Risk Sharing and Private Transfers: What about Urban Households?*

Donald Cox
Boston College

Emmanuel Jimenez
World Bank

I. Introduction

Do private interhousehold transfers facilitate risk sharing in an urban environment? While recent evidence suggests that rural households can at least partially protect themselves against income fluctuations and other risks, far less is known about urban households. Further, there are reasons to believe that advantages and obstacles to risk sharing might be different in the city than in the countryside. We explore a unique data set for poor households in Cartagena, Colombia, that contains information on private interhousehold transfer networks and intrayear income variability. Such networks are large, and private transfers respond to income, income variability, and other household characteristics in ways that suggest households pool risks.

Knowing about informal risk-sharing institutions is important for understanding the determinants of the well-being of poor people in developing countries. Formal credit and insurance markets are less widespread in developing countries than in developed countries, and publicly provided social insurance programs are often lacking in low-income countries. Thus, poor households may not adequately protect themselves from events that threaten their well-being, such as drought, pestilence, unemployment, or illness. If so, there may be strong justification for policy intervention. For example, R. Townsend conjectures that improved legal systems that allow more complete contracts across and within villages—between, say, borrowers and lenders or insurers and those insured—could improve farmer welfare in Thailand.[1]

Recently, economists have begun to investigate attempts that households make to mitigate the effects of these negative shocks, and nearly

all of the empirical work pertains to rural families. Townsend rejects the strongest form of complete risk sharing for rural households in India, since household income sometimes enters significantly in consumption regressions even after controlling for village consumption.[2] But the coefficients for household income are never very large, leading him to conclude that the risk-pooling model is a good approximation of household behavior. Using United States consumption data, B. Mace also finds some evidence favoring the risk-sharing hypothesis, though results are sensitive to functional form and measurement error.[3] J. Cochrane, also using United States data, finds that some events (e.g., unemployment, strikes) appear to be insured, while others (e.g., long illness) are not.[4] In contrast, H. Alderman and C. Paxson's review notes that consumption-based tests are difficult to interpret because the small estimated effects of current household income on consumption are consistent with the permanent income hypothesis as well as risk pooling, and they might also be the artifact of errors in measurement of household income.[5] Further, a strong relationship between changes in village and household consumption could just be an affirmation of the permanent income hypothesis.[6]

The literature that addresses the means by which households smooth income and consumption (e.g., diversification of activities, private transfers, capital markets) provides a valuable complement to the consumption-based studies of risk pooling.[7] There is some direct evidence that family-related private-transfer networks mitigate the impact of income variability. For example, M. Rosenzweig, using data from rural Indian villages, finds that fluctuations in household income are in part offset by private transfers, which supports the household risk-sharing model. Further, he advances the idea that household networks are superior to some other modes of mitigating consumption shortfalls, such as local credit markets.[8] Rosenzweig and O. Stark suggest that a primary reason for marriage-related migration is that it nurtures the long-distance, interhousehold linkages needed to diversify risk due to, say, local weather conditions.[9] In case of drought, for example, it helps to be linked to a farm household located far away and unaffected by local weather conditions. Using data from six south Indian farming villages, they find that the variance of consumption is lower for households that have long-distance linkages. Further, they find that households facing exogenously higher income risks because of inferior land or uncertain rainfall patterns tend to form longer-distance linkages. These presumably are the ties that facilitate private, insurance-related income transfers between households. R. Lucas and Stark provide evidence that migrant remittances to households in Botswana function in part as insurance. Being afflicted by drought and having drought-sensitive assets such as cattle is associated with higher remittance receipts.[10]

Despite these advances, more research is needed. Although the articles mentioned above suggest that private transfers function in part as insurance, in the words of Alderman and Paxson, ''there is simply not enough literature from enough countries to draw general conclusions about the scope and importance of household risk-pooling.''[11] Moreover, most evidence about risk sharing pertains to rural families; much less is known about urban networks. A priori, it is not clear whether the urban setting facilitates or hinders risk sharing. Moral hazard problems and the availability of alternative jobs suggest that urban risk sharing will be lower than rural risk sharing, ceteris paribus. Lower income variability suggests the opposite. We elaborate on these ideas below.

Insurance schemes might not function as well in the city because the source of an income shortfall may be more difficult to verify for urban workers. Households might agree ex ante to insure only against exogenous calamities and not to compensate for events that appear to be within the households' control, since otherwise the scheme would be fraught with moral hazard problems. But knowing which shocks are beyond the households' control is likely to be easier in rural settings. A farmer who took reasonable precautions but whose crops were destroyed by drought, insects, or fire can credibly claim to be the victim of the vagaries of nature, but someone who sells lottery tickets on a street corner might have a bad month because business was sour or because he did not work hard enough. Further, workers in rural insurance networks are likely to share the same general knowledge about requisite safety measures in the face of weather- or price-related shocks, but a cook and a mechanic may know so little about one another's work that neither can tell if the other is falling prey to moral hazard. If insurers cannot identify the exogenous component of income shocks, insurance schemes are likely to break down.

One way to distinguish fate from negligence is by monitoring, which is likely to require propinquity. But if network members live too close to one another, they will not be able to insure against area-specific shocks, and monitoring costs may be too high to take advantage of spatial risk diversification. However, one advantage of the urban setting is that income covariation between adjacent city dwellers is probably less than between adjacent farmers. Since the city has more occupational diversity than does the countryside, the value of far-flung networks for reducing income covariance should be lower.

Occupational diversity might also reduce the value of networks if households diversify their portfolios of human capital by having individual members work at different jobs or hold multiple jobs. Having multiple jobs is analogous to crop diversification in the rural setting. Each option is a substitute for interhousehold insurance. One prediction that emerges from these considerations is that urban insurance networks, if

they exist, should be more tightly clustered geographically compared to rural networks. A second prediction is that opportunities for multiple jobs reduce the need for interhousehold risk sharing.

Another important condition for risk sharing, emphasized by Rosenzweig, is that network members need to know each other's risks and expected liabilities in order to form insurance contracts.[12] Farmers who use a fixed technology are likely to experience weather-related shocks coming from a fixed probability distribution that is stable over time. In contrast, workers in a fast-changing urban environment are not likely to face the same distribution of shocks year in and year out. Economists know little about wage and employment dynamics among the urban poor because the requisite panel data are lacking. So we have little evidence on the nature of the distribution of unobservable earnings in the urban informal sector. Whether this distribution is stationary is not known, which leaves the viability of urban risk-sharing networks an open question.

One of the reasons why there is so little evidence concerning urban risk sharing is that the urban counterparts of data sets used by economists to study the rural sector are scarce. In this article, we use a data set from Colombia that contains extensive information about private transfers, household networks, and intra-annual earnings fluctuations that can be used to examine the risk-sharing hypothesis. The data set has two advantages. First, survey respondents were asked detailed questions about private transfers received and given and they were also questioned about household networks and private transfers. Detailed questionnaires are far superior to the usual summary ones for measuring the extent and magnitude of private transfers. Second, although the data set is cross-sectional, it contains information about intra-annual income variability that is useful for investigating the connection between stability in incomes and transfer behavior. One of the salient findings in the empirical work below is that the stabilization of household incomes would likely be associated with a sharp reduction in private transfer activity.

II. The Cartagena Data Set

The data set is based on a survey specifically designed to measure the extent and magnitude of interhousehold transfers for a sample of 507 low-income households.[13] What is significant about the survey is the explicit and detailed focus on private transfers. Household survey data sets that contain private transfer information are somewhat scarce, and those that contain responses to lengthy, detailed questionnaires are scarcer still. The large number of questions that the Cartagena data set contains is valuable for investigating interhousehold risk sharing because the more respondents are prompted, the more likely they will recall and report transfers they gave or received. Surveys containing only summary questions are likely to be affected by severe underreporting, which opens the

possibility for misleading conclusions about the effectiveness of private transfers for stabilizing income.[14]

Another feature is that survey respondents were asked questions dealing with intrayear earnings fluctuations. Workers were asked to report their highest and lowest monthly earnings along with their current monthly earnings. Thus, we can determine which households have earnings flows that are stable throughout the year. In addition, the Cartagena data set contains indicators of marital status, female headship, education, and other demographic variables that are useful for investigating private transfer behavior.

Many studies of private transfers are affected by problems of omitted variables related to the source of transfers received and the destination of transfers given. A unique feature of the Cartagena data set is that it contains information about the number of network members by relative financial status. Each household reported the number of network members who were better off financially, members who were worse off, and members with equal financial status. Such a network is defined as ''a set of individuals or households who regularly assist each other through the provision of money, goods, services or the provision of accommodation.''[15] Household members were queried in detail about financial and in-kind private transfers that they received from members of their network.

The sample is not a representative cross-section. It covers households who reside in a $6.5 \times .5$ kilometer site in southeastern Cartagena. These households tend to be poor; their incomes are about 60% or less of Cartagena's average income and are about a third of that of urban Colombia. Thus, one caveat to note is that our results may not apply to rural areas or upper-income urban areas. However, our sample would not cause a potential truncation bias under the assumption that income is exogenous to the model, since in this case our sample would be stratified according to a right-hand-side variable in the empirical model.

III. Specification

The primary question we investigate is, how do private transfers respond to the level and variability of household income? We express the incidence of private transfers as a function of indicators of household resources, their variability over time, and a variety of additional controls. Households with stable earnings are expected to be less likely to receive transfers than households with unstable earnings. Those that have especially low earnings or experience unemployment would be more likely to be recipients. Conversely, those who have especially high earnings would be less likely to receive transfers. Further, transfers should be targeted to households with low average incomes.

Listed below is the set of variables related to the risk-sharing hypothesis:

Stable income—a dummy that takes on a value of 1 if the earnings of each worker in the household is constant throughout the year.

Unemployment—a dummy that takes on a value of 1 if a worker in the household is not continuously employed throughout the year.

High income—a dummy that takes on a value of 1 if the highest monthly earnings in the past year exceed twice the value of current monthly earnings.

Low income—a dummy that takes on a value of 1 if the lowest earnings of the year are less than half of the value of current monthly earnings.

In addition to these variables, we enter the number of workers in the household. This variable indicates the household's ability to self-insure and is expected to be inversely related to transfer receipts. We also include the log of household income, which is also expected to be inversely related to transfer receipts. The estimating equation for transfer receipts also contains a set of network variables. Households with many members who are well off are expected to be more likely to receive a transfer; those with more members who are worse off would be less likely to receive.

Other controls included in the equation for private transfer receipt are years of schooling, a quadratic in age of household head, a dummy indicating whether the household was headed by a single female, and household size. If private transfers respond to liquidity constraints, they should be positively related to schooling. The reason is that more schooling is associated with higher permanent income—with current income held constant, higher permanent income implies higher desired consumption and hence private transfers. The liquidity-constraints hypothesis also implies that life-cycle timing should matter a great deal for private transfers. Alternatively, with perfect capital markets, only the discounted value of net transfers is identified, and the age profile of these transfers is indeterminate.

Nearly all empirical studies from a variety of countries such as Botswana, El Salvador, Italy, Peru, the Philippines, and the United States indicate that private transfers are targeted toward women.[16] The relationship between female headship and transfers is partly mechanical. For example, one reason women receive a disproportionate share of transfers in the Philippines is that much of what they receive comes from international remittances from their husbands. But the female-status—transfer relationship can also be based on behavioral patterns, and it is found in studies of purely intergenerational—as opposed to interspousal—transfers. For example, if transfers are in part payments for in-kind services, an activity in which women tend to be disproportionately involved, we would expect financial transfers to be targeted to women. Further, if women face more limited opportunities in the labor market, they are likely to be subject to higher income risk than male-headed households

are. So the relationship between female headship and transfers could be consistent with the risk-sharing hypothesis.

IV. Main Results

We first provide an overview of the characteristics of the sample according to transfer status. Observations from the original sample of 507 were dropped if they had missing values for transfers given or received or for earnings, age, or low or high monthly earnings for the past year. These sample selection criteria leave us with a sample of 369 households.

   We divided the households according to net transfer status (table 1). A household is labeled a net transfer recipient if gross transfers received exceed gross transfers given. Gross transfers received, in turn, are defined as the sum of cash plus in-kind transfers received throughout the year from other households. Households report the average monthly value of the transfer received and the number of months that the transfer was received. Gross transfer receipts are defined as the product of the average transfer receipt and the number of months the transfer was received. We divide this product by 12 so that all monetary values are on a monthly basis. Gross transfers given are derived in the same way.

   Of the 369 households in the sample, a third were net transfer recipients, 39% were givers, and the rest were not involved with transfers (table 1). We refer to the last group as ''others.''[17] Private transfers are an important source of income for households. Among recipients, for example, net transfer receipts are 15% of pre–private-transfer income (table 1). These basic figures appear to fulfill a necessary condition for private transfers to perform a consumption-smoothing role—they must be a large enough potential income supplement.

   Donors make net gifts of 6% of their pretransfer income. The average gift is smaller than the average receipt, which appears to cast doubt on the idea that households downplay their dependency or exaggerate their generosity. At the same time, however, recall that the sample is a low-income one, so an excess of aggregate receipts over gifts is to be expected if transfers flow from high- to low-income households.[18]

   The rank order of pretransfer incomes is consistent with the idea that transfers do indeed flow from better- to worse-off households (table 1). Average pretransfer income of donors is highest, that of recipients is lowest, and that of others is in between. Recipients are, on average, 8 years younger than donors and have less schooling. A higher proportion of recipients have experienced unemployment during the year prior to the survey, and female-headed households are overrepresented among recipients and underrepresented among donors.

   Figure 1 illustrates the variability of monthly earnings for individual workers in the sample. Recall that workers were asked to report their highest and lowest monthly earnings, in addition to their current earn-

TABLE 1

SELECTED CHARACTERISTICS OF CARTAGENA HOUSEHOLDS BY PRIVATE-TRANSFER STATUS

Variable	Net-Transfer Recipients	Net-Transfer Donors	Others*	All Households
Total income before transfers (pesos)	14,086	23,161	15,919	18,090
Total income after transfers (pesos)	16,195	21,742	15,919	18,253
Years of schooling	4.040	4.458	3.922	4.168
Age of household head (years)	48.040	40.937	44.136	44.217
Proportion of households with:				
Low income†	.105	.106	.058	.092
High income‡	.065	.070	.097	.076
Stable income	.065	.169	.136	.125
At least one member unemployed	.427	.282	.388	.360
No. of network members better off	3.202	1.951	2.282	2.463
No. of network members equally well off	2.557	3.310	3.320	3.060
No. of network members worse off	.935	1.676	.971	1.230
Female-headed	.419	.119	.252	.257
Married head	.556	.859	.738	.724
Head is single female	.387	.092	.214	.225
No. of:				
Household members	6.661	6.401	6.786	6.596
Workers in the household	1.258	1.444	1.398	1.369
Transfers:				
Proportion giving net transfers	.000	1.000	.000	.385
Net transfers given (pesos)	.000	1,419.1	.000	546.11
Proportion receiving net transfers	1.000	.000	.000	.336
Net transfers received (pesos)	2,109.4	.000	.000	708.85
Proportion giving gross transfers	.363	1.000	.029	.515
Gross transfers given (pesos)	116.98	1,503.8	1.049	618.30
Proportion receiving gross transfers	1.000	.289	.029	.455
Gross transfers received (pesos)	2,226.4	84.676	1.049	781.03
No. of cases	124	142	103	369

* Neither a net-transfer recipient nor a net-transfer donor.
† Aggregate low income less than half monthly earnings.
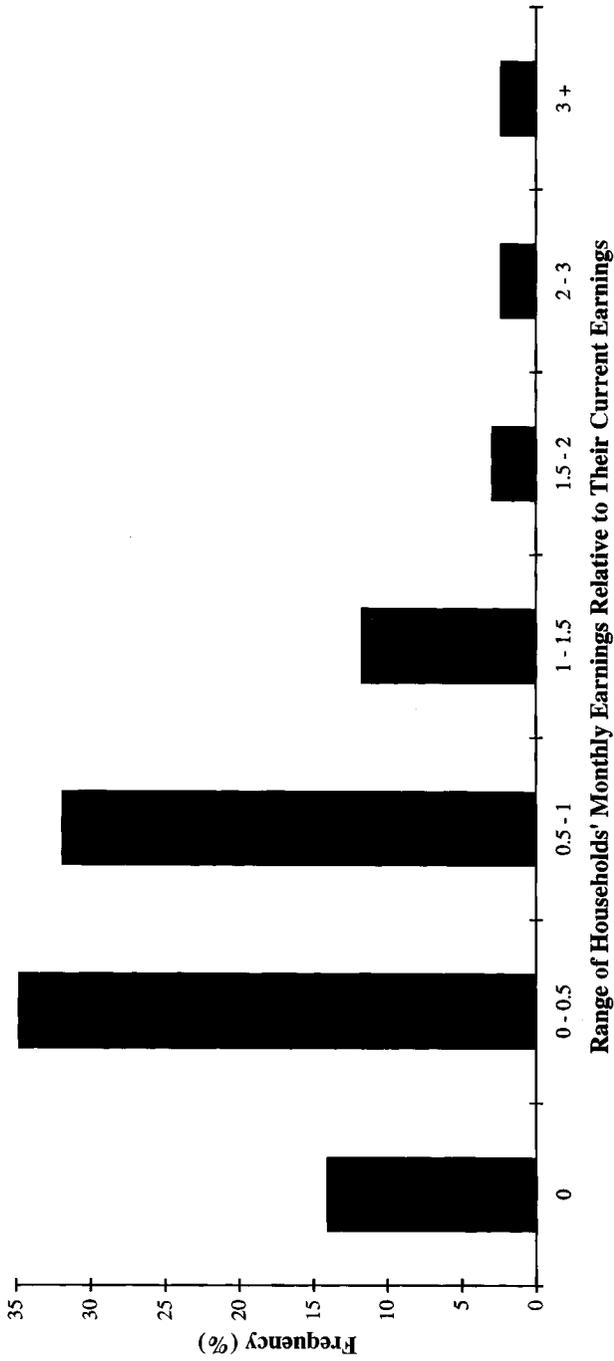‡ Aggregate high income greater than twice monthly earnings.

FIG. 1.—How variable are household earnings? Note that the range is defined as a household's highest monthly earnings reported for the year, minus its lowest monthly earnings in the same year. The X-axis shows (high earnings − low earnings)/current earnings.

TABLE 2

ORDERED PROBIT ESTIMATES: DEPENDENT VARIABLE—NET TRANSFER RECEIPT

Variable	Coefficient	Asymptotic *t*-Ratio	Variable Mean
Constant	2.277	2.964	1.0000
Log of income	−.1605	−3.531	9.2423
Years of schooling	.0355	1.683	4.1680
Age of household head	−.0692	−2.315	44.217
Age squared	.0009	2.969	2,139.1
Low income	.1733	.808	.0921
High income	−.4574	−1.683	.0759
Stable income	−.3922	−1.728	.1247
Unemployment	.3558	2.474	.3604
Network members better off	.0690	2.755	2.4634
Network members worse off	−.0812	−2.115	1.2304
Network members equal	.0020	.092	3.0596
Head is single female	.8789	5.043	.2249
Household size	.0536	2.056	6.5962
No. of workers	−.1336	−1.751	1.3686
Mu(1)	.8780	11.182	· · ·

	Count	Frequency
Dependent variable:		
Net givers (dep. var. = 0)	142	.38
Nongivers, nonrecipients (dep. var. = 1)	103	.28
Net recipients (dep. var. = 2)	124	.34
Log-likelihood	−346.51	
Restricted (slopes = 0) log-likelihood	−402.26	
$\chi^2(14)$	111.51	

ings. Low earnings equal zero if a worker experiences a month or more of unemployment. We calculated the range of earnings (i.e., high earnings minus low earnings) and expressed it relative to current earnings, for the 342 households with positive current earnings. Fourteen percent of these households had monthly earnings that were completely stable throughout the year. About half of the households had a range of earnings that exceeded 50% of their current earnings. (The median value of the range of monthly earnings relative to current earnings for this sample is 0.52, and the mean is 0.71.) A substantial fraction of households, more than 19%, had a range of monthly earnings that exceeded their current monthly earnings. So the income swings that some of these households experience can be quite wide.[19]

To explore in detail the patterns of giving and receiving, we estimated an ordered probit model of net receipts (table 2).[20] Consider the latent variable t_i, which determines the transfer status of household i. The latent variable is a linear function of observables X plus an error term e, assumed to be distributed standard normal, so that

$$t_i = X_i\beta + \epsilon_i.$$

The transfer-status variable T is determined as follows:

$$T_i = 0 \text{ if } t \leq \mu_1,$$
$$T_i = 1 \text{ if } \mu_1 < t \leq \mu_2,$$
$$T_i = 2 \text{ if } t > \mu_2,$$

where $T = 2$ denotes net-recipient status, $T = 1$ denotes neither giving nor receiving, and $T = 0$ denotes net giver status.[21]

The ordered probit results are presented in table 2. Each of the variables associated with the risk-sharing hypothesis enters with the expected sign. The coefficient on log income indicates that an increase in monthly income from 5,000 to 20,000 pesos reduces the probability of transfer receipt 8 percentage points. Most important, having stable income reduces the probability of transfer receipt by nearly 13 percentage points. (The estimated coefficient is significant at the .1 level.) Experiencing unemployment sometime during the year prior to the survey raises the probability of transfer receipt by about the same amount. Having an extremely high draw for earnings reduces the probability of receiving a transfer, and having an extremely low draw for earnings increases it, though the latter coefficient estimate is not significant at even the .1 level. The Wald test statistic for the joint significance test for low, high, and stable income yields a χ^2 test statistic that is significant at the .08 level. With aggregate income of the household constant, an increase in the number of workers in the household is associated with a smaller probability of transfer receipt, but an increase in household size is associated with a larger probability of transfer receipt. Together, these results indicate substantial support for the risk-sharing hypothesis.

We also find that having more network members who are relatively better off raises the probability of transfer receipt and having more members who are worse off lowers the probability of receipt. The number of household members with equal financial status has a negligible effect on transfers. The probability of transfer receipt increases with schooling and then falls and rises over the life cycle, hitting a trough at age 37. Elsewhere, we find a similar, U-shaped age pattern for transfers in Peru.[22] We concluded in that study that the pronounced life-cycle pattern for transfers indicated that they were responsive to liquidity constraints. Finally, we find strong evidence that private transfers are targeted to female-headed households. All else being equal, female status raises the probability of transfer receipt by almost 33 percentage points.

There is some reason to suspect that the ordered-probit specification might be too restrictive. For example, we find that the effects of having a stable source of earnings throughout the year reduces the probability

TABLE 3

PROBIT ESTIMATES: DEPENDENT VARIABLE—NET TRANSFER RECEIPT

Variable	Coefficient	Asymptotic *t*-Ratio	Variable Mean
Constant	1.0137	1.150	1.0000
Log of income	−.1555	−2.866	9.2423
Years of schooling	.0461	1.852	4.1680
Age of household head	−.0551	−1.573	44.217
Age squared	.0008	2.176	2,139.1
Low income	.3866	1.461	.0921
High income	−.5782	−1.867	.0759
Stable income	−.5482	−1.999	.1247
Unemployment	.3374	1.981	.3604
Network members better off	.0799	2.743	2.4634
Network members worse off	−.0526	−1.132	1.2304
Network members equal	−.0048	−.187	3.0596
Head is single female	.8365	4.502	.2249
Household size	.0403	1.410	6.5962
No. of workers	−.1023	−1.140	1.3686
Dependent variable count:			
Net recipients	124		
Nonrecipients	245		
Dependent variable mean	.34		
Log-likelihood	−192.00		
Restricted (slopes = 0) log-likelihood	−235.56		
$\chi^2(14)$	87.116		

of transfer receipt, but the ordered-probit specification forces the effect of having stable income on the probability of giving a transfer to be positive. Nevertheless, those with stable earnings sources might have less incentive to participate in insurance networks, raising the possibility that having stable earnings reduces both the probability of receiving and giving a transfer.

We informally investigated the symmetry of the transfer relationship by running separate probit equations for receiving and giving a transfer, and these probit results are presented in tables 3 and 4, respectively. We find that giving and receiving do appear strongly symmetrical. Every coefficient, except two that are each statistically insignificant, changes sign in moving from the probit for transfer receipt to transfer giving.

The four coefficients associated with earnings fluctuations—low income, high income, stable income, and unemployment—are jointly significant at the .01 level in the probit equation in table 3. We can use these estimates to answer the following question: if earnings were stabilized, how much would the incidence of private transfers fall? The average predicted probability implied by the probit estimates of table 3 is 0.336. Setting the low-income, high-income, and unemployment dummies to 0 and the income stability dummy to 1, the estimates in table 3

TABLE 4

PROBIT ESTIMATES: DEPENDENT VARIABLE—NET TRANSFER GIVER

Variable	Coefficient	Asymptotic t-Ratio	Variable Mean
Constant	−2.9132	−2.991	1.0000
Log of income	.1550	2.227	9.2423
Years of schooling	−.0242	−.982	4.1680
Age of household head	.1052	2.652	44.217
Age squared	−.0014	−3.106	2,139.1
Low income	.0089	.034	.0921
High income	.3894	1.309	.0759
Stable income	.3198	1.354	.1247
Unemployment	−.3846	−2.220	.3604
Network members better off	−.0570	−1.875	2.4634
Network members worse off	.1063	2.502	1.2304
Network members equal	−.0019	−.081	3.0596
Head is single female	−.9270	−4.332	.2249
Household size	−.0760	−2.408	6.5962
No. of workers	.1704	1.883	1.3686
Dependent variable count:			
Net givers	142		
Nongivers	227		
Dependent variable mean	.38		
Log-likelihood	−201.94		
Restricted (slopes = 0) log-likelihood	−245.89		
$\chi^2(14)$	87.897		

imply an average predicted probability of 0.180—a reduction of more than 40%. Our estimates indicate that income stabilization could well be a quantitatively important determinant of private transfers.

Distinguishing altruistic from nonaltruistic motives for interhousehold insurance by looking at receiving is difficult, since each theory predicts, for example, that income and transfers move inversely. Looking at giving may be more informative about the root motivation of private transfers. For example, if having a stable income reduced the probability of giving transfers, we might be led to conclude that transfers were selfishly motivated—those with secure sources of income refrain from contributing to transfer networks. The opposite finding, a positive relationship between income stability and propensity to contribute, would be more consistent with altruism. The probit estimates for net giving in table 4 are compatible with the latter motive. Although the coefficient is not significant at conventional levels, the point estimates indicate that having a stable income raises the probability of giving transfers.

The measures of low income draws used in tables 2–4 are each based on the sum of low earnings across all earners in the household. If this sum is less than the sum of monthly earnings across all earners in the household, the low-income dummy equals 1. The high-income-draw dummy in these tables is also based in this case on the sum of high earn-

ings of household members. But with multiple workers it is possible that, even with such high and low draws, households have a stable income if one worker's high coincides with another's low. We also use a different, though again imperfect, measure of low and high draws. The low-draw dummy takes on a value of one if any worker's low earnings are less than half of his monthly earnings. These results, which are not presented here but which are available from us, are much the same as those in tables 2–4. Having stable income or an especially high earnings draw is inversely associated with receiving a transfer.

We also experimented with alternative definitions of income stability. We used a less restrictive definition of stability, in which low and high earnings stayed within a specified band (e.g., ±10% of monthly earnings) instead of requiring that earnings remain constant throughout the year. These specifications produced results similar to those reported in tables 2–4. We also replaced the stable dummy with the coefficient of variation calculated from high, low, and current earnings. The results from this specification also reinforce the results reported here.

The log-income measure pertains only to the survey month, but transfer status is, in a sense, an annual concept—so there is a mismatching of time periods. Accordingly, we also use an average of monthly earnings, low earnings, and high earnings in the log-income measure. Again, the basic results are little changed with this modification. In particular, the log-income measure is inversely related to the probability of transfer receipt.

We treat the number of network members as exogenous, and this assumption may well be defensible, since much of the network is likely determined by number of siblings, something over which the household has little control. But it may be that networks themselves are endogenous. A simple tabulation comparing the number of network members by stability status or unemployment status reveals an interesting pattern. Households with stable earnings have 5.9 other households in their network. Those with nonstable earnings have 7.1. Similarly, those that experience a spell of unemployment have more network members. These numbers suggest that perhaps those facing more uncertain earnings have an incentive to form larger networks.

Earlier we suggested that households making transfers have an incentive to monitor one another, and the Cartagena data set has some information relevant for this issue, the frequency of visits. The visiting pattern supports the idea that households involved in financial transfers monitor each other more often. Visits per household were nearly 40% higher for those giving or receiving transfers than for those not involved with transfers. While this pattern must be carefully interpreted, since it is also consistent with other behavior (e.g., reciprocal altruism, exchange), it does support the insurance idea.

The Cartagena data set gives the location of network members (in

the immediate neighborhood, city, and so on). About 60% of network members live in Cartagena, which contrasts with Rosenzweig's findings for rural networks.[23] Therefore, the prediction discussed earlier, that urban networks would be tend to be more tightly clustered, appears to be borne out in the data.

V. Conclusion

There is an emerging consensus that rural households engage in inter-household risk sharing through inter vivos transfer networks, but little previous evidence exists for urban settings, where both costs and benefits of such insurance are likely to be very different. We find evidence that transfers are insurance motivated. Variables that proxy household income risk have a large impact on private transfer behavior. The results suggest that private and public transfers are likely to interact in a way that might thwart the effectiveness of government policies designed to protect households from the effects of income shortfalls. Our estimates suggest that policies such as unemployment insurance or social security could affect private incentives to provide insurance and support. Public provision of insurance to protect incomes of the poor, for example, could weaken family networks and dilute the impact of the policy. On the other hand, public transfer programs expand the risk-sharing pool and facilitate its insurance function. Further, government stabilization of incomes of some households could leave them in a better position to provide private insurance for others in their network, reinforcing the effects of social insurance. No matter how private and public transfers interact, however, it is unlikely that private transfer networks would be impervious to changes in social insurance policy.

Notes

1. Robert M. Townsend, ''Financial Systems in Northern Thai Villages,'' *Quarterly Journal of Economics* 110 (November 1995): 1011–46.

2. Robert M. Townsend, ''Risk and Insurance in Village India,'' *Econometrica* 62 (May 1994): 539–92.

3. Barbara Mace, ''Full Insurance in the Presence of Aggregate Uncertainty,'' *Journal of Political Economy* 99 (October 1991): 928–56; Julie Nelson, ''On Testing for Full Insurance Using Consumer Expenditure Survey Data,'' *Journal of Political Economy* 102 (April 1994): 384–94.

4. John Cochrane, ''A Simple Test of Consumption Insurance,'' *Journal of Political Economy* 99 (October 1991): 957–76.

5. Harold Alderman and Christina Paxson, ''Do the Poor Insure? A Synthesis of the Literature on Risk and Consumption in Developing Countries,'' in *Economics in a Changing World,* ed. E. Bacha (New York: St. Martin's, 1994), pp. 48–78.

6. In addition to these problems, we have little guidance for determining empirically the set of households belonging to the risk-sharing pool. Townsend, in ''Risk and Insurance in Village India,'' argues that this errors-in-variables problem poses little problem for testing, since it stacks the cards against the risk-sharing hypothesis. He assumes that households in the same village comprise the risk-sharing pool. But Mark Rosenzweig, in ''Risk, Implicit Contracts and the Family in Rural Areas of Low Income Countries,'' *Economic Journal* 98 (December 1988): 1148–70, and others argue that in rural settings it is better for households to participate in far-flung networks than to pool their fortunes with nearby neighbors.

7. Jonathan Morduch's survey, ''Income Smoothing and Consumption Smoothing,'' *Journal of Economic Perspectives* 9 (Summer 1995): 103–14, emphasizes the distinction between trying to cope with shocks before they happen (income smoothing) and reacting to shocks after they occur (consumption smoothing).

8. See n. 6 above. Positive within-village correlation in weather-related shocks, for example, induce large swings in the price of credit if the supply of loanable funds is village based. Unlike credit, interhousehold transfers can cross village boundaries and provide more effective insurance against village-specific shocks. Rosenzweig finds that 60% of the value of private transfers come from outside the village, suggesting that spatial diversification is important for mitigating income risk.

9. Mark Rosenzweig and Oded Stark, ''Consumption Smoothing, Migration, and Marriage: Evidence from Rural India,'' *Journal of Political Economy* 97 (August 1989): 905–27.

10. See Robert E. B. Lucas and Oded Stark, ''Motivations to Remit: Evidence from Botswana,'' *Journal of Political Economy* 93 (October 1985): 901–18. Implicit in much of the analysis of private transfer behavior is the condition that capital markets do not work well for many households. Evidence for the United States (Donald Cox, ''Intergenerational Transfers and Liquidity Constraints,'' *Quarterly Journal of Economics* 105 [February 1990]: 187–217) indicates that private familial transfers appear to be targeted to those who face borrowing constraints. Some developing-country evidence suggests a similar connection. For example, in Peru, the incidence of private transfers mirrors almost exactly the age-earnings profile and is sensitive to current household income (Donald Cox, Zekeriya Eser, and Emmanuel Jimenez, ''Motives for Private Income Transfers Over the Life-Cycle: An Analytical Framework and Evidence for Peru,'' *Journal of Development Economics* [in press]. Impediments to savings (e.g., severe inflation, theft) may exist as well, heightening the importance of old-age support from children (as in John Hoddinott, ''Rotten Kids or Manipulative Parents: Are Children Old-Age Security in Western Kenya?'' *Economic Development and Cultural Change* 40 [April 1992]: 545–65). Without capital market imperfections, it is unlikely that the timing of private transfers would be so important.

11. Alderman and Paxson, p. 28.

12. Rosenzweig.

13. The data are described in detail in Michael Bamberger, Daniel Kaufmann, Eduardo Velez, and Scott Parris, ''Interhousehold Transfers and Survival Strategies of Low-Income Households: Experiences from Latin America, Asia, and Africa'' (World Bank, Washington, D.C., July 1992, mimeographed). An-

other useful reference is Daniel Kaufmann, ''Social Interaction as a Strategy for Survival among the Poor: A Theory and Evidence'' (Ph.D. diss., Harvard University, 1982).

14. An example of how sensitive private-transfer reporting is to survey design can be found in a recent wave of the Panel Study of Income Dynamics in the United States (Survey Research Center, Institute for Social Research, *A Panel Study of Income Dynamics,* Wave XXI [Ann Arbor: University of Michigan Institute for Social Research, 1991]). For each year, the income module of the survey contains a summary question on private transfers received from relatives and friends. About 4% of households report receiving such transfers. The 1988 wave contained a special module with far more detail about the nature of such transfers (their origin, characteristics of donors, and so on). The reporting rate for private transfers in this wave was about four times higher, even though the definition of private transfer receipt was roughly the same.

15. Bamberger, Kaufmann, Velez, and Parris, p. 2-1.

16. See, respectively, Lucas and Stark; Daniel Kaufmann and David Lindauer, ''A Model of Income Transfers for the Urban Poor,'' *Journal of Development Economics* 22 (July/August 1986): 337–50; Luigi Guiso and Tullio Jappelli, ''Intergenerational Transfers and Capital Market Imperfections: Evidence from a Cross-Section of Italian Households,'' *European Economic Review* 35 (January 1991): 103–20; Cox, Eser, and Jimenez; Donald Cox, Bruce Hansen, and Emmanuel Jimenez, ''Are Households Altruistic? Private Transfers in a Laissez-Faire Economy'' (Boston College, Boston, 1996, mimeographed); and Donald Cox, ''Motives for Private Income Transfers,'' *Journal of Political Economy* 95 (June 1987): 508–46.

17. Three of the ''others'' were actually involved with transfers, but they gave and received exactly the same amounts. Since for the most part we are looking at net receipts and gifts, we include these households in the ''other'' category.

18. Strictly speaking, risk sharing need not imply that transfers flow from rich to poor, only that they flow from those with windfalls to those with shortfalls. But if transfers are altruistic, not only would they mitigate risks (as discussed in Gary Becker, ''A Theory of Social Interactions,'' *Journal of Political Economy* 82 [November/December 1974]: 1063–94) but would also help equalize lifetime incomes. So, if risk sharing is rooted in altruism transfers would flow from rich to poor.

19. Another way to characterize the variability in earnings is to consider the coefficient of variation (*cv*) calculated from the three earnings values (current, low, and high). The median *cv* is 0.33, and the *cv*'s at the 25th and 75th quantiles are 0.12 and 0.64, respectively.

20. We focus on a model of transfer incidence, rather than amounts. As is the case with many studies of private transfers, the qualitative choice models are much better determined than those used to explain transfer amounts. For example, ordinary least squares estimated on nonlimit observations for net transfer receipts indicates that the set of covariates X_i is not even close to statistically significant at conventional levels. In what follows, we concentrate on transfer events rather than amounts.

21. Implicit in the ordered probit specification is the idea, supported by our empirical work, that giving and receiving are symmetrical. One instance in which such symmetry would occur is if private transfers were motivated by mutually altruistic feelings so that they flowed from better- to worse-off households. We explore less restrictive specifications below.

22. Cox, Eser, and Jimenez (n. 10 above).

23. Rosenzweig (n. 6 above).