SAMUEL FREIJE
ROSANGELA BANDO
FERNANDA ARCE

Conditional Transfers, Labor Supply, and Poverty: Microsimulating Oportunidades

This article summarizes a microsimulation exercise for the Mexican human development program, Oportunidades, and presents a series of simulations of its actual and potential impact on poverty at the national, urban, and rural levels. The simulations were partly based on a program that we designed and developed under the sponsorship of the Mexican Secretary of Social Development (*Secretaría de Desarrollo Social,* or SEDESOL). Our goal is to provide ex ante estimates of the potential effects that changes in program design, in terms of the selection of beneficiaries or the amount of benefits, may have on poverty in Mexico.[1]

Microsimulations differ from impact evaluations in that they allow analysts to assess different possible policy measures before the measures are implemented. Several ex post impact evaluations of Oportunidades highlight the success of this program in terms of beneficiaries' health and school enrollment. Oportunidades, which was initiated in late 1997 under the name Progresa, has become one of the most important and well-known programs in Latin America. By 2005, it covered five million recipient families with an

Freije is with the Universidad de las Américas; both Bando and Arce were with the Centro de Investigación en Economía y Políticas Públicas when the paper was written. Bando is currently a doctoral student at the University of California at Berkeley; Arce is with the Secretaría de Desarrollo Social (SEDESOL).

The results of this paper are partly based on a microsimulation software funded by SEDESOL's Subsecretaría de Prospectiva, Planeación y Evaluación, under contract SSPPEP-AD-011-04. We thank Raquel Bernal, François Bourguignon, Francisco Ferreira, Gonzalo Hernández Licona, Luis Felipe López-Calva, Carlos Medina, Isidro Soloaga, and Amadeo Spadaro for their comments at different stages of this research. Special thanks go to Rafael Ramírez for his programming skills.

1. The simulation tool also performs microsimulations for the SEDESOL programs LICONSA and DICONSA. See IPD (2005).

annual budget of U.S.$2.1 billion.[2] The program served only rural families through 2002, when it began to incorporate the urban sector. Given the program's current size and reputation, administration officials cannot rely exclusively on ex post evaluations for judging the potential effects of changes in the program. A prospective instrument for steering the future development of Oportunidades is needed, and microsimulation techniques, such as those presented here, complement and expand what we have learned from the experimental and quasi-experimental program evaluations.

Microsimulation exercises are a common instrument of policy analysis. Several literature surveys review the evolution and extension of these techniques applied to different fields.[3] Bourguignon and Pereira survey different techniques for evaluating the distributional impact of economic policies.[4] They classify these different techniques according to several perspectives: ex ante or ex post, accounting or behavioral, and partial equilibrium or general equilibrium. Ex ante simulations are performed before the program is enacted to predict the likely impact on chosen variables, whereas ex post simulations show what a program's effect would have been had it been designed differently. The accounting simulation simply modifies the variable under study (for example, transfers or subsidized prices) for the selected sample observations that comply with certain eligibility requirements. Distribution indicators for the database with modifications are then compared with the corresponding measures for the original data. The comparison of the real and counterfactual distributions allows analysts to infer the distributional impact of the hypothesized change in policy. However, individuals affected by policy changes are generally likely to modify their behavior (for example, labor supply or consumer demand) in response to the policy. A behavioral simulation, therefore, is one that seeks to predict the changes in economic behavior by the individuals subject to changes in the policy under study. Finally, a general equilibrium simulation takes into account the effects of changes in beneficiaries' behavior on other variables, such as the price of related goods, as well as changes in the behavior of individuals who are not subject to the policy but are affected by changes in the market as a whole.

2. For a description of Oportunidades and some impact evaluation studies, see Skoufias (2005); Gertler and Boyce (2001); Skoufias and McClafferty (2001); Parker and Skoufias (2000); and Gertler (2000). For the most recent external evaluation, see Hernández Prado and Hernández Ávila (2005).

3. For reviews, see Bourguignon and Pereira da Silva (2003); Conte, Hegselmann, and Terna (1997); Harding (1996); and Merz (1991). For an earlier attempt at simulating Oportunidades, see Attanasio, Meghir, and Székely (2003).

4. Bourguignon and Pereira da Silva (2003).

This paper presents results from both accounting and behavioral partial equilibrium microsimulation exercises that allow us to estimate the effects of Oportunidades on poverty. Oportunidades is a conditional cash and benefits program that promotes the accumulation of human capital in terms of education, health, and nutrition among Mexican families in extreme poverty. Selected families receive cash transfers of up to 1,540 Mexican pesos a month (around U.S.$150), on the condition that children attend school and family members go to health clinics provided by the program. The program does not explicitly aim to raise family income in the short term.[5] The cash transfers do change households' disposable income, however, and they can raise the beneficiaries' families above the poverty line.

Given that most evaluations of Oportunidades address its impact on health and schooling, we instead gauge the program's impact on current poverty. To measure poverty, we use the Mexican government's official poverty line and definition. Poverty is thus defined as the percentage of households whose current income per member is below an absolute (food) poverty line.[6] The microsimulation exercise also provides measures of the poverty gap to support additional inferences about the cost of expanding the program by increasing either benefits or beneficiaries.

This paper addresses one counterfactual question and two prospective simulation questions. First, what would have happened to poverty if the program had been abolished? Second, what would happen to poverty if the amount of transfers changed? Finally, what would happen to poverty if the coverage of the program was extended in urban areas? Our simulations suggest the following answers. The national poverty headcount would have been 2 percentage points higher if Oportunidades had not been in place, while the rural headcount would have been 5 percentage points higher. This implies that up to a third of the reduction in rural poverty between 1998 and 2002 can be associated with the program. Doubling benefits under the current selection criteria would reduce the rural poverty headcount 7 percentage points further, but it would have no significant effect on urban poverty. Doubling the number of urban beneficiaries would reduce poverty only if better targeting were adopted.

5. The amount of the cash transfer per child was originally calculated so that it would approximately substitute for the forgone income of working children who start attending school because of the program. We thank John Scott Andretta for this remark.

6. The Mexican phrase is *línea de pobreza alimentaria.* In 2002 (the reference year), the poverty lines for the rural and urban sector, respectively, were 494.77 and 672.25 Mexican pesos per capita per month (around U.S.$45 and U.S.$65). See Comité Técnico para la Medición de la Pobreza (2002) for technical details on the construction of poverty lines for Mexico.

These results are derived from accounting simulations. The behavioral simulations render almost identical results because labor supply is largely insensitive to the changes simulated in these exercises. Moreover, our models show that noticeable changes in the distribution of hours of work would require implausibly large changes in cash transfers.

The paper develops as follows. We start by describing the structure of the microsimulation tool and then undertake an accounting exercise to answer the questions posed above. A subsequent section describes the Mexican labor market and the econometric models estimated for simulating labor supply responses to changes in cash transfers. We then run a behavioral exercise and test some of our simulated results by comparing them with some recent experimental and quasi-experimental studies performed on Oportunidades. The final section concludes.

Design of the Microsimulation Exercise

We assemble two microsimulation exercises: one accounting and one behavioral. The two exercises make use of the same database and follow very similar procedures.[7] Figure 1 shows the flowchart for the behavioral exercise.[8] The different computations performed on the data can be summarized in four stages. The first stage identifies the new beneficiaries of the program according to the analyst's instructions. It can add or remove beneficiaries based on two criteria: expansion or contraction of the program using the current selection mechanism; and selection of beneficiaries based on new criteria. Under the former criterion, two probit models (one for rural and one for urban areas) mimic the program's operating rules. These models allow us to estimate each household's probability of being selected and to rank households by these probabilities. Expansion (contraction) of the program is simulated by including (discarding) households whose fitted probabilities are immediately below (above) the critical index that defines the status quo of the program. This procedure assumes that the program runs as is and simply changes the number of beneficiary families. In the case of selection of beneficiaries based on new criteria, households are chosen as beneficiaries based

7. The procedure is akin to the one suggested by Bourguignon and Ferreira (2003).

8. The flowchart for the accounting exercise is very similar, but excludes the loop that allows for changes in school attendance and labor supply. The accounting simulation flowchart is available on request.

FIGURE 1. Flowchart for Behavioral Simulations

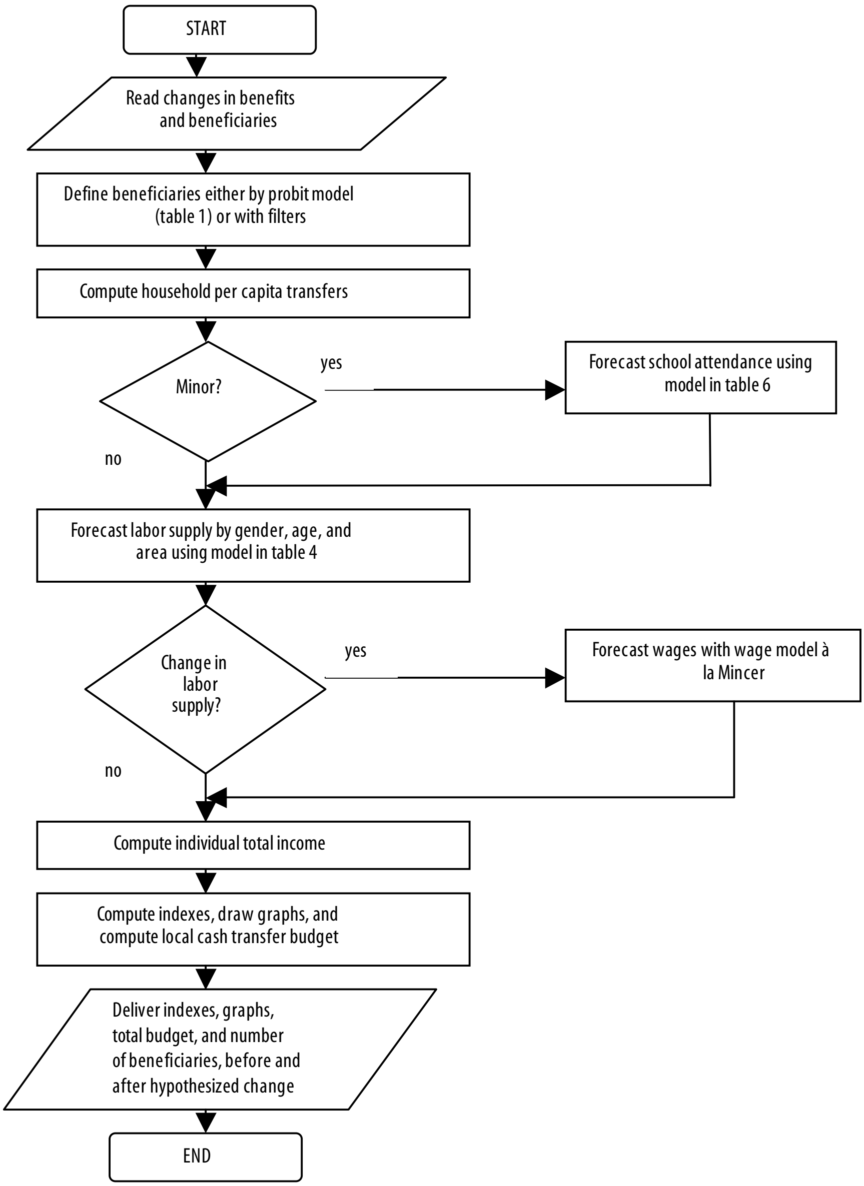

on characteristics that may, or may not, be the same as the current selection mechanism. This procedure further assumes that the agency is 100 percent effective in targeting beneficiaries. Thus, instead of using fitted probabilities to select beneficiaries, all households in the database with given characteristics are designated beneficiaries. For instance, if we wanted to simulate the effect of a program that only benefits poor families with two daughters aged four to fourteen, this procedure would select only this type of families, and it would select all of them.

The second stage defines new transfers to the simulated beneficiaries. The program allows for changes in the annual transfer for books, the monthly scholarship per child by gender, and the monthly household bonus. Only cash transfers are simulated.

The third stage simply adds the transfers defined in the second stage to the households selected in the first stage. If households were beneficiaries originally, the initial transfer is subtracted before the new transfer is added. Formally, household monetary income is defined as follows:

$$M_h^n = \sum_{i=1}^{N_h} w_{ih} + t_{ih}^o + d_{ih}\left(t_{ih}^n - t_{ih}^o\right) + r_{ih} + s_{ih}, \tag{1}$$

where d_{ih} equals one if the individual is a beneficiary under simulation and zero otherwise; M_h^n represents total family income after the change in transfers; N_h is the number of members in family h; w_{ih} stands for individual wages; and t_{ih}^o and t_{ih}^n are old and new individual transfers from Oportunidades. Other transfers and income from other sources are represented by r and s, respectively.

The fourth stage takes total household income from all families and computes poverty indexes following Foster, Greer, and Thorbecke.[9] It also calculates the total budget of cash transfers under the simulated program design.

In our behavioral simulations, the third stage incorporates models of school attendance and labor supply. The school attendance model is run for individuals aged twelve to seventeen, and it makes transfers conditional on the minor attending school. Thus, the scholarship transfer is given to the family only if the minor in question is forecast to attend school. Labor supply is forecast for both minors and adults, given the cash transfers received by the family. The third stage then follows three additional steps.

9. Foster, Greer, and Thorbecke (1984).

First, if the individual is between twelve and seventeen years of age, then we simulate changes in school assistance as a consequence of the changes in benefits introduced in stage two. We estimate the probability of school attendance using a probit model for the urban and rural sectors. If the individual is over 17 years old, the model skips this step.

Second, the model simulates changes in the labor supply as a consequence of changes in benefits introduced in stage two. Labor supply models and econometric estimation are explained below.

Finally, if labor supply changes, we estimate the new wage using an earnings model corrected for selection.[10] Otherwise, the original wage is kept for each individual.

The Data

We use the 2002 National Household Income and Expenditures Survey (*Encuesta Nacional de Ingresos y Gastos de los Hogares,* or ENIGH) as the basis for creating the simulated populations. This survey is representative at the national, urban, and rural levels, which allows us to compute poverty indexes at the three levels. It is the best survey for our microsimulation exercise because it provides the basis for official poverty figures and because it includes questions on individual and household characteristics that allow us to estimate several models of labor supply, school attendance, and program participation. In particular, it is the only national survey in Mexico that includes a range of questions on different sources of family income, including transfers.

The 2002 ENIGH also has several drawbacks, however. First, the survey was two years old when we performed the simulations.[11] This may reduce its similarity with the current Mexican socioeconomic reality. Second, the survey was not designed to be representative of social programs. For this reason, the survey significantly underestimates the number of recipients of Oportunidades compared with the official number of beneficiaries reported

10. Six selection correction models of wages were estimated for men, women, and minors in the urban and rural areas. These models have the expected results in terms of signs and individual and joint significance for the coefficients. The null hypothesis of no correlation between the wage and participation equation error terms is rejected in five of the six models (with the exception of rural minors). The estimation results for these models are available on request from the authors. See also IPD (2005).

11. The 2004 ENIGH was not publicly available when we built this microsimulation exercise.

by SEDESOL: the survey reports 161,355 urban beneficiaries (versus 508,446 according to SEDESOL) and 2,976,816 rural beneficiaries (versus 3,731,554 from SEDESOL), for a national total of 3,138,171 beneficiaries (versus 4,240,000 from SEDESOL).

We address this problem by calibrating the database to the official numbers. This makes the comparisons between the 2002 ENIGH and the simulated populations useful for policy analysis. The calibration exercise is based on the program participation models described in the first stage. The idea is that a probit model is estimated for program participation (see table 1). Using the fitted probabilities from this model (that is, the *p* score), we ranked households by whether they reported being beneficiaries. We then found the *p* score above which the cumulative number of households equals the official number of household beneficiaries.

Given that the calibrated number of beneficiaries is different from the number of beneficiaries in the sample, some households that reported no transfers from Oportunidades are fitted to be recipients. To maintain official poverty figures and to make comparable simulations, we redefine household monetary income as

$$M_h^n = \sum_{i=1}^{N_h} w_{ih} + t_{ih}^o + d_{ih}\left(t_{ih}^n - \tilde{t}_{ih}^o\right) + r_{ih} + s_{ih}, \tag{2}$$

where $\tilde{t}_{ij}$ represents the transfer that would correspond to the household based on its characteristics, using the current set of rules.[12]

In the participation models (see table 1), we try to replicate the discriminating formula used by Oportunidades personnel to select beneficiaries. The formula is confidential, so we estimate the probability of being a beneficiary in urban and rural areas, using as explanatory variables a series of household characteristics that are known to be part of the selection criteria.[13] The requisite characteristics for being a beneficiary are the same in both areas, but the selection mechanisms are different. In rural areas, Oportunidades was initially randomly assigned to poor households, whereas in urban areas, it was assigned to households that applied for the program and that had qualifying characteristics. Given that the targeting of beneficiaries is also affected

12. This formula is based on the assumption that fitted beneficiaries misreport their received transfers from Oportunidades as another source of income. Assuming that they do not report these transfers at all would imply that official poverty measures are overestimated, which we are not willing to accept because of the need to calibrate the database to official figures.

13. Here we closely follow Skoufias, Davis, and Behrman (2001).

TABLE 1. Probit Model for Selection of Oportunidades Beneficiaries in Rural and Urban Areas[a]

	Rural area			Urban area		
Explanatory variable	*df / dx*	*Std. error*	P > |z|	*df/dx*	*Std. error*	P > |z|
Household demographics						
Age of household head	0.0001	0.0005	0.868	0.0000	0.0000	0.608
Gender of household head (male)	−0.0245	0.0185	0.192	0.0004	0.0007	0.519
Schooling of household head	−0.0153	0.0021	0.000	−0.0002	0.0001	0.005
Proportion of women in household	0.0752	0.0308	0.015	0.0019	0.0013	0.092
Children aged 0–5	0.0158	0.0093	0.088	0.0001	0.0003	0.741
Children aged 6–11	0.0704	0.0086	0.000	0.0012	0.0005	0.000
Children aged 12–14	0.1227	0.0117	0.000	0.0014	0.0006	0.000
Children aged 14–16	0.0969	0.0182	0.000	−0.0003	0.0007	0.703
Proportion of dependents in household	0.0093	0.0306	0.761	−0.0007	0.0014	0.612
Presence of disabled person	0.0400	0.0284	0.150	0.0004	0.0012	0.730
Household labor						
Number of day workers	0.0264	0.0116	0.023	0.0009	0.0007	0.189
Number of salaried workers	−0.0308	0.0091	0.001	0.0001	0.0003	0.857
Number of self-employed workers	0.0788	0.0101	0.000	0.0005	0.0004	0.236
Number of family workers	0.0076	0.0099	0.446	0.0009	0.0005	0.049
Access to social security				−0.0008	0.0006	0.138
Dwelling characteristics						
Dwellers per room	0.0049	0.0048	0.305	−0.0001	0.0002	0.680
Earthen floor	0.0421	0.0164	0.009	0.0038	0.0024	0.003
Piped water	−0.0437	0.0154	0.004	−0.0009	0.0011	0.330
Latrine	−0.0169	0.0154	0.268	−0.0009	0.0014	0.453
Electricity	0.1313	0.0177	0.000	0.0009	0.0009	0.465
Roof				−0.0001	0.0006	0.847
Walls				−0.0006	0.0014	0.595
Refrigerator	−0.0772	0.0159	0.000	−0.0011	0.0009	0.109
Stove	−0.1336	0.0180	0.000	−0.0023	0.0017	0.042
Washing machine				0.0004	0.0005	0.493
Car				−0.0004	0.0007	0.516
Video recorder				−0.0014	0.0008	0.031
Water heater				−0.0013	0.0008	0.072
Telephone				−0.0008	0.0007	0.225
Cable television				0.0115	0.0136	0.078
Blender	0.0715	0.0166	0.000			
Radio	−0.0147	0.0152	0.328			
Television	0.0253	0.0173	0.139			
School attendance						
Children aged 6–11 not attending school	−0.2123	0.0223	0.000			
Children aged 12–14 not attending school	−0.1184	0.0244	0.000	−0.0008	0.0007	0.361
Children aged 15–16 not attending school	−0.0729	0.0246	0.005	−0.0010	0.0006	0.223
Intercept	−0.0001	0.0000	0.000	0.0000	0.0000	0.007
Summary statistics						
Observed p		0.3545091			0.0110358	
Predicted p (at x bar)		0.2958774			0.0016715	
No. observations		6753			10330	
Likelihood ratio chi squared (28)		2083.75			287.79	
Probability > chi squared		0				
Pseudo R squared		0.2373			0.2295	

Source: Authors' calculations, based on data from the 2002 ENIGH.

a. The dependent variable is a dummy that equals 1 if a household is a beneficiary of Oportunidades and 0 otherwise. The regression model is a probit estimation using Stata8 dprobit instruction.

by the usual administrative problems, the probit is capturing both the selection mechanisms and the problems of finding beneficiaries in the field.[14] We run the probit model on the 2002 ENIGH database that was used as an anchor for the microsimulation exercise.[15]

The estimation for rural areas shows that households with children under sixteen years of age and without basic services (such as piped water or a finished floor) and household appliances (such as refrigerators) are most likely to be beneficiaries. The household head's education level and the number of salaried family members (that is, proxies for family income) diminish the probability that a rural household will be a beneficiary. The estimation for urban areas shows similar results: households with children aged twelve to sixteen and without basic services such as piped water or a finished floor are most likely to be beneficiaries.

Most of the coefficients have the expected signs, in both urban and rural areas. The model for the rural area has a good forecasting ability (the predicted probability at the mean was 0.296, which is close to the actual probability of 0.353). The urban model, however, has a low forecasting ability (the predicted probability at the mean was 0.001, while the actual probability is 0.011). The model remains an imperfect way to identify urban beneficiaries because of the incipient implementation of Oportunidades in urban areas in 2002. Nevertheless, the results for both models are quite similar, and we keep the urban model despite its limitations, because no other database is available.

An Accounting Exercise

The microsimulation program helps us address questions on what would happen to poverty in Mexico if the selection of beneficiaries or the size of benefits changed. We focus on three questions, out of the many that could be postulated. What would have happened to poverty if the program had been cancelled? What would happen to poverty if transfers were doubled? What would happen to poverty if the number of urban beneficiaries were doubled?

14. For a description of the selection procedures, see "Así se contruye Oportunidades: Informe 2002," available online at www.oportunidades.gob.mx/htmls/informe_2002/sitiofinal/index.html (accessed September 2004). More formally, the procedures for selecting beneficiaries and distributing benefits are closely detailed in several issues of the *Diario Oficial.* See, for instance, SEDESOL (2003).

15. The original microsimulation tool used probit models run on other databases: namely, the Encuesta de Características Socioeconómicas de los Hogares (ENCASEH) and the Encuesta de Características Socioeconómicas de los Hogares Urbanos (ENCASURB).

These three questions allow us to gauge the actual and potential impact of the program on short-term poverty.

Table 2 summarizes the results of the accounting simulations. The first column simply states the conditions of the program in 2002, for national, urban, and rural areas, as reported by our calibration of the 2002 ENIGH. This is the default, or the official levels of poverty and beneficiaries according to SEDESOL. The second column shows that abolishing Oportunidades would be associated with a rise in national poverty of two percentage points (from 20.24 to 22.48 percent in the headcount index). The effect would be stronger among rural households, where the poverty headcount index would rise from 34.71 to 39.96 percent, whereas urban households would show a slight increase from 11.43 to 11.83 percent. Comparing these figures with the official poverty levels between 1998 and 2002 indicates that the program accounts for nearly a third of the reduction in rural poverty. In fact, rural poverty declined by 17.4 percentage points. This means that, according to our simulation, 5.25 points of the reduction (around 28 percent) could be ascribed to Oportunidades. The effect on urban poverty is negligible since Oportunidades only started operations in urban areas in 2002.[16] These results should be considered an upper bound, for two reasons. First, accounting simulations do not consider behavioral effects. If Oportunidades did not exist, for example, households' labor supply might well be larger, and family income (net of transfers) could thus be correspondingly higher. Second, we did not calibrate other income sources. A large chunk of individuals in the sample did not report being beneficiaries, and their cash transfer income was adjusted as part of the calibration exercise explained above. Other income sources may also be misreported, but they have not been adjusted. We thus assume that only changes in cash transfers affect changes in total income, which may overestimate the total impact of Oportunidades.

Doubling benefits would have a strong impact on rural poverty and almost no effect on urban poverty (see column 3 of table 2). Rural and urban headcount indexes would fall to 27.22 percent and 10.69 percent, respectively. Doubling the number of beneficiaries in urban areas, while keeping benefits

16. Official figures for poverty in Mexico are taken from Cortés Cáceres and others (2002). That report states that the extreme poverty headcount in 1992 was situated at 0.135 for urban areas, 0.356 for rural areas, and 0.225 at the national level. It peaked in 1996 at 0.265 (urban), 0.524 (rural), and 0.371 (national). In 1998 the figures were 0.213 (urban), 0.521 (rural), and 0.339 (national); they then fell to their 2002 levels of 0.114 (urban), 0.347 (rural), and 0.203 (national). Finally, preliminary figures for year 2004, as reported by the Comité Técnico para la Medición de la Pobreza en México, are 0.110 (urban), 0.276 (rural), and 0.173 (national).

TABLE 2. Poverty Indexes for Accounting Simulations[a]

			Selection Criteria			
			Current selection (Probit)		*Perfect targeting*	
Sample group and index	*Default (2002 ENIGH)*	*Cancellation of Oportunidades*	*Doubling benefits*	*Doubling urban beneficiaries*	*Current benefits*	*Doubling benefits*
National						
FGT(0)	0.2024	0.2248	0.1695	0.1983	0.1755	0.1300
	(0.0056)	(0.0058)	(0.0052)	(0.0056)	(0.0121)	(0.0090)
FGT(1)	0.0627	0.0847	0.0451	0.0612	0.0548	0.0346
	(0.0021)	(0.0028)	(0.0016)	(0.0021)	(0.0043)	(0.0025)
No. households	4,238,277	—	4,238,277	4,747,393	4,318,328	4,318,328
Budget (pesos)	2,074,221,793	—	4,148,443,587	2,353,282,073	1,861,364,407	3,722,728,813
Urban						
FGT(0)	0.1143	0.1183	0.1069	0.1077	0.0868	0.0708
	(0.0055)	(0.0056)	(0.0052)	(0.0054)	(0.0101)	(0.0086)
FGT(1)	0.0278	0.0321	0.0245	0.0254	0.0213	0.0160
	(0.0017)	(0.0019)	(0.0015)	(0.0016)	(0.0026)	(0.0020)
No. households	507,652	—	507,652	1,016,768	1,371,907	1,371,907
Budget (pesos)	289,948,484	—	579,896,968	569,008,764	468,964,138	937,928,277
Rural						
FGT(0)	0.3471	0.3996	0.2722	0.3471	0.3211	0.2273
	(0.0111)	(0.0114)	(0.0104)	(0.0111)	(0.0220)	(0.0171)
FGT(1)	0.1200	0.1710	0.0788	0.1200	0.1098	0.0652
	(0.0046)	(0.0063)	(0.0034)	(0.0046)	(0.0087)	(0.0052)
No. households	3,730,625	—	3,730,625	3,730,625	2,946,421	2,946,421
Budget (pesos)	1,784,273,309	—	3,568,546,618	1,784,273,309	1,392,400,268	2,784,800,537

Source: Authors' calculations, based on data from the 2002 ENIGH.

a. Poverty indexes (Foster, Greer, and Thorbecke 1984) and their standard errors (in parentheses) were calculated using Stata sepov instruction.

the same, would also have no effect on urban poverty (see column 4). Given that the urban poverty gap is very small (nearly 3 percent, in columns 2 and 3), the first exercise suggests that most selected beneficiaries are not poor.

These results indicate that the selection of beneficiaries in urban areas may be seriously mistargeted in the simulation exercise. We identify three possible reasons for this, two ascribed to our modeling and one to the program itself. First, the urban program began in 2002, so the database may not contain enough information to identify urban beneficiaries. Second, the econometric model used to identify beneficiaries (see table 1) could be affected by other econometric problems, resulting in a failure to adequately identify potential beneficiaries. Finally, the program administrators may have problems targeting beneficiaries in the field as a result of the self-selection explained above.[17]

In rural areas, Oportunidades served 40 percent of families in 2002. Despite this, poverty still hovered at about one-third of the rural population. In addition, the rural poverty gap is 12 percent of the poverty line, which means that some families may receive cash transfers and remain in poverty. The average monthly cash transfer among rural beneficiaries is 478 pesos per family, while the rural poverty line is 494 pesos per head. The average transfer thus represents 19 percent of the poverty line (assuming an average five-member family), which makes the hypothesis of still-poor beneficiaries reasonably likely for those in the bottom of the distribution. We further posit that the process of targeting beneficiaries continues to be problematic. Families in communities with a very high marginality index and poor families in localities with a low marginality index may still be excluded from the program.[18]

Given these drawbacks, we undertake our simulations assuming perfect targeting. That is, we use the alternative procedure explained earlier, in which every household under the poverty line is assumed to be a beneficiary of Oportunidades, no matter its demographic composition or location. Admittedly, this is not the objective of the program, which aims to promote human capital accumulation. However, the purpose of the simulation exercise is to ascertain the program's impact on short-term poverty. Column 5 of table 2 shows the

17. Urban beneficiaries actually increased by 35 percent (nearly 174,000 urban families) from 2002 to 2004, but official records show only a very slight reduction in urban poverty (from 11.4 to 11.0 percent).

18. According to the operating rules, households are selected after localities are targeted based on their marginality index (see, for instance, SEDESOL 2003). The marginality index is computed using a set of economic and human development statistics for the locality. For a description of the marginality index in Mexico, see CONAPO (2001).

simulation results for an exercise of prefect targeting on the poor, maintaining current benefits. In this case, rural and urban poverty would have fallen to 32.11 percent and 8.68 percent, respectively. That is, poverty would have declined around 2.6 percentage points more in both areas than it did under the current program, although the difference is statistically significant only for the urban sector (compare columns 5 and 1). Doubling benefits under perfect targeting would have further reduced urban poverty to 7.08 percent and rural poverty to 22.73 percent; both figures are significantly different from the 2002 default level (see columns 6 and 1). This exercise suggests that better targeting is necessary for reducing urban poverty, whereas larger transfers are needed for reducing rural poverty.

These exercises also allow us to estimate the cost effectiveness of the program. The average cash transfer cost for each percentage point reduction in the national poverty headcount in 2002 was around 926 million pesos a month.[19] This amount is equivalent to 1.1 percent of the central government's total expenditures in 2002 and 2.1 percent of the central government's social spending (that is, expenditures on health, education, and social protection). This average hides important differences between urban and rural sectors. One percentage point of rural poverty reduction costs 339 million pesos a month, but urban poverty reduction costs 710 million pesos. The latter figure reflects the almost null effect of the program in urban areas in 2002. Doubling the benefits to current beneficiaries would reduce the national average cost to 750 million pesos a month, basically as a result of a reduction in the average cost for rural areas to 280 million pesos a month.

Since the urban numbers may be affected by problems in selecting urban beneficiaries, the cost effectiveness derived from perfect-targeting simulations could be informative. Perfect targeting would have a cost of 380 million pesos a month per percentage point reduction of the national poverty headcount. This is due to a reduction of average monthly costs to 148 million and 178 million pesos in urban and rural areas, respectively. Additionally, the difference between the poverty headcounts for the perfect-targeting simulation and the official figures for 2002 is statistically significant for the urban sector, but not for the rural sector. Targeting thus does not appear to be a serious

19. Cost is computed by dividing the total budget for transfers by the difference between the poverty headcount considered and the headcount simulated under termination of Oportunidades. For example, from table 2 the national budget for transfers in 2002 was 2.074 billion pesos a month. This figure divided by the difference between the poverty headcount under the simulated termination of the program (that is, 0.2248) and the official figure for year 2002 (that is, 0.2024) renders a ratio of 925.9 million pesos a month.

problem in rural areas. Improved targeting of urban beneficiaries, however, reduces poverty in this area and improves the cost effectiveness of the whole program by more than 50 percent. This conclusion calls for additional efforts in targeting beneficiaries in urban areas.[20]

Doubling benefits under perfect targeting would reduce average costs to 162 million pesos a month in rural areas, but it would increase average monthly costs to 197 million pesos in urban areas. This stems from the much larger additional fall in rural than in urban poverty: 9.4 and 1.6 percentage points, respectively. This result is similar to the simulation of doubling benefits under the current selection mechanisms. Hence, expanding benefits would further reduce national poverty by generating a large reduction in rural poverty.

The difference in the impact of expanding benefits on rural and urban poverty reflects the different demographic composition and official poverty lines of the two sectors. Consequently, the relative size of transfers differs between rural and urban households. The average transfer under perfect targeting and doubled benefits would be 945 pesos for rural beneficiaries versus 684 pesos for urban households. These transfers represent 38 percent and 20 per-cent, respectively, of the poverty line for a five-member family. Despite this smaller relative transfer, a larger share of poor urban households leaves poverty after targeting than do poor rural households because of the smaller urban poverty gap: 26 percent of the urban poor rise out of poverty, versus only 19 percent of the rural poor. Doubling benefits has less of an effect in urban areas, since many urban households are already out of poverty, but it helps many rural households that were left behind: a further 24 percent of rural households (but only 14 percent of urban) leave poverty after benefits are expanded. In sum, targeting would be very effective in helping the urban poor because these households are closer to the poverty line; once identified, they cross the poverty line sooner than rural households, even though their transfers are smaller, on average. Expanding benefits, however, is more helpful for the rural poor because they have a larger poverty gap than the urban poor and because there are more rural than urban poor.

Perfect targeting and doubling benefits would have a cost-effectiveness ratio of 393 million pesos a month (that is, 0.5 percent of the central govern-

20. While the average cost in transfers would be smaller, we do not consider the administrative costs of perfect targeting. A full cost-effectiveness analysis of the program therefore cannot be derived from this exercise. Nevertheless, recent studies report that the administrative and operative costs of the program have been stable at around 3 percent and 6 percent, respectively, of total expenditures for several years (see Meneses González and others 2005). Whether further efforts for targeting would represent the same proportion of costs is an open question.

ment's total expenditures in 2002). This is less than half the ratio for 2002, but its total cost would also double to 3.7 billion pesos a month (that is, 4.4 percent of the central government's total expenditures in 2002).

Econometric Analysis of Labor Supply

The above results refer to accounting simulation exercises. This section explains the econometric models of labor supply that support the behavioral simulations. Labor supply in Mexico is difficult to model for several reasons: The demographic structure of households is diverse, the rural sector is large, the urban market is dualistic, and the distribution of hours is disperse and bimodal. This heterogeneity makes it difficult to propose a one-size-fits-all model of Mexican labor supply. The manner in which we address these difficulties, which is by no means perfect, is to estimate several models separately, and explain the implications for the microsimulation exercise.

Mexican households are very heterogeneous, ranging from the "conventional" family with two parents and their children to nontraditional families including singles, singles with children, and extended families. Conventional families represent 51.11 percent of Mexican households; this pattern is observed both in the rural and urban areas. Most households have between two and five members, of whom one or two are active participants in the labor market. More than half of all households have two or more members in the labor market.

These data suggest that Mexican families often rely on earnings from multiple household members for sustenance, so a model of joint labor supply would be ideal. The literature on developed countries includes several examples of these models.[21] However, all these models refer to the case of two-parent, or two-adult, households, whereas the Mexican labor market has a large share of households with two or more active members. In many Mexican households minors also work, making a model of joint labor supply for two adults less applicable as a general representation.[22] Estimating a model of

21. See, for example, van Soest (1995); Hoynes (1996); Keane and Moffitt (1998); Blundell and others (2000). For the case of developing countries, see Pradhan and van Soest (1997). For the case of México, see Hernández Licona (1997).

22. Freije and López-Calva (2001) report that around 10 percent of children aged fourteen to sixteen and 4 percent of those aged twelve to thirteen worked for a wage in Mexican urban areas in the late 1990s.

joint labor supply with more than two participants, whether adults or minors, is a daunting task. We know of no such studies to date. For simplicity, we therefore estimate separate individual labor supply models for men, women, and minors.[23]

In terms of geographic distribution, 38 percent of the Mexican population and 36.6 percent of the labor force live in rural areas.[24] Rural areas thus represent a sizeable component of the Mexican labor market and require special and separate attention.[25] The nonagricultural wages of individuals who are partially or not at all occupied in agricultural activities provide an adequate representation of marginal productivity and can be used as an explanatory variable for labor supply models. Many rural individuals, however, work only on farms, often for subsistence, and their wages are not observable. To address this problem, we need to estimate shadow wages based on models of the agricultural production function.[26] Estimating these models requires data on farm inputs and outputs. The ENIGH has a farm production module, but its data are incomplete and unreliable. We therefore forgo estimating labor supply for rural workers who are fully engaged in agricultural production and instead estimate models for the rest of the rural population. As a result, our estimation of labor supply models leaves out around 40 percent of the adult and 50 percent of the minor labor force in rural areas. As discussed, we estimate separate labor supply functions for urban and rural areas.

Another important aspect of the Mexican labor market is its dualism. Around half of the urban labor force does not contribute to social security programs (either public or private). Roughly one-fifth of the adult labor force works independently, that is, is self-employed, and one-quarter of the minor labor force works without payment. These data imply that a large share of the working population is characterized by low productivity and lack of social protection. We do not deal with this issue explicitly, however.[27] We simply

23. For the rest of the paper, we use the term minors to refer to individuals aged twelve to seventeen. Although there is evidence of working children under the age of twelve, the database we used does not allow us to identify the labor status of these children.

24. The National Institute of Statistics, Geography, and Information (INEGI) defines rural areas as localities with fewer than 15.000 inhabitants.

25. Behrman (1999) highlights the importance of rural labor markets in developing countries and calls for a separate analysis of their behavior.

26. Skoufias (1994) and Jacoby (1993) develop the methodology for estimating labor supply models of agricultural workers. DeJanvry and Sadoulet (2001) and Yúnez-Naude and Taylor (2001) make investigations on this subject for rural Mexico.

27. The wage models mentioned in footnote 10 include controls for working in the covered and noncovered (formal and informal) sectors.

exclude unpaid workers (minors and adults) when estimating labor supply models.

The Mexican labor market is also characterized by long working hours. The ENIGH reports that 18 percent of minors, 21 percent of females, and 39 percent of urban males in the labor force work more than 48 hours a week. In the rural sector, these figures are 24 percent, 29 percent, and 50 percent, respectively. However, part-time employment is an important chunk of labor supply, particularly among minors and females. Given this distribution of hours of work, we chose to estimate a categorical model of labor supply.[28] We experimented with several categories and several cut-off points for defining each category. We also tried simple probit, participation-corrected probit, and multinomial logit models. We ultimately chose the multinomial logit because it had a good forecasting capacity and allowed for different levels of work hours. Table 3 shows the percentage of accurately predicted observations for each model.[29]

The multinomial model starts from a linear random utility function,

$$U_{ij}^{A} = \alpha \mathbf{I}_{ij}^{A} + \beta \mathbf{H}_{ij}^{A} + \delta \mathbf{Y}_{ij}^{A} + \lambda t_{ij}^{A} + \gamma w_{ij}^{A} + \varepsilon_{ij}^{A},$$

and a latent variable structure,

$$z_i^{A} = j, \qquad \text{if } U_{ij}^{A} = \max\left(U_{i0}^{A}, U_{i1}^{A}, U_{i2}^{A}, U_{i3}^{A}\right),$$

where $A = 1, 2, 3$ and $j = 0, 1, 2, 3$ and where the index A denotes the group (that is, adult man, adult woman, or minor), the index j denotes labor force participation (inactive, part-time, full-time, or overtime), vector $\mathbf{I}$ includes individual characteristics, vector $\mathbf{H}$ represents household characteristics, vector $\mathbf{Y}$ encompasses family income sources per head (rest of the family salaries and other income), w_{ij} is individual wages, t_{ij} is individual transfers (from Oportunidades and other sources), and ε_{ij} comprises unobservables affecting preferences. The same model is estimated separately for urban and rural areas, so six different types of individuals are modeled; to facilitate notation, we drop the super index A in what follows.

28. Alternatively, we could have estimated a continuous hours model of labor supply, as described by Creedy and Duncan (2002) or Blundell and McCurdy (1999). However, we preferred a categorical model more in line with the microsimulation exercise carried out by Bourguignon, Ferreira, and Leite (2003) because of its simplicity and applicability to our subject.

29. Full results for the models are available on request from the authors.

TABLE 3. Forecasting Ability of Different Econometric Models for Mexican Labor Supply
Percentage of accurate predictions

		Urban		
Model and sample group	*Rural*[a]	*Total*	*Employed workers*	*Independent workers and employers*
Multinomial Logit				
Men	66.02	65.01	70.29	73.49
Women	82.47	83.86	80.98	89.52
Minors	91.24	95.19	91.03	98.55
Probit				
Men	80.45	83.59	—	—
Women	74.05	65.03	—	—
Minors	88.98	86.60	—	—
Probit with sample selection				
Men	47.17	43.28	—	—
Women	13.88	59.46	—	—
Minors	53.18	83.57	—	—

Source: Authors' calculations, based on data from the 2002 ENIGH.
a. Excluding agricultural workers.

If ε_{ij} is independently distributed as a type I extreme value distribution, then

$$\text{prob}\left(U_{ij} > U_{ik}\right) = \text{prob}\left(z_i = j\right) = \frac{\exp\left(\alpha\mathbf{I}_{ij} + \beta\mathbf{H}_{ij} + \delta\mathbf{Y}_{ij} + \gamma w_{ij} + \lambda t_{ij}\right)}{\sum_{j=0}^{3}\exp\left(\alpha\mathbf{I}_{ij} + \beta\mathbf{H}_{ij} + \delta\mathbf{Y}_{ij} + \gamma w_{ij} + \lambda t_{ij}\right)},$$

where $k \neq j$.[30] When no data are available on the characteristics of different alternatives and all the variation is across individuals, this model (known as the conditional logit) can be recast as the multinomial logit:

$$\text{prob}\left(z_i = j\right) = \begin{cases} \dfrac{1}{1+\sum_{j=1}^{3}\exp\left(\alpha_j\mathbf{I}_i + \beta_j\mathbf{H}_i + \delta_j\mathbf{Y}_i + \gamma_j w_i + \lambda_j t_i\right)}, & j = 0 \\[2ex] \dfrac{\exp\left(\alpha_j\mathbf{I}_i + \beta_j\mathbf{H}_i + \delta_j\mathbf{Y}_i + \gamma_j w_i + \lambda_j t_i\right)}{1+\sum_{j=1}^{3}\exp\left(\alpha_j\mathbf{I}_i + \beta_j\mathbf{H}_i + \delta_j\mathbf{Y}_i + \gamma_j w_i + \lambda_j t_i\right)}, & j = 1,2,3 \end{cases}$$

30. See McFadden (1974).

After estimating the multinomial logit model, we can get three vectors of estimates: $\{\hat{\alpha}_j, \hat{\beta}_j, \hat{\delta}_j, \hat{\gamma}_j, \hat{\lambda}_j\}$, where $j = 1, 2, 3$. These estimates allow us to compute fitted probabilities for each labor participation category after changes in Oportunidades benefits (Δt_i) and assume the individual will choose the one with the highest simulated utility. Formally:

$$\text{given} \quad \hat{U}_{ij} = \hat{\alpha}_j \mathbf{I}_i + \hat{\beta}_j \mathbf{H}_i + \hat{\delta}_j \mathbf{Y}_i + \hat{\lambda}\left(t_i + \Delta t_i\right) + \hat{\gamma}_j w_i + \hat{\varepsilon}_{ij},$$

$$\text{then} \quad \hat{z}_i = j \quad if \quad \hat{U}_{ij} > \hat{U}_{ik}, \qquad \forall k \neq j,$$

where, in addition, the simulated residuals are randomly selected so that they comply with the following conditions:

$$\hat{\varepsilon}_{i,j} = -\ln\left[-\ln\left(u_{i,j}\right)\right], u \sim U\left(0,1\right), \text{ and}$$

$$\hat{U}_{i,j} > \hat{U}_{i,k}, \qquad \forall k \neq j, \qquad \text{if} \quad z_i = j.$$ [31]

The results for the six multinomial logit models, corresponding to the odds ratios for family income derived from transfers, are presented in table 4. Individual and household characteristics, as well as different per capita family income sources, are the explanatory variables. The specification is the same for every model, with the exception of minors, who have a slightly different set of explanatory variables.[32] The main concern with multinomial logit models is the assumption of independence of irrelevant alternatives. We applied a Hausmann test to all six models and failed to reject the null hypothesis in three of them (rural men, rural minors, and urban minors). In the other three, the null hypothesis was rejected in some, but not all, possible exclusion restrictions.

Since the estimated parameters from a multinomial logit do not have a direct interpretation, coefficients are always presented as odds ratios with respect to the comparison case (being out of the labor force). For instance, in the case of transfers from Oportunidades, the odds ratios are as follows:

31. That is, residuals are such that they are distributed following an extreme value type I (also known as Gumbel distribution) and the ranking of original simulated utilities agrees with the observed labor supply. We thank Phillippe Leite for providing us with a Stata program that produces randomly selected residuals according to these criteria. The original microsimulation tool did not make use of these calibrated residuals. See IPD (2005).

32. Full econometric results are available from the authors. Also see IPD (2005).

TABLE 4. Multinomial Logit Results for Transfers: Odds Ratios[a]

	Part-time (under 20 hours)			Full-time (20–48 hours)			Overtime (over 48 hours)		
Sample group	*Coefficient*	*Z statistic*	P > |z|	*Coefficient*	*Z statistic*	P > |z|	*Coefficient*	*Z statistic*	P > |z|
Rural									
Men	0.696	1.239	0.215	0.934	0.183	0.855	0.686	1.234	0.217
Women	0.792	1.419	0.156	0.637	2.307	0.021	0.693	1.827	0.068
Minors	0.217	2.138	0.032	0.223	2.041	0.041	0.254	1.669	0.095
Urban									
Men	0.932	0.621	0.534	0.705	4.053	0.000	0.594	5.155	0.000
Women	0.798	2.953	0.003	0.731	3.211	0.001	0.783	2.197	0.028
Minors	0.982	0.053	0.958	0.493	1.107	0.268	0.978	0.032	0.975

Source: Authors' calculations, based on data from the 2002 ENIGH.

a. The table presents results from a multinomial logit estimation, using Stata8 mlogit instruction, which includes marital status, age, education, other sources of family income (such as individual wages), household size, and location as explanatory variables. Transfers are defined as monetary transfers received by all family members and are measured in per capita thousands of pesos. Odds ratios are relative to inactive individuals, defined as all those who are unwilling or unable to work and are thus unemployed. Full econometric results are available on request.

$$\frac{P\left(z_i = j \middle| \mathbf{I}, \mathbf{H}, \mathbf{Y}, w, t\right)}{P\left(z_i = 0 \middle| \mathbf{I}, \mathbf{H}, \mathbf{Y}, w, t\right)} = \exp\left(\alpha_j \mathbf{I}_i + \beta_j \mathbf{H}_i + \delta_j \mathbf{Y}_i + \gamma_j w_i + \lambda_j t_i\right).$$

Then, assuming $\Delta t_i = 1$, the estimated odds ratios for categories $j = 1, 2, 3$ are

$$e^{\hat{\lambda}_j} = \frac{\exp\left[\hat{\alpha}_j \mathbf{I}_i + \hat{\beta}_j \mathbf{H}_i + \hat{\delta}_j \mathbf{Y}_i + \hat{\gamma}_j w_i + \hat{\lambda}_j \left(t_i + \Delta t_i\right)\right]}{\exp\left(\hat{\alpha}_j \mathbf{I}_i + \hat{\beta}_j \mathbf{H}_i + \hat{\delta}_j \mathbf{Y}_i + \hat{\gamma}_j w_i + \hat{\lambda}_j t_i\right)}$$

$$= \frac{\hat{P}\left(z_i = j \middle| t = t_i + \Delta t_i\right) \Big/ \hat{P}\left(z_i = 0 \middle| t = t_i + \Delta t_i\right)}{\hat{P}\left(z_i = j \middle| t = t_i\right) \Big/ \hat{P}\left(z_i = 0 \middle| t = t_i\right)}.$$

The effects of individual and household characteristics on labor supply display several regularities among all groups. Being single reduces the likelihood that both rural and urban men will work full-time or overtime, relative to being inactive, whereas women show the opposite effect. Age very slightly reduces the likelihood of working full-time and overtime for men and women, in both urban and rural areas, whereas minors show the opposite effect. The effect of education is less clear. Several coefficients are insignificant, and no marked patterns appear. However, very low levels of education (that is, no formal schooling and incomplete primary education) are often associated with odds ratios below unity for males and females in urban and rural areas. In other words, individuals with little or no education are more likely to be inactive than employed. Higher education, whether complete or incomplete, presents odds ratios above unity for urban males and females and for rural females. In the case of rural and urban minors, school attendance has odds ratios well below unity for full-time and overtime employment, whereas the odds ratios are not significantly different from unity for part-time work. This indicates that school attendance clearly excludes minors from full-time work, but not from part-time work. The odds ratio for number of children per household is usually above one, but often insignificant. This odds ratio is significantly greater than unity for urban women working part-time and for rural women working overtime. This means that family size increases the relative likelihood of part-time work among urban women and of overtime hours among rural women. Finally, regional dummies show no clear pattern, either. Some are individually significant and others are not, but they are jointly significant, indicating that geographical factors influence the Mexican labor supply.

The most important coefficient for our purposes is the one corresponding to transfers. We use this partial effect to simulate the behavioral response to simulated changes in Oportunidades benefits. The odds ratios for this variable are always below unity, often significant, and, in some cases, monotonically decreasing from part-time to overtime work. This means that transfers are associated with a lower relative likelihood of being employed and often also with a decreasing relative likelihood of working additional hours. This is consistent with theoretical predictions of a pure income effect on labor supply. The clearest example is the case of urban men (see table 4). The odds ratio declines from 0.932 for part-time work (although it is not significantly different from one) to 0.705 for full-time employment and 0.594 for overtime hours (both of which are significant). This means that when a discrete change in transfers occurs, urban men are relatively less likely to work part-time (as compared with being inactive), relatively less likely to work full-time (as compared with part-time), and less likely to work overtime than full-time. The pattern for other groups is not as neat, however: odds ratios for rural men and urban minors are not significant, and they are significant but not monotonic for rural and urban women.

Although indicative, the odds ratios do not show the size and trend of the partial effect of the explanatory variable on the dependent variable. The partial effect of Oportunidades transfers in our multinomial logit model is as follows:

TABLE 5. Marginal Effect of Transfers on Predicted Hours Category[a]

Sample group	*Part-time*	*Full-time*	*Overtime*	*Out of labor force*
Urban				
Men	0.0089	0.0339	−0.0428	2.21E−15
Women	0.0106	−0.0180	0.0075	1.72E−10
Minors	0.1310	−0.1623	0.0313	7.52E−09
Rural				
Men	−0.0027	0.0771	−0.0743	5.59E−20
Women	0.0373	−0.0387	0.0014	6.53E−07
Minors	−0.0111	−0.0008	0.0119	2.15E−05

Source: Authors' calculations, using the models in table 4.

a. Marginal effect is computed on the average individual (that is, following equation 3 from the text).

$$(3) \qquad \frac{\partial P\left(z = j \middle| \mathbf{I}, \mathbf{H}, \mathbf{Y}, w, t\right)}{\partial t} = P\left(z = j \middle| \mathbf{I}, \mathbf{H}, \mathbf{Y}, w, t\right) \left[\lambda_j - \sum_{j=1}^{J} \lambda_j P\left(z = j \middle| \mathbf{I}, \mathbf{H}, \mathbf{Y}, w, t\right)\right],$$

which depends on the level of transfers and the level of all the other explanatory variables. These marginal effects are usually computed for the average individual. That is, fitted probabilities are computed evaluating the control variables at their means. Table 5 shows the marginal effects for each work category for the six groups. Oportunidades transfers induce a reduction in the probability of working overtime, in favor of full-time, for urban and rural men, and a reduction in the probability of working full-time, in favor of part-time (and even some overtime), for all women and urban minors. Rural minors show a pattern of lower probability of full- and part-time employment in favor of inactivity and overtime.

These results are not fully satisfactory, however, because they refer to the partial effect on the average individual. The average individual may not be representative of all groups in the sample or may not even exist.[33] Besides, even if we accept the average individual, the marginal effect of transfers at the average transfer might not be of interest because this effect may vary according to the transfer level. Consequently, for the microsimulation exercise, we need to know how the distribution of hour categories for the total labor supply varies with different transfers.

33. For instance, the individual whose marital status is the share of singles in the sample (say, 0.32, instead of the zero-or-one dummy values) does not exist.

The effect of transfers on labor supply can be illustrated by the method of recycled predictions, which consists in varying the characteristic of interest (namely, Oportunidades transfers) across the whole dataset and then taking the average of predictions across all observations. Formally, for each work category, $j = 0, 1, 2, 3$, we have

$$S\left(j, \Delta t\right) = \frac{1}{n}\sum_{i}^{n} \hat{P}\left(z_i = j \middle| I = \mathbf{I}_i, \mathbf{H} = \mathbf{H}_i, \mathbf{Y} = \mathbf{Y}_i, w = w_i, t = t_i + \Delta t\right), \quad (4)$$

where Δt refers to an equal change in transfers for every individual within the group under consideration and $S(j, \Delta t)$ refers to the share of the group that is forecast to be in category j if there is a Δt change in transfers. This is based on the fact that the average predicted frequency equals the actual frequency under a multinomial logit model.[34]

Figures 2 and 3 show simulations of the distribution of the urban labor force. The partial effects shown are compatible with the odds ratios and the marginal effects reported above, but they tell a more interesting story. In the case of urban men (see figure 2), the share of men working overtime decreases as transfers increase, while the share of those working part-time increases. The share of urban men working full-time first increases and then declines monotonically. The share of men who are out of the labor force shows a slight but steady increase. Surprisingly large transfers are required to obtain substantial effects. A family transfer of 2,000 pesos per capita per month is needed for the share of men working overtime to fall from 32 to 25 percent. The share of men working full-time does not decline until monthly transfers reach 5,000 pesos per family member. This indicates that plausible changes in transfers (for example, an increase of 100 pesos a month per family member) would have a negligible effect on the distribution of urban men's working hours. Urban women (see figure 3) display a similar pattern, but the distribution of hours of work is even less sensitive to changes in transfers (full-time work falls from 30 to 28 percent for the same increase in transfers).

In the case of rural men, the change in transfers would reduce the share of overtime workers from 37 to 26 percent, mostly in favor of full-time work, which rises from 36 to 46 percent. Rural women present a smaller reduction in full-time workers (from 17 to 14 percent), but the change is mostly reflected in a rise in inactivity (from 68 to 71 percent). As in the case of urban

34. See Maddala (1983, pp. 22–37).

FIGURE 2. Simulation of the Labor Force Distrubution for Urban Men

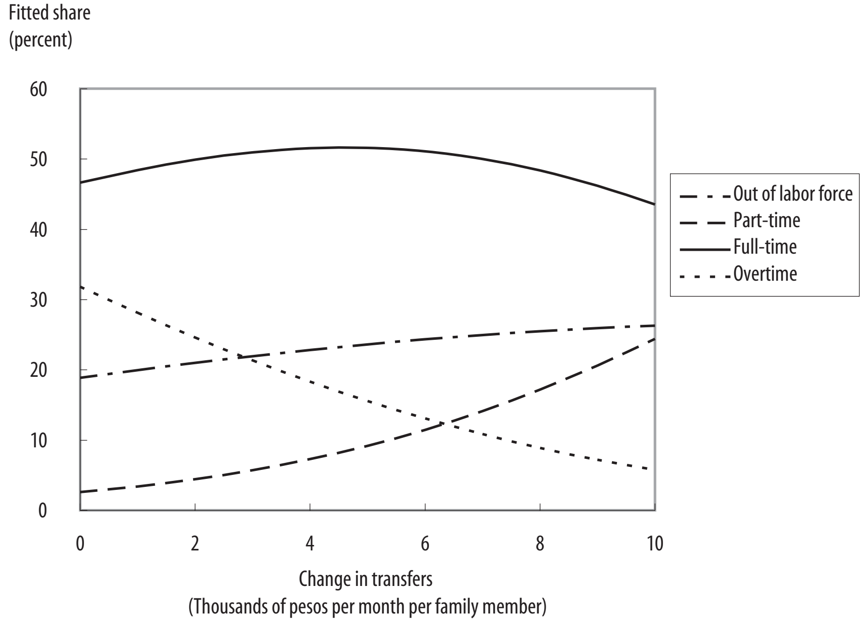

FIGURE 3. Simulation of the Labor Force Distrubution for Urban Women

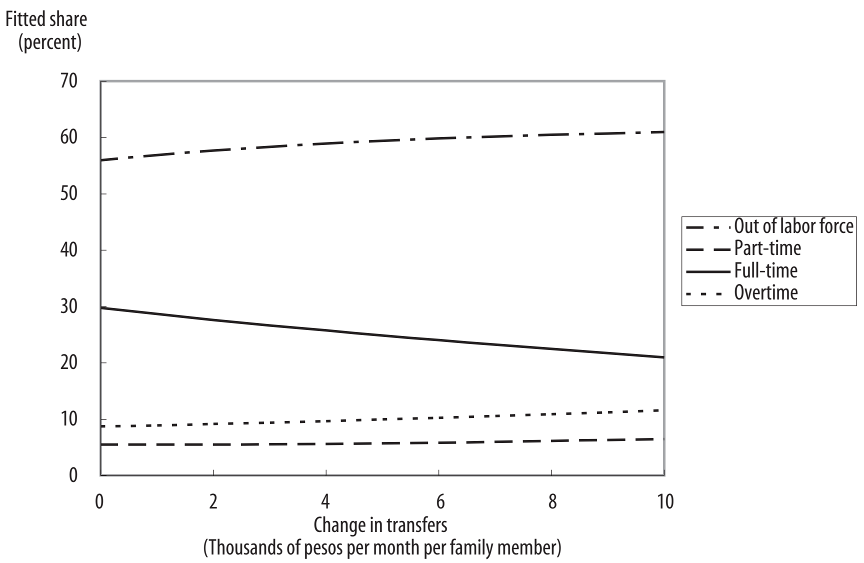

families, these results are for implausible transfers (namely, 2,000 pesos a month per family member). They suggest, however, that rural women are more likely to leave the labor force than rural men for a given change in transfers.[35]

In the case of rural minors, the hypothesized transfers would cause a 5 percent increase in inactivity, as 3 percent of part-time workers and 2 percent of full-time workers withdraw from the labor force. The results for urban minors are unexpected: while full-time work would fall 3.9 percent, overtime would increase by 2.5 percent and part-time work by 1.4 percent.[36]

All the above evidence suggests that Oportunidades transfers, at their current level, have almost no effect on labor supply. Plausible changes in these transfers would have negligible effects on labor supply. These exercises must be interpreted only as illustrations, however. As explained in equation 4, these distributions are simulated given changes in the whole sample, whereas the policy simulations of interest are such that changes in transfers correspond only to the selected beneficiaries. The microsimulation exercise resolves this issue by computing the effect on each individual, depending on whether he or she is a beneficiary, and then adding up the individual results to determine the effect on the whole population. Before we discuss the results of the behavioral simulations, we present the results from our estimation of a model of school attendance.

The School Attendance Model

As explained earlier, stage three of the behavioral simulation predicts labor supply for minors after their school attendance is simulated. The idea is that after the family receives a transfer, it decides, individual by individual, whether each minor attends school. The results of this simulation are then introduced in the labor supply model for minors as an explanatory variable.

The school attendance model returns very similar results for urban and rural minors (see table 6). The probability of school attendance declines

35. The study of the effect of transfers on adult labor supply is an ongoing research agenda in Mexico. Our results are compatible with Parker and Skoufias (2000), who find that only rural women change the use of their time when benefiting from Oportunidades. They show that changes occur in time allocated to household activities, but not in farm or labor market work. Cox-Edwards and Rodríguez-Oreggia (2006) also find a null effect of transfers from remittances on the labor supply. Amuedo-Dorantes and Pozo (2006), however, find no effect of remittances on male labor supply and a significant drop in female labor supply.

36. Figures showing the simulations of the labor force distribution for urban minors and for the rural population are available on request.

TABLE 6. Probit Model of School Attendance for Urban and Rural Minors

	Rural areas			*Urban areas*		
Explanatory variable	*Coef.*	*Std. error*	P > | z |	*Coef.*	*Std. error*	P > | z |
Individual characteristics						
Age	–0.059	0.044	0.178	–0.004	0.045	0.927
Man	–0.400	0.014	0.000	–0.390	0.015	0.000
Household characteristics						
Woman over 18	0.010	0.028	0.711	0.105	0.029	0.000
Minor	–0.007	0.013	0.564	–0.142	0.017	0.000
Household head characteristics						
Married	0.210	0.106	0.049	0.321	0.101	0.001
Male	–0.079	0.110	0.476	–0.215	0.103	0.037
Schooling						
No schooling	–0.452	0.067	0.000	–0.517	0.080	0.000
Incomplete primary	–0.249	0.064	0.000	–0.154	0.062	0.013
Complete primary (omitted)						
Incomplete secondary	0.380	0.166	0.022	0.228	0.102	0.025
Complete secondary	0.136	0.102	0.183	0.392	0.071	0.000
Incomplete high school	0.390	0.279	0.162	0.675	0.146	0.000
Complete high school	0.653	0.221	0.003	0.853	0.127	0.000
Tertiary	1.029	0.264	0.000	1.018	0.119	0.000
Employment						
Informal (omitted)						
Formal	0.094	0.079	0.235	0.162	0.053	0.002
Part-time	0.149	0.087	0.088	–0.134	0.106	0.205
Full-time (omitted)						
Overtime	–0.002	0.052	0.972	0.057	0.050	0.249
Family income[a]						
Labor income[b]	0.142	0.056	0.011	0.038	0.033	0.253
Transfers[c]	0.763	0.136	0.000	0.201	0.095	0.034
Nonmonetary transfers	–0.097	0.095	0.306	–0.057	0.070	0.416
Other income[d]	0.112	0.081	0.167	0.282	0.075	0.000
Constant	6.361	0.234	0.000	6.529	0.250	0.000

Source: Authors' calculations, based on data from the 2002 ENIGH.
a. In thousands of pesos per capita.
b. Includes all family wage incomes except the individual's own.
c. Total monetary transfers received by the family.
d. Dummy variable that equals 1 if the family has other sources of income and 0 otherwise.

with the age of the minor, but it increases with the household head's level of education, formal employment, and marriage. More important, an increase in transfers of 1,000 pesos a month per family member is associated with a 76 percent increase in the likelihood of participating in formal education in rural areas and a 20 percent increase in urban areas. This implies that transfers have a large and significant effect on school attendance in Mexico.

School attendance is also the most important determinant of child labor in the labor supply models.[37] For both rural and urban minors, the odds ratios for school attendance are significantly and monotonically well below unity for full-time and overtime work. This means that once the child is in school, the likelihood of working full-time relative to being inactive declines significantly; the effect is even stronger for overtime hours. The odds ratios are not significantly different from unity for part-time work, which means that school attendance does not conflict with minors working part-time.

A Behavioral Exercise

In this section, we replicate our earlier accounting experiments using a behavioral microsimulation exercise. That is, we simulate labor supply changes of individuals from households whose transfers were changed in the simulation exercises. Table 7 displays the results for these behavioral simulations, while tables 8 and 9 show the distribution of fitted work categories for the accounting and behavioral exercises, respectively. Given the small partial effect of transfers on labor supply reported earlier, we expect the changes in poverty under the behavioral simulations to be very similar to those calculated with accounting exercises.

The behavioral simulation of canceling Oportunidades renders the same headcount index in the urban area and a slightly higher, though not statistically different, index in the rural area (see column 2 of table 7). Doubling benefits under the current selection mechanism does not generate changes in labor supply relative to the fitted distribution of hours worked (compare columns 2 and 3 of table 9). The simulated poverty indexes are thus almost the same in the behavioral and accounting exercises.

These simulations are compatible with the results from the labor supply models, which show that the labor supply is insensitive to small changes in transfers. In fact, table 9 shows that labor supply has negligible changes in the rural sector and no changes at all in the urban sector. For instance, doubling the current average transfers per household from around 500 pesos a month to 1,000 pesos a month implies a change in per capita transfers of 100

37. The odd ratios for full-time and overtime work, relative to inactivity, are 0.249 and 0.156, respectively, for urban minors and 0.335 and 0.185 for rural minors. All are highly significant, which means that attending school significantly reduces the probability of working full-time or overtime (relative to being inactive). Full econometric results are available from the authors on request. See IPD (2005).

TABLE 7. Poverty Indexes for Behavioral Simulations[a]

			Selection criteria			
			Current selection (probit)		*Perfect targeting*	
Sample group and index	*Default (2002 ENIGH)*	*Cancellation of Oportunidades*	*Doubling benefits*	*Doubling urban beneficiaries*	*Current benefits*	*Doubling benefits*
National						
FGT(0)	0.2024	0.2267	0.1639	0.1982	0.1755	0.1300
	(0.0056)	(0.0058)	(0.0051)	(0.0056)	(0.0121)	(0.0090)
FGT(1)	0.0627	0.0889	0.0425	0.0610	0.0548	0.0346
	(0.0021)	(0.0030)	(0.0015)	(0.0021)	(0.0043)	(0.0025)
No. households	4,238,277	—	4,238,277	4,747,393	4,318,328	4,318,328
Budget (pesos)	2,402,635,619	—	4,805,271,238	2,705,284,826	1,861,364,407	3,722,728,813
Urban						
FGT(0)	0.1143	0.1184	0.1066	0.1075	0.0868	0.0708
	(0.0055)	(0.0056)	(0.0052)	(0.0054)	(0.0101)	(0.0086)
FGT(1)	0.0278	0.0327	0.0241	0.0251	0.0213	0.0160
	(0.0017)	(0.0019)	(0.0015)	(0.0016)	(0.0026)	(0.0020)
No. households	507,652	—	507,652	1,016,768	1,371,907	1,371,907
Budget (pesos)	322,740,115	—	645,480,230	625,389,322	468,964,138	937,928,277
Rural			0.0893			
FGT(0)	0.3471	0.4046	0.2578	0.3471	0.3211	0.2273
	(0.0111)	(0.0114)	(0.0100)	(0.0111)	(0.0220)	(0.0171)
FGT(1)	0.1200	0.1810	0.0727	0.1200	0.1098	0.0652
	(0.0046)	(0.0066)	(0.0032)	(0.0046)	(0.0087)	(0.0052)
No. households	3,730,625	—	3,730,625	3,730,625	2,946,421	2,946,421
Budget (pesos)	2,079,895,504	—	4,159,791,008	2,079,895,504	1,392,400,268	2,784,800,537

Source: Authors' calculations, based on data from the 2002 ENIGH.

a. Poverty indexes and their standard errors (in parentheses) were calculated using Stata sepov instruction.

TABLE 8. Distribution of Labor Force from Accounting Simulations
Percent (except for Beneficiaries totals)

Sample group, index, and employment status	*Default*	*Accounting simulation*			
		Cancel	*Double*	*Target*	*Target + double*
National					
FGT(0)	0.2024	0.2248	0.1695	0.1755	0.1300
Urban					
FGT(0)	0.1143	0.1183	0.1069	0.0868	0.0708
Beneficiaries[a]	507,652	—	507,652	1,371,907	1,371,907
Male					
Inactive	18.93	18.93	18.93	18.93	18.93
Part-time	2.98	2.98	2.98	2.98	2.98
Full-time	46.45	46.45	46.45	46.45	46.45
Overtime	31.65	31.65	31.65	31.65	31.65
Female					
Inactive	54.18	54.18	54.18	54.18	54.18
Part-time	6.07	6.07	6.07	6.07	6.07
Full-time	30.38	30.38	30.38	30.38	30.38
Overtime	9.37	9.37	9.37	9.37	9.37
Minors					
Inactive	94.72	94.72	94.72	94.72	94.72
Part-time	1.21	1.21	1.21	1.21	1.21
Full-time	3.14	3.14	3.14	3.14	3.14
Overtime	0.92	0.92	0.92	0.92	0.92
Rural					
FGT(0)	0.3471	0.3996	0.2722	0.3211	0.2273
Beneficiaries[a]	3,730,625	—	3,730,625	2,946,421	2,946,421
Male					
Inactive	14.05	14.05	14.05	14.05	14.05
Part-time	3.29	3.29	3.29	3.29	3.29
Full-time	44.76	44.76	44.76	44.76	44.76
Overtime	37.90	37.90	37.90	37.90	37.90
Female					
Inactive	58.73	58.73	58.73	58.73	58.73
Part-time	11.37	11.37	11.37	11.37	11.37
Full-time	20.26	20.26	20.26	20.26	20.26
Overtime	9.64	9.64	9.64	9.64	9.64
Minors					
Inactive	90.92	90.92	90.92	90.92	90.92
Part-time	2.58	2.58	2.58	2.58	2.58
Full-time	4.54	4.54	4.54	4.54	4.54
Overtime	1.96	1.96	1.96	1.96	1.96

Source: Authors' calculations, based on data from the 2002 ENIGH.
a. Beneficiaries correspond to the number of families receiving cash transfers from Oportunidades.

TABLE 9. Distribution of Labor Force from Behavioral Simulations

Percent (except for Beneficiaries totals)

Sample group, index, and employment status	*Default*	*Behavioral simulation*			
		Cancel	*Double*	*Target*	*Target + double*
National					
FGT(0)	0.2024	0.2267	0.1639	0.1755	0.1300
Urban					
FGT(0)	0.1143	0.1184	0.1066	0.0868	0.0708
Beneficiaries[a]	507,652	—	507,652	1,371,907	1,371,907
Male					
Inactive	18.93	18.93	18.93	18.93	18.93
Part-time	2.98	2.98	2.98	2.98	2.98
Full-time	46.45	46.45	46.45	46.45	46.45
Overtime	31.65	31.64	31.65	31.65	31.65
Female					
Inactive	54.18	54.18	54.18	54.18	54.18
Part-time	6.07	6.07	6.07	6.07	6.07
Full-time	30.38	30.38	30.38	30.38	30.38
Overtime	9.37	9.37	9.37	9.37	9.37
Minors					
Inactive	94.72	94.72	94.72	94.72	94.72
Part-time	1.21	1.21	1.21	1.21	1.21
Full-time	3.14	3.14	3.14	3.14	3.14
Overtime	0.92	0.92	0.92	0.92	0.92
Rural					
FGT(0)	0.3471	0.4046	0.2578	0.3200	0.2200
Beneficiaries[a]	3,730,625	—	3,730,625	2,946,421	2,946,421
Male					
Inactive	14.05	14.05	14.06	14.05	14.05
Part-time	3.29	3.29	3.29	3.29	3.29
Full-time	44.76	44.72	44.77	44.76	44.76
Overtime	37.90	37.94	37.89	37.90	37.90
Female					
Inactive	58.73	58.73	58.73	58.73	58.73
Part-time	11.37	11.35	11.40	11.37	11.37
Full-time	20.26	20.26	20.26	20.26	20.26
Overtime	9.64	9.66	9.61	9.64	9.64
Minors					
Inactive	90.92	90.85	90.92	90.92	90.92
Part-time	2.58	2.65	2.58	2.58	2.58
Full-time	4.54	4.54	4.54	4.54	4.54
Overtime	1.96	1.96	1.96	1.96	1.96

Source: Authors' calculations, based on data from the 2002 ENIGH.

a. Beneficiaries correspond to the number of families receiving cash transfers from Oportunidades.

to 200, which hardly has any effect on anyone's labor decisions, according to the results of the multinomial models (see figures 2 and 3).

Simulating perfect targeting of Oportunidades also renders similar head-count indexes in both urban and rural areas for behavioral and accounting simulations (compare the last two columns of tables 7 and 2). The lower poverty indexes, relative to the selection of beneficiaries, are due to the assumption of perfect targeting (which is obviously an extreme assumption), while the stability of labor supply responses reflects the insensitivity of the labor supply to the simulated changes in cash transfers.

Validation

A valuable feature of recent program evaluation studies is to compare simulation results derived from structural models with the results of earlier studies that make use of experimental or quasi-experimental data.[38] This procedure is not straightforward in our case, however, because the aim of our simulation exercise was to forecast changes in poverty and the adult labor supply, whereas most experimental evaluations deal with changes in school attainment, health, and other demographics. We therefore compare the results of experimental studies measuring the impact on child labor with our labor supply simulations for minors (that is, only individuals aged twelve to seventeen). We also compare the school attendance results from experimental studies with the forecasts derived from our model on school attendance (see table 10).

Table 11 summarizes the comparison. With respect to school attendance, our simulation exercise shows effects that are compatible with those found by experimental studies. For urban areas, both the available experimental study and our simulation show no effect on school attendance for minors aged twelve to seventeen.[39] This result may reflect the fact that the Oportunidades program was recently implemented in urban areas (both the experimental study and our simulation use data collected in the year 2002), as well as the difficulty of inducing school attendance among urban youth.

For rural areas, our simulations render a positive effect, as do the selected experimental studies.[40] The difference is that our simulation finds small

38. See, for instance, Todd and Wolpin (2003).
39. The experimental study is Todd and others (2005).
40. Parker and Skoufias (2000); Schultz (2004); Buddelmeyer and Skoufias (2004).

TABLE 10. School Attendance for Minors (Forecast[a])
Percent

	With Oportunidades		*Without Oportunidades*		*Difference*	
Region and age	*Girls*	*Boys*	*Girls*	*Boys*	*Girls*	*Boys*
Rural						
12	93.7	94.7	93.4	94.4	0.4	0.3
13	87.3	89.3	86.7	88.7	0.6	0.6
14	78.6	80.1	77.7	79.2	0.9	0.9
15	65.3	67.0	64.3	66.0	1.1	1.0
16	49.8	53.1	48.6	52.2	1.2	0.9
17	39.9	38.8	39.2	38.2	0.8	0.6
Urban						
12	96.6	96.8	96.6	96.8	0.0	0.0
13	93.2	93.3	93.2	93.3	0.0	0.0
14	86.6	87.7	86.6	87.7	0.0	0.0
15	77.7	79.5	77.7	79.5	0.0	0.0
16	68.1	69.2	68.1	69.2	0.0	0.0
17	55.5	56.2	55.5	56.1	0.0	0.0

Source: Authors' calculations.
a. Using probit of school attendance (see table 6).

effects (between 0.3 and 1.2 percentage points), whereas the experimental studies show larger effects (between 3.0 and 9.0 percentage points). This stems from the use of different reference databases. The experimental studies usually compare treatment and control groups, which by definition are very similar, and do not necessarily have a sample selection mechanism that makes the sample representative of the national population. Our microsimulation exercise makes use of a nationally representative sample, so we are measuring the effects at the national level. The national average effects, after weighting by the relative size of the poor and non-poor populations, are bound to be smaller than in experimental studies.

With respect to child labor, the comparisons reveal more divergent results, even among the experimental studies. For urban areas, Todd and others find significant negative effects for boys and insignificant effects for girls, while we find no effect at all. For rural areas, our simulations and Buddelmeyer and Skoufias both find insignificant effects for boys and girls, but Parker and Skoufias, as well as Schultz, find negative effects with varying significance by gender and age. In this case, our simulations coincide with the experimental studies in terms of the direction of the effect (either null or negative), but no uniform validation can be made in terms of the significance or size of the effects.

TABLE 11. Comparison of Microsimulation Results with Previous Experimental and Quasi-Experimental Studies[a]

Variable, sample, and study	Sample period	Summary of findings: Boys	Summary of findings: Girls
School attendance			
Urban			
Microsimulation	1998–2002	No change	No change
Todd and others (2005)	2002–03	Insignificant effect for ages 12–18 (except age 16: +33.9 percent)	Insignificant effect for ages 12–18
Rural			
Microsimulation		Positive effect for ages 12–17 (0.003, 0.010)	Positive effect for ages 12–17 (0.004, 0.012)
Buddelmeyer and Skoufias (2004)	1997–99	Significant positive effect for ages 12–16 (0.07)	Significant positive effect for ages 12–16 (0.09)
Skoufias and Parker (2001)	1997–99	Significant positive effect for ages 12–17 (0.06)	Significant positive effect for ages 12–15 (0.10); insignificant positive effect for ages 16–17
Schultz (2004)[b]	1997–99	Insignificant positive effect for 8 or more years of schooling	Insignificant positive effect for 7 or more years of schooling
Child labor			
Urban			
Microsimulation[c]		No change	No change
Todd and others (2005)[c]	2002–03	Significant negative effect for ages 12–14 (–5.3) and 15–17 (–6.5)	Insignificant effect for ages 12–16; significant negative effect for ages 17 (–17.5)
Rural			
Microsimulation[c]	1998–2002	Insignificant negative effect	Insignificant negative effect
Buddelmeyer and Skoufias (2004)[d]	1997–99	Insignificant effect for ages 12–16	Insignificant effect for ages 12–16
Skoufias and Parker (2001)[d]	1997–99	Significant negative effect for ages 12–13 (–0.04); insignificant negative effect for ages 14–17	Significant negative effect for ages 14–15 (–0.04); insignificant negative effect for ages 12–13 and 16–17
Schultz (2004)[e]	1997–99	Significant negative effect for secondary school (–0.02)	Insignificant negative effect for secondary school

a. Microsimulation results for school attendance are presented in table 10; results for child labor are presented in table 8.

b. These are difference-in-differences estimates. When using a probit model, Schultz finds a positive and significant effect on secondary enrollment.

c. Refers to paid child labor.

d. Refers to paid and unpaid child labor.

e. Refers to paid child labor. These are difference-in-differences estimates. When using a probit model and IV instruments for program participation, Schultz finds a significant negative effect for girls (–005) and an insignificant negative effect for boys.

Conclusions

In this paper, we have performed several microsimulation exercises to gauge the impact of cash transfers from Oportunidades on poverty. The first exercise consisted of simulating what would have happened if no cash was transferred (that is, the termination of the program). In this case, the national poverty headcount would have been 2 percentage points higher than in the base year (2002), the rural poverty headcount would have been 5 percentage points higher, and the urban poverty headcount would have had no significant change. The second exercise simulates doubling cash transfers to existing beneficiaries. In this case, the national headcount would have been 3 percentage points lower, the rural headcount would have been 7 percentage points lower, and the urban headcount would again show no significant change. The third exercise keeps cash transfers the same but doubles the number of urban beneficiaries. This leaves the urban poverty headcount practically unchanged.

The simulations reveal a serious targeting problem in urban areas. This may stem from the limitations of our simulation tool or from actual problems with the program in the field. To address this issue, we undertook a fourth simulation in which we assume that every poor household is covered by the program in both urban and rural areas. In this case, national poverty would have been nearly 2 percentage points lower than in 2002, owing to a reduction of approximately the same size in both areas. The fifth exercise doubles the amount of transfers to every poor household: poverty would have fallen 7, 4, and 12 percentage points at the national, urban, and rural areas, respectively.

We performed the five exercises using both accounting and behavioral simulations. We found no significant difference in the results, because our models of labor supply are insensitive to the changes simulated in cash transfers.

The microsimulation exercise presented in this paper has some additional implications. First, it shows that Oportunidades is associated with nearly a third of the reduction in rural poverty in Mexico for the period since the program was established. The reduction of two percentage points in the national poverty headcount for 2002 is almost entirely due to the 4.9 percentage point reduction in rural poverty. This reduction ascribed to Oportunidades had an average cost in cash transfers of nearly 326 million pesos a month per percentage point (that is, 1.2 percent of the central government's total spending).

Second, despite this relatively large impact, further reductions in poverty would require bigger transfers in rural areas and better coverage in urban areas. Even if the expansion were effective, it would impose additional costs on the Mexican social policy budget. An expansion of coverage in the urban sector would increase the average cost of poverty reductions unless targeting is improved in this area. Larger benefits for rural beneficiaries would reduce the average cost per percentage point reduction. These differential effects stem from different phenomena: lack of targeting in urban areas and a large poverty gap in rural areas.

Third, labor supply reactions to changes in cash transfers are almost negligible for most groups. Changes in men's, women's, or minors' labor supply would require huge subsidies that are not likely to occur. This implies that the transfers that Oportunidades is distributing among its beneficiaries do not have sizeable effects on adult labor supply.

Several limitations need to be addressed in future attempts to simulate the potential impact of Oportunidades. First, a more recent database ought to be used as the base year. This would contribute to making the simulations more appropriate and easier to interpret. Additionally, more recent data on the urban recipients of Oportunidades will provide information on the selection of beneficiaries and the heterogeneity of transfers. This additional information would improve the consistency of estimators in both the participation and the labor supply models for the urban area. Second, more complex models of labor supply, such as joint family labor supply and agricultural labor supply, need to be estimated. Given the variety of family structures and the large agricultural labor supply in Mexico, these are obvious gaps to fill in future research. These complex models could also deal with the thorny issue of simultaneous decisions on school attendance, instead of individual by individual, when families have more than one school-age child.