

The Absorption Illusion: Mass Housing Construction and School Quality in Brazil

APEP Autonomous Research* @ailscl

March 31, 2026

Abstract

Brazil's Minha Casa Minha Vida program delivered 1.5 million lottery-allocated housing units through the FAR modality to 1,028 municipalities between 2009 and 2023, creating sudden population shocks in peripheral urban neighborhoods. Using municipality-level IDEB school quality scores in a staggered difference-in-differences design with the [Callaway and Sant'Anna \(2021\)](#) estimator, I find a precise null effect on primary school quality: the ATT is -0.018 ($SE = 0.032$), ruling out effects larger than 0.08 IDEB points. Beneath this null lies a compositional resorting: municipal schools—which directly serve new residents—show a negative point estimate (-0.033), while state schools improve ($+0.130$, $p < 0.05$). The aggregate null masks governance-specific responses. Standard two-way fixed effects overstates the negative effect by $2.5\times$, illustrating heterogeneity bias in staggered adoption designs.

JEL Codes: I21, R23, H75

Keywords: housing policy, school quality, neighborhood effects, staggered DiD, Brazil

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 37m).

1. Introduction

In 2009, Brazil launched the largest subsidized housing program in Latin American history. Minha Casa Minha Vida (MCMV) delivered over six million units to low-income families, transforming the physical landscape of cities from Manaus to Porto Alegre. For the families who won the lotteries, the program offered a path out of informal settlements. For the schools that absorbed their children, the question was different: what happens when thousands of new students arrive at once?

The neighborhood effects literature has focused almost exclusively on movers. The Moving to Opportunity experiments (Chetty et al., 2016; Ludwig et al., 2013) and their intellectual descendants ask what happens to families who leave disadvantaged neighborhoods. The inverse question—what happens to the places that receive them—has received far less attention. Yet from a welfare perspective, the receiving-community externality matters: if housing lotteries improve outcomes for beneficiaries while degrading schools for incumbents, the net social calculus is more ambiguous than the mover-focused literature suggests.

This paper asks whether MCMV’s large-scale housing construction affected school quality in receiving municipalities. I exploit the staggered rollout of MCMV’s Faixa 1 program, which allocated fully subsidized units through municipal lotteries under the FAR (Fundo de Arrendamento Residencial) modality. Between 2009 and 2023, 4,642 FAR projects delivered approximately 1.5 million housing units to 1,028 municipalities. The lottery-based allocation and staggered delivery across municipalities provide plausibly exogenous variation in the timing of large population shocks to local school systems.

I measure school quality using the IDEB (*Índice de Desenvolvimento da Educação Básica*), a composite index of standardized test scores (Prova Brasil) and school pass rates reported biennially at the municipality level from 2005 to 2023. The two pre-program waves (2005 and 2007) provide baseline measures before MCMV’s launch. Using the Callaway and Sant’Anna (2021) doubly robust estimator—which avoids the well-documented biases of two-way fixed effects in staggered adoption designs (Goodman-Bacon, 2021; Sun and Abraham, 2021; De Chaisemartin and d’Haultfoeuille, 2020)—I estimate the average treatment effect on the treated.

The main finding is a precise null. The Callaway–Sant’Anna ATT for primary schools (*anos iniciais*, grades 1–5) is -0.018 (SE = 0.032), with a 95% confidence interval of $[-0.081, 0.045]$. This rules out effects larger than 7% of one standard deviation in IDEB scores. For lower secondary schools (*anos finais*, grades 6–9), the point estimate is $+0.017$ (SE = 0.030)—also a precise null. The event study shows no pre-trends ($p = 0.068$) and flat post-treatment dynamics: MCMV housing construction did not detectably move school quality in either

direction.

The aggregate null, however, conceals a divergent pattern across school governance types. Brazil’s public education system operates through two parallel networks: municipal schools, managed and funded by city governments, and state schools, managed by state governments with broader catchment areas. Municipal schools are the primary point of absorption for MCMV beneficiaries, as new housing projects are sited within municipal boundaries and Faixa 1 families enroll children locally. When I estimate the ATT separately by network, municipal schools show a negative point estimate (-0.033 , $SE = 0.038$) while state schools show a statistically significant positive effect ($+0.130$, $p < 0.05$). This pattern is *suggestive* of compositional resorting—MCMV housing may redirect students into nearby municipal schools while state schools benefit from reduced crowding—though the municipal school estimate is not individually significant, so the decomposition should be interpreted with caution.

I call this the *absorption illusion*: the aggregate effect looks like nothing happened, but the sign pattern across governance types is consistent with differential network-specific responses to the same population shock. The municipal school decline and state school improvement roughly cancel when averaged. Without direct enrollment data, however, this mechanism remains a hypothesis rather than a confirmed channel.

Three additional findings support this interpretation. First, the dose-response relationship is flat: municipalities receiving more housing units show no larger effects, suggesting the null is not driven by insufficient treatment intensity. Second, the result is robust to alternative estimators (stacked DiD, not-yet-treated controls), sample restrictions (excluding 2021 COVID wave), and a leave-one-state-out jackknife across all 27 Brazilian states. Third, standard TWFE overstates the negative effect by a factor of 2.5 (-0.046 , $p < 0.01$), illustrating the practical importance of heterogeneity-robust estimators when treatment effects vary across cohorts—a methodological point that reinforces [Goodman-Bacon \(2021\)](#) and [De Chaisemartin and d’Haultfœuille \(2020\)](#).

This paper contributes to three literatures. First, it extends the neighborhood effects literature by studying receiving communities rather than movers, joining [Aliprantis and Richter \(2019\)](#) and [Chyn \(2018\)](#) in asking about the spatial equilibrium consequences of place-based policies. Second, it adds to the growing evidence on education supply responses to demand shocks, complementing work on immigration and school quality ([Duflo, 2001](#); [Hoxby, 2000](#); [Hunt, 2017](#)) with evidence from a housing-lottery setting. Third, it contributes a practical illustration of TWFE bias in a policy-relevant application with 1,028 treated units and seven treatment cohorts, extending the methodological contributions of [Callaway and Sant’Anna \(2021\)](#) and [Sun and Abraham \(2021\)](#) to a developing-country education context.

2. Institutional Background

Minha Casa Minha Vida. Brazil launched MCMV in March 2009 as part of its counter-cyclical response to the global financial crisis. The program operates through multiple income bands: Faixa 1 serves households earning up to three minimum wages (approximately R\$4,400/month in 2023), while higher bands serve middle-income families through subsidized mortgages. Faixa 1 units are fully subsidized—beneficiaries pay a nominal monthly fee of R\$80–R\$270 depending on location—making selection into the program effectively universal among eligible families (Cardoso, 2013).

The FAR modality and lottery allocation. Within Faixa 1, the FAR modality is the dominant delivery channel for urban areas. Under FAR, the federal government contracts with private developers through the Caixa Econômica Federal (the state development bank) to build residential complexes. Municipalities select beneficiaries through public lotteries from a pre-registered demand list. The lottery mechanism is central to the program’s design and to this paper’s identification: conditional on a municipality receiving a FAR project, the allocation of specific families to units is randomized. My identification exploits variation in *when* municipalities receive their first FAR project, not in *who* within a municipality is allocated housing.

Scale and geography. Between 2009 and 2023, the FAR modality contracted 4,642 projects totaling approximately 1.5 million housing units across 1,028 municipalities. The first wave of contracts (2009–2010) covered 401 municipalities, with subsequent waves in 2011–2013 (380 municipalities) and later cohorts through 2025. Projects are concentrated in medium-to-large urban areas: the median municipality has a population above 50,000. Housing complexes are typically located in peripheral neighborhoods where land is cheaper, often at the urban fringe (Rolnik, 2015).

Brazil’s dual public school system. Brazilian public education operates through two parallel networks managed by different levels of government. Municipal schools (*escolas municipais*) are funded and operated by city governments; they typically serve elementary and lower secondary grades and admit students from the immediate neighborhood. State schools (*escolas estaduais*) are managed by state education secretariats, often with broader geographic coverage and more selective enrollment patterns. This institutional distinction is crucial for understanding how MCMV housing shocks propagate through the education system: new residents primarily enroll in nearby municipal schools, while state schools may be less directly affected.

IDEB as a quality measure. The IDEB was created in 2007 by INEP (Brazil’s National Institute for Educational Studies and Research) as a composite index combining two components: standardized test performance on the Prova Brasil assessment (mathematics and Portuguese) and school-level pass rates. The index ranges from 0 to 10, with the national average rising from 3.5 in 2005 to 5.3 in 2023. IDEB is computed biennially at the school, municipality, state, and national levels for three school stages: *anos iniciais* (grades 1–5), *anos finais* (grades 6–9), and *ensino médio* (high school). The biennial frequency defines the panel structure: the analysis uses 10 waves from 2005 to 2023.

3. Data

MCMV project data. I obtain project-level MCMV data from the Ministry of Cities’ open data portal.¹ The dataset contains 25,888 projects across all modalities, of which 4,642 are FAR. Each record includes the municipality IBGE code, contract signature date, number of contracted and delivered units, project status, and developer information. I define treatment timing as the year of the first FAR contract in each municipality.

IDEB scores. Municipality-level IDEB data come from INEP’s official download portal.² The 2023 release contains all waves (2005–2023) in a single file, providing a balanced panel of 5,568 municipalities observed over 10 biennial waves. I use the “Pública” (all public schools) aggregate as the primary outcome and separately extract municipal and state school breakdowns for the mechanism analysis.

Panel construction. I merge the two datasets on the six-digit IBGE municipality code, which is common to both sources. The resulting analysis panel contains 53,119 municipality-wave observations for *anos iniciais* and 53,068 for *anos finais*. Of the 5,568 municipalities in the IDEB data, 1,028 (18.5%) received at least one FAR project and are classified as treated; the remaining 4,540 serve as the never-treated control group.

Table 1 presents summary statistics. In the pre-treatment period (2005–2007), treated and control municipalities have similar mean IDEB scores (3.90 vs. 3.80), though treated municipalities score slightly higher on average—reflecting the urban-size gradient in both MCMV project allocation and school resources. The full-panel averages (2005–2023) show both groups improving over time, consistent with Brazil’s long-run trend in educational quality.

¹Available at <https://www.gov.br/cidades/pt-br/aceso-a-informacao/acoes-e-programas/habitacao/programa-minha-casa-minha-vida/bases-de-dados-do-programa-minha-casa-minha-vida>, accessed March 2026.

²Available at <https://download.inep.gov.br/ideb/resultados/>, accessed March 2026.

Table 1: Summary Statistics

	Treated		Never-Treated		Full Sample	
	Mean	SD	Mean	SD	Mean	SD
<i>Panel A: Pre-Treatment (2005–2007)</i>						
IDEA Score	3.90	0.90	3.80	0.93	3.82	0.92
Municipalities	1026		4465		5491	
<i>Panel B: Full Panel (2005–2023)</i>						
IDEA Score	5.05	1.16	5.01	1.21	5.01	1.20
Municipality-waves	10,192		42,927		53,119	
Municipalities	1028		4530		5558	
<i>Panel C: MCMV FAR Program</i>						
Total FAR projects			4,642			
Total units contracted			1,494,586			
Treated municipalities			1028			

Notes: Treated municipalities received at least one MCMV Faixa 1 (FAR) housing project. IDEB is the *Índice de Desenvolvimento da Educação Básica*, a composite index of standardized test scores (Prova Brasil) and pass rates ranging from 0 to 10. Panel A shows pre-treatment means (2005–2007 IDEB waves, before MCMV launch in 2009). Panel C reports total FAR modality projects from the Ministry of Cities open data.

4. Empirical Strategy

4.1 Identification

I exploit the staggered rollout of MCMV FAR projects across municipalities in a difference-in-differences framework. The treatment indicator D_{it} equals one if municipality i has received its first FAR contract by IDEB wave t , and zero otherwise. The identifying assumption is that, absent the MCMV housing shock, school quality in treated and control municipalities would have followed parallel trends.

Two features of the setting support this assumption. First, the timing of FAR contracts is driven by supply-side factors—federal budget cycles, developer capacity, land availability, and Caixa’s contracting pipeline—rather than by municipality-level school quality trends. While municipality characteristics (population, urbanization) affect whether a city is eligible for FAR, the *timing* within the eligible set is plausibly orthogonal to IDEB trajectories. Second, the pre-treatment period (2005–2007) predates the MCMV program entirely, so no anticipation effects contaminate the baseline.

4.2 Estimation

I estimate the average treatment effect on the treated using the [Callaway and Sant’Anna \(2021\)](#) doubly robust estimator. This method addresses the well-documented problems with standard two-way fixed effects in staggered adoption settings ([Goodman-Bacon, 2021](#); [Sun and Abraham, 2021](#); [De Chaisemartin and d’Haultfœuille, 2020](#)): when treatment effects vary across cohorts or over time, TWFE can produce biased estimates by using already-treated units as implicit controls. The CS estimator computes group-time average treatment effects $ATT(g, t)$ for each cohort g (defined by first treatment year) and period t , then aggregates using appropriate weights.

I map each municipality’s first FAR contract year to the nearest IDEB wave to define treatment cohorts: 2009, 2011, 2013, 2015, 2017, 2019, and 2023. The primary control group is never-treated municipalities. Standard errors are clustered at the municipality level. I report both the simple aggregate ATT and dynamic (event-study) estimates with event time measured in IDEB waves (two-year intervals).

4.3 Threats to validity

Selection into treatment. Municipalities that receive FAR projects are not randomly selected—they tend to be larger, more urban, and initially higher-performing. However, the DiD design requires parallel *trends*, not parallel *levels*. The event study directly tests for differential pre-trends.

Concurrent policies. MCMV was part of Brazil’s broader Growth Acceleration Program (PAC), which included infrastructure investments in treated municipalities. To the extent that PAC investments independently affected school quality (e.g., through improved transportation or sanitation), my estimates capture the bundled effect of the housing-plus-infrastructure package rather than housing alone. I note this limitation but cannot separately identify the housing channel.

Contract versus delivery timing. Treatment is defined using the FAR contract signature date, but housing units typically require 2–4 years to build and deliver. The actual enrollment shock occurs upon occupancy, not contracting. This timing lag likely attenuates the estimated impact effect: the event-study $t = 0$ coefficient captures municipalities where contracts were signed but construction may still be underway. The biennial IDEB frequency partially mitigates this concern—by $t + 2$ (the next wave), most early projects were delivered—but short-run effects may be understated.

Table 2: Effect of MCMV FAR Housing on School Quality (IDEB)

	Anos Iniciais (Grades 1–5)		Anos Finais (Grades 6–9)	
	TWFE (1)	CS (2021) (2)	TWFE (3)	CS (2021) (4)
ATT	-0.0456*** (0.0149)	-0.0179 (0.0321)	0.0019 (0.0141)	0.0172 (0.0302)
95% CI	[-0.075, -0.016]	[-0.081, 0.045]	[-0.026, 0.029]	[-0.042, 0.076]
Pre-trend p -value	—	0.068	—	—
Control group	—	Never-treated	—	Never-treated
Estimation	DR	DR	DR	DR
Observations	53,106	53,106	53,065	53,065
Municipalities	5,558	5,558	5,567	5,567
Treated municipalities	1028	1028	1028	1028
Municipality FE	Yes	—	Yes	—
Wave FE	Yes	—	Yes	—

Notes: Columns (1) and (3) report two-way fixed effects estimates with municipality and wave fixed effects. Columns (2) and (4) report Callaway and Sant’Anna (2021) doubly robust estimates using never-treated municipalities as the control group. Standard errors clustered at the municipality level in parentheses. The pre-trend p -value tests the null of parallel trends across all pre-treatment periods. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Biennial measurement. IDEB is observed every two years, limiting the event-study window to two pre-treatment waves for the earliest cohort and constraining the precision of pre-trend tests. The marginal pre-trend p -value (0.068) reflects this power limitation rather than clear evidence of differential trends. Later treatment cohorts (2013 onward) have 4+ pre-treatment waves, and the Callaway–Sant’Anna estimator uses all available pre-periods per cohort.

5. Results

5.1 Main results

Table 2 presents the main estimates. The Callaway–Sant’Anna ATT for *anos iniciais* (column 2) is -0.018 (SE = 0.032), statistically indistinguishable from zero. The 95% confidence interval $[-0.081, 0.045]$ rules out effects larger than 0.09 standard deviations of IDEB. For *anos finais* (column 4), the estimate is $+0.017$ (SE = 0.030), also a precise null. Neither school level shows a detectable effect of MCMV housing on educational quality.

The TWFE estimates (columns 1 and 3) tell a different story for primary schools: the conventional estimator produces a significant negative coefficient of -0.046 ($p < 0.01$). This

Table 3: Dynamic Treatment Effects (Event Study)

Event Time	Anos Iniciais		Anos Finais	
	ATT	SE	ATT	SE
$t - 4$	-0.0141	0.0181	-0.0326*	0.0171
$t - 2$	—	—	—	—
$t = 0$	0.0230	0.0172	0.0052	0.0163
$t + 2$	-0.0329	0.0257	0.0032	0.0251
$t + 4$	-0.0019	0.0390	0.0413	0.0351
$t + 6$	-0.0043	0.0440	0.0297	0.0369
$t + 8$	-0.0379	0.0501	-0.0013	0.0442

Notes: Dynamic treatment effects from Callaway and Sant’Anna (2021), aggregated by event time relative to first MCMV FAR contract in the municipality. Event time measured in IDEB waves (biennial). $t = 0$ is the first wave in which the municipality had received MCMV FAR housing. Doubly robust estimation with never-treated control group. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

discrepancy—a factor of 2.5 between TWFE and CS estimates—arises from the heterogeneity bias inherent in staggered adoption designs. Early-treated cohorts (2009–2010, 401 municipalities) serve as implicit controls for later cohorts under TWFE, contaminating the estimate when treatment effects evolve dynamically. The CS estimator corrects this by restricting comparisons to never-treated units.

5.2 Event study

Table 3 reports the dynamic treatment effects. The pre-treatment coefficient at $t - 4$ is -0.014 ($SE = 0.018$), close to zero and insignificant. The impact effect at $t = 0$ is positive ($+0.023$) but imprecise. Subsequent post-treatment estimates oscillate around zero with no systematic trend—the largest magnitude is -0.038 at $t + 8$, still well within the confidence band. The flat post-treatment trajectory indicates that MCMV housing did not produce either immediate or cumulative effects on school quality.

5.3 Mechanism: municipal versus state schools

Table 4 decomposes the aggregate null by school governance type. Municipal schools—the primary absorbers of MCMV families—show a negative ATT of -0.033 ($SE = 0.038$), while state schools show a positive and statistically significant effect of $+0.130$ ($SE = 0.057$).

Table 4: Mechanism: Municipal vs. State School Networks

	All Public (1)	Municipal (2)	State (3)
CS ATT	-0.0179 (0.0321)	-0.0327 (0.0376)	0.1299** (0.0574)
Control	Never-treated	Never-treated	Never-treated
School level	Grades 1–5	Grades 1–5	Grades 1–5

Notes: Callaway and Sant’Anna (2021) ATT estimates by school network governance type, for *anos iniciais* (grades 1–5). Column (1) uses the aggregate public school IDEB. Column (2) uses municipal schools only, which are managed by city governments and serve local residents—the primary absorbers of MCMV Faixa 1 beneficiaries. Column (3) uses state schools, managed by state governments with broader catchment areas. The divergence between municipal (negative) and state (positive) estimates is consistent with compositional resorting: MCMV housing may redirect lower-performing students into nearby municipal schools while reducing pressure on state schools. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The divergence is consistent with compositional resorting: MCMV housing concentrates new (and potentially lower-performing) students in nearby municipal schools, while state schools—with broader catchment areas and sometimes competitive enrollment—experience favorable composition shifts or reduced crowding pressure.

This pattern explains why the aggregate effect is null: the municipal school decline and state school improvement approximately cancel in the population-weighted average. What appears as “no effect” is in fact a reallocation across governance types. The absorption illusion arises because the aggregate IDEB masks offsetting network-specific responses to the same population shock.

Two caveats temper this interpretation. First, the municipal school estimate is not individually significant, so the decomposition is suggestive rather than definitive. Second, I cannot directly observe enrollment flows between networks—the mechanism is inferred from the sign pattern rather than directly measured.

5.4 Robustness

Table 5 summarizes five robustness checks. The main finding is insensitive to the choice of estimator, sample, and treatment definition.

Table 5: Robustness Checks

Specification	ATT	SE
<i>Panel A: Estimator Variation</i>		
Main: CS, never-treated	-0.0179	(0.0321)
CS, not-yet-treated	-0.0389	(0.0257)
Stacked DiD	-0.0374***	(0.0133)
TWFE (biased baseline)	-0.0456***	(0.0149)
<i>Panel B: Sample Variation</i>		
Excluding 2021 (COVID)	0.0110	(0.0309)
Leave-one-state-out [range]	[-0.0277, -0.0054]	—
<i>Panel C: Treatment Definition</i>		
All MCMV modalities (TWFE)	0.0555***	(0.0104)

Notes: All specifications use the IDEB *anos iniciais* (grades 1–5) panel. Panel A varies the estimator. Panel B varies the sample: excluding the 2021 COVID-affected wave, and leave-one-state-out jackknife (27 states, range of ATT estimates). Panel C redefines treatment to include all MCMV modalities (FAR, Rural, Oferta Pública, Entidades), not just FAR. Standard errors clustered at the municipality level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Alternative estimators. The stacked DiD estimate (-0.037 , $SE = 0.013$) and CS with not-yet-treated controls (-0.039 , $SE = 0.026$) are somewhat more negative than the main specification but remain small in magnitude relative to the IDEB scale (0–10). The stacked estimator produces a significant coefficient because it avoids the attenuation that arises from averaging across many heterogeneous group-time effects in the CS aggregation.

Sample variation. Excluding the 2021 wave—which captured COVID-disrupted learning—shifts the ATT to $+0.011$ ($SE = 0.031$), suggesting that the modest negative bias in the main estimate is partly driven by the pandemic wave. The leave-one-state-out jackknife produces ATT estimates ranging from -0.028 to -0.005 across 27 states, with no single state driving the result.

Treatment definition. Expanding treatment to include all MCMV modalities (FAR, Rural, Oferta Pública, Entidades) yields a positive TWFE coefficient ($+0.055$, $p < 0.001$), likely reflecting the different populations served by non-FAR modalities. Rural housing projects, for example, serve smaller communities with different school quality dynamics. The FAR-specific estimate is the relevant causal object for understanding urban housing lotteries.

6. Discussion

The central finding is that Brazil’s largest housing program had no detectable effect on school quality in receiving municipalities—but this null conceals a real reallocation. The absorption illusion teaches a general lesson about place-based policy evaluation: aggregate outcomes can mask governance-specific responses that matter for institutional design.

The mechanism is consistent with a growing literature on how administrative boundaries shape the incidence of demand shocks. When MCMV delivers housing to a municipality, the immediate enrollment burden falls on municipal schools because they serve local neighborhoods. State schools, with broader geographic scope, are insulated from the hyperlocal population shock and may even benefit from reduced competitive pressure for students or teachers. This governance-channel heterogeneity would be invisible in a design that pooled all public schools.

Several limitations warrant caution. First, the municipality is a relatively large unit of analysis. MCMV projects are sited in specific peripheral neighborhoods, but effects are averaged across all schools in the municipality—many of which may be far from any housing project. This aggregation likely attenuates localized enrollment shocks, making the precise null partly an artifact of dilution rather than a true zero at the neighborhood level. Future work exploiting school-level IDEB data with geocoded project locations could sharpen identification. Second, the IDEB is a coarse composite that may miss quality dimensions relevant to MCMV recipients—class sizes, teacher turnover, or infrastructure strain are not captured by test-score composites. Third, without enrollment microdata from the Censo Escolar, the compositional resorting mechanism remains indirect—inferred from sign patterns across governance types rather than measured through student flows. Fourth, MCMV was part of Brazil’s broader Growth Acceleration Program (PAC), so the estimates capture the bundled effect of housing-plus-infrastructure rather than housing alone.

7. Conclusion

Mass housing construction reshapes neighborhoods, but not necessarily their schools—at least not on aggregate. Brazil’s MCMV program delivered 1.5 million subsidized housing units to over a thousand municipalities without detectably moving school quality. The null, however, is an illusion of aggregation: beneath it lies a systematic resorting of students across governance types. For policymakers designing housing programs with educational co-benefits in mind, the lesson is not that schools are resilient but that the burden is unequally distributed—and the institutions that absorb it are precisely those with the fewest resources to respond.

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: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Aliprantis, Dionissi and Francisca GC Richter**, “What is a neighborhood? An examination using the Los Angeles Family and Neighborhood Survey,” *Journal of Urban Economics*, 2019, *112*, 90–110.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cardoso, Adauto Lucio**, *O Programa Minha Casa Minha Vida e seus Efeitos Territoriais*, Letra Capital, 2013.
- Chaisemartin, Clément De and Xavier d’Haultfœuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Chetty, Raj, Nathaniel Hendren, 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.
- Chyn, Eric**, “Moved to opportunity: The long-run effects of public housing demolition on children,” *American Economic Review*, 2018, *108* (10), 3028–3056.
- Duflo, Esther**, “Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment,” *American Economic Review*, 2001, *91* (4), 795–813.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Hoxby, Caroline M**, “Does competition among public schools benefit students and taxpayers?,” *American Economic Review*, 2000, *90* (5), 1209–1238.
- Hunt, Jennifer**, “The effect of immigration on the educational attainment of natives,” *Journal of Human Resources*, 2017, *52* (4), 1060–1118.
- Ludwig, Jens, Greg J Duncan, Lisa A Gennetian, Lawrence F Katz, Ronald C Kessler, Jeffrey R Kling, and Lisa Sanbonmatsu**, “Long-term neighborhood effects on low-income families: Evidence from Moving to Opportunity,” *American Economic Review*, 2013, *103* (3), 226–231.

Rolnik, Raquel, *Guerra dos lugares: a colonização da terra e da moradia na era das finanças*, Boitempo, 2015.

Sun, Liyang and Sarah Abraham, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
IDEB (Grades 1–5)	-0.0179	0.0321	0.903	-0.0198	0.0355	Small negative
IDEB (Grades 6–9)	0.0172	0.0302	0.687	0.0250	0.0439	Small positive
<i>Panel B: Heterogeneous</i>						
IDEB Municipal Schools	-0.0327	0.0376	0.903	-0.0362	0.0417	Small negative
IDEB State Schools	0.1299	0.0574	0.903	0.1438	0.0636	Moderate positive
IDEB (Grades 1–5, excl. 2021)	0.0110	0.0309	0.903	0.0122	0.0342	Small positive
IDEB Dose-Response	-0.0309	0.0666	0.903	-0.0342	0.0738	Small negative

Notes: **Country:** Brazil. **Research question:** Does large-scale subsidized housing construction (MCMV Faixa 1 FAR) affect school quality in receiving municipalities, as measured by the composite IDEB index? **Policy mechanism:** The Minha Casa Minha Vida program delivers lottery-allocated housing to low-income families in peripheral urban neighborhoods, potentially straining local school infrastructure through sudden enrollment increases. **Outcome definition:** IDEB (Índice de Desenvolvimento da Educação Básica), a composite of Prova Brasil standardized math and Portuguese scores with school pass rates, ranging 0–10. **Treatment:** Binary indicator for municipality receiving first MCMV FAR contract. **Data:** INEP IDEB school-level data aggregated to municipality level (2005–2023, biennial, 5,568 municipalities) merged with Ministry of Cities MCMV open data (4,642 FAR projects, 1,028 treated municipalities). **Method:** Callaway and Sant’Anna (2021) doubly robust staggered DiD with never-treated control group, standard errors clustered at the municipality level. **Sample:** All Brazilian municipalities with IDEB scores; treatment defined as first FAR modality contract date; restricted to *anos iniciais* (grades 1–5) for primary outcome and municipal/state school breakdowns for mechanism analysis. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment (2005–2007) standard deviation of IDEB among treated municipalities. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

A. Standardized Effect Sizes