The World Bank Economic Review, 00(0), 2026, 1–21
https://doi.org10.1093/wber/lhaf044
Article

OXFORD

Conditional Cash Transfers, Schooling Decisions, and Labor-Market Conditions

Teresa Molina and Joaquim Vidiella-Martin

Abstract

Does the effectiveness of an education policy depend on the job opportunities in the local labor market? This paper provides a theoretical and empirical investigation of how schooling decisions respond to conditional cash transfer programs, across areas with different exposure to export manufacturing. Results show that Mexico's PROGRESA program, documented to have increased educational attainment, was less effective in areas with more export-oriented manufacturing jobs. A theoretical model, combined with empirical evidence, suggests this is because these jobs generate more convex opportunity costs of schooling. Consistent with this, the heterogeneity documented is strongest among those old enough to be working in factory jobs. In addition, this heterogeneity is primarily driven by jobs that directly influence schooling opportunity costs: low-wage jobs and jobs for school-aged workers.

JEL classification: I38, F16, O14, I28

Keywords: conditional cash transfers, export manufacturing, Mexico, opportunity costs

1. Introduction

Reducing the cost of schooling is a common policy lever used to promote human capital investment. The effects of such policies, like private school vouchers and financial aid, vary widely across settings (Herbaut and Geven 2020; Epple, Romano, and Urquiola 2017). Conditional cash transfer (CCT) programs, which lower the cost of schooling by conditioning transfers on school attendance, are no exception (Millán et al. 2019; Glewwe and Muralidharan 2016): Fiszbein and Schady (2009), for example, document impacts on attendance rates ranging from −3 to 31 percentage points. Though some variation might be due to differences in program characteristics and their implementation, these are unlikely to be the only explanations.

To better understand potential sources of heterogeneity in education policy effectiveness, we begin with a model of the optimal schooling decision, which suggests that labor-market conditions (specifically, the types of job opportunities available) could be important in determining the success of these education policies. Our parameter of interest is the response of optimal schooling to a price reduction, as this captures

Teresa Molina (corresponding author) is an associate professor at the University of Hawaii at Manoa, USA, and a research fellow at the IZA Institute of Labor Economics; her email address is tmolina@hawaii.edu. Joaquim Vidiella-Martin is a senior research fellow at the University of Oxford, UK, and affiliated with Erasmus University Rotterdam, Netherlands; his email address is joaquim.vidiellamartin@phc.ox.ac.uk. The authors thank Achyuta Adhvaryu, Carolina Arteaga, David Atkin, Chris Karbownik, Anant Nyshadham, Owen O'Donnell, Tom Vogl, and seminar participants at the World Bank, UC Irvine, UGA, NEUDC, BREAD, UVA, NYU, Emory, JEES, UWA, UB, EALE, and UH Manoa for helpful comments. A supplementary online appendix is available with this article at *The World Bank Economic Review* website.

how programs like CCTs—which effectively reduce the price of schooling—affect eventual educational attainment. We show that the magnitude of this response depends on the convexities of the opportunity cost and wage functions, both of which are determined by the types of jobs that are available in a community. Adult jobs can affect perceptions about future returns to schooling, while jobs for school-aged youths can affect the forgone wage component of the opportunity cost function. Across countries, job opportunities have been substantially affected by globalization and trade liberalization (Nallari et al. 2012; Autor, Dorn, and Hanson 2016), raising the question of what these changes might mean for education policy.

Our empirical investigation focuses on the Programa de Educación, Salud y Alimentación (PROGRESA), a Mexican CCT program that is documented to have increased average educational attainment. PROGRESA began during the later part of a period of rapid trade liberalization in Mexico, which brought about a large increase in export-oriented manufacturing jobs.[1] Atkin (2016) shows that export manufacturing reduced schooling levels in this setting, which implies that export manufacturing increased marginal opportunity costs by more than it increased the returns to schooling. However, because this finding does not inform us about the convexities of the wage and cost functions, it remains unclear whether the expansion of export manufacturing would reduce or enhance the schooling impact of PROGRESA.

To answer this question, we take advantage of the initial PROGRESA evaluation, which randomized villages to receive PROGRESA in 1998 (treatment) or 2000 (control). We combine this with variation in the intensity of export-oriented manufacturing across subdelegations, geographic units that are substantially larger than villages but smaller than states, used for administrative reporting by the Mexican Social Security Institute (IMSS).[2] Given differences in the specific types of export manufacturing jobs taken on by men and women, we calculate gender-specific counts of export jobs for each subdelegation and estimate heterogeneity in the PROGRESA treatment effect along this dimension. Specifically, we regress attendance rates and eventual educational attainment on a PROGRESA treatment dummy and the interaction between PROGRESA treatment and subdelegation-level gender-specific export jobs (controlling for subdelegation fixed effects and a rich set of demographic controls). We estimate a negative and significant interaction coefficient, which indicates that the impact of PROGRESA was smaller for those with greater exposure to export jobs.

To understand why export manufacturing jobs are associated with a smaller PROGRESA impact, we explore data on wages and opportunity costs in areas with higher vs. lower concentrations of export jobs. This descriptive analysis suggests that opportunity costs are more convex in areas with more export manufacturing jobs, perhaps because the wages at these export jobs increase faster with schooling than in other jobs (e.g., agriculture). In conjunction with the model, this provides one reason why PROGRESA was less effective in these areas. Of course, this is not the only way in which export manufacturing jobs differ from others, and we are unable to provide definitive evidence that this is the most important explanation.

That said, investigating what types of individuals and what types of jobs are driving the heterogeneity we document provides support for the importance of opportunity costs. Heterogeneity is stronger for those who are old enough to be working in factory jobs (at least 15 years old), and it is driven primarily by the types of export jobs likely to factor into the opportunity cost of schooling: low-wage jobs and those held by young workers. These findings suggest export-oriented manufacturing jobs reduce the PROGRESA impact because they translate into more rapidly increasing forgone wages for children who are (or whose parents are) deciding on the optimal level of schooling.

1 Throughout the paper we use “export-oriented manufacturing jobs” and “export jobs” interchangeably.

2 As detailed in the data section, we use subdelegations as our geographic unit of analysis because they strike the necessary balance between granularity and sample size: aggregating to larger units (delegations or states) would leave too few observations for inference, while aggregating to smaller units, such as municipalities, often eliminates within-unit variation in PROGRESA treatment status.

Although we are unable to directly link opportunity cost convexities to the heterogeneity we document, we are able to explore and rule out alternative explanations. We show that our interaction coefficient is not simply picking up gender differences in the PROGRESA treatment effect. We also show that the heterogeneity is not driven by correlations between export jobs and other characteristics, like subdelegation-level educational attainment, urban shares, or average income; a child's baseline educational attainment; or household-level migration, income, or occupation types. Our results are also robust to the use of an alternate export manufacturing variable: predicted export-job growth generated using a shift-share strategy.

These findings speak to a broader empirical literature showing how schooling levels are influenced by opportunity costs and (perceived) returns to schooling (Jensen 2010, 2012; Shah and Steinberg 2019; Cascio and Narayan 2022), especially those focusing on trade-related changes to the labor market (Atkin 2016; Blanchard and Olney 2017; Greenland and Lopresti 2016; Edmonds, Topalova, and Pavcnik 2009). Unlike these studies, our focus is not on schooling *levels*, but schooling *responses* to a price reduction, a policy-relevant parameter of interest that captures the effectiveness of a policy at increasing educational attainment. Importantly, knowing how schooling levels are affected by a particular shock or labor-market characteristic is not enough to predict how that characteristic will affect the schooling response to a price reduction (i.e., the schooling impact of a program like PROGRESA).

This study also expands our understanding of the interactions between different development policies. Economic development is a multifaceted phenomenon, which often requires the simultaneous pursuit of a variety of different goals. Increasing educational attainment is one goal often prioritized by governments and international organizations (United Nations 2016). The creation of a strong manufacturing sector, and in particular one that is export oriented, is another goal that has featured prominently in the development path of many nations (Lederman, Olarreaga, and Payton 2010; Lustig 2001; Page 2016). Both targets play an important role in government policy, but little is known about how the pursuit of one goal affects progress towards the other.

2. Theoretical Framework

We begin by outlining a simple theoretical framework that sheds light on how labor-market conditions can influence the schooling impact of policies that reduce the price of education. Suppose parents maximize discounted future wages minus the cost of schooling:

$$\beta W(S) - c(S) - pS, \tag{1}$$

where wages are a function of schooling ($W(S)$), opportunity costs are forgone wages $c(S)$, and the price of one year of school is p. The optimal level of schooling is determined by the expression

$$\beta \frac{\partial W}{\partial S} = \frac{\partial c}{\partial S} + p.$$

Labor-market conditions—specifically, the types of jobs that are available to an individual—affect this expression in two ways. First, jobs can affect perceptions about the future returns to schooling ($\frac{\partial W}{\partial S}$).[3] In addition, certain jobs, which are available to school-aged youth, can affect the marginal opportunity cost of schooling ($\frac{\partial c}{\partial S}$).

3 The wage function depends on the jobs and income that will be available when these youths eventually enter the labor market, which could be informed by the conditions in the labor market at the time of the decision. For example, 70 percent of survey respondents in the Dominican Republic report that people in their community were their primary source of information about expected income (Jensen 2010). In Madagascar, Nguyen (2008) finds that expectations about future returns to schooling are influenced by information about current labor-market conditions.

In this paper, more than the optimal level of schooling, we are interested in the response of optimal schooling to a decrease in p, which is given by

$$-\frac{dS}{dp} = \left(\frac{\partial^2 c}{\partial S^2} - \beta\frac{\partial^2 W}{\partial S^2}\right)^{-1}. \tag{2}$$

Assuming that the term inside the brackets is positive (i.e., that the second-order condition for a maximum holds), this predicts what has been documented empirically—reducing the price of schooling typically increases educational attainment.

More importantly, however, this expression shows that the magnitude of the impact of a price reduction depends on the second derivatives of the opportunity cost and wage functions. In particular, in labor markets with more convex opportunity costs (larger $\frac{\partial^2 c}{\partial S^2}$), the magnitude of the response will be smaller. In labor markets with smaller $\frac{\partial^2 W}{\partial S^2}$ (that is, marginal benefits that are either increasing slower or decreasing faster), the schooling response will also be smaller.

This expression can also be interpreted in terms of the gap between benefits and forgone wages $(W(S) - c(S))$, or net benefits. In areas where net benefits decrease faster with schooling (i.e., where the marginal net benefits are more negative), the schooling response to a price reduction will be smaller.

Notably, these predictions hinge on the second derivatives rather than the first derivatives (the marginal opportunity cost, $\frac{\partial c}{\partial S}$, or the return to schooling, $\frac{\partial W}{\partial S}$), though these could also matter—in ambiguous ways—due to their role in determining the optimal level of schooling, which in turn influences the magnitude of the expression in (2). This means that knowing how a particular industry or occupation composition affects the optimal level of schooling (determined by equation (1)) does not allow us to predict how it will affect the schooling response to a price reduction (equation (2)), as we discuss in the supplementary online appendix. The importance of the second derivatives makes it difficult to predict which types of labor markets will enhance or reduce the schooling impact of these types of education policies.

We now turn to an empirical analysis of this question, using export manufacturing jobs as our source of labor-market variation. Because the rapid and substantial expansion of export-oriented manufacturing has been a primary source of economic change in much of the developing world, it is important to understand how it may have impacted the effectiveness of education policies. We focus on Mexico's CCT program, PROGRESA, which we describe in the next section.

3. Background

3.1. PROGRESA

CCTs are now widely used across the globe (World Bank Group 2017), but one of the first CCT programs, PROGRESA, began in Mexico in 1997. The program provided cash transfers to poor families that satisfied certain education and health-related requirements.

The education component of PROGRESA, which is the focus of this paper, consisted of cash payments made to mothers whose children had school attendance rates of at least 85 percent. When the program first started, it covered children in third to ninth grade, but this was expanded to include high-school students starting in 2001. Grant amounts increased with grade level, with higher amounts for girls than boys, and ranged from 105 to 660 pesos per month in 2003.[4] Since its inception, PROGRESA has been expanded and renamed several times. It changed its name to Oportunidades in 2002 and was further restructured and renamed Prospera in 2015 (Ordóñez-Barba and Silva-Hernández 2019). In 2019, Prospera was discontinued and replaced by the Benito Juárez scholarship program for education, providing grants

4 See Skoufias and Parker (2001), Skoufias (2005), Behrman, Parker, and Todd (2009a), and Behrman, Parker, and Todd (2011) for more program details.

to enrolled students and eliminating the health and nutrition components of the program (Diario Oficial de la Federación 2019).

PROGRESA was implemented experimentally in 506 rural villages in seven states: Guerrero, Hidalgo, Michoacán, Puebla, Queretaro, San Luis Potosí, and Veracruz. Villages were randomized into either treatment or control: the treatment group (320 villages) started receiving benefits in the spring of 1998, and the control group (186 villages) did not receive benefits until the end of 1999.

The randomized variation has allowed for rigorous evaluations of the program's effects on a wide range of outcomes, summarized in Parker et al. (2017). The most relevant findings for our study are those related to educational outcomes. Short-run evaluations of the program compare treatment and control villages in 1998 and 1999 (when PROGRESA had not yet been rolled out to the control group), and find PROGRESA increased school attendance, enrollment, and grade progression, and reduced dropout (Skoufias and Parker 2001; Behrman, Sengupta, and Todd 2005; Schultz 2004). Medium-run evaluations compare educational attainment in treatment and control villages in a 2003 follow-up survey, and show higher educational attainment and grade progression in treatment villages (Behrman, Parker, and Todd 2011, 2009b). Because the control group was already exposed to PROGRESA by this time, these estimates capture the effect of being exposed to PROGRESA 18 months earlier. More recently, researchers have sought to explore the program's long-run effects on education and labor-market outcomes using data collected 10 to 20 years after the program first started (Araujo and Macours 2021; Parker and Vogl 2023).

Previous work has examined heterogeneity in the effect of PROGRESA (and other CCTs) across a number of other dimensions—child gender (Manley, Gitter, and Slavchevska 2013; Lee and Shaikh 2014), early-life circumstances (Adhvaryu et al. 2018), household and village poverty levels (Maluccio and Flores 2005; Dammert 2009; Galiani and McEwan 2013), and other household characteristics (Djebbari and Smith 2008; Handa et al. 2010; Angelucci et al. 2010). The focus of our study, however, is on heterogeneity driven by labor-market conditions.

3.2. Export Manufacturing

The beginning of the PROGRESA program coincided with the tail end of a period of rapid trade liberalization in Mexico. After pursuing an import substitution strategy for decades, Mexico sharply reversed course by joining the General Agreement on Trade and Tariffs in 1986, followed by the North American Free Trade Agreement (NAFTA) in 1994. The manufacturing sector in Mexico was considered to be the key driver of economic growth and industrial development since the 1980s (Cámara de Diputados 2004), and these free trade agreements were part of a deliberate strategy to improve Mexico's economy using the manufacturing industry (Moreno-Brid 2007).

As a result of this shift in policy, Mexico saw a large increase in manufacturing jobs at factories producing goods for export. From 1986 to 2000, the number of formal sector jobs in export manufacturing sectors more than tripled, from less than 900,000 to over 2.7 million (Atkin 2016). Notably, employment growth was concentrated primarily in the manufacturing industry: agricultural employment declined substantially in the decade following NAFTA, which meanwhile had little effect on employment in the services sector (Polaski 2003). Several studies have documented how these changes affected employment, wages, schooling levels, and inequality across genders and skill levels (Revenga 1997; Hanson and Harrison 1999; Juhn, Ujhelyi, and Villegas-Sanchez 2014; Aguayo-Tellez et al. 2013; Atkin 2016). We expand on this work by investigating how the resulting changes in the labor market influenced the education effects of the PROGRESA program.

The expansion of export manufacturing certainly affected opportunity costs for school-aged youths. Using the IMSS data, we estimate that the monthly wage of a factory worker under the age of 20, in our

PROGRESA subdelegations of interest, was approximately 2,200 pesos per month in 2003, about three times as large as the monthly PROGRESA education transfer for the oldest beneficiaries.[5]

Atkin (2016) shows the expansion of export-oriented jobs increased the marginal cost of schooling more than the marginal benefit, subsequently reducing average schooling levels. In the context of the model in our theoretical framework, this finding informs us about the expression in equation (1), which determines the optimal level of schooling, but does not allow us to predict whether export jobs will increase or decrease the schooling impact of CCTs, captured by equation (2). The simulations described in supplementary online appendix S2 help illustrate this point. We provide one example of a cost function that reduces optimal schooling (as export jobs have done) while increasing the schooling response to a CCT, and another cost function that reduces optimal schooling while reducing the schooling response.

An important advantage of focusing on export-oriented manufacturing jobs is that they tend to be driven in large part by external demand, not just by local demand and supply. Perhaps because of this, shares of export manufacturing jobs in our setting are not strongly correlated with other socioeconomic characteristics. For example, the correlation between subdelegation-level export manufacturing jobs and average income in our sample is 0.09, while the correlation between overall manufacturing shares and average income is 0.45. For education, these correlations are 0.06 for export manufacturing and 0.47 for overall manufacturing. In addition, in Atkin (2016), which uses an instrumental variables strategy to account for the potential endogeneity of export manufacturing job growth, naive OLS and IV estimates do not differ substantially. Although he uses year-to-year changes in export jobs (while our focus is on the total stock of jobs), this provides some evidence that many determinants of export manufacturing jobs are arguably exogenous.

4. Data

We merge the data collected for the evaluation of the PROGRESA program with employment data from the IMSS, both of which we describe below. In addition to these two sources, we use Mexican census data collected by the National Institute of Statistics, Geography, and Informatics (INEGI) and provided by IPUMS (Minnesota Population Center 2015).

4.1. PROGRESA Data

The data collected for the evaluation of the PROGRESA program include a baseline survey of all households in PROGRESA villages in October 1997 and three years of follow-up surveys every six months, from 1998 to 2000. Another follow-up survey was carried out in 2003 in all 506 villages that were part of the original evaluation sample. These surveys collected detailed information on household composition and demographics, education, health, employment status, and income. In our analysis, we use the 1997 baseline survey, three surveys that took place in 1998–1999 before the control group received PROGRESA, and the 2003 follow-up.

We define a treatment dummy that is equal to 1 for households in one of the 320 villages placed in the treatment group, and to 0 for the 186 villages in the control group. In our main analysis, we use two education outcome variables: years of educational attainment, measured using the 2003 wave, and an indicator for school attendance, measured in each of the 1998–1999 waves. Age, gender, household size, household head age, household head gender, parental education, and parental language variables

5 To estimate the potential wages of PROGRESA beneficiaries, we take the monthly IMSS data from 2003 and restrict our analysis to employees below 20 years of age. As described in the data section, salaries are reported as multiples of the minimum wage. The average salary for our sample of interest is 2.6 times the minimum wage, which was set at 40 daily pesos in 2003 in the subdelegations of interest. Assuming employees in the manufacturing industry work for 22 days a month, the average monthly wage equals approximately 2,200 pesos. We compare this to the PROGRESA monthly transfers for the oldest beneficiaries, which amount to 660 pesos (Behrman, Parker, and Todd 2011).

are used as controls. The PROGRESA survey also collects information on the employment status, labor-market income, and migration status of other household members, which we use in some of the analyses.

Our main sample consists of all children aged 5–16 in the original survey (October 1997), living in households eligible for PROGRESA and with non-missing educational attainment in 2003.[6] This consists of over 23,000 individuals, belonging to over 8,000 households in 506 villages.

4.2. IMSS Data

We also use data on all formal private-sector jobs from the IMSS, from 1997 until 2003. The IMSS data include monthly records of the number of insured workers in each category, where a category is defined by location, industry, employer size, employee age, employee gender, and employee salary range. For example, one observation of this dataset provides the number of formal sector female workers employed in a particular month in a particular municipality, aged between 20 and 25, earning between 2 and 3 times the minimum salary, and working at a firm that hires between 51 and 250 employees, in a specific industry.

The IMSS data assign each firm to one of 276 industry categories, without indicating whether firms are export-oriented or not. Following Atkin (2016), we define export-oriented manufacturing firms as those which belong to a three-digit International Standard Industrial Classification (ISIC) industry where more than 50 percent of output was exported for at least one-half of the study's sample years (1986–2000).[7] Comparing export-oriented jobs to all other (formal sector) jobs in the IMSS in August 1997 (the first month of publicly available data) and focusing on the states in which PROGRESA was implemented, we find that export-oriented jobs are held by younger workers (31.1 compared to 35.8 years old on average), pay similar wages (3.3 times the minimum salary compared to 3.2 times the minimum salary), and have larger shares of female workers (0.39 compared to 0.30). Firms are also larger in export-oriented industries, with 52 percent of jobs (compared to 35 percent in non-export industries) found in firms with more than 250 workers.

For each month, we calculate the number of export-oriented manufacturing jobs, relative to the size of the working-age population (between 15 and 49 years of age), obtained from Mexico's 1990 census. Supplementary online appendix fig. S1.1 illustrates the geographic distribution of these jobs at the municipality level, in each of the seven PROGRESA states. Because the PROGRESA evaluation took place in rural areas, the majority of villages are located in areas with low exposure to export manufacturing, though a non-trivial number are located in municipalities or neighboring municipalities with export shares greater than 10 percent.

In our main specification, we separate jobs by gender because export manufacturing employment opportunities do appear to be somewhat gender segmented. For example, 69 percent of female export manufacturing jobs in PROGRESA states are in textiles (compared to only 24 percent of male jobs), while 25 percent of male jobs are in metals (compared to only 5 percent of female jobs). In some of the analyses, we also categorize jobs based on the salary range and age of the insured individual. We define low-wage jobs as those with a salary up to two times the statutory minimum salary, and high-wage jobs as those with a salary above this threshold. Similarly, we define young (old) export jobs as those with registered ages below (above) 25 years old.

Our analysis requires the use of a geographic unit within which there is sufficient variation in PROGRESA treatment status. The IMSS divides the country into delegations, which are further divided into subdelegations. Subdelegations are regional offices that serve as local branches of the IMSS, providing various administrative services for the region, such as enrollment and registration of affiliated individuals, collection of contributions from employers, and coordination of hospital admissions. Publicly available

6 The last restriction applies only to the educational attainment regressions, not the attendance regressions.

7 The resulting export industries are Apparel, Footwear, Leather and Leather Products, Wood and Cork Products, Petrochemical Refinement, Metal Products, Electronic and Mechanical Machinery, Electrical Machinery, Transport Equipment, Scientific and Optical Equipment.

Figure 1. Histogram of Subdelegation-Level Export Jobs Shares

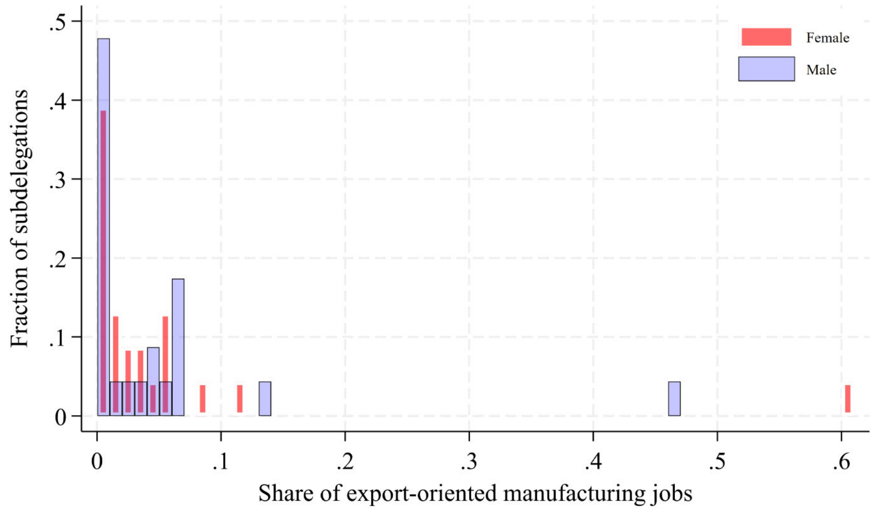

Source: Authors' analysis based on Mexican Social Security Institute (IMSS) and census data.
Note: Distribution of subdelegation-level export-job shares, calculated as the number of export-oriented manufacturing jobs held by men and women in each subdelegation, divided by the size of the working-age population aged 15–49.

health and employment data from the IMSS are provided at the state, delegation, subdelegation, and municipality levels (e.g., Doubova et al. 2021). Because over half of the municipalities in the sample were comprised of either all treatment or all control villages, we aggregate municipality-level counts of export jobs to the subdelegation level. We choose to use subdelegations over the larger delegations and states because there are only 7 states and 9 delegations in the PROGRESA evaluation sample but 23 subdelegations. Figure 1 illustrates the distribution of subdelegation-level export-job shares, separately for female and male jobs.

4.3. Summary Statistics

In supplementary online appendix table S1.1 we report summary statistics of individual and household characteristics in our sample of interest using data from the first available wave (which is the baseline survey in most cases). We split the sample into high export subdelegations (with above-median export-job shares) and low export subdelegations (with below-median export-job shares). At baseline, treated individuals are comparable (in terms of age, gender, school attendance, years of schooling, household composition, and parental characteristics) to those in the control group, within both types of regions.[8]

Because only households classified as poor were considered eligible for PROGRESA, we restrict our analysis, as most existing studies do, to this subset of the population. In addition, we restrict to children of school-going age during the experimental period—specifically, those aged 5 to 16 in 1997.[9]

8 For all parental variables, we assign children the values belonging to the household head and the household head's spouse. Over 90 percent of the children in our sample are sons or daughters of the household head.

9 If we assume that children start first grade at age 6 and do not repeat grades, children aged 5 to 13 in 1997 would have been in PROGRESA-eligible grades during the first 18 months of the program, while only the treatment group was exposed. We include three older age cohorts as they might also have been eligible due to schooling interruptions and grade repetitions. In the baseline survey, if we assume that 7-year-olds should have completed at least 1 year of school, 8-year-olds 2 years of school, and so on, we calculate that 44 percent of our main sample is at least 2 years behind and 24 percent is at least 3 years behind.

5. Empirical Strategy

5.1. Educational Attainment

Our first outcome of interest is educational attainment in 2003. By this time, PROGRESA was operating in both treatment and control villages, but treatment villages had been exposed to the program for 18 additional months. To estimate the heterogeneous effects of this additional exposure, we estimate the following specification:

$$E_{igjs} = \beta_1 T_j J_{sg} + \beta_2 T_j + \beta_3 J_{sg} + \beta_4 X_{ig} + \mu_s + \epsilon_{igjs}, \tag{3}$$

where E_{igjs} is the educational attainment of child i of gender g in village j and subdelegation s, as of 2003. Also, T_j is an indicator equal to 1 for the randomly assigned treatment villages, and J_{sg} is the number of export-oriented jobs in subdelegation s for gender g in 1997 (as a fraction of the subdelegation's working-age population according to the 1990 census). We use 1997 because it is the earliest year of publicly available IMSS data and the only year before the rollout of PROGRESA (which ensures it could not have been affected by the program itself).[10] By this time, a large portion of the expansion of export manufacturing jobs had already taken place. We use a gender-specific jobs variable to allow for female outcomes to be more affected by female jobs (and vice versa for male outcomes), as Atkin (2016) found to be the case.

To facilitate the interpretation of coefficient magnitudes, we standardize this variable. This means that β_2 represents the effect of PROGRESA for a subdelegation with the average number of export jobs. Here, our coefficient of interest is β_1, which captures heterogeneity in the PROGRESA effect across varying levels of export-job availability. A positive coefficient would indicate that PROGRESA is more effective in areas with more export jobs, while a negative coefficient would indicate that PROGRESA is less effective in these areas.

We include subdelegation fixed effects (μ_s), which control for subdelegation-specific unobservables that are fixed over time. Even with these fixed effects, we are relying on variation across subdelegations as well as variation across genders to estimate the interaction coefficient of interest. Even if J_{sg} were not gender specific, the inclusion of subdelegation fixed effects would still allow for the estimation of the interaction coefficient β_1, but not the coefficient on the main effect, β_3. In fact, although we use a gender-specific jobs variable to allow for men and women to respond to jobs that they are more likely to go into, our results are very similar when we use the total number of jobs instead of gender-specific jobs (see supplementary online appendix tables S1.2 and S1.3).

The variable X_{ig} is a vector of child-level controls. In our baseline specification, we include age and gender dummies. We later add demographic controls from the baseline survey: household size, age of household head, gender of household head, maternal and paternal education (dummies for secondary-school attendance), and maternal and paternal language dummies.[11] Because our export jobs variable (J_{sg}) is gender specific, to ensure that β_1 is not capturing gender differences in PROGRESA's education effects, we also add a treatment-by-female interaction in subsequent specifications. Finally, for all regressions, we cluster our standard errors at the village level, which was the level of treatment assignment.[12]

10 In robustness checks, to further alleviate concerns about the endogeneity of export jobs, we also use an alternate export manufacturing variable: predicted export-job growth generated using a shift-share strategy.

11 For continuous variables, we replace missing values with the sample mean. For parental education and language, we include a dummy for missing values.

12 We base this decision on Abadie et al. (2022). There are two "treatments" to consider in our case: the randomized PROGRESA treatment at the village level, and the effect of export jobs interacted with treatment (at the subdelegation level). At the village level, there is clustering in the sampling as well as in the treatment assignment, and we are arguing that there is heterogeneity in the effect of PROGRESA across villages: according to Abadie et al. (2022), we must therefore at least cluster at the village level. At the subdelegation level, however, because we are including subdelegation fixed effects, the conditions that determine whether to cluster are slightly different. In particular, if there are no heterogeneous

In order to ensure that our estimate of β_1 is not being confounded by PROGRESA treatment heterogeneity due to other variables potentially correlated with export jobs, we also estimate specifications that include interactions between treatment and other subdelegation-level and household-level characteristics, denoted C_{ijs} in the following regression:

$$E_{igjs} = \beta_1 T_j J_{sg} + \beta_2 T_j + \beta_3 J_{sg} + \beta_4 X_{ig} + \beta_5 T_j C_{ijs} + \beta_6 C_{ijs} + \mu_s + \epsilon_{igjs}. \tag{4}$$

We run separate regressions using different definitions of C_{ijs}: for example, subdelegation-level average schooling, income, and urban shares from the 1990 census (in which case the main effect is absorbed by the subdelegation fixed effects). We also use the following (all taken from the 1997 survey): child *i*'s baseline educational attainment, a vector of father and mother occupation category dummies, older sibling work status, household per capita labor income, and a vector of proxies for the temporary migration of household members (separate dummies indicating if a father or mother is not living at home, as well as the continuous share of household members not living at home).

5.2. Attendance

We next explore the contemporaneous effect of PROGRESA on school attendance, using all waves before 2003 (the October 1997 baseline survey, October 1998, October 1999, and November 1999). Specifically, for child i of gender g in village j and subdelegation s, observed in wave w, we estimate

$$\begin{aligned} A_{igjsw} = {} & \alpha_1 T_j P_w J_{sg} + \alpha_2 T_j P_w + \alpha_3 T_j J_{sg} + \alpha_4 P_w J_{sg} \\ & + \alpha_5 T_j + \alpha_6 J_{sg} + \alpha_7 X_{ig} + \mu_s + \delta_w + \epsilon_{igjsw}, \end{aligned} \tag{5}$$

where A_{igjsw} is a school-attendance dummy variable. As before, J_{sg} captures the number of export-oriented jobs in subdelegation s for gender g (as a fraction of the working-age population and standardized, as above) in 1997, and P_w is a dummy for post-treatment waves (all waves except the 1997 baseline).

The main coefficient of interest is α_1. This captures heterogeneity in the PROGRESA treatment effect across areas with varying export-job exposure. Including the baseline wave helps improve statistical precision and also builds in some validity checks. For example, we would expect α_5 (the difference between treatment and control villages prior to the rollout of PROGRESA) and α_3 (heterogeneity in this difference by export jobs) to be equal to zero.

As in the first specification, we include a vector of child and household controls (X_{ig}). We also estimate versions of this regression that add female interactions: a female dummy interacted with T_j, P_w, and $T_j P_w$.

We conduct a similar robustness exercise to the one described above, outlined in the regression below:[13]

$$\begin{aligned} A_{igjsw} = {} & \alpha_1 T_j P_w J_{sg} + \alpha_2 T_j P_w + \alpha_6 J_{sg} + \alpha_7 X_{ig} + \alpha_8 T_j P_w C_{ijs} + \alpha_9 C_{ijs} \\ & + \mu_s + \delta_w + \epsilon_{igjsw}. \end{aligned} \tag{6}$$

We use C_{ijs} to denote the same subdelegation and household-level characteristics described above.

5.3. Attrition

To explore whether differential attrition due to PROGRESA or export-job exposure could be complicating the interpretation of the coefficients defined above, we begin with the full sample of eligible children aged 5 to 16 at baseline. In supplementary online appendix table S1.4, we estimate equation (3) using a dummy

treatment effects at the subdelegation level then we should not be clustering at the subdelegation level. The relevant "treatment" to consider here is the export jobs interaction with treatment, and we argue this variable affects outcomes in the same way across individuals.

13 Note that we drop $T_j J_{sg}$, $P_w J_{sg}$, and T_j for parsimony after having established that these coefficients are small in magnitude and not statistically significant. In this specification, $T_j P_w$ can be thought of as a dummy variable equal to 1 if village j is treated in wave w.

equal to 1 for those who have a non-missing educational attainment variable in 2003 (and are therefore included in our analysis) as our dependent variable. The coefficient on the treatment indicator is negative and statistically significant, indicating that—for a village with the average level of export jobs—treatment individuals were 3 percentage points less likely to be included in our sample. In addition, the interaction coefficient is negative and significant, which means this gap was significantly larger in areas with more export jobs. While there are a number of potential reasons for this result, one explanation is that the cash transfers allowed treatment households or individuals to migrate away from their village (therefore leaving the sample), primarily in subdelegations with promising export-related job opportunities in nearby areas (not within commuting distance).

Regardless of the reason, this could complicate the interpretation of our coefficient estimates if the children who dropped out of the sample had systematically different education outcomes from those who remained. To address this concern, we also estimate a specification that uses an inverse probability weighting procedure similar to that of other studies on PROGRESA (Behrman, Parker, and Todd 2011, 2009a). Specifically, we estimate a probit regression using the indicator for sample inclusion as our outcome variable and the following set of independent variables: PROGRESA treatment, export jobs, their interaction, and all three of these variables interacted with the full set of demographic controls. This allows us to predict each individual's probability of being included in the sample given their treatment status, export-job level, and observable covariates (which are allowed to contribute to the prediction in different ways for different treatment groups and export-job levels). We can then estimate weighted regressions using the inverse of these predicted probabilities as weights, therefore accounting for changes in sample composition due to differential attrition.

6. Results

6.1. Main Results

Table 1 reports the results of equation (3). Across all columns, there is a positive and significant coefficient on the treatment dummy and a negative and significant coefficient on the interaction term, which indicates that PROGRESA improved educational attainment, but less so in areas with many export jobs. This pattern of results is robust to the inclusion of additional controls (columns 4 through 6) and treatment-by-female interactions (columns 2, 3, 5, and 6). The latter indicates that the Treat-by-ExportJobs coefficient is not simply picking up gender differences in the PROGRESA impact. In addition, our results are similar when we use inverse probability weighting to account for differential attrition (columns 3 and 6). The ExportJobs coefficient is statistically insignificant across all specifications.[14]

Estimates in column 6 reveal that PROGRESA increased educational attainment by 0.21 years for areas with the average number of export jobs, which is similar to estimates of around 0.2 years for the full sample (Behrman, Parker, and Todd 2011).[15] For an area that lies one standard deviation above the mean in terms of export jobs, the interaction coefficient of −0.27 implies no PROGRESA effect at all (while the sum of the two coefficients is in fact negative, it is not significantly different from zero). In supplementary online appendix fig. S1.2, we show the entire distribution of treatment effect magnitudes: the vast majority are positive, and only a small share are negative.

Because these regressions use educational attainment in 2003, when PROGRESA was available in both treatment and control villages, the estimated treatment effects can be interpreted as the effect of being

14 In these results, the lack of statistical significance could be a result of subdelegation fixed effects absorbing most of the variation in this variable. However, all coefficient estimates are very similar even without subdelegation fixed effects (results available upon request). Export jobs may not exhibit a significant relationship with educational attainment because of opposing forces canceling out: export firms may have targeted areas with better education outcomes ex ante but may have also increased the opportunity cost of schooling.

15 The ExportJobs variable is standardized to facilitate the interpretation of the main effect.

Table 1. Heterogeneous Effects of PROGRESA on Educational Attainment

	(1) Educational attainment	(2) Educational attainment	(3) Educational attainment	(4) Educational attainment	(5) Educational attainment	(6) Educational attainment
Treat × ExportJobs	−0.37	−0.37	−0.33	−0.27	−0.28	−0.27
	(0.15)**	(0.15)**	(0.12)***	(0.13)**	(0.13)**	(0.10)***
Treat	0.16	0.21	0.23	0.16	0.21	0.21
	(0.096)*	(0.11)**	(0.10)**	(0.087)*	(0.097)**	(0.095)**
ExportJobs	0.081	0.068	0.10	0.00070	−0.012	0.040
	(0.15)	(0.15)	(0.14)	(0.14)	(0.14)	(0.14)
Treat × Female	–	−0.095	−0.093	–	−0.098	−0.094
		(0.085)	(0.085)		(0.083)	(0.083)
Observations	23,272	23,272	23,272	23,272	23,272	23,272
Mean of DV	6.89	6.89	6.87	6.89	6.89	6.87
Controls	Basic	Basic	Basic	All	All	All
Weighted	No	No	Yes	No	No	Yes
p-value for sum	0.28	0.42	0.52	0.51	0.71	0.69

Source: Authors' analysis based on PROGRESA and Mexican Social Security Institute (IMSS) data.

Note: These regressions use the 2003 survey wave, restricting to children aged 5 to 16 at baseline (in 1997). Treat is an indicator for PROGRESA treatment villages. ExportJobs is the ratio of the gender-specific number of export-oriented jobs in the subdelegation in 1997, over the subdelegation's working-age population according to the 1990 census, standardized. "Basic" controls include gender, cohort fixed effects, and subdelegation fixed effects. "All" controls add household size, household head age, household head gender, as well as parental education and language dummies (including dummies for missing values). Weighted regressions use the attrition weights described in empirical strategy section, and "*p*-value for sum" reports the *p*-value testing the null hypothesis that $\beta_1 + \beta_2 = 0$. Standard errors (clustered at village level) in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

exposed to the program 18 months earlier. We now move on to investigate the intermediate changes leading up to these increases in educational attainment—that is, the contemporaneous effects of PROGRESA on school attendance during the years in which the control group had not yet received the program.

The attendance results, from equation (5), are reported in table 2. We estimate that PROGRESA increased attendance rates by approximately 3 percentage points for the average subdelegation. However, for subdelegations one standard deviation above the mean, the effect is 2 percentage points smaller (and not significantly different from zero). As was the case with educational attainment, attendance improved due to PROGRESA, but less so for areas with many export jobs. Results are robust to the inclusion of additional demographic controls, female interaction terms, and inverse probability weighting. The histogram of PROGRESA attendance effects (in supplementary online appendix fig. S1.3) reveals the majority of subdelegations demonstrated positive effects (most of which are significantly different from zero) and only a small share saw negative (but insignificant) effects.

Because treatment was randomly assigned and the program was not rolled out until after the baseline survey, we would expect to see no differences across treatment and control during the baseline survey. The small and statistically insignificant coefficient on Treat shows this is true. For similar reasons, we would not expect any job-related heterogeneity in the treatment-control gap in the baseline survey, which is confirmed by the statistically insignificant coefficient on Treat-by-ExportJobs. This regression also shows that control villages exhibit no change in the export jobs relationship after PROGRESA was implemented in the treatment villages (indicated by Post-by-ExportJobs coefficients that are statistically indistinguishable from zero). Having established this, we exclude Treat, Treat-by-ExportJobs, and Post-by-ExportJobs from the remaining attendance regressions for parsimony, as we continue to add new interaction terms. In this simpler specification, Treat-by-Post is equivalent to a dummy variable equal to 1 if the individual is in a village that is currently treated (as of the relevant wave), which we label as "Treated" in the tables.

Table 2. Heterogeneous Effects of PROGRESA on School Attendance

	(1) School attendance	(2) School attendance	(3) School attendance	(4) School attendance	(5) School attendance	(6) School attendance
Treat × Post × ExportJobs	−0.022	−0.021	−0.022	−0.022	−0.021	−0.022
	(0.0082)***	(0.0083)**	(0.0067)***	(0.0083)***	(0.0083)**	(0.0067)***
Treat × Post	0.029	0.026	0.027	0.029	0.026	0.026
	(0.0059)***	(0.0078)***	(0.0078)***	(0.0059)***	(0.0078)***	(0.0078)***
ExportJobs	−0.012	−0.012	−0.0097	−0.013	−0.013	−0.010
	(0.017)	(0.017)	(0.018)	(0.017)	(0.017)	(0.018)
Treat × ExportJobs	0.0012	0.00061	0.014	0.0022	0.0016	0.013
	(0.011)	(0.011)	(0.0096)	(0.011)	(0.011)	(0.0092)
Post × ExportJobs	0.00082	0.00078	−0.00072	0.00070	0.00067	−0.00092
	(0.0035)	(0.0035)	(0.0035)	(0.0035)	(0.0035)	(0.0035)
Treat	0.0058	0.0066	0.010	0.0051	0.0059	0.0091
	(0.0077)	(0.0086)	(0.0087)	(0.0076)	(0.0084)	(0.0084)
Observations	95,705	95,705	95,705	95,705	95,705	95,705
Mean of DV	0.83	0.83	0.83	0.83	0.83	0.83
Controls	Basic	Basic	Basic	All	All	All
Additional treatment interactions	None	By female	By female	None	By female	By female
Weighted	No	No	Yes	No	No	Yes
p-value for sum	0.49	0.64	0.70	0.53	0.69	0.68

Source: Authors' analysis based on PROGRESA and Mexican Social Security Institute (IMSS) data.

Note: These regressions use the 1997, 1998, and both 1999 survey waves, restricting to children aged 5 to 16 at baseline (in 1997). Treat is an indicator for PROGRESA treatment villages. ExportJobs is the ratio of the gender-specific number of export-oriented jobs in the subdelegation in 1997, over the subdelegation's working-age population according to the 1990 census, standardized. Post is an indicator for all waves after 1997. "Basic" controls include gender, cohort fixed effects, wave fixed effects, and subdelegation fixed effects. "All" controls add household size, household head age, household head gender, as well as parental education and language dummies (including dummies for missing values). "By female" treatment interactions include a female indicator interacted with Treat-by-Post (in all columns), in addition to a female indicator interacted with Treat and Post. Weighted regressions use the attrition weights described in the empirical strategy section, and "p-value for sum" reports the p-value testing the null hypothesis that $\alpha_1 + \alpha_2 = 0$. Standard errors (clustered at village level) in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6.2. Mechanisms

According to the model in our theoretical framework, the finding that export manufacturing jobs reduce the schooling impact of PROGRESA implies these jobs result in net benefits that decrease faster with each additional year of schooling. This could be due to export jobs changing the convexity of the perceived future wage function or the convexity of the forgone wage function. While second derivatives are generally difficult to measure, we present some descriptive analysis to shed light on these parameters. First, in fig. 2 A, we plot the quadratic relationship between total income and schooling among adults. We do this separately for subdelegations in the top quartile and those in the bottom quartile in terms of export jobs, in order to shed light on how export jobs might affect the convexity of the wage function. The solid red line, which represents high-export areas, has a steeper and more rapidly increasing slope compared to the dotted blue line, which represents low-export areas. In other words, the marginal benefits of schooling appear to be increasing faster in high-export areas. This implies a larger $\frac{\partial^2 W}{\partial S^2}$, which would predict larger CCT schooling effects for high-export areas—the opposite of what our results show. It therefore appears that the convexity of the wage function is not the driving force behind our results.

However, fig. 2 B shows that a comparison of opportunity costs leads to a different prediction—lower CCT schooling effects for high-export areas. Here, we plot the quadratic relationship between total income and schooling among out-of-school youths. While this is only a rough proxy for the relationship between forgone wages and schooling, it provides some insight into the convexity of the opportunity cost function. This curve is steeper for areas with many export jobs, suggesting that opportunity costs are more

Figure 2. Income-Schooling Relationship, by Export Job Quartiles

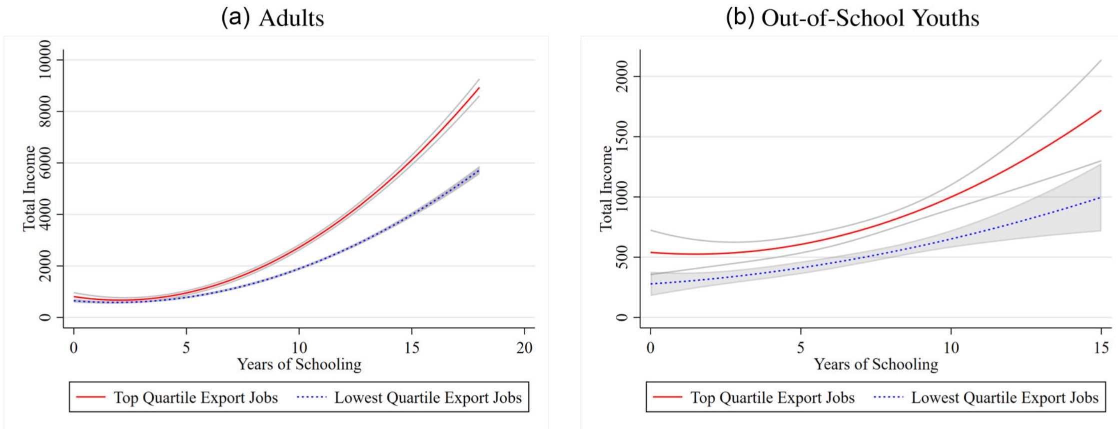

Source: Authors' analysis based on Mexican Social Security Institute (IMSS) and census data.
Note: Solid red and dotted blue lines depict the predicted quadratic relationship between income and schooling using the 2000 Mexican census, among individuals in the seven PROGRESA states. Panel A restricts to adults aged 25–55. Panel B restricts to out-of-school youths aged 13–20. Gray lines/regions represent 95 percent confidence intervals. Quartiles are defined by classifying subdelegations according to the number of export-oriented jobs (as a share of the working-age population) in 2000.

convex in these areas. This translates into a larger $\frac{\partial^2 c}{\partial S^2}$ for export areas, which would predict lower CCT education effects.

Our findings of smaller CCT schooling effects in areas with more export jobs suggest that the opportunity cost channel (specifically, more convex costs) dominates over the wage function channel. We provide further evidence for this claim by exploring what types of individuals and what types of jobs are driving the heterogeneity documented.

We first show the heterogeneity is stronger for those old enough to be actually working a factory job. By the time PROGRESA had begun, substantial growth in export manufacturing had already taken place, which means that the older children in our sample would have been exposed to the opportunity cost channel from the beginning of our study period. We use 15 as the cutoff age, as this is the median of the official minimum working age at the time (14) and the minimum working age without parental consent (16) (Atkin 2016). The first two columns of table 3 show that while the interaction term is negative and significant for those who would have been aged 15 for at least one year in the sample period (those 16 and older in 2003), it is smaller and insignificant for those who would have been too young.[16] This is made even clearer in supplementary online appendix fig. S1.4, which plots the entire distribution of treatment effects for each group, revealing substantially greater variance for the working-age group.[17]

We document a similar result for attendance effects, for which we split the sample into younger than 15 and those 15 and older at the time of the survey. As with educational attainment, the last two columns reveal a significant negative interaction term only for the working-aged and not the younger sample.[18]

16 The p-value for the test of the difference between these two coefficients is 0.11.
17 What is important here is the extent of the heterogeneity and not the size of the main effect. It has been documented previously that PROGRESA's effect on educational attainment was larger among those who were in later grades, and this can be seen by comparing the size of the Treated coefficient in columns 1 and 2. This is likely due to the fact that more students were on the margin of dropping out at these ages and does not necessarily speak to the argument that the opportunity cost channel was the dominant force driving heterogeneity by export jobs. It is the fact that the *heterogeneity* (by labor-market conditions) was larger for working-age individuals that adds support for this argument.
18 The difference between these coefficients is statistically significant at the 1 percent level.

Table 3. Heterogeneous Effects of PROGRESA for Working-Age Versus Younger Cohorts

	(1) Educational attainment	(2) Educational attainment	(3) School attendance	(4) School attendance
Treated × ExportJobs	−0.14	−0.39	−0.0045	−0.080
	(0.094)	(0.19)**	(0.0079)	(0.020)***
Treated	0.12	0.31	0.031	0.026
	(0.069)*	(0.14)**	(0.0069)***	(0.020)
ExportJobs	0.12	−0.12	−0.024	0.057
	(0.12)	(0.23)	(0.014)*	(0.041)
Observations	10,906	12,366	80,149	15,556
Mean of DV	5.785	7.871	0.916	0.410
Controls	All	All	All	All
Additional treatment interactions	By female	By female	By female	By female
Sample	Not working age	Working age	Not working age	Working age

Source: Authors' analysis based on PROGRESA and Mexican Social Security Institute (IMSS) data.

Note: Columns 1 and 3 use the 2003 survey wave, columns 2 and 4 use the 1997, 1998, and both 1999 survey waves, and all columns restrict to children aged 5 to 16 at baseline (in 1997). In columns 1 and 3, Treated = 1 for PROGRESA treatment villages; in columns 2 and 4, Treated = 1 if a village has PROGRESA at the time of the survey. ExportJobs is the ratio of the gender-specific number of export-oriented jobs in the subdelegation in 1997, over the subdelegation's working-age population according to the 1990 census, standardized. "All" controls include gender, cohort fixed effects, wave fixed effects, subdelegation fixed effects, household size, household head age, household head gender, as well as parental education and language dummies (including dummies for missing values). "Working age" is defined as those older than 15 (for educational attainment regressions) or those currently aged 15 or older (for attendance regressions). Standard errors (clustered at village level) in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Moreover, supplementary online appendix fig. S1.5 reveals a much larger variance of treatment effects for the working-age group.

Table 4 provides further support for the opportunity cost channel. Here, we examine whether the negative interaction coefficients reported above are being driven by the types of export jobs that would actually factor into the opportunity costs of school. Specifically, we differentiate between export jobs that are low wage (vs. high wage) and held by younger (vs. older) workers. These jobs are more obtainable for someone who drops out of school before graduating high school and are therefore more relevant to the opportunity cost function. While jobs held by workers with low education levels would also be a good proxy for this, education information is unfortunately not available in the IMSS. Though we run separate analyses for the age and wage distinctions, we note that the number of low-wage jobs in a subdelegation is highly correlated with the number of young jobs.

The results of table 4 reveal that the negative interaction coefficients reported above are indeed being driven by low-wage and young jobs. In columns 1 and 2, we include one interaction between treatment and export jobs among low-wage workers (earning less than double the minimum salary), and one interaction between treatment and export jobs among high-wage workers. For both educational attainment and school attendance, it is only the low-wage job interaction that generates a negative and significant coefficient. In columns 3 and 4, we repeat the exercise, this time including treatment interactions with young export jobs (25 years old and under) and older export jobs. In both columns, it is only the young export jobs variable that generates a negative interaction coefficient.

In sum, this evidence supports the idea that PROGRESA was less effective in areas with more export manufacturing because these types of jobs increase the convexity of the opportunity cost function. Although fig. 2 showed that export manufacturing jobs also increase the convexity of the wage function (which should lead to larger PROGRESA effects), our results indicate the marginal cost channel appears to dominate over the marginal benefits channel.

One possible reason for the importance of opportunity costs relative to future wages is migration, which could weaken the relationship between local (subdelegation-level) labor-market conditions and

Table 4. Heterogeneous Effects of PROGRESA using Different Types of Export Jobs

	(1) Educational attainment	(2) School attendance	(3) Educational attainment	(4) School attendance
Treated × ExportJobs (Type 1)	−0.21	−0.018	−0.37	−0.047
	(0.11)*	(0.0080)**	(0.34)	(0.023)**
Treated × ExportJobs (Type 2)	0.023	0.0079	0.14	0.034
	(0.12)	(0.012)	(0.37)	(0.025)
Treated	0.20	0.030	0.20	0.029
	(0.095)**	(0.0084)***	(0.096)**	(0.0083)***
ExportJobs (Type 1)	0.068	0.0085	0.13	0.036
	(0.094)	(0.0079)	(0.27)	(0.020)*
ExportJobs (Type 2)	−0.080	−0.026	−0.16	−0.054
	(0.16)	(0.017)	(0.35)	(0.028)*
Observations	23,272	95,705	23,272	95,705
Mean of DV	6.894	0.833	6.894	0.833
Controls	All	All	All	All
Type 1	Low wage	Low wage	Young	Young
Type 2	High wage	High wage	Old	Old
Additional treatment interactions	By female	By female	By female	By female

Source: Authors' analysis based on PROGRESA and Mexican Social Security Institute (IMSS) data.
Note: Columns 1 and 3 use the 2003 survey wave, columns 2 and 4 use the 1997, 1998, and both 1999 survey waves, and all columns restrict to children aged 5 to 16 at baseline (in 1997). In columns 1 and 3, Treated = 1 for PROGRESA treatment villages; in columns 2 and 4, Treated = 1 if a village has PROGRESA at the time of the survey. ExportJobs is the ratio of the gender-specific number of export-oriented jobs (defined by the specified type) in the subdelegation in 1997, over the subdelegation's working-age population according to the 1990 census, standardized. "All" controls include gender, cohort fixed effects, wave fixed effects, subdelegation fixed effects, household size, household head age, household head gender, as well as parental education and language dummies (including dummies for missing values). Standard errors (clustered at village level) in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

perceived future wages. Individuals might form their expectations about the future wage function using information from areas outside their subdelegation, especially in a setting like Mexico where 18 percent of residents (in 2000) were living in a state different from their state of birth; the share who have migrated across subdelegations is likely much larger.[19] Another possibility is that parental preferences might play an important role in the optimal schooling decision. If parents value current income more than their child's future income, this would result in marginal costs receiving a heavier weight in the maximization problem.

6.3. Robustness

Taken together, these results support the argument that export jobs reduce PROGRESA schooling effects by changing the shape of the opportunity cost function (as opposed to the wage function), consistent with the discussion in our theoretical framework. The validity of this interpretation, however, requires that the heterogeneity we document is not caused by some other correlate of the export jobs variable.[20] If exporting firms make decisions about where to locate or where to expand based on characteristics of a subdelegation, these characteristics might be generating the heterogeneity we document. Another possibility is that export jobs are correlated with household or individual characteristics and that PROGRESA treatment effects vary across these characteristics rather than export jobs.

19 We calculate this share using state-level migration data provided by INEGI. The data include the total population and the number of individuals living in a state different from the one they were born in. Data are available at https://www.inegi.org.mx/temas/migracion/.

20 Given the results discussed in table 3, any problematic correlate would have to generate heterogeneity for certain age groups and not others.

To evaluate this, we estimate regressions (4) and (6), which allow for heterogeneous effects of PROGRESA based on subdelegation and household characteristics, and report results in supplementary online appendix tables S1.5 (for educational attainment) and S1.6 (for attendance). Each column represents a different regression that controls for treatment interacted with a different subdelegation, household, or individual characteristic. Subdelegation characteristics are taken from the 1990 census, while household characteristics are taken from the baseline survey, which ensures that these variables were not responding to either PROGRESA treatment or the export jobs variable (both determined in 1997). We allow for heterogeneity with respect to subdelegation-level schooling, income, and urban shares (columns 1 to 3). In column 4, we allow for differential effects based on the child's educational attainment as of 1997. At the household level, we allow for heterogeneity by temporary migration proxies, household per capita labor income, father's occupation type, mother's occupation type, and sibling work status (columns 5 to 9). In the final column, we include all interaction terms from the previous columns. Reported coefficients can be interpreted as the effects for the average child (for continuous variables) or modal child (for categorical variables).[21]

In both tables, all specifications reveal treatment main effects and export-job interactions that are almost identical to those estimated in tables 1 and 2. Even in column 10 of supplementary online appendix tables S1.5 and S1.6, which include the entire set of interaction terms, the magnitudes of the coefficients of interest are similar to baseline estimates (though not statistically significant in supplementary online appendix table S1.5). In short, the treatment effect heterogeneity we document does appear to be driven by the availability of export jobs, and not by any of these other characteristics. It is worth noting that some of these characteristics do drive treatment heterogeneity. For example, attendance effects are smaller for children with higher baseline schooling and educational attainment effects are smaller for children with mothers who are employees (coefficients not reported but available upon request). Importantly, however, these other dimensions of heterogeneity do not appear to be confounding the estimates in our main specifications, which seem to be capturing what they were intended to—heterogeneity based on export-job availability.

Our next robustness check explores alternatives to the export jobs variable. We construct variables similar to a Bartik instrument, which combines industry composition in a baseline period with national-level industry growth rates to create a predicted employment growth variable arguably uncorrelated with location-specific changes that could be generating endogeneity problems. Specifically, for each subdelegation, we calculate the employment share in each export-oriented industry in a baseline period, multiply this by the national growth rate of the industry from the baseline period to period t, and sum across all export manufacturing industries to predict growth in export manufacturing from baseline to period t. We do this using the IMSS data in columns 1 and 2 of table 5 (where the only possibility for a baseline year is 1997) and census data in columns 3 and 4 (for which we use the 1990 census as our baseline). In the educational attainment regressions, we use the shift-share variable for $t = 2003$. In the attendance regressions, we use a time-varying variable, assigning each survey wave to the predicted year from the year before the survey. All of these regressions are reduced form regressions, where we simply replace our original export jobs variable with these shift-share variables. These results yield similar conclusions to our main regressions: positive PROGRESA treatment effects that are smaller in areas with lower predicted growth in export manufacturing.

21 Continuous variables are standardized so that the other coefficients can be interpreted as effects for an individual with average levels of the particular variable. For categorical variables, where interactions with several dummy variables are included in the regression, the omitted category is the model category, which means that coefficients represent effects for the modal individual. For example, most fathers are employees, which means that this is used as the omitted category and the coefficients reported in the table represent the effect of PROGRESA (and export-job heterogeneity) for children whose fathers are employees.

Table 5. Heterogeneous Effects of PROGRESA Using Alternative Export Jobs Variables

	(1) Educational attainment	(2) School attendance	(3) Educational attainment	(4) School attendance	(5) Educational attainment	(6) School attendance
Treated × ExportJobs	−0.38	−0.010	−0.24	−0.016	−0.55	−0.028
	(0.19)**	(0.0067)	(0.13)*	(0.0086)*	(0.28)**	(0.018)
Treated	0.15	0.036	0.18	0.029	0.13	0.027
	(0.12)	(0.0084)***	(0.11)*	(0.0083)***	(0.11)	(0.0092)***
ExportJobs	0.23	0.0061	−0.048	0.0023	0.28	−0.011
	(0.100)**	(0.0045)	(0.094)	(0.0069)	(0.28)	(0.024)
Observations	23,272	95,705	23,272	95,705	23,272	95,705
Mean of DV	6.894	0.833	6.894	0.833	6.894	0.833
Controls	All	All	All	All	All	All
Additional treatment interactions	By female	By female	By female	By female	By female	By female
Export job variable	Predicted growth (Bartik) from IMSS	Predicted growth (Bartik) from IMSS	Predicted growth (Bartik) from Census	Predicted growth (Bartik) from census	1997 export job share, winsorized	1997 export job share, winsorized

Source: Authors' analysis based on PROGRESA, Mexican Social Security Institute (IMSS), and census data.
Note: Columns 1 and 3 use the 2003 survey wave, columns 2 and 4 use the 1997, 1998, and both 1999 survey waves, and all columns restrict to children aged 5 to 16 at baseline (in 1997). In educational attainment regressions, Treated = 1 for PROGRESA treatment villages; in attendance regressions, Treated = 1 if a village has PROGRESA at the time of the survey. ExportJobs represents the specified export-job variable, standardized. "All" controls include gender, cohort fixed effects, wave fixed effects, subdelegation fixed effects, household size, household head age, household head gender, as well as parental education and language dummies (including dummies for missing values). Standard errors (clustered at village level) in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5 also explores the sensitivity of our results to outliers—specifically, the one subdelegation with substantially higher export manufacturing exposure than the rest, as documented in fig. 1. In columns 5 and 6, we winsorize our original export jobs variable at the 95th percentile, essentially replacing values from the top subdelegation with the next-highest values (for each gender). For educational attainment, this adjustment yields an even larger interaction coefficient than in our main specification. For attendance, the interaction term is no longer statistically significant, though it is slightly larger in magnitude than our baseline results.

7. Conclusion

In this paper, we highlight two theoretical channels through which labor-market conditions can influence the effectiveness of policies that reduce the cost of schooling. Job types can alter the convexity of an individual's perceived wage function, or the convexity of their opportunity cost function. The relative importance of these channels can determine the effectiveness of educational policies across labor markets. Empirically, we focus on Mexico, which rolled out its landmark CCT program, PROGRESA, during a period of trade liberalization that substantially increased the availability of export-oriented manufacturing jobs. This allows us to examine whether PROGRESA was more or less effective in areas with greater exposure to these export-oriented jobs. This exercise sheds light on the broader issue of how export promotion and CCTs—two common development policies—interact.

Although previous literature showed that export manufacturing jobs led to lower educational attainment in Mexico, whether these jobs reduce or enhance CCT effectiveness at increasing educational attainment remained an open question. Our empirical analysis answers this question, showing that PROGRESA was less successful at improving schooling outcomes in areas with greater exposure to export manufacturing. These results, combined with insights from our model and additional descriptive evidence, demonstrate that the opportunity cost channel dominates over the wage benefits channel in this context. Consistent with this, we show that the heterogeneous effects of PROGRESA are driven primarily by jobs that are likely to factor into the opportunity cost of schooling—specifically, low-wage jobs and jobs for younger workers.

Our results highlight and explain why conclusions from individual evaluations may not generalize to settings with different labor markets. This echoes the lessons of recent work highlighting the difficulties involved in generalizing from the results of individual well-identified studies: treatment effects often vary widely across settings, individuals, and over time (Card, Kluve, and Weber 2018; Rosenzweig and Udry 2020; Dehejia, Pop-Eleches, and Samii 2021; Meager 2022). Given the widespread popularity of CCTs across the developing world, it is important to understand what drives variation in the success of these programs, and our findings show that the types of jobs available to program beneficiaries play an important role. More generally, this paper provides evidence that labor-market conditions influence the effectiveness of government policies, which could be one understudied explanation for why the effects of minimum wage policy, health insurance expansions, financial aid programs, and other government policies differ drastically across settings.

Data availability statement

This paper utilizes four publicly available data sources.

- Data from the PROGRESA program can be accessed at the following link: https://dataverse.harvard.edu/dataset.xhtml?persistentId=hdl:1902.1/18235
- Data from the Mexican Institute of Social Security (IMSS) can be accessed at https://datos.imss.gob.mx/. The IMSS offers data on various aspects of social security, including healthcare, pensions, and workplace injury insurance.
- Replication files for Atkin (2016), used to characterize export-oriented sectors, can be accessed via https://www.aeaweb.org/articles?id=10.1257/aer.20120901. These files contain detailed data on export-oriented industries and their economic impact.
- Mexican census data can be accessed at IPUMS International (https://international.ipums.org/international/)

Code used for the analysis can be found at https://osf.io/f5uhq/.

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge. 2023. "When Should You Adjust Standard Errors for Clustering?." *Quarterly Journal of Economics* 138(1): 1–35.

Adhvaryu, A., A. Nyshadham, T. Molina, and J. Tamayo. 2024. "Helping Children Catch Up: Early Life Shocks and the PROGRESA Experiment." *The Economic Journal* 134(657): 1–22.

Aguayo-Tellez, E., J. Airola, C. Juhn, and C. Villegas-Sanchez. 2013. "Did Trade Liberalization Help Women? The Case of Mexico in the 1990s." *Research in Labor Economics* 38.

Angelucci, M., G. De Giorgi, M. A. Rangel, and I. Rasul. 2010. "Family Networks and School Enrolment: Evidence from a Randomized Social Experiment." *Journal of Public Economics* 94(3–4): 197–221.

Araujo, M. C., and K. Macours. 2021. "Education, Income and Mobility: Experimental Impacts of Childhood Exposure to PROGRESA After 20 Years." IDB Working Paper Series No. IDB-WP-01288.

Atkin, D. 2016. "Endogenous Skill Acquisition and Export Manufacturing in Mexico." *American Economic Review* 106(8): 2046–85.

Autor, D. H., D. Dorn, and G. H. Hanson. 2016. "The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade." *Annual Review of Economics* 8: 205–40.

Behrman, J. R., S. W. Parker, and P. E. Todd. 2009a. "Medium-Term Impacts of the Oportunidades Conditional Cash Transfer Program on Rural Youth in Mexico." *Poverty, Inequality and Policy in Latin America*: 219–70.

———. 2009b. "Schooling Impacts of Conditional Cash Transfers on Young Children: Evidence from Mexico." *Economic Development and Cultural Change* 57(3): 439–77.

———. 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.

Behrman, J. R., P. Sengupta, and P. Todd. 2005. "Progressing Through PROGRESA: An Impact Assessment of a School Subsidy Experiment in Rural Mexico." *Economic Development and Cultural Change* 54(1): 237–75.

Blanchard, E. J., and W. W. Olney. 2017. "Globalization and Human Capital Investment: Export Composition Drives Educational Attainment." *Journal of International Economics* 106: 165–83.

Cámara de Diputados. 2004. "Evolución del Sector Manufacturero de México, 1980-2003." Technical report, Mexico City: Centro de Estudios de las Finanzas Públicas.

Card, D., J. Kluve, and A. Weber. 2018. "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations." *Journal of the European Economic Association* 16(3): 894–931.

Cascio, E. U., and A. Narayan. 2022. "Who Needs a Fracking Education? The Educational Response to Low-Skill-Biased Technological Change." *ILR Review* 75(1): 56–89.

C Dammert, A.. 2009. "Heterogeneous Impacts of Conditional Cash Transfers: Evidence from Nicaragua." *Economic Development and Cultural Change* 58(1): 53–83.

Dehejia, R., C. Pop-Eleches, and C. Samii. 2021. "From Local to Global: External Validity in a Fertility Natural Experiment." *Journal of Business & Economic Statistics* 39(1): 217–43.

Diario Oficial de la Federación. 2019. "Decreto por el que se crea la Coordinación Nacional de Becas para el Bienestar Benito Juárez." https://www.dof.gob.mx/nota_detalle.php?codigo=5561693&fecha=31/05/2019. Accessed: 2021-01-02.

Djebbari, H., and J. Smith. 2008. "Heterogeneous Impacts in PROGRESA." *Journal of Econometrics* 145(1-2): 64–80.

Doubova, S. V., H. H. Leslie, M. E. Kruk, R. Pérez-Cuevas, and C. Arsenault. 2021. "Disruption in Essential Health Services in Mexico During COVID-19: An Interrupted Time Series Analysis of Health Information System Data." *BMJ Global Health* 6(9): e006204.

Edmonds, E. V., P. Topalova, and N. Pavcnik. 2009. "Child Labor and Schooling in a Globalizing World: Some Evidence from Urban India." *Journal of the European Economic Association* 7(2-3): 498–507.

Epple, D., R. E. Romano, and M. Urquiola. 2017. "School Vouchers: A Survey of the Economics Literature." *Journal of Economic Literature* 55(2): 441–92.

Fiszbein, A., N. R. Schady and F. Ferreira. 2009. *Conditional Cash Transfers: Reducing Present and Future Poverty*. 1. Washington D.C.: World Bank Publications. Report Number 47603.

Galiani, S., and P. J. McEwan. 2013. "The Heterogeneous Impact of Conditional Cash Transfers." *Journal of Public Economics* 103: 85–96.

Glewwe, P., and K. Muralidharan. 2016. "Improving Education Outcomes in Developing Countries: Evidence, Knowledge Gaps, and Policy Implications." In *Handbook of the Economics of Education*, 5 653–743. Elsevier.

Greenland, A., and J. Lopresti. 2016. "Import Exposure and Human Capital Adjustment: Evidence from the US." *Journal of International Economics* 100: 50–60.

Handa, S., B. Davis, M. Stampini, and P. C. Winters. 2010. "Heterogeneous Treatment Effects in Conditional Cash Transfer Programmes: Assessing the Impact of PROGRESA on Agricultural Households." *Journal of Development Effectiveness* 2(3): 320–35.

Hanson, G. H., and A. Harrison. 1999. "Trade Liberalization and Wage Inequality in Mexico." *ILR Review* 52(2): 271–88.

Herbaut, E., and K. Geven. 2020. "What Works to Reduce Inequalities in Higher Education? A Systematic Review of the (Quasi-) Experimental Literature on Outreach and Financial Aid." *Research in Social Stratification and Mobility* 65: 100442.

Jensen, R. 2010. "The (Perceived) Returns to Education and the Demand for Schooling." *Quarterly Journal of Economics* 125(2). 515–548.

———. 2012. "Do Labor Market Opportunities Affect Young Women's Work and Family Decisions? Experimental Evidence from India." *Quarterly Journal of Economics* 127(2): 753–92.

Juhn, C., G. Ujhelyi, and C. Villegas-Sanchez. 2014. "Men, Women, and Machines: How Trade Impacts Gender Inequality." *Journal of Development Economics* 106: 179–93.

Lederman, D., M. Olarreaga, and L. Payton. 2010. "Export Promotion Agencies: Do They Work?" *Journal of Development Economics* 91(2): 257–65.

Lee, S., and A. M. Shaikh. 2014. "Multiple Testing and Heterogeneous Treatment Effects: Re-Evaluating the Effect of PROGRESA on School Enrollment." *Journal of Applied Econometrics* 29(4): 612–26.

Lustig, N. 2001. "Life Is Not Easy: Mexico's Quest for Stability and Growth." *Journal of Economic Perspectives* 15(1): 85–106.

Maluccio, J., and R. Flores. 2005. *Impact Evaluation of a Conditional Cash Transfer Program: The Nicaraguan Red de Protección Social*. Washington, D.C.: International Food Policy Research Institute.

Manley, J., S. Gitter, and V. Slavchevska. 2013. "How Effective Are Cash Transfers at Improving Nutritional Status?" *World Development* 48: 133–55.

Meager, R. 2022. "Aggregating Distributional Treatment Effects: A Bayesian Hierarchical Analysis of the Microcredit Literature." *American Economic Review* 112(6): 1818–47.

Millán, T. M., T. Barham, K. Macours, J. A. Maluccio, and M. Stampini. 2019. "Long-Term Impacts of Conditional Cash Transfers: Review of the Evidence." *World Bank Research Observer* 34(1): 119–59.

Minnesota Population Center. 2015. *Integrated Public Use Microdata Series, International: Version 6.4 [Machine-readable database]*. Minneapolis: University of Minnesota.

Moreno-Brid, J. C. 2007. "Economic Development and Industrial Performance in Mexico post-NAFTA." *Taller Nacional sobre "Migración interna y desarrollo en México: Diagnóstico, perspectivas y políticas* 16. https://www.cepal.org/sites/default/files/courses/files/jcmoreno.pdf

Nallari, R., B. Griffith, Y. Wang, S. Andriamananjara, D. H. Chen, and R. Bhattacharya. 2012. "*World Bank Publications* The World Bank Group 2390 A Primer on Policies for Jobs"

Nguyen, T. 2008. "Information, Role Models and Perceived Returns to Education: Experimental Evidence from Madagascar." MIT.

Ordóñez-Barba, G. M., and A. L. Silva-Hernández. 2019. "Progresa-Oportunidades-Prospera: Avatares, alcances y resultados de un programa paradigmático contra la pobreza." *Papeles de Población* 25(99): 77–111.

Page, J. 2016. "*The East Asian Miracle and Development Policy: A Twenty-Year Retrospective*" In:H. ,J. , andY. (eds). Japan's Development Assistance. London: Palgrave Macmillan.

Parker, S. W., and P. E. Todd 2017. "Conditional Cash Transfers: The Case of Progresa/Oportunidades." *Journal of Economic Literature* 55(3): 866–915.

Parker, S. W., and T. Vogl. 2023. "Do Conditional Cash Transfers Improve Economic Outcomes in the Next Generation? Evidence from Mexico." *Economic Journal* 133(655): 2775–806.

Polaski, S. 2003. "Jobs, Wages, and Household Income." *NAFTA's Promise and Reality: Lessons from Mexico for the Hemisphere*: 11–38.

Revenga, A. 1997. "Employment and Wage Effects of Trade Liberalization: The Case of Mexican Manufacturing." *Journal of Labor Economics* 15(S3): S20–S43.

Rosenzweig, M. R., and C. Udry. 2020. "External Validity in a Stochastic World: Evidence from Low-Income Countries." *Review of Economic Studies* 87(1): 343–81.

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

Shah, M., and B. M. Steinberg. 2019. "Workfare and Human Capital Investment: Evidence from India." *Journal of Human Resources* 56(2): 380–405.

Skoufias, E. 2005. "PROGRESA and Its Impacts on the Welfare of Rural Households in Mexico." Technical Report 139, International Food Policy Research Institute.

Skoufias, E., and S. W. Parker. 2001. "Conditional Cash Transfers and Their Impact on Child Work and Schooling: Evidence from the PROGRESA Program in Mexico." *Economia* 2(1): 45–96.

United Nations. 2016. "Transforming our World: The 2030 Agenda for Sustainable Development."

World Bank Group. 2017. "Closing the Gap: The State of Social Safety Nets 2017." Technical report. Washington, DC.