

Albarran, Pedro; Attanasio, Orazio P.

Working Paper

Do public transfers crowd out private transfers? Evidence from a randomized experiment in Mexico

WIDER Discussion Paper, No. 2002/06

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Albarran, Pedro; Attanasio, Orazio P. (2002) : Do public transfers crowd out private transfers? Evidence from a randomized experiment in Mexico, WIDER Discussion Paper, No. 2002/06, ISBN 9291901318, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki

This Version is available at:

<https://hdl.handle.net/10419/52739>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



United Nations
University

WIDER

World Institute for Development Economics Research

Discussion Paper No. 2002/6

Do Public Transfers Crowd Out Private Transfers?

Evidence from a Randomized Experiment
in Mexico

Pedro Albarran¹
and Orazio P. Attanasio²

January 2002

Abstract

The paper will use the data from the PROGRESA evaluation to show the extent to which such a programme crowds out private transfers. An interesting aspect of the paper is the fact that it can use the randomization (one-third of the villages in the sample were randomized in the programme for evaluation purposes).

Keywords: private transfers, crowding-out, public programmes

JEL classification: O16, I30

Copyright © UNU/WIDER 2002

¹CEMFI, Spain; ²University College London, United Kingdom

This study has been prepared within the UNU/WIDER project on Insurance Against Poverty, which is directed by Dr Stefan Dercon.

UNU/WIDER gratefully acknowledges the financial contribution to the project by the Ministry for Foreign Affairs in Finland.

Acknowledgements

We would like to thank Patricia Muniz and Ana Santiago for many conversations about the data and Susan Parker for some useful suggestions. This paper would not have been possible without the inspiration of Jose Gomez de Leon who first stimulated our interest in the PROGRESA programme.

UNU World Institute for Development Economics Research (UNU/WIDER) was established by the United Nations University as its first research and training centre and started work in Helsinki, Finland in 1985. The purpose of the Institute is to undertake applied research and policy analysis on structural changes affecting the developing and transitional economies, to provide a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and to promote capacity strengthening and training in the field of economic and social policy making. Its work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

UNU World Institute for Development Economics Research (UNU/WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by Jaana Kallioinen at UNU/WIDER
Printed at UNU/WIDER, Helsinki

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

ISSN 1609-5774
ISBN 92-9190-130-X (printed publication)
ISBN 92-9190-131-8 (internet publication)

1 Introduction

An important issue that is often neglected in the design of programmes targeted to the poor in developing countries, is that of the interaction of the proposed programme with existing private arrangements. This issue is relevant for poverty alleviation programmes, for insurance programmes, for education subsidies and virtually any intervention and has implications not only for the private sector activities that are close substitute for the programme intervention, but also for other aspects. The neglect of these effects is lamentable, as a proper cost benefit analysis, and the consideration of alternative programmes should recognize that such interventions do not happen in a vacuum. The ‘side’ effects of a public intervention programme do not need to be negative: while it is possible that some private activities might be ‘crowd out’ by the intervention it is possible that others (and possibly unrelated ones) would grow as a consequence of it.

Many studies have considered whether the introduction of public transfers affects private transfers among the households targeted by the public scheme. From a theoretical point of view, as clearly stated by Cox (1987), there are several reasons to expect an effect of public transfers on private ones. As we discuss below, several models predict a negative relationship between public and private transfers. There are also models that predict a negative effects of the introduction of insurance schemes on private transfers: the intuition here is that the latter play an insurance role that can be crowded out, for a variety of reasons, by the introduction of public insurance.

From an empirical point of view, the challenge to test for the presence and the empirical relevance of these effects is that the recipients of many of the programmes on which data are available are not a random sample of the population. It is therefore very tricky to compare recipients and non-recipients to identify the effect of the programme.

Our paper is a contribution to the analysis of the interaction of public programmes and private transfers. We analyse a large welfare programme in rural Mexico: PROGRESA. As we discuss in detail below, PROGRESA is a programme aimed at fostering the accumulation of human capital by increasing school enrolment, improving nutrition and health practices. It should be stressed that the PROGRESA is not a pure transfer programme. The subsidies poor households receive are not unconditional, but depend on a number of other choices. Indeed, the programme is better described as a change in the relative prices of education and health services rather than a pure income subsidy. However, the programme has a sizeable unconditional component and is, overall, perceived as having a positive income effect. These considerations, however, should be kept in mind when evaluating and interpreting our findings.

Even more importantly, our findings should not be interpreted as an evaluation of the programme: PROGRESA’s main objective is the improvement of the process of human capital accumulations among poor families in rural Mexico: in what follows we do not evaluate how successful PROGRESA was in achieving its stated goals but only whether the programme has affected other aspects of the life of the households living in the villages where the programme is affected. Neither we consider the effect that the programme might have had on other aspects of these people’s lives.

From an empirical point of view, the PROGRESA data set has a big advantage: the evaluation sample we use has a randomized component, which introduces genuinely exogenous variation that can be used to identify the parameters of interest.

2 A theoretical framework

Whether the presence of public transfers affects private transfers has been the topic of a large number of studies. As pointed out by Cox (1987), there are several reasons why public transfers could affect private ones. Some models, initiated by Becker (1974), explain transfers with altruism. Other models, such as Bernheim et al. (1985), instead, appeal to intra-households informal exchanges whereby a transfer is motivated by the provision of some kind of service that the recipient performs for the donor. Finally, it is possible, that transfers are part of insurance schemes by which participants share idiosyncratic risk. All these models predict that the occurrence and size of private transfers are affected by the presence of public transfers that change the income of the recipient or of the donor. However, different models have different implications and stress different mechanisms for the interaction of private and public transfers.

The implications of altruism are quite unambiguous: if public transfers are directed to the recipient of a transfer, both the probability of observing a transfer and its size conditional on occurrence are likely to decrease. The opposite is true for a public transfer targeted to the net donor of a private transfer.

If, instead, private transfers occur as a payment for some sort of service, it can be shown that an increase in the income of the recipient, could give rise either to a decrease or to an increase in the amount transferred (see Cox 1987). In particular, it is possible that transfers are the outcome of a Nash bargaining mechanism involving both sides of the transfer and possibly some services given from the recipient to the donor. In such a situation, it is possible that the relationship between the amount of transfers, conditional on a positive transfer, and the recipient's income is non-monotonic. Cox et al. (1998), for instance, show the results of some simple simulations of a parent-child model where at low levels of income, an increase in the children's income raises her threat point and therefore, raises the transfer. Past a certain level of income, however, further increases will cause the effect to become negative.

Private transfers can also be part of insurance schemes through which households are linked in order to share idiosyncratic risk. If this is the case, both under perfect risk sharing and under imperfect risk sharing (due to the presence of imperfect enforceability or asymmetric information), public transfers will be partly undone by private transactions, even though the mechanisms at play might be slightly different and depends on whether the government transfer targets a segment of the population and whether they are contingent on some observable state (such as is the case for government insurance). Under perfect risk sharing, which effectively implies income pooling, targeted public transfers will enter the pool of resources available to the agents sharing risks and, therefore, are likely to be shared. They are also bound to affect the size and direction of private transfers. If they do not change the Pareto weights in the social planner problem (or to be more concrete the particular competitive equilibrium that is selected) and if they are targeted towards the net receivers of private transfers, they are likely to reduce them. In models with imperfect information, if the public transfers are observed and the targeting is done on the basis of observable variables, private transfers will be reduced via essentially the same mechanism at play under perfect risk sharing.

Under imperfect enforceability, on the other hand, the introduction of a public unconditional transfer that will move agents away from situations where the marginal utility of consumption is very high, will induce a reduction in the amount of equilibrium risk sharing and, therefore, the amount of private transfers that happen in equilibrium for any given income shock. These effects are discussed at length and illustrated with numerical simulations by Attanasio and Rios-Rull (2000a, b). In general the properties of this model depend on the specific utility function and the features of the economic environment. With power utility, one gets that the introduction of non-contingent public transfers, of the kind studied in the empirical application below, reduce the amount of private transfers.

Many studies have looked at the relationship between private transfers and recipient income. In particular, Cox (1987) and his co-authors (Cox et al. 1997, 1998) have tested in a variety of contexts both the presence of crowding out and alternative models of transfers. In most situations, Cox and his co-authors find that while the likelihood of receiving a transfer is negatively related to income, as predicted both by altruistic and exchange models, more often than not, the relationship between transfers conditional on a positive transfer and income is non-monotonic, being positive at low levels of income and negative at larger levels. That is transfers first increase with income and then decrease. The fact that transfers increase as a function of income is consistent with the exchange model where the transfer is provided in exchange of a service: as the income of the recipient of the transfer increases, this is likely to increase the implicit price of the recipient's 'services' (see Cox 1987, for example).

Schoeni (1997), has looked at transfers in the US PSID to find that, while poorer households are more likely to receive transfers both in money and in time, the altruistic model does not fully explain the patterns observed in the data.

More recently, Jensen (1999) has looked directly at the effects on migrant remittances of the introduction of a public pension scheme in South Africa. Jensen's is one of the few studies that looks at the crowding out effect directly.

Foster and Rosenzweig (2001) consider the altruistic motive in a model in which transfers across individuals constitute the only mechanisms for sharing idiosyncratic risk in a situation in which borrowing is limited and contracts cannot be enforced perfectly. Foster and Rosenzweig (2001) test one of the implications of a model with imperfect enforceability, namely that, conditional on the current shock an individual receives, current transfers are negative related to the cumulative level of past transfers. They find support for this hypothesis in data from Bangladesh and Pakistan.

Most studies of the crowding out effect suffer from an important endogeneity problem: public policy programmes are typically targeted towards households that are in particular need of transfers. It is therefore difficult to identify the effect that public transfer programmes have on private transfers and in particular to assess what the level of private transfer would have been in the absence of a given programme comparing beneficiaries to non-beneficiaries of the programme. In this respect our study is unique as it exploits the randomization component in our data: a set of villages in our sample were randomly excluded from the programme we are studying for two years. This allows us a direct evaluation of the effect of the programme on private transfers. We now turn to a description of the data set we use and of the programme it refers to.

3 The PROGRESA programme and the evaluation data set

3.1 The PROGRESA programme

In 1997 the Mexican government decided to start a new and large welfare programme targeted to rural Mexico. The programme had three components: health, nutrition and education. The health component consists of a number of initiatives aimed at improving information about vaccination, nutrition, contraception and hygiene and of a programme of visits for children and women to health centres. Participation into the health component is a pre-condition for participating into the nutrition component that, in addition to a basic monetary subsidy received by all beneficiary households, gives some in kind transfers to households with very young infants and pregnant women. The largest component of the programme is the education one. Beneficiary households with school age children receive grants conditional on school attendance. The size of the grant increases with the grade and, for secondary education, is slightly higher for girls than for boys. Finally, all the transfers are received by the mother in the household.

The programme first targeted the poorest communities in rural Mexico. Roughly speaking, the two criteria communities had to satisfy to qualify for the programme were a certain degree of poverty (as measured by what is called an 'index of marginalization') and access to certain basic structures (schools and health centres). Once a locality qualifies, individual households could qualify or not for the programme, depending on a single indicator that is affected by a number of poverty variables (income, house type and so on). Eligibility was determined in two steps. First, a general census of the Progresia localities measured the variables needed to compute the indicator and each households was defined as 'poor' or 'not-poor' (where 'poor' is equivalent to eligibility). Subsequently, in March 1998, an additional survey was carried out and some households were added to the list of beneficiaries. This second set of households is called 'densificados'.

The programme was phased in slowly and is currently very large: at the end of 1999 its budget was US\$ 777 m and was implemented in more than 50,000 localities. At that time, about 2.6 million households, or 40 per cent of all rural families and one ninth of all households in Mexico, were included in the programme. The cost of the programme is about 0.2 per cent of Mexican GDP. The programme has received a considerable amount of attention and publicity and similar programmes are currently being implemented in Honduras, Nicaragua and Argentina. (See IFPRI 2000 for additional details on the programme and its evaluation).

The programme represents a substantial help for the beneficiaries. The nutritional component was 100 pesos per month (or US\$ 10) in the second semester of 1998, which corresponds to 8 per cent of the beneficiaries' income in the evaluation sample. The education component, which is conditional to school enrolment of children between third and ninth grade, can be added up to a maximum total amount of 625 pesos per month (or US\$ 62.5) or 52 per cent of the beneficiaries' income. The average grant in the sample we use was 348 pesos per month for households with children and 250 for all beneficiaries or 21 per cent of the beneficiaries' income. In addition to the (bi) monthly payments, beneficiaries with children in school age receive a small annual grant for school supplies.

It should be stressed that the education component, which is the largest component in the programme, is not a pure transfer, as it is conditional on school attendance. It is therefore better described as a change in the relative price of education. The education component, however, is likely to have a substantial wealth effect. For instance, according to IFPRI (2000) there is not much evidence of a decrease in beneficiaries' income as a consequence of a reduction of child labour. This wealth effect and the presence of an unconditional component (the nutritional supplement) justify our attempt to identify the crowding out of private transfers.

The agency running the programme used the fact that, for logistic reasons, the programme could not be started everywhere simultaneously, to start an evaluation sample. Among the beneficiaries localities, 506 were chosen randomly and included in the evaluation sample. Among these, 320 randomly chosen were assigned to the communities where the programme started early, while 186 were assigned to the communities where the programme started almost two years later (December 1999 rather than May 1998). An extensive survey was carried out in the evaluation sample: after the initial data collection between the end of 1997 and the beginning of 1998, an additional four instruments were collected in November 1998, March 1999, November 1999 and April 2000. Within each village in the evaluation sample, the survey covers all the households and collects extensive information on consumption, income, labour supply, school enrolment, transfers, and a variety of other issues. While each instrument contained a core questionnaire, they differed in that some of them also contained some additional modules. There are also some minor differences in the way some questions are formulated across waves.

The randomisation through which the villages were assigned to the treatment and control groups seems well executed: Behrman and Todd (1999) present evidence in this respect. In particular, most variables seem to be not statistically different between the treatment and control villages. As one should expect to get around 5 per cent of rejections, we feel comfortable in using explicitly the randomization to identify the crowding out of private transfers.¹

3.2 The data

The data set we use is the wave collected in October/November 1998 and contains information on 25,846 households. After loosing, for various reasons, some observations, we are left with 23,268 households.²

As we mention in section 3.1, not all households within a village qualify for Progresa and eligibility was determined in two steps. Fortunately, we observe eligibility status (in both steps) both in the treatment and in the control variables. We therefore have four

¹ For studies that consider the effect of the programme on enrollment, it is a bit worrying that one of the few variables that appears to be significantly different between control and treatment villages is pre-programme school enrollment.

² We discarded 2,361 households who did not fully answer the questionnaire, so that relevant information (on transfers, consumption, household characteristics, etc.) was missing; additionally, we found and dropped 223 households having more than one household head and/or household head's spouse: these are multiple families living in a single household.

groups of households: poor and non-poor living in treatment villages and poor and non-poor living in control villages. Here and in the econometric application we do not distinguish between households that gained eligibility in the first and second round, that is we do not distinguish between ‘densificado’ poors and households that were designed as poor in the first survey. In Table 1, we show how many of the households included in the sample live in treatment (PROGRESA) and control villages and how they split between ‘poor’ and ‘non-poor’, where ‘poor’ is a household that was defined as such by PROGRESA and therefore entitled to the programme.

Unfortunately, matters are further complicated by the fact that not all poor in the treatment villages receive the programme. It is not clear why this was the case. It was probably due to a combination of administrative delays, non-compliance with the programme and so on. In Table 1, therefore, as in the econometric specifications we estimate below, we divide eligible (or poor) households between those who actually receive the programme and those who do not.³

Figure 1 summarizes the structure of our evaluation sample: as it should be clear from the final branches of the tree, there are a total of 5 groups to consider: in control villages there are eligible and non eligible households (which we call poor and non-poor); in the treatment villages there are the same two groups, but the eligible are divided into beneficiaries and non-beneficiaries. In what follows, for consistency with the evaluation literature, we will be occasionally referring to this last group as ‘non-compliers’ with a slight abuse of language. From an econometric point of view, the existence of such a group is a problem because we cannot identify ‘non-compliers’ in control villages.

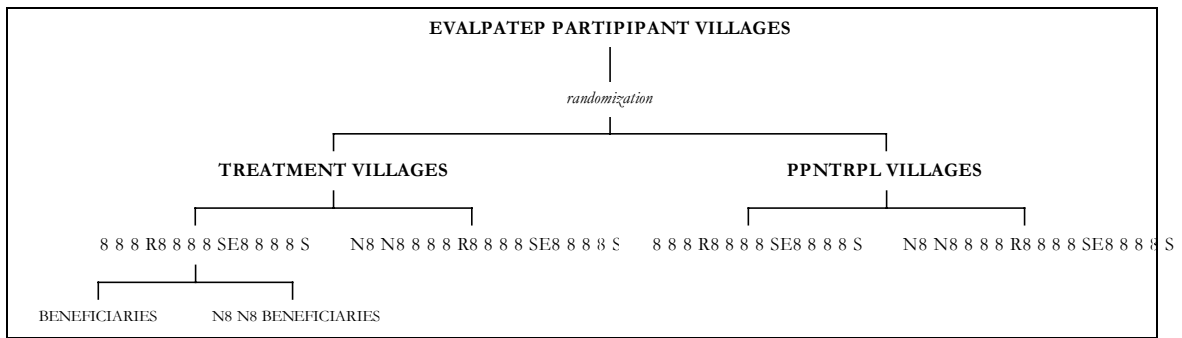
Table 1
Distribution of households, October 1998

		Treatment 320 villages	Control 186 villages	All 506 villages
Non-poor		3,176 (22.12)	2,011 (22.57)	5,187 (22.29)
Poor	Benef.	8,304 (57.83)	6,898 (77.43)	18,081 (77.71)
	Non-benef.	2,879 (20.05)		
		14,359 (100)	8,909 (100)	23,268 (100)

Note Percentage of households in parentheses.

³ Most of these poor non-beneficiaries are ‘densificado’ households (so, they might have not fully incorporated to the programme yet); see Table A1 in the Appendix, where we extend Table 1 to distinguish the ‘densificado’ poor.

Figure 1



That is, in control villages we do not know which of the ‘poor’ households would fail to become beneficiaries were the programme implemented.

If the ‘non-compliers’ were a random subset of the poor households, the problem would not be too serious. Unfortunately, this is not the case. It turns out that, among poor households in treatment villages, ‘non-compliers’ are significantly (but not perfectly!) predicted by several observable variables. For instance, ‘densificado’ households, or households with no children are much more likely to be non-compliers. This prevents a direct comparison of poor, beneficiary households in the treatment villages with the poor households in the control villages.

The questionnaire in the November 1998 wave included two sections on private transfers received by each household member. In the first, the respondent was asked about monetary and non-monetary transfers obtained during the month preceding the interview from friends, neighbours and relatives living either in or outside the village. The second section, instead, included information about support provided during the last six months by former family members who left the household within the last five years (‘migrants’). In both sections, we know if each transfer is in money or in kind (food, clothes, etc.), but the latter are not given a monetary value. In what follows, we aggregate this individual information, thus focusing on households as our unit of analysis. We build two types of household variables on transfers: an indicator that takes the value one if any family member has received some transfers, and, for monetary transfers, the total amount of transfers (by adding up individual transfers received by each member). We build such variables separately for transfers from friends and transfers from migrants. We also convert migrant’s monetary transfers (reported for the last six months) into a monthly basis, so as to be comparable with monetary transfers from friends.

Table 2 shows the relative importance of these inter-household transfers in our sample by computing the percentage of households receiving some transfers for each of the five groups we consider. We found remarkably few households receiving private transfers, as compared to other data sets (see, for instance, Foster and Rosenzweig 2001). The proportion of households receiving transfers from friends and relatives is particularly low. This might be due to the fact that the question focuses on the month preceding the interview. Notice that the proportion of ‘non-poor’ households in treatment and control villages is not statistically different. Notice also that, while ‘poor’ households are more likely to receive transfers from family and friends (at least in the control villages), ‘non-poor’ are (both in treatment and control villages) more likely to receive remittances from migrants.

Table 2
Percentage of households receiving transfers

From friends, neighbours and relatives (but not migrants)			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	0.0532	0.0442	0.0785
Control villages	0.0462	0.0523	0.0567
t-statistic for difference	1.1328	-1.0640	
(p-value)	(0.257)	(0.287)	
From migrants			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	0.0784	0.0497	0.0660
Control villages	0.0766	0.0539	0.0560
t-statistic for difference	0.2393	-0.5964	
(p-value)	(0.811)	(0.551)	
From either of them			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	0.1228	0.0890	0.1344
Control villages	0.1154	0.1006	0.1051
t-statistic for difference	0.8074	-0.9670	
(p-value)	(0.420)	(0.334)	

The large majority of transfers received by households were in money: around ninety per cent of the households supported by migrants and three quarters of those obtaining help from anyone else were being given monetary transfers. In Table 3, we show the average monetary amount received by households: that is, the average among those actually getting some monetary transfer. Remarkably, conditional on receiving a transfer, 'non-poor' households receive much larger transfers than poor ones. This is true both in treatment and control village and is particularly true for transfers from friends and relatives.

Table 3
Average monthly amount of received monetary transfers
(in pesos, conditional on positive transfer of each type)

From friends, neighbours, and relatives (but not migrants)			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	784.23	403.18	422.05
Control villages	704.35		410.47
t-statistic for difference	0.5074		465.21
(p-value)	(0.613)		-1.1363
			(0.256)
From migrants			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	259.53	191.40	202.13
Control villages	254.59		194.61
t-statistic for difference	0.0929		225.31
(p-value)	(0.926)		-1.4404
			(0.150)
From either of them			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	472.54	292.17	329.81
Control villages	433.20		305.21
t-statistic for difference	0.5220		359.37
(p-value)	(0.602)		-1.9441
			(0.052)

If all the poor households in the treatment villages were receiving grants, Tables 2 and 3 could give a first simple test of crowding out. One could simply compare ‘poor’ households in treatment and control villages and check whether those in treatment villages receive smaller and less frequent transfers. Unfortunately, as we discussed above, the fact that a non-random subset of households in treatment villages does not receive (for a reason or another) the treatment makes such a straightforward test unfeasible. While we discuss this problem and our proposed solution below, at this stage it is nonetheless interesting to notice that both among the poor and non-poor and for both transfers from migrants and others, the treatment and control villages do not exhibit any differences in the percentage of households receiving a transfer and in the quantity received. The only significant difference is shown between the amount of transfer received by poor households from anyone (i.e. aggregating migrants and others). In interpreting these results, however, it should be considered that among the

poor in treatment villages, only some households receive transfers, and that the tables indicate that these households receive less often and smaller amounts. Interestingly, the poor who do not receive the programme in the treatment villages (the ‘non-compliers’) seem to receive more often and larger amounts. In our more formal analysis below we will use these differences to unravel the effect of the programme.

Before we start our formal analysis, it is worth considering some of the features of the households in the villages, distinguishing programme participants and household receiving transfers.

Families in the survey were asked about their total consumption on different goods in the last week before the survey; we know how many of these consumed goods have been bought and how many have been produced by the household (i.e. ‘self-consumption’). Thus, we compute the weekly expenditure on food for households in our sample by summing up across the different items; ‘self-consumption’ is valued at market prices.⁴ Summary statistics are reported in Table 4. Table 4 presents the weekly expenditure on food for households in our sample; we have included consumption from goods produced by the household valued at market prices.

Table 4
Average (and median) weekly household expenditure in food (in pesos)

Households not receiving transfers			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	156.39 (133)	145.57 (127)	130.62 (112)
Control villages	168.37 (143)	137.64 (119)	
Households receiving transfers			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	156.97 (130)	142.56 (116)	128.96 (99)
Control villages	146.50 (129)	134.35 (109)	
All households			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	156.46 (133)	145.23 (126)	130.40 (111)
Control villages	165.84 (142)	137.29 (118)	

⁴ The market value of self-consumption was computed using information from a locality survey, carried out by PROGRESA and which is parallel to the main household survey. Amongst other local information, it provides the prices of consumption goods in the two closest shops.

Table 5 reports the average household income in our sample. We added up self-reported income (by each household member) from the following sources: wages, net profits from self-employment, pensions, interests, communitary earnings and income from rental (of land, animals or machinery);⁵ we have also included self-consumption. Therefore, our measure of household income does not include either the private transfer received or the programme support. Notice that there seems to be no important differences between treatment and control villages. However, we can reject the hypothesis that beneficiaries in treatment villages and poor households in control villages have the same average income (for households not receiving transfers and for all households, but not for households receiving transfers): the income of beneficiaries seems to be lower than those of the control households. It is tempting to interpret this as an effect of the programme arising from a reduction in child labour supply. However, Gomez de Leon and Parker (1999) find no strong effects of the programme on child labour supply.

Table 5
Average (and median) weekly household income (in pesos)

Households not receiving transfers			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	370.36 (231)	262.33 (189)	289.18 (188)
Control villages	395.68 (258)		326.89 (200)
Households receiving transfers			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	321.92 (214)	299.10 (177)	474.19 (180)
Control villages	325.72 (211)		258.82 (175)
All households			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	365.78 (230)	264.68 (189)	306.45 (188)
Control villages	389.15 (254)		321.54 (195)

⁵ As the respondent can report daily, weekly, biweekly, monthly or annual income, all reported income is first converted into weekly income; we used information on the number of days worked in the previous week to compute this weekly income when daily earnings were reported.

Table 6
Transfers receipts relative to income and consumption

A) Median of ratio monthly monetary transfer/food expenditure			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	0.4021	0.3127	0.4130
Control villages	0.3302	0.3708	
B) Median of ratio monthly monetary transfer/household income			
	Non-poor	Poor	
		Benef.	Non-benef.
Treatment villages	0.2901	0.1633	0.1832
Control villages	0.2239	0.2189	

In Table 6, we show that transfers are quite important relative to household food expenditure and household income. Transfers can expand the budget by 15 per cent to 30 per cent and help to buy over one third of the food consumed in the household.

4 An econometric specification and econometric problems

In this section we describe our econometric specification and some of the econometric problems we had to face.

4.1 Econometric specification

As we discussed in section 3, we have information both on monetary transfers and on transfers in kind. However, for the latter, we do not have information on the value of the transfer. For this reason, we specify a Probit model for the probability of receiving any transfers, be it in kind or monetary.⁶ Such a variable does not consider the size of the transfer. Moreover, as some households might be giving *and* receiving transfers, we cannot assess whether such households are net givers or receivers. On the other hand, as shown in Table A2 in the Appendix, most of the families involved in inter-households transfers (around 10 per cent out of all households, both in treatment and in control villages) have just received some transfer, whereas less than 0.25 per cent have received but also given transfers. Likewise, it is interesting that the non-poor households are more likely to give transfers without receiving them than the poor ones.

For net monetary transfers, instead, we estimate a simple Tobit, truncated at zero. Again, few households in our sample are giving monetary transfers (25 out of 23,268), and only six of them are also receiving transfers from friends or migrants. So, we have

⁶ In our Probit specification, the dependent variable is the indicator for receiving transfers commented in the previous section.

finally only 21 net givers and the difference of analysing gross and net transfers is negligible.

As we found it difficult to think of a variable that would affect the probability of receiving a transfer while not affecting the size of the grant, we did not use a generalized Tobit, as it would only be identified through functional form assumptions.

4.2 Some econometric problems

As mentioned above, not everybody entitled to the programme gets it in the Progresa village. With a slight abuse of language, we call these households, which correspond to the ‘treatment-poor-non-beneficiary’ branch of the tree in Figure 1, ‘non-compliers’. Our understanding of this phenomenon is limited. However, it seems that it is due to a combination of factors, including administrative delays, non-compliance with the terms of the programme and outright refusal of the programme. Whatever the reason for the existence of non-compliers, two things are evident. First, we cannot identify them in the control villages. Second, they are not a random subset of the beneficiaries. If we try to predict ‘non-compliers’ among the beneficiaries in the treatment villages, we find many significant factors, even though, no combination of variables can predict them perfectly.⁸ The issue is potentially a serious one and has been discussed in the evaluation literature, see Heckman et al. (1998) and Angrist et al. (1996). If we were willing to assume that the effect of the programme on the non-compliers is zero, it is easy to show, (see Angrist et al. 1996) that the presence of non-compliers would not imply a problem, given the randomized nature of the treatment. However, we feel such an assumption is too strong in our context. As the non-compliers (and the non-beneficiaries) live in villages where other individuals receive the programme, it might be that, for these households, the probability of receiving a private transfer increases as a consequence of the programme. Such a scenario would violate the assumption that the programme has no effect on non-compliers and therefore invalidate the procedure proposed by Heckman et al. (1998) and by Angrist et al. (1996). In other words, general equilibrium effects are very likely to be operative in our context. It is therefore worthwhile to see if we could identify the effects of interest using alternative assumptions to those proposed in the literature. To make our point it is not necessary to focus on the non-linearities implied by the Probit and Tobit structure. To keep the discussion simple, therefore, we will discuss it within a linear system.

Let Y_t and Y_c denote the outcomes of interest in treatment and control villages, respectively. Let p take the value 0, 1 in both control and treatment villages indicating poor household entitled (in the treatment villages) or eventually entitled (in control villages) in the programme. For $p=1$, let b take the value 0 or 1; $b=1$ are beneficiaries and $b=0$ are non-beneficiaries. Finally let k be the fraction of beneficiaries. Proper randomization will guarantee that k is the same in treatment and control villages, even if complying (p) is a decision variable.

⁷ It turns out that the non-compliers are much more frequent among the densificado households and among those with few or no children.

We are interested in estimating the effect of the programme, that is, the average effect of the treatment on the treated:⁸ $E[Y_t | p = 1, b = 1] - E[Y_c | p = 1, b = 1]$. Notice that we do not observe or measure the second element because we do not know who would be a beneficiary in the control villages. Notice also that:

$$E[Y_t | p = 1] = kE[Y_t | p = 1, b = 1] + (1 - k)E[Y_t | p = 1, b = 0]$$

$$E[Y_c | p = 1] = kE[Y_c | p = 1, b = 1] + (1 - k)E[Y_c | p = 1, b = 0]$$

From this it follows that:

$$(1) \quad E[Y_t | p = 1, b = 1] - E[Y_c | p = 1, b = 1] =$$

$$\frac{E[Y_t | p = 1] - E[Y_c | p = 1]}{k} - \frac{1 - k}{k} \{E[Y_t | p = 1, b = 0] - E[Y_c | p = 1, b = 0]\}$$

If one assumes that the second term on the right-hand side of this expression is zero, which would be the case if one assumes that the treatment has no effect on non-beneficiaries, then you obtain what Heckman et al. (1998) or Angrist et al. (1996) suggest.

However, if one not willing to make this assumption, it is necessary to rely on something else. The alternative and weaker assumptions we propose is the following:

$$(2) \quad E[Y_t | p = 1, b = 0] - E[Y_c | p = 1, b = 0] = E[Y_t | p = 0] - E[Y_c | p = 0]$$

that is that the difference in the outcomes of non-compliers in treatment and control villages is the same as the difference in the outcomes for non-beneficiaries. That is, if one thinks that the treatment has an effect on non-beneficiaries because of general equilibrium effects, we assume that these effects are the same for the non-poor and for the non-beneficiaries. The right hand side of equation (2) can be estimated.

In the next section, we present two estimates for the effect of the programme along the lines discussed here. We call Effect 1 to the effect of the programme estimated using equation 1 under the assumption that the treatment has no effect on non-compliers (thus the second term on the right-hand side of (1) vanishes). On the other hand, by substituting (2) into (1) you obtain what we call Effect 2 in the tables containing our empirical results.

In addition to these issues, one needs to worry about the computation of standard errors. First, one would like to use methods that are robust. Second, one would like to take into account the presence of correlation among observations at the village level. For this reason we use bootstrapped standard errors with cluster effects at the village level. In other words, we first estimate conventional Probit and Tobit models. Then we use bootstrapping allowing for village cluster effects, such that the sample drawn during each replication is a bootstrap sample of clusters. Therefore we obtain with this

⁸ It may also be of some interest to assess $E[Y_t | b = 1] - E[Y_c | b = 1]$, that is ignoring the compliers issue altogether. Given the randomization, this is trivial.

procedure standard error for the parameters in the Probit and Tobit models which are robust to village cluster effects. Furthermore, we report the bootstrapped mean of the two effects and their corresponding bootstrapped standard errors.

4.3 Crowding-out and grant size

If the grant was uniform across households (as is the case of the nutritional supplement), one would only be interested in the coefficient on the dummy variable indicating beneficiary households in treatment villages or in the estimated effect that takes into account the non-complier problem discussed above. In reality, however, different households are entitled to different amounts of grant and, ex-post, receive different levels of support. It is therefore interesting to check whether private transfers are affected by the level of the public grant, that is, whether households receiving a larger public grant are less likely to receive a private transfer and, conditionally on getting it, receive smaller amounts.

As mentioned above, the amount each household is entitled to depends on the number of school age children living in the household and the last approved grade of each child. The amount they do receive depends on whether each child is enrolled or not. Therefore, what the household receives in addition to the basic nutritional subsidy is not an unconditional grant but is linked to specific actions (and depends on the household demographic structure). More than a transfer, a large proportion of the programme we consider should be interpreted as a change in the price of education for these households. However, for households that would have enrolled their children in school regardless of the presence of the grant, the grant represents a pure income transfer.

To have an idea of the effect of the size of the grant on private transfers, while side-stepping the issue of the possible endogeneity of the grant level, we add to our Probit and Tobit specifications the potential grant that a ‘poor’ households (in a treatment or control village) would be entitled to receive. In particular, we consider six different levels of the potential grant and interact the corresponding dummies with the treatment-poor and control-poor dummies as well as with the ‘treatment-poor-non-beneficiary’ dummy to take into account the non-complier issue discussed in the previous section. Having computed the effect of the grant for each of the six grant classes, we then impose either a linear or a quadratic effect of the grant by minimum distance.

5 Results

In this section, we present our estimation results for the probability of receiving a transfer and its amount. We show separate estimates for those transfers coming from ‘migrants’ and from friends or relatives.

First, we report the results for the transfers received from anyone but migrants. In particular, Table 7 reports the results of a Probit model for the probability of receiving any (in columns 1-3) and the results of a Tobit model for net monetary transfers (columns 4-6). All specifications in these tables include a set of controls (state dummies, shock dummies, individual illness and days lost) whose coefficients are not reported to save space. In columns (1) and (4), we check the general effect of the programme on private transfers, that is, we introduce a dummy for the treatment villages but we do not distinguish between beneficiaries and non-beneficiaries. We find a very small negative effect that it is not significantly different from zero.

Table 7
Transfers from friends

	Probit			Tobit		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment villages	-0.0143			-46.3592		
	(0.0457)			(67.2524)		
Treatment villages		-0.1152	-0.1020		-188.4291	-174.7272
Poor benef.		(0.0641)	(0.0569)		(97.1155)	(75.1016)
Treatment villages		0.1461	0.1278		174.0266	154.2646
Poor non-benef.		(0.0651)	(0.0703)		(112.1027)	(101.3515)
Treatment villages		-0.0709	-0.0480		-71.0503	-43.6969
Non-poor		(0.0637)	(0.0663)		(110.1851)	(113.8477)
Control villages		-0.1356	-0.1178		-126.0521	-102.7882
Non-poor		(0.0651)	(0.0717)		(102.2041)	(98.7407)
Log consumption			-0.2270			-239.3304
			(0.0343)			(48.7671)
Effect 1		-0.0646	-0.0577		-128.1521	-121.2951
		(0.0636)	(0.0647)		(106.9990)	(102.7401)
Effect 2		-0.0870	-0.0819		-147.2029	-141.7623
		(0.0788)	(0.0798)		(86.0295)	(105.8298)
Observations	23,247	23,247	23,247	23247	23247	23247

Note Robust Standard errors in parentheses.

Constant, state dummies, shock dummies, days of illness, days of work lost are always present. Effect 1: treatment effect under the assumption that the effect on non-compliers in treatment and control villages is the same.

Effect 2: treatment effect under the assumption that the difference in the effect on non-compliers in treatment and control villages is the same as the difference between non-beneficiaries. These effects and their standard errors have been computed by bootstrapping

In columns (2)-(3), we distinguish the programme's effects on households according to their entitlement status and also take into account if they are actually beneficiaries or not. In particular, we substitute the treatment villages dummy with a set of dummies that indicate the poor beneficiaries, the poor non-beneficiaries and the non-poor in the 'treatment' villages and the non-poor in the control villages. The excluded group, therefore, is that of the 'poor' in the control villages.

In the absence of ‘non-compliers’, one could read from the coefficient on the treatment poor (the beneficiary of the programme relative to similar households in control villages) the effect of the programme. However, as discussed in section 4, the presence of non-compliers requires additional computations to unravel the effect of the programme. The effect estimated making the assumptions suggested by Heckman et al. (1998) and by Angrist et al. (1996) are reported in the row labelled ‘Effect 1’. The effect estimated the alternative assumption we suggest and discuss in section 4 are in the row labelled ‘Effect 2’.

The results in column (2) indicate that beneficiaries are less likely to get private transfers as compared with poor households in control villages; on the other hand, non-beneficiaries appears to receive more transfers. Although our alternative assumption concerning the importance of general equilibrium effects seems to be relevant; the negative impact is not too strong in both cases. These results are basically unchanged when, in column (3), we add to our basic specification the log of consumption. Such a variable is strongly significant and takes a negative sign: households who enjoy a higher level of consumption receive less transfers. The coefficients on the programme dummies, however, are substantially unaffected.

When considering net monetary transfers (columns 4-6), the negative effect appears strongly significant. The amount by which the public grant crowds out is less than one for one (140 is below the average grant obtained from PROGRESA) but is substantial.

As discussed in section 4.3, we explicitly take into account the possibility that the effect of the programme on private transfers depends not only on whether a household receives the public transfer or not, but also on the ‘intensity of the treatment’, that is, on how much a household can (potentially) be paid. Thus, we fit a linear or quadratic relationship between the programme effect and the potential grant, whose coefficients are reported in Table 8. In particular, the coefficients we report in this table, take into account the non-compliers issue, using both assumptions discussed in section 2. As we mentioned in Section 4.3, the parameters are fitted by minimum distance on the coefficients corresponding to those we label ‘Effect 1’ and ‘Effect 2’ in Table 7, for each of six grant categories. A complete set of estimates is available upon request. In Table 8, we also report the effect of the programme, implied by the estimated coefficients in each specification, evaluated at the average grant in our sample (250 pesos).

Table 8
Treatment Effects on transfers from friends, as a function of the grant size

	Probit				Tobit			
	Effect 1		Effect 2		Effect 1		Effect 2	
Intercept	-0.0647 (0.1053)	0.4658 (0.1946)	-0.0801 (0.1025)	0.4468 (0.1922)	-184.159 (156.133)	795.567 (320.841)	-161.037 (150.083)	803.471 (314.109)
Grant	-0.0002 (0.0002)	-0.0044 (0.0013)	-0.0002 (0.0002)	-0.0044 (0.0013)	-0.034 (0.377)	-7.738 (2.236)	-0.034 (0.377)	-7.738 (2.236)
Grant ²		6.5E-06 (2.0E-06)		6.5E-06 (2.0E-06)		0.012 (0.003)		0.012 (0.003)
Effect at mean Grant	-0.1076 (0.0810)	-0.2163 (0.0877)	-0.1230 (0.0785)	-0.2353 (0.0858)	-192.724 (112.460)	-397.093 (126.751)	-169.602 (107.912)	-389.189 (124.866)

Both in the Probit and Tobit case, the coefficients in the linear specifications do not appear to be individually significant, although the programme effect evaluated at the average grant are comparable with those displayed in Table 7 (in particular, they are negative and marginally significant). On the other hand, we find that the quadratic specification fits quite well: all coefficients are individually significant. As for the reported effects evaluated at the average grant, we find again that the programme has a significant negative effect on the probability of receiving a transfer and on the amount received conditional on receiving a private transfer. In this case, our point estimate for the programme effect implies a large crowding-out. Although our estimates are higher than those in many existing studies, some papers also find that, for very poor households in developing countries, public transfers can induce a strong crowding out. Cox and Jimenez (1995), for instance, report evidence that public pension programmes for the Philippines would result in a 92 per cent reduction of private transfers. In another recent paper, Cox et al. (1999) claim that, once the non-monotonicity in the transfer-income relationship is taken into account, public transfers are found to have a substantial negative effect on private transfers.

We now move to the consideration of the results that we obtain for transfers received from migrants. Table 9 reports the results of a Probit model for the probability of receiving any transfer and the results of a Tobit model for net monetary transfers. The set of controls common to all specifications and the different specifications (in each column) are the same as in Table 7. However, in estimating these equations we use only households that report having migrants. This explains the reduction in the number of observations. As we are estimating the effect of the programme off the variation induced by the randomization, we do not need to worry about the selection issues connected with the consideration of only households with migrants.

Table 9
Transfers from migrants (only for households with migrants)

	Probit		Tobit			
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment villages	-0.0364 (0.0462)			-25.2468 (22.8879)		
Treatment villages		-0.1134 (0.0639)	-0.1113 (0.0729)		-59.2782 (22.3539)	-61.2460 (30.4710)
Poor benef.						
Treatment villages		0.0083 (0.0780)	0.0046 (0.0810)		-28.4244 (34.8725)	-24.8774 (28.6178)
Poor non-benef.						
Treatment villages		0.0026 (0.1054)	0.0047 (0.0907)		-6.0889 (32.2626)	-8.8712 (33.4399)
Non-poor						
Control villages		-0.0664 (0.0932)	-0.0629 (0.0667)		-43.3385 (30.6215)	-46.4576 (36.6859)
Non-poor						
Log consumption			-0.0347 (0.0344)			33.7289 (15.6499)
Effect 1		-0.1105 (0.0778)	-0.1097 (0.0800)		-69.1235 (35.2458)	-69.8627 (34.2884)
Effect 2		-0.1344 (0.0929)	-0.1331 (0.0778)		-82.0255 (35.7559)	-82.8814 (36.7581)
Observations	3721	3721	3721	3721	3721	3721

Note: Robust Standard errors in parentheses. See also note in Table 7.

As in Table 7, the coefficients associated with the plain ‘Treatment villages’ dummy in column (1) and (4) are negative but not statistically different from zero. Once again, however, once we control for who receives the programme, we found significant crowding out effects, both in the Probit and in the Tobit specifications. Notice however, that the crowding out of this kind of transfers seems to be less strong than that for transfers received from friends and family. This result is consistent with the idea that people living in the village or near by, have better information on the nature of the grant and therefore can react more to its introduction.

Finally, we present in Table 10 the results of fitting a linear or quadratic relationship between the programme effect and the potential grant, whose coefficients are reported in Table 8; also we present the effect of the programme. In this case, both linear and quadratic specification fits very poorly, being insignificant all the estimated coefficients in the relationship. Nevertheless, the estimated effects, evaluated at the average grant, are similar to those in Table 9.

Table 10
Treatment effects on transfers from migrants as a function of the grant size

	Probit				Tobit			
	Effect 1		Effect 2		Effect 1		Effect 2	
Intercept	-0,1058 (0,1362)	-0,3491 (0,2804)	-0,1353 (0,1344)	-0,3751 (0,2764)	-60,189 (56,619)	-119,842 (110,850)	-70,744 (57,150)	-130,140 (110,769)
Grant	5,0E-05 (3,4E-04)	0,0018 (0,0018)	0,0001 (0,0003)	0,0018 (0,0018)	-0,009 (0,133)	0,409 (0,681)	-0,009 (0,133)	0,409 (0,6810)
Grant ²		-2,5E-06 (2,5E-06)		-2,5E-06 (2,5E-06)		-0,001 (0,001)		-0,001 (0,001)
Effect at mean Grant	-0,0932 (0,0889)	-0,0623 (0,0942)	-0,1227 (0,0863)	-0,0883 (0,0930)	-62,438 (36,463)	-55,704 (38,017)	-72,992 (36,732)	-66,002 (38,392)

6 Conclusions

In this paper we have analysed the effect of a large welfare programme in rural Mexico, PROGRESA, on the private transfers received by the beneficiaries of such a programme. While PROGRESA is not a pure transfer programme, as most of the subsidies the beneficiaries household receive are conditional on school attendance or other specific actions the programme is meant to encourage, the evaluation sample we use has the big advantages of having a pure randomization component. In particular, in 186 of the 506 villages included in the evaluation sample, the implementation of the programme was delayed for two years. This allows us to overcome the endogeneity problem that plagues many of the papers that have looked at the crowding out of private transfers induced by public programmes. While it is true that PROGRESA is not a pure transfer, it is fair to say that it does have a positive income effect for the beneficiary families, both because it has an unconditional component and because the education subsidy probably more than compensate for the loss of income from child labour. Moreover, the programme does not change the observed enrolment behaviour of many households.

From a methodological point of view, we have to deal with the fact that some beneficiary households in treatment villages do not receive the programme. This feature of the programme prevents the simple comparison of treatment and control villages' poors, because some of the treatment villages' poors do not receive the programme. On the other hand, those poor households who do receive the programme in treatment villages are not strictly comparable to the poor in control villages. We use two different sets of assumptions that allow us to identify the effect of interest.

Our results indicate that the programme does crowd out private transfers. Both the likelihood to receive a transfer and the amount received conditional on receiving private transfers are significantly and negatively affected by the programme. This result is consistent with the implications of several theoretical models that have been considered in the literature. We consider separately the transfers received from friends and relatives other than migrant children and those received from migrants and find that the crowding out effects, especially in the amount received, are weaker for the migrant transfers. This might be an indication that the mechanism through which private transfers are crowded out by the public programme is different for the two classes of transfers.

This evidence is particularly important, as the variation we use is, by construction, exogenous. Our data therefore avoid some of the problems that have affected the empirical literature on this subject. The general conclusion we draw from this exercise is that this particular programme has a negative effect on private transfers. Some of our estimates indicate that the crowding out effects can be quite large. Having said this, however, in interpreting our results it should be remembered that the main objective of PROGRESA is that of fostering the accumulation of human capital among poor rural households. Our results do not say anything about the effectiveness of the programme in achieving such a goal. It should be clear, however, that when evaluating a public programme, such as PROGRESA, one has to take into account the fact that such programmes do not occur in a vacuum but interact with existing mechanisms within a society.

Appendix

Table A1
Distribution of Households, October 1998

Treatment villages				
	Non-poor	Benef.	Poor Non-benef.	Total
Non-poor in 1997	3,176	1,380	1,977	6,533
Poor in 1997		6,924	902	7,826
	3,176	8,304	2,879	14,359
Control villages				
	Non-poor		Poor	Total
Non-poor in 1997	2,011		2,248	4,259
Poor in 1997			4,650	4,650
	2,011		6,898	8,909

Table A2
Households involved in transfers (from friends and migrants) by entitlement

Treatment villages				
	Non-poor	Benef.	Poor Non-benef.	Total
Received not given	380 (11.96)	724 (8.72)	382 (13.27)	1486 (10.34)
Received given	10 (0.31)	15 (0.18)	5 (0.17)	30 (0.21)
Not received given	37 (1.16)	33 (0.40)	21 (0.73)	91 (0.63)
Not received not given	2749 (86.56)	7531 (90.70)	2471 (85.83)	12751 (88.81)
Total	3176 (100.00)	8303 (100.00)	2879 (100.00)	14358 (100.00)
Control villages				
	Non-poor		Poor	Total
Received not given	229 (11.39)		712 (10.32)	941 (10.56)
Received given	3 (0.15)		13 (0.19)	16 (0.18)
Not received given	25 (1.24)		33 (0.48)	58 (0.65)
Not received not given	1754 (87.22)		6140 (89.01)	7894 (88.61)
Total	2011 (100.00)		6898 (100.00)	8909 (100.00)

References

- Angrist, J.D., Imbens, G.W. and Rubin, Donald B. (1996), 'Identification of Causal Effects Using Instrumental Variables', *Journal of the American Statistical Association*, 91, 444–55.
- Attanasio, O.P. and J.V. Rios Rull (2000a), 'Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?' *European Economic Review*, 44, 1,225–58
- Attanasio, O.P. and J.V. Rios Rull (2000b), 'Consumption Smoothing and Extended Families', forthcoming in Dewatripoint and Hansen (eds) *Advances in Economic Theory—The World Congress of the Econometric Society—Seattle 2000*.
- Becker, G. (1974), 'A Theory of Social Interactions', *Journal of Political Economy*, 82, 1,063–93.
- Behrman and Todd (1999), 'Randomness in the Experimental Samples of PROGRESA', IFPRI Mimeo.
- Bernheim, D., Schleifer, A. and L.H. Summers (1985), 'The Strategic Bequest Motive', *Journal of Political Economy*, 93, 1,045–76.
- Cox, D. (1987), 'Motives for Private Transfers', *Journal of Political Economy*, 95, 508–46.
- Cox, D., Hansen, B.E. and E. Jimenez (1999), 'How Responsive are Private Transfers to Income? Evidence from a *laissez-faire* Economy', Manuscript.
- Cox, D., Eser, Z. and E. Jimenez (1998), 'Motives for Private Transfers over the Life Cycle: An Analytical Framework and Evidence from Peru' *Journal of Development Economics*, 55, 57–80.
- Cox, D. and E. Jimenez (1995), 'Private Transfers and the Effectiveness of Public Income Redistribution in the Philippines', in van de Walle, D. and K. Nead (eds) *Public Spending and the Poor: Theory and Evidence*, Baltimore and London: Johns Hopkins University Press for the World Bank, pp.226-58.
- Cox, D., Jimenez, E. and W. Okrasa (1997), 'Family Safety Nets and Economic Transition: A Study of Worker Households in Poland', *Review of Income and Wealth*, 43, 191–209.
- Foster, A. and M. Rosenzweig (2001), 'Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas', *Review of Economics and Statistics*. 83(3), August, pp.389–407.
- Gomez de Leon, J. and S. Parker (1999), 'The Impact of Anti-Poverty Programs on Labor Force Participation', Progres, Mimeo.
- Heckman, J., Smith, J. and C. Taber (1998), 'Accounting for Dropouts in Evaluations of Social Programs', *Review of Economics and Statistics*, 80, 1–14.
- IFPRI (2000), 'Is PROGRESA Working? Summary of the Results of an Evaluation by IFPRI', Mimeo, Washington DC.
- Jensen, R. (1999), 'Public Transfers, Private Transfers, and the 'Crowding Out' Hypothesis: Evidence from South Africa', Mimeo.
- Schoeni, R.F. (1997), 'Private Intrahousehold Transfers of Money and Time: New Empirical Evidence', *Review of Income and Wealth*, 43, 423–64.