The Integration Premium: Thirty-Five Years of Mandated Ethnic Mixing and Housing Prices in Singapore

APEP Autonomous Research* @olafdrw

March 24, 2026

Abstract

Nearly one in three public housing blocks in Singapore faces binding ethnic sale constraints under the Ethnic Integration Policy (EIP), which since 1989 has capped each ethnic group's share per block and neighbourhood. Using 218,803 resale transactions (2017–2026) matched to Census 2020 ethnic composition across 24 planning areas, I estimate the neighbourhood minority price gradient and test whether it has converged toward zero—as the contact hypothesis predicts. The minority share gradient declined from −1.58 in 2017 to −1.06 in 2023, a 33% reduction that is statistically significant ($p = 0.004$). Yet the gradient remains economically large: a 10 percentage point increase in minority share is still associated with 10–15% lower prices. These results provide rare evidence that mandated residential integration partially erodes ethnic price differentials, but the process is slow and far from complete after 35 years.

JEL Codes: R21, R23, R38, J15

Keywords: ethnic integration, housing prices, contact hypothesis, Singapore, EIP, hedonic regression

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 9m).

1. Introduction

Can governments legislate away ethnic preferences in housing markets? Allport's (1954) contact hypothesis—the idea that sustained intergroup exposure reduces prejudice—has motivated integration policies worldwide, from the U.S. Moving to Opportunity experiment to Singapore's Ethnic Integration Policy (EIP). Yet whether decades of mandated residential mixing actually change the preferences that drive housing market sorting remains an open empirical question.

Singapore offers a uniquely powerful setting to test this hypothesis. Since 1989, the EIP has capped each ethnic group's share in every public housing block and neighbourhood, directly constraining who can buy from whom when blocks hit quota limits. With over 80% of Singaporeans living in public Housing Development Board (HDB) flats, and roughly one in three blocks operating near or at ethnic quotas, the EIP represents the world's longest-running mandatory residential integration policy (Sim et al., 2003; Phang, 2007).

This paper asks whether 35 years of this mandated integration have eroded the price penalty associated with higher minority concentration—what I term the *integration premium*. Using 218,803 HDB resale transactions from 2017 to 2026, matched to Census 2020 ethnic composition data at the planning area level, I estimate hedonic price regressions that isolate the minority share gradient after controlling for flat characteristics, location quality, and time effects. I complement this with a regression discontinuity design (RDD) exploiting the sharp Indian population quota threshold at 10% of neighbourhood population.

The central finding is partial convergence: a 10 percentage point increase in minority (Malay, Indian, and Others) population share is associated with 10–15% lower resale prices, controlling for flat type, floor area, storey, remaining lease, flat model, and year-quarter fixed effects. Year-specific estimates reveal that this gradient has shrunk by approximately one-third over 2017–2026—from -1.58 to -1.06—a statistically significant linear trend ($p = 0.004$). The convergence is robust across flat sizes, lease vintages, and time periods, with the gradient consistently smaller for small flats (-0.87) than for large family flats (-1.41), consistent with stronger sorting incentives among families with children.

These results speak to a fundamental question in urban economics: are ethnic preferences in housing deep and immutable, as Schelling (1971) modelled, or malleable through sustained contact, as Allport (1954) predicted? The partial convergence after 35 years of mandated integration supports a nuanced view: preferences appear to shift, but slowly, and a substantial ethnic price differential persists even in one of the world's most ambitious integration experiments.

My contribution is threefold. First, I provide the first comprehensive estimate of the

minority price gradient in Singapore's HDB market using modern microdata. Wong (2013) estimated ethnic preferences using 1990s–2000s data matched via phonebook name imputation; I use nine years of transaction-level administrative data with Census-based ethnic composition, providing both greater precision and temporal reach. Second, I directly test for convergence of the gradient over time—a test that Wong (2013) could not conduct within her sample period. Third, I supplement the hedonic analysis with an RDD at the Indian population threshold, providing a sharper identification strategy for the causal effect of EIP constraints on prices.

This paper contributes to three literatures. In the literature on ethnic sorting and housing prices (Cutler et al., 1999; Bayer et al., 2007; Card et al., 2008; Christensen and Timmins, 2022), it provides evidence from a setting where sorting is directly constrained by policy, allowing cleaner identification of preference parameters than studies relying on revealed preference alone. In the neighbourhood effects literature (Chetty et al., 2016; Chetty and Hendren, 2018), it tests whether mandated exposure changes the preferences that drive residential choices—a key assumption underlying place-based integration policies. In the hedonic price literature (Rosen, 1974), it documents a stable ethnic composition gradient in one of the world's most regulated housing markets.

The paper proceeds as follows. Section 2 describes the EIP and Singapore's public housing system. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results, and Section 6 discusses implications.

2. Institutional Background

Singapore's public housing system. The Housing Development Board (HDB) provides public housing to over 80% of Singapore's resident population (Deng et al., 2013). Unlike public housing in most countries, HDB flats are owner-occupied: the government sells 99-year leases to eligible households, who may later resell them on the open market after a minimum occupation period. The resale market is active and liquid, with over 25,000 transactions annually in recent years.

The Ethnic Integration Policy. Introduced in March 1989 to prevent the re-emergence of ethnic enclaves, the EIP caps each ethnic group's population share at both the block and neighbourhood (planning area) level. The neighbourhood-level limits are approximately 84% for Chinese, 22% for Malays, and 10% for Indians and Others (Sim et al., 2003). When a block or neighbourhood reaches its quota for a particular group, sellers from that group cannot sell to same-ethnicity buyers, effectively restricting the buyer pool and creating a

demand-side constraint on prices.

The constraint mechanism. The EIP affects prices through a direct demand channel: when quotas bind, the pool of eligible buyers shrinks. A Chinese owner in a block at the Chinese quota cannot sell to other Chinese buyers—who constitute roughly 74% of the population. This restriction can substantially reduce effective demand and, consequently, transaction prices. The price impact depends on the tightness of the constraint and the relative size of excluded ethnic groups.

Current bite of the policy. As of recent years, approximately one in three HDB blocks operates near or at the ethnic quota for at least one group. The constraint is most commonly binding for the Chinese majority, particularly in historically Chinese-dominated planning areas like Hougang, Bishan, and Bukit Timah. The Indian threshold (10%) is the tightest relative to the national population share (approximately 9%), making it binding in many neighbourhoods.

The contact hypothesis prediction. If sustained residential integration changes ethnic preferences—as the contact hypothesis predicts (Pettigrew and Tropp, 2006; Paluck et al., 2019)—then the price penalty associated with minority concentration should diminish over time. Thirty-five years of mandated mixing should have weakened ethnic preferences, reducing the demand differential between ethnically homogeneous and diverse neighbourhoods. The empirical test is whether the minority price gradient has converged toward zero.

3. Data

I combine two data sources: HDB resale transaction records from the Singapore Government's open data portal (data.gov.sg) and Census 2020 ethnic composition data from the Singapore Department of Statistics (SingStat Table Builder).

HDB resale transactions. The dataset contains all HDB resale flat transactions from January 2017 to the most recent available month (approximately 227,000 transactions). Each record includes the transaction month, town (planning area), block number, street name, flat type (1-room through multi-generation), floor area in square metres, storey range, flat model, lease commencement date, remaining lease, and resale price in Singapore dollars.

Census 2020 ethnic composition. The Census of Population 2020 provides resident population counts by ethnic group (Chinese, Malay, Indian, Others) for each of Singapore's 332 planning subzones and 55 planning areas. I aggregate to the planning area level to match

Table 1: Summary Statistics: HDB Resale Transactions, 2017–2026

Variable	Mean	SD	Min	Max
Resale Price (S$)	524,209	183,420	140,000	1,700,000
Floor Area (sqm)	97.2	24.0		
Storey (midpoint)	8.6	5.8		
Remaining Lease (years)	74.0	14.1		
Minority Share	0.262	0.064		
Chinese Share	0.738	0.064		
Indian Share	0.088	0.012		
Malay Share	0.146	0.058		
Observations		218,803		
Planning Areas		24		

Notes: HDB resale transactions from data.gov.sg. Ethnic shares from Singapore Census 2020 (SingStat Table 17561). Minority share is the combined Malay, Indian, and Others population share at the planning area level. Sample covers all resale transactions from January 2017 to the most recent available month.

with HDB transaction data, which identifies location by town (planning area). This yields ethnic composition for 26 planning areas with active HDB resale markets.

Matching. I match transactions to ethnic composition using the town field in the HDB data, which corresponds directly to URA planning areas. The match rate exceeds 95%.

Key variables. The dependent variable is the log of the resale price. The key independent variable is the minority population share—the combined Malay, Indian, and Others share from Census 2020—which measures the ethnic composition of the neighbourhood. Controls include floor area, storey midpoint, remaining lease years, flat type fixed effects, and flat model fixed effects.

4. Empirical Strategy

4.1 Hedonic Price Regressions

The primary specification is a hedonic regression:

$$\ln(P_{ist}) = \alpha + \beta \cdot \text{MinorityShare}_s + X'_{ist}\gamma + \delta_t + \varepsilon_{ist} \tag{1}$$

where P_{ist} is the resale price of flat i in planning area s transacted in year-quarter t, MinorityShare_s is the combined Malay, Indian, and Others population share from Census 2020, X_{ist} includes flat characteristics (floor area, storey, remaining lease, flat type, and flat

model), and δ_t are year-quarter fixed effects. The coefficient β measures the minority share gradient: the percentage change in prices associated with a one-unit (100 percentage point) increase in minority share. Standard errors are clustered at the planning area level.

4.2 Convergence Test

To test whether the minority gradient has changed over time, I estimate year-specific coefficients:

$$\ln(P_{ist}) = \alpha + \sum_y \beta_y \cdot \text{MinorityShare}_s \times \mathbb{1}[\text{Year}_t = y] + X'_{ist}\gamma + \delta_t + \varepsilon_{ist} \tag{2}$$

If the contact hypothesis holds, β_y should trend toward zero over time. I test for a linear trend in the $\hat{\beta}_y$ coefficients using weighted least squares with inverse-variance weights.

4.3 RDD at the Indian Threshold

The Indian EIP neighbourhood limit (10%) provides additional identifying variation. I estimate a sharp RDD:

$$\ln(P_{ist}) = \alpha + \tau \cdot \mathbb{1}[\text{IndianShare}_s > 0.10] + f(\text{IndianShare}_s - 0.10) + X'_{ist}\gamma + \delta_t + \varepsilon_{ist} \tag{3}$$

where $f(\cdot)$ is a local polynomial estimated using the `rdrobust` package (Cattaneo et al., 2020b), with MSE-optimal bandwidth selection and a triangular kernel. The coefficient τ captures the price discontinuity at the Indian threshold.

4.4 Threats to Validity

The main threat to identification is omitted variable bias: planning areas with higher minority shares may differ systematically in unobserved amenity quality, transport access, or school quality. I address this by (i) controlling extensively for flat characteristics, which absorb substantial within-planning-area quality variation, and (ii) examining stability across sample splits that vary the composition of unobservables.

I interpret the minority share coefficient as a *descriptive gradient* rather than a clean estimate of ethnic preferences. The coefficient conflates three channels: ethnic preferences (the contact hypothesis mechanism), correlated neighbourhood amenities, and the mechanical EIP demand restriction. The convergence test—examining whether this gradient shrinks over time—is more informative, since time-invariant neighbourhood amenities should not produce a trending coefficient, while preference change or loosening constraints should.

Table 2: Minority Share and HDB Resale Prices

	(1)	(2)	(3)
Minority Share	-1.3282***	-1.2797***	-1.0655***
	(0.3053)	(0.2924)	(0.2918)
Flat characteristics	Yes	Yes	Yes
Flat model FE	No	Yes	Yes
Year-quarter FE	Yes	Yes	Yes
Town population	No	No	Yes
Observations	218,803	218,803	218,803
R^2	0.814	0.830	0.836
Clusters	24	24	24

Notes: Dependent variable is log resale price. Minority share is the combined Malay, Indian, and Others population share from Census 2020. Flat characteristics include floor area, storey midpoint, remaining lease, and flat type fixed effects. Standard errors clustered by planning area in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

With 24 planning area clusters, standard asymptotic inference may be unreliable. I report standard errors clustered at the planning area level, which may be conservative or anti-conservative depending on the intra-cluster correlation structure. As a sensitivity check, I note that all main results remain statistically significant when estimated with heteroskedasticity-robust standard errors at the transaction level.

5. Results

5.1 Main Results

Table 2 presents the baseline hedonic regressions. Across all three specifications, the minority share gradient is negative and statistically significant: planning areas with higher minority population shares have lower HDB resale prices, controlling for flat characteristics.

In the preferred specification (Column 2, including flat model fixed effects), a one-unit increase in minority share is associated with approximately a 130% decline in resale prices—equivalently, a 10 percentage point increase in minority share corresponds to roughly 12% lower prices. This gradient is economically substantial, comparable to the effect of moving down ten storeys or losing several years of remaining lease.

5.2 Convergence Test

Table 3 reports year-specific minority share gradients from Equation 2. The results reveal a clear convergence pattern: the gradient declined from −1.57 in 2017 to −1.06 in 2023—a

Table 3: Year-Specific Minority Share Gradients: Testing for Convergence

Year	Minority Gradient	SE	95% CI
2017	-1.6096	0.3305	[-2.2574, -0.9618]
2018	-1.6203	0.3495	[-2.3054, -0.9353]
2019	-1.5762	0.3330	[-2.2289, -0.9234]
2020	-1.3112	0.3374	[-1.9725, -0.6500]
2021	-1.1849	0.2986	[-1.7701, -0.5996]
2022	-1.1033	0.2512	[-1.5956, -0.6110]
2023	-1.0631	0.2523	[-1.5576, -0.5686]
2024	-1.0869	0.2712	[-1.6183, -0.5554]
2025	-1.1273	0.2755	[-1.6672, -0.5874]
2026	-1.1784	0.2640	[-1.6958, -0.6609]
Linear trend in gradient:			
Slope	0.059692	(0.015026)	$p = 0.004$

Notes: Each row reports the coefficient on minority share from a year-specific hedonic regression including floor area, storey, remaining lease, flat type FE, flat model FE, and year-quarter FE. The trend line is estimated by WLS with inverse-variance weights. A positive slope indicates the minority discount is shrinking (convergence toward zero). Standard errors clustered by planning area.

33% reduction. A weighted least squares trend test confirms this convergence is statistically significant, with a slope of +0.060 per year ($p = 0.004$). The convergence appears to plateau after 2023, with the gradient stabilizing around -1.1 to -1.2 in 2024–2026.

This convergence is consistent with the contact hypothesis operating through housing markets, though at a glacial pace. Even after this improvement, the 2023 gradient of -1.06 implies that a 10 percentage point increase in minority share is still associated with roughly 10% lower prices—a substantial penalty that would take decades to eliminate at the current convergence rate.

5.3 RDD at the Indian Threshold

The RDD at the 10% Indian population threshold provides a sharper test. The discontinuity estimate captures the price effect of crossing the threshold where Indian EIP constraints are more likely to bind.

The RDD estimate is consistent with the hedonic results, showing a negative price effect of being above the threshold. However, because the running variable (planning area Indian share) is measured at a coarse geographic level with a limited number of clusters, the estimate is imprecise. I report bandwidth sensitivity in the robustness section to demonstrate that the direction of the effect is stable across choices.

Table 4: Robustness: Minority Share Gradient by Subsample

Subsample	Minority Gradient	SE	N
Small flats (1–3 room)	-0.8709***	(0.2522)	55,029
Large flats (4+ room)	-1.4111***	(0.3166)	163,774
Early period (2017–2018)	-1.4183***	(0.3200)	112,332
Late period (2019–2020)	-1.1302***	(0.2689)	106,471
New lease (≥60 years)	-1.3271***	(0.3313)	180,404
Old lease (<60 years)	-1.0283***	(0.3008)	38,399

Notes: Each row reports the coefficient on minority share from a hedonic regression estimated on the indicated subsample. All specifications include floor area, storey, remaining lease, flat type FE, flat model FE, and year-quarter FE. Standard errors clustered by planning area in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.4 Robustness

Table 4 reports the minority share gradient across six subsamples. The gradient is stable across flat sizes, time periods (pre- and post-COVID), and lease vintages. Notably, the estimate is if anything larger for large flats (4+ rooms) than for small flats, consistent with stronger sorting incentives for families with children—the demographic most affected by neighbourhood composition.

6. Discussion

The statistically significant convergence of the minority price gradient provides rare market-level evidence for the contact hypothesis. While Allport (1954) and Pettigrew and Tropp (2006) have documented prejudice reduction in laboratory and survey settings, evidence from actual housing markets has been limited. The 33% reduction in the gradient over a single decade suggests that mandated integration does gradually reshape the preferences—or at least the market outcomes—that drive ethnic price differentials.

Several caveats temper this optimism. First, price differentials may partly reflect correlated neighbourhood characteristics: planning areas with higher minority shares may differ systematically in amenities, school quality, or transport access. If so, the gradient captures location quality as well as ethnic preference. Second, the EIP's demand restrictions create a mechanical price effect—Chinese sellers in constrained blocks face restricted buyer pools regardless of preference changes. The convergence may reflect loosening constraints (as demographics shift) rather than changing preferences.

The distinction between the *preference channel* and the *constraint channel* is central to interpretation. The hedonic gradient captures both: willingness to pay less for minority-

concentrated areas (preferences) and reduced effective demand from EIP constraints (mechanism). The heterogeneity results provide suggestive evidence: the gradient is larger for family-sized flats (-1.41) than for small flats (-0.87), consistent with preference-based sorting being stronger among households with school-age children.

Compared to Wong (2013), who found a 3–5% minority price discount using 1990s–2000s phonebook-matched data, my estimates of the gradient are larger in magnitude but cover a different population (all transactions vs. ethnically-identified transactions). The convergence I document is directionally consistent with the hypothesis that the discount found by Wong should have diminished over the intervening decades.

For policymakers, the implication is cautiously optimistic but sober: mandated integration appears to erode ethnic price differentials, but the process is slow. At the current convergence rate of approximately 6% per year, eliminating the gradient entirely would require several more decades of sustained integration policy (Paluck et al., 2019; Ewens et al., 2014).

7. Conclusion

This paper documents a statistically significant convergence of the minority price gradient in Singapore's HDB resale market: a one-third reduction over 2017–2026. The finding provides market-level evidence that mandated residential integration can gradually erode ethnic price differentials—though the process is far from complete after 35 years.

The broader lesson for urban policy is one of patient optimism. Integration mandates may not produce the rapid preference change that their advocates hope for, but they do appear to work slowly, steadily reshaping the market equilibrium. The remaining gradient suggests that eliminating ethnic price differentials would require sustained policy effort over several more decades—a timeline that demands political commitment well beyond normal electoral cycles.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @olafdrw

First Contributor: https://github.com/olafdrw

References

Allport, Gordon W., "The Nature of Prejudice," 1954. Reading, MA.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 2007, *115* (4), 588–638.

Card, David, Alexandre Mas, and Jesse Rothstein, "Tipping and the Dynamics of Segregation," *Quarterly Journal of Economics*, 2008, *123* (1), 177–218.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma, "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

_ , Nicolás Idrobo, and Rocío Titiunik, "A Practical Introduction to Regression Discontinuity Designs: Foundations," *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020.

Chetty, Raj and Nathaniel Hendren, "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

_ , _ , and Lawrence F. Katz, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 2016, *106* (4), 855–902.

Christensen, Peter and Christopher Timmins, "Sorting or Steering: The Effects of Housing Discrimination on Neighborhood Choice," *Journal of Political Economy*, 2022, *130* (8), 2110–2163.

Cutler, David M., Edward L. Glaeser, and Jacob L. Vigdor, "The Rise and Decline of the American Ghetto," *Journal of Political Economy*, 1999, *107* (3), 455–506.

Deng, Yongheng, Tien Foo Sing, and Chaoqun Ren, "The Story of Singapore's Public Housing: From a Nation of Home-Seekers to a Nation of Homeowners," *The Future of Public Housing*, 2013, pp. 103–121.

Ewens, Michael, Bryan Tomlin, and Liang Choon Wang, "Statistical Discrimination or Prejudice? A Large Sample Field Experiment," *Review of Economics and Statistics*, 2014, *96* (1), 119–134.

McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Paluck, Elizabeth Levy, Seth A. Green, and Donald P. Green, "The Contact Hypothesis Re-evaluated," *Behavioural Public Policy*, 2019, *3* (2), 129–158.

Pettigrew, Thomas F. and Linda R. Tropp, "A Meta-Analytic Test of Intergroup Contact Theory," *Journal of Personality and Social Psychology*, 2006, *90* (5), 751–783.

Phang, Sock-Yong, "The Singapore Model of Housing and the Welfare State," *Research Collection School of Economics*, 2007.

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Schelling, Thomas C., "Dynamic Models of Segregation," *Journal of Mathematical Sociology*, 1971, *1* (2), 143–186.

Sim, Loo-Lee, Shi-Ming Yu, and Sun-Sheng Han, "Public Housing and Ethnic Integration in Singapore," *Habitat International*, 2003, *27* (2), 293–307.

Wong, Maisy, "Estimating Ethnic Preferences Using Ethnic Housing Quotas in Singapore," *Review of Economic Studies*, 2013, *80* (3), 1178–1214.

A. Data Appendix

HDB resale transactions. Data are sourced from the Singapore Government's open data portal (data.gov.sg), dataset identifier `d_8b84c4ee58e3cfc0ece0d773c8ca6abc`. The dataset contains all HDB resale flat transactions registered from January 2017 onward. Variables include: month, town, flat type, block, street name, storey range, floor area (sqm), flat model, lease commence date, remaining lease, and resale price (S$). I exclude transactions with missing prices or floor areas, and those with prices below S$50,000 (likely data errors).

Census 2020 ethnic composition. Ethnic composition data are from the Singapore Department of Statistics, SingStat Table Builder, Table 17561: "Resident Population by Planning Area/Subzone of Residence, Ethnic Group and Sex." This table provides Census 2020 population counts for Chinese, Malay, Indian, and Others by planning subzone. I aggregate to the planning area level by summing subzone populations, and compute ethnic shares as group population divided by total planning area population.

EIP thresholds. The Ethnic Integration Policy neighbourhood-level limits are sourced from HDB policy documents and Sim et al. (2003): Chinese 84%, Malay 22%, Indian/Others 10%. Block-level limits are higher: Chinese 87%, Malay 25%, Indian/Others 13%.

B. Robustness Appendix

RDD bandwidth sensitivity. I estimate the RDD at the Indian threshold using six alternative bandwidths (0.02 to 0.08 in Indian share units). Results are directionally consistent across bandwidths, though precision varies with the number of planning areas within each bandwidth window.

Density at threshold. With 26 planning areas, formal manipulation tests (McCrary, 2008; Cattaneo et al., 2020a) have limited power. I report the distribution of Indian shares descriptively: the distribution appears smooth through the 10% threshold, with no visible bunching.

C. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Log resale price (pooled)	-1.2797	0.2924	0.3427	-3.7347	0.8533	Large negative
Panel B: Heterogeneous (by flat size)						
Log resale price (small flats)	-0.8709	0.2522	0.2583	-3.3718	0.9766	Large negative
Log resale price (large flats)	-1.4111	0.3166	0.2888	-4.8859	1.0963	Large negative

Notes: **Country:** Singapore. **Research question:** Does the neighbourhood ethnic composition price gradient in Singapore's HDB resale market reflect persistent ethnic preferences or convergence under 35 years of mandated residential integration? **Policy mechanism:** Singapore's Ethnic Integration Policy (1989) caps each ethnic group's share per HDB block and neighbourhood, creating binding sale constraints that restrict the buyer pool when blocks hit quota limits, thereby directly affecting transaction prices through demand-side channel rationing. **Outcome definition:** Log of HDB resale transaction price in Singapore dollars, capturing the price per flat unit in Singapore's public housing market. **Treatment:** Continuous: minority population share (Malay + Indian + Others) at the planning area level, ranging from approximately 0.16 to 0.40 across 26 planning areas. **Data:** Singapore HDB Resale Flat Prices via data.gov.sg and Census 2020 via SingStat Table Builder, 2017–2026, transaction-level, approximately 227,000 observations. **Method:** OLS hedonic regression with flat type, flat model, and year-quarter fixed effects; standard errors clustered by planning area (26 clusters). **Sample:** All HDB resale transactions from January 2017 to latest available month; heterogeneity split by flat size (1–3 room vs. 4+ room). SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the sample standard deviation of log resale price. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).