Chapter 9
Panel Data

9.1 Introduction

Most household survey data come from a single cross-section of households surveyed at a single point in time. This is useful if the purpose is to get a snapshot of income or poverty, and it does allow for a detailed analysis – for instance, of the proximate determinants of health or malnutrition or income. However, it is rarely possible to get an adequate appreciation of dynamic effects – how incomes in a household rise and fall, how households form and re-form – without panel data.

When a survey is repeated over time, so we have multiple observations for the same person (or household, or firm, or village), we have panel data. In this chapter we first summarize the types of panel data, and review their strengths and weaknesses. We then look at what can be learned from transition matrices, review an application to the growth of household enterprises, and discuss the econometric issues that arise when working with panels. The last example in the chapter uses panel data to measure the effects of microcredit in Thailand.

9.2 Types of Panel Data

The commonest form of panel is based on interviewing households over more than one round of a sample survey. For instance, the 1993 Vietnam Living Standards Survey interviewed 4,800 households. The 1998 round tried to contact 4,704 of these, and succeeded in interviewing 4,305 of the original households. It also added 1,694 new households, so that a total of 5,999 were surveyed in 1998.

This case illustrates several of the features of panels. Some of the initial households were not contacted again, in a bid to improve the efficiency of the sampling. This in turn required some adjustments in the sample weights. Of the 4,704 households to be resurveyed in 1998, 399 could not be contacted, representing an annual attrition rate of 1.8%. And the addition of new households to the sample in 1998 was, in part, due to a desire to maintain the representativeness of the sample.

D. Haughton and J. Haughton, *Living Standards Analytics*, Statistics for Social and Behavioral Sciences, DOI 10.1007/978-1-4614-0385-2_9,

Occasionally, a sample of households is surveyed several times. The most famous example, in developing countries at least, may be the 240 households that were surveyed by the Institute for Crop Research in the Semi-Arid Tropics (ICRISAT) in southwest India annually from 1975 to 1985. A more recent important multi-year panel is that collected by Robert Townsend and his collaborators in central Thailand: the initial survey of households was undertaken in 1997, on the eve of the Asian financial crisis, and there have been annual rounds of surveys since then (see http://cier.uchicago.edu/data/ for details, and access to much of the data).

It is much more common to find rotating panels, where some fraction of the households are surveyed for two rounds, and then rotate off. For instance, the first living standards survey in the Côte d'Ivoire covered 1,600 households in 1985; the following year, about 800 of these households were resurveyed, and 800 new households were added to the roster of those interviewed. The use of rotating panels helps reduce respondent fatigue, and limits the loss of representativeness.

Although the household is nearly always the unit that is (re-) sampled, there have been some exceptions. The Peru Living Standards Survey of 1990 returned to the 1,280 *dwellings*, rather than households, that had been surveyed in Lima in 1985–1986. Of the 1,052 interviews, 745 were for the same households. And the much-used Panel Survey on Income Dynamics in the USA began with a sample of 4,800 *individuals* in 1968, and has resurveyed them regularly ever since.

9.3 Why Panel Data?

The most important reason for collecting panel data is to be able to measure transitions over time (Haughton and Khandker 2009). Only with panel data could we, for example, accurately determine how many people move into and out of poverty between one year and the next. Thus panel data are essential for the understanding of poverty dynamics.

Table 9.1 illustrates this nicely (Haughton et al. 2001). It is based on the panel data component of the Vietnam Living Standards Surveys of 1993 and 1998. Households are sorted into expenditure per capita quintiles for each year, and

Table 9.1 Expenditure quintile transition matrix, Vietnam, 1993–1998

	No. of households, by expenditure/capita quintile, 1998					
	Poor	Poor-mid	Middle	Mid-upper	Upper	Total
Quintile 1993						
Poor	**384**	216	127	54	9	**791**
Poor-mid	193	**264**	223	120	32	**832**
Middle	100	183	**234**	254	85	**856**
Mid-upper	38	127	217	**301**	205	**888**
Upper	12	35	100	209	**550**	**906**
Total	**727**	**825**	**901**	**938**	**881**	**4,272**

Source: Vietnam Living Standards Surveys of 1993 and 1998
Note: Totals in each quintile are unequal because of differential attrition in the sample between 1993 and 1998. Bold values give totals (bottom row) and diagonal elements

Table 9.2 Chronic, persistent, and transient poverty, Chinese Provinces, 1985–1990

	Chronically poor				
	Persistently poor $c_t \leq z, \forall t$	Not persistently poor $\bar{c} \leq z, c_t > z$ for some t	Transient poor (and not chronically poor) $\bar{c} > z, c_t < z$ for some t	Never poor $c_t > z, \forall t$	Percentage of poverty (measured by P_2) that is transient
Total sample	6.2	14.4	33.4	46.0	49
Guangdong	0.4	1.0	18.3	80.3	84
Guangxi	7.1	16.1	37.4	39.4	49
Guizhou	11.9	21.2	40.2	26.7	43
Yunnan	4.9	18.0	35.6	41.5	57

Source: Jalan and Ravallion (1998)
Notes: c_t is the consumption by a person in time t, z is the poverty line, and $\bar{c}$ is mean consumption of the person over the time period under study. P_2 is the squared poverty gap measure of poverty (see Chap. 10 for details)

these are then cross-tabulated to create a *transition matrix*. For this we may see, for instance, that 550 of the 4,272 households were in the top quintile in both years.

Taken at face value, this matrix shows considerable economic mobility. Of the 790 households that were in the poorest quintile in 1993, more than half (406) had moved a higher quintile by 1998, and 200 of these had jumped two quintiles or more.

Another interesting example comes from Jalan and Ravallion (1998), who have information on a sample of households in four large provinces in China for each year from 1985 to 1990. From these data, one may classify households into:

- The *persistently poor*, who are poor every year, so consumption (c_t) is always below the poverty line (z)
- The *chronically but not persistently poor*, who are poor on average ("chronically"), a but occasionally rise above the poverty line
- The *transient poor*, who are not poor on average, but who dip into poverty from time to time
- The *never poor*

Table 9.2 summarizes the main findings. A striking feature of these numbers is that many people are only poor from time to time. These households may need insurance more urgently than income support, unlike the persistently poor.

If economic mobility is high, it becomes more difficult to target just the chronically poor. It also paints a picture of income distribution that is highly dynamic. However, these measures almost certainly overstate the extent of economic mobility, possibly quite seriously. The problem arises because of measurement error. We know (see Chap. 10) that income and expenditure are measured imperfectly, and the ratio of noise to signal is only amplified when we look at changes in these magnitudes.

By definition, we cannot know the precise extent of measurement error. Nonetheless, there is some evidence that it is important. Breen and Moisio (2004), using data from the European Community Household Panel for ten countries, estimate a latent class model that tries to correct for measurement error. They find that more people are persistently poor, and fewer people move into and out of poverty, than the survey data suggest. The effects can be large: using a poverty rate that is 60% of median net income, they find, for example, that 7.3% of the Danish population is persistently poor, and not 3.7% (as a straightforward reading of the survey data would imply). Similarly, Lee et al. (2010), using a simulation approach, estimate that the rate at which people move out of poverty from 1 year to the next in South Korea is only three-quarters as large as the raw survey data would indicate.

Panel data have some other advantages. Econometrically, they help us reduce the effects of unobserved heterogeneity, an issue we return to below. They may allow for more powerful measures of impact evaluation when they permit double differencing.

Panel data also allow one to compare the means of variables more precisely. Let X_1 be a vector of observations on a variable, such as income per capita, from a first sample, and X_2 be the observations from a second sample. We want to see whether $\bar{X}_1$ and $\bar{X}_2$ are statistically significantly different, and for this we need $\mathrm{var}(\bar{X}_1 - \bar{X}_2)$. Ignoring other complications, such as clustering, we have $\mathrm{var}(\bar{X}_1 - \bar{X}_2) = \mathrm{var}(\bar{X}_1) + \mathrm{var}(\bar{X}_2) - 2\mathrm{cov}(\bar{X}_1, \bar{X}_2)$. If the samples are independent – for example, from two unrelated surveys – then $\mathrm{cov}(\bar{X}_1, \bar{X}_2) = 0$. However, if the data are drawn from a panel, it is likely that $\mathrm{cov}(\bar{X}_1, \bar{X}_2) > 0$, which reduces $\mathrm{var}(\bar{X}_1 - \bar{X}_2)$, and makes the test of differences more powerful.

9.4 Why Not Panel Data?

Helpful as panel data can be, there are some drawbacks. The two biggest problems are attrition, and nonrepresentativeness.

Attrition occurs when some of those who were interviewed in the first round drop out of the panel in the second or subsequent rounds. There are a number of possible reasons: the household may have dissolved, members may have died, the family may have moved, the household is no longer willing to respond to the survey, and so on. The trouble is that those who drop out are unlikely to be representative of the initial sample; older households are more likely to die off, for instance. The 1998 Vietnam Living Standards Survey was able to interview only 4,305 out of the 4,704 they wanted to resample from 1993; this represents an attrition rate of 1.8% per year, which is actually rather modest. It requires good record keeping, and some persistence, to keep attrition low.

The second problem with a panel is that as it ages, it becomes less representative of the sampling frame, even if there is no attrition. The reason is that a panel does not, by definition, add newly formed households, and so gradually reflects only older, and more stable, households. This is not an insuperable problem, provided that one is willing to “top up” the panel from time to time, or use rotating panels.

9.5 Application: The Birth and Growth of NFHEs

During the 1990s the Vietnamese economy grew rapidly, with GDP rising by about 7% per year. Yet some observers (e.g., Perkins 1994) worried about the potential lack of private enterprises, and particularly of non-farm household enterprises (NFHEs). Vijverberg and Haughton (2004) address the issues of whether NFHEs in this period were up to the task of spawning enough firms with promise, of creating jobs in their own right, and of fostering upward income mobility.

The analysis is based on the Vietnam Living Standards Surveys (VLSS) of 1993 and 1998. A good place to start is by noting that the proportion of adults working in an NFHE fell from 25.7% in 1993 to 23.7% in 1998, although the proportion of those for whom this was the sole source of income rose from 9.5% to 10.4% over the same period.

An affluent household is more likely to participate in an NFHE than a poor one. This shows clearly in Table 9.3, where we see that about one in three chronically poor households, but more than one in two affluent ones, had someone working in an NFHE. Here, the chronically poor are those who were in the bottom 60% of the expenditure per capita distribution in 1993 and the bottom 40% in 1998; affluent households are those who were in the top two quintiles in both years. The computations are based on the panel of 4,304 households surveyed in both 1993 and 1998.

An interesting feature of this study is the way in which the authors created a panel *of non-farm household enterprises*. Both VLSS surveys asked households about the NFHEs in which they participated, including the age of the enterprise, the sector in which it operated, turnover, and so on. In principle, by matching households and their associated enterprises it should be straightforward to construct a panel of enterprises. In practice it was much more difficult; the authors write (p. 103), "the 1997–98 round uses a different set of industrial codes. The respondents are decidedly imprecise about the enterprise's age. There are changes in the identity of the person who is most knowledgeable ... a household could list up to three enterprises in 1992–93 and up to four in 1997–98." Enterprises are matched over time based on the survey information on enterprise age, industry

Table 9.3 Reliance on non-farm household enterprises, Vietnam

	% of households with a non-farm household enterprise	
	in 1993	in 1998
Chronically poor households	35.6	35.0
Affluent households	58.0	54.9

Source: Vijverberg and Haughton (2004), based on Vietnam Living Standards Surveys of 1993 and 1998

Table 9.4 Accounting for the panel enterprises

	1993	1998	Type of ent.
Total enterprises surveyed	2,795	3,493	
– not included in the 1998 sample	47		
– not included in the 1993 sample		1,042	
– dropped out of the sample in 1997 (attrition)	267		Attrited
= **Enterprises potentially matchable**	2,481	2,397	
– household had no enterprise in 1998	764		Terminated
– household had no enterprise in 1993		701	Startup
= **Enterprises potentially in panel**	1,717	1,696	
– household has another enterprise in 1993 but not in 1998	83		Terminated
– household has another enterprise in 1998 but not in 1993		96	Startup
– no match at all among industry code, entrepreneur, age	321		
–		306	Startup
– manual inspection found no possible match	345		Terminated
–		326	Startup
= **Matched**	968	968	Panel
of which: automatic match between 1993 and 1998 enterprises	514	514	
manual match between 1993 and 1998 enterprises	454	454	

Source: Vijverberg and Haughton (2004), Table 6

code, and the identity of the entrepreneur, but many of the matches had to be made by hand, given the imprecision in the data.

Of the 2,795 NFHEs reported in the 1993 survey, and 3,439 in the 1998 survey, only 968 could be matched to form a panel. Table 9.4 accounts for the problem. Thus attrition bias is a potential concern, although it does not appear to have been too damaging for the analysis in practice.

Vijverberg and Haughton first ask what predicts whether a firm survived between 1993 and 1998, using a logistic regression model (where the dependent variable is 1 if an enterprise survives) to tease out the effects. Overall, 39% of the enterprises surveyed in 1993 survived, in that they were identified as still operating in 1998. An NFHE was more likely to survive if it was larger and older in 1993, run by a woman, or operated by a prime-age entrepreneur.

The performance of NFHEs over time may be summarized with the help of Table 9.5, where the firms are divided into quintiles by enterprise income ("net revenue"). About half of the top-quintile firms in 1993 were still operating in 1998, and over half of these survivors were still in the top net-income quintile. Nothing succeeds like success! Firms with low net income in 1993 – in the bottom two quintiles – were only half as likely to be operating still in 1998.

The analysis suggests that NFHEs play a role during the transition from an agrarian to an urban society. NFHEs are relatively unimportant in poor rural areas, where there is a shortage of credit, weak infrastructure, low levels of education, and a limited amount of local demand for goods and services. NFHEs are also relatively unimportant in the most affluent urban areas, where wage labor may offer a better alternative to the family business. In between these extremes, NFHEs are important. Vijverberg and Haughton conclude (p. 94) that, "between 1993 and 1998 . . . the

Table 9.5 Dynamics of enterprise income in Vietnam: what happened to the 1993 enterprises by 1998?

Quintile of 1998 enterprise net income	Quintile of 1993 enterprise net income				
	Low	Low-mid	Middle	Mid-upper	Upper
Low	5.90	6.80	6.80	3.04	1.25
Low-mid	5.19	7.51	8.59	7.33	3.40
Middle	4.65	6.26	10.20	10.02	6.98
Mid-upper	2.50	3.76	7.69	12.88	10.91
Upper	0.89	1.97	4.65	8.41	25.58
Enterprise terminated	69.23	64.22	52.59	45.26	39.36
Household attrited	8.05	8.41	8.23	11.09	11.99
Household dropped	3.58	1.07	1.25	1.97	0.54
Total (%)	100.00	100.00	100.00	100.00	100.00
No. of observations	559	559	559	559	559

Source: Vijverberg and Haughton (2004), Table 12

proportion of adults working in NFHEs fell, as did the proportion of households with such an enterprise. The growth in NFHE sales, expenditures, and income lagged behind GDP growth . . . based on the experience of recent history, non-farm household enterprises [can be expected to] play only a modest supporting role in fostering rapid economic growth in Vietnam."

9.6 Statistical Analysis of Panel Data

Suppose we are interested in measuring the determinants of rice production. We have some measure of output for the ith farmer in time t, given by y_{it}, and observations on a set of regressors $\mathbf{x}_{it}$ that might include the area sown, fertilizer and water use, and the like. Based on data for year $t = 1$ we might estimate

$$y_{i1} = \alpha + \mathbf{x}'_{i1}\beta + \varepsilon_{i1}. \tag{9.1}$$

Among the problems that we are likely to face here is omitted variable bias, which arises because we are unable to observe farmers' abilities (A_i), which prevents us from estimating a truer model like

$$y_{i1} = \alpha + \mathbf{x}'_{i1}\beta + \gamma A_i + \varepsilon_{i1}. \tag{9.2}$$

Now suppose that we have a panel with data for 2 (or more) years. Then we may difference (9.2) and estimate

$$\Delta y_i = \Delta \mathbf{x}'_i\beta + \Delta\varepsilon_i. \tag{9.3}$$

We have now swept away the effects of ability, and presumably removed much of the omitted variable bias from our estimate of β.

Quite generally we have

$$y_{it} = \alpha_i + \mathbf{x}'_{it}\beta + \varepsilon_{it}, \quad t = 1, 2, \ldots \tag{9.4}$$

We could simply pool the data and apply ordinary least squares (OLS), as in

$$y_{it} = \alpha + \mathbf{x}'_{it}\beta + (\varepsilon_{it} + \alpha_i - \alpha). \tag{9.5}$$

Note, however, that the errors are now likely to be correlated over time, because they include a component (the α_i) that recurs for a given farmer year after year.

A *fixed effects* model would allow for a different intercept (i.e., α_i) for each individual. We may think of the error term as being $\varepsilon_{i1} + \alpha_i$, consisting of a time-invariant component (α_i) and an idiosyncratic part (ε_{it}). For unbiasedness, we require that the $\boldsymbol{x}_{it}$ be uncorrelated with ε_{it}, but one can allow $\mathbf{x}_{it}$ to be correlated with α_i. Cameron and Trivedi (2009, p. 231) view this as a limited form of endogeneity.

The fixed effects model produces a "within" estimator. The estimate of the coefficients β is identified on the basis of variation in the $\mathbf{x}_i$ variables over time; if these variables do not change over time, $\hat{\beta}$ cannot be computed. This is problematic if the variable of interest is time invariant, or nearly so. For instance, suppose we want to estimate the effect of education on earnings, and we have a panel data of adults. In most cases, the level of education of an adult does not change over time, in which case the influence of education will be subsumed into the individual fixed effects, and cannot be identified separately.

In a *random effects* model we assume that the α_i from (9.4) are independently and identically distributed (iid), in addition to the usual assumption that the idiosyncratic error ε_{it} is iid. Given

$$y_{it} = \mathbf{x'}_{it}\beta + (\varepsilon_{i1} + \alpha_i), \tag{9.6}$$

this means we are assuming $\alpha_i \sim (\alpha, \sigma^2_\alpha)$ and $\varepsilon_{it} \sim (0, \sigma^2_u)$. This is a stronger assumption than with fixed effects, where we allowed α_i to be correlated with the $\mathbf{x}_{it}$ regressors. Put another way, for consistency:

- Fixed effects requires the assumption $E(\varepsilon_{it}|\alpha_i, \mathbf{x}_{it}) = 0$
- Random effects requires the assumption $E(\varepsilon_{it}|\mathbf{x}_{it}) = 0$.

In choosing between the two models, the key question is whether the individual effects (α_i) are correlated with the regressors ($\mathbf{x}_{it}$); if they are, we cannot use random effects.

The standard way to determine whether we should use a fixed effects (FE) or random effects (RE) estimator is using a Hausman test. If the FE estimates $\hat{\beta}_{\mathrm{FE}}$ are close to the RE estimates $\hat{\beta}_{\mathrm{RE}}$, then the RE specification is acceptable. Unfortunately, the test requires the RE estimator to be efficient, which will not be the case if, for instance, the survey data are clustered. One can handle this problem using bootstrapped standard errors; Cameron and Trivedi (2009, p. 262) provide Stata code, and we discuss bootstrapping more completely in Chap. 11.

It is much more difficult to apply panel data to nonlinear models. In the fixed effects case we cannot usually avoid estimating every α_i – whereas in the linear regression case these could be differenced out – and this is problematic in short panels, which are the norm when working with household survey data. The choice in such cases is typically between a RE and a pooled model. This is a relatively specialized field, nicely handled in Chap. 18 of Cameron and Trivedi (2009), and the references therein.

9.7 Illustration: Thai Microcredit

In order to illustrate the power of panel data, we return to the case of the Thailand Village Fund, which is discussed in the context of impact evaluation in Chap. 12.

Starting in 2001, the Thaksin government in Thailand began to provide a million baht (almost US$25,000) to every village in the country, to be used as working capital for locally run rotating credit associations. By 2004 the funds were operating in most villages, extending short-term (up to 1 year) loans that averaged US$402 each. In the same year, the Thailand Village Fund was easily the largest microcredit scheme in the world, with a gross loan portfolio of $3.00 billion, and 7.5 million active borrowers, 48% of whom were women.

Naturally, we would like to know what effect the introduction of the Village Fund had on measures such as household income or expenditure.

Because the Village Fund was put in place so quickly, there are almost no villages that could serve as controls (because they lack a local fund), and even if there were, such villages are likely to be atypical. Thus, if one is to measure the impact of Village Fund loans, there is no choice but to make comparisons between borrowers and nonborrowers within Village Fund villages, or between periods of time when a given person does, and does not, borrow.

Our interest is in the effect of T_{it}, a binary "treatment" variable that is equal to 1 if household i borrows from the Village Fund in time t, and to 0 otherwise. We would like to know whether, and to what extent, Village Fund borrowing affected household per capita spending (y_{it}).

One could estimate a cross-sectional regression, or use propensity score matching – as done in Chap. 12 – but this does not fully resolve the problem of bias due to unobserved heterogeneity: those who borrow may be more dynamic than, or somehow different from, their peers, possibly in ways that we cannot effectively measure. Hence the attraction of panel data, which would allow us to remove at least some of these effects.

It turns out that the 2004 round of the Thailand Socio-Economic Survey included 5,755 rural households that were also surveyed in 2002. This is the panel that we use in all the estimations described below. But before performing any regressions, it is helpful to generate some summary statistics that show the extent to which variables

Table 9.6 "Within" and "between" variances for Thailand data, 2002 and 2004

Variable		Mean	Std. Dev.	Min	Max	Observations
vfborrow	overall	.4130392	.4924041	0	1	N = 10108
	between		.4190994	0	1	n = 5054
	within		.2585244	−.0869608	.9130392	T = 2
malehead	overall	.7189355	.4495411	0	1	N = 10108
	between		.4262321	0	1	n = 5054
	within		.1429384	.2189355	1.218935	T = 2
nadultm	overall	1.147111	.7150896	0	6	N = 10108
	between		.648015	0	5.5	n = 5054
	within		.3024422	−1.352889	3.647111	T = 2

Notes: Output generated by the `xtsum` command in Stata. Data are from the panel component of the 2002 and 2004 Thailand Socio-Economic Surveys. Each observation represents one adult. "vfborrow" is equal to 1 if the person borrows from the Village Fund, and to 0 otherwise. "malehead" is equal to 1 for a male head of household, 0 otherwise. "nadultm" is the number of adult males in the household

show variation over time ("within" variation) and across observations at a point in time ("between" variation).

The overall variance of observations on a variable x in the dataset is given by

$$s^2_{\text{overall}} = \frac{1}{NT-1}\sum_{i=1}^{N}\sum_{t=1}^{T}(x_{it}-\bar{x})^2, \tag{9.7}$$

where N is the number of observations in a given year, T is the number of years (here just 2002 and 2004), $\bar{x}$ is the grand mean of all the observations, and x_{it} is the value of variable x for individual i at time t. The between variance measures how much an individual's value, averaged over the time periods – i.e., $\bar{x}_i = (1/T)\sum_{t=1}^{T} x_{it}$– varies from the overall mean, so

$$s^2_{\text{between}} = \frac{1}{N-1}\sum_{i=1}^{N}(\bar{x}_i-\bar{x})^2. \tag{9.8}$$

The within variance measures how much the values for an individual vary from the temporal mean – in other words, how much the values for a single individual vary over time. We note that

$$s^2_{\text{overall}} \simeq s^2_{\text{within}} + s^2_{\text{between}}. \tag{9.9}$$

The `xtsum` command in Stata computes these variances, and a few lines of output are shown in Table 9.6. Note how the mean value of vfborrow (i.e., T_{it}) is 0.41, so 41% of those sampled borrowed from the Village Fund. The within variance is 0.259, which implies that a considerable number of people borrowed in one, but not both, of 2002 and 2004. This is important, because without such variation we

Table 9.7 Panel estimates of the effect on (the log of) spending per capita of Village Fund borrowing ($\hat{\gamma}$) and earners per household ($\hat{\beta}$)

	Village Fund borrowing			Earners per household		
	$\hat{\gamma}$	SE	*t*-stat	$\hat{\beta}$	SE	*t*-stat
Pooled	0.016	0.014	1.17	0.119	0.063	1.89
Random effects	0.017	0.012	1.36	0.134	0.050	2.66
Fixed effects	0.035	0.015	2.26	0.110	0.035	3.15
IV (differenced)	0.142	0.064	2.20	0.106	0.034	3.10

Source: Boonperm et al. (2011), based on Thailand Socio-Economic Surveys of 2002 and 2004. Based on 5,054 usable observations from the (rural) panel component of these surveys. "SE" is standard error; "*t*-stat" refers to the *t*-statistic. Other variables included in the equations, but not shown here, include age, education, gender of head of household; number of adults, women, in household; number of men, women, working in agriculture, industry, trading, services; one-adult, two-parent, one-parent households; number of earners in the household; whether head is self-employed. Instruments for the IV ("instrumental variables") model include inverse of number of households in the village interacted with education of head; number of adult men, women, in household; age of head; and number of men, women, working the agricultural sector

would not be able to identify the effects of Village Fund Borrowing on income using the fixed effects ("within") estimator.

The model we wish to estimate looks like this:

$$y_{it} = \alpha_i + \mathbf{x}'_{it}\beta + \gamma T_{it} + \varepsilon_{it}. \tag{9.10}$$

We are particularly interested in the sign and magnitude of γ, which would give us the average treatment effect of Village Fund borrowing.

In Table 9.7 we report a selection of the results of estimating a variety of specifications of (9.10); similar results appear in Boonperm et al. (2011). In each case we adjust for the clustering of observations at the village level, which increases the size of the standard errors relative to the unadjusted case.

The first row in Table 9.7 shows the coefficient estimates and standard errors for the pooled regression, where $\alpha_i = \alpha$. We have $\hat{\gamma} = 0.016$ (and $t = 1.17$); the estimate is not statistically significant, but if it were, it would tell us that Village Fund borrowing is associated with about 1.6% more spending, other things being equal. The random effects results, shown in the second row, tell a similar story, with $\hat{\gamma} = 0.017$ and $t = 1.36$.

We get somewhat different results when we employ fixed effects or, equivalently, estimate a differenced equation [as in (9.3)]. Now $\hat{\gamma} = 0.035$ and $t = 2.26$ (with a p-value of 0.024). It appears that Village Fund borrowing is associated with a 3.5% increase in spending. This is not implausible; for Village Fund borrowers the loans represented, on average, 11.6% of income in 2004.

At this point, it is worth emphasizing again that our estimate of the effect of Village Fund borrowing on household spending is based on *changes in borrowing* by individuals between 2002 and 2004. Some households borrowed in 2002 but not in 2004, others in 2004 but not 2002. In effect, our estimates hinge on the behavior of these individuals, and not on a comparison between borrowing adults and different nonborrowing adults.

Which results are more plausible, fixed or random effects? Using a Hausman test, we reject the null hypothesis of no differences between the estimated random effects and fixed effects coefficients; it follows that the fixed effects specification is more appropriate.[1]

The use of fixed effects goes a long way toward addressing concerns about the endogeneity of borrowing. However, it is worth asking whether it would be better to use instrumental variables estimation. For this to be credible, we would need to find at least one instrument that plausibly affects whether one borrows, but is not correlated with the errors ε_{it}. One possibility is the size of the village; each village got an initial million baht in funds, irrespective of size. So households in larger villages are likely to have a smaller probability of getting a loan. On the other hand, we have no reason to believe that the size of the village would influence how a loan affects household spending.

We only have a measure of village size for 2004, so this is a time-invariant variable, and cannot serve directly as an instrument for T_{it}. However, we can interact village size (or its inverse) with time-varying variables, and use these hybrids in the first-stage regression.

The results are shown in the final row of Table 9.7. They show a much larger effect of Village Fund borrowing on household spending per capita, at a statistically significant 14% (with a p-value of 0.028). But is this a more plausible estimator than the basic fixed effects specification? We note first that our chosen instruments are relevant, in that they have a statistically significant influence on whether someone borrows from the Village Fund ($F(8{,}674) = 33.5$; p-value $= 0.00$). The GMM C-statistic tests whether the treatment variable (T_{ij}) should be considered endogenous, and provides only weak support for the endogeneity of Village Fund borrowing, given the instruments available.

In short, one can make a plausible case that the best specification here is fixed effects. We thus conclude that, until more information suggests otherwise, our best estimate is that Village Fund borrowing raised spending by about 3.5% in 2004.

References

Boonperm, Jirawan, Jonathan Haughton, and Shahidur R. Khandker. 2011. Does the Village Fund matter in Thailand? Suffolk University, Boston.

Breen, Richard, and Pasi Moisio. 2004. Poverty dynamics corrected for measurement error. *Journal of Economic Inequality* 2: 171–191.

[1] We used a robust version of the Hausman test, which allows for clustering in the design. Cameron and Trivedi (2009, p. 262) provide details.

Cameron, Colin, and Pravin Trivedi. 2009. *Microeconometrics using Stata*. College Station: Stata Press.

Haughton, Dominique, Jonathan Haughton, Le Thi Thanh Loan, and Nguyen Phong. 2001. Shooting stars and sinking stones. In *Living standards during an economic boom: The case of Vietnam*, ed. Dominique Haughton, Jonathan Haughton, and Nguyen Phong. Hanoi: Statistical Publishing House.

Haughton, Jonathan, and Shahidur Khandker. 2009. *Handbook on poverty and inequality*. Washington, DC: World Bank.

Jalan, Jyotsna, and Martin Ravallion. 1998. Transient poverty in Postreform rural China. *Journal of Comparative Economics* 26: 338–357.

Lee, Nayoung, Geert Rider, and John Strauss. 2010. *Estimation of poverty transition matrices with noisy data*. Shatin: Chinese University of Hong Kong.

Perkins, Dwight. 1994. Industrialization. *In Search of the dragons' trail: Economic reform in Vietnam [Viet Nam cai cach kinh te cheo huong rong bay]*, ed. Dapice David, Haughton Jonathan, and Perkins Dwight. Hanoi: Political Publishing House [Nha Xuat Ban Chinh Tri Quoc Gia]. 89–101.

Vijverberg, Wim, and Jonathan Haughton. 2004. Household enterprises in Vietnam: Survival, growth, and living standards. In *Economic growth, poverty, and household welfare in Vietnam*, ed. Paul Glewwe, Nisha Agrawal, and David Dollar. Washington, DC: World Bank.