation gain is specification dependent. It is small for individual prices (0.23%), absent under H2 (a slight deterioration), and smaller in the later period (0.63% in 2022–2025 against 2.38% in 2018–2021 in within-period checks). Forcing both logged areas into H1 retains a 1.70% basket gain and minimum-error selection retains 1.84% (Online Resource D). Robustness to the saturation measure was tested by re-estimating th e dependence parameters and all six scenarios with the above-ground-only index of Secti on 3.3. With the alternative measure λ_s falls from 0.168 to 0.147 (SE 0.024) and λ_e is 0. 593, close to the gross-area estimate of 0.591; holding the participant lists fixed, the Gan gnam-centered minus citywide difference in Gangnam share rises from 10.6 to 12.2 point s (19.4% to 31.6%), the outside share under Gangnam-centered receiving is 68.4%, and r eceiver–undesignated reconfiguration is 38.4% against 19.8% under citywide receiving. R e-deriving the eligibility pools and allocations on the alternative index (2,583 sending an d 3,044 receiving parcels) gives a difference of 11.3 points and an outside share of 67.6 %. The conclusions are therefore not an artifact of counting below-ground floor area (On line Resource F, Table S8). The concentration contrast is also robust to the random parti cipant draw and to the specification alternatives above, which give Gangnam-centered mi nus citywide differences of 10.6–10.9 points (Online Resource D, Table S4).

5. Implications for TDR appraisal

For TDR appraisal the results show that receiving-district placement is associated with th e geography of modeled change, while the within-district selection rule changes the reorg anization of receiver–undesignated relationships. Changes in the standard deviation of the fixed-group average residuals remain small under the four receiving rules. Five uses foll ow, each tied to a measure the framework reports. First, the geography of modeled chan ge can be compared across allocations that are indistinguishable in volume and in overall reconfiguration. In Seoul, moving the same 72,812 m^2 through district-balanced rather t han Gangnam-centered receiving halves Gangnam's share of modeled change (15.5% aga inst 30.7%) at essentially the same citywide total. A planner weighing a concentrated rec eiving area for infrastructure or market-depth reasons can see what that choice implies fo r where relational consequences fall.

Second, the benchmarks identify comparison properties for further validation rather than a monitoring boundary. In this sample, the observed Gangnam-centered allocation assigns 69.3% of modeled reorganization outside Gangnam; the matched random Gangnam-recei ver benchmark has a mean of 69.2%, and the district-and-size permutation benchmark ha s a mean of 70.6%. These are model-based results, distinct from the receiver-only and se nder-inclusive mechanical expectations of 39.5% and 62.5%. The practical reading is that the saturation graph names the comparable-size, comparable-saturation parcels outside th e receiving area whose prices should be tracked alongside it if the modeled relationships are to be tested against realized outcomes.

Third, sender selection is a design margin in its own right. The 7.3–9.9% range at fixed receivers shows that which parcels relinquish capacity affects how much the similarity network moves, independently of where the capacity lands. Programs that select sending areas for conservation reasons can report this consequence alongside the conservation benefit.

Fourth, the fixed-group result bounds what the reorganization means in aggregate. Because the median change in group-level uncertainty lies within ±1% under every receiving rule, with the sender group under high-eligible sending the only larger movement, and annual changes across all six allocations range from −2.22% to +2.10%, the covariance measures should be presented as a map of which relationships change, not as evidence that any rule raises or lowers the uncertainty of district or peer-group price levels. Beyond these four uses, receiving demand, infrastructure capacity, transaction costs and conservation objectives remain separate inputs to the decision; the covariance measures add a relational dimension to a comparative screen.

Fifth, the framework can be extended to analyze TDR-induced price increases and their potential transmission. While this study holds buildings and model parameters fixed and varies only development-rights allocations, future research could incorporate a price model to predict receiving-site value gains from redevelopment or expansion and examine how these gains may be reflected in the prices of other properties connected through the covariance network. Properties with similar development conditions are potential competitors for future development rights, so development and value gains at receiving sites may alter their expectations of obtaining rights and realizing development returns. Analyzing this pathway requires updating saturation and the similarity matrix as construction occurs and specifying assumptions that link allocation expectations to price responses.

6. Limitations

The scenarios describe conditional differences among allocations rather than causal effects of an implemented transferable development rights (TDR) program. They simulate the reallocation of nominal development entitlements while holding buildings, mean-equation coefficients, dependence parameters and innovation variance fixed. In practice, receiving additional rights may induce construction or expansion, changing both floor area and development saturation. The resulting similarity matrix could therefore differ from the matrix reconstructed here. Historical predictive performance supports the baseline model but does not establish its validity after actual transfers. Saturation-based links also represent attribute similarity rather than observed market interactions. Saturation and unit price share lot area in their denominators, and property characteristics may be endogenous.

Development saturation and participation eligibility approximate development position rather than directly measure verified legal entitlements or demand. Receivers were selected

from properties with development saturation of at least 0.8, using proximity to or exceed ance of the standardized capacity benchmark as a proxy for potential demand for additio nal development rights. This selection does not establish willingness to acquire rights, no r does saturation above one necessarily indicate unlawful development. The capacity ben chmark does not fully account for property-specific incentives and allowances under distr ict ordinances, area and district designations, district-unit plans and joint development arr angements. Capacity is also calculated using a standardized zoning schedule across transa ction years, while registered gross floor area includes below-ground space. The above-gr ound alternative substantially changes network membership, but the principal geographica l and counterpart contrasts persist after model re-estimation. Rights-market prices, infrastr ucture constraints, construction responses and equilibrium adjustment remain outside the analysis, and the commercial and mixed-use sample does not directly evaluate a particula r residential redevelopment program.

Validation and inference are also conditional. Spatial validation uses held-out property att ributes and overlapping property baskets. Predictive gains vary across specifications and periods, and prediction-interval coverage remains below its nominal level. Reported inter vals exclude model-selection uncertainty and cross-year dependence for the same parcels. Nearest-neighbor selection can generate link replacement. Continuous weights preserve t he geographical concentration contrast. Randomization and permutation benchmarks also have limited scope, and some comparisons remain sensitive to pair membership and base line covariance. District shares depend on the sample and the rule used to attribute pair endpoints to districts, while absolute covariance changes may conceal cancellation betwe en positive and negative changes. These measures identify neither welfare effects nor opt imal allocations. Direct testing of the proposed pathway requires verified entitlements, ob served transfers, subsequent construction activity and repeated price observations that dist inguish changes in mean prices from changes in residual dependence.

7. Conclusion

This study makes two contributions. First, it develops a framework for analyzing transfer able development rights (TDR) allocations by linking a development-saturation similarity matrix, (W), to the covariance structure of a spatial error model. Reallocating developm ent rights changes individual properties' saturation, thereby altering connections among p roperties with similar development conditions and their modeled price covariance. Secon d, the study uses this framework to compare alternative allocations that transfer the same volume of development capacity. The Seoul results show that sender and receiver selecti on changes the extent and geographical distribution of network reconfiguration, as well a s the composition of the property groups involved. Even when receivers are concentrated in a particular area, modeled network reconfiguration extends beyond that area. These fi ndings indicate that TDR appraisal should consider relationships involving senders and u

ndesignated properties with similar development conditions alongside receiving sites. A p olicy-relevant next step is to combine this framework with a price model that distinguish es receiving-site gains from changes in development expectations among similar properti es. Updating saturation and W as construction occurs would allow conditional price resp onses elsewhere to be examined under explicit assumptions. Covariance changes alone es tablish neither actual price increases nor causal spillovers.

References

Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2018). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, 126(6), 2224–2276. https://doi.org/10.1086/700073

Bathelt, H., and Glückler, J. (2003). Toward a relational economic geography. *Journal of Economic Geography*, 3(2), 117–144. https://doi.org/10.1093/jeg/3.2.117

Boschma, R. (2005). Proximity and innovation: A critical assessment. *Regional Studies*, 39(1), 61–74. https://doi.org/10.1080/0034340052000320887

Bruno, E., Falco, E., Shahab, S., and Geneletti, D. (2023). Integrating ecosystem services in transfer of development rights: A literature review. *Land Use Policy*, 131, 106694. https://doi.org/10.1016/j.landusepol.2023.106694

Cheshire, P. C., and Hilber, C. A. L. (2008). Office space supply restrictions in Britain: The political economy of market revenge. *Economic Journal*, 118(529), F185–F221. https://doi.org/10.1111/j.1468-0297.2008.02149.x

Dubin, R. A. (1988). Estimation of regression coefficients in the presence of spatially autocorrelated error terms. *Review of Economics and Statistics*, 70(3), 466–474.

Elhorst, J. P., Lacombe, D. J., and Piras, G. (2012). On model specification and parameter space definitions in higher order spatial econometric models. *Regional Science and Urban Economics*, 42(1–2), 211–220. https://doi.org/10.1016/j.regsciurbeco.2011.09.003

Glaeser, E. L., and Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, 32(1), 3–30. https://doi.org/10.1257/jep.32.1.3

Gneiting, T., and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378. https://doi.org/10.1198/016214506000001437

Hilber, C. A. L., and Vermeulen, W. (2016). The impact of supply constraints on house prices in England. *Economic Journal*, 126(591), 358–405. https://doi.org/10.1111/ecoj.12213

Kelejian, H. H., and Piras, G. (2014). Estimation of spatial models with endogenous weighting matrices, and an application to a demand model for cigarettes. *Regional Science and Urban Economics*, 46, 140–149. https://doi.org/10.1016/j.regsciurbeco.2014.03.001

Linkous, E. R. (2016). Transfer of development rights in theory and practice: The restructuring of TDR to incentivize development. *Land Use Policy*, 51, 162–171. https://doi.org/10.1016/j.landusepol.2015.10.031

Linkous, E. R. (2017). Transfer of development rights and urban land markets. *Environment and Planning A*, 49(5), 1122–1145. https://doi.org/10.1177/0308518X16686794

Machemer, P. L., and Kaplowitz, M. D. (2002). A framework for evaluating transferable

development rights programmes. *Journal of Environmental Planning and Management*, 45(6), 773–795. https://doi.org/10.1080/0964056022000024334

McConnell, V., and Walls, M. (2009). Policy monitor: U.S. experience with transferable development rights. *Review of Environmental Economics and Policy*, 3(2), 288–303. https://doi.org/10.1093/reep/rep008

McConnell, V., Walls, M., and Kopits, E. (2006). Zoning, TDRs and the density of development. *Journal of Urban Economics*, 59(3), 440–457. https://doi.org/10.1016/j.jue.2005.12.005

Pinkse, J., Slade, M. E., and Brett, C. (2002). Spatial price competition: A semiparametric approach. *Econometrica*, 70(3), 1111–1153. https://doi.org/10.1111/1468-0262.00320

Pruetz, R., and Standridge, N. (2009). What makes transfer of development rights work? Success factors from research and practice. *Journal of the American Planning Association*, 75(1), 78–87. https://doi.org/10.1080/01944360802565627

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., and Dormann, C. F. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8), 913–929. https://doi.org/10.1111/ecog.02881

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1), 34–55. https://doi.org/10.1086/260169

Seoul Metropolitan Government (2022). Ordinance on Urban Planning, Ordinance No. 8435, effective 11 July 2022, Article 55 [in Korean]. Enhanced Local Laws and Regulations Information System.

Seoul Metropolitan Government (2025). Unused floor area to places that need it: Proposed introduction of transferable development rights and policy conference on 25 February [in Korean]. My Hand Seoul, 24 February.

Seya, H., Yamagata, Y., and Tsutsumi, M. (2013). Automatic selection of a spatial weight matrix in spatial econometrics: Application to a spatial hedonic approach. *Regional Science and Urban Economics*, 43(3), 429–444. https://doi.org/10.1016/j.regsciurbeco.2013.02.002

Shahab, S., Clinch, J. P., and O'Neill, E. (2018). Estimates of transaction costs in transfer of development rights programs. *Journal of the American Planning Association*, 84(1), 61–75. https://doi.org/10.1080/01944363.2017.1406816

Statistics Korea (2013). Administrative district boundaries of Seoul, municipal level.

Walls, M., and McConnell, V. (2007). Transfer of development rights in U.S. communities: Evaluating program design, implementation, and outcomes. Resources for the Future, Washington, DC.

Yang, X., Gao, Y., and Li, Q. (2025). Beyond geography: House price covariance and socio-economic distance in homeownership decisions. *Empirical Economics*, 69(5), 2927–2971. https://doi.org/10.1007/s00181-025-02731-2

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- TDROnlineResource.docx