Intergenerational Education Mobility in the Philippines: Sources of Within-Country Variation*

Teresa Molina† Cara Tan‡

February 13, 2026

Abstract

This paper uses full count census data from the Philippines to generate measures of intergenerational education mobility by municipality and birth cohort. We document substantial variation in both relative and upward mobility across municipalities and explore potential sources of this variation. Regressions that control for basic socioeconomic variables leave 20% of the variation in upward mobility and 67% of the variation in relative mobility unexplained. We examine the extent to which government policy and environmental shocks contribute to these within-country mobility differences, exploiting municipality-cohort variation in exposure to the "4Ps" conditional cash transfer program and to severe storms. We document that the 4Ps program significantly increased both relative and upward mobility. The effect of typhoon exposure was negative and statistically significant for relative mobility, and negative but insignificant for upward mobility. For relative mobility, a year of exposure to the conditional cash transfer program is equivalent to the effect of 1.65 fewer months of typhoon exposure during childhood.

JEL Codes: J62, I24, I38

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

*Many thanks to seminar participants at UP Diliman, USC CESR, UCSD, University of Tokyo, Keio-UH Workshop on Family Economics, and SEA Annual Meeting for helpful feedback on the paper. We also thank Pablo Acosta and Yasuhiro Kawasoe for sharing typhoon data, and 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' socioeoconomic outcomes are independent of those of their parents' – is a key metric reflecting equality of opportunity in a society. In addition to being an important societal goal in its own right, a high level of mobility can contribute to economic growth and social stability (Narayan et al., 2018). Early work focused primarily on generating country-level estimates of intergenerational mobility, which are now synthesized in global databases (Neidhöfer et al., 2018; Van der Weide et al., 2024). Thanks to dramatic improvements in the availability and quality of data that links outcomes across generations, more recent work has been able to calculate sub-national measures of mobility. Using administrative tax records in the United States, Chetty et al. (2014) construct county-level measures of intergenerational income mobility and document substantial variation across counties. Due to data limitations in developing countries and historical contexts, researchers in these settings often focus instead on intergenerational mobility in terms of educational attainment, a strong predictor of lifetime income (Alesina et al., 2021; Asher et al., 2024; Bailey et al., 2026).

We extend this line of work by examining educational mobility across municipalities in the Philippines. Using full count census data from 2000, 2007, 2010, and 2015, we calculate municipality-cohort-level measures of education mobility. The first measure is the slope in a regression of child education rank on parental education rank, which captures decreasing relative mobility (or increasing relative persistence). We capture upward mobility using the predicted education rank of a child born to parents in the bottom half of the education distribution: specifically, the average percentile rank (in the national age-specific education distribution) of children born to parents in the 25th percentile of the national education distribution of adults.

Like previous work focused on sub-national measures of mobility (Alesina et al., 2021; Asher et al., 2024; Chetty et al., 2014), we document substantial variation in mobility across Filipino municipalities. This variation is highly correlated with socioeconomic conditions – average educational attainment in particular. However, in regressions that control for average educational attainment, assets, and other demographic characteristics and spatial attributes, 67% of the variation in relative mobility and 20% of the variation in upward mobility remains unexplained. The fact that some municipalities are more effective at raising educational outcomes overall, and at improving outcomes for children from less-educated families in particular, appears to be driven by unobserved municipality-level characteristics.

To further explore this, we examine the extent to which municipalities' exposure to positive and negative shocks might be responsible for some of this unexplained variation in mobility. We begin with government policy, for which there is often sub-national variation in exposure and intensity. Given that 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.

To estimate the effects of the 4Ps program on 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 across municipalities and across cohorts in the number of years an individual has been exposed to the 4Ps program when we observe them in 2015. Using a regression that includes both 2015 and 2007 (as a pre-program year) and controls for municipality and age-by-census-year fixed effects, we estimate the effect of 4Ps exposure on three outcomes of interest: average educational attainment, the intergenerational slope, and upward mobility, all measured at the municipality-cohort level.

One additional year of 4Ps exposure led to a 0.03 year increase in average municipality-level educational attainment, a 0.34 percentage point reduction in the rank-rank slope (i.e., an increase in relative mobility), and a 0.47 unit increase in the average percentile rank of children born to parents at the 25th percentile (i.e., an increase in upward mobility). Scaling by the average length of 4Ps exposure in our 2015 sample (4.5 years), the effect on educational attainment translates into an effect that is 12% of the inter-quartile range, while the effects on the intergenerational slope and upward mobility translate into an effect approximately 20% the size of the inter-quartile range. Effects are similar for boys and girls, and they are largely driven by municipalities with higher 4Ps participation, which provides support for the validity of our empirical strategy.

While government policies like CCTs have the potential to improve educational mobility (to varying degrees in different municipalities), there are also negative shocks that may have opposite effects. The Philippines, as one of the most disaster-prone countries in the world, is frequently hit by severe typhoons that generate serious damage to infrastructure and often force the closure of schools. We examine how typhoon exposure during childhood affects intergenerational mobility across municipalities, taking advantage of quasi-exogenous variation in exposure across municipalities and birth cohorts. Consistent with a large literature documenting that childhood exposure to extreme weather events has detrimental effects on human capital accumulation, we find typhoons significantly reduce average educational attainment (Adhvaryu et al., 2024; Aguilar and Vicarelli, 2022; Duque et al., 2018). We also find typhoons reduce relative mobility, suggesting that children of parents with less education are hit harder. In regressions on upward mobility, the typhoon exposure coefficient is negative but not statistically significant. Comparing the effects of both shocks on relative mobility, we find that a year of exposure to the conditional cash transfer program is equivalent to the effect of 1.65 fewer months of typhoon exposure during childhood.

This paper contributes to the literature on the effects of government policies on intergenerational mobility. Early work on this topic relied on cross-sectional comparisons of intergenerational mobility across countries or states with different government spending patterns (Ichino et al., 2011; Mayer and Lopoo, 2008). More recently, fueled in part by the growing availability of high-quality administrative data in developed countries, researchers have been able to take advantage of quasi-exogenous policy variation in these countries and have documented policies like school tracking 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. We expand on this nascent literature by providing evidence on CCTs, one of the most common forms of social protection in the developing world, and by separately examining relative mobility, the focus of most of the aforementioned papers, and upward mobility, which has received less attention.[1]

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, as described in more detail in section 3.3. 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.

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

[1] All of the aforementioned papers, on both developed and developing countries, estimate effects on relative mobility by regressing a child outcome on a parent outcome, policy variable, and their interaction, while we take a different approach that uses measures of mobility as our outcome variables of interest.

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.

Though linkage rates are higher for younger children, educational attainment measures are still incomplete for them. Keeping this tradeoff in mind, we focus our analysis on children aged 15-18: we have parental education information for over 80% of this cohort, who are old enough to have completed at least nine years of schooling. In Appendix section C, we examine an alternative approach yielding higher linkage rates, linking children (a younger generation) to the previous generation (as opposed to parents specifically), as Alesina et al. (2021) do.

2.2 National and Local Mobility

Our first measure of mobility is 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.[2] This captures increasing intergenerational persistence, or decreasing relative mobility, as higher values indicate parent and child outcomes that are more tightly linked. 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 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).[3]

Figure 1 plots the relationship between child education rank and parent education rank for each of the four census years, focusing on children aged 16. Due to the discrete nature of the educational attainment variable, substantial 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

[2]This 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).

[3]As 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.

et al. (2014) does with income – a linear relationship between child and parent rank, with an R^2 close to 1, in all years.[4] From 2000 to 2010, the rank-rank slopes hovered close to 0.5. By 2015, however, the rank-rank slope had declined to 0.43, indicating a reduction in intergenerational persistence or an increase in relative mobility. Because, at the national level, the rank-rank slope also determines the intercept and therefore the predicted rank at the parent's 25th percentile, this also implies an increase in upward mobility.

The focus of this paper will be on municipality-cohort levels of 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. Because the national rank-rank slope is the weighted average of unit-level rank-rank slopes, plus a cross-unit covariance term, any improvements in local mobility (holding cross-unit covariances constant) will generate improvements in national mobility.

In each census year, we obtain the rank-rank slope for municipality j and age cohort a by regressing child education rank (in the age-specific national education distribution) onto parental education rank (in the national adult distribution), restricting to children of age a living in municipality j. We also calculate upward mobility for each municipality-cohort 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, with local measures of mobility, our upward mobility measure does provide additional information on top of the rank-rank slope. 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.

Local variation in relative and upward mobility is illustrated in Figures 2 and 3 respectively. 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, after averaging across cohorts aged 15-18 in the 2015 census. The inter-quartile ranges for all measures reflect substantial within-country variation. Table 1 also reports all pairwise correlations between these measures. Mean educational attainment is highly correlated with upward mobility ($r = 0.87$), but only weakly correlated with the rank-rank slope ($r = -0.19$). 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 higher, at $r = -0.35$.

These correlations are also apparent in a comparison of Figures 2, 3, and an analogous map displaying

[4]This 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.

average educational attainment across municipalities (Figure 4). Similar geographic patterns appear in all maps, with the northern island of Luzon and areas near large cities like Metro Manila and Cebu tending to have 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. We regress each mobility measure on a range of municipality-level variables, first one by one (panel A of Figures 5 and 6) and then all in a single regression (panel B of Figures 5 and 6). We first examine average educational attainment and an asset index, generated by extracting the principal 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 and drinking water source, and lot ownership. We also examine two characteristics related to municipality demographic composition: the share of Muslim residents and the share of overseas workers. 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, below a municipality). Because these data are only available in 2015, we restrict this analysis to the 2015 census. We calculate the share of barangays that are urban and the share of barangays within 2 kilometers of a market, elementary school, high school, university, library, hospital, and health center. In the regressions reported in Figures 5 and 6, we standardize all outcome and predictor variables so that the coefficient magnitudes 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.

In Figure 5, which examines the rank-rank slope, our measure of relative mobility, we find that educational attainment and assets are significantly negatively correlated with the rank-rank slope in panels A and B, indicating they are positively associated with relative mobility. Individually, however, they only explain between 1 to 3% of the variation in relative mobility. Muslim shares and overseas worker shares are also negatively correlated with the rank-rank slope, though the overseas worker relationship is insignificant in panel B. Only in panel B, after conditioning on all the other predictors, urban shares have a significant positive relationships with the rank-rank slope, indicating that holding all other variables constant, urban areas have lower relative mobility. Finally, although some of the proximity variables are statistically significant in their individual regressions, none are significant in panel B. Even in the regression that controls for all of these

variables, the adjusted R^2 is only 0.33, indicating that the majority of municipality-level variation in relative mobility remains unexplained.

While there are some similarities, patterns are slightly different in Figure 5. Educational attainment is the single most important predictor of upward mobility. It explains 77% of the variation in upward mobility when included alone, and all variables together explain only slightly more of the variation (adjusted R^2=0.8). All predictor variables have statistically significant coefficients in their individual regressions but become much smaller and occasionally change signs in the regression with all predictors included. For example, Muslim shares are a significant negative predictor of upward mobility in panel A but a significant positive predictor in panel B. The asset index, likely because of its high correlation with educational attainment, is a strong predictor in panel A but statistically insignificant in panel B.

Ultimately, a non-trivial amount of the variation in educational mobility is not well-explained by basic socioeconomic variables. For relative mobility, in fact, around two-thirds of the variation remains unexplained. In the remainder of the paper, we explore how positive and negative shocks at the municipality level might contribute to some of the local variation in mobility documented above.

We first 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 possibly also within municipalities. Whether this is true, however, depends on the existence and nature of spillovers from beneficiaries to non-beneficiaries, which have been documented in both Mexico and the Philippines (Angelucci and De Giorgi, 2009; Filmer et al., 2023). Moreover, even without spillovers, the magnitudes of these municipality-level effects are not clear, and magnitudes are important for understanding geographic variation in educational 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"). 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 in 2010, and so on. We take advantage of the program's staggered rollout 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 as having at least 10 set 1 households in 2010, we record the municipality's 4Ps introduction year as 2008. If not, but if it had at least 10 set 2 households in 2010, 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.[5] 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.[6]

[5]We use the report closest to each set's implementation year to reduce classification errors due to migration.
[6]The data include set 7 and 8 households, but only in municipalities that had already received the program.

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 from the rest of the municipalities we study.

Appendix Figure A2 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-cutoff results in a few municipalities classified as first receiving the program in 2014. The pattern for the 1-household cutoff is somewhat similar as well, but it shows a noticeably more rapid increase in the number of municipalities exposed to the 4Ps from 2009 to 2011.

3.3 Existing Evidence

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) 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).[7]

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.

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

Despite these generally positive effects, some evidence of negative unintended consequences have 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.[8] 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 the municipality-level approach we take to estimating effects on educational attainment and intergenerational education mobility.

4 Empirical Strategy

4.1 Estimating the effect of the 4Ps

Unlike the evaluations of the 4Ps program described above, which rely on household-level comparisons, our empirical strategy involves comparing municipalities that received the 4Ps program in different years. The main motivation for this municipality-level approach is that our main outcome of interest – intergenerational education mobility – is a municipality-level variable. Even when we examine educational attainment (an individual-level variable) as an outcome, because the lack of 4Ps status information in the census precludes an individual-level comparison of beneficiaries to non-beneficiaries, we examine municipality-level average education. 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 the past 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.

The staggered rollout of the program generates variation in the number of years a municipality-cohort has been exposed to the program as of a particular census year. In the 2015 census, for those aged 15 to

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

18 (our age cohort of interest), exposure ranges from two to eight years (see Appendix Figure A3). Relying only on cross-sectional exposure variation would be problematic, however, because exposure is likely to be correlated with both observed and unobserved municipality-level characteristics. We therefore also use the 2007 census, prior to the introduction of the 4Ps, which will allow us to control for municipality fixed effects. More specifically, for outcome Y_{jya} of cohort aged a in municipality j observed in census year y (either 2010 or 2015), we estimate the following regression:

$$Y_{jya} = \beta \text{Exposure}_{jya} + \mu_j + \gamma_{ay} + \epsilon_{jya}, \tag{1}$$

where Y_{jya} represents one of three outcome variables: average educational attainment, the rank-rank slope, and upward mobility measured among ages a ranging from 15 to 18 in municipality j and census year y. We first pool boys and girls then estimate separate regressions for each. Exposure_{jya} captures the number of years an individual of age a in municipality j was potentially exposed to the 4Ps program by census year y, based on the municipality's program introduction year and the ages eligible for the education component of the 4Ps (6-14 until 2014, when this was expanded to 6-18). This varies across municipalities, census years, and age cohorts, which allows us to control for both municipality fixed effects (μ_j) and year-by-age fixed effects (γ_{ay}). The latter allows for an age trend that varies across years. Under the assumption that any unobserved municipality-level determinants of outcomes correlated with exposure are time- and age-invariant and therefore captured by municipality fixed effects, β provides the effect of an additional year of exposure to the 4Ps. We cluster standard errors at the municipality level and weight by population size.

Potential exposure, which Exposure_{jya} captures, 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 if all eligible and enrolled 4Ps households (given the existing targeting rules) were exposed to an additional year of it. 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. That said, β in equation (1) will represent the average intent-to-treat effect across municipalities with varying shares of the population participating in the 4Ps. We therefore also estimate this regression separately for those in the top quartile and bottom 75% of municipality-level 4Ps enrollment rates to verify that our estimates are indeed larger in municipalities with a larger share of participating households.

In order to interpret β as the causal effect of an additional year of potential 4Ps exposure, we need to

assume there are no time-varying municipality-specific variables that might be correlated with exposure and the outcome variables. If outcome Y_{jya} was growing faster in municipalities where the 4Ps was introduced earlier (likely of lower socioeconomic status) than later, for reasons other than the 4Ps program, this would generate a spurious positive coefficient in the above regression. To address this concern, we also estimate specifications that allow for differential trends for municipalities with different poverty levels by controlling for a 2015 year dummy interacted with dummies for each decile of baseline poverty (measured in 2003, obtained from National Statistical Coordination Board (2009)). In addition to 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 run the following regression, restricting to years 2000 and 2007:

$$Y_{jya} = \alpha \text{PlaceboExposure}_{jya} + \mu_j + \gamma_{ay} + \epsilon_{jya}. \tag{2}$$

A statistically significant α in these regressions would indicate that outcome Y_{jya} was trending differently in areas that received the 4Ps program at different times – even prior to program implementation. This would cast doubt on the validity of our approach in equation (1).

4.2 Estimating the effect of typhoons

In order to estimate the effect of typhoon exposure on our outcomes of interest, we rely on data from Skoufias et al. (2020), which calculates typhoon exposure at the barangay-year level. Using this dataset, we calculate Typhoons_{jya}, the total number of months an individual was exposed to a typhoon from age 0 to 14 (i.e., before the youngest individuals in our sample are observed in the data). This value varies across municipalities as well as across cohorts. We include this variable in our original specification to estimate the effects of the 4Ps as well as typhoon exposure, as follows:

$$Y_{jya} = \beta_1 \text{Exposure}_{jya} + \beta_2 \text{Typhoons}_{jya} + \mu_j + \gamma_{ay} + \epsilon_{jya}. \tag{3}$$

We also conduct a similar placebo analysis using data from the 2000 and 2007 censuses. In addition to calculating placebo 4Ps exposure, we also calculate placebo typhoon exposure: the number of months an individual would have been exposed to typhoons if they remained in their same municipality from age 20

to 24.[9] As these typhoons took place after we observe their educational attainment in the data, α_2 in the following regression should not be significantly different from zero:

$$Y_{jya} = \alpha_1 \text{PlaceboExposure}_{jya} + \alpha_2 \text{PlaceboTyphoons}_{jya} + \mu_j + \gamma_{ay} + \epsilon_{jya}. \quad (4)$$

For all typhoon-related analysis, we use Conley standard errors to account for spatial correlation in error terms up to a 200km radius, which is important due to strong spatial correlations in typhoon exposure.

5 Results

5.1 Preliminary Evidence

Before reporting the results of our main specification, we begin by comparing trends in outcomes across municipalities that received the 4Ps in different years. Using our main sample (cohorts aged 15 to 18 in the 2007 and 2015 censuses), we estimate the following regression and plot the β_t coefficients:

$$Y_{jya} = \sum_{t=2008}^{2013} \beta_t 1(\text{4Ps Introduction Year} = t) \times 1(y = 2015) + \mu_j + \gamma_{ay} + \epsilon_{jya}. \quad (5)$$

Each β_t coefficient reflects the difference in outcome Y_{jya} measured among 15-18-year-olds in 2007 and 2015, for each 4Ps introduction year group after netting out municipality and age-by-year fixed effects. Because early recipient municipalities had longer exposure to the 4Ps program by 2015 (as shown in Appendix Figure A3), if the 4Ps had a causal impact on outcomes, we would expect larger changes for early recipient than late recipient municipalities.

In panel A of Figure 7, this appears to be the case. Municipalities that received the program last, in 2013 (the omitted category), show the smallest gains in educational attainment. Individuals in these areas were exposed to the program for 2–3 years by 2015, compared to no exposure in 2010. Gains are larger for the 2012 group (up to 4 years of exposure) and greater still for the 2011 and 2010 groups (up to 5 and 6 years, respectively). The 2008 and 2009 groups saw the longest exposure (up to 8 years), and consistent with this, the 2009 group showed the largest improvement in attainment. The 2008 group, however, saw more muted gains, only slightly higher than the 2010 and 2011 groups. Overall, the pattern of coefficients suggests a positive effect of the 4Ps program on educational attainment. Similarly, in panels B and C, we find that the largest decreases in the intergenerational slope (representing an improvement in mobility) took place in

[9] Given that the youngest individual in these regressions was born in 1992, and the typhoon data are available until 2016, 24 is the oldest age for which we can obtain typhoon data.

the 2009 group, with magnitudes decreasing with later introduction years. For upward mobility, there is a pattern similar to that of educational attainment.

5.2 Main Results

Table 2 reports the effects of 4Ps exposure on educational attainment and our two measures of education mobility, beginning with our baseline specification in panel A. Panel B controls for a 2015 dummy interacted with separate indicators for baseline poverty deciles, in order to allow for areas with different levels of poverty (a determinant of 4Ps implementation timing) to trend differently over time. Across all outcomes, estimates are similar though slightly smaller in panel B. We focus on this specification for the remainder of the paper, as it ensures our estimates are not being driven by different trends in high- and low-income areas.

Based on the results in panel B, 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 4.5 years among 15-18 year olds in 2015, this corresponds to an effect of 0.135 years, which is approximately 16% 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 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.[10]

The intergenerational rank-rank slope – a measure of relative mobility – was also significantly affected by the 4Ps program. An additional year of potential 4Ps exposure led to a reduction of 0.34 percentage points (an improvement in mobility): average exposure (4.5 years) corresponds to an effect size approximately 15% of the inter-quartile range. As 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 0.47 positions. For the average level of exposure (4.5), this translates into an effect that is about 27% of the inter-quartile range.

In Table 3, we repeat our analysis for different groups, using the specification controlling for trends by baseline poverty status (though Appendix Table A1 reveals very similar results when we drop these controls). We first repeat our regressions using the same municipality-level outcome variables calculated separately for each gender. Coefficient estimates are almost identical for girls (panel A) and boys (panel B).

We also explore heterogeneity by municipality characteristics. Municipalities vary in the share of the population actually eligible for and enrolled in the 4Ps, which means the coefficients estimated in the previous tables likely mask substantial heterogeneity. Because our outcomes use the entire municipality, holding all

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

else constant, effect sizes should be larger in municipalities with larger 4Ps shares. To verify this, we split the sample into the top quartile and bottom three-quarters of the 4Ps enrollment share distribution. In panels C and D of Table 3, we indeed find effects are larger in above-median enrollment municipalities.

5.3 Robustness Checks

To what extent are these results driven by differential trends in outcomes unrelated to 4Ps exposure? While our preferred specification controls flexibly for differential trends by baseline poverty status, there may be other sources of bias. To explore this, we run the placebo test described by equation (2), where we assign municipalities in 2000 and 2007 to the exposure they would have had if the program started in 2001 instead of 2008. 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, the estimates in Table 4 are small in magnitude and statistically insignificant, whether or not we control for trends by baseline poverty rates. Differences in outcome trends between early and late recipients only appear after the 4Ps had been begun, not before, which is in line with our interpretation of the significant coefficients in Table 2 as driven by the 4Ps program itself.

All previous results used an exposure variable generated using 10 households as the cutoff for categorizing a municipality as having received the 4Ps program in a particular year. In Appendix Table A2, we show results are robust to using 1 and 100 as alternative cutoffs.

Our measures of intergenerational mobility can only be calculated using linked parent-child pairs. In Appendix section B, we show that our estimates of the intergenerational correlation are similar to those obtained from a dataset that collects information about parental education from all respondents. In Appendix section C, 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. Nevertheless, even given the reassuring similarities when it comes to the intergenerational correlations in Appendix section B and the results in Appendix section C, if the 4Ps program changes the number and types of pairs we are able to link, this could complicate the interpretation of our education mobility effects discussed above.

To investigate this possibility further, we repeat our main regression using the percent of children successfully linked to their parents as the outcome variable. The result, reported in column 1 of Appendix Table A3, 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 0.42 percentage points (compared to a mean of 80.82 percent). Perhaps unsurprisingly given the size of this effect, when we control for the percent linked and repeat our regressions on the intergenerational correlation, intergenerational slope, and upward mobility, our results are very similar to those reported in columns 2 to 4 of Table 2. 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.

5.4 Typhoons

Having established that the 4Ps program generated significant improvements in education mobility, we next explore the effects of typhoon exposure. Table 5 reports the results of equation (3) in panel A and placebo specification (4) in panel B, which include variables representing the 4Ps as well as typhoons. In panel A, the coefficients on 4Ps Exposure are similar to those estimated in Table 2, indicating that previous estimates were not biased by the omission of typhoon exposure from the regressions. In addition, we find that storm exposure had significant negative effects on average educational attainment and relative mobility. An additional month of storm exposure from ages 0 to 14 yielded a reduction of 0.02 years of schooling on average, and an increase in the rank-rank slope of 0.20 units, both of which are slightly lower than the (opposite-signed) estimates for effects of a year of 4Ps exposure. Although the coefficient on typhoon exposure in the upward mobility regression is negative, consistent with the educational attainment and rank-rank slope results, it is not statistically significant.

6 Conclusion

In this paper, we calculate and demonstrate substantial variation in municipality-level measures of intergenerational education mobility across the Philippines. Given that much of this variation is not explained by basic socioeconomic, demographic, and spatial variables, we examine how positive and negative shocks – that are municipality specific – might generate within-country variation. Taking advantage of multiple waves of census data and the rollout of the 4Ps program across municipalities, we find that the 4Ps program improved average educational attainment, relative mobility, and upward mobility. In contrast, exposure to typhoons led to significant reductions in educational attainment and relative mobility. The effect of an additional year of the 4Ps program on relative mobility was equivalent to the effect of 1.65 fewer months of typhoon exposure during childhood.

References

Adhvaryu, A., Molina, T., Nyshadham, A., and Tamayo, J. (2024). Helping children catch up: Early life shocks and the progresa experiment. *The Economic Journal*, 134(657):1–22.

Aguilar, A. and Vicarelli, M. (2022). El niño and children: Medium-term effects of early-life weather shocks on cognitive and health outcomes. *World Development*, 150:105690.

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.

Angelucci, M. and De Giorgi, G. (2009). Indirect effects of an aid program: how do cash transfers affect ineligibles' consumption? *American economic review*, 99(1):486–508.

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.

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.

Card, D. (1999). The causal effect of education on earnings. *Handbook of labor economics*, 3:1801–1863.

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.

Duque, V., Rosales-Rueda, M., Sanchez, F., et al. (2018). How do early-life shocks interact with subsequent human-capital investments? evidence from administrative data. In *IZA world of labor conference*.

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.

Ichino, A., Karabarbounis, L., and Moretti, E. (2011). The political economy of intergenerational income mobility. *Economic Inquiry*, 49(1):47–69.

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.

Mayer, S. E. and Lopoo, L. M. (2008). Government spending and intergenerational mobility. *Journal of Public Economics*, 92(1-2):139–158.

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 progresa/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.

Skoufias, E., Kawasoe, Y., Strobl, E., and Acosta, P. (2020). Identifying the vulnerable to poverty from natural disasters: The case of typhoons in the philippines. *Economics of Disasters and Climate Change*, 4(1):45–82.

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.

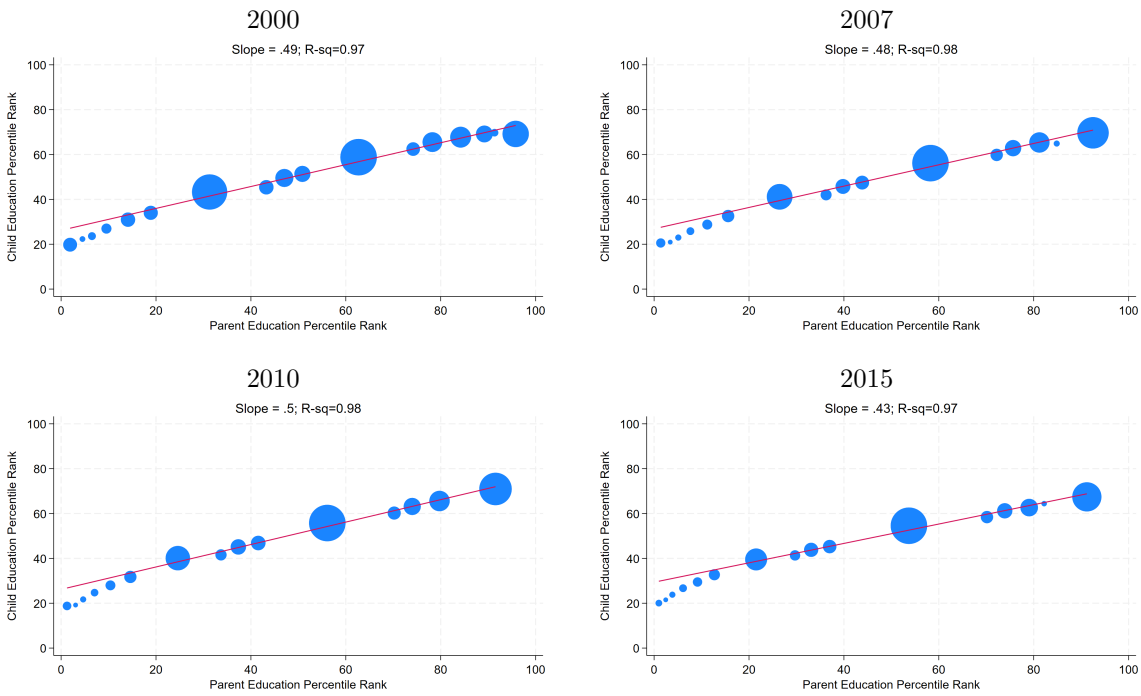
World Bank Group (2012). Philippines conditional cash transfer program: Impact evaluation 2012. Technical report, World Bank Group, Washington, D.C. Accessed: 2024-08-05.

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.

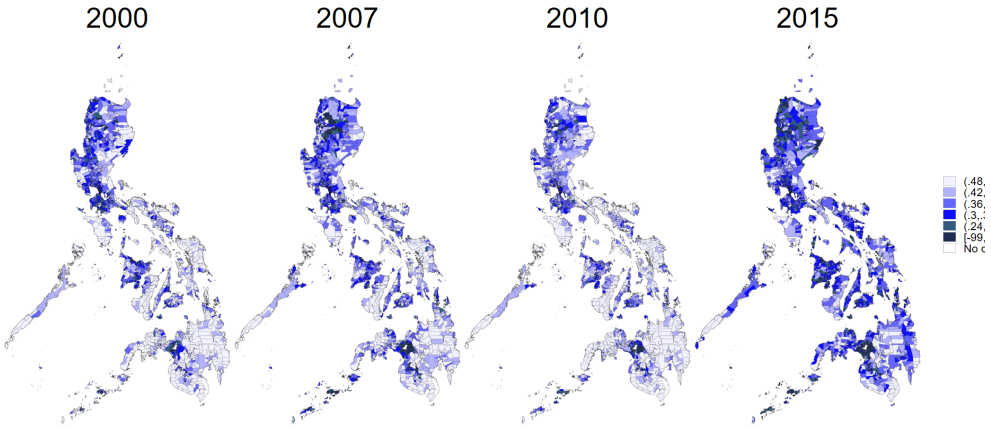
Tables and Figures

Figure 1: Rank-Rank Slope, by Year



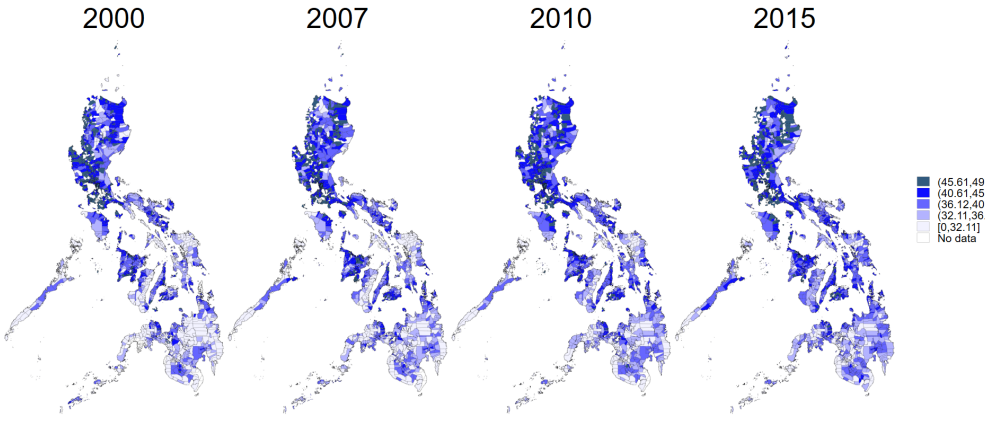
Notes: Figure restricts to the age 16 cohort.

Figure 2: Relative Mobility Across Municipalities, by Year



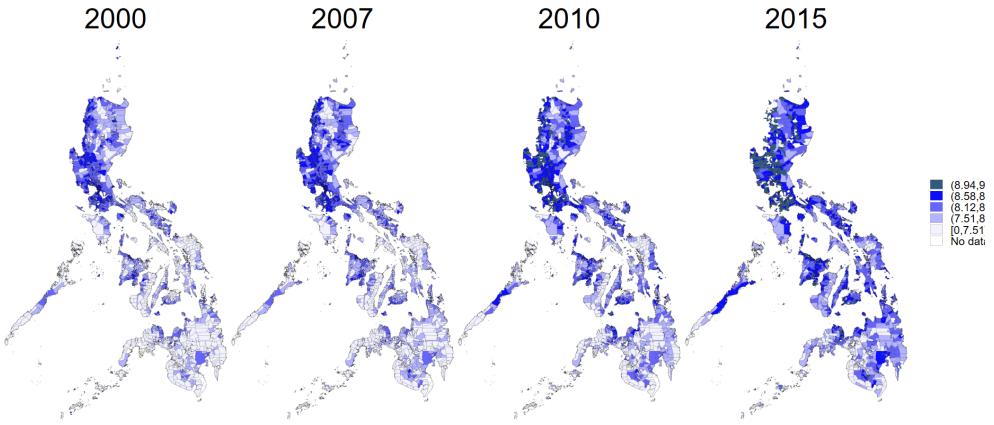
Notes: Maps display the rank-rank slope for each municipality: the slope coefficient in a regression of a child's age-specific educational rank on the maximum of both parents' educational rank, restricting to the age 16 cohort.

Figure 3: Upward Mobility Across Municipalities, by Year



Notes: Maps display upward mobility for each municipality: 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, restricting to the age 16 cohort.

Figure 4: Educational Attainment Across Municipalities, by Year



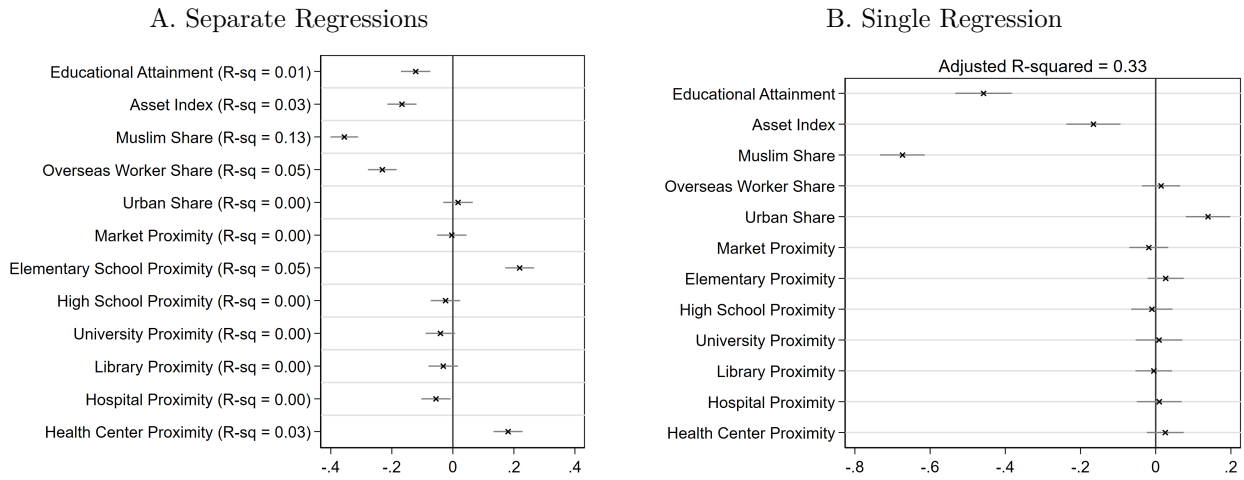
Notes: Maps display average educational attainment among 16-year-olds in each municipality.

Table 1: Descriptive Statistics

	(1) Educational Attainment	(2) Intergenerational Slope	(2) Upward Mobility
Mean	8.61	0.40	40.14
10th percentile	7.65	0.30	32.73
25th percentile	8.28	0.35	36.28
Median	8.77	0.40	40.04
75th percentile	9.14	0.45	44.19
90th percentile	9.37	0.50	47.73
Correlation with Educational Attainment	1	-0.19	0.87
Correlation with Intergenerational Slope		1	-0.35

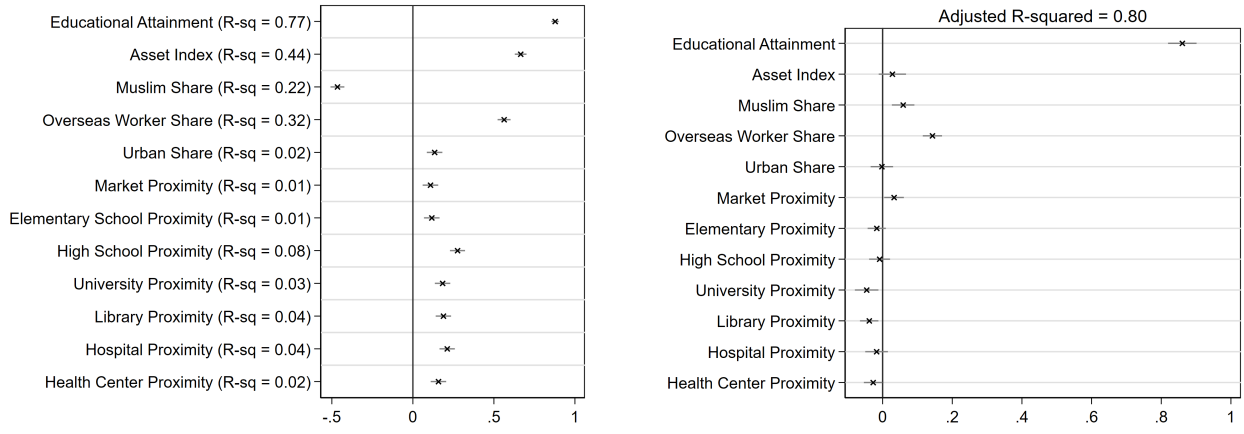
Notes: Statistics calculated after averaging cohort-municipality-level measures across all cohorts aged 15 to 18 in 2015.

Figure 5: Correlates of Relative Mobility



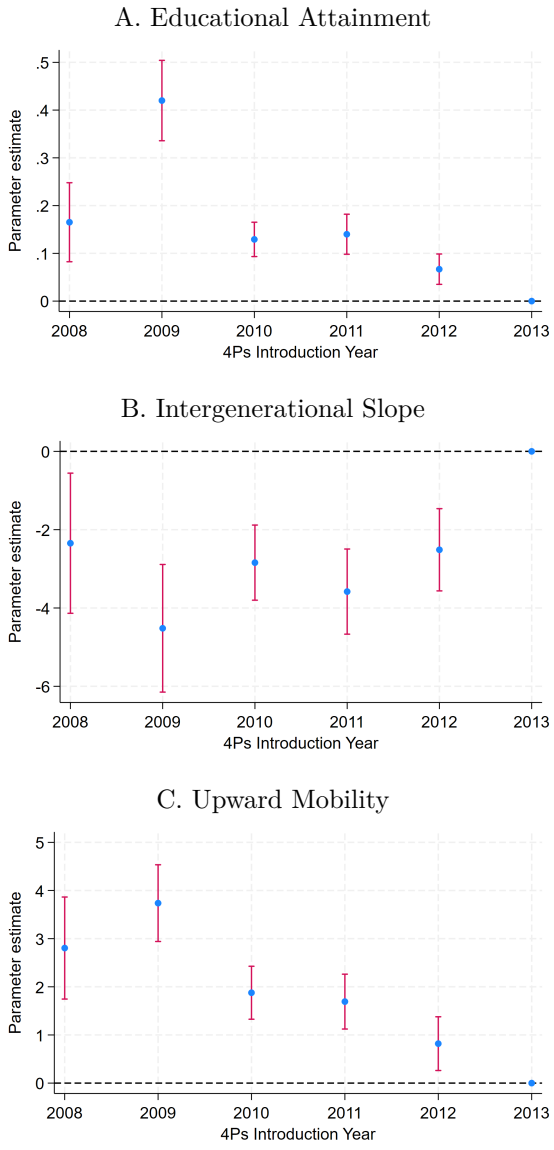
Notes: Regressions use municipality-level data from 2015. Mobility and educational attainment measures restrict to the age 16 cohort. "Proximity" is measured as the share of barangays in the municipality that are within 2km of the stated item. All outcome and predictor variables are standardized to a mean of 0 and standard deviation of 1.

Figure 6: Correlates of Upward Mobility



Notes: Regressions use municipality-level data from 2015. Mobility and educational attainment measures restrict to the age 16 cohort. "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 7: Outcome Trends by 4Ps Introduction Year



Notes: Coefficients and 95% confidence intervals for the β_t coefficients in equation (5) are reported. Standard errors are clustered at the municipality level. The intergenerational slope is multiplied by 100.

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

	(1) Educational Attainment	(2) Intergen. Slope	(3) Upward Mobility
A. Baseline specification			
4Ps Exposure	0.04***	-0.30*	0.53***
	(0.01)	(0.18)	(0.09)
B. Controlling for trends by baseline poverty			
4Ps Exposure	0.03***	-0.34**	0.47***
	(0.01)	(0.16)	(0.10)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	15,731,155	12,587,156	12,587,156
Outcome mean	8.53	41.81	40.53

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2010 and 2015 censuses and control for municipality and birth cohort fixed effects. Panel B also controls for a year dummy interacted with dummies for baseline poverty share deciles. The intergenerational slope is multiplied by 100 for readability.

Table 3: Effects of 4Ps Exposure on Educational Attainment and Mobility, by Group

	(1) Educational Attainment	(2) Intergen. Slope	(3) Upward Mobility
A. Females			
4Ps Exposure	0.03***	-0.46**	0.46***
	(0.01)	(0.19)	(0.09)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	7,691,888	5,868,235	5,868,235
Outcome mean	8.88	36.91	47.19
B. Males			
4Ps Exposure	0.03***	-0.26*	0.48***
	(0.01)	(0.15)	(0.11)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	8,039,267	6,718,921	6,718,921
Outcome mean	8.19	45.34	34.90
C. Top Quartile Eligibility Shares			
4Ps Exposure	0.10***	-0.51*	0.58***
	(0.02)	(0.27)	(0.20)
Observations (Unweighted)	3,244	3,244	3,244
Observations (Weighted)	2,145,533	1,793,319	1,793,319
Outcome mean	7.47	41.04	34.31
D. Bottom Three Quartiles Eligibility Shares			
4Ps Exposure	0.03***	-0.33*	0.46***
	(0.01)	(0.18)	(0.10)
Observations (Unweighted)	9,836	9,836	9,836
Observations (Weighted)	13,585,622	10,793,837	10,793,837
Outcome mean	8.69	41.94	41.56

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2010 and 2015 censuses and control for municipality and birth cohort fixed effects, along with a year dummy interacted with dummies for baseline poverty share deciles. The intergenerational slope is multiplied by 100 for readability.

Table 4: Placebo Tests

	(1) Educational Attainment	(2) Intergen. Slope	(3) Upward Mobility
A. Baseline specification			
4Ps Placebo Exposure	0.00	-0.27	-0.09
	(0.01)	(0.19)	(0.10)
B. Controlling for trends by baseline poverty			
4Ps Placebo Exposure	-0.00	-0.08	-0.11
	(0.01)	(0.13)	(0.09)
Observations (Unweighted)	12,860	12,860	12,860
Observations (Weighted)	13,869,829	11,264,275	11,264,275
Outcome mean	8.19	43.68	40.17

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2000 and 2007 censuses and control for municipality and birth cohort fixed effects. Panel B also controls for a year dummy interacted with dummies for baseline poverty share deciles. 4Ps Placebo exposure is the number of years of 4Ps exposure a municipality-cohort would have had if the first year of implementation was in 2000 instead of 2008. The intergenerational slope is multiplied by 100 for readability.

Table 5: Effects of 4Ps Exposure and Typhoons on Educational Attainment and Education Mobility

	(1) Educational Attainment	(2) Intergen. Slope	(3) Upward Mobility
A. Main specification			
4Ps Exposure	0.03***	-0.33**	0.46***
	(0.01)	(0.13)	(0.11)
Storm Exposure	-0.02***	0.20***	-0.07
	(0.01)	(0.05)	(0.10)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	15,731,155	12,587,156	12,587,156
Outcome mean	8.53	41.81	40.53
B. Placebo Specification			
4Ps Placebo Exposure	0.00	-0.08	-0.10
	(0.01)	(0.11)	(0.11)
Placebo Storm	-0.01	0.03	0.12
	(0.01)	(0.14)	(0.11)
Observations (Unweighted)	12,860	12,860	12,860
Observations (Weighted)	13,869,829	11,264,275	11,264,275
Outcome mean	8.19	43.68	40.17

Notes: Conley standard errors, using a cutoff of 200km, are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2000 and 2007 censuses and control for municipality and birth cohort fixed effects, as well as a year dummy interacted with dummies for baseline poverty share deciles. Storm Exposure is the number of months a municipalit-cohort was exposed to a typhoon from age 0 to 14. 4Ps Placebo exposure is the number of years of 4Ps exposure a municipality-cohort would have had if the first year of implementation was in 2000 instead of 2008. Placebo storm is the number of months of a municipality-cohort would have been exposed to typhoons from age 20 to 24. 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

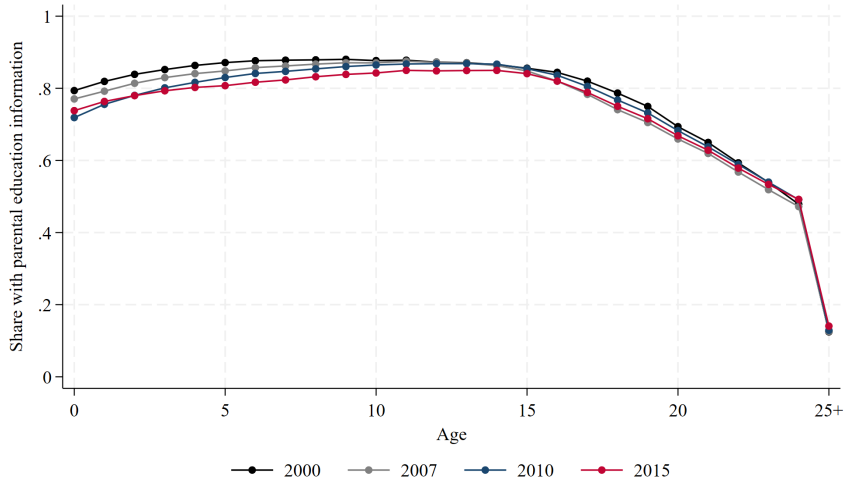
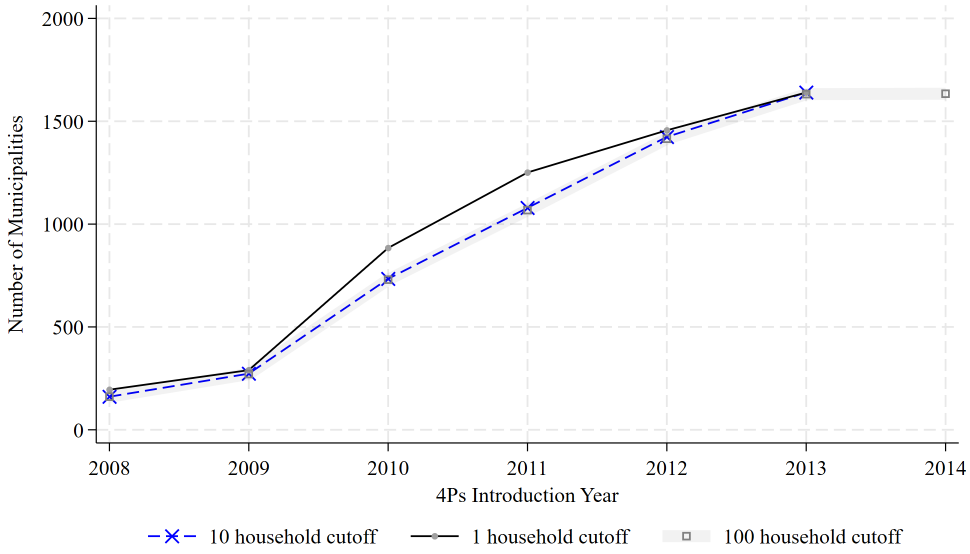
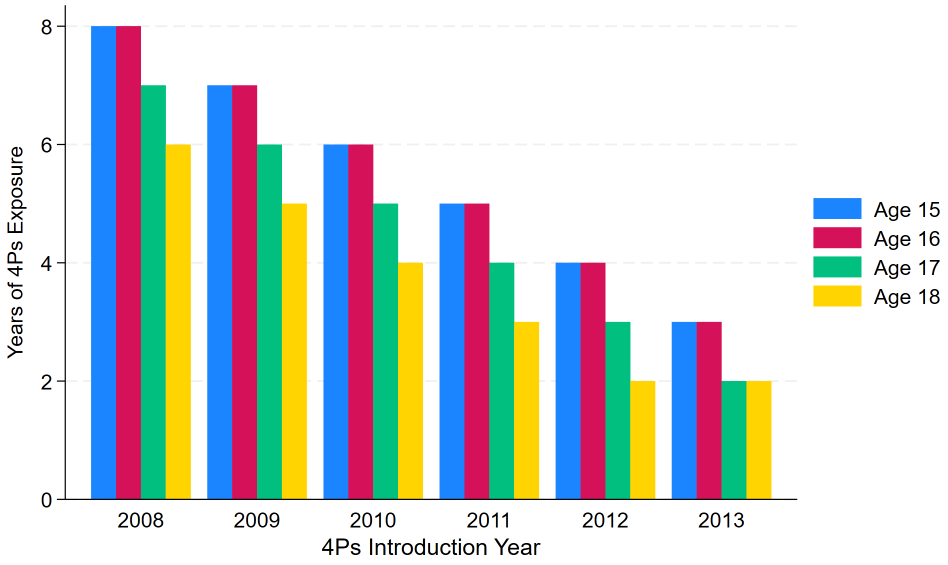


Figure A2: Program Rollout Across Municipalities



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 A3: Years of 4Ps Exposure as of 2015, by Age and Introduction Year



Notes: Exposure calculated using the 10-household cutoff rule.

Table A1: Effects of 4Ps Exposure on Educational Attainment and Education Mobility, by Group (No Trends by Baseline Poverty)

	(1) Educational Attainment	(2) Intergen. Slope	(3) Upward Mobility
A. Females			
4Ps Exposure	0.04***	-0.43**	0.54***
	(0.01)	(0.22)	(0.10)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	7,691,888	5,868,235	5,868,235
Outcome mean	8.88	36.91	47.19
B. Males			
4Ps Exposure	0.04***	-0.22	0.53***
	(0.01)	(0.16)	(0.10)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	8,039,267	6,718,921	6,718,921
Outcome mean	8.19	45.34	34.90
C. Top Quartile Eligibility Shares			
4Ps Exposure	0.09***	-0.43	0.57***
	(0.02)	(0.29)	(0.20)
Observations (Unweighted)	3,244	3,244	3,244
Observations (Weighted)	2,145,533	1,793,319	1,793,319
Outcome mean	7.47	41.04	34.31
D. Bottom Three Quartiles Eligibility Shares			
4Ps Exposure	0.02***	-0.23	0.48***
	(0.01)	(0.20)	(0.11)
Observations (Unweighted)	9,836	9,836	9,836
Observations (Weighted)	13,585,622	10,793,837	10,793,837
Outcome mean	8.69	41.94	41.56

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2010 and 2015 censuses and control for municipality and birth cohort fixed effects. The intergenerational slope is multiplied by 100 for readability.

Table A2: Robustness to Other Cutoff Definitions

	(1) Educational Attainment	(2) Intergen. Slope	(3) Upward Mobility
A. Cutoff of 1			
4Ps Exposure	0.03***	-0.35**	0.43***
	(0.01)	(0.17)	(0.10)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	15,731,155	12,587,156	12,587,156
Outcome mean	8.53	41.81	40.53
B. Cutoff of 100			
4Ps Exposure	0.03***	-0.35**	0.47***
	(0.01)	(0.16)	(0.10)
Observations (Unweighted)	13,032	13,032	13,032
Observations (Weighted)	15,726,508	12,583,991	12,583,991
Outcome mean	8.53	41.81	40.52

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2010 and 2015 censuses and control for municipality and birth cohort fixed effects, along with a year dummy interacted with dummies for baseline poverty share deciles. The intergenerational slope is multiplied by 100 for readability. "Cutoff" refers to the number of households used to define whether a municipality had received the 4Ps by a particular year.

Table A3: Investigating the Role of Parent-Child Linkage Rates

	(1) Percent Linked	(2) Intergen. Slope	(3) Upward Mobility
4Ps Exposure	0.45***	-0.29*	0.49***
	(0.12)	(0.17)	(0.09)
Percent linked		-0.12***	-0.06***
		(0.02)	(0.02)
Observations (Unweighted)	13,080	13,080	13,080
Observations (Weighted)	15,731,155	12,587,156	12,587,156
Outcome mean	80.01	41.81	40.53

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2010 and 2015 censuses and control for municipality and birth cohort fixed effects, along with a year dummy interacted with dummies for baseline poverty share deciles. The intergenerational slope is multiplied by 100 for readability. Percent linked is the percentage of the municipality-cohort with parental education information.

B 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 and/or spouse), we attempt to compare our 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.

C Alternative Linking Strategy

In this section, we describe the methods and results from an alternative strategy for calculating intergenerational correlations and upward mobility. 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 C1). 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 C1: 88% of 15 to 18-year-olds in 2015 (compared to the 80% of 15 to 18-year-olds we were able to link to their parents in 2015). It is therefore reassuring that we come to similar conclusions with this wider coverage.

Table C1: Effects of 4Ps Exposure on Education Mobility (Using Alternative Linking Strategy)

	(1) Intergen. Slope (Older Generation)	(2) Upward Mobility (Older Generation)	(3) Intergen. Slope (Older Generation)	(4) Upward Mobility (Older Generation)
4Ps Exposure	-0.29*	0.51***	-0.33**	0.45***
	(0.17)	(0.09)	(0.15)	(0.09)
Observations (Unweighted)	13,080	13,080	13,080	13,080
Observations (Weighted)	13,607,782	13,607,782	13,607,782	13,607,782
Outcome mean	40.28	40.58	40.28	40.58
Additional Controls	None	None	Year x Baseline Poverty	Year x Baseline Poverty

Notes: Standard errors (clustered at the municipality level) are in parentheses * p< 0.1, ** p< 0.05, *** p< 0.01. All regressions use the 2010 and 2015 censuses and control for municipality and birth cohort fixed effects. "Year x Baseline Poverty" indicates a 2015 year dummy interacted with dummies for baseline poverty share deciles. The intergenerational slope is multiplied by 100 for readability.

Figure C1: Share of children with education information from older generation, by age

