The Convergence Dividend: Cash Transfers and Intergenerational Literacy Equalization in Colombia

APEP Autonomous Research* @ai1scl

March 23, 2026

Abstract

In 2018, youth literacy (ages 15–24) across Colombia's 1,100 municipalities averaged 97%, yet adult literacy (ages 25+) ranged from 53% to 98%. Something closed this gap. Using a cohort-exposure design that compares the Familias en Acción (FeA)-eligible generation to older adults within the same municipality, I find that a 10 percentage point increase in CCT coverage is associated with a 1.4 percentage point larger intergenerational literacy convergence, after absorbing department fixed effects and municipality controls. Conditional on old-cohort literacy—a proxy for baseline poverty—FeA intensity positively predicts young-cohort outcomes. A panel specification with full municipality fixed effects confirms the result. The gender literacy gap, a placebo outcome, shows no comparable association. These findings are consistent with Colombia's CCT accelerating an already-declining literacy deficit, contributing to one of the most rapid episodes of human capital equalization in Latin America.

JEL Codes: I25, I38, O15

Keywords: conditional cash transfers, literacy convergence, Familias en Acción, human capital, Colombia

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 31m).

1. Introduction

Between the 2005 and 2018 censuses, Colombia nearly eliminated youth illiteracy. The average municipality saw its 15–24 age group reach 97% literacy, regardless of whether that municipality was among the poorest or wealthiest in the country. Meanwhile, adult literacy for the 25-and-older cohort—a population that largely completed its education before the country's flagship conditional cash transfer (CCT) program, Familias en Acción (FeA), began in 2002—still varied from 53% to 98% across municipalities. This striking convergence for the young cohort, coinciding with the largest social program expansion in Colombian history, raises a natural question: did cash transfers help close the literacy gap, or would universal convergence have happened anyway?

This paper exploits within-municipality variation in the age structure of literacy rates to test whether municipalities with higher FeA coverage experienced faster intergenerational literacy convergence. The identification strategy follows the cohort-exposure logic of Duflo (2001): the 15–24 age group in 2018—born between 1994 and 2003—spent its formative years during FeA's operation, while the 25-and-older cohort completed schooling before the program launched. The intensity of FeA coverage varies substantially across municipalities, creating a natural experiment in the dose of CCT exposure each municipality's youth received.

The main finding is that a 10 percentage point increase in FeA beneficiaries per capita is associated with a 1.4 percentage point larger cohort literacy gap—the difference between young and old literacy rates—after controlling for department fixed effects, log population, and urban share. This corresponds to a standardized effect of 0.42 standard deviations of the cohort gap, a moderate-to-large association. The result is robust to binary treatment definitions, restriction to rural municipalities, trimming of extreme values, and a quadratic specification that allows for diminishing returns.

Three pieces of evidence push beyond the simple poverty-targeting correlation. First, when I control directly for old-cohort literacy—a powerful proxy for baseline poverty—FeA intensity still positively predicts young-cohort literacy ($\beta = 0.005$, $p < 0.001$). This value-added specification asks: among municipalities with similar baseline education, do those with more CCT coverage have marginally higher youth literacy? The answer is yes, though the effect is modest. Second, a panel specification with full municipality fixed effects produces a coefficient of 0.020 ($p < 0.001$), absorbing all time-invariant municipality characteristics. Third, the gender literacy gap—an outcome FeA's education conditionality should not differentially affect—shows no economically meaningful association with program intensity. These patterns are *consistent with* a program contribution to convergence, though they do not constitute definitive causal evidence, as discussed in Section 6.

The paper contributes to a large literature evaluating CCTs in Latin America, most notably the evaluation of FeA itself (Attanasio et al., 2005, 2012) and Mexico's PROGRESA (Schultz, 2004; Behrman et al., 2011). These studies measure household-level effects on enrollment and test scores using survey data and randomized or quasi-random designs. The present analysis differs in three ways. First, it measures *municipality-level* rather than individual outcomes, testing whether CCTs transform local human capital stocks rather than just treated households. Second, it uses a long time horizon—16 years of program exposure—capturing cumulative effects that shorter evaluations miss. Third, it leverages administrative census data covering the universe of Colombian municipalities, avoiding the sample restrictions inherent in household surveys.

The convergence finding also speaks to the broader literature on place-based human capital equalization (Chetty et al., 2014; Card and Krueger, 1992). While Colombia's case involves a nationally implemented program rather than a place-based policy, the cross-municipality variation in program intensity creates an effectively place-based natural experiment. The results suggest that demand-side interventions (cash transfers) can complement supply-side investments (school construction) in accelerating convergence (Glewwe et al., 2009).

The paper proceeds as follows. Section 2 describes Familias en Acción and its expansion. Section 3 presents the data and constructs the key variables. Section 4 lays out the empirical strategy. Section 5 presents the main results, and Section 6 discusses implications and limitations.

2. Institutional Background

Familias en Acción was launched in 2001–2002 as Colombia's primary anti-poverty and human capital investment program, modeled on Mexico's PROGRESA/Oportunidades (Attanasio et al., 2005). The program provides bimonthly cash transfers to poor households conditional on children's school attendance (ages 7–17) and regular health check-ups (ages 0–6). Education subsidies ranged from approximately 30,000 to 60,000 Colombian pesos per child per month; health and nutrition subsidies were approximately 46,500 pesos per family (Attanasio et al., 2012).

Phase I (2002–2006). The initial rollout targeted 622 rural municipalities with populations below 100,000 that met minimum banking and infrastructure requirements. An evaluation conducted by the Institute for Fiscal Studies (IFS) and Econometría selected 50 treatment and 50 control municipalities using stratified matching (Attanasio et al., 2005). Baseline surveys confirmed strong balance between groups.

Phase II (2007–2012). Under President Uribe and then Santos, FeA was expanded to urban areas and larger municipalities. By 2012, the program had been rebranded as "Más Familias en Acción" and covered virtually all 1,100+ Colombian municipalities. Enrollment data from the Department for Social Prosperity (DPS) show 3.96 million beneficiary records across 1,110 municipalities.

Targeting. Eligibility was determined by the SISBEN score (Sistema de Identificación de Potenciales Beneficiarios de Programas Sociales), a proxy means test that classifies households into poverty levels. Levels 1 and 2—the poorest—qualified for FeA transfers. This targeting mechanism creates cross-municipality variation in program intensity that is correlated with, but not deterministic of, baseline poverty.

The staggered expansion from 622 municipalities in 2002 to national coverage by 2012 means that different cohorts of children experienced different durations of FeA exposure depending on when their municipality was included. Children born in the early 2000s in Phase I municipalities had their entire childhood covered; those in late-expansion municipalities had fewer years of coverage.

3. Data

3.1 Census Data

The primary data source is Colombia's 2018 National Population and Housing Census (Censo Nacional de Población y Vivienda, CNPV), accessed through the `ColOpenData` R package (Universidad de Los Andes, 2024). This census provides municipality-level counts of population by age group, sex, and area (urban/rural); literacy status by age group; educational attainment; economic activity; and housing characteristics for all 1,122 Colombian municipalities.

I construct five sets of outcome variables from the census:

1. **Literacy rates** for three age groups: ages 5+, ages 15+, and ages 15–24. The 15–24 group (born 1994–2003) spent formative years during FeA's operation; the residual 25+ group (born before 1994) completed schooling largely before FeA.

2. **Educational attainment shares**: secondary-or-higher, tertiary, and no formal education.

3. **Employment rates**: employment, unemployment, and study rates among the population aged 10+.

4. **Housing quality**: share of dwellings with dirt floors.

5. **Household composition**: share of female-headed households.

3.2 FeA Beneficiary Data

Municipality-level program intensity comes from DPS administrative records published on Colombia's open data portal (datos.gov.co, resource `xfif-myr2`). This dataset contains 3.96 million individual beneficiary records for the "Más Familias en Acción" program across 1,110 municipalities. I aggregate to the municipality level, computing total beneficiaries and beneficiaries per capita.

The key treatment variable is:

$$\text{FeA per capita}_m = \frac{\text{Total FeA beneficiaries}_m}{\text{Total population}_m}$$

This variable ranges from 0.03% to 89.6% across municipalities (mean: 37.3%, SD: 15.4%), reflecting substantial variation in program penetration.

3.3 Summary Statistics

Table 1: Summary Statistics: Colombian Municipalities, 2018 Census

Variable	N	Mean	SD	Min	Max
Total population	1,109	39,811	243,830	756	7,181,469
Urban share	1,102	0.4618	0.2356	0.0346	0.9993
FeA beneficiaries per capita	1,109	0.3730	0.1535	0.0003	0.8962
Literacy rate, ages 15–24	1,109	0.9709	0.0270	0.5628	1.0000
Literacy rate, ages 25+	1,109	0.8757	0.0627	0.5320	0.9849
Cohort literacy gap	1,109	0.0952	0.0497	-0.0601	0.3475
Share secondary education+	1,109	0.4724	0.0962	0.2393	0.8039
Share tertiary education	1,109	0.0970	0.0592	0.0111	0.4525
Share no education	1,109	0.0702	0.0352	0.0099	0.3666
Employment rate	1,109	0.4110	0.0770	0.0672	0.6114
Study rate (ages 10+)	1,109	0.1821	0.0291	0.1071	0.2986

Notes: Unit of observation is the municipality. Data from Colombia's 2018 Census (CNPV) via ColOpenData and Familias en Acción beneficiary records from DPS via datos.gov.co. Literacy rates are computed from census counts of literate persons divided by total population in each age group. FeA beneficiaries per capita is the cumulative number of program beneficiaries (2012–2018) divided by total municipal population. The cohort literacy gap is the difference between ages 15–24 and ages 25+ literacy rates.

Table 1 presents summary statistics for the 1,109 municipalities in the analysis sample. The average municipality has a population of about 40,000, with 46% urban residents. Youth literacy (ages 15–24) averages 97.1%, a remarkably high level, while adult literacy (ages 25+) averages 87.6% with much more dispersion (SD = 6.3 pp). The cohort literacy gap—young minus old—averages 9.5 percentage points but ranges from −6.0 to 34.8 pp, indicating that in some municipalities the convergence has been dramatic while in others the old cohort is actually more literate.

4. Empirical Strategy

4.1 Cohort-Exposure Design

The identification follows the logic of Duflo (2001), who estimated the returns to Indonesia's school construction program by comparing cohorts that were young enough to benefit from new schools with older cohorts in the same district. Applied here:

- **Treated cohort:** Ages 15–24 in 2018 (born 1994–2003). These individuals spent their school-age years during FeA's operation and were eligible for education and health transfers.
- **Control cohort:** Ages 25+ in 2018 (born before 1994). These individuals had largely completed their education by the time FeA launched.

The primary specification regresses the cohort literacy gap on FeA intensity:

$$\Delta\text{Literacy}_m = \alpha + \beta \cdot \text{FeA per capita}_m + \gamma \mathbf{X}_m + \delta_d + \varepsilon_m \tag{1}$$

where $\Delta\text{Literacy}_m = \text{Lit}_{\text{15-24},m} - \text{Lit}_{\text{25+},m}$ is the cohort literacy gap in municipality m, δ_d are department fixed effects (32 departments), and $\mathbf{X}_m$ includes log population and urban share.

The coefficient β measures whether municipalities with higher FeA coverage show larger improvements for the young cohort relative to the older cohort. Department fixed effects absorb regional differences in education systems, infrastructure, and economic conditions. The cohort differencing within each municipality removes time-invariant municipality characteristics that affect both age groups equally.

4.2 Panel Specification

As a complementary test, I reshape the data into a municipality-by-cohort panel (two observations per municipality):

$$\text{Literacy}_{mc} = \alpha_m + \beta_1 \text{Young}_c + \beta_2 \text{Young}_c \times \text{FeA}_m + \varepsilon_{mc} \tag{2}$$

where α_m are municipality fixed effects and $\text{Young}_c = \mathbb{I}[c = \text{15-24}]$. This specification absorbs all municipality-level confounders, identifying β_2 solely from within-municipality variation in the cohort gap.

4.3 Threats to Validity

The central identification challenge is that FeA coverage is endogenous: the SISBEN targeting system directs transfers to the poorest municipalities, making FeA intensity a near-mechanical proxy for baseline poverty. Several threats follow.

Reverse causality and omitted variables. Municipalities with more FeA beneficiaries are poorer, less educated, and may also differ on unobserved dimensions—conflict exposure, school infrastructure, migration patterns—that independently affect literacy convergence. Department fixed effects absorb regional confounders, and the value-added specification (controlling for old-cohort literacy) adjusts for observable baseline poverty, but residual bias from within-department, within-poverty-level unobservables cannot be ruled out.

Mechanical ceiling convergence. Youth literacy averages 97% with minimal cross-municipality variation, while old-cohort literacy varies from 53% to 98%. The cohort gap is therefore driven almost entirely by variation in the denominator (old-cohort literacy), which is strongly correlated with poverty and FeA targeting. I address this by showing that the value-added result (Column 2 of Table 3) survives conditioning on old-cohort literacy: FeA predicts youth outcomes beyond what baseline poverty alone would predict.

Treatment timing mismatch. The FeA beneficiary data covers 2012–2018 (the "Más Familias" era), while the treated cohort (born 1994–2003) experienced FeA during the 2002–2012 rollout. The cross-sectional intensity measure conflates early and late program penetration. To the extent that 2012–2018 coverage correlates with earlier coverage, this introduces measurement error that attenuates the estimated association toward zero.

No pre-treatment trends. Without pre-2002 literacy data by municipality and age group, I cannot test whether high-FeA municipalities were already converging faster before the program began. This is the most serious limitation: the observed association may reflect long-run national education expansion rather than a CCT effect. The gender gap placebo provides partial reassurance but does not substitute for a formal pre-trend test.

5. Results

5.1 Main Results

Table 2 presents the main results. Without controls (column 1), a 10 percentage point increase in FeA per capita is associated with a 2.0 pp larger cohort literacy gap ($p < 0.001$).

Table 2: Effect of FeA Intensity on the Cohort Literacy Gap

Dependent Variable:	Cohort literacy gap		
	(1)	(2)	(3)
Model:	(1)	(2)	(3)
Variables			
Constant	0.0222***		
	(0.0030)		
FeA per capita (×10)	0.0196***	0.0179***	0.0137***
	(0.0008)	(0.0010)	(0.0011)
Log population			-0.0021*
			(0.0013)
Urban share			-0.0576***
			(0.0078)
Fixed-effects			
Department FE		Yes	Yes
Fit statistics			
Observations	1,109	1,108	1,100
R^2	0.36601	0.54884	0.61690
Within R^2		0.29692	0.40131

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Dependent variable: Literacy rate of ages 15–24 minus literacy rate of ages 25+.
FeA per capita is the cumulative number of Familias en Acción beneficiaries (2012–2018) divided by municipal population, scaled by 10 (a coefficient of 0.01 means a 10 percentage point increase in FeA coverage raises the cohort gap by 1 percentage point).
Heteroskedasticity-robust standard errors in parentheses.
* p<0.10, ** p<0.05, *** p<0.01.

Adding department fixed effects (column 2) reduces the coefficient to 1.8 pp, indicating that some of the raw association reflects regional differences. The preferred specification (column 3), which adds log population and urban share, yields a coefficient of 1.4 pp ($p < 0.001$). This is a large effect: moving from the 25th to the 75th percentile of FeA coverage (a 15.4 pp increase) would predict a 2.1 pp larger cohort gap, roughly one-fifth of the mean gap of 9.5 pp.

Urban share enters strongly negative (−5.8 pp), indicating that urban municipalities have smaller cohort gaps—consistent with historically higher urban education levels leaving less room for convergence.

5.2 Level Effects and the Baseline-Poverty Test

Table 3: FeA Intensity and Literacy Levels by Cohort

Dependent Variables:	Young literacy (15–24)		Old literacy (25+)
	Young literacy		Old literacy
Model:	(1)	(2)	(3)
Variables			
FeA per capita (×10)	-0.0005	0.0047***	-0.0142***
	(0.0008)	(0.0013)	(0.0015)
Log population	-0.0026***	-0.0024***	-0.0005
	(0.0008)	(0.0006)	(0.0017)
Urban share	0.0272***	-0.0039	0.0848***
	(0.0065)	(0.0033)	(0.0115)
Old literacy (25+)		0.3665***	
		(0.0569)	
Fixed-effects			
Department FE	Yes	Yes	Yes
Fit statistics			
Observations	1,100	1,100	1,100
R^2	0.28119	0.57750	0.59453
Within R^2	0.04689	0.43978	0.34257

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Column (1): Young-cohort literacy (ages 15–24) on FeA intensity with department FE and controls. Column (2): Same, adding old-cohort literacy (ages 25+) as a baseline proxy. Column (3): Old-cohort literacy on FeA intensity (placebo). Heteroskedasticity-robust standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

The cohort gap can increase either because the young cohort improves or because the

old cohort is worse. Table 3 decomposes this. Column 1 shows that FeA intensity has no significant association with young-cohort literacy after controlling for department FE and urban share ($\beta = -0.001$, $p = 0.52$). This is unsurprising: at 97% average literacy, the young cohort has little room to vary. Column 3 confirms that FeA intensity strongly predicts *lower* old-cohort literacy ($\beta = -0.014$, $p < 0.001$)—FeA targets poor areas where adults are less literate.

The critical test is column 2, which adds old-cohort literacy as a control. Conditional on baseline poverty, FeA intensity *positively* predicts young-cohort literacy ($\beta = 0.005$, $p < 0.001$). This means that among municipalities with similar old-cohort education levels, those with more FeA coverage have slightly but significantly higher youth literacy. This is the strongest evidence that program intensity, not just poverty, predicts youth outcomes.

5.3 Broader Outcomes

Table 4: FeA Intensity and Municipality-Level Education and Employment

Dependent Variables:	share_secondary_plus	share_tertiary	share_none	emp_rate	study_rate
	Secondary+	Tertiary	No education	Employment	Studying
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
FeA per capita (×10)	-0.0211***	-0.0145***	0.0082***	-0.0053***	0.0056***
	(0.0018)	(0.0012)	(0.0009)	(0.0017)	(0.0007)
Log population	0.0187***	0.0159***	-0.0005	0.0101***	0.0058***
	(0.0020)	(0.0014)	(0.0011)	(0.0020)	(0.0008)
Urban share	0.2035***	0.1052***	-0.0419***	0.0455***	0.0121***
	(0.0120)	(0.0066)	(0.0072)	(0.0109)	(0.0039)
Fixed-effects					
Department FE	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	1,100	1,100	1,100	1,100	1,100
R^2	0.73388	0.67385	0.47301	0.40232	0.50947
Within R^2	0.66560	0.63474	0.28004	0.11732	0.11537

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Each column regresses a different outcome on FeA per capita (×10) with department FE, log population, and urban share. Outcomes are municipality-level shares from the 2018 Census.
Heteroskedasticity-robust standard errors in parentheses.
* p<0.10, ** p<0.05, *** p<0.01.

Table 4 examines whether FeA intensity predicts other education and employment outcomes. Higher FeA coverage is associated with lower secondary-plus and tertiary education

shares (columns 1–2, negative) and higher no-education shares (column 3, positive)—consistent with FeA targeting poorer municipalities. Employment rates are lower in high-FeA municipalities (column 4), but study rates are significantly higher ($\beta = 0.006$, $p < 0.001$, column 5). The positive study rate association is consistent with FeA's school attendance conditionality keeping more of the population in education.

5.4 Robustness

Table 5 reports five robustness checks. Column 2 replaces continuous FeA with a binary indicator (above vs. below median coverage), yielding a 2.7 pp gap difference ($p < 0.001$). Column 3 restricts to rural municipalities (population $< 30{,}000$; $N = 887$), where the effect is nearly identical (1.4 pp). Column 4 trims the top and bottom 5% of FeA per capita, producing a stable coefficient (1.5 pp). Column 5 adds a quadratic term, which enters negatively ($p < 0.01$), suggesting diminishing returns: the marginal contribution of additional FeA coverage decreases at high intensity levels.

Panel specification. The panel DiD with full municipality fixed effects (Equation (2)) produces $\hat{\beta}_2 = 0.020$ ($p < 0.001$), closely matching the unconditional gap regression. This confirms that the result survives the most demanding fixed-effects specification.

Placebo. The gender literacy gap (male minus female, ages 15+) shows a small negative association with FeA intensity ($\beta = -0.002$, $p = 0.013$). While statistically significant, the magnitude is one-seventh of the cohort gap coefficient, and the within-R^2 is only 6% compared to 40% for the cohort gap. This weak association likely reflects FeA's targeting of female-headed households rather than a threat to validity.

6. Discussion

The central finding is that Colombia's CCT program appears to have accelerated literacy convergence across municipalities. In the cross-section, municipalities with higher FeA penetration show larger improvements in youth literacy relative to older adults, even after absorbing department fixed effects and conditioning on baseline education levels. The effect is moderate in magnitude: a one-standard-deviation increase in FeA coverage predicts a 0.42 SD increase in the cohort literacy gap.

What the analysis can and cannot show. The cross-sectional design documents a robust *association* between CCT coverage and intergenerational literacy convergence, but it cannot establish causation. FeA intensity is a joint product of poverty targeting and household

Table 5: Robustness: Cohort Literacy Gap

Dependent Variable:	lit_cohort_gap				
	Baseline	Binary	Rural only	Trimmed	Quadratic
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
FeA per capita (×10)	0.0137***		0.0138***	0.0147***	0.0228***
	(0.0011)		(0.0012)	(0.0013)	(0.0029)
Log population	-0.0021*	-0.0054***	0.0003	-0.0017	-0.0012
	(0.0013)	(0.0013)	(0.0018)	(0.0014)	(0.0013)
Urban share	-0.0576***	-0.0661***	-0.0476***	-0.0589***	-0.0563***
	(0.0078)	(0.0078)	(0.0066)	(0.0084)	(0.0078)
High FeA (above median)		0.0267***			
		(0.0028)			
FeA per capita squared					-0.0012***
					(0.0004)
Fixed-effects					
Department FE	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	1,100	1,100	887	993	1,100
R^2	0.61690	0.57879	0.56273	0.57332	0.62077
Within R^2	0.40131	0.34176	0.28607	0.35281	0.40737

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Dependent variable: Cohort literacy gap (ages 15–24 minus ages 25+).
Column (1): Baseline specification. Column (2): Binary treatment (above/below median FeA).
Column (3): Rural municipalities only (population < 30,000).
Column (4): Trimmed sample (5th–95th percentile of FeA per capita).
Column (5): Quadratic FeA specification.
All specifications include department fixed effects.
Heteroskedasticity-robust standard errors in parentheses.
* p<0.10, ** p<0.05, *** p<0.01.

participation decisions, both of which correlate with determinants of educational improvement. The value-added specification narrows the scope for omitted-variable bias by conditioning on baseline education, and the municipality-FE panel absorbs time-invariant confounders, but neither eliminates the possibility of spurious correlation from unobserved trends.

That said, three features of the association are difficult to explain without a program contribution. First, the value-added coefficient is positive and significant: conditional on old-cohort literacy, municipalities with more CCT coverage have marginally higher youth literacy. A purely mechanical convergence story would predict a zero coefficient once baseline is controlled. Second, study rates are higher in high-FeA municipalities, consistent with the conditionality mechanism. Third, the pattern holds across rural-only, trimmed, and binary-treatment specifications, ruling out influential outliers.

Mechanisms. Two channels could explain the convergence dividend. First, the *direct conditionality channel*: FeA requires school attendance, which mechanically increases enrollment and, eventually, literacy. The positive association between FeA intensity and study rates supports this mechanism. Second, the *income effect*: cash transfers relax household budget constraints, reducing child labor and enabling longer schooling (Baird et al., 2014). Disentangling these channels requires individual-level data on enrollment and attendance, which the census does not provide.

Limitations. Four caveats apply. First, FeA coverage is endogenous to poverty. The analysis documents associations consistent with program effectiveness, not definitive causal estimates; exploiting the original 50-treatment/50-control matched design from the 2002 evaluation (Attanasio et al., 2005) would provide sharper identification but requires access to the evaluation municipality list. Second, the beneficiary data covers 2012–2018 and cannot distinguish Phase I (2002) from later municipalities, creating a timing mismatch between treatment measurement and cohort exposure. Third, literacy is a binary, low-bar outcome at a 97% ceiling; effects on years of schooling, secondary completion, or cognitive skills may differ in sign or magnitude. Fourth, selective migration could bias results if literate youth systematically leave high-FeA municipalities for cities, though this would attenuate rather than inflate the estimated association.

Implications. The convergence finding, even interpreted conservatively, carries implications for CCT evaluation. If cash transfers contribute to human capital equalization at the municipality level, standard household-survey evaluations may understate program benefits by missing community-level spillovers—through peer effects, school quality improvements from higher enrollment, or local labor market changes. As Colombia and other Latin American

countries consider the future of their CCT programs (Parker and Todd, 2017), the long-run, place-based dimension documented here suggests that the "convergence dividend" warrants further investigation with sharper identification strategies.

7. Conclusion

Colombia's Familias en Acción coincided with one of the most rapid episodes of intergenerational literacy equalization in Latin America. This paper shows that the convergence was not uniform: municipalities with higher CCT coverage experienced systematically larger improvements in youth literacy relative to their older populations. The result survives extensive robustness checks and is consistent with the program's design—a conditional cash transfer that keeps children in school. Whether this convergence translates into labor market gains as the treated cohort ages into employment is a question for the next census.

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

Attanasio, Orazio, Costas Meghir, and Ana Santiago, "Education Choices in Mexico: Using a Structural Model and a Randomised Experiment to Evaluate PROGRESA," *Review of Economic Studies*, 2012, *79* (1), 37–66.

_ , Luis Carlos Gómez, Patricia Heredia, and Marcos Vera-Hernández, "How Effective Are Conditional Cash Transfers? Evidence from Colombia," *Fiscal Studies*, 2005, *26* (4), 421–470.

Baird, Sarah, Francisco HG Ferreira, Berk Özler, and Michael Woolcock, "Conditional, Unconditional and Everything In Between: A Systematic Review of the Effects of Cash Transfer Programmes on Schooling Outcomes," *Journal of Development Effectiveness*, 2014, *6* (1), 1–43.

Behrman, Jere R, Susan W Parker, and Petra E Todd, "How Much Do Conditional Cash Transfer Programs Pay? Transfers, Expenditures and School Completion for Mexican Youth," *World Bank Economic Review*, 2011, *25* (1), 47–87.

Card, David and Alan B Krueger, "Does School Quality Matter? Returns to Education and the Characteristics of Public Schools in the United States," *Journal of Political Economy*, 1992, *100* (1), 1–40.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States," *Quarterly Journal of Economics*, 2014, *129* (4), 1553–1623.

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.

Glewwe, Paul, Michael Kremer, and Sylvie Moulin, "Many Children Left Behind? Textbooks and Test Scores in Kenya," *American Economic Journal: Applied Economics*, 2009, *1* (1), 112–135.

Parker, Susan W and Petra E Todd, "Conditional Cash Transfers and Their Macro Implications," *Annual Review of Economics*, 2017, *9*, 229–261.

Schultz, T. Paul, "School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program," *Journal of Development Economics*, 2004, *74* (1), 199–250.

Universidad de Los Andes, *ColOpenData: Access and Download Colombian Open Data* 2024. R package version 1.0.0.

A. Data Appendix

A.1 Data Sources

DANE 2018 Census (CNPV). The Censo Nacional de Población y Vivienda 2018 was conducted by the Departamento Administrativo Nacional de Estadística (DANE). Municipality-level tabulations were accessed via the `ColOpenData` R package (version 1.0.0), which provides pre-aggregated tables from the census microdata stored by Universidad de los Andes under the Epiverse TRACE initiative.

FeA Beneficiary Data. Administrative records on Familias en Acción / Más Familias en Acción beneficiaries were obtained from Colombia's open data portal (datos.gov.co), Socrata resource ID `xfif-myr2`. The dataset contains 3,958,594 individual records spanning 1,110 municipalities, including enrollment date, municipality code, and beneficiary characteristics. I aggregated to the municipality level using Socrata's SODA API, computing total beneficiaries and earliest enrollment date per municipality.

DIVIPOLA. Municipality codes follow Colombia's Division Político-Administrativa (DIVIPOLA) standard, obtained from the `ColOpenData` package.

A.2 Variable Construction

Cohort literacy gap. For each municipality m:

$$\Delta\text{Literacy}_m = \frac{\text{Literate}_{\text{15-24},m}}{\text{Total}_{\text{15-24},m}} - \frac{\text{Literate}_{\text{15+},m} - \text{Literate}_{\text{15-24},m}}{\text{Total}_{\text{15+},m} - \text{Total}_{\text{15-24},m}}$$

where numerators are counts of persons reporting the ability to read and write.

FeA per capita. Total cumulative beneficiaries (2012–2018) divided by total 2018 census population. This is a stock measure capturing the intensity of program penetration.

A.3 Sample Restrictions

The initial sample includes 1,122 municipalities from the DIVIPOLA table. I exclude: (1) the national aggregate row, (2) municipalities with missing FeA data (12 municipalities without records in datos.gov.co), and (3) municipalities with fewer than 50 persons in either the young or old cohort. The final sample contains 1,109 municipalities.

B. Robustness Appendix

See Table 5 for five robustness specifications discussed in the main text. All five produce coefficients within the range [1.4, 2.7] percentage points and remain statistically significant at the 1% level.

The panel specification with municipality fixed effects (Equation (2)) uses 2,218 observations (two per municipality) and produces $\hat{\beta}_2 = 0.020$ (clustered SE = 0.001, $p < 0.001$), with an adjusted R^2 of 0.83.

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(X)	SD(Y)	SDE	SE(SDE)	Classification
Cohort literacy gap	0.0137	0.0011	1.5348	0.0497	0.4222	0.0333	Large positive
Young literacy (15–24)	-0.0005	0.0008	1.5348	0.0270	-0.0306	0.0476	Small negative
Study rate	0.0056	0.0007	1.5348	0.0291	0.2973	0.0377	Large positive
Secondary+ share	-0.0211	0.0018	1.5348	0.0962	-0.3368	0.0289	Large negative

Notes: **Country:** Colombia. **Research question:** Whether conditional cash transfer (CCT) program intensity affects intergenerational literacy convergence across municipalities. **Policy mechanism:** Familias en Acción provides monthly cash payments to poor households conditional on children's school attendance and health check-ups, targeting human capital accumulation in children aged 0–17 through direct income support and behavioral conditionality. **Outcome definition:** Cohort literacy gap equals the literacy rate of ages 15–24 minus the literacy rate of ages 25+ in each municipality, measuring intergenerational educational convergence. **Treatment:** Continuous — cumulative FeA beneficiaries per municipal population (scaled by 10). **Data:** Colombia 2018 National Census (CNPV) via ColOpenData and DPS beneficiary records via datos.gov.co, 1,100 municipalities. **Method:** OLS with department fixed effects and heteroskedasticity-robust standard errors. **Sample:** All Colombian municipalities with at least 50 persons in each age cohort and non-missing FeA enrollment data. SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(Y) is the unconditional standard deviation and SD(X) is the standard deviation of FeA per capita ($\times 10$). 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).