J Popul Econ
DOI 10.1007/s00148-008-0207-6

ORIGINAL PAPER

Are parents altruistic? Evidence from Mexico

Christian Schluter · Jackline Wahba

Received: 13 April 2007 / Accepted: 1 April 2008

Abstract Are parents altruistic or selfish? We contribute to the continuing debate of this question by proposing a simple test which is implemented using experimental data from the Mexican anti-poverty programme PROGRESA. Benefit eligibility is randomised. Our estimation strategy explicitly addresses potentially confounding factors and selection bias problems. We reject selfishness of parents in non-urban Mexico as PROGRESA beneficiaries spend more on child-related goods and do not increase spending on adult-related goods compared to parents in the control group. At the same time, we reject some rival theories.

Keywords Parental altruism · Poor households · PROGRESA

JEL Classification D12 · D64 · I38

1 Introduction

Are parents altruistic or selfish? The empirical and theoretical literatures both appear to be divided on this point, and so far, no consensus has emerged. One side in this continuing debate might be labelled the "neo-Marshallian view", according to which children are mainly considered as a source of income and a provider of services. As Arrondel and Masson (2003) put it: "Indeed, for a

Responsible editor: Deborah Cobb-Clark

C. Schluter · J. Wahba (✉)
Economics Division, University of Southampton,
Highfield, Southampton SO17 1BJ, UK
e-mail: j.wahba@soton.ac.uk

"neo-Marshallian" view, parental altruism should have risen over time with "civilization" and the increase of well-being, the development of the welfare state and higher investments in education; presumably, it was much lower in old Europe when children were mainly considered as a resource for retirement, education remained low, and child work was quite widespread, and remains limited in poor LDCs". For a detailed exposition of this view, see, e.g. Caldwell (1978). The other side of the debate is perhaps best represented by the self-perception of many parents: "When they give it [monetary benefits] to me, sometimes [my children] have shoes and clothes, I save it for food and if I see that my children need a pair of pants, or something like it, I buy a pair of pants for my children. But it is always for them." (a PROGRESA stakeholder quoted in Skoufias 2000). Or in the words of Becker and Barro (1988): "*Obviously* many parents are altruistic towards their children in the sense that the utility of the parents depends positively on the utility of the children" (p.4, our italics). The lack of consensus is reflected in the theoretical literature, some authors assuming altruism, others parental selfishness. Meanwhile, the empirical literature often tests excessively narrow interpretations (e.g. whether adult children and parents, both economically active but geographically separated, operate on a common budget constraint). Therefore, it is perhaps unsurprising that the empirical literature typically rejects the parental altruism hypothesis. Whether parents are "selfish" (Becker 1981a) or altruistic is an important question, and plays a key role in the economics of fertility and child labour and for the design of aid and transfer programmes. Child poverty and child labour remain endemic in many countries, and international organisations such as the World Bank deploy vast resources to alleviate both. Such aid programmes are often designed on the premise of parental altruism. If parents are not altruistic but selfish, social programmes using cash benefits rather than (assignable) benefit in kinds will not yield the sought outcome. We seek to add to this debate by proposing a new test for parental altruism based on experimental data about the Mexican poverty alleviation programme known by its acronym as PROGRESA (*Programa de Educación, Salud y Alimentación*). Given the "neo-Marshallian view", we examine the parental altruism hypothesis on a tough testing ground: poor parents in a less developed country (LDC).

Our approach to testing parental altruism has three key features. First, we follow the literature and interpret altruism as utility interdependence: altruistic parents care about the welfare of their children. Selfish parents do not, use children as a source of income now or later, and maintain the consumption of child-related goods at some minimum level. Testing for altruism vs selfishness is a strong hypothesis, but it arises naturally from interpreting altruism as utility interdependence and captures precisely the conflict between the neo-Marshallian view and the parental self-perception quoted above. Finally, utility interdependence vs independence is precisely the set of hypotheses which are examined in the well-known literature discussed below in Section 1.1.

Second, we focus on children during childhood, and the relationship between child outcomes and parental resources. Most altruism tests in the literature study inter vivos transfers: children have already reached adulthood,

are wage earners, and have typically left the parental home, while parents are still economically active and typically rich enough to leave bequests. These altruism tests typically test for income pooling as altruistically linked units operate on a common budget constraint, and the principal policy concern is Ricardian equivalence.

Third, to the best of our knowledge, ours is the first paper to base a test of parental altruism on exogenous variations in incomes. All other tests use cross-sectional variation in incomes, which may be potentially endogenous. See, e.g. Lundberg et al. (1997) for an extensive criticism of basing tests which focus on expenditure on nonexogenous variations in income. They exploit exogenous variations in (maternal) income, but the authors focus on income pooling. The income variation arises as a natural experiment in the UK of switching child benefit payment from the father to the mother. Our altruism test is based on experimental panel data on PROGRESA which exploits variations across time (before and after the introduction of the benefit) and across households (in treatment and control groups). We describe the programme in detail in Section 2. Briefly, the principal element of PROGRESA is to give cash transfers to poor parents conditional on the regular school attendance of their children. Communities were randomly assigned into treatment and control areas, and only poor households in the treatment group have initially received the PROGRESA benefit. Our test for parental altruism is based on the marginal income effect of the exogenously assigned PROGRESA benefit on assignable expenditures.

The strength of the randomised experiment results in a simple test of the parental altruism hypothesis. In particular, we consider differences-in-differences in child-related expenditures between subgroups of PROGRESA beneficiaries and control households, as well as assignable adult goods. Food consumption is not assignable in our data, and therefore, it is treated as a household public good. We also consider parental leisure. Increases in item expenditures relating to meeting eligibility rules (such as child clothing for school attendance) are differenced out. Essentially, our procedure is to test the income effect stemming from the exogenously assigned PROGRESA benefit on these goods. Exploiting the experimental nature of the data allows us to make causal inferences about the implicit structural altruism parameter.

The test explicitly confronts the confounding impact of various factors. First, if parents are not altruistic but selfish, and use children as sources of income or to provide services, they still need to incur expenses which enable the children to carry out these tasks. Such services are considered explicitly in Cox (1987), and include 'market' activities (i.e. home production) and 'non-market' activities (such as companionship and obedience). In our model, if children are sent to school for selfish reasons, they still need to be clothed. Or parents might have to incur costs to meet eligibility conditions. Second, the schooling decision is taken facing a subsistence constraint since we are studying poor households in a LDC. If subsistence is threatened, children will have to work. Moreover, the PROGRESA benefit is conditional on school attendance. Hence, we do not seek to interpret the schooling decision or the

incidence of child labour in terms of parental altruism (so schooling does not enter the parental utility function as in Ravallion and Wodon 2000). Finally, an observed positive marginal income effect on child-related goods could be rationalised with rival models in which parents are selfish. For instance, in the Cox (1987) model, parents decide on both transfer to children and the level of service provision. There, non-altruistic parents provide compensation for child services. For this reason, we also consider adult-related goods including parental leisure.

The outline of the paper is as follows: In Section 1.1, we briefly review some tests for altruism presented in the literature. Aspects of the benefit programme and the data are briefly described in Section 2. The altruism test is set out and discussed in Section 3. The test for parental altruism is implemented in Section 3.3, while in Section 3.4, we carry out an extensive sensitivity analysis. We conduct various tests to eliminate further potential confounding factors. Section 4 concludes.

1.1 The related literature

We briefly review some prominent altruism tests proposed in the empirical literature in order to highlight the similarities and differences of our approach. Specifically, we focus on the age of the children and the exogeneity of income variations. Interpreting altruism as utility interdependence is standard (see, e.g. the surveys by Arrondel and Masson 2003 or Bergstrom 1997. The classic reference is Becker 1981b).

We focus on children during their childhood. Most altruism tests in the literature study inter vivos transfers: children have already reached adulthood, are wage earners, and have typically left the parental home, while parents are still economically active and typically rich enough to leave bequests. This class of altruism test examines income pooling as altruistically linked units operate on a common budget constraint. In Cox (1987), the outcome studied is the parental demand for services provided by children. If parents are altruistic, their demand depends only on aggregate income, not its distribution between children and parents. Non-altruistic parents use transfers to provide compensation for child services (and do not equalise marginal utilities of consumption). Altonji et al. (1992) test whether the distribution of food consumption is independent of the distribution of resources within extended families controlling for total resources, based on Panel Study of Income Dynamics (PSID) child split offs data for the USA. A similar test was conducted by Hayashi (1995), who estimated food Engel curves from a cross section of Japanese two-generation households. A different test for altruism is implemented in Altonji et al. (1997), who study inter vivo transfers directly by estimating transfer derivatives. Under income pooling, redistributing one dollar from a recipient child to donor parents should lead to a one-dollar increase in the parents' transfer to the child. The test is again carried out on PSID data. Cox et al. (1998) following Cox (1987) also focus on transfer derivatives and estimate a transfer equation controlling for selection (i.e. transfer receipt) on US and Peruvian data. They

juxtapose the predictions of the altruism model and a bargaining model. The principal motivation underlying all these studies is Ricardian equivalence: an altruistically linked extended family will neutralise most of a government's redistribution. All these papers reject income pooling. A first test for parental altruism for young children is provided in Bhalotra (2001), who, as us, interprets consumption allocations in terms of altruism. As in Hoddinott and Haddad (1995), goods considered are child and adult clothing, food, and tobacco. However, our empirical strategy is completely different, as we exploit truly exogenous variations in income arising from the experimental nature of the programme. Moreover, we do not estimate parametric demand functions but make non-parametric comparisons of means. Bhalotra tests the hypothesis of parental altruism by investigating whether "adult consumption is decreasing in child labour (...) Using a large household survey for Pakistan, m-demands for adult consumption are estimated, distinguishing tobacco, tea and coffee and adult clothing and footwear."[1] In her specification of the m-demands, child labour is taken to be the reference good, and the estimation is carried out using IV and semi-parametric techniques. We do not focus on child labour. Moreover, as Lundberg et al. (1997) argue extensively, basing tests which focus on expenditure on non-exogenous variations in income is very problematic since the cross-sectional income variations are potentially endogenous. These authors exploit exogenous variations in (maternal) income, but they focus on income pooling (and hence inter-vivo transfers). The income variation arises in time as a natural experiment in the UK of switching child benefit payment from the father to the mother. Given our panel structure, we observe an exogenous variation in income across time (before and after PROGRESA started) and across households (in treatment and control groups).

Income pooling is also tested in Attanasio and Lechene (2002) and Rubalcava et al. (2002), who use, as we do, PROGRESA data. These authors perform the well-known test of the 'unitary' model of household decision making by running cross-sectional regressions of budget shares on total expenditure and female income shares (which relate to PROGRESA since the benefit is paid to the mother). Income pooling is rejected as female income share coefficients are significant, and PROGRESA is inferred to have increased female bargaining power in the household. Moreover, Rubalcava et al. (2002) find evidence of improved diets and better child clothing.

We also consider expenditure items (child goods, adult goods and food), but our estimation strategy is different. We do not estimate parametric demand functions. Rather, we perform non-parametric comparisons of mean outcomes

[1]Denote by C_P adult consumption, and L_C child leisure. Her test of parental altruism is based on $\partial C_P/\partial L_C$. "If parents are altruistic then $\partial C_P/\partial L_C > 0$. Altruism is defined as the appearance of L_C, child leisure (or schooling), in the utility function of parents. If L_C does not appear in [the utility function], or if parents are not altruistic, then (...) $\partial C_P/\partial L_C = 0$." (p.10) In a personal communication received after completing this paper, Bhalotra has alerted us to a revised version of her paper, in which she has lessened the focus on child labour. She now tests whether the demand for child goods is increasing in adult consumption, and the altruism coefficient is $(\partial C_P/\partial L_C)^{-1}$.

in two periods between treatment and control households, i.e. a difference-in-difference computation. This estimation strategy also controls for item expenditures relating to meeting eligibility rules (such as child clothing for school attendance). Moreover, we also examine the parental good leisure (which is not considered by these authors). A link to Rubalcava et al. (2002) can be made by assuming in their model that the mother is more altruistic than the (possibly selfish) father. However, it then still needs to be argued why PROGRESA has not led to any increase in paternal clothing or tobacco. Our interpretation is parental altruism. Seeking to identify the actual household decision making process is beyond the scope of our paper.

Our testing strategy exploits many features of the PROGRESA data set. This data set has been used typically in evaluating the impact on the programme on child-related outcomes [among others, Skoufias (2001), Skoufias and Parker (2001) and Schultz (2004)]. In this literature, the focus is on whether treatment impacts on the chosen outcome. We do not carry out a programme evaluation but interpret child outcomes in terms of parental altruism.

2 The PROGRESA programme and the data

Our investigation is based on a unique large scale data set from PROGRESA, which is a major welfare programme run by the Mexican government in rural areas. See Skoufias (2001) for an extensive description and a summary of the programme's evaluations carried out by IFPRI, PROGRESA's contracted evaluator.[2] Overall, these evaluations find that PROGRESA had an effect on the examined outcomes (e.g. child labour, health and schooling). We use data from the evaluation surveys collected between November 1997 and November 1999. At the end of 1999, PROGRESA covered approximately 2.6 million families or about 40% of all rural families in Mexico and had a budget of approximately $777 million, the equivalent of 0.2% of Mexico's GDP. Since then, the programme has grown in coverage and size, which has been reflected in the change of its name from PROGRESA to Oportunidades. It now covers a larger set of rural households and expanded in 2002 into urban areas. Five million families currently benefit from Oportunidades. The programme has served as a model for other Latin American countries. Since we use the evaluation surveys 1997–1999, we refer to the programme as PROGRESA.

PROGRESA is a multifaceted programme made up of three closely linked components related to education, health and nutrition. The programme provides educational grants to families, designed to promote school enrolment and attendance, provides resources for improving the quality of schools,

[2]The data have recently been made publically available at http://evaloportunidades.insp.mx/en/index.php. The programme's web site is http://www.oportunidades.gob.mx/. See also http://www.ifpri.org/themes/progresa.htm. We have used the data distributed by IFPRI.

provides resources for increasing the quality and availability of health care, and gives direct monetary transfers and nutritional supplements to families. PROGRESA's monetary educational grants for each child (less than 18 years of age) are made conditional on school enrolment and a minimum 85% attendance rate. The size of the grant increases as children progress to higher grades, in order to give incentives for continuing education (especially moving from primary to secondary school). Grants are slightly higher for girls than for boys.[3] The benefit is paid to the mother and constitutes a substantial proportion of the recipient's household income. In November 1998, the benefit could range between 100 and 625 pesos per month. Based on administrative records, Rubalcava et al. (2002) report that the average benefit was 275 pesos per month, which constituted 29% of average household income.

A baseline survey (the PROGRESA Census ENCASEH97) was undertaken in October/November 1997 and was followed up with subsequent surveys (ENCEL) at approximately half-year intervals. The survey was conducted in 320 randomly selected treatment localities (in which treatment was initiated soon after the baseline survey) and 186 control localities (in which there has been no treatment during our observation window). All 506 of these localities were selected in a stratified random selection procedure (with stratification by populations of localities) from the localities identified by PROGRESA administrators as being eligible to participate in the programme. Our statistical procedures take account of this clustering. We focus on households classified as eligible for PROGRESA by the programme administrators based on a marginalisation index. We refer to these households as poor and have excluded non-eligible households from the analysis.[4] Eligible households in control localities started to receive benefits after November 1999 (which was not known to control households prior to inclusion). We use the evaluation surveys 1997–1999.

Our altruism test exploits two features of the data: we use the experimental nature of the programme by focussing on the exogenous change in income due to PROGRESA, and we also employ the panel aspects of the data set. In order to minimise the confounding impact of secular changes, we concentrate on child- and adult-related outcomes before and immediately after the implementation of PROGRESA. Rounds 1 and 2 (November 1997 and March 1998) are referred to below as pre-PROGRESA and round 3 (October 1998) and round 5 (November 1999) as PROGRESA rounds.

2.1 The data

In the subsequent examination, we refer to data from the ENCEL1998 November survey (or synonymously round 3) as the PROGRESA period.

[3]For specific amounts, see Table 1 in Skoufias (2001).

[4]We have also excluded the larger set of so-called *densificados*, who gained eligibility after the programme has started. These are somewhat richer households with older heads and fewer children.

Table 1 Assuming assumption 1 for our sample: differences in mean expenditures pre-PROGRESA

	Expenditures			Budget share [%]		
	T	C	d_g	T	C	d_g
Boys' clothing	24.52	24.06	0.28	2.69	2.56	0.92
Girls' clothing	22.55	22.30	0.23	2.46	2.36	0.69
Toys	2.07	2.16	−0.20	0.18	0.19	−0.13
Food	624.48	621.86	0.11	71.87	71.68	0.23
Tobacco	3.25	3.28	−0.05	0.31	0.36	0.79
Mens' clothing	15.53	14.77	0.62	1.61	1.52	0.82
Womens' clothing	13.22	12.43	0.67	1.32	1.22	1.24
Log expending	6.63	6.67	−1.05			
N	3,346	1,884		3,346	1,884	

d_g is the (period t cross-sectional) difference-in-means statistic based on robust standard errors, adjusted for clustering at the community level, which has an asymptotic standard normal distribution. Bold statistics denote significance at the 5% level (here, none of the differences are significant). T refers to treatment group and C to control group.

In the sensitivity analysis of Section 3.4, we also consider the ENCEL1999 November survey (round 5). All households in our sample are classified as poor by the PROGRESA administrators, so they potentially qualify for the benefit. Table 8 in the Appendix reports some summary statistics for demographics and parental and village-level characteristics for the pre-PROGRESA period. Our base sample consists of 5,230 households, of which 64% are in the treatment group.

The PROGRESA survey questions relating to food and tobacco expenditures refer to a period of the last 7 days prior to the interview. Questions relating to expenditure on clothing and toys cover the last 6 months. This question design should reduce problems arising from recall error and potential infrequency of some purchases. Table 1 above reports budget shares and monthly household expenditures on goods examined in detail in the next section. The high food budget share is an indicator of the poverty experienced by these children. All nominal quantities are expressed in 1997 prices. One reported imperfect measure of household resources is monthly household income from paid work (this excludes income from self-employment, which we could not estimate in both periods). Table 2 below reports incomes and distinguishes between fathers and mothers.

Table 2 Testing changes in parental labour supply: differences in mean incomes from work

	Pre-PROGRESA			PROGRESA		
	T	C	d_g	T	C	d_g
Household income from work	1,164	1,207	−0.86	979.24	953.18	0.57
Fathers' work income	778.49	841.12	−1.67	651.12	632.12	0.65
Mothers' work income	54.00	38.81	1.44	28.50	17.31	**2.17**

See the footnote of Table 1.

Table 3 The altruism test: differences in mean expenditures in PROGRESA period

	Expenditures				Budget shares [%]			
	T	C	d_g	D_g	T	C	d_g	D_g
Boys' clothing	14.15	10.07	**3.70**	**3.62**	1.95	1.41	**4.43**	**0.40**
Girls' clothing	13.29	9.88	**3.41**	**3.84**	1.38	1.81	**3.54**	**0.34**
Toys	0.27	0.23	0.41	0.13	0.03	0.03	0.24	0.001
Food	535.95	497.02	**2.08**	**36.30**	74.09	73.97	0.13	−0.07
Tabacco	2.34	2.54	−0.36	−0.17	0.36	0.29	1.07	−0.02
Mens' clothing	6.44	6.13	0.41	−0.46	0.77	0.76	0.14	−0.08
Womens' clothing	5.14	5.44	−0.43	−1.09	0.72	0.64	1.08	−0.18

See the footnote of Table 1. D_g is defined in Eq. 1.

Moving from the pre-PROGRESA to the PROGRESA period, the children in our sample apparently have experienced, on average, a deterioration of total household real resources (see Tables 2 and 3): real income from paid work has fallen, as has real total expenditure and the budget share of food have increased (see also Handa et al. 2000). This worsening of resources does not affect our test since we compare mean outcomes for beneficiaries and the control group.

3 Testing for parental altruism

Our test of parental altruism focuses on the income effect of PROGRESA on child-related expenditures. We interpret parental altruism in the standard way as utility interdependence: Altruistic parents care about the welfare of their children and selfish parents do not. In a formal model, the extent of parental altruism could be captured in the structural equation describing parental preferences by a scalar altruism parameter, say η, which weighs the utility of children derived from child-related expenditures.[5] Altruistic parents exhibit $\eta > 0$, whereas selfish parents exhibit $\eta = 0$. This is a strong hypothesis, but arises naturally from interpreting altruism as utility interdependence. Moreover, utility interdependence vs independence is precisely the set of hypotheses which are examined in the well-known literature discussed briefly in Section 1.1, and this captures precisely the conflict between the neo-Marshallian view and the parental self-perception discussed in Section 1 (we return to the interpretation of altruism as utility interdependence below in Section 3.1.2).

Our test statistics for parental altruism are differences-in-differences in child-related expenditures between subgroups of PROGRESA beneficiaries and control households. We only consider households which sent their children to school in the PROGRESA period in order to eliminate the confounding influence of subsistence-constrained parents. For altruistic parents, the

[5]The Working Paper version of this paper proposes such a model.

difference-in-difference is positive; for selfish parents, it is zero. Additional evidence is provided by considering also adult-specific goods such as adult clothing or tobacco and the household public good food (recall that individual food consumption shares are not observable). We also examine adult leisure via its complement adult labour. Throughout, we explicitly distinguish in our estimations for child-related goods between boys and girls and for adult-related goods between men and women.

3.1 The estimating equations

More formally, let $I = 1$ indicate PROGRESA beneficiaries and $I = 0$ control households. $c_{g,\tau}$ refers to expenditure on good g in period t, which is either the pre- or PROGRESA period, $t \in \{\text{pre}, \text{post}\}$. The change in expenditure on good g is denoted by $\Delta c_g = c_{g,\text{post}} - c_{g,\text{pre}}$. The specific goods are child-specific goods (indexed by K), adult-specific goods (A) and household public good food (F), i.e. $g \in \{K, A, F\}$. The non-parametric difference-in-difference is then given, for each good, by

$$D_g \equiv E\{\Delta c_g | I = 1\} - E\{\Delta c_g | I = 0\}. \tag{1}$$

In the pre-PROGRESA period, the cross-sectional difference in child-related expenditures between PROGRESA beneficiaries and control households equals the systematic difference, if any, between the two groups. We label this difference the selection bias SB. In the PROGRESA period, the cross-sectional expenditure difference equals the PROGRESA income effect and the selection bias. If either the selection bias is zero or time invariant, the difference-in-difference non-parametrically identifies the altruism effect. Hence, for altruistic parents, we have $D_K > 0$.

Our extended econometric specification explicitly captures a potential time-invariant selection bias, and also allows for household heterogeneity as follows. Expenditures on good g before and during PROGRESA are assumed to be given by

$$c_{g,\text{pre}} = \alpha_{g,0} + \alpha'_{g,1} X_{\text{pre}} + \varepsilon_{g,1},$$

and

$$c_{g,\text{post}} = \alpha_{g,0} + \tau_g + \alpha'_{g,1} X_{\text{post}} + D_g \times I + \varepsilon_{g,\text{post}}.$$

The variates X_t collect family characteristics, and τ_g represents a growth term (a common parallel trends assumption). The error terms satisfy $E\left\{\varepsilon_{g,t} | X_t\right\} = 0$. Since no benefit is received pre-PROGRESA, I is excluded from the $c_{g,1}$ equation. The selection bias is, if present, assumed to be time-invariant, so $\text{SB}_g \equiv E\left\{c_{g,\text{pre}} | I = 1\right\} - E\left\{c_{g,\text{pre}} | I = 0\right\}$. Adding and subtracting the selection bias in the cross-sectional expenditure equations, we have

$$c_{g,\text{pre}} = \alpha_{g,0} + E\left\{\varepsilon_{g,\text{pre}} | I = 0\right\} + \alpha'_{g,1} X_{\text{pre}} + \text{SB}_g \times I + \theta_{g,\text{pre}} \tag{2}$$

and

$$c_{g,\text{post}} = \alpha_{g,0} + \tau_g + E\left\{\varepsilon_{g,psot} | I = 0\right\} + \alpha'_{g,1} X_{\text{post}} + \text{SB}_g \times I + D_g \times I + \theta_{g,\text{post}} \tag{3}$$

with $E\left\{\theta_{g,t}|I\right\} = 0$. Combining these two equations, we obtain our estimating equation

$$\begin{aligned} c_{g,t} &= c_{g,1} + \left(c_{g,P} - c_{g,1}\right) 1\,(t = P) \\ &= \alpha_{g,0} + E\left\{\varepsilon_{g,1}|I = 0\right\} + \alpha'_{g,1} X_t + \mathrm{SB}_g I \\ &\quad + \tau_g \times 1\,(t = P) + D_g \times I \times 1\,(t = P) + \theta_{g,t}. \end{aligned} \tag{4}$$

Thus, D_g is estimable without bias from the regression of $c_{g,t}$ on X_t, I, $1\,(t = P)$ and $I \times 1\,(t = P)$ as the coefficient of the interaction term $I \times 1\,(t = P)$. The regression constant is not identified, but this is not problematic since we have no inherent interest in this object.

Note that the selection bias is zero if the PROGRESA benefit is randomly assigned in our data. We refer to this as assumption A1 and explicitly test it below in Section 3.2. Our test focuses on parents with children who attend school during PROGRESA. Benefit eligibility is randomised. However, receiving the benefit is conditional on school attendance, school attendance is an endogenous variable and the sample used in the analysis is potentially non-random. This gives rise to the potential selection bias. We address this concern in several ways. First, the altruism test is unaffected if the selection bias is time-invariant. Second, we consider subsamples for which the selection bias is absent. We discuss these issues in greater detail in Section 3.1.1 below. One final concern is the role of parental leisure and its complement parental labour supply. It is conceivable that selfish parents appropriate the PROGRESA benefit while reducing their labour supply. Given the dual nature of this adult good, we examine parental labour supply separately in Section 3.2.

3.1.1 Randomised benefit eligibility, schooling choice and selection bias

Our test focuses on parents with children who attend school during PRO GRESA. Benefit eligibility is randomised (and we verify the empirical validity of this assumption in Section 3.2 below). However, receiving the benefit is conditional on school attendance, school attendance is an endogenous variable and the sample used in the analysis is potentially non-random. A concern then is the presence of a potential selection bias affecting our estimates. We address this concern in two ways.

First, we consider a subgroup of parents with children for whom the PRO GRESA benefit certainly constitutes an exogenous income increase. School enrollment in primary school is almost universal, so the benefit does not affect school attendance. For this group, the benefit thus constitutes an exogenous income increase, which is randomly assigned.

Second, the selection bias, if present, can only be associated with a group of children whose school attendance is affected by PROGRESA: in the absence of the benefit, they would have sent their children to work, but with the benefit, they send their children to school. We refer to these children as 'switchers'. Children who attend school irrespective of the PROGRESA benefit are labelled 'stayers'. In period t, a cross-sectional selection bias is potentially

present if stayers and switchers are systematically different. Since the majority of children who attend school PROGRESA attend school pre-PROGRESA as well, any such difference must be small. Taking the difference-in-differences, the cross-sectional bias terms cancel out provided they are time-invariant. This is very plausible in the current context since pre- and PROGRESA periods are separated by only 1 year, and we have conditioned on school attendance in the pre-PROGRESA period.

We consider this second argument more precisely. First, consider how the schooling indicators before and during PROGRESA $(s_{\text{pre}}, s_{\text{post}})$ partition children. $(0, 0)$ indicates children who do not attend school, and $(1, 0)$ indicates 'drop outs'. By conditioning on $s_P = 1$, these subgroups are dismissed in particular households that are subsistence constrained during PROGRESA. Therefore, the children in our sample are indicated by $(0, 1)$, the 'returnees', or by $(1, 1)$. This group consists of 'stayers' (who attend irrespective of any PRO GRESA benefit) and 'switchers'. Thus, by conditioning on $s_{\text{post}} = 1$, children in the control group are all stayers, and children in the treatment group consists of stayers and switchers. To what extent do switchers and stayers differ? The selection bias at time t is zero if being a switcher is purely determined by random variations, e.g. in child wages (ceteris paribus, switchers have slightly higher wage offers during PROGRESA than stayers, but the wage offer is sufficiently low so that, with the PROGRESA benefit, the children stay at school). The time t selection bias is present if switchers in the treatment group and stayers in the control group are systematically different (in terms of observables or unobservables). For instance, if child wages increase in age, and this is the only source of variation, then switchers are slightly older than stayers. However, since both stayers and switchers attend school in the PROGRESA period, and the pre-PROGRESA and PROGRESA period are separated by only 1 year, any systematic difference is likely to be small. Moreover, given this timing, we also expect the per-period selection bias to be time-invariant, so that by taking the difference-in-differences, the bias terms cancel out.

3.1.2 Utility interdependence and altruism

Modelling parental altruism in the standard way as utility interdependence leads to the rejection of the parental selfishness hypothesis if $D_K > 0$ is observed. However, other models are conceivable, which predict $D_K > 0$ but in which altruism is not the underlying mechanism. In these situations, looking jointly at D_K and D_A can provide circumstantial evidence against non-altruistic interpretations of $D_K > 0$. For instance, consider a world in which parents are selfish and have children because they provide services (along the 'neo-Marshallian' view cited in Section 1). The marginal income effect associated with PROGRESA increases the demand for these services (if they are normal goods). Children are paid by parents not with money but in kind, such as clothing. Alternatively, selfish parents might value the social prestige from the public perception of their well-clothed children. These models also

generate the prediction that $D_K > 0$. However, if parents are selfish, we would also expect to observe $D_A > 0$ (unless the implicit relative weight of child services is implausibly high). Therefore, observing jointly $D_K > 0$ and $D_A = 0$ (including parental leisure), we interpret as strong evidence in favour of parental altruism.

3.2 Testing the validity of assumption 1 and changes in parental labour supply

Our test of parental altruism is based on assumption A1, namely that programme assignment is perfectly random. The quality of the randomisation procedure has been assessed in Behrman and Todd (2000), who conclude that, in almost all dimensions, programme assignment in the entire PROGRESA sample is random.

We turn to testing the empirical validity of this assumption in our specific sample of households with children in school during PROGRESA, by examining several variables: first, we consider item expenditure and budget shares, as well as total expenditure (Table 1), then several income measures (Table 2), and finally household demographics, parental characteristics, and village level characteristics (Appendix Table 8). For each variate, we report its pre-PROGRESA mean for treatment (T) and control (C) groups, and compute a difference-of-means test statistic, denoted by d_g for good g. More precisely, d_g is the difference in means divided by its standard error, and inference for this is adjusted for clustering at the community level. Note that d_g has an asymptotic standard normal distribution.

We turn to specific results. Table 1 considers item expenditures and budget shares in the pre-PROGRESA period. None of the individual expenditure items exhibit statistically different means pre-PROGRESA: d_g is never significant. This conclusion also applies to log total expenditure (row 8).[6] Next, the left panel (labelled pre-PROGRESA) of Table 2 considers several income variables: total household income from work, as well as the separate work incomes for father and mother. Again, we find no significant differences between treatment and control groups. Finally, in the Appendix Table 8, we also consider household demographics, parental characteristics, and village-level characteristics. Again, none of the means are significantly different pre-PROGRESA between treatment and control households. To summarise this evidence, we conclude that assumption A1 is valid for our sample.

Next, we assess the first potential channel of parental selfishness by examining the (robust) difference in mean work incomes during PROGRESA. The results are reported in the right panel (labelled PROGRESA) of Table 2. Measuring labour supply and, thus, indirectly, leisure by work income, we

[6]For completeness, we also examine log total household expenditure in the PROGRESA period. As expected, the introduction of the benefit increases the relative total expenditure of the treatment group: the means are 6.45 for the treatment group and 6.38 for the control group, and d equals 2.18.

reject this channel of parental selfishness: parents in beneficiary households have not reduced their labour supply by appropriating the PROGRESA benefit. Neither total household income from work nor its principal component, the father's income from work is statistically different between treatment and control groups. During PROGRESA, mean income from work for mothers is statistically higher for beneficiary households. This difference, however, is likely to have a negligible effect since the mother's income from work constitutes only around 2% of total work income, and the incidence of wage income is 3% (4%) in control (treatment) households. Anyway, we address the potentially confounding factor (extra spending on child-related goods financed by higher mother's work income) in our robustness analysis in Section 3.4.

3.3 The altruism test: empirical results

We proceed to implement the test for parental altruism. In this section, we confine our attention to data from the PROGRESA round 3 and focus on the entire sample of children at school in this period; in Section 3.4, we consider specific subsamples, in particular, children in primary school (for which no selection bias is present in any period), and we repeat the experiments with PROGRESA round 5 data. The child-related expenditures we focus on are clothing for boys and girls and toys. We also examine changes in expenditure on the household public good food and the parental goods tobacco and adult clothing. We explicitly distinguish between boys and girls since the benefit is slightly larger for girls than for boys, and parental altruism might be asymmetric.

Our principal set of tests are tests based on the difference-in-differences D_g for good g, given in Eq. 1, which is free from any selection bias under the stated assumptions. For greater transparency, we report the per-period difference in mean expenditures between PROGRESA beneficiaries and households in the control group separately, before subtracting the per-period differences. In Table 1 above, we have reported the differences in means in the pre-PROGRESA period. In Table 3, we repeat the exercise for the PROGRESA period. As before, d_g is the clustering adjusted robust difference in mean test statistics associated with good g. Pre-PROGRESA, none of the differences were significant. Thus, if D_g turns out to be significant, this significance can only stem from the causal effect of PROGRESA arising in the PROGRESA period. Boys' and girls' clothing expenditures and budget shares are significantly higher in the PROGRESA period for PROGRESA beneficiaries, as is the expenditure on food. This significance translates into a significant difference-in-difference. Poor households spend, not surprisingly, very little on toys; hence, we ignore this expenditure item in the subsequent analysis. Importantly, PROGRESA beneficiaries and control households do not differ in their expenditure on adult goods, nor in terms of parental leisure as shown in Table 2 above.

We interpret this as strong evidence in favour of the parental altruism hypothesis and, thus, against parental selfishness. In the light of our discussion

Table 4 The altruism test: estimates of D_g in linear model with household heterogeneity and a non-paramteric control function to account for any selection bias

	Expenditure (robust SE)	Budget shares (robust SE)
Boys' clothing	**4.02**	**0.46**
	(1.60)	(0.15)
Girls' clothing	**4.08**	**0.33**
	(1.67)	(0.14)
Toys	0.17	−0.01
	(0.33)	(0.03)
Food	**43.08**	−0.05
	(22.00)	(0.82)
Tobacco	−0.11	−0.012
	(0.77)	(0.09)
Men's clothing	−0.78	−0.12
	(0.60)	(0.11)
Women's clothing	−0.92	−0.18
	(1.14)	(0.09)

The estimating equation is given by Eq. 4. Robust standard errors. Bold coefficients are significant at the 5% level. Budget shares regression: coefficients and robust SE $\times 100$.

of alternative explanations of utility interdependence, we reject service-for-transfers alternatives: the benefit has no impact on adult-related goods, but selfish parents would be expected to exhibit $D_A > 0$. Moreover, the benefit effects on boys and girls are similar, so parental altruism is 'symmetric'.

Next, we estimate Eq. 4 in order to explore whether a selection bias is present and the effects of household heterogeneity. The selection bias arises if switchers are systematically different from stayers. Our regressors include controls for the demographic structure of the household, the educational background of the parents and controls for the village.[7] It turns out that none of the coefficients of I are significant (not reported here explicitly), so we conclude that the selection biases are absent. In Table 4, we report, for the sake of brevity, the estimates of D_g.[8] Table 4 reproduces the pattern of Table 3 and the point estimates are very similar. We thus proceed to use this specification in all our subsequent experiments.

3.4 Sensitivity tests

We conduct a variety of sensitivity tests in order to check the robustness of our results. First, we consider specific subsamples, then we use data from PROGRESA round 5 instead of round 3. Given that only one additional year has elapsed, it is likely that the parallel trends assumption underlying the difference-in-difference approach is still valid.

[7]Specifically, the vector of controls includes the following: the demographic controls are household size, the number of girls and boys and the number of children at school; the parental controls are the ages of the father and mother pre-PROGRESA, individual indicators for whether the parents are literate, and the age the father entered the labour market; the village controls are indicators for the presence of a pre-school, a primary school and a secondary school. See Appendix Table 8 for descriptive statistics.

[8]Reports of all the coefficients are available on request.

First, we have argued above that for households with children in primary school, the PROGRESA benefit constitutes a randomly assigned exogenous income increase since enrollment is nearly universal. There is no selection bias by construction. Panel 1 of Table 5 reports the results. The point estimates are very similar to the earlier results.

Next, we address a potentially confounding factor associated with mothers' labour supply. Our preliminary sample analysis reported in Table 2 has revealed that, whilst mothers' work income is no different between treatment and control groups pre-PROGRESA, the difference is statistically significant during PROGRESA. In particular, work income has fallen for both groups, but women in beneficiary households earn more (selfish mothers would be expected to reduce their labour supply). A concern is that the observed difference in child-related expenditure could be financed from comparatively higher mothers' labour supply. To investigate this issue, we restrict the sample to households in which women do not earn a wage both pre- and during PROGRESA (the majority of households). The results are reported in panel 2 of Table 5. We find that this concern can be dismissed, as the point estimates, compared to Table 4, hardly change.

Third, we focus on the costs associated with meeting the eligibility conditions. Our difference-in-difference approach controls for necessary school-related expenditures. However, when switching children from child labour to schooling, some households might incur higher than necessary costs on child clothing and uniforms. This potentially confounding factor is eliminated by considering only households who send their children to school both pre- and during PROGRESA. Panel 3 of Table 5 reports the results. Again, the point estimates are very similar to our earlier results. In summary, we conclude that our results based on data for PROGRESA round 3 are robust.

We now turn to PROGRESA round 5 data. We have repeated all the experiments. Panel 1 of Table 6 is directly comparable to Table 4, and all remaining panels have their analogues in Table 5. Overall, the pattern of Table 4 is reproduced in all experiments: Boys' and girls' clothing increase significantly for PROGRESA beneficiaries, as well as food expenditure; adult-related goods are unaffected. Boys' and girls' clothing are affected similarly so that altruism is symmetric. Compared to the results based on round 3 data, the point estimates have increased slightly for boys' and girls' clothing and have doubled for food expenditure.

In our final set of experiments, we reconsider the parental good leisure which we examine, as before, via its complement labour supply. We consider fathers and mothers separately. Selfish parents could be expected to reduce their labour supply by appropriating the PROGRESA benefit. In Table 2, we have considered the per-period differences in parental labour supply between treatment and control households. We now consider directly the differences-in-differences. Panel 1 of Table 7 reports the results. The difference-in-difference estimates lead to the inference that parental labour supply and, thus, parental leisure has not changed. As in the previous experiments, we proceed

Table 5 Further sensitivity tests: specific sub-samples for PROGRESA round 3 data

	1. Households with children in primary school		2. Non-working mothers		3. Households with children in school	
	Expenditure	Budget shares	Expenditure	Budget shares	Expenditure	Budget shares
Boys' clothing	**4.55**	**0.46**	**3.91**	**0.43**	**3.47**	**0.37**
	(1.79)	**(0.17)**	**(1.65)**	**(0.15)**	**(1.64)**	**(0.15)**
Girls' clothing	**3.58**	**0.34**	**4.07**	**0.30**	**4.03**	**0.34**
	(1.68)	**(0.15)**	**(1.78)**	**(0.15)**	**(1.76)**	**(0.14)**
Toys	−.100	−0.02	0.31	0.02	0.26	0.01
	(0.42)	(0.03)	(0.34)	(0.03)	(0.34)	(0.03)
Food	**43.61**	0.39	**43.46**	−0.098	**40.33**	0.32
	(24.80)	(0.97)	**(22.29)**	(0.084)	**(22.20)**	(0.84)
Tobacco	−0.35	0.004	−0.28	−0.02	0.26	0.02
	(1.03)	(0.12)	(0.80)	(0.09)	(0.77)	(0.09)
Men's clothing	−0.56	−0.14	−0.58	−0.10	−0.76	−0.12
	(1.66)	(0.13)	(1.35)	(0.10)	(1.40)	(0.12)
Women's clothing	−0.27	−0.16	−0.69	−0.17	−0.87	−0.17
	(1.53)	(0.11)	(1.21)	(0.09)	(1.20)	(0.09)
Number of households	2,966		4,604		4,520	

See the footnote of Table 4.

Table 6 Robustness checks II for PROGRESA round 5 data

	1. Pre-PROGRESA and PROGRESA		2. Households with children in primary school		3. Non-working mothers		4. Households with children in school	
	Expenditure	Budget shares	Expenditure	Budget shares	Expenditure	Budget shares	Expenditure	Budget shares
Boys' clothing	**5.44**	**0.43**	**6.19**	**0.56**	**5.75**	**0.53**	**5.15**	**0.41**
	(1.69)	**(0.19)**	**(2.01)**	**(0.23)**	**(1.76)**	**(0.18)**	**(1.67)**	**(0.18)**
Girls' clothing	**7.14**	**0.62**	**5.57**	**0.50**	**7.60**	**0.69**	**7.10**	**0.64**
	(1.77)	**(0.17)**	**(1.94)**	**(0.22)**	**(1.92)**	**(0.18)**	**(1.81)**	**(0.17)**
Toys	−0.04	−0.01	−0.34	−0.04	0.09	−0.002	−0.01	−0.01
	(0.32)	(0.03)	(0.37)	(0.03)	(0.33)	(0.02)	(0.33)	(0.03)
Food	**88.95**	0.80	**79.35**	0.57	**78.90**	0.24	**85.18**	0.88
	(23.21)	(0.92)	**(28.74)**	(1.10)	**(23.60)**	(0.95)	**(23.60)**	0.92
Tobacco	−0.65	−0.06	−1.22	−0.11	−0.96	−0.08	−0.61	−0.06
	(0.93)	(0.11)	(1.27)	(0.15)	(0.99)	(0.11)	(0.95)	(0.11)
Men's clothing	0.90	0.07	1.43	−0.01	1.08	0.10	0.99	0.08
	(1.21)	(0.13)	(1.57)	(0.16)	(1.27)	(0.14)	(1.24)	(0.13)
Women's clothing	1.26	0.07	2.57	0.04	1.47	0.09	1.40	0.09
	(1.22)	0.10	(1.91)	(0.14)	(1.34)	(0.11)	(1.27)	(0.10)
Number of households	4,487		2,157		3,951		4,296	

See the footnotes of Tables 4 and 5.

Table 7 Parental leisure

	PROGRESA Round 3			PROGRESA Round 5		
	1. Pre-PROGRESA and PROGRESA	2. Households with children in school	3. Households with children in primary school	4. Pre- PROGRESA and PROGRESA	5. Households with children in school	6. Households with children in primary school
Father's work income	67.77	66.09	79.81	81.80	84.59	**153.35**
	(36.28)	(37.89)	(43.33)	(45.94)	(47.35)	**(56.02)**
Mother's work income	−3.83	−4.78	8.03	6.57	8.35	8.39
	(8.96)	(9.18)	(11.40)	(13.16)	(13.32)	(18.05)
Number of households	4984	4520	2966	4487	4296	2157

See the footnotes of Tables 4 and 5.

to exclude potentially confounding factors by restricting the sample. We close down the potential channel of substitution between adult and child labour, and expenditures linked to meeting eligibility conditions, by considering only households with children in school pre- and during PROGRESA. We then close down further substitution effects by restricting the sample further to households with children only in primary school, so the PROGRESA benefit induces only income effects. The results are reported in panels 2 and 3 of Table 7. In neither experiment does parental labour supply fall (so this channel of selfishness is rejected). In fact, all the differences are insignificant. We repeat these experiments with PROGRESA round 5 data. The results are reported in panels 4–6 in Table 7. The results repeat the previous pattern. The only exception is panel 6. By construction, these households only experience income effects. It is rather surprising then that fathers in beneficiary households work *more* relative to the control group.

To summarise the results of our sensitivity analyses, we conclude that our altruism test results remain robust across all subsamples and periods considered. In particular, the pattern of Table 4 is reproduced throughout and parents have not increased their consumption of leisure. Moreover, the coefficients for boys' and girls' clothing remain similar in all experiments.

4 Conclusions

We have developed a new test of parental altruism, based on experimental data from the Mexican anti-poverty programme PROGRESA. Exploiting the randomised programme assignment enabled us to propose a simple test. Specifically, we find significant and symmetric positive effects on expenditure on boys' and girls' clothing. The effects on adults' clothing, tobacco, and parental leisure are insignificant. These results are robust across all the sensitivity tests, varying subsamples and periods, we have conducted. In particular, we have eliminated potentially confounding effects associated with benefit conditionality and substitution effects and have addressed the issue of selection bias.

We have interpreted these findings as strong evidence for parental altruism and against parental selfishness. The empirical results make service-for-transfer alternatives implausible. This empirical evidence thus conforms to the self-perception of the PROGRESA beneficiary quoted in Section 1 – 'it [the monetary benefit] is always for them [the children]' – and contradicts the 'neo-Marshallian' view. Our empirical finding is an important contribution to the continuing debate on parental altruism (in particular given the often negative results in the empirical literature) and has focused on poor non-urban Mexican households. Moreover, we have also shown that, at least for these households, cash benefits intended for children actually reach them, so that the design of transfer and aid programmes does not have to be rethought from the perspective of parental altruism.

Acknowledgements We are grateful for comments from Deborah Cobb-Clark and our referees, which helped to improve the paper, and several seminar participants.

Appendix

Table 8.

Table 8 Some summary statistics for the Pre-PROGRESA period

	Treatment	Control	d_g
Household level variables			
Household size	7.12	7.16	0.45
	0.05	*0.07*	
Number of boys	2.27	2.26	0.14
	0.03	*0.03*	
Number of girls	2.14	2.18	0.72
	0.03	*0.04*	
Number of children at primary schooling	2.48	2.50	0.43
	0.03	*0.05*	
Number of children at secondary schooling	0.31	0.32	0.02
	0.02	*0.02*	
Number of boys at primary schooling	1.27	1.27	0.08
	0.02	*0.03*	
Number of boys at secondary schooling	0.18	0.17	0.92
	0.01	*0.01*	
Number of girls at primary schooling	1.21	1.23	0.56
	0.02	*0.03*	
Number of girls at secondary schooling	0.13	0.15	0.97
	0.015	*0.01*	
Parental variables			
Father's age	40.90	40.94	0.12
	0.219	*0.27*	
Father's age when first worked	13.02	12.92	0.70
	0.08	*0.11*	
Father literate	0.76	0.76	0.11
	0.01	*0.02*	
Mother's age	36.39	36.23	0.65
	0.15	*0.20*	
Mother literate	0.65	0.64	0.35
	0.02	*0.02*	
Village level variables:			
Has pre-school	0.91	0.90	0.18
	0.02	*0.03*	
Has primary school	0.99	0.98	0.82
	0.01	*0.01*	
Has secondary school	0.27	0.29	0.21
	0.03	*0.04*	
Male agricultural wage	26.93	27.55	0.60
	0.64	*0.83*	
Number of households	3,346	1,884	

See the footnote of Table 1.

References

Altonji JG, Hayashi F, Kotlikoff LJ (1992) Is the extended family altruistically linked? Direct test using micro data. Am Econ Rev 82:1177–1198

Altonji JG, Hayashi F, Kotlikoff LJ (1997) Parental altruism and inter vivos transfers: theory and evidence. J Polit Econ 105:1125–1166

Arrondel L, Masson A (2003) Altruism, exchange or indirect reciprocity: what do the data on family transfers show? In: Gérard-Varet LA, Mercier-Ythier J, Kolm S (eds) Economics of giving, reciprocity and altruism. North-Holland Elsiever, Amsterdam

Attanasio O, Lechene V (2002) Tests of income pooling in household decisions. Rev Econ Dyn 5:720–748

Becker GS (1981a) Altruism in the family and selfishness in the market place. Economica 48(189):1–15

Becker GS (1981b) A treatise on the family. Harvard University Press, Cambridge

Becker GS, Barro RJ (1988) A reformulation of the economic theory of fertility. Q J Econ 103(1):1–25

Behrman J, Todd P (2000) Randomness in the experimental samples of PROGRESA. Mimeo, IFPRI, Washington, DC

Bergstrom T (1997) A survey of theories of the family. In: Rosenzweig M, Stark O (eds) Handbook of population and family economics. North-Holland, Amsterdam

Bhalotra S (2001) On selfish parents and child labour. Working paper, University of Bristol

Caldwell JC (1978) A theory of fertility: from high plateau to destabilization. Popul Dev Rev 4(4):553–577

Cox D (1987) Motives for private income transfers. J Polit Econ 95(3):508–546

Cox D, Eser Z, Jimenez E (1998) Motives for private transfers over the life cycle: an analytical framework and evidence for Peru. J Dev Econ 55:57–80

Handa S, Huerta M-C, Perez R, Straffon B (2000) Poverty inequality and spillover in Mexicos education. Health and nutrition program. International Food Policy Research Institute, Washington, DC

Hayashi F (1995) Is the Japanese extended family altruistically linked? A test based on Engel curves. J Polit Econ 103:661–674

Hoddinott J, Haddad L (1995) Does female income share influence household expenditures? Evidence from Cote d'Ivoire. Oxf Bull Econ Stat 57:77–95

Lundberg SJ, Pollak RA, Wales TJ (1997) Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. J Hum Resour 32(3):463–480

Ravallion M, Wodon Q (2000) Does child labor displace schooling? Econ J 110:C158–C175

Rubalcava L, Teruel G, Thomas D (2002) Weflare design, women's empowerment, and income pooling. CIDE Working Paper

Schultz TP (2004) School subsidies for the poor: evaluating the Mexican PROGRESA poverty program. J Dev Econ 74:199–250

Skoufias E (2000) Is PROGRESA working? Summary of the results of an evaluation by IFPRI. IFPRI, Washington, DC

Skoufias E (2001) PROGRESA and its impacts on the human capital and welfare of households in rural Mexico: a synthesis of the results of an evaluation by IFPRI. Mimeo, IFPRI, Washington, DC

Skoufias E, Parker SW (2001) Conditional cash transfers and their impact on child work and schooling: evidence form the PROGRESA program in Mexico. IFPRI, FCND no 123 Washington, DC