Does poverty entrap?

Stefan Dercon Joseph S. Shapiro [*]

Abstract

We test whether income dynamics over 30 years in rural Indian villages suggest the presence of a poverty trap. A weighted dynamic panel estimate using rainfall as an instrument for lagged income addresses measurement error, attrition bias, and the inherent endogeneity of lagged income while allowing the steady state of income to vary across individuals. Estimates of the model provide robust evidence of a low-level steady state equilibrium but, consistent with other studies, no evidence of a poverty trap. Schooling substantially increases a person's steady state income, however, implying that education provides one secure escape from poverty.

[*]Authors contact: stefan.dercon@economics.ox.ac.uk and j.s.shapiro@lse.ac.uk. We thank Reena Badiani for extensive and insightful help in preparing and processing the data.

1 Introduction

With about 2.8 billion people receiving less than two dollars per day in income (Chen and Ravallion 2001), the issue of whether and why the destitute escape poverty constitutes a central question in economic research. Theories of poverty traps explain why living in poverty at some time causes a person to remain poor in the future, or why a country's poverty causes the country to remain in future poverty (Galor and Zeira 1993, Azariadis and Stachurski Forthcoming). These theories imply stark conclusions: a positive income shock could prevent a person from living in poverty for the indefinite future, while a sufficiently grave negative shock to income could prevent a person from ever escaping poverty. Some such theories assume that a person requires a fixed and indivisible investment to purchase a good like education or credit (Banerjee and Newman 1993); others assume increasing returns to income via nutrition or another means (Dasgupta and Ray 1986); while still others show how leaving the poor without bargaining power can cause the poor not to save (Mookherjee and Ray 2002).

Admittedly imperfect tests of these elegant models have offered little empirical support, however, leaving Dasgupta (1997) to describe that they reside "awkwardly" in development thinking. A model of nutrition poverty traps has received empirical criticism from several studies (Bliss and Stern 1982, Swamy 1997, Rosenzweig 1988), though Dasgupta (1997) argues that they use flawed tests. A theory of fixed costs to entering businesses has received similarly little summport(McKenzie and Woodruff 2003).

Several recent studies have proposed that a poverty trap could arise through a combination of mechanisms, or through some unstudied mechanism. These studies essentially examine whether a regression of some welfare measure (income, consumption, or assets) on its lag has shape that could indicate the presence of a poverty trap. Nonparametric kernel regressions of current on lagged assets using small samples from Kenya, Ethiopia, Madagascar, and South Africa show unstable equilibria over some low values of income that suggest the presence of a possible poverty trap. These studies ignore the endogeneity of lagged income in a dynamic panel models, however, and the potential bias in data obtained from many-year recall questions, limited generalizability of sample sizes under 200 individuals, and bias of bivariate kernel regressions at discontinuities (Fan 1992) give their conclusions limited scope. In higher income areas, studies applying methods with corrections for various econometric challenges in estimating income dynamics to data from China, Eastern Europe, and Urban Mexico have found evidence for some stable low-level equilibria but no evidence of a poverty trap. (Antman and McKenzie Forthcoming, Lokshin and Ravallion 2004, Jalan and Ravallion 2003)

The econometric challenges involved in testing for the presence of poverty traps are not trivial, and most create bias towards failing to reject the hypothesis that poverty does not entrap people. Hence one could reasonably conclude that existing literature fails to establish whether poverty traps actually do not exist or whether available data and methods have inadequate power to detect them. Econometric problems abound. Panel data with short duration–typically less than five years (Dercon and Shapiro 2006)–may not capture the dynamics that ensure poverty's persistence. The nature of a dynamic panel model ensures that regression of income on its one-period lagged value will inflate the effect of lagged income on current income. Measurement error in income creates a mirage of income mobility, so a person whose true income remains constant over time may appear to enter then escape poverty. Individuals who attrit from a panel may have substantial differences from individuals who remain in a panel, and ignoring attrition may overstate or understate mobility.

Existing studies address some but not all of these concerns. Jalan and Ravallion (2003) and Lokshin and Ravallion (2004) use the Arellano and Bond (1991) GMM estimator to identify the association of a cubic polynomial of lagged income with current income. But if measurement error has serial correlation, as at least one U.S. comparison of survey-reported income with independent income reports suggests (Bound and Krueger 1991), then using distant lags of income as instruments for once-lagged income, as the GMM methods do, will overstate mobility. (Antman and McKenzie Forthcoming) for this and other reasons condemn the possibility of using panels for identifying nonlinear income dynamics, and propose instead the use of pseudo-panels to average out measurement error across individuals.

The present study shows how panel methods can address these econometric criticisms and consistently test for the presence of a poverty trap. We test whether a poverty trap characterizes the income dynamics of individuals in an unusually long 30-year panel from six villages in India's semi-arid tropics. We interact rainfall shocks with household characteristics to provide strong instruments for a cubic polynomial of lagged income in a dynamic panel model, obviating the need for GMM methods and addressing the critical problem of measurement error in income. In the absence of exclusion restrictions that would allow reliable estimate of a selection model, we address attrition using weighted least squares (WLS), where a weight equals the inverse of an observation's fitted probability of appearing in the sample. Under a non-trivial identification assumption, WLS can consistently identify regression parameters even in the presence of severe attrition. We also recover individual effects, allowing for individual heterogeneity wherein income dynamics may differ across individuals. Finally, we identify the correlates of these individual effects, revealing individual characteristics which lead to sustained increases in the trajectory of incomes.

After addressing the many statistical problems inherent in testing for the presence of a poverty trap, we find evidence of low-level equilibria which have abated over time, but no evidence of a poverty trap. Changes in permanent individual and household characteristics – education, geographic location, and household head characteristics – can substantially change the trajectory of a person's income. Shocks to income may cause medium-term poverty or wealth due to slow adjustment back to an individual's steady state, but we do not find evidence that positive or negative shocks permanently change a person's income. Given the robust positive association of individual and household education with an individual's fixed effect – interpretable as an individual-specific income trend – it remains very likely that shocks which discourage human capital accumulation (e.g., (Jacoby and Skoufias 1997)) may cause permanent and not transient poverty.

The paper proceeds as follows. Section 2 outlines the econometric obstacles inherent in testing those models of poverty traps which only consider income dynamics. Section 3 describes the 30-year panel data set. Section 4 presents the main results, and section 5 concludes.

2 Econometric Method

An estimate of how lagged income affects current income must address five statistical problems: the endogeneity of lagged income in a dynamic panel model; measurement error in income; individual heterogeneity non-random attrition; and short panel duration. We discuss solutions for each, and we combine responses to these potential biases in the final estimator that we propose and implement.

2.1 Dynamic panels and measurement error

We estimate an AR(1) model where the income y_{it} of person i at time t depends on a linear function of a polynomial of person i's lagged income,[1] and a composite error term with time-invariant and idiosyncratic components ρ_i and v_{it}:

$$y_{it} = \alpha y_{i,t-1} + \rho_i + v_{it} \tag{1}$$

Since lagged income correlates positively with the composite error term $\rho_i + v_{it}$, estimating equation (1) by OLS generates inconsistent estimates of α. A first-differenced version of equation (1) eliminates the individual effect ρ_i:

$$\Delta y_{it} = \alpha \Delta y_{i,t-1} + \Delta v_{it} \tag{2}$$

where $\Delta x_t = x_t - x_{t-1}$. Since y_{t-1} correlates with $v_{i,t-1}$, OLS estimation of equation α in (2) also produces inconsistent estimates (Nickell 1981).

General method of moments (GMM) estimators use lagged income as instruments for $\Delta y_{i,t-1}$. If the error term Δv_{it} in equation (2) lacks second-order serial correlation, and if equation (1) is dynamically complete, then further lags of income and their first-differences can serve as valid instruments for $\Delta y_{i,t-1}$. Under these assumptions, use of all such available

[1] Some studies describe such an estimate as a test of "non-linear income dynamics" (Antman and McKenzie Forthcoming, Jalan and Ravallion 2003). While specifying lagged income as a higher-order polynomial allows current income to vary nonlinearly with lagged income, the regression function itself is linear in the higher-order terms of lagged income.

instruments provides a consistent and efficient general method of moments estimator of α, the parameter of interest (Anderson and Hsiao 1982, Arellano and Bond 1991, Blundell, Bond, and Windmeijer 2000, Bond 2002).

Several existing studies use these GMM estimators with short panels to test whether a poverty trap characterizes income. One of these studies shows overidentification tests of the instruments' validity and a test that the residuals have second-order autocorrelation (Jalan and Ravallion 2003); the other studies do not mention these tests (Lokshin and Ravallion 2004, Antman and McKenzie Forthcoming), and none of the three studies evaluates whether the GMM instruments are weak, an important problem when a dynamic panel has a near-unit root (Stock, Wright, and Yogo 2002, Blundell, Bond, and Windmeijer 2000), and a problem which large samples do not eliminate (Bound, Jaeger, and Baker 1995).

Measurement error also creates a more substantial problem in these papers, since some U.S. data suggest that measurement error in an individual's income has positive autocorrelation across waves of a panel (Bound and Krueger 1991). Antman and McKenzie (Forthcoming) show that in the presence of such measurement error, GMM estimators provide inconsistent estimates of the paramters in equation (2), and severity of the problem persists upon specifying lagged income as a higher-order polynomial.

As an alternative to GMM methods, we propose use of other exogenous excluded instruments for lagged income: lagged rainfall. Rainfall provides a useful instrument for income in analysis of poor agricultural areas.[2]. Since the economies of agricultural villages heavily depend on weather, flood and drought sharply affect the income of most households in the village: agricultural households have lower yield in seasons of extreme weather, households that earn income from agricultural labor find less work in times of extreme weather, and most individuals in these communities depend on good weather for strong income.

Although we only have rainfall data at the village data, we expect that rain affects different households differentially: landowners may benefit most from good rain, while households with more potential workers may also benefit most from rain. We obtain monthly village-level

[2]See, for example, Paxson (1992), Miguel (2005), Miguel, Satyanath, and Sergenti (2004), and the review in **?**

rainfall measures from several mandal-level collection stations, and for the rainfall instruments we interact measures of rainfall shocks with household characteristics. Comparing a variety of specifications of rainfall shows that the most robust relationship between rainfall and contemporaneous income appears when rainfall has positive deviations from its village fixed-effect, representing either bounty crops or flooding, depending on the rainfall's severity. We interact rainfall with household landholdings and the number of children aged 2-8 in the household, both of which show independent, large, and significant relationships with contemporaneous income. Landholding makes intuitive sense, since households with larger plots will receive more benefit from years with good rainfall and more harm from years with flood. The association with children is more surprising. It may however arise because children leave school to work in times of duress (Jacoby and Skoufias 1997).

For each rainfall instrument $Z_{i,t-1}$, we require two conditions:

$$cov(Z_{i,t-1}, \Delta y^*_{i,t-1}) \neq 0 \tag{3}$$

$$cov(Z_{i,t-1}, \Delta v_{i,t}) = 0 \tag{4}$$

where $y^*_{i,t-1}$ denotes true, unobserved income. Condition (3) requires that the instrument strongly correlate with true lagged income, while condition (4) requires that the instrument not correlate with the measurement error of lagged income or with other components of the structural equation error. One could interpret invalidity of the GMM estimators as failure to meet these two criteria.

2.2 Individual heterogeneity

It is possible that individuals have multiple equilibria for income, or that poverty creates a trap for some but not all individuals. Fixed individual factors – education, geographic location, and others – may affect the trajectory of an individual's income. Since these fixed factors may correlate with income and hence bias regression estimates, equation (2) uses first-differencing to eliminate these fixed effects.

But the effects themselves have economic interest, and recovering these parameters allows us to observe the correlation between observable individual fixed characteristics and the part of an individual's income trajectory which does not depend on short-term income dynamics. Put another way, some individuals may have exogenous reasons that cause their incomes to increase by a certain amount each year, and the individual effect contains these factors. **?** include a fixed effect in their model of income dynamics but eliminate it by first differencing and leave it unestimated, though they do depict recursion diagrams for different percentiles of the income distribution. Antman and McKenzie (Forthcoming) highlight the economic importance of the fixed effect and show how it shifts a person's income trajectory, but do not examine the correlates of the fixed effect.

Since the idiosyncratic errors have mean zero across the population, we estimate the individual effect by the deviation of an individual's mean outcome from the predicted mean (Antman and McKenzie Forthcoming):

$$\hat{\alpha}_i = \overline{Y}_i - \hat{\beta}_1 \overline{Y}_{i,t-1} - \hat{\beta}_2 \overline{Y}^2_{i,t-1} - \hat{\beta}_3 \overline{Y}^3_{i,t-1}$$

where we average the dependent and independent variables across the years in which they would appear if we had not first-differenced the model.

To estimate the correlates of these fixed effects, we regress them on a vector Z_i of fixed individual characteristics:

$$\alpha_i = \phi_0 + Z_i \phi_1 + \epsilon_i$$

The parameters ϕ_1 show the correlation of individual characteristics with the fixed effects. A positive association $\phi_j > 0$ for some element j of the vector Z_i implies that phi_j given an individual continuously increasing income regardless of shocks. Since we observe income and several fixed characteristics each wave only at the household level, all regression estimates in the paper use standard errors robust to heteroskedasticity and serial correlation within household-years.

2.3 Attrition and panel duration

If attrition randomly removed observations from each wave of a survey, then attrition would only decrease the precision of estimated parameters. But attrition occurs for non-random reasons: individuals leave villages due to fixed and time-variant characteristics like shocks and job opportunities that cause a person or household to move. Since attrition may correlate with observed and unobserved characteristics which influence income, estimating equation (2) by any method without addressing attrition can produce an inconsistent estimate of α.

The problem has similarity to selection models where an econometrician observes a response variable for only a subset of a cross-sectional survey, and indeed the first selection models explicitly discussed their potential for addressing panel attrition (Heckman 1979). Lokshin and Ravallion (2004), for example, simultaneously estimate a GMM regression with an equation where baseline household composition, education, and location variables serve as instruments for selection in a regression of income on its lag.

But it is difficult to argue that these or any variables affect selection but not income, as one requires for consistent estimates of the parameters in equation (2). Other authors suggest more detailed procedures which estimate selection models for each time period, but these too require exclusion restrictions (Wooldridge 2002, 581-586). Any selection model requires observation of factors which vary across individuals, affect the probability of disappearing from the panel, and are independent of income. Fitzgerald, Gottschalk, and Moffitt (1998) propose a cost-benefit model wherein individuals consider the net value of participating in a survey, and interview duration or interview payments affect the value of the survey. ICRISAT offers no such variation of interview payments across respondents. Furthermore, ICRISAT attrition occurs rarely due to refusal and more often due to migration and death. We conclude that no variables from available data can credibly satisfy the required exclusion restriction.

Weighted least squares (WLS), sometimes called inverse probability weighting, eliminates the need for exclusion restrictions, though WLS does require the model to satisfy non-trivial identification assumptions (Fitzgerald, Gottschalk, and Moffitt 1998, Wooldridge 2000, Wooldridge 2002). In some cross-sectional surveys where individuals refuse to participate, surveyors construct weights to represent an individual's probability of participating in the survey, and inference using these surveys weights responses by the inverse of these probabilities.[3] WLS in the present context plays a similar role.

The attrition-corrected results use a random population sample at time $t = 1$ and define the selection variable s so an observation appears in a wave if and only if $s_{it} = 1$. We treat attrition as an absorbing state, in that an individual who attrits from the sample at time t does not reappear, so $s_{it} = 1 \rightarrow s_{rt} = 1 \quad \forall\ r < t$. Although such an approach forces us to drop observations that vanish for one or more rounds then reappear, this loss of precision and information allows us to use a potentially consistent estimator of regression parameters even in the face of substantial attrition.

For this correction to provide a consistent estimator, we must assume that a set of baseline covariates z_{i1} has enough predictive power that outcomes and covariates at any future time are independent of selection:

$$P(s_{it} = 1|y_{it}, x_{it}, z_{i1}) = P(s_{it} = 1|z_{i1}) \tag{5}$$

Writers generally describe assumption (5) as selection on observables or ignorability of selection. To consistently estimate equation (**??**) while assuming selection on observables, for each time period, we estimate a probit of s_{it} on z_{i1} using all observations that appear in the baseline survey. We obtain estimated probabilities $\hat{p}_{it}$ for each time period and individual, then weight the regression by the inverse of these fitted probabilities, equivalent to minimizing the following function:

$$\sum_{i=1}^{N}\sum_{t=1}^{T}\left(\frac{s_{it}}{\hat{p}_{it}}\left(\Delta y_{it} - \alpha\Delta y_{i,t-1}\right)\right)^2 \tag{6}$$

An analogy argument can show that, under assumption (5), equation (6) produces a consistent estimator which has a probability limit identical to an unweighted regression if the data had no attrition (Wooldridge 2000, Wooldridge 2002).

The dataset that we use has the advantage of unusually long duration: 30 years, a

[3] A separate reason for constructing and using weights arises when survey design and not respondent refusal or absence causes individuals to have unequal probabilities of appearing in the data.

length parallelled by only a small handful of existing datasets (Dercon and Shapiro 2006). Unfortunately these data have a gap of about fifteen years: we have annual surveys for every year between 1975 and 1983, then surveys again for 2001-2005 (see the following section).[4] In the effort to examine the long-run factors that influence poverty and welfare, such long-term panel duration provides critical information on income dynamics. But given our focus on income dynamics, ignoring this gap in the middle and treating 1983 as if it preceded 2001 will yield problematic estimates for later years.

To address the 1984-2000 gap, we use one-year lags of variables for all years. Taking the first difference of a model which includes uses a one-period lag of the dependent variable as a regressor forces us to use the first two waves of the panel only in calculating first differences and lags, but not for outcomes. Hence we use the first difference of the dependent variable (current income) from nine waves of the panel (1977, 1978, 1979, 1980, 1981, 1982, 1983, 2003, 2004), while we use the first difference of the independent variables (lagged income) from a different set of nine waves (1976, 1977, 1978, 1979, 1980, 1981, 1982, 2002, 2003). Since we use lagged rainfall as an instrument rather than the many lags of income which the GMM estimators use, the 1984-2001 gap creates no other obstacles in estimating the dynamic panel model.

3 Data: the 30-year ICRISAT Panel

The International Crops Research Institute for the Semi-Arid Tropics (ICRISAT) near Hyderabad, India, collected annual surveys between 1975 and 1984, then for the same households in the period 2001-2005. The core data included 240 households from six villages in India's semi-arid topics: the villages of Aurepalle and Dokur in the Mahbubnagar District of the Indian state of Andhra Pradesh; the villages Shirapur and Kalman in the Sholapur District of the state of Maharashtra, and the villages Kanzara and Kinkheda in the Akola District of Maharashtra. Villagers generally work in dryland farming, with limited irrigation (Badiani, Dercon, Krishnan, and Rao 2006).

[4] Income data in year 1984 included only a small subset of individuals. A 1992 round of income data included few individuals and had different methodology than other years, while 2005 data are still being processed.

For the early data collection, interviewers lived in the villages and interviewed households every 3-4 weeks to obtain income information. The more recent data use one interview per year for 2001-2003 and two per year for 2004. A tracking survey allowed followup of individuals interviewed in the 1975-84 rounds. Walker and Ryan (1990) provide detailed description of the early survey rounds and research stemming from them, while Badiani, Dercon, Krishnan, and Rao (2006) provide an appendix with further detail on the recent data collection.

4 Results

4.1 Income trends

Descriptive statistics show substantial attrition but large increases in income for remaining households (Table 1). From an initial set of 238 households in 1427, the number 30 years later decreased to 1283. This final total includes newborns and individuals in splitoff households, but not completely new households added to the survey.

The early part of the panel had a slight upward trend in income, partly due to particularly low income in the first two years of the survey. Mean income increased substantially for the second part of the panel and increased by 70 percent between 1983 and 2001, and finding that coheres with the regression results of **?** using the same data. Attrition affected the number of individuals in the survey, but did not substantially change their representation across the six villages, as each village had between 15 and 18 percent of respondents in every round (Table 1).

4.2 Regression results

We discuss first-stage and second-stage TSLS estimates, then examine the correlates of individual fixed effects. We then present probits predicting the probability of attriting from the sample.

We measure rainfall shocks as the deviation of December rainfall from its village fixed effect, defined only for positive deviations. Although the Kharif monsoon represents the largest source of annual rain in these villages and December falls during the smaller Rabi monsoon, a variety of specification checks showed that December rainfall displays the strongest relationship with contemporaneous income. The correlation may arise because flooding most often occurs in December, because household least often expect substantial rainfall in December, or for other reasons.

We regress lagged annual income, its square and cube on rainfall in December of the same year, and on the interaction of December rainfall with the household's landholding and number of children. For each variable we estimate one specification which controls for year dummies and another which does not. As in the main equation, the first-stage equation is first differenced to remove fixed effects.

All three variables have strong correlations with lagged income: each additional millimeter of rainfall correlates with a 6 rupee drop in income, suggesting that these shocks measure unexpected harmful excess rainfall rather than a positive and adequate rain. Owning land shields individuals from the rainfall shock, perhaps because the shock has the greatest effect on day laborers. Having young children in the household exacerbates the effect of flood. Adding year indicators slightly decreases the coefficients but does not affect their statistical significance. For explaining the square and cube of income, regression coefficients maintain their signs and statistical significance but increase substantially in magnitude.

For regressing income on these instruments, the excluded instruments have a joint F-statistic of 7.6 to 8.5 for income, 6.25 for income squared, and 3.9 to 4.3 for income cubed. Conventional critical values would suggest that these constitute strong instruments and avoid the potential bias that can arise from weak correlation between excluded instruments and endogenous variables in reduced form equations (**?**, Stock, Wright, and Yogo 2002). Stock and Yogo (2003), however, emphasize that strength of instruments can have different meaning than merely statistically significant positive association with the endogenous variable. Their earlier work proposed a critical value for the joint significance of excluded instruments of 5, while their later work suggests that the critical value must depend on the number of endogenous variables: for two endogenous variables, they suggest critical values between 9 and 12. Based on the early criteria, the instruments for income and its square are strong while the instruments for income cubic are marginally weak; based on the newer criteria, the instruments for linear income are marginally weak while instruments for income's square and cube are more weak.

Given the availability of rainfall data at the village rather than individual level, these constitute fairly strong predictors of a polynomial of income. In our estimates of the structural equations, the endogenous regressors have joint statistical significance at 97 percent confidence without including year indicators and 92 percent confidence including year indicators, implying that even given the conservative inference of adjusting standard errors for correlation within household-year cells, implying that these instruments are strong enough to allow lagged income to maintain fairly strong and robust association with current income. While stronger instruments might provide more precise results, we obtain extremely consistent results across a variety of specifications, offering some confidence in the validity of inference based on these instruments.

The second stage results offer the surprising implication that for the average individual, an increase in current income slightly decreases future income. The associated figure shows the explanation: for the average individual, income has a low steady state (Figure 1). Individuals above the steady state in a given period on average move towards it in the following period, making high incomes anomalous and leading to a negative coefficient on lagged income. The result persists upon adding year controls or specifying lagged income as a higher-order polynomial.

The figure shows distributions of the recursion diagram for the 5th, 50th, and 95th percentile of the individual fixed effect. Given the low value of the steady state, and impossibility of persistent incomes below zero,[5], low percentiles of the fixed effect have values only slightly below the median. The fixed effect, however, has a long right tail, and the upper values have extremely high incomes. Most individuals have a steady state income just below the poverty line, while the more fortunate individuals have steady state incomes far above the poverty line.[6]

[5]A few observations in the data have negative values of income due to borrowing in bad years, but these values become positive subsequently

[6]In 1993 an Expert Group of the Government of India suggested a consumption poverty line of 49 Rs/month in 1973-74 prices, equivalent to 630 Rs/year in 1975 prices. Given our use of income rather than income, we

To retrieve some idea of the causes of these different steady states, we recover the fixed effects and regress them on time-invariant observable characteristics, as described in the previous section. The results are generally consistent across specifications (Table 4). An individual's schooling has about ten times as large an association with persistent income increases than a household head's schooling does. Every year of a person's schooling associates with an additional Rs 90-118 increase in income, while every year of a household head's schooling associates with only an additional 9-15 Rs additional annual increase in income. Despite their geographic proximity, these villages have substantial heterogeneity in soil and other characteristics (Walker and Ryan 1990). Correspondingly, individuals in different villages have different income trajectories: Village A, the reference, improved least, while village D improved comparatively the most, and the effects had large magnitude.

To address attrition, we estimate the probability of attriting from the sample in each year as a function of baseline covariates: we consider landholding, income, household head age, household head education, and the individual's age in 1975. Income and household head characteristics strongly predict attrition in most years, while landholding has a statistically significant effect only in later years. The marginal effects have varying magnitude–every year of household head age increases probability of attrition in later years by 2 to 6 percent, while every year of household head schooling increases the probability of attrition by only 1 percent. A likelihood ratio test soundly rejects the hypothesis that the regressors lack joint statistical significance (Table 5).

We obtain each individual's fitted probability of appearing in the sample according to these probits, then reestimate results using WLS with weights equal to the inverse of these fitted probabilities.

The results change very little upon correction for attrition bias (Tables 6-8). The average individual still has a low steady-state income, though the right tail of the individual effects has less skewness, shown by the decrease in the 95th percentile of the fixed effects upon controlling for attrition. Education maintains its robust relationship with the trend of income, and a

---

follow **?** in using a poverty line of 500 Rs/year.

household head's education maintains its somewhat weaker relationship.

5 Conclusions

A variety of theories suggest why a person who becomes poor at any time will remain poor indefinitely. Most such theories focus on a technology with increasing returns to scale which arises from a particular social mechanism–nutrition, education, fixed costs to entering a business, or another. The ideas of poverty traps that arise from these theories constitute a central theory of development economics at both the micro and macro levels. But these theories have received extremely little empirical support, possibly due to econometric pitfalls in the methods underlying the relevant empirical studies, as Dasgupta (1997) argues occurs for tests of the nutrition-efficiency wage theory, or possibly because no poverty trap in fact exists.

The large number of people in extreme penury constitutes only one reason underpinning the importance of understanding whether and why the destitute escape poverty. The presence of poverty traps would also implies a startling policy conclusion: a small transfer to a poor individual or household could change that person from low- to high-level equilibrium and permanently remove a person from poverty.

Since most existing theories of poverty traps assume some form of fixed investment cost, or increasing returns to assets or income, we examine whether income dynamics give evidence of increasing returns. A variety of econometric problems arise in this analysis: lagged income is inherently endogenous in a dynamic panel model; measurement error in income will cause OLS or GMM estimates to understate income's persistence; individual heterogeneity may disguise the fact that some individuals face a poverty trap even though the average individual does not; non-random attrition may remove individuals who escape poverty from the survey biasing estimates in favor of finding a poverty trap; and short panel duration may give inadequate time to observe sufficient movement in income.

The bivariate kernel regressions or GMM methods that existing papers use address some but not all of these pitfall. We instead interact rainfall with household landholding and composition variables to obtain valid and fairly strong and instruments for a cubic polynomial of contemporaneous income, with a first-stage F statistic between 5 and 9 for the excluded instruments. Given the implausibility of variables that might satisfy an exclusion restriction necessary for estimating a selection model, we estimate the probability that an observation attritts in each year as a function of baseline observable variables, then weight regression estimates by the inverse of these probabilities. Under a non-trivial ignorability assumption, this method consistently identifies regression parameters even in the presence of severe and non-random attrition. The limited availability of covariates for predicting attrition allows us to partly though not completely erase its effects.

We apply to estimator an unusually long 30-year panel dataset from several villages in India's semi-arid tropics. The early rounds of this ICRISAT data underpinned some of the most influential papers in development economics, and annual revisits to these villages between 2001 and 2005 has allowed followup of the original households.

Our results consistently show the presence of low-level equilibria, but no evidence of a poverty trap. Regardless of the specification of income, we identify a steady state value of income slightly below the poverty line of 500 Rupees. An individual who experiences a positive income shock will eventually return to this steady state, while a person affected by negative income shock will eventually return as well. While we do not estimate the speed of transition dynamics, our results do not eliminate the possibility that shocks have long-lasting effects; we merely find that the effects are not permanent.

Individual fixed effects contain all time-invariant factors which give income a positive trend that is independent of shocks. By retrieving these fixed effects and regressing them on individual characteristics, we identify correlates of income's slope. Although we use few regressors, a person's education has large and robust association with a positive income slope: each year of schooling associates with a statistically significant 100 Rupee per year increase in the *trend* of income. As in most associations of education with income, this result may represent unobserved factors like ability which correlate with both education and income. Village effects also have large magnitude, and village of residence combined with a person's and household head's education explain 16-17 percent of variation in the individual fixed effects. Overall, these results do not support theories of poverty traps based in short-term income dynamics. But we do find robust evidence of low-level equilibria in these villages, reinforced by low levels of education.

References

Anderson, T., and C. Hsiao (1982): "Formulation and Estimation of Dynamic Models Using Panel Data," *Journal of Econometrics*, 18(1), 47–82.

Antman, F., and D. McKenzie (Forthcoming): "Poverty traps and Nonlinear Income Dynamics with Measurement Error and Individual Heterogeneity," *Journal of Development Studies*.

Arellano, M., and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economic Studies*, 48(2), 277–297.

Azariadis, C., and J. Stachurski (Forthcoming): "Poverty Traps," *Handbook of Economic Growth*.

Badiani, R., S. Dercon, P. Krishnan, and K. Rao (2006): "Changes in Living Standards in Villages in India 1975-2004: Revisiting the ICRISAT village level studies," .

Banerjee, A., and A. Newman (1993): "Occupational Choice and the Process of Development," *Journal of Political Economy*, 101(2), 274–298.

Bliss, C. J., and N. H. Stern (1982): *Palanpur: The economy of an Indian village*. Oxford University Press, Oxford.

Blundell, R., S. Bond, and F. Windmeijer (2000): "Estimation in dynamic panel data models: improving on the performance of the standard GMM estimators," .

Bond, S. R. (2002): "Dynamic panel data models: a guide to micro data methods and practice," *Portuguese Economic Journal*, 1(2), 141–162.

Bound, J., D. A. Jaeger, and R. M. Baker (1995): "Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak," *Journal of the American Statistical Association*, 90(430), 443–450.

Bound, J., and A. B. Krueger (1991): "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?," *Journal of Labor Economics*, 9(1), 1–24.

Chen, S., and M. Ravallion (2001): "How did the world's poorest fare in the 1990s?," *Review of Income and Wealth*, 47(3), 238–300.

Dasgupta, P. (1997): "Nutritional Status, the Capacity for Work, and Poverty Traps," *Journal of Econometrics*, 77.

Dasgupta, P., and D. Ray (1986): "Inequality as a Determinant of Malnutrition and Unemployment: Theory," *The Economic Journal*, 96(384).

Dercon, S., and J. Shapiro (2006): "Moving On, Staying Behind, Getting Lost: Lessons on poverty mobility from longitudinal data.," CSAE Working Paper, University of Oxford.

Fan, J. (1992): "Design-adaptive Nonparametric Regression," *Journal of the American Statistical Association*, 87(420).

Fitzgerald, J., P. Gottschalk, and R. Moffitt (1998): "An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics," *Journal of Human Resources*, 33(2), 251–299.

Galor, O., and J. Zeira (1993): "Income Distribution and Macroeconomics," *Review of Economic Studies*, 60(1), 35–52.

Heckman, J. J. (1979): "Formulation and Estimation of Dynamic Models Using Panel Data," *Econometrica*, 47(1), 153–162.

Jacoby, H., and E. Skoufias (1997): "Estimating the Return to Schooling: Progress on Some Persisitent Econometric Problems," *Review of Economic Studies*, 64.

Jalan, J., and M. Ravallion (2003): *Insurance Against Poverty*chap. Household Income Dynamics in Rural China. Oxford University Press.

Lokshin, M., and M. Ravallion (2004): "Household Income Dynamics in Two Transition Economies," *Studies in Nonlinear Dynamics & Econometrics*, 8(3).

McKenzie, D., and C. Woodruff (2003): "Do entry costs provide an empirical basis for poverty traps?," BREAD Working Paper No. 020.

Miguel, E. (2005): "Poverty and Witch Killing," *Review of Economic Studies*, 72, 1153–1172.

Miguel, E., S. Satyanath, and E. Sergenti (2004): "Economic Shocks and Civil Conflict: An Instrumental Variables Approach," *Journal of Political Economy*, 112, 725–753.

Mookherjee, D., and D. Ray (2002): "Contractual Structure and Wealth Accumulation," *American Economic Review*, 92(4), 818–849.

Nickell, S. (1981): "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49(6), 1417–1426.

Paxson, C. H. (1992): "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," *The American Economic Review*, 82(1), 15–33.

Rosenzweig, M. R. (1988): *Handbook of Development Economics*chap. Labour markets in low-income countries. North-Holland, Amsterdam.

Stock, J. H., J. H. Wright, and M. Yogo (2002): "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics*, 20(4), 518–529.

Swamy, A. V. (1997): "A simple test of the nutrition-based efficiency wage model," *Journal of Development Economics*, 53, 85–98.

Walker, T. S., and J. G. Ryan (1990): *Village and Household Economies in India's Semi-Arid Tropics*. Johns Hopkins Press, Baltimore.

WOOLDRIDGE, J. (2000): "Inverse probability weighted M-estimators for sample selection, attrition, and stratification," Institute for Fiscal Studies Working Paper CWP11/02.

WOOLDRIDGE, J. M. (2002): *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA.