

Conditional Cash Transfers and the Geography of Intergenerational Education Mobility in the Philippines*

Teresa Molina[†]

Cara Tan[‡]

August 21, 2026

Abstract

Using population census data, this paper documents substantial variation in municipality-level measures of upward and relative education mobility across the Philippines. Regressions that control for a rich set of covariates leave one-fifth of the variation in upward mobility and two-thirds of the variation in relative mobility unexplained. To explore the extent to which these measures capture meaningful variation that can be changed by government policy, we exploit municipality-level variation in exposure to a national conditional cash transfer program. Exposure to the program significantly increased average municipality education and upward mobility but had no significant effects on relative mobility.

JEL Codes: J62, I24, I38

Keywords: intergenerational mobility, conditional cash transfers, education, Philippines

*Many thanks to seminar participants at UP Diliman, USC CESR, UCSD, University of Tokyo, Keio-UH Workshop on Family Economics, SEA Annual Meeting, Hitotsubashi University, ADB, UH Manoa, and Center for Global Development for helpful feedback on the paper. We also thank Julien Labonne for sharing the municipality-level 4Ps household counts.

[†]University of Hawaii at Manoa, IZA, UHERO; tmolina@hawaii.edu

[‡]California Policy Lab

1 Introduction

Intergenerational mobility – the degree to which individuals’ socioeconomic outcomes are independent of and able to surpass those of their parents’ – is an important metric that reflects equality of opportunity in a society. Levels of intergenerational mobility vary widely across countries (Neidhöfer et al., 2018; Van der Weide et al., 2024) and, strikingly, also within countries (Alesina et al., 2021; Asher et al., 2024; Bailey et al., 2026; Chetty et al., 2014). This within-country variation has important implications due to the importance of neighborhood-level mobility: children who spend more time in high-mobility areas end up with better socioeconomic outcomes in the long run (Alesina et al., 2021; Chetty and Hendren, 2018).

In this paper, we examine intergenerational education mobility across municipalities in the Philippines. Using full count census data from 2000, 2007, 2010, and 2015, we calculate municipality-level measures of education mobility. One measure is the slope in a regression of child education rank on parental education rank, which captures decreasing relative mobility or increasing relative persistence: steeper slopes indicate that a child’s educational attainment is tightly linked with their parents’. We capture the phenomenon of upward mobility by using the same regression to predict the education rank of a child born to parents in the bottom half of the education distribution. Specifically, our upward mobility measure is the predicted percentile rank of children born to parents in the 25th percentile of the national education distribution of adults. This captures the extent to which children born to parents on the lower end of the education distribution are able to rise up the ranks.

Like previous work that examines within-country variation in regional measures of mobility (Alesina et al., 2021; Bailey et al., 2026; Chetty et al., 2014), we document substantial variation in upward and relative mobility across Filipino municipalities. In regressions that predict upward mobility using a rich set of municipality characteristics, socioeconomic conditions explain a large share of the municipality-level variation, but approximately one-fifth of the variation remains unexplained – even after controlling for average child education in the municipality. In other words, some municipalities appear to have better upward mobility simply because households are better-off socioeconomically and children are able to achieve higher levels of education overall, but there remain many unobserved factors driving the variation in upward mobility across municipalities. This is true to an even greater degree for relative mobility: around two-thirds of the variation in the rank-rank slope remains unexplained after controlling for average educational attainment, assets, and other demographic conditions and spatial attributes.

The large amount of unexplained variation in mobility across municipalities could be an indication of two different possibilities: either that other meaningful omitted factors – like government policy – play a key role in generating this geographic variation, or that the variation is driven primarily by random noise. If the

latter were true, this would leave little room for government policy to shape the geographic distribution of mobility in a country. It would also indicate that these municipality-level measures of mobility, which are becoming increasingly popular in the economics literature, may have limited policy relevance. We therefore investigate whether government policy has any effect on these local measures of mobility.

Because social protection policies have been theorized to be effective tools for increasing intergenerational mobility (Alvaredo et al., 2018; Black and Devereux, 2011; Narayan et al., 2018; World Bank, 2022; World Bank Group, 2022), we examine a Filipino conditional cash transfer (CCT) program called *Pantawid Pamilyang Pilipino Program*, also known as the 4Ps program. CCT programs, which provide transfers to poor households who satisfy certain education and health requirements, are a common form of social protection, now used widely in the developing world. Many studies find significant improvements in educational attainment for those exposed to CCTs, typically those of lower socioeconomic status, during childhood (Araujo and Macours, 2021; Millán et al., 2019; Parker and Todd, 2017). Given this, along with the fact that program intensity typically varies across regions, CCTs certainly have the potential to alter the geographic distribution of education mobility – but by how much is unclear.

To estimate the effects of the 4Ps program on local intergenerational mobility, we use the fact that the program was rolled out to different municipalities in different years, from 2008 to 2013. This generates variation in the number of years a municipality has been exposed to the program in any given year. We are interested in the average causal response (ACR), as defined by Callaway et al. (2026), to this multi-valued treatment variable (years of exposure to the 4Ps program). We begin with a two-way fixed effects (TWFE) regression that includes both 2015 and 2007 (as a pre-program year) and controls for municipality and year fixed effects. In addition, because Callaway et al. (2026) have shown that this TWFE estimate is a weighted average of causal responses that uses counter-intuitive weights, we also use the methods proposed in their paper to recover a more intuitively weighted ACR. The validity of both estimates relies on the strong parallel trends assumption: that average potential outcomes among all municipalities would have evolved in parallel with those that were actually exposed to a given level of exposure by 2015, for all possible values of exposure.

We estimate the effect of 4Ps exposure on three outcomes of interest: average educational attainment, upward mobility, and relative mobility as measured by the intergenerational slope, all calculated at the municipality level. Based on our preferred ACR estimates, one additional year of 4Ps exposure led to a 0.03 year increase in average municipality-level educational attainment and a 0.31 unit increase in the average percentile rank of children born to parents at the 25th percentile (i.e., an increase in upward mobility). Our preferred specification detects no statistically significant effects on the intergenerational slope, which could reflect the noisier nature of this measure or its lack of sensitivity to changes that affect a narrow window of the parental education distribution. Scaling by the average length of 4Ps exposure in our 2015 sample (5.2

years), the effect on educational attainment translates into an effect that is 19% of the inter-quartile range, while the effect on upward mobility translates into an effect size approximately 17% of the inter-quartile range. Effects are similar for boys and girls.

Our main estimates allow for different outcome trajectories in municipalities with different poverty levels at baseline. This is important because the 4Ps was targeted to poorer municipalities first. If poorer municipalities would have improved faster even in the absence of the program, this would result in an overestimation of the effect of the 4Ps. In addition, we examine whether municipalities with greater 4Ps exposure were on different trajectories compared to those with less 4Ps exposure prior to the start of the program. The results of this analysis suggest that our main results – improvements in educational attainment and upward mobility – are unlikely to be driven by differential pre-trends continuing into the program period.

This paper contributes to the literature on the effects of government policies on intergenerational mobility. While many papers explore correlations between local mobility measures and various characteristics, they do not investigate the causal determinants of this local variation (Alesina et al., 2021; Asher et al., 2024; Chetty et al., 2014; Fan et al., 2021). That said, there is substantial causal evidence on the effects of government policy on national-level mobility. Using high-quality administrative data in high-income countries, researchers have documented policies like school tracking reforms, school spending reforms, and universal childcare can generate improvements in intergenerational mobility, measured using both income and education (Biasi, 2023; Havnes and Mogstad, 2015; Lange and von Werder, 2017; Pekkarinen et al., 2009). In developing countries, where high quality administrative income data are rare, studies have focused primarily on education mobility. Both compulsory schooling laws (Cornelissen and Dang, 2022) and sanitation programs (Wang and Zou, 2023) have been found to generate improvements. In addition to national mobility, however, government policy can also affect the spatial distribution of mobility within a country, which is important given the persistent effects of childhood neighborhood-level mobility. This paper’s novelty, therefore, lies in its focus on local mobility, paired with a causal investigation of the effects of CCTs, one of the most common forms of social protection in the developing world.

The finding that the 4Ps program significantly improves a municipality’s level of upward mobility has two important implications for the literature on within-country variation in mobility. First, it indicates that the observed variation in upward mobility across municipalities is unlikely to be predominantly driven by selection, or the sorting of different types of families into different municipalities, because government policy is capable of changing these measures. Other papers have taken a different approach to exploring selection issues by asking whether neighborhoods have causal effects on children (comparing the outcomes of those who move neighborhoods at different ages) and also conclude that selection is not the full story (Alesina et al., 2021; Chetty and Hendren, 2018). A related implication is that these local measures of upward mobility are

useful, in the sense that they capture variation that can actually be moved by policy. This was not obvious ex-ante, as random shocks or measurement error could in theory be the primary drivers of variation. While this could still be the case for the intergenerational slope, for which we obtained imprecise estimates of the 4Ps effect, for upward mobility it is clear that policies like the 4Ps have the potential to alter its geographic distribution.

Finally, this paper also contributes to the literature on the 4Ps program, which has been found to have had generally positive effects on health and education outcomes despite some unintended negative consequences (De Hoop et al., 2019; Filmer et al., 2023; Kandpal et al., 2016; World Bank Group, 2012). Despite the existence of many evaluations, however, it is still not clear what the program has done for intergenerational education mobility in the Philippines – primarily because few studies examine educational attainment (as opposed to enrollment or attendance) as an outcome.

In the next section, we describe the data and methods used to construct our various measures of mobility and provide a descriptive analysis of these measures. Section 3 describes the institutional background and rollout data related to the 4Ps program. In section 4, we describe the empirical strategy used to estimate the effects of the 4Ps on municipality-level mobility. Section 5 presents the results and section 6 concludes.

2 Education Mobility Measures

2.1 Census Data

This paper relies on full population censuses of the Philippines, conducted by the Philippines Statistics Authority in 2000, 2007, 2010, and 2015. For each household, basic demographic information – including gender, age, relationship to the household head, and educational attainment – is recorded for each household member. Because the census does not collect information about income, we focus on intergenerational education mobility, which is positively correlated with intergenerational income mobility across countries, especially in developing countries (Narayan et al., 2018). Although educational attainment may be a more narrow measure of socioeconomic status than lifetime income or consumption, it is an important causal driver of these more comprehensive outcomes of interest (Card, 1999). Focusing on education mobility also has several advantages, including less measurement error and the ability to use data from younger individuals who have completed (most of) their schooling.

Measuring intergenerational education mobility requires knowledge about the educational attainment of the parents of the individuals in the generation of interest. Unfortunately, the census does not explicitly ask about parental education levels, but we are able to obtain parental education for children of household heads

by linking within households using the “relationship to household head” variable. This allows us to link a large share of children to their parents. As we show in Appendix Figure A1, linkage rates are over 80% for children aged 16 and younger but less than 50% for 24-year-olds.

Based on this approach, a given child may be missing parental education for two main reasons: either they are not the child of the household head (but their parents are still in the household), or they do not live in the same household as either of their parents (because both parents have a different permanent residence or have passed away).¹ Older children may be unlinked if they move out of the house for school or work. Moving out of the household for university is particularly problematic as this results in the systematic exclusion of children at the higher end of the education distribution. Because the last census in our analysis (2015) took place before the Philippines’ shift from 10 years to 12 years of primary and secondary schooling, the typical student in our study period starts university at 17.² Therefore, our main analysis focuses on the age 16 cohort, the oldest cohort for whom school-related moves are unlikely. We restrict to the oldest feasible cohort because they have the most complete – though admittedly still incomplete – measures of educational attainment. As reported in Figure A1, an average of 84% of 16-year-olds are linked across census years.

In Appendix section D, we describe an alternative approach that yields higher linkage rates, by linking children (a younger generation) to anyone in the previous generation (as opposed to parents specifically), as Alesina et al. (2021) do. This captures all children who are unlinked simply because they are not children of the household head (the first group described above) and a subset of those whose parents do not live in the same household (the second group). As reported in Figure D1, 88% of 16-year-olds are matched using this method, which we use in robustness checks.

2.2 National Education Mobility

We begin with the intergenerational rank-rank slope: the slope coefficient in a regression of a child’s age-specific educational rank on the maximum of both parents’ educational rank.³ This captures increasing intergenerational persistence, or decreasing relative mobility. A larger positive slope indicates parent and child outcomes that are more tightly linked, which implies a larger gap between the average ranks of children born to parents at the top and bottom of the distribution. Lower values imply a smaller gap between the top and bottom, which reflects higher relative mobility. A flatter rank-rank slope is not necessarily a desirable outcome for all, as it may reflect worse outcomes for those at the upper end of the distribution. However, a flat rank-rank slope does not necessarily imply that outcomes of children of high-rank parents are worsening

¹Overseas Filipino workers (OFWs) are included in the census household roster and therefore can be linked to their children.

²Across the four census waves, an average of 17% of 17-year-olds had completed one year of college, compared to only 5% of 16-year-olds.

³This measure is highly correlated with measures that use the mean of the parents’, only the father’s, or only the mother’s educational attainment instead (all correlations greater than 0.95).

in absolute terms; they may be improving in absolute terms, just less than children of lower-rank parents.

Our second measure captures upward mobility, or the ability of children from lower-educated parents to rise up the ranks. Specifically, we calculate the average educational percentile rank (in the age-specific national education distribution) of children born to parents at the 25th percentile of the adult education distribution, predicted using the rank-rank regression described above. This measure is the education analog of the “absolute upward mobility” measure used by Chetty et al. (2014) in their examination of income mobility and similar in spirit to “upward interval mobility” in education used by Asher et al. (2024).⁴

Figure 1 plots the relationship between child education rank and parent education rank for each of the four census years. Due to the discrete nature of the educational attainment variable, large shares of both parents and children have the exact same rank. In these cases, we assign everyone in the same category to the mean rank. Despite our use of a discrete variable, we find – like Chetty et al. (2014) do with income – a linear relationship between child and parent rank, with an R^2 (from a linear regression of the scatterplot points – average child rank on parental rank) close to 1, in all years.⁵ Rank-rank slopes hovered between 0.46 and 0.49 from 2000 to 2007, and then fell to 0.42 in 2015. This indicates an increase in relative mobility in these last five years. Upward mobility, which hovered between 38.1 and 38.4 for the first three census waves, also improved in the last five years, increasing to 39.7.

As described above, one key limitation is the need to restrict to coresident parent-child pairs, which is the reason we focus on the age 16 cohort. By age 16, many children have not yet reached their final level of education, which means we may not be capturing true mobility in terms of completed education. To a certain extent, the use of percentile ranks instead of raw schooling levels helps alleviate these concerns, provided that the ranking of children according to their education level at age 16 is similar to the ranking according to their final educational attainment. To generate some insight into whether we are likely to be over- or under-estimating education mobility due to our age 16 restriction, Figure A2 illustrates the same rank-rank slope figures for each age from 15 to 18 (restricting to the 2015 census). Both relative and upward mobility decrease with age: the intergenerational slope is 0.37 at age 15 and 0.53 at age 18, while upward mobility falls from 41.1 to 38.3. This suggests that our primary measures of education mobility, based on attainment at age 16, are likely to be higher than measures that use the final level of completed schooling, though it is important to note that the age gradient in this figure could also be driven by the decreasing share of linked children at higher ages. In our empirical analysis of the 4Ps program, we use the age 16 cohort for our main

⁴As Chetty et al. (2014) argue with respect to their absolute upward mobility measure, even though education rank is technically a relative outcome, the municipality-level average education rank in the national distribution is essentially an absolute outcome given that changes in the education distribution of one municipality are unlikely to affect the national distribution. To avoid confusion, however, we refrain from using the word “absolute” to describe our upward mobility measure.

⁵This was also documented by Bailey et al. (2026), who examine intergenerational education mobility in the United States for birth cohorts in the first and second half of the 20th century.

analysis but show robustness to expanding the sample to ages 15-18.

2.3 Within-Country Variation in Local Education Mobility

The focus of this paper is on municipality-level education mobility. Theoretically, education mobility might vary across municipalities and cohorts if they face different conditions that weaken the link between parent and child outcomes (thereby making it easier for children of lower-ranked parents to rise up the ranks). Changes in mobility at the local level have implications for national mobility. As outlined in Appendix B, the national rank-rank slope can be decomposed into two components: one that involves a weighted average of municipality-level rank-rank slopes and one that involves a cross-unit covariance term. Any improvements in local relative mobility (holding cross-unit covariances constant) will generate improvements in national relative mobility.

In each census year, we obtain the rank-rank slope for municipality j by regressing child education rank (in the age-specific national education distribution) onto parental education rank (in the national adult distribution), restricting to children aged 16 living in municipality j . Based on the decomposition in Appendix B, we can calculate what share of the national rank-rank slope is explained by these municipality-specific measures of relative mobility (as opposed to cross-municipality variation in mean child and parent education). The within-municipality component constitutes between 70-74% of the national rank-rank slope, as reported in Figure 1.

We also calculate upward mobility for each municipality using this same regression. We predict the average percentile rank in the national age-specific education distribution of children born to parents at the 25th percentile of the national adult education distribution. Unlike with national measures (where the rank-rank slope essentially determines the level of upward mobility), our local measure of upward mobility does provide additional information on top of the rank-rank slope.⁶ Because y-intercepts can differ across municipalities, two municipalities with the same rank-rank slope may have different measures of upward mobility if one municipality is somehow better at improving educational attainment for the entire distribution.

Although there is a tight linear relationship between parent and child education rank at the national level (Figure 1), it is not clear whether this still holds at the municipality level. There are over 1,600 municipalities in the Philippines, and while the average size of the age 16 cohort in each municipality is close to 1,200 individuals, many municipalities are relatively small: around 30 have fewer than 100 16-year-old residents. In Figure A3, we plot the relationships between child education rank and parent education rank for four specific

⁶Note that if no individuals shared the same education percentile rank, the national rank-rank slope would completely determine the intercept and therefore the level of upward mobility. In our setting, the existence of shared percentile ranks means that a given rank-rank slope does not imply a unique level of upward mobility, though in practice it fixes a narrow range for upward mobility.

municipalities (representing the 10th, 25th, 50th, and 75th percentiles of the population size distribution). The R^2 in these regressions (of average child rank on parental rank), particularly for municipalities with smaller populations, are lower than those in the national-level figures above, though still reasonably high (ranging from 0.78 to 0.94). These figures illustrate, however, that especially for sparsely populated lower parental ranks, a linear regression may not provide the best fit, and a single slope coefficient may not be sufficient to summarize the intergenerational education relationship for all municipalities. Reassuringly, however, both the linear (red solid line) and quadratic (dotted black line) regressions provide similar predictions for the 25th percentile, our measure of upward mobility.

Local variation in upward and relative mobility is illustrated in Figures 2 and 3 respectively. In both figures, darker areas have higher mobility (either a higher level of upward mobility or a lower intergenerational slope). There is substantial variation in both of these measures across municipalities. In Table 1, we report the mean and various percentiles of the distribution of municipality-level outcomes. The inter-quartile ranges for all measures reflect substantial within-country variation: 0.83 years of schooling, approximately 10 percentile ranks for upward mobility, and 0.12 units for the intergenerational slope. To place the variation in relative mobility into perspective, we compare these numbers to the global quintiles of a slightly different measure of relative mobility, reported in [Van der Weide et al. \(2024\)](#): the slope in a regression of child education on parental education (in years, not ranks).⁷ The 10th percentile of our intergenerational rank-rank slope distribution (0.24) would rank in the top quintile of the global relative mobility distribution in [Van der Weide et al. \(2024\)](#), whereas the 90th percentile (.48) would rank in the 4th quintile of the global distribution. Although the measures are not identical, this comparison makes it clear that the within-country variation in the Philippines is not trivial. For upward mobility, for which comparable global estimates are not yet available, we compare to the [Chetty et al. \(2014\)](#) analysis of upward income mobility in the United States. Their study reports the standard deviation of upward income mobility is 5.68, whereas we obtain a standard deviation of 7.12 (reported in Table 1).

Given that Figure A3 suggests there is variation in the linear fit of the rank-rank relationship at the municipality level, some of the variation documented in Table 1 could be due to a noisier relationship in some municipalities than others, due to small population sizes. However, even when we repeat this exercise at the province level, where population sizes are on average almost 20 times larger, we still find a non-trivial amount of variation (see Table A1).

Table 1 also reports all pairwise correlations between these measures. Mean educational attainment is highly correlated with upward mobility ($r = 0.88$), but only weakly correlated with the rank-rank slope

⁷Figure C.3 of [Van der Weide et al. \(2024\)](#) maps the global distribution of what they refer to “1-BETA,” or 1 minus the intergenerational (level) slope.

($r = -0.11$). The correlation between upward mobility and the rank-rank slope, both of which take into account child-parent links as opposed to just education levels of children, is slightly stronger, at $r = -0.26$. These associations are also apparent in a comparison of Figures 2, 3, and an analogous map displaying average educational attainment across municipalities (Figure 4). Similar geographic patterns appear in all maps, with the northern island of Luzon, particularly areas near Metro Manila, generally exhibiting better outcomes across all three measures. The correlations between each mobility measure and educational attainment suggests that some of the geographic patterns in mobility might be driven by the same municipality characteristics that contribute to determining average educational attainment.

We next explore what factors predict the within-country variation in mobility documented above. First, we calculate an asset index, generated by extracting the first component from a principal component analysis of the following household characteristics and then averaging by municipality: building type, roof material, wall material, fuel for lighting, cooking water source, and drinking water source. We also examine four characteristics related to municipality demographics: municipality population size, the share of indigenous residents, the share of Muslim residents, and the share of overseas workers.⁸ Given that the Philippines is predominantly Catholic, with distinct Muslim-majority regions, these religious differences may be associated with cultural norms and institutions that could contribute to differences in mobility patterns. The Philippines also has a large population of overseas workers, whom studies have shown to play a role in driving human capital development among sending households and regions (Khanna et al., 2022; Yang, 2008), which could also lead to differences in mobility. Finally, we examine a set of variables obtained from census data collected for each barangay (the smallest administrative unit, one level below the municipality). For each municipality, we calculate the share of barangays within 2 kilometers of a market, elementary school, high school, university, library, hospital, and health center. Because these data are only available in 2015, we restrict this analysis to the 2015 census.

In Figures 5 and 6, we first regress the relevant mobility measure on each of these municipality-level variables in separate regressions, standardizing both the outcome and predictor variables to a mean of 0 and standard deviation 1 so that the coefficients in these univariate regressions can be interpreted as correlations. We then run a multivariate regression including all municipality predictors described above, where coefficients can be interpreted as the predicted increase, in terms of standard deviations of the outcome variable, associated with a one-standard-deviation increase in the predictor variable (holding all other variables constant). Finally, we estimate another multivariate regression, including all municipality predictors as well as average child

⁸We use the ethnicity variable in the individual-level census data to identify indigenous residents. Ethnic groups are identified as indigenous using the list of Indigenous Peoples groups published by the National Commission on Indigenous Peoples (NCIP). There are some Indigenous Peoples groups that are also recognized as Muslim tribes by the National Commission on Muslim Filipinos (NCMF). In this analysis, we do not count these groups as indigenous, but results are similar when we do.

education, in order to capture the extent to which these relationships may be operating through levels of child education. For instance, municipalities with high levels of child educational attainment tend to have high upward mobility (as documented in Table 1), and assets may be predictive of better mobility because they are predictive of higher education levels.⁹

In Figure 5, assets and overseas worker shares are positively correlated with upward mobility in univariate regressions, but these relationships become much weaker when additional controls, especially educational attainment, are added. Muslim shares and indigenous shares are negatively correlated with upward mobility in univariate regressions, but the sign of this relationship flips with inclusion of child educational attainment as a control. That is, Muslim and indigenous areas have lower upward mobility on average, but compared to other areas with the same level of education, upward mobility in these areas is higher. The adjusted R^2 in the complete multivariate regression is 0.81, much of which is driven by the high explanatory power of average child educational attainment. Without this, only 60% of the variation can be explained by the remaining variables.

While there are some similarities, patterns are slightly different for the intergenerational slope in Figure 6. Here, the asset index, Muslim shares, and overseas worker shares are strongly associated with better relative mobility (lower rank-rank slopes). This holds in the univariate regressions and multivariate regressions without education-related controls. For assets and overseas worker shares, the relationship becomes much weaker (and in fact, statistically insignificant for overseas worker shares) once child education is included in the regression, suggesting that these variables are correlated with relative mobility primarily because of their relationship with average levels of child education. For Muslim shares, on the other hand, the relationship grows stronger with the inclusion of child education. Indigenous shares in both multivariate regressions are significantly associated with better relative mobility, although these coefficients are much smaller in magnitude than those on Muslim shares. Finally, although some of the proximity variables are statistically significant in their individual regressions, most are statistically insignificant or small in magnitude in the multivariate regressions. Notably, even in the regression that controls for all variables, the adjusted R^2 is only 0.32, indicating that more than two-thirds of municipality-level variation in relative mobility remains unexplained. Without average educational attainment, the adjusted R^2 is only 0.24. The lower explanatory power could be in part due to variation in the linearity of the intergenerational education relationship across municipalities: the slope may be a more accurate representation of the intergenerational relationship in some municipalities than others.

⁹Similarly, the levels of and variation in parental education may be mechanically related to these measures of mobility, and municipality characteristics may be correlated with mobility simply because of their relationship with parental education. We have also conducted the same exercise controlling for the mean and standard deviation of parental education, and our main conclusions remain the same.

Ultimately, a non-trivial amount of the variation in educational mobility is not well-explained by basic socioeconomic and demographic variables. In Figures A4 and A5, we also explore a wider set of controls, drawing from [Dulay \(2022\)](#), and find a similar portion of the variation in upward and relative mobility remains unexplained (slightly less than one-fifth and two-thirds, respectively). The expanded set of variables includes the historical presence of various religious institutions, some of which have been found to have persistent effects on contemporary municipality outcomes: for example, Catholic missions during Spanish colonial rule ([Dulay, 2022](#)). We also include weather, geographic, and land characteristics. Several variables demonstrate a strong correlation with upward and relative mobility in univariate regressions, but the majority are no longer significant predictors once assets and average education are also included as covariates. Some exceptions include distance to Jolo (a major center of Islamic governance before and during Spanish colonial rule), which is positively correlated with upward mobility in both multivariate regressions, and one of the soil quality metrics, which is positively correlated with relative mobility (negatively correlated with the rank-rank slope) in both multivariate regressions.

There are two possible interpretations of the large variation in local education mobility and the non-trivial unexplained share. One possibility is that this variation reflects the influence of important factors we have not yet accounted for, including heterogeneous intensity of or exposure to government policy across regions. Another is that it is simply noise. The distinction between these two scenarios has implications for the question of whether local mobility measures can serve as meaningful targets for policy intervention, an important question given the growing use of such measures in the economics literature. In the remainder of the paper, we address this issue by exploring how a national government policy affects the geographic distribution of local education mobility.

More specifically, we examine a national CCT program, described in more detail in the following section. A large literature documents that CCT programs, which primarily target the poor, can improve various socioeconomic outcomes for their beneficiaries ([Molina-Millán et al., 2019](#); [Parker and Todd, 2017](#)). In addition, several studies that examine heterogeneous effects have documented larger effects in poorer villages or households ([Dammert, 2009](#); [Galiani and McEwan, 2013](#); [Maluccio and Flores, 2005](#)). Based on these two points, it might seem clear that CCTs generate improvements in socioeconomic mobility at the national level, and likely also within municipalities. However, the magnitudes of these municipality-level effects are not obvious, and magnitudes are important for understanding the extent to which government policy can change the landscape of intergenerational mobility.

3 4Ps Program

3.1 Program Details

The *Pantawid Pamilyang Pilipino Program*, commonly known as the 4Ps, is a CCT program that provides cash to eligible households who meet a set of education and health requirements. It began in 2008 and has since become the primary social protection program of the Philippines. Households are eligible for the 4Ps if they have at least one child (or a pregnant woman) and if their poverty score falls below a pre-determined cutoff.

As part of the education component of the program, households receive a cash transfer for each child (up to a maximum of three children) enrolled in school and attending at least 85% of school days. Initially, this incentive applied to children aged 6 to 14, but this was expanded to include children up to age 18 in 2014. The health component of the program provides a cash transfer to households who fulfill a set of health-related conditions, including regular health visits for young children and pregnant women in the household. Based on the transfer amounts in the first year of the program, a household who met all program conditions would receive a maximum transfer of PhP 1400, estimated to be approximately one quarter of beneficiary monthly income at the time.

When it was first launched in 2008, the 4Ps program was targeted to the poorest households in the 20 poorest provinces (known as “Set 1”). According to the selection criteria, Set 1 included the poorest municipalities in the 20 poorest provinces, based on 2003 poverty rates, and “pockets of poverty in urban areas” ([Asian Development Bank, 2015](#), p.3). Almost 300,000 households were included in Set 1, including close to 4,000 households that were part of a randomized evaluation of the program, described in section 3.3. In each subsequent year, the program was expanded to new “sets” of households according to a list of targeting criteria based on region-specific poverty score cutoffs, indigenous population shares, and other vulnerability indicators. Set 2 households were added in 2009, set 3 households in 2010, and so on. We take advantage of the program’s staggered rollout, which generated variation across municipalities in the length of exposure to the program by the 2015 census, to estimate its impact on our outcomes of interest.

3.2 Program Rollout Data

Information about the 4Ps program rollout comes from the Department of Social Welfare and Development (DSWD). We use their reports of the number of 4Ps households in each municipality, as of December 2010, March 2013, and March 2016. Importantly, the counts are broken down by set number, which helps us predict when a municipality first received the 4Ps program.

To identify municipalities that received the program in 2008 and 2009, we use the earliest available report (December 2010). The presence of any set 1 household does not necessarily mean the municipality received the program in 2008, as set 1 households may have moved in after enrollment. We therefore use a threshold of 10 households for our main analysis (and test cutoffs of 1 and 100 in the appendix). If a municipality is recorded in the 2010 report as having at least 10 set 1 households, we record the municipality’s 4Ps introduction year as 2008. If not, but if it had at least 10 set 2 households, we record its 4Ps introduction year as 2009.

For introduction years 2010–2012 (sets 3–5), we apply the same method using the 2013 report. For the municipalities still unclassified after this, we use the 2016 report.¹⁰ Using this approach, 2013 (associated with set 6) is the latest year in which any of the municipalities in the report first received the 4Ps program.¹¹ Less than 1% of municipalities are not included in the DSWD reports, presumably because they did not have any 4Ps beneficiaries in 2010, 2013, or 2016, and we exclude them from our analysis because they are likely to be systematically different, both in terms of levels and trends, from the rest of the municipalities we study.

Appendix Figure A6 shows the cumulative number of municipalities exposed to the 4Ps by year, using our main cutoff (10 households) and the alternatives (1 and 100). Patterns for the 10- and 100-household cutoffs are similar (both revealing the majority of municipalities received the program between 2009 and 2012), though the 100-household cutoff results in a few municipalities classified as first receiving the program in 2014. The pattern for the 1-household cutoff is similar but shows a noticeably more rapid increase in the cumulative number of municipalities exposed to the 4Ps from 2009 to 2011.

This approach identifies when a municipality first received the program, even if additional households were added to the program in subsequent years (in later sets). For most municipalities, the vast majority of 4Ps households were introduced to the program in its initial introduction year, but for large cities, this is not always the case. As described above, Set 1 targeted small areas of high poverty in urban areas, which means that some urban municipalities are classified as having received the program in 2008 even though only a small number of households were enrolled at the time and a larger share of households were included in the program later. To identify these municipalities, we flag those where the share of Set 1 households is less than 1% of the total number of households in the municipality. This approach identifies seven municipalities, all highly urbanized cities, that we drop from the analysis.

3.3 Existing Evidence on the 4Ps

There exist several evaluations of the 4Ps program, which we categorize into three groups: short-run experimental evaluations (Kandpal et al., 2016; World Bank Group, 2012), regression discontinuity (RD)

¹⁰We use the report closest to each set’s implementation year to reduce classification errors due to migration.

¹¹The data include set 7 and 8 households, but only in municipalities that had already received the program.

evaluations (Onishi et al., 2019; Orbeta et al., 2014, 2021b), and longer-run evaluations using experimental variation (Dervisevic et al., 2021; Orbeta et al., 2021a). The first group uses a randomized controlled trial (RCT) conducted for evaluation purposes in 2008. 65 pilot villages were assigned to a treatment group that received the program 2.5 years earlier than 65 control villages. The papers that use this RCT estimate the effect of the 4Ps program by comparing eligible households in villages exposed to the program for approximately 2.5 years to eligible households in villages yet to be introduced to the program.

The second group uses RD methods, comparing households just below and just above the poverty score cutoff used to determine eligibility. Three separate RD evaluations used surveys conducted in 2011-2012 (Onishi et al., 2019), 2013 (Orbeta et al., 2014), and 2017-2018 (Orbeta et al., 2021b), capturing the effect of 2-9 years of eligibility (relative to none).

The last group of studies aims to estimate long-run effects of the program by taking advantage of the experimental variation from the original RCT. Because the control group eventually received the program, both studies focus on specific cohorts at critical ages during the RCT. Orbeta et al. (2021a) examine children born April 2009–April 2013 to capture exposure during the first 1000 days of life, relative to later exposure. Dervisevic et al. (2021) focus on those aged 12.5-14 years old, who aged out of the program before the control group received it, thus capturing the effect of exposure during adolescence (relative to none).¹²

Overall, these studies document positive effects of the 4Ps program on various health and education outcomes: for example, higher school enrollment and attendance, along with lower rates of severe stunting. Despite these generally positive effects, some evidence of negative unintended consequences has been found. For example, De Hoop et al. (2019) finds that the 4Ps program increased child labor while increasing school participation, likely because the cash transfer only partially offset education costs and children needed to work to make up the shortfall. Filmer et al. (2023) find that, by raising prices, the 4Ps program increased stunting among ineligible children, particularly in remote villages with larger shares of eligible households.

Despite many evaluations, it is still not clear what the program has done for intergenerational education mobility in the Philippines. One reason is that few studies examine educational attainment (as opposed to enrollment or attendance) as an outcome.¹³ Another reason is that most rely on household-level comparisons – e.g., between those currently eligible and ineligible – which may over- or under-estimate broader area-level effects if there are spillovers from beneficiaries to non-beneficiaries. In the next section, we describe our municipality-level approach to estimating effects on educational attainment and intergenerational education mobility.

¹²Interpretation is complicated by the fact that the control group received back payments for grants missed during the treatment period (Orbeta et al., 2021a).

¹³The exceptions are World Bank Group (2012) and Dervisevic et al. (2021), neither of which find significant effects when using the initial experimental variation.

4 Empirical Strategy

Unlike the evaluations of the 4Ps program described above, which rely on household-level comparisons, our empirical strategy relies on comparing municipalities. The 4Ps was introduced to different municipalities in different years, which generates variation across municipalities in the number of years of exposure to the program at any given time. The main motivation for this municipality-level approach is that our main outcome of interest is intergenerational education mobility measured at the municipality level. Even when we examine educational attainment (an individual-level variable) as an outcome, we use the municipality-level average. It is important to note, however, that this municipality-level approach allows us to estimate the effect of 4Ps exposure on the municipality as a whole. This parameter is different from what has been estimated in previous work and is only possible because we have information on the entire population, unlike existing work that has focused on samples collected with an RCT or RD design in mind. Given the [Filmer et al. \(2023\)](#) evidence on negative spillovers to ineligible households (for health outcomes), our estimate may be different from – and if so, is potentially more policy relevant than – the result of an individual-level approach.

Central to this approach is a multi-valued treatment variable: municipality j 's years of exposure to the 4Ps in a given year y (Exposure_{jy}). Our parameter of interest is the overall average causal response (ACR) to 4Ps exposure, as defined by [Callaway et al. \(2026\)](#). In this setting, the ACR is the difference between the potential outcome of municipality j under d years of exposure compared to under $d - 1$ years of exposure, averaged across all municipalities and levels of exposure d .

In order to recover this parameter, we use data from 2007, when $\text{Exposure}_{jy}=0$ for all municipalities, and 2015. In 2015, for our 16-year-old cohort of interest, Exposure_{jy} ranges from three years (for set 6 municipalities) to eight years (for set 1 municipalities). A key required assumption is strong parallel trends over this time period. In other words, we must assume that the evolution of average potential outcomes (among all municipalities) would have been similar to the evolution of outcomes among those that actually received d years of exposure by 2015, for all possible values of d (from 3 to 8).

We begin with a two-way fixed effects (TWFE) regression for outcome Y_{jy} , measured among 16-year-olds in municipality j observed in census year y (either 2007 or 2015):

$$Y_{jy} = \beta \text{Exposure}_{jy} + \delta_y + \mu_j + \epsilon_{jy}, \quad (1)$$

where Y_{jy} represents one of three outcome variables: average educational attainment, upward mobility, and the rank-rank slope in municipality j and census year y . Exposure_{jy} captures the number of years an

individual in municipality j was potentially exposed to the 4Ps program by census year y , based on the municipality’s program introduction year. This varies across municipalities and census years, which allows us to control for both municipality fixed effects (μ_j) and year fixed effects (δ_y), essentially a dummy for the year 2015. We cluster standard errors at the municipality level and weight by population size.

Potential exposure, captured by Exposure_{jy} , is not the same as actual exposure because not everyone in a municipality where the 4Ps program is available is eligible for or enrolled in the program. Our estimates should therefore be interpreted as intent-to-treat estimates. We are not asking what would happen if everyone in a municipality were exposed to the 4Ps for an additional year. Rather, we are asking what would happen to overall municipality outcomes if all eligible and enrolled 4Ps households (given the existing targeting rules) were exposed to an additional year of the 4Ps. In some ways, this is the more policy relevant question because the 4Ps program, like most CCT programs, was never intended to be universal and will continue to be targeted.

Because the timing of the rollout of the 4Ps program was based largely on poverty levels in 2003, one concern is that poorer municipalities would have demonstrated different outcome trajectories even in the absence of the 4Ps program. If outcome Y_{jy} was growing faster in poorer municipalities (where the 4Ps was introduced earlier) for reasons other than the 4Ps program, this would generate a spurious positive coefficient in the above regression. In order to ensure that β is not simply picking up differential trends for municipalities with different poverty rates, we also explore specifications that control for baseline poverty (measured in 2003, obtained from [National Statistical Coordination Board \(2009\)](#)) interacted with a dummy for the year 2015. This allows the change in outcomes to vary based on baseline poverty levels and ensures that these differential trends are not being picked up by β .

[Callaway et al. \(2026\)](#) show that, under the strong parallel trends assumption, β in equation (1) recovers a weighted average of dose-specific causal responses. However, it implicitly uses counter-intuitive weights with a hump-shaped distribution centered around the average dose, which does not reflect the distribution of exposure in our setting (shown in [Figure A7](#)). We therefore also use the methods proposed by [Callaway et al. \(2026\)](#) to estimate an ACR that uses more intuitive weights: the aggregate population size (of the 16-year-old cohort in the base year, 2007) across all municipalities at each exposure level.

To implement this analysis, we first calculate the change in each municipality-level outcome, from 2007 to 2015, and regress this change on dummy variables for each level of exposure (again clustering by municipality and weighting by population size):

$$\Delta Y_j = \alpha + \sum_{d=4}^8 \beta_d 1(\text{Exposure}_{j2025} = d) + \epsilon_j. \quad (2)$$

Under slightly weaker assumptions than the strong parallel trends assumption, estimates of β_d provide consistent estimates of the average treatment effect of each level of exposure, relative to three years of exposure (which we are implicitly classifying as our “untreated” group because they had the least 4Ps exposure in our sample, given that we drop the small share of municipalities without any 4Ps exposure by 2015). More specifically, β_d provides the average treatment effect of having d relative to three years of exposure under the standard parallel trends assumption (that in the absence of the program, outcomes in municipalities with d years of exposure in 2015 would have trended in parallel with those that had three years of exposure in 2015). Under the strong parallel trends assumption, estimates of $(\beta_d - \beta_{d-1})$ are consistent estimators of the ACR for each level of exposure d . We aggregate these exposure-level-specific ACRs by averaging across all $(\beta_d - \beta_{d-1})$, weighting by the total size of each group, and report the aggregated overall ACR for each outcome.

5 Results

We begin with estimates from the TWFE specification in equation (1), reported in Table 2. In the baseline specification in panel A, we report statistically significant coefficients for all three outcomes, suggesting that the 4Ps program generated relatively large improvements in educational attainment, upward mobility, and relative mobility. However, once we allow for differential trends by baseline poverty rate – by controlling for a year dummy interacted with the municipality-level poverty rate in 2003 – coefficient estimates are much smaller. The estimate for educational attainment in panel B is half the size of the estimate in panel A; the estimate for the intergenerational slope regression is no longer statistically significant in panel B. Allowing for these differential trends is important because municipalities with higher poverty in 2003 received the program earlier; if outcomes of poorer municipalities would have demonstrated larger improvements from 2007 to 2015 even in the absence of the 4Ps program, estimates in panel A could be over-estimating program effects.

As discussed in section 4, even if the strong parallel trends assumption holds, TWFE estimates reflect a weighted average of exposure-level-specific causal responses that involve potentially counter-intuitive weights. We therefore use the methods of Callaway et al. (2026) to estimate an intuitively weighted ACR. This involves first regressing the municipality-level change in each outcome (from 2007 to 2015) on dummies for each 4Ps exposure category, as in equation (2). The results of these preliminary regressions, reported in Figure 7, are instructive, as each coefficient can be interpreted as the average treatment effect of having d years of exposure as opposed to the minimum of three years of exposure, under the assumption that outcomes in each exposure level group would have evolved in parallel with outcomes in the minimum exposure group, had they also received the minimum amount of exposure (the standard parallel trends assumption). Municipalities

where 16-year-olds had only three years of exposure to the 4Ps program by 2015 (the omitted category) showed the smallest gains in educational attainment from the 2007 to the 2015 censuses; gains increase with years of exposure, such that the largest gains in educational attainment were demonstrated by municipalities with eight years of exposure (panel A). Upward mobility in panel B yields a pattern similar to that of educational attainment. In contrast, there is no clear pattern for the intergenerational slope in panel C: while all municipalities with more than three years of exposure saw larger reductions in the intergenerational slope than the omitted category, among the remaining groups, coefficients exhibit a flat trend. Taken together, these patterns suggest improvements in educational attainment and upward mobility, but not in relative mobility as measured by the intergenerational slope.

Table 3 reports estimates of the ACR, which calculates the differences between adjacent coefficient estimates in Figure 7 and averages across these differences, weighting by the size of each group in the baseline year (2007). Panel A includes no covariates in the preliminary regression, whereas panel B controls for the baseline poverty rate (allowing changes in outcomes to vary by baseline poverty, as in panel B of Table 2), but results are similar across the two specifications.¹⁴ Consistent with the graphical patterns in Figure 7, the 4Ps program generated significant improvements in educational attainment and upward mobility, but we do not detect any statistically significant effects on the intergenerational slope.

Because upward mobility is calculated using the intergenerational slope, these results – taken at face value – suggest that the program increased upward mobility not by flattening the intergenerational rank-rank slope but rather by shifting up the intercept of the rank-rank regression – in other words, increasing educational attainment for all children in a municipality. However, given that only those under the poverty line were eligible for the 4Ps, this seems unlikely: it would require positive spillovers to affect (untargeted) children across the entire distribution of parental education. A more likely explanation for the null effects on relative mobility despite improvements in upward mobility is that the program changed the slope of the regression only for a small section of the parental education distribution, yielding little change in the overall municipality slope. This highlights the limitations of the intergenerational slope as a municipality-level measure of relative mobility: it assumes the same linear relationship across the entire distribution and may not be able to pick up meaningful changes if they only affect the slope at specific portions of the parental education distribution.

Coefficient magnitudes suggest one additional year of potential 4Ps exposure at the municipality level led to a statistically significant 0.03 year increase in average municipality educational attainment. At the average level of 4Ps exposure, which was 5.2 years among 16 year-olds in 2015, this corresponds to an effect of 0.156 years, which is approximately 19% of the inter-quartile range in educational attainment as reported in Table 1. To put this magnitude into perspective, an additional 18 months of exposure to the widely studied

¹⁴Panel B is the specification that corresponds to Figure 7.

Mexican CCT program, PROGRESA, led to an increase of about 0.2 years of schooling among those eligible for the program (Behrman et al., 2011). Our estimates would imply an increase of 0.045 years of schooling in response to 18 months of 4Ps exposure, averaged over the entire population, not just those eligible for the 4Ps.¹⁵ For upward mobility, an additional year of 4Ps exposure increased the average percentile rank of a child born to parents at the 25th percentile by 0.31 positions. For the average level of exposure (5.2), this translates into an effect that is about 17% of the inter-quartile range.

In Table 4, we repeat our analysis for boys and girls separately, using the specification controlling for trends by baseline poverty status. ACR estimates are almost identical for girls (panel A) and boys (panel B).

5.1 Robustness Checks

To what extent are these results driven by differential trends in outcomes unrelated to 4Ps exposure? While our preferred specification controls for differential trends by baseline poverty status, there may be other sources of bias. To explore this, we run a falsification test using the 2000 and 2007 censuses, treating 2007 as if it were 2015 and assigning exposure accordingly. That is, we create a “false” 4Ps exposure variable that shifts back the date of each municipality’s initial 4Ps implementation by 8 years. This results in municipalities being assigned their 2015 exposure value in 2007, even though all municipalities had zero exposure in 2007. We then calculate the ACR (by first estimating equation (2) and then aggregating across all coefficients) using these census years and placebo exposure variable. This allows us to evaluate whether the outcomes of early and late recipients of the 4Ps program were already trending differently prior to the program’s inception. Reassuringly, for the two outcomes for which we documented significant effects (educational attainment and upward mobility), the estimates in panel A of Table 5 are small in magnitude and statistically insignificant. Differences in outcome trends between early and late recipients only appear after the 4Ps had begun, not before, which is in line with our attribution of the significant coefficients in Table 3 to the 4Ps program.

These placebo tests do reveal a negative and statistically significant ACR for the intergenerational slope. This suggests a potential violation of the identifying assumption, with the direction of bias unclear.¹⁶ While this particular result highlights another reason to interpret the relative mobility results with caution, the absence of significant placebo effects in columns 1 and 2 provides support for the causal interpretation of the educational attainment and upward mobility results.

As outlined in section 2.1, we focus on age 16 because it is the oldest age at which children are unlikely to

¹⁵In 2016, the 4Ps program had reached 4.2 million households. There were about 25 million households in the 2015 census.

¹⁶Municipalities with greater exposure to the program by 2015 had been demonstrating faster improvements in relative mobility (declines in the intergenerational slope) even prior to the introduction of the program. These pre-trends could be a signal that the high 4Ps exposure municipalities would have continued to demonstrate larger improvements than low exposure municipalities in the absence of the program, or that high exposure municipalities would have eventually seen a slowdown in these improvements relative to the low exposure ones.

have moved out of their house for university, reducing the possibility of a systematic relationship between parent-child linkage rates and the child’s level of educational attainment. However, we also show robustness to expanding our age range to 15-18, which adds one cohort with higher linkage rates (but less time to reach their final level of completed schooling) and two cohorts with the opposite – lower linkage rates but more time to complete their schooling. To implement this analysis, we estimate the following regression:

$$\Delta Y_{ja} = \sum_{k=15}^{18} \sum_{d=4}^8 \alpha_k + \beta_d^k 1(\text{Exposure}_{j2025} = d) \times 1(\text{Age}_a = k) + \epsilon_{ja}, \quad (3)$$

which is equivalent to estimating equation (2) separately for each age cohort a . We then average $(\beta_d^a - \beta_{d-1}^a)$ across all age cohorts a and levels of exposure d , weighting by the size of each group. The ACR estimates in panel B of Table 5 are similar to our main ones, though slightly smaller for educational attainment and slightly larger for upward mobility.

All reported results thus far use a 10-household cutoff rule to determine the 4Ps introduction year for each municipality: if there are at least 10 households reported in a particular municipality and a particular set (and not in any earlier sets), we assign the municipality the introduction year associated with that set. Figure A6 illustrates an almost identical rollout pattern when we use 100 households as the cutoff, although there are a few municipalities identified as having received the program for the first time in 2014 as opposed to 2013, which is the latest introduction year under the 10-household cutoff rule. This results in two municipalities being classified as having 2 years of exposure for 16-year-olds in 2015, which we combine with the “untreated” group (those with 3 years of exposure) due to its small sample size. Estimates that use the 100-household cutoff, reported in panel C, are essentially identical to the baseline 10-household estimates. Figure A6 reveals larger differences in the rollout pattern when we use the 1-household cutoff, which leads to a larger share of municipalities classified into earlier introduction years. Perhaps because much of this reclassification is erroneous, reflecting earlier-set households migrating to later-set municipalities, we find smaller and statistically insignificant effects on educational attainment when we use the 1-household cutoff in panel D. Upward mobility effects, however, are the same as our baseline estimates.

Our measures of intergenerational mobility can only be calculated using children who can be linked to at least one of their parents. Although we are able to link the majority (84%) of 16 year-olds, we conduct two exercises to help alleviate concerns about the selected nature of our sample. First, in Appendix section C, we show that estimates of the intergenerational correlation in our data are similar to those obtained from the International Social Survey Program (ISSP), which collects information about parental education from all respondents and therefore does not suffer from the same selection issue. Second, in Appendix section D,

we explore an alternative approach to linking generations, similar to the one used by [Alesina et al. \(2021\)](#). Instead of linking children to parents, we link younger generations to older generations, which reduces the number (and perhaps also the extent of selection) of children we have to drop from our analysis. Our results are almost identical when we use this approach, which succeeds in linking 88% of children.

Even given the reassuring similarities when it comes to the intergenerational correlations in Appendix section C and the results in Appendix section D, if the 4Ps program changes the number and types of children we are able to link, this could complicate the interpretation of our education mobility effects discussed above. To investigate this possibility further, we estimate the ACR for the percent of children successfully linked to their parents as the outcome variable. The result, reported in column 1 of Appendix Table A2, confirms that 4Ps exposure does indeed appear to affect municipality-level linkage rates. Specifically, 4Ps exposure appears to increase linkage rates, perhaps because the program reduces the need for parents to migrate away for better job opportunities or because the conditions of the program are easier to satisfy when at least one parent is living at home. Reassuringly, however, this effect is very small: an additional year of exposure increases linkage rates by 0.28 percentage points (compared to a mean of 84 percent). Perhaps unsurprisingly given the size of this effect, when we control for the percent linked in the preliminary regression used to calculate the ACR and re-calculate the ACR for the rank-rank slope and upward mobility, our results are very similar to those reported in columns 2 and 3 of Table 3. Although this is not our preferred specification because it involves the potentially problematic inclusion of an endogenous outcome as a control, the fact that the results are so similar to our main results helps alleviate concerns about compositional effects.

6 Conclusion

This paper examines the geographic distribution of intergenerational education mobility across the Philippines and investigates the role of government policy in shaping this distribution. Using population census data, we document substantial variation in both upward and relative mobility across municipalities. While socioeconomic conditions and average educational attainment explain a meaningful portion of this variation, large shares remain unexplained – approximately one-fifth for upward mobility and two-thirds for relative mobility.

We examine the extent to which government policy is able to move these municipality-level measures of mobility. If the unexplained variation is driven primarily by noise, such measures would offer little value for policy evaluation. Specifically, we evaluate the effects of the 4Ps CCT program by exploiting its staggered rollout across municipalities from 2008 to 2013. Using the methods developed by [Callaway et al. \(2026\)](#) to estimate ACRs, we find that one additional year of 4Ps exposure increased average municipality-level

educational attainment by 0.03 years and increased upward mobility by 0.31 percentile points. Scaling by average exposure length, these represent magnitudes of approximately 17-19% of the inter-quartile range. We find no statistically significant effect on relative mobility as measured by the intergenerational slope.

One important implication of these findings is that the observed geographic variation in municipality-level upward mobility is neither driven solely by selection nor noise. Social protection policy, which is commonly seen as a key policy lever for changing intergenerational mobility, is indeed capable of shifting this measure, suggesting it may be a valid target for policy evaluation. For the intergenerational slope, however, the implications are less clear. The absence of statistically significant effects may not necessarily indicate that the program had no impact on relative mobility. The intergenerational slope may not be sufficiently sensitive to changes that affect only a subset of children in a narrow window of the parental education distribution. These insights are particularly important given the growing use of local mobility measures in the economics literature.

References

- Alesina, A., Hohmann, S., Michalopoulos, S., and Papaioannou, E. (2021). Intergenerational mobility in africa. *Econometrica*, 89(1):1–35.
- Alvaredo, F., Chancel, L., Piketty, T., Saez, E., and Zucman, G. (2018). *World inequality report 2018*. Belknap Press.
- Araujo, M. C. and Macours, K. (2021). Education, income and mobility: Experimental impacts of childhood exposure to Progresa after 20 years.
- Asher, S., Novosad, P., and Rafkin, C. (2024). Intergenerational mobility in india: New measures and estimates across time and social groups. *American Economic Journal: Applied Economics*, 16(2):66–98.
- Asian Development Bank (2015). The design, expansion, and impact of pantawid. Supplementary document, Asian Development Bank. Linked document to Additional Financing of Social Protection Support Project, RRP PHI 43407-014.
- Bailey, M. J., Mohammed, S., and Mohnen, P. (2026). The evolution of us educational mobility over the 20th century and the role of public education.
- Behrman, J. R., Parker, S. W., and Todd, P. E. (2011). Do conditional cash transfers for schooling generate lasting benefits?: A five-year followup of progresa/oportunidades. *Journal of Human Resources*, 46(1):93–122.
- Biasi, B. (2023). School finance equalization increases intergenerational mobility. *Journal of Labor Economics*, 41(1):1–38.
- Black, S. E. and Devereux, P. J. (2011). Recent developments in intergenerational mobility. *Handbook of labor economics*, 4:1487–1541.
- Callaway, B., Goodman-Bacon, A., and Sant’Anna, P. H. (2026). Difference-in-differences with a continuous treatment. *American economic review*. forthcoming.
- Card, D. (1999). The causal effect of education on earnings. *Handbook of labor economics*, 3:1801–1863.
- Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The quarterly journal of economics*, 133(3):1107–1162.
- Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The quarterly journal of economics*, 129(4):1553–1623.


- Cornelissen, T. and Dang, T. (2022). The multigenerational impacts of educational expansion: Evidence from vietnam. *Labour Economics*, 78:102243.
- Dammert, A. C. (2009). Heterogeneous impacts of conditional cash transfers: Evidence from Nicaragua. *Economic Development and Cultural Change*, 58(1):53–83.
- De Hoop, J., Friedman, J., Kandpal, E., and Rosati, F. C. (2019). Child schooling and child work in the presence of a partial education subsidy. *Journal of Human Resources*, 54(2):503–531.
- Dervisevic, E., Perova, E., and Sahay, A. (2021). Long-term impacts of a short exposure to ccts in adolescence: evidence from the philippines. *World Bank Policy Research Working Paper*, 9617.
- Dulay, D. (2022). The search for spices and souls: Catholic missions as colonial state in the philippines. *Comparative Political Studies*, 55(12):2050–2085.
- Fan, Y., Yi, J., and Zhang, J. (2021). Rising intergenerational income persistence in china. *American Economic Journal: Economic Policy*, 13(1):202–230.
- Filmer, D., Friedman, J., Kandpal, E., and Onishi, J. (2023). Cash transfers, food prices, and nutrition impacts on ineligible children. *Review of Economics and Statistics*, 105(2):327–343.
- Galiani, S. and McEwan, P. J. (2013). The heterogeneous impact of conditional cash transfers. *Journal of Public Economics*, 103:85–96.
- Havnes, T. and Mogstad, M. (2015). Is universal child care leveling the playing field? *Journal of public economics*, 127:100–114.
- Kandpal, E., Alderman, H., Friedman, J., Filmer, D., Onishi, J., and Avalos, J. (2016). A conditional cash transfer program in the philippines reduces severe stunting. *The Journal of nutrition*, 146(9):1793–1800.
- Khanna, G., Murathanoglu, E., Theoharides, C. B., and Yang, D. (2022). Abundance from abroad: Migrant income and long-run economic development. Technical report, National Bureau of Economic Research.
- Lange, S. and von Werder, M. (2017). Tracking and the intergenerational transmission of education: Evidence from a natural experiment. *Economics of Education Review*, 61:59–78.
- Maluccio, J. and Flores, R. (2005). *Impact evaluation of a conditional cash transfer program: The Nicaraguan Red de Protección Social*. International Food Policy Research Institute.
- Millán, T. M., Barham, T., Macours, K., Maluccio, J. A., and Stampini, M. (2019). Long-term impacts of conditional cash transfers: Review of the evidence. *The World Bank Research Observer*, 34(1):119–159.

- Molina-Millán, T., Barham, T., Macours, K., Maluccio, J. A., and Stampini, M. (2019). Long-term impacts of conditional cash transfers: Review of the evidence. *The World Bank Research Observer*, 34(1):119–159.
- Narayan, A., Van der Weide, R., Cojocaru, A., Lakner, C., Redaelli, S., Mahler, D. G., Ramasubbaiah, R. G. N., and Thewissen, S. (2018). *Fair progress?: Economic mobility across generations around the world*. World Bank Publications.
- National Statistical Coordination Board (2009). 2003 city and municipal level poverty estimates. Technical report.
- Neidhöfer, G., Serrano, J., and Gasparini, L. (2018). Educational inequality and intergenerational mobility in latin america: A new database. *Journal of development economics*, 134:329–349.
- Onishi, J., Filmer, D. P., Friedman, J., Kandpal, E., Avalos, J. E., and Chaudhury, N. (2019). Evaluating pantawid pamilya using regression discontinuity design: Key results and lessons. Technical report, World Bank Group, Washington, D.C. Accessed: 2024-08-05.
- Orbeta, A., Abdon, A., del Mundo, M., Tutor, M., Valera, M. T., and Yarcia, D. (2014). Keeping children healthy and in school: Evaluating the pantawid pamilya using regression discontinuity design. second wave impact evaluation results. Technical report, Department of Social Welfare and Development (DSWD).
- Orbeta, A. C., Melad, K. A. M., and Araos, N. V. V. (2021a). Longer-term effects of the pantawid pamilyang pilipino program: evidence from a randomized control trial cohort analysis:(third wave impact evaluation). Technical report, PIDS Discussion Paper Series.
- Orbeta, A. C., Melad, K. A. M., and Araos, N. V. V. (2021b). Reassessing the impact of the pantawid pamilyang pilipino program: Results of the third wave impact evaluation. Technical report, PIDS Discussion Paper Series.
- Parker, S. W. and Todd, P. E. (2017). Conditional cash transfers: The case of progres/oportunidades. *Journal of Economic Literature*, 55(3):866–915.
- Pekkarinen, T., Uusitalo, R., and Kerr, S. (2009). School tracking and intergenerational income mobility: Evidence from the finnish comprehensive school reform. *Journal of public Economics*, 93(7-8):965–973.
- Van der Weide, R., Lakner, C., Mahler, D. G., Narayan, A., and Gupta, R. (2024). Intergenerational mobility around the world: A new database. *Journal of Development Economics*, 166:103167.
- Wang, D. and Zou, E. W. (2023). Public wash programs, long-run child development, and intergenerational mobility: New evidence from rural people’s republic of china. Technical report, ADBI Working Paper.

- World Bank (2022). *Taxes, Transfers, and Subsidies: Improving Progressivity and Reducing the Cost to the Poor*, chapter 5, pages 151–186. World Bank, Washington, DC.
- World Bank Group (2012). Philippines conditional cash transfer program: Impact evaluation 2012. Technical report, World Bank Group, Washington, D.C. Accessed: 2024-08-05.
- World Bank Group (2022). *Charting a Course Towards Universal Social Protection: Resilience, Equity, and Opportunity for All*. World Bank Group, Washington, DC.
- Yang, D. (2008). International migration, remittances and household investment: Evidence from philippine migrants' exchange rate shocks. *The Economic Journal*, 118(528):591–630.


Tables and Figures

Figure 1: Relationship between Child Education Rank and Parent Education Rank, by Year




Notes: Figure restricts to the age 16 cohort. Each point represents the average child education percentile rank for each parental education percentile rank, where parental education is defined as the maximum of both parents. Dot size reflects the number of children in each group. R^2 calculated from linear regression of *average* child education rank on parent education rank (depicted by solid red line). p25 represents this paper's measure of upward mobility: the predicted child rank for parents at the 25th percentile.

Figure 2: Upward Mobility Across Municipalities, by Year




Notes: Maps display average upward mobility for each municipality among the 16-year-old cohort: the average percentile rank (in the national age-specific education distribution) born to parents in the 25th percentile of the national education distribution of adults. Thresholds between categories represent the 10th, 25th, 50th, 75th, and 90th percentiles of the 2015 distribution.

Figure 3: Intergenerational Slope Across Municipalities, by Year



Notes: Maps display the rank-rank slope for each municipality among the 16-year-old cohort, obtained from a regression of child education rank on parental education rank. Thresholds between categories denote the 10th, 25th, 50th, 75th, and 90th percentiles of the 2015 distribution.

Figure 4: Educational Attainment Across Municipalities, by Year




Notes: Maps display average educational attainment in each municipality among the 16-year-old cohort. Thresholds between categories denote the 10th, 25th, 50th, 75th, and 90th percentiles of the 2015 distribution.

Table 1: Descriptive Statistics

	(1) Educational Attainment	(2) Upward Mobility	(3) Intergenerational Slope
Mean	8.43	40.82	0.36
10th percentile	7.49	32.07	0.24
25th percentile	8.12	36.10	0.30
Median	8.58	40.62	0.36
75th percentile	8.95	45.74	0.42
90th percentile	9.18	49.70	0.48
Standard Deviation	0.76	7.12	0.10
Correlation with Educational Attainment	1	0.88	-0.11
Correlation with Upward Mobility		1	-0.26


Notes: Statistics calculated using the 2015 census and age 16 cohort.

Figure 5: Correlates of Upward Mobility



Notes: Regressions use municipality-level data from 2015. “Proximity” is measured as the share of barangays in the municipality that are within 2km of the stated item. All outcome and predictor variables are standardized to a mean of 0 and standard deviation of 1.

Figure 6: Correlates of Intergenerational Slope




Notes: Regressions use municipality-level data from 2015. “Proximity” is measured as the share of barangays in the municipality that are within 2km of the stated item. All outcome and predictor variables are standardized to a mean of 0 and standard deviation of 1.

Table 2: Effects of 4Ps Exposure on Educational Attainment and Education Mobility: TWFE Estimates

	(1)	(2)	(3)
	Educational Attainment	Upward Mobility	Intergen. Slope
A. Baseline specification			
4Ps Exposure TWFE	0.06*** (0.01)	0.50*** (0.08)	-0.65*** (0.14)
B. Controlling for differential trends by baseline poverty rate			
4Ps Exposure TWFE	0.03*** (0.01)	0.47*** (0.12)	-0.08 (0.18)
Observations (Unweighted)	3,255	3,255	3,255
Observations (Weighted)	3,479,545	2,874,146	2,874,146
Mean of Y	8.28	40.59	38.85

Notes: Standard errors (clustered at the municipality level) are in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions use the 2007 and 2015 censuses, control for municipality and year fixed effects, and weight by the number of 16-year-olds in each municipality-year. Panel B also controls for a year dummy interacted with the municipality poverty rate in 2003. The intergenerational slope is multiplied by 100 for readability.

Figure 7: Changes in Outcomes by Years of 4Ps Exposure



Notes: Coefficients and 95% confidence intervals for the β_d coefficients in equation (2), which regresses the change in the outcome variable from 2007 to 2015 on dummies for each exposure level, are reported. Regressions are weighted using number of 16-year-olds in each municipality in the baseline year (2007) and control for baseline poverty rate. Standard errors are clustered at the municipality level. The intergenerational slope is multiplied by 100.

Table 3: Effects of 4Ps Exposure on Educational Attainment and Education Mobility: ACR Estimates

	(1)	(2)	(3)
	Educational Attainment	Upward Mobility	Intergen. Slope
A. Baseline specification			
4Ps Exposure ACR	0.03*** (0.01)	0.31** (0.13)	-0.25 (0.24)
B. Controlling for baseline poverty rate			
4Ps Exposure ACR	0.03** (0.01)	0.31** (0.14)	-0.10 (0.25)
Observations (Unweighted)	1,625	1,625	1,625
Observations (Weighted)	1,718,503	1,421,891	1,421,891
Mean of ΔY	0.61	1.44	-5.45

Notes: Standard errors (clustered at the municipality level) are in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The ACR is estimated from a regression of the change in the outcome variable from 2007 to 2015 on dummies for each exposure level, weighting by the number of 16-year-olds in each municipality in the base year. The ACR is a group-size weighted average of the differences between adjacent exposure-level dummy coefficients. The intergenerational slope is multiplied by 100 for readability.

Table 4: Effects of 4Ps Exposure on Educational Attainment and Mobility, by Gender: ACR Estimates

	(1)	(2)	(3)
	Educational Attainment	Upward Mobility	Intergen. Slope
A. Girls			
4Ps Exposure ACR	0.03** (0.01)	0.35** (0.16)	-0.42 (0.30)
Observations (Unweighted)	1,625	1,625	1,625
Observations (Weighted)	835,858	663,208	663,208
Mean of ΔY	0.55	0.75	-6.05
B. Boys			
4Ps Exposure ACR	0.03** (0.01)	0.32** (0.15)	0.14 (0.28)
Observations (Unweighted)	1,625	1,625	1,625
Observations (Weighted)	882,645	758,683	758,683
Mean of ΔY	0.65	1.70	-4.51

Notes: Standard errors (clustered at the municipality level) are in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The ACR is estimated from a regression of the change in the outcome variable from 2007 to 2015 on dummies for each exposure level, weighting by the number of individuals in each municipality in the base year and controlling for baseline poverty rate. The ACR is a group-size weighted average of the differences between adjacent exposure-level dummy coefficients. The intergenerational slope is multiplied by 100 for readability.

Table 5: Robustness Checks

	(1)	(2)	(3)
	Educational Attainment	Upward Mobility	Intergen. Slope
A. Placebo			
4Ps Placebo Exposure ACR	-0.00 (0.01)	0.14 (0.13)	-0.60** (0.28)
Observations (Unweighted)	1,604	1,604	1,604
Observations (Weighted)	1,443,371	1,224,806	1,224,806
Mean of ΔY	0.16	0.52	-1.90
B. Ages 15-18			
4Ps Exposure ACR	0.02** (0.01)	0.36*** (0.11)	-0.23 (0.19)
Observations (Unweighted)	6,500	6,500	6,500
Observations (Weighted)	6,781,273	5,485,111	5,485,111
Mean of ΔY	0.65	1.28	-4.53
C. Cutoff of 1			
4Ps Exposure ACR	0.03** (0.01)	0.30** (0.14)	-0.08 (0.25)
Observations (Unweighted)	1,619	1,619	1,619
Observations (Weighted)	1,718,034	1,421,556	1,421,556
Mean of ΔY	0.61	1.43	-5.46
D. Cutoff of 1			
4Ps Exposure ACR	0.01 (0.01)	0.30** (0.15)	-0.19 (0.26)
Observations (Unweighted)	1,625	1,625	1,625
Observations (Weighted)	1,718,503	1,421,891	1,421,891
Mean of ΔY	0.61	1.44	-5.45

Notes: Standard errors (clustered at the municipality level) are in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The ACR is estimated from a regression of the change in the outcome variable on dummies for each exposure level, weighting by the number of individuals in each municipality in the base year and controlling for baseline poverty rate. The ACR is a group-size weighted average of the differences between adjacent exposure-level dummy coefficients. In panel A, the change is calculated over 2000 to 2007; in remaining panels, the change is calculated over 2007 to 2015. All panels restrict to the age 16 cohort, except panel B, which expands to the 15-18 age range. Panels A and B use the standard 10-household cutoff to determine a municipality's 4Ps introduction year; panel C uses a 1-household cutoff, and panel D uses a 100-household cutoff. The intergenerational slope is multiplied by 100 for readability.

A Appendix Tables and Figures

Figure A1: Share of children with parental education information, by age





Table A1: Descriptive Statistics, Province-Level

	(1) Educational Attainment	(2) Upward Mobility	(3) Intergenerational Slope
Mean	8.47	40.24	0.37
10th percentile	7.79	33.21	0.27
25th percentile	8.08	35.79	0.33
Median	8.58	39.63	0.37
75th percentile	8.94	44.77	0.42
90th percentile	9.16	47.85	0.46
Standard Deviation	0.61	5.90	0.07
Correlation with Educational Attainment	1	0.91	-0.32
Correlation with Upward Mobility		1	-0.41

Notes: Statistics calculated using the 2015 census and age 16 cohort.

Figure A2: Relationship between Child Education Rank and Parent Education Rank, by Age




Notes: All figures use the 2015 census and restrict to the stated age cohort. Each point represents the average child education percentile rank for each parental education percentile rank, where parental education is defined as the maximum of both parents. Dot size reflects the number of children in each group. R^2 calculated from linear regression of *average* child education rank on parent education rank (depicted by solid red line). p25 represents this paper's measure of upward mobility: the predicted child rank for parents at the 25th percentile.

Table A2: Investigating the Role of Parent-Child Linkage Rates

	(1)	(2)	(3)
	Percent Linked	Upward Mobility	Intergen. Slope
4Ps Exposure ACR	0.28** (0.14)	0.29** (0.15)	-0.05 (0.25)
Observations (Unweighted)	1,625	1,625	1,625
Observations (Weighted)	1,718,503	1,421,891	1,421,891
Mean of ΔY	-0.37	1.44	-5.45
Additional Controls		Δ Percent Linked	Δ Percent Linked


Notes: Standard errors (clustered at the municipality level) are in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The ACR is estimated from a regression of the change in the outcome variable from 2007 to 2015 on dummies for each exposure level, weighting by the number of 16-year-olds in each municipality in the base year and controlling for baseline poverty rate and the controls listed in each column. The ACR is a group-size weighted average of the differences between adjacent exposure-level dummy coefficients. Coefficients are aggregated using the size of each category of exposure level. The intergenerational slope is multiplied by 100 for readability. Percent linked is the percentage of the municipality-cohort with parental education information.

Figure A3: Relationship between Child Education Rank and Parent Education Rank, by Municipality, 2015




Notes: Figure restricts to the age 16 cohort in a single municipality: panels A, B, C, and D restrict to municipalities at the 10th, 25th, 50th, and 75th percentile of the municipality-cohort size distribution, respectively. Each point represents the average child education percentile rank for each parental education percentile rank, where parental education is defined as the maximum of both parents. Dot size reflects the number of children in each group. R^2 calculated from linear regression of *average* child education rank on parent education rank (depicted by solid red line). Dotted black line represents quadratic fit. p25 represents this paper's measure of upward mobility: the predicted child rank for parents at the 25th percentile.

Figure A4: Correlates of Upward Mobility, Expanded Variable Set




Notes: Regressions use municipality-level data from 2015. “Proximity” is measured as the share of barangays in the municipality that are within 2km of the stated item. All outcome and predictor variables are standardized to a mean of 0 and standard deviation of 1. Expanded variables (from *Encomienda* to *Miles from Jolo*) are obtained from [Dulay \(2022\)](#). All soil variables are coded such that higher values represent conditions more beneficial for crop production.

Figure A5: Correlates of Intergenerational Slope, Expanded Variable Set




Notes: Regressions use municipality-level data from 2015. “Proximity” is measured as the share of barangays in the municipality that are within 2km of the stated item. All outcome and predictor variables are standardized to a mean of 0 and standard deviation of 1. Expanded variables (from *Encomienda* to *Miles from Jolo*) are obtained from [Dulay \(2022\)](#). All soil variables are coded such that higher values represent conditions more beneficial for crop production.

Figure A6: Program Rollout Across Municipalities



Notes: Cumulative number of municipalities classified into each 4Ps introduction year. “Household cutoff” refers to the number of households in a set required to classify a municipality as having received the program in the year associated with that set.

Figure A7: Distribution of 4Ps Exposure



Notes: Histogram displays distribution of municipality exposure to 4Ps program, weighted by size of 16-year-old cohort in 2007 (consistent with the weights used in equation (2)).

B Decomposition of National Rank-Rank Slope

The national rank-rank slope, which is the slope coefficient in a regression of child education rank (R^c) on parent education rank (R^p), is given by:

$$\beta_{national} = \frac{\text{Cov}(R^c, R^p)}{\text{Var}(R^p)}.$$

The numerator can be decomposed using the law of total covariance as follows:

$$\text{Cov}(R^c, R^p) = \sum_{j=1}^J w_j \text{Cov}_j(R^c, R^p) + \text{Cov}(\bar{R}_j^c, \bar{R}_j^p), \quad (4)$$

where w_j is the population share of municipality j , $\text{Cov}_j(\cdot, \cdot)$ denotes the covariance computed within municipality j , and $\bar{R}_j^c$ and $\bar{R}_j^p$ are the mean rank across all children and parents in municipality j , respectively. The first term is the weighted average of within-municipality covariances, and the second term is the covariance of the municipality means. Substituting the decomposed covariance into the original expression yields

$$\beta_{national} = \sum_{j=1}^J w_j \frac{\text{Cov}_j(R^c, R^p)}{\text{Var}(R^p)} + \frac{\text{Cov}(\bar{R}_j^c, \bar{R}_j^p)}{\text{Var}(R^p)}. \quad (5)$$

Because $\frac{\text{Cov}_j(R^c, R^p)}{\text{Var}(R^p)}$ is simply the slope in a municipality-specific regression of child rank on parent rank, or the rank-rank slope of municipality j (β_j), it is now clear that the national rank-rank slope consists of the weighted average of municipality-specific slopes and a cross-unit covariance component:

$$\beta_{national} = \sum_{j=1}^J w_j \beta_j + \frac{\text{Cov}(\bar{R}_j^c, \bar{R}_j^p)}{\text{Var}(R^p)}. \quad (6)$$

C Alternative Source for the Intergenerational Correlation

Because our measures of intergenerational mobility can only be calculated for a subset of the population (children of the household head), we compare basic intergenerational measures with those obtained from a dataset that collects information about parental education from all respondents: the 1999 International Social Survey Program (ISSP).

Using the 1999 ISSP and restricting to their youngest age category (18-24), which consists of 219 respondents in the Philippine survey, we calculate an intergenerational correlation of 0.42. Using the closest census (from 2000), restricting to respondents aged 18-24, we obtain an intergenerational correlation of 0.46. This discrepancy is not large, and we argue that it would be even smaller for our population of interest (the age 16 cohort), which has much higher parent-child linkage rates than the age 18-24 group: we are able to link 84% of 16-year-olds and only 65% of those aged 18-24.

D Alternative Linking Strategy

In this section, we describe the methods and results from an alternative strategy for calculating the intergenerational rank-rank regression. Following [Alesina et al. \(2021\)](#), instead of linking children with their parents, we link children with anyone from the older generation in their household. There are four generations that can potentially be identified in a household using the “relationship to household head” variable, which means there are three possible ways to define an older-younger pair. The first “older” generation consists of the household head’s father, mother, uncles, and aunts. This generation’s educational attainment was linked to the household head. Because of our age restrictions, the children in our sample are unlikely to be linked to their older generation using this method. More relevant to our analysis, the second “older” generation consists of the household head, household head’s spouse(s), and the household head’s siblings, who are linked to a “younger” generation of the household head’s biological children, stepchildren, nephews, and nieces. The third “older” generation consists of the household head’s sons, daughters, sons-in-law, and daughters-in-laws, who are linked to the “younger” generation consisting of the household head’s grandchildren. In this analysis, for each child, we calculate the maximum years of education across all of their older generation household members.

The correlations between this new measure and our original measure of the rank-rank slope and upward mobility are very high – 0.984 and 0.996, respectively. Given this, it is unsurprising that our main results look almost identical when we repeat our regressions using the new measures as our outcomes (Table D1). One attractive feature of these new measures is that we are able to link a larger share of children to someone from the older generation, as shown in Figure D1: on average 88% of 16-year-olds compared to 84% of 16-year-olds using the original parent-based approach. It is therefore reassuring that we come to similar conclusions with this wider coverage.

Table D1: Effects of 4Ps Exposure on Education Mobility (Using Alternative Linking Strategy)

	(1)	(2)
	Upward Mobility (Older Generation)	Intergen. Slope (Older Generation)
4Ps Exposure ACR	0.31** (0.14)	-0.03 (0.23)
Observations (Unweighted)	1,625	1,625
Observations (Weighted)	1,421,891	1,421,891
Mean of ΔY	1.49	-5.16

Notes: Standard errors (clustered at the municipality level) are in parentheses * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The ACR is estimated from a regression of the change in the outcome variable from 2007 to 2015 on dummies for each exposure level, weighting by the number of 16-year-olds in each municipality in the base year and controlling for baseline poverty rate. The ACR is a group-size weighted average of the differences between adjacent exposure-level dummy coefficients. The intergenerational slope is multiplied by 100 for readability.

Figure D1: Share of children with education information from older generation, by age

