Assessing vulnerability to poverty: concepts, empirical methods and illustrative examples

Shubham Chaudhuri
Department of Economics
Columbia University

June 2003

Abstract

A household's observed poverty level is an *ex-post* measure of a household's well-being (or lack thereof). But poverty is a stochastic phenomenon and the current poverty level of a household, may not necessarily be a good guide to the household's expected poverty in the future. For thinking about appropriate forward-looking anti-poverty interventions (i.e., interventions that aim to go beyond the alleviation of current poverty to prevent or reduce future poverty), the critical need then is to go beyond a cataloging of who is currently poor and who is not, to an assessment of households' vulnerability to poverty. In this paper, we make the case for broadening the scope of poverty assessments to take account of vulnerability to poverty and outline a conceptual and empirical approach for doing so. The paper has two broad aims: first, to provide a conceptual and methodological overview of the uses and empirical implementation of vulnerability assessments using household-level data; and second, to demonstrate, through a number of illustrative examples as well as two more detailed country studies, how the general methodological approach can be usefully applied and tailored to particular contexts and data, to yield policy-relevant insights about the nature and extent of vulnerability.

Table of Contents

1. Introduction

Poverty reduction has long been recognized as the implicit objective of development policy. For more than a decade now, national poverty assessments have been used on a routine basis to inform policy discussions on poverty alleviation in numerous developing economies. These poverty assessments have drawn on cross-sectional household surveys to provide a detailed profile of the poor, and to document the incidence of poverty in various segments of the population.

But poverty is a stochastic phenomenon. Today's poor may or may not be tomorrow's poor. Currently non-poor households who face a high probability of a large adverse shock, may, on experiencing the shock, become poor tomorrow. And among the currently poor households there may be some who are only transitorily poor as well as other who will continue to be poor (or poorer) in the future. In other words, a household's (or an individual's) observed poverty level or status—defined in most cases simply in terms of the household's observed level of consumption expenditure relative to a pre-selected poverty line—is an ex-post measure of a household's well-being (or lack thereof). But for policy purposes, what really matters is the ex-ante risk that a household will, if currently non-poor, fall below the poverty line, or if currently poor, will remain in poverty. And the current poverty level of a household, may not necessarily be a good guide to the household's vulnerability to poverty in the future. For thinking about appropriate forward-looking anti-poverty interventions (i.e., interventions that aim to go beyond the alleviation of current poverty to prevent or reduce future poverty), the critical need then is to go beyond a cataloging of who is currently poor and who is not, to an assessment of households' vulnerability to poverty.

In this paper, we make the case for broadening the scope of poverty assessments to take account of vulnerability to poverty and outline a conceptual and empirical approach for doing so. The paper has two broad aims:

- first, to provide a conceptual and methodological overview of the uses and empirical implementation of vulnerability assessments using household-level data

- and second, to demonstrate, through a number of illustrative examples as well as two more detailed country studies, how the general methodological approach can be usefully applied and tailored to particular contexts and data, to yield policy-relevant insights about the nature and extent of vulnerability.

The paper is organized as follows. The next section provides a conceptual overview. We begin by clarifying the links between the concepts of poverty, risk and vulnerability. We then detail the rationale for vulnerability assessments. An assessment of vulnerability, we argue, is necessary and desirable not only because vulnerability is an inherently important dimension of well-being, but also because such an assessment serves other important instrumental functions: it informs the design of forward-looking poverty reduction strategies, it highlights the distinction between poverty prevention and poverty alleviation interventions, and it clarifies the role of risk in the dynamics and persistence of poverty.

We then outline a simple taxonomy for thinking about the multiple interlinked factors that make households vulnerable to poverty and use this taxonomy to suggest ways in which a vulnerability assessment might be organized. The section ends with an operational definition of vulnerability to poverty that can, in principle, be taken to the data.

The third section outlines a general and fairly flexible methodology for empirical implementing vulnerability assessments using household data. After summarizing the basic approach, we turn to a detailed discussion of the various steps involved and the econometric issues that arise at each step.

The fourth section contains a series of illustrative examples demonstrating the ways in which vulnerability estimates might be used to inform policy. The examples are drawn from three different studies: a study of vulnerability in rural south and southwestern China using longitudinal household-level data for the six years from 1985 to 1990 (Chaudhuri and Jalan (2003)); a vulnerability assessment for Indonesia using data from the mini-SUSENAS collected in December 1998 and August 1999 (Chaudhuri, Jalan and Suryahadi (2003)); and a study using household level data from the Philippines for 1997 and 1998 (Chaudhuri and Datt (2002)). Further details about the setting, the data and the econometric strategy pursued in each of these studies are summarized in the appendix.

2. Vulnerability to poverty: a conceptual overview

2.1. Poverty, risk and vulnerability

Poverty is an *ex-post* measure of a household's well-being (or lack thereof). It reflects a current state of deprivation, of lacking the resources or capabilities to satisfy current needs. *Vulnerability*, on the other hand, may be broadly construed as an *ex-ante* measure of well-being, reflecting not so much how well off a household currently is, but what its future prospects are. What distinguishes the two is the presence of *risk*—the fact that the level of future well-being is uncertain. The uncertainty that households face about the future stems from multiple sources of risk—harvests may fail, food prices may rise, the main income earner of the household may become ill, etc. If such risks were absent (and the future were certain) there would be no distinction between ex-ante (vulnerability) and ex-post (poverty) measures of well-being.

2.2. Why should we worry about vulnerability?

The case for considering the role of risk in the design and implementation of social policy has been made eloquently elsewhere (see, for instance, Holzmann (2001), Holzmann and Jørgensen (2000), and Heitzmann, Canagarajah and Siegel (2002)). The claim that the nature and magnitude of the risks that households face, and the scope of the risk-management mechanisms they have access to, given the environments in which they operate, potentially play a central role in the dynamics and scale of poverty is also supported by both theoretical analyses and empirical evidence. Drawing on these arguments, we least four reasons why broadening the scope of poverty assessments to include an analysis of vulnerability to poverty is both desirable and necessary.

First, and most obviously, for thinking about appropriate forward-looking anti-poverty interventions, it clearly is necessary to go beyond a cataloging of who is currently poor, how poor they are, and why they are poor to an assessment of households' vulnerability to poverty—who is likely to be poor, how likely are they to be poor, how poor are they likely to be, and why are they likely to be poor. An atemporal or static approach to well-being, if strictly adhered to, is of limited use in thinking about policy interventions to improve well-being that can only occur in the future. Of course, in practice, poverty assessments, even if couched in atemporal terms, are used in the process of policy formulation. But in doing so, implicit assumptions are being made about the extent to which the situation recorded in the poverty assessment will be reproduced over time. A reconceptualization in terms of vulnerability to poverty, which, by definition has to be forward-looking, forces us to make these assumptions explicit and consider the potential role and effects of risk.

Second, a focus on vulnerability to poverty serves to highlight the distinction between ex-ante poverty prevention interventions and ex-post poverty alleviation interventions. A simple public health analogy makes this distinction clear. Just as efforts to combat a disease outbreak include both treatment of those already afflicted as well as preventive measures directed at those at risk, poverty reduction strategies need to incorporate both alleviation and prevention efforts.

Third, addressing vulnerability also has instrumental value. Because of the many risks households face, they often experience shocks leading to a wide variability in their income. In the absence of sufficient assets or insurance to smooth consumption, such shocks may lead to irreversible losses, such as distress sale of productive assets, reduced nutrient intake, or interruption of education that permanently reduces human capital (Jacoby and Skoufias, 1997), locking their victims in perpetual poverty. Aware of the potential of such irreversible outcomes, vulnerable people often engage in risk mitigating strategies to reduce the probability of such events occurring. Yet, these strategies yield typically low average returns. Thus, when people lack the means to smooth consumption in the face of variable incomes, they are often trapped in poverty through their attempts to steer clear of irreversible shocks (Morduch, 1994; Barrett, 1999). In a similar vein it is being observed at the macro-level that economic growth slows down in the face of downward risks resulting from structural phenomena such as climatic vagaries, fluctuations in the terms of trade and political insecurity (Guillaumont, Guillaumont, Brun, 1999). Policies directed at reducing vulnerability—both at the micro and macro level—will be instrumental in reducing poverty.

Last but not least, vulnerability is an intrinsic aspect of well-being. That exposure to risk and uncertainty about the future adversely effect current well-being is one of the central tenets of the basic economic theory of human behavior, embodied in the assumption that individuals and households are risk averse. And as the World Development Report 2000/2001 on *Attacking Poverty* documents, this presumption is echoed by findings from worldwide consultations that indicate that risk and uncertainty are a central preoccupation of the poor.

2.3. What makes a household vulnerable to poverty? A taxonomy

A household's vulnerability to poverty at any point in time depends on how its livelihood prospects and well-being is likely to evolve over time. And that in turn depends on its future income prospects, the degree of income volatility it faces, its ability to smooth consumption in the face of income or other livelihood shocks. These in turn depend on the complex dynamic interlinkages between the environment—macroeconomic, institutional, sociopolitical and physical—in which the household operates, the resources, human, physical and financial it commands, and its behavioral responses. Such a dynamic perspective on household well-being suggests that the proximate causes of poverty and vulnerability to poverty are:

- exposure to adverse aggregate shocks (e.g. macroeconomic shocks or commodity price shocks) and/or adverse idiosyncratic shocks (e.g., localized crop damage or illness of the main income-earner in the household). A household may have a high level of exposure for one of both of two reasons:

  - it faces high levels of underlying risk
  - it has a limited ability to maintain its well-being in the face of adverse livelihood shocks, i.e., limited ability to cope with risks

- low long-term income generating capacity

Those who are vulnerable to transitory poverty suffer primarily from exposure to adverse shocks. On the other hand the structurally or chronically poor are those who are both exposed to adverse shocks and have limited long-term income generating capacity. Poverty reduction efforts must protect the former and assist the latter.

Both the long-term poor and those who find themselves in poverty because of an adverse shock, may adopt a variety of coping strategies to meet basic essential needs. Some of these coping strategies, while they might enable the household to meet critical short-term needs, can be costly in terms of the future well-being of the household, and in particular may condemn the children of the household to a lifetime of poverty as well. Measures to prevent the transmission of poverty from one generation to the next must be an essential component of any sustainable poverty reduction strategy.

Poverty prevention efforts that aim to reduce vulnerability to poverty and prevent the transmission of poverty must go beyond the proximate causes of poverty and vulnerability to address the multiple underlying causes of poverty. Any categorization of the underlying causes of poverty is ultimately somewhat arbitrary given the numerous complicated ways in which the various factors that lead to poverty are intertwined. A household is more likely to be exposed to adverse shocks and have limited earnings prospects and income-generating capacity if it:

- has low levels of human capital, know-how and access to information
- suffers from physical and psychological disabilities
- has few productive and financial assets
- suffers from social exclusion or inadequate networks of social support
- has limited access to credit and risk-management instruments

- lives in a setting with adverse agroclimatic conditions and limited natural resources
- lives in a community where there is insufficient entrepreneurial activity and job creation
- works in a sector that is particularly sensitive to macroeconomic volatility and sectoral shocks

The multiple interlocking paths to poverty are illustrated schematically in Figure 1. The conceptual approach outlined above is both comprehensive and general. To be fruitfully applied in a particular context, the relative importance (the weights) of the various paths to poverty has to be ascertained through careful empirical analyses. And that is what vulnerability assessments are about.

2.4. Organizing vulnerability assessments

Clearly, given the complexity of the multiple dynamic interlinkages illustrated in Figure 1, vulnerability assessments should take place at multiple levels, and should be directed at multiple issues. No single empirical methodology or approach can simultaneously encompass all these issues. Rather, what is needed is research and analysis on a number of fronts using different data sources and possibly, somewhat different empirical methods.

Nevertheless, the taxonomy sketched in Figure 1 provides an useful basis for organizing vulnerability assessments. Mirroring the nested hierarchical structure of the taxonomy, vulnerability assessments can be structured into a series of hierarchically related questions, with the questions at each step being progressively more and more narrowly focused.

To begin with, vulnerability assessments need to be able to say something about the extent of vulnerability in the population. How widespread is vulnerability to poverty—how many face a non-negligible risk of poverty, how likely are they to be poor and how poor are they likely to be? Are vulnerability and poverty distinct phenomena? Is the scale of poverty reduction interventions appropriate?

At the next step, we need to know more about who the vulnerable are and how concentrated vulnerability is within different segments of the population. If interventions seem necessary, where should they be directed?

Following naturally from this comes the next question, what types of interventions are necessary? Vulnerability assessments can address this issue at multiple levels. At the highest level, they can shed light on the proximate causes of vulnerability: are households vulnerable to poverty primarily because their consumptions are volatile, which would imply they are mostly vulnerable to transitory poverty, or are they structurally poor? How do the proximate causes of vulnerability vary across various segments of the population?

Consumptions are volatile because households are exposed to risk. Even for structurally poor households, consumption volatility may contribute significantly to vulnerability. So at the next step, we need to better understand why consumptions are volatile? Is it because households face high levels of underlying risk? Or is it that they have a limited ability to cope with even moderate levels of risk?

Lastly, these latter two questions immediately suggest the need to identify the sources of risk that households are most exposed to, as well as the need for a better understanding of the risk management instruments, public or private, to which households have access.

2.5. Defining vulnerability to poverty

Poverty and vulnerability (to poverty) are two sides of the same coin. The observed poverty level or status of a household (defined simply in terms of a household's observed level of consumption expenditure relative to a pre-selected poverty line) is the ex-post realization of a random variable, the ex-ante *expectation* of which can be taken to be the household's level of vulnerability.

To operationalize this conceptualization of vulnerability to poverty in terms of *expected* poverty, we begin with the following general formulation of the poverty index for a household h at time t:

$$p_{ht} = \frac{u(z) - u(c_{ht})}{|u(z)|} \tag{2.1}$$

Here z is a pre-specified poverty line, c_{ht} is the consumption level of household h at time t and $u(.)$ is an increasing function. If we take as the form for $u(.)$ the following functional form:

$$u(c) = z^{\alpha} - (\max\{0, z - c\})^{\alpha} \tag{2.2}$$

with α taking on integer values, 0, 1, 2, etc., the poverty index, (2.1), reduces to the familiar Foster-Greere-Thorbecke (1984) family of poverty measures:

$$p_{\alpha,ht} = \left(\max\left\{0,\ \frac{z - c_{ht}}{z}\right\}\right)^{\alpha} \tag{2.3}$$

When $\alpha = 0$, the poverty index is simply a binary indicator of whether a household is poor. When $\alpha = 1$, the index becomes the poverty gap ratio, and with $\alpha = 2$, the squared poverty gap.

There is no reason, however, to necessarily limit ourselves to the specific form of $u(.)$, (2.2), implied by the FGT poverty measures. We could as well consider alternative forms more familiar from expected utility theory such as:

$$u(c) = \frac{c^{1-\rho}}{1-\rho}$$

and this is in fact what Ligon and Schecter (2002) suggest.

However, because the Foster-Greere-Thorbecke measures are familiar from their widespread use in poverty assessments, and because they are more easily interpreted and exposited than the utility-based measures, we adopt the FGT measures as the base for constructing our measures of vulnerability. The corresponding definitions of vulnerability are therefore:

$$\begin{aligned} v_{\alpha,ht} &= E[p_{\alpha,h,t+1}(c_{h,t+1}) \mid F(c_{h,t+1})] \\ &= \int (\max\{0, \frac{z - c_{h,t+1}}{z}\})^{\alpha} dF(c_{h,t+1}) \\ &= F(z) \int_{\underline{c}}^{z} \left(\frac{z - c_{h,t+1}}{z}\right)^{\alpha} \frac{f(c_{h,t+1})}{F(z)} dc_{h,t+1} \end{aligned} \tag{2.4}$$

where $F(c_{h,t+1})$ and $f(c_{h,t+1})$ respectively denote the cumulative distribution and density functions of $c_{h,t+1}$.

A number of other measures or definitions of vulnerability have been proposed in the emerging literature on vulnerability assessment. A few comments about the relative advantages and drawbacks of the measures we propose are therefore in order. Vulnerability has sometimes been defined in terms of a household's ability to smooth consumption in the face of income shocks with those households whose consumptions are more sensitive to income shocks being considered more vulnerable. There are two problems with this definition, both of which are absent in the poverty-based measures we propose.

First, defining vulnerability solely in terms of a household's consumption smoothing ability ignores the variation across households in levels of exposure to income shocks. A household may well have a lower ability to smooth consumption but it may also be subject to fewer income shocks. Defining vulnerability, as we do, in terms of expected poverty takes both these components of vulnerability into account. Second, measures that focus on the ability to smooth consumption ignore the asymmetry that, we would argue, is crucial to the notion of vulnerability, namely the importance of exposure to downside risk. The advantage of the expected poverty formulation of the vulnerability measure is the built-in asymmetry of poverty measures that implicitly gives more weight to downside risks.

Vulnerability has also been defined in terms of exposure to adverse *shocks* to welfare, rather than in terms of exposure to poverty. This definition also differs substantively from ours in that our definition would include among the vulnerable, households who are currently poor and have a high probability of remaining poor even if they do not experience any large adverse welfare shocks. On the other hand, our definition would exclude those households among the non-poor who face a high probability of a large adverse shock but are currently well-off enough so that even were they to experience the shock, they would still remain non-poor.

With the central role that the notion of poverty, and more generally, the level of household welfare, plays in policy discussions, a measure of vulnerability that takes account of welfare levels—in particular, poverty levels—seems preferable. However, it should be recognized that to the extent that defining welfare relative to a prespecified poverty line is considered somewhat arbitrary, that arbitrariness carries over to our vulnerability measures. A potentially more serious concern is the fact that, as Ligon and Schecter (2002) point out, one of the vulnerability measures we propose, the expected poverty status or future likelihood of poverty of a household, has the perverse implication that increases in risk would reduce the vulnerability level of those with mean consumption levels below the poverty line. But, as we demonstrate below, that does not negate its usefulness in informing certain policy decisions. It is only in quantifying the contribution of risk to vulnerability that we need to be careful not to rely on this measure, and instead use one of the other measures, which do not suffer from this shortcoming.

3. Empirical methods for assessing household vulnerability to poverty

This section provides an overview of the measurement and econometric issues that arise in carrying out vulnerability assessments, and outlines a general and fairly flexible methodology for organizing and carrying out such assessments.

3.1. The basic approach

Whatever the precise measure of vulnerability one chooses to work with, the starting point has to be an explicit specification of the underlying data-generating process for consumption. Vulnerability assessments, by definition, have to be explicitly forward-looking. No matter how rich the data, the vulnerability of households is never directly observable. In contrast, most poverty assessments are couched in atemporal terms and, given the right data, it is possible to actually observe the current poverty level or status of the household.

From this it naturally follows that the observed consumption expenditures at a point in time (i.e., from a single cross-section survey) should be viewed as the outcome (snapshot) of a dynamic process that is occurring in real time. And this means that vulnerability assessments (again, in contrast to poverty assessments which remain largely atheoretical) have to be rooted in explicit models of intertemporal household behavior.

How general and flexible the specification of the consumption process can be depends first and foremost on the data that are available. Given the limitations of most data sets, a priori restrictions on the consumption process will almost certainly need to be made. And in this, it is important to be clear about the assumptions implicit in whatever specification is ultimately adopted.

Once a specification has been chosen, the next step is to estimate the parameters of the process using the household data. In general it will be possible to estimate the key parameters in a fairly flexible way without making too many stringent distributional assumptions. However, in going from estimates of the consumption process to estimates of vulnerability, the problem of estimating the distribution of consumption will need to be faced. Here they are two possible approaches. The first is to work with a pre-specified parametric distribution. The second is to use non-parametric techniques to get at the distribution of future consumption.

Lastly, the estimates of the consumption process and the estimates of vulnerability can be used in number of different ways to inform the design of poverty reduction policies. To summarize, then, the basic approach consists of four steps:

- Step 1: specify the data generating process for consumption
- Step 2: use survey data on household consumption expenditures and characteristics to estimate the relevant parameters of the consumption process
- Step 3: make the necessary distributional assumptions needed to draw inferences about future consumption prospects—i.e., to go from estimates of the consumption process to estimates of vulnerability

- STEP 4: use the vulnerability estimates, decompositions of the vulnerability estimates, and a variety of counterfactuals constructed using the estimates of the consumption process, to address various policy-relevant questions

3.2. Specifying the consumption process

The level of vulnerability at time t is defined in terms of the household's consumption prospects at time $t+1$. The difference is noteworthy because it reflects an important distinction between the notion of vulnerability and the concept of poverty. Vulnerability is a forward-looking or ex-ante measure of a household's well-being, whereas poverty is an ex-post measure of a household's well-being (or lack thereof). This implies that while the poverty status of a household is concurrently observable-i.e., with the right data we can make statements about whether or not a household is *currently* poor-the level of vulnerability is not. We can *estimate* or *make inferences* about whether a household is *currently* vulnerable to *future* poverty, but we can never directly observe a household's current vulnerability level.

An assessment of vulnerability is, therefore, innately a more difficult task than assessing who is poor and who is not. To assess a household's vulnerability to poverty we need to make inferences about its future consumption prospects. And in order to do that, we need a framework for thinking explicitly about both the inter-temporal aspects and cross-sectional determinants of consumption patterns at the household level.

Over the last two decades, a large literature has developed which addresses precisely these issues (See Deaton(1992) and Browning & Lusardi(1995) for excellent overviews). This literature suggests that a household's consumption in any period will, in general, depend on a number of factors. Among them its wealth, its current income, its expectations of future income (i.e., lifetime prospects), the uncertainty it faces regarding its future income and its ability to smooth consumption in the face of various income shocks. Each of these will in turn depend on a variety of household characteristics, those that are observable and possibly some that are not, as well as a number of features of the aggregate environment (macroeconomic and socio-political) in which the household finds itself. At a general conceptual level, this suggests the following reduced form expression for consumption:

$$c_{ht} = c(X_h, \beta_t, \alpha_h, e_{ht}) \tag{3.1}$$

where X_h represents a bundle of observable household characteristics, β_t is a vector of parameters describing the state of the economy at time t, and α_h and e_{ht}

represent, respectively, an unobserved time-invariant household-level effect, and any idiosyncratic factors (shocks) that contribute to differential welfare outcomes for households that are otherwise observationally equivalent.

Substituting from (3.1) into (2.4) we can rewrite the expression for the vulnerability level of a household as:

$$v_{ht} = E[p_{\alpha,h,t+1}(c_{h,t+1}) \mid F(c_{h,t+1} | \, X_h, \beta_t, \alpha_h, e_{ht})] \tag{3.2}$$

The expression above makes clear that a household's vulnerability level derives from the stochastic properties of the inter-temporal consumption stream it faces, and these in turn depend on a number of household characteristics and characteristics of the environment in which it operates. And at a conceptual level, the expression is very general in a number of respects.

First, it allows for the possibility of complicated interactions between the multiple cross-sectional determinants of a household's vulnerability level. For instance, X_h could include variables such as the educational attainment of the head of the household, presence of a government poverty scheme in the community in which the household resides, as well as interactions between the two to capture potential inequities in the level of access to public programs.

Second, because a household's vulnerability is defined in terms of its future consumption prospects *conditional* on its current characteristics, both observed and unobserved, the possibility of poverty traps and other non-linear poverty dynamics is implicitly built in.

And third, the possible contribution of aggregate shocks and unanticipated structural changes in the macro-economy to vulnerability at the household level is also incorporated through inclusion of the time-varying set of parameters, β_t.

In practice, as will be clear in the next section, data constraints will usually not permit estimation of vulnerability at the level of generality embodied in expression (3.2). Nevertheless the formulation is useful in providing a basis for thinking through the possible implications of the various restrictions that will need to be imposed in any attempt to estimate vulnerability with the sorts of data that are usually available.

3.3. Econometric issues

A number of econometric issues arise in implementing the basic approach outlined above, many of them driven by data constraints. To illustrate some of the issues that arise we begin by laying out what in some respects is an almost ideal specification of the consumption process for the purposes of estimating vulnerability:

$$\ln C_{hjt} = X_h\alpha_j + X_h P_t\beta_j + X_h R_{jt}\gamma_j + X_h M_{hjt}\delta_j + v_{jt} + \eta_h + e_{hjt}\sqrt{g(X_h, \theta_j} \quad (3.3)$$

Here X_h is a vector of observable characteristics of household h, P_t, a vector of observable macro shocks in year t, for instance, commodity price shocks, R_{jt} captures observable locally covariate shocks in area j in year t, for instance weather shocks, M_{hjt} denotes an observable idiosyncratic shock experienced by household h in area j in year t, e.g., illness of the main income earner, v_{jt} represents unobserved area-specific shocks, η_h, an unobserved time-invariant household effect, and e_{hjt} an idiosyncratic time-varying disturbance term.

Estimation of (3.3) would clearly impose significant demands on the data, demands that are unlikely to be met in most instances. If as is quite common, panel data are not available, we cannot control for unobserved household-level effects. Not only does this potentially bias the estimates of the coefficients on the variables we do observe, it also raises the possibility that unobserved heterogeneity in the cross-section will be confounded with inter-temporal variation in consumption levels. The lack of panel data also eliminates the possibility of exploring poverty dynamics. Nevertheless, as we illustrate in a later section, cross-sectional data can be useful for certain purposes.

The other main constraint that must often be faced is the lack of information on area-specific variables, R_{jt} and on the idiosyncratic shocks experienced by individual households, M_{hjt}. In the former case, it becomes that much more difficult to distinguish time-invariant area-specific characteristics from time-varying locally covariate shocks. Even with panel data, this problem cannot always be overcome if there are seasonal effects that need to be considered as well. Without information on M_{hjt} it may still be possible to consistently estimate (3.3) but the analysis will be much less informative.

3.4. Estimating household vulnerability

From (3.2) it is clear that because a household's vulnerability to poverty is a *non-linear* function of its future consumption levels, it will depend, not just on its expected (i.e., mean) consumption looking forward, but also on the volatility (i.e., variance, from an inter-temporal perspective) of its consumption stream, and possibly on higher moments of the consumption process as well. A salaried low-level government employee with an expected level of consumption roughly similar to that of a self-employed proprietor of a small business may nevertheless be much less vulnerable to poverty because of the relative stability of the former's consumption stream.[1]

To go from estimates of the consumption process to an estimate of the household's vulnerability to poverty we need to therefore not only estimate its expected consumption in the future but also to be able to say something about the distribution of its future consumption. At a minimum, even we are willing to make the parametric assumption that consumption is log-normally distributed and hence the entire distribution of consumption is captured by the mean and variance, this implies that we need to estimate the variance of its future consumption. In this section we describe a couple of different ways of doing this. But first we highlight a key element of the estimation strategy, whether or not a parametric approach is adopted, which is the need to allow for heteroskedasticity in the specification of the consumption process.

3.4.1. Allowing for heteroskedasticity

The specification of the consumption process discussed in the previous section is not a new one. Similar specifications have been estimated in a number of studies exploring household consumption behavior and the determinants of consumption. These have included a number of poverty assessments. However, in most previous studies, the disturbance term is implicitly thought of as stemming from measurement error or some unobserved factor that is incidental to the main focus of the analysis. And thus, it is usually assumed that the variance of the disturbance term is the same for all households.

There are two problems with this assumption when the specification of the consumption provides the basis for estimating vulnerability to poverty. First, within this framework the variance of the disturbance term is interpreted in *economic* terms as the *inter-temporal variance of log consumption.* Viewed from this perspective, the assumption that the variance of log consumption is the same for all households seems quite restrictive, regardless of its statistical import. That is because it forces the estimates of the mean and variance of consumption to be monotonically related across households, ruling out the possibility that a household with a lower mean consumption may nevertheless face greater consumption volatility than a household with a higher average level of consumption. Both formal and anecdotal evidence points to high levels of income and consumption volatility for poor households.

[1] Of course at times of macroeconomic crises accompanied by rapid inflation, the situations may easily be reversed.

Moreover, in purely statistical terms, unlike in other settings where failure to account for heteroskedasticity results in a loss of efficiency but need not bias the estimates of the main parameters of interest, here, the standard deviation of the disturbance term enters directly in generating an estimate of vulnerability (see (3.4) below). A biased estimate of this parameter will therefore lead to a biased estimate of vulnerability.

To address this problem we need to allow the variance of the disturbance term, e_{ht}, to depend upon the particular characteristics of the household. A simple way of doing so would be to begin by specifying a functional form such as:

$$\ln \sigma^2_{\ln c_{h,t}} = X_{ht}\gamma + Z_h\delta$$

where X_{ht} and Z_h are, respectively, vectors of observable time-varying and time-invariant household characteristics. Under the interpretation of the disturbance term, e_{ht}, in the consumption equation as the shock to consumption, the log of the squared estimated residuals from the consumption equation provides an estimate of $\ln \sigma^2_{\ln c_{h,t}}$. Estimates of γ and δ can then be obtained from the following regression:

$$\ln \widehat{e}^2_{ht} = X_{ht}\gamma + Z_h\delta + u_{ht}$$

3.4.2. Parametric estimates of vulnerability

Imagine that we have obtained estimates of the mean and variance of one-period ahead log consumption, where these are denoted $\widehat{\mu}_{\ln c_{h,t+1}}$ and $\widehat{\sigma}^2_{\ln c_{h,t+1}}$ respectively. If we are willing to assume that consumption is log-normally distributed (i.e., that $\ln c_{h,t+1}$ is normally distributed), the estimates of vulnerability can be straightforwardly generated using the properties of the normal distribution.

Specifically, letting $\Phi(.)$ denote the cumulative density of the standard normal, the estimate of $v_{0,ht}$—the vulnerability to poverty defined as the likelihood of poverty of household h at time t—will be given by:

$$\widehat{v}_{0,ht} = \widehat{\Pr}\left(\ln c_{h,t+1} < \ln z \mid \widehat{\mu}_{\ln c_{h,t+1}}, \widehat{\sigma}^2_{\ln c_{h,t+1}}\right) = \Phi\left(\frac{\ln z - \widehat{\mu}_{\ln c_{h,t+1}}}{\widehat{\sigma}_{\ln c_{h,t+1}}}\right) \qquad (3.4)$$

The expressions for vulnerability to poverty defined in terms of the expected poverty gap ratio or the expected squared poverty gap are a bit more complicated. Even with the assumption of log-normality, these cannot be evaluated analytically. However, estimates of vulnerability under these two definitions are easily obtained from Monte Carlo simulations using the estimates of $\widehat{\mu}_{\ln c_{h,t+1}}$ and $\widehat{\sigma}^2_{\ln c_{h,t+1}}$.

3.4.3. Non-parametric approaches

Kamanou and Morduch (2002) propose a non-parametric approach to estimating the distribution of future consumption. The approach uses a Monte Carlo design to simulate the future distribution of consumption, where the simulations are based on bootstrapping the empirical distribution of observable shocks and estimated residuals.

This approach is very promising. The only drawback with this particular implementation is that it implicitly assumes that the shocks to consumption experienced by different households are drawn from the same distribution. This clearly goes against what we argue for above, that households in different circumstances, facing different risks and with differing access to risk-management instruments should be presumed to experience different levels of consumption volatility, i.e., that we ought to allow for heteroskedasticity.

One way of addressing this shortcoming is to combine the parametric and non-parametric approaches by first flexibly (i.e., allowing for heteroskedasticity) estimating the variance of consumption for each household, and then using these estimated variances as propensity scores in constructing kernel-based re-sampling weights for the bootstrapping procedure under the non-parametric approach.

4. Interpreting and using vulnerability estimates to inform policy: illustrative examples

To demonstrate how household-level vulnerability estimates, generated applying the methodology outlined above, may be interpreted and used to inform poverty-reduction policies, in this section, we go through a series of illustrative examples. The examples are drawn from three different studies: a study of vulnerability in rural south and southwestern China using longitudinal household-level data for the six years from 1985 to 1990 (Chaudhuri and Jalan (2003)); a vulnerability assessment for Indonesia using data from the mini-SUSENAS collected in December 1998 and August 1999 (Chaudhuri, Jalan and Suryahadi (2003)); and a study using household level data from the Philippines for 1997 and 1998 (Chaudhuri and Datt (2002)). Further details about the setting, the data and the econometric strategy pursued in each of these studies are summarized in the appendix.

4.1. Documenting aggregate vulnerability and poverty

The natural first step in a vulnerability assessment is to obtain a sense of the overall level of vulnerability in the population of interest by plotting the aggregate distribution of vulnerability. While this can, in principle, be done for each of the specific measures of vulnerability to poverty described earlier—those based on the FGT poverty measures as well as those derived from utility theory—in practice, such a plot is probably easiest to interpret and grasp when vulnerability is defined as the probability of future poverty, i.e., in terms of a household's expected poverty status.

Figure 2 plots the distribution of vulnerability defined in this way for rural south and southwest China in 1985. If we imagine labeling a household as vulnerable if its estimated vulnerability level exceeded some threshold, what Figure 2 depicts is the estimated incidence of vulnerability for vulnerability thresholds ranging from 0 to 1—measured along the horizontal axis—for the population as a whole as well for sub-samples sorted by observed poverty status in 1985. By construction, as the threshold increases, the incidence of vulnerability (the fraction of the population that has an estimated probability of being poor higher than the threshold) declines. Thus, at a threshold of zero, everyone is vulnerable while no one is vulnerable at the threshold of one. Perhaps not surprisingly, for any given threshold, the incidence of vulnerability is higher for the poor than for the population as a whole, which in turn is higher than the incidence of vulnerability amongst the nonpoor. More significantly, Figure 2 suggests that for a wide range of thresholds, poverty and vulnerability are significantly different from each other. Not all the poor are vulnerable while a significant proportion of the nonpoor are vulnerable.

By depicting the entire distribution of vulnerability, Figure 2 provides a wealth of information about the vulnerability of the population. But, in many instances it will not be feasible to directly compare entire distributions and we will need to summarize the key properties of the underlying distribution through some well-chosen summary measures. One such measure, one that is particularly useful for expositional purposes, is the fraction of the population that has a vulnerability level above some threshold and can therefore be deemed vulnerable. The choice of a vulnerability threshold is of course, ultimately somewhat arbitrary. However, a threshold of 0.50 stands out as a possible focal point in that a household whose vulnerability level exceeds 0.50 is more likely than not to end up poor. Even then there remains the question of the time horizon over which a household's vulnerability to poverty should be assessed. Here again, a certain degree of arbitrariness is unavoidable. We consider two possibilities—a time horizon of one year, which can be thought of in terms of the likelihood of poverty in the near future (or short-term), and a time horizon of three years, which roughly corresponds to the likelihood of poverty in the medium-term. We classify as *vulnerable* all households who we estimate to be more likely than not to be poor at least once in the next three years. Of these households, we label as those who ith estimated vulnerability levelsand that is a threshold of 0.50. so for ease of discussion

Table 1 uses this classification scheme to summarize the distributions depicted in Figure 2. At the aggregate level, while 26% of the population is observed to be poor, we estimate that 37% of the population is vulnerable to poverty. Hence, there clearly are households who are observed to be currently (i.e., in 1985) non-poor whose ex-ante probability of poverty is nevertheless estimated to be quite high, so much so that they are more likely than not to be poor at some point in the next three years. In fact, of the 74% of the population that is observed to be non-poor, over 21% are estimated to be vulnerable. This implies that nearly 16% of the population, though not currently poor is vulnerable to poverty. These estimates therefore appear to support the often-stated (and vaguely defined) claim that the observed incidence of poverty underestimates the fraction of the population that is vulnerable to poverty.Amongst the poor, 78% are estimated to be vulnerable.

On the other hand, from the third column of Table 1 it is also apparent that there are some households who are observed to be poor, whose vulnerability level is, nevertheless, low enough for them to be classified as non-vulnerable. In particular we estimate that 20% of the observed poor is non-vulnerable. And while that may, at first glance, seem surprising, it simply reflects the stochastic nature of the relationship between poverty and vulnerability that underlies the distinction between the two concepts.

Amongst those we classify as vulnerable, 61% are estimated to be highly vulnerable implying that the highly vulnerable make up nearly 23% of the overall population. And of the highly vulnerable, only 70% are observed to be poor, which implies that nearly 7% of the population is highly vulnerable but currently non-poor.

The main message that emerges from considering these aggregate numbers is that while poverty and vulnerability are closely related concepts, there remain important distinctions between the two and neither notion nests the other. And this, in turn, has two important implications for policy. First, the fraction of the population that faces a non-negligible risk of poverty (and hence, by definition, is taken to be vulnerable) may be quite different from the fraction that is observed to be poor in any given period. In this particular population the former is estimated to be higher than the latter but that finding is less important than the fact that the two are quite different. Second, and more importantly, these numbers suggest that the characteristics of those who are observed to be poor at any given point in time may well differ from the characteristics of those who are estimated to be vulnerable to poverty, or equivalently, that the relative incidences of vulnerability and poverty may differ across segments of the population. Interventions and programs that aim to reduce the level of vulnerability in the population may therefore need to be targeted differently from those aimed at poverty alleviation. We illustrate this point in the next section with data from Indonesia.

4.2. Comparing poverty and vulnerability profiles

Table 2 presents the poverty and vulnerability profiles for Indonesia in December 1998. We report both the overall estimates for rural and urban Indonesia and also disaggregated by regions and certain select demographic and community characteristics. Table 3 provides us with some insights on average, about the geographical location of the vulnerable as well as their socio-economic characteristics.

We begin by detailing the spatial distribution of poverty and vulnerability. Poverty and vulnerability in Indonesia are largely rural phenomena. Relative to their share in the population, rural households are over-represented among the poor and the vulnerable. While 61% of Indonesia's population is rural, 80% of the observed poor live in rural areas as do 82% of those we estimate to be vulnerable. The highly vulnerable are even more disproportionately rural, with 91% of this group located in rural areas. The disproportionate contribution of rural households to overall poverty and vulnerability stems from the much higher incidence of poverty and vulnerability in rural areas. About 30% of the rural population is observed to be poor, whereas in urban areas, the observed poverty rate is 12%. Similarly, while we estimate that 20% of the urban population is vulnerable, 60% of the rural population is estimated to be vulnerable.

The imbalances in the contributions of rural and urban areas to overall poverty and vulnerability are reproduced at the regional level. Urban areas, regardless of region, are under-represented among the poor and the vulnerable, relative to their shares in the population. With the exception of rural Sumatra, rural areas tend to be over-represented. In absolute terms, rural areas of Java, Kalimantan and Sulawesi contribute the largest numbers to the populations of the poor and vulnerable. And of the 9% of the population that we estimate to be highly vulnerable, a fifth are found in rural areas of Kalimantan and Sulawesi and another 20% live in rural areas of West Java.

The tremendous variation in the poverty rates across the far-flung regions of Indonesia has been documented elsewhere (see Pradhan et. al (2000)). The fifth column of Table 3 confirms the presence of these regional disparities. The fraction of the population that is observed to be poor ranges from a low of 2% in Jakarta to a high of 56% in rural areas of West and East Nusa Tengarra, Papua and Maluku (which have collectively been labeled "Rest of Indonesia"). Except for Central Java and Yogyakarta, where 22% of the urban population is observed to be poor, urban areas have lower observed poverty rates than rural areas.

Inter-regional differences in the estimated incidence of vulnerability are even more pronounced than the regional disparities in poverty rates. The fraction of the population estimated to be vulnerable ranges from a low of 2% in Jakarta to a high of 77% in rural Central Java and Yogyakarta. Again, while urban areas generally have lower vulnerability rates, Central Java and Yogyakarta are exceptional in that 46% of the urban population in these two provinces is estimated to be vulnerable.

A comparison of the observed poverty rates and the estimated incidences of vulnerability across the 13 geographic domains we have defined reveals two points, both indicative of the ways in which the distribution of vulnerability can differ across regions.

First, in keeping with our findings at the national level, in each of the domains, the estimated incidence of vulnerability is at least as high and in most cases higher, than the observed incidence of poverty. However, there is considerable variation in the ratio of the fraction of the population that is vulnerable to the fraction that is poor. The vulnerability to poverty ratio is 1.00 in Jakarta and 1.27 in urban Sumatra indicating that vulnerability to poverty is quite concentrated in these two regions. In contrast, in several other regions, mostly rural, vulnerability to poverty is dispersed in the population, with the fraction that is vulnerable more than the double the fraction that is poor.

Second, two regions with roughly similar observed poverty rates may have very different incidences of vulnerability. For instance, in both East Java and Bali and what we term the "Rest of Indonesia", about 8% of the urban population is observed to be poor. However, we estimate that only 10% of the population of urban East Java and Bali is vulnerable, whereas in the "Rest of Indonesia," over 21% of the urban population is vulnerable.

Turning next to the other correlates of poverty and vulnerability, the one that stands out is the educational attainment of the household head. Of the 69% of the population that lives in households headed by individuals with at most a primary school education-who comprise 88% of the poor and an overwhelming 95% of the vulnerable-nearly 30% are poor while 63% are vulnerable to poverty.

Within this group, households headed by individuals with no schooling are particularly at risk-28% of the population in such households is estimated to be highly vulnerable. In sharp contrast, within the populations in the two highest educational attainment categories, which together make up 21% of the overall population, the observed poverty rate is only 5%, the vulnerability rate is 2% and the fraction that is vulnerable is less than 1%. Even among households headed by individuals with at most junior schooling, the poverty rate, at 12%, is less than half that for households just one step down in the educational attainment hierarchy. The drop in the incidence of vulnerability to just 14% from 61% is even more striking.

If we divide up the sample according to the employment status of the household head we do not get such a clear trend though the incidence of vulnerability is understandably lower for salaried workers in the public and private sectors than it is for those in other employment categories. Somewhat surprisingly, the group with the highest rates of poverty and vulnerability is those who are self-employed with some help from family and hired workers. Of the 31% of the population belonging to this group, more than half are vulnerable.

When the population is split along other demographic characteristics, there is, surprisingly, hardly any difference in the poverty and vulnerability rates for different groups. So for instance, households with high dependency ratios are as likely to be poor and vulnerable as households with low dependency ratios, and households headed by females are as likely to be poor and vulnerable as male-headed households. Perhaps the only difference of note is the higher fraction of female headed households that is estimated to be highly vulnerable.

Community characteristics such as the availability of transport facilities, the presence of a bank or cooperative in the community, industrial activity and access to clean water are all associated with lower levels of vulnerability and poverty. Of these, access to clean water is associated with the sharpest drops in poverty and especially vulnerability.

4.3. Using vulnerability assessments in geographic targeting

The targeting of poverty alleviation resources is often based on the geographic distribution of poverty. For instance, in China counties that are classified as national poor or provincial poor counties (based on assessments of the extent of poverty) selectively receive additional government support. And in India, plan allocations to the states at least partially reflect the degree of need as captured by the level of poverty. Targeting on the basis of geography is also implicit in the various formulae used to determine the allocation of funds under the numerous devolution schemes that have been introduced in recent years.

In the Philippines, the main block grant from the central government to local government units (LGUs), the Internal Revenue Allocation (IRA), is not explicitly based on a poverty criterion, but the argument has often been made that it should be (World Bank, 2000). While this is plausible from an ex-post redistributive perspective, there is the further consideration whether these allocations ought to be linked to the extent of vulnerability if the aim is to reach those most prone to poverty in the near future. The issue is important insofar as those prone to being poor differ from those currently observed to be poor.

Figure 3 plots the observed incidence of poverty (on the horizontal axis) against the estimated incidence of vulnerability for each of the 77 provinces in the Philippines. For most provinces, the estimated incidence of vulnerability is higher (often considerably higher) than the observed incidence of poverty. This can be seen from the fact that most of the points lie above the 45-degree line. In some provinces, the ratio of the vulnerable to the poor is over 4.

More noteworthy still is the substantial re-ranking that takes place when provinces are ordered in terms of the incidence of vulnerability rather than the observed incidence of poverty. Because the provinces are ordered along the horizontal axis in terms of increasing incidence of poverty, the re-ranking is reflected in the non-monotonicity of the scatter plot. Note the three bottom (poorest) and the two top (richest) provinces have the same poverty and vulnerability rankings. But between the two tail-ends there is a lot of re-sorting. The re-rankings are particularly stark for provinces that appear in the upper left and lower right quadrangles defined by the vertical and horizontal lines indicating, respectively, the poverty and vulnerability rates at the national level. The poverty rate in these provinces is below the national rate, and so any poverty-targeting scheme based on poverty rates would allocate relatively fewer funds, on a per-capita basis for these provinces. However, in terms of the incidence of vulnerability to poverty, these provinces are above the national rate, and should, in principle, receive, on per-capita basis, proportionally more funds for poverty programs. Thus, vulnerability-based allocations could differ significantly from observed poverty-based allocations.

The key to resolving this apparent dilemma lies in distinguishing ex-ante poverty prevention interventions from ex-post poverty alleviation interventions. An example drawn from public health makes this distinction clearer. Consider a situation where public health interventions are aimed at reducing the incidence of some disease. Suppose information is available on both the incidence of disease in different regions, as well as on the fraction of the population in different regions that is at high risk of contracting the disease. Funds for treatment of those already afflicted should clearly be directed to regions where the incidence of the disease is highest. But funds for preventive measures (such as vaccinations) ought to be directed to regions where the fraction of the population at risk is the largest. And the two sets of regions need not coincide. Regions with a higher incidence of the disease may also be regions where the risk of contracting the disease is concentrated among those afflicted. So the fraction of the population at risk may well be lower than in other regions where the incidence of the disease is lower.

The analogy with our treatment of vulnerability should be clear. The incidence of poverty, like the incidence of the disease, should determine the allocation of funds for treatment, which in the case of poverty means funds for ex-post poverty alleviation programs. The allocation of funds for preventive interventions-ex-ante interventions aimed at poverty prevention-should however be guided by the incidence of vulnerability to poverty.

In practice, the difference between ex-ante and ex-post interventions will most likely be realized in terms of the particular line agencies through which resources are channeled. The funds for focused ex-post interventions such as food-for-work schemes or means-tested transfer programs are likely to be disbursed through very different channels than funds for ex-ante interventions. The latter will in general be much more varied in nature, and depending on the context may range from vocational training schemes, agricultural extension programs, social investment funds to major irrigation projects.

4.4. Exploring the proximate causes of vulnerability

Consider Figure 4, which shows the simulated consumption streams (over a 50-period time horizon) for two different households.[2] The consumption streams of the two households look very different. Household A, on average, enjoys a much higher level of consumption, but its consumption is quite volatile. Household B, on the other hand, has a relatively stable inter-temporal consumption profile, but with much lower levels of consumption, on average. What is special about these two households is that despite the obvious differences in their mean levels of consumption and in the volatility of their consumption streams, the simulations have been constructed so that *their vulnerability levels are the same.*

Figure 4 illustrates, rather starkly, the general point that households with similar levels of vulnerability may be vulnerable for very different reasons.[3] For some,

[2]The simulations are based on actual estimates of mean consumption and consumption variance for two households in the rural southern China panel.

[3]Conversely, two households with the same mean level of consumption may have very dif-

vulnerability may stem primarily from low long-term consumption prospects (household B above). For others, consumption volatility may be the main source of vulnerability to poverty (household A above). From a policy perspective it will be important to distinguish between these two possibilities. For instance, vulnerability due to high volatility may call for ex-ante interventions that reduce the risks faced by households or insure them against such risks. On the other hand, to address vulnerability due to low endowments what might be needed are transfer programs. Clearly, a decomposition of the sources of vulnerability at the household level into the two components described above can help inform that choice.

At the same time it should be recognized that the two possibilities represent stylized extremes which are potentially interconnected in subtle ways. For instance, it may be that with inadequate risk management instruments at their disposal, households forego risky but, on average, high return earnings opportunities in favor of lower risk but lower return income streams. And in that case while the vulnerability of the household may appear to be due to low endowments, the true source of vulnerability may lie in an inability to adequately deal with risk.

Figure 4 also illustrates another important point, which is the mean and standard deviation of consumption need not be monotonically related across households. In the case of rural southern China that is clearly not the case, as Figure 5 illustrates. Though, for each of the four provinces, there appears to be a strong positive association between the estimated mean and the estimated variance of consumption, also visible are numerous instances where a household has both a higher estimated standard deviation of consumption as well as a lower estimated mean level of consumption than several of the other households. This possibility for a household with a lower mean level of consumption to face greater consumption volatility is, as we noted earlier, not allowed in the methods used in most poverty assessments. The standard there is to implicitly force the estimated variance of consumption to always be higher for households with higher estimated mean consumptions. Figure 5 therefore highlights the importance of keeping the estimation strategy adequately flexible for the mean and variance of consumption to be separately estimated.

To facilitate the discussion of the proximate causes of vulnerability, amongst those we classify as vulnerable (because they are more likely than not to be poor in the medium term) we distinguish between those who would not be vulnerable in the absence of consumption volatility and those who are *structurally poor.*

ferent levels of vulnerability if the degree of consumption volatility they are subject to, differs substantially.

For the former group, who might be said to be *vulnerable to transitory poverty*, interventions that reduce consumption volatility by either reducing their exposure to risk or enhancing their ex post coping capacity would be sufficient to reduce vulnerability. For the latter group, however, risk-reducing interventions alone may be inadequate, and must most likely be accompanied by interventions that improve their mean livelihood prospects.

There is an obvious parallel between the classification we propose above and the more familiar distinction between the transient poor and the chronic poor. Loosely speaking, households who are vulnerable to transitory poverty are in a sense more likely to be only transitorily poor, whereas households who are structurally poor are more likely to be chronically poor. But the parallel should not be taken too far because there are important distinctions between the two classification schemes. Households that are, under our classification, vulnerable to transitory poverty have very high levels of vulnerability and may therefore be poor more often than not. Should these households be included among the transient poor? Ultimately, the two taxonomies differ fundamentally because of the different questions they pose. The distinction between the transient poor and the chronic poor is based on the question: how often is the household poor? On the other hand the distinction we propose is based on the question: why is the household poor?

Table 4 provides a breakdown of the proximate cause of vulnerability in rural southern China as of 1985. Of the population as a whole, we estimate that 16% is vulnerable due solely to consumption volatility, while 21% are vulnerable because they are structurally poor. Thus, of the 37% of the population that is vulnerable, over a half are so due to structural poverty. Consumption volatility is also the main source of vulnerability for those currently poor.

Of the 80% of the poor whom we estimate to be vulnerable, nearly a quarter are vulnerable because their consumptions are volatile. Put another way, 23% of the poor would not be poor if ways could be found to stabilize their consumption streams, while maintaining their mean consumption levels.[4]

Table 4 reveals several interesting patterns in the way the proximate causes of vulnerability vary across various segments of the population. For instance, higher road density lowers the incidence of vulnerabilty to structural poverty from 0.24 to 0.17, but does not affect the incidence of vulnerability to transitory poverty.

[4] This last qualifier is important because, even without any public intervention there might well have been ways in which these households could have reduced the volatility of their consumption streams. That they ”chose” not to do so suggests that the cost incurred in terms of a reduction in mean consumption, in stabilizing consumption may have been too high.

The most striking pattern recorded in Table 4 is the geographic variation in the relative importance of the two proximate causes of poverty. Figure 6 graphically illustrates this variation.

4.5. Estimating the contribution of risk to vulnerability

Exposure to risk is obviously the primary determinant of vulnerability for those who are vulnerable solely due to consumption volatility. However, risk can potentially be a significant factor in the vulnerability of even those we estimate to be structurally poor. Using the estimates of the mean and variance of the consumption process at the household level, it is possible to estimate the contribution of risk to the vulnerability levels of individual households. The basic step involved is in estimating the counterfactual vulnerability level of a household in the absence of risk, that is, if the household's consumption in every period were to be fixed at its mean level of consumption.

Table 5 displays the average fraction of estimated vulnerability levels—where vulnerability is defined in terms of the expected poverty gap—that is attributable to consumption volatility for various segments of the population. What is striking is that, except in a few instances, notably, hilly areas of Guizhou province, 25% or more of the vulnerability level of even structurally poor households can be attributed to risk.

4.6. Predicting future poverty using current vulnerability estimates

Lastly, we demonstrate the usefulness of vulnerability estimates, even those generated from a single cross-section, in predicting future poverty. We use data from a single cross-section of a two-year panel, 1997 in the case of the Philippines, 1998 for Indonesia, to obtain estimates of vulnerability for each household. We order and group households into quintiles (Philippines) or deciles (Indonesia) based on these vulnerability estimates. We then compare the predicted poverty rate for each quintile or decile with the *actual incidence of poverty* in the next year, 1998 for the Philippines and 1999 for Indonesia.

Figure 7 illustrates the results we obtained. For each of the two countries, this figure presents a comparison of the predicted poverty rate (i.e., mean estimated vulnerability level from the earlier cross-section) and the actual poverty rate in the later year for each decile or quintile of the vulnerability distribution estimated using the earlier cross-section. Keeping in mind that the period in question was one where the Philippine economy was beginning to feel the ripple effects of the Asian financial crisis and the El-Nino induced drought, and the Indonesian economy was recovering from the Asian crisis, it is striking that our vulnerability estimates, by and large, reproduce the ordinal properties of the true distribution of vulnerability in the population.

5. Appendix

5.1. The data and the setting in the three country studies

Philippines

The data for the Philippines are from the 1997 Family Income and Expenditure Survey (FIES) and the 1998 Annual Poverty Indicators Survey (APIS). These two surveys span a period when the Philippine economy was beginning to feel the ripple effects of the Asian financial crisis, the effects of which were compounded by an El Nino-induced drought beginning around September 1997.

The FIES, which is conducted by the National Statistics Office (NSO) of the Government of the Philippines, is the main survey used for generating poverty and income distribution statistics in the Philippines. It is conducted every three years. The 1997 round sampled 39,520 households using urban and rural areas of each province as principal domains. The survey provides detailed household income and consumption data, and some basic information on household attributes. APIS provides data, not only on incomes and expenditures, but also on a wide range of variables such as health, education, family planning, and family access to housing, water and sanitation, credit, and a number of community (barangay) characteristics. The 1998 APIS covered 38,710 sample households. Because APIS was designed to be a longitudinal survey forming a panel with the 1997 FIES, 23,150 households were common to both surveys. These panel households constitute my sample.

There are however some issues of comparability between the FIES and the APIS data. The APIS uses a much shorter consumption module of just two pages (27 expenditure lines), compared to over 20 pages (over 400 expenditure lines) in the FIES. To sidestep this problem—though perhaps not entirely satisfactorily—we normalize the household-level consumption aggregates from each of the surveys by the relevant poverty lines at the two dates.

We use the poverty lines developed by Balisacan (1999) that correspond to a nutritional norm of 2000 calories per person per day and allow for basic nonfood expenditure. Balisacan (1999) estimated a set of provincial poverty lines which provides estimates of spatial cost of living differentials. We use the Manila poverty line of P10,577 per person per year for 1997 (and P11,677 per person per year for 1998), while the provincial poverty lines are used to express nominal consumption of all households into 1997 Manila prices (see Balisacan (1999) for further details).

Indonesia

The data for Indonesia come from two sources. The main data on household characteristics and consumption expenditures come from the Mini-SUSENAS, which is a smaller version of the SUSENAS (National Socio-Economic Survey) that is the primary household expenditure survey in Indonesia.[5] We combine these with data from the 1996 "Village Potential" (PODES) Survey which provides a wide range of information on the characteristics of the villages/communities ("desa") in which these households reside.

The Mini-SUSENAS survey was first conducted in December 1998 and again in August 1999, using the same sample frame, and moreover, with about 75% of the original 10,000 or so households being surveyed on both occasions. The Mini-SUSENAS therefore provides a 2-period panel for roughly 7,500 households. The time period spanned by the two rounds of the panel was one during which the Indonesian economy was recovering from the financial crisis. By December 1998 the rupiah had stabilized and by the middle of 1999 democratic elections had been held.

To normalize the consumption levels I used the poverty lines used by Chaudhuri, Jalan and Suryahadi (2001). The poverty lines were constructed starting with the set of regional poverty lines for February 1999 calculated by Pradhan et al. (2000). These were then deflated to December 1998 and August 1999, using as deflators, a set of re-weighted provincial CPIs (Pradhan et al. (2000)). The Indonesian CPI has a food share of 0.4, while the food share of the poverty lines is 0.8, reflecting the importance of food to the poor. So for each province a re-weighted CPI with a food share of 0.8 was re-calculated. Another weakness of the CPI is that it is based solely on urban prices. Unfortunately, this weakness carried over to the re-weighted CPI. Moreover the same deflator was used for urban and rural poverty lines within a province, which amounts to assuming that the inflation rates in urban and rural areas in a province, during the period of interest—December 1998 to August 1999—were the same.

Rural southern China

We use a panel data set constructed from the Rural Household Budget Surveys (RHS) implemented by China's State Statistical Bureau (SSB).The RHS is a well-designed and executed budget survey of a random sample of households drawn from a sample frame spanning rural China (including small-medium towns), and with unusual effort made to reduce non-sampling errors. Sampled households keep a daily record of all transactions, and log books on production. Interviewing

[5] Details about the Mini-SUSENAS survey, and the procedure used to construct the consumption aggregates that we use are available in BPS (2000).

assistants visit each sampled household every two weeks to check on their progress and collect the data. Checks are made at the county statistical office, with return visits to the households when necessary. The consumption data obtained from such an intensive survey process are almost certainly more reliable than those obtained by the common cross-sectional surveys in which the consumption data are based on recall at a single interview.[6]

The household data are collated with geographic data at the village, county and the province levels. At the village level, we have data on topography (whether village is in plains, or in hills, or in mountains), on location (whether it is in a coastal area), ethnicity (whether it is a minority village or not), and whether the village is in a revolutionary base area (areas where the Communist Party had established its bases prior to 1949). At the county level we have a much larger database drawn from county administrative records. At the province level we simply include dummy variables for the province. All nominal values have been normalized by 1985 prices.

We use a sample of 5,820 households observed over the six-year period 1985-90 from four contiguous provinces in southern China, namely Guangdong, Guangxi, Guizhou, and Yunnan (with roughly equal numbers of sampled households in each) having a total population of 176 million in 1990. The region is a fairly good representation of the current regional disparities in rural China. Three of the provinces (Guangxi, Guizhou and Yunnan) form a region of south-west China which is widely regarded as one of the poorest regions in the country. Guangdong, on the other hand, is a relatively rich coastal region. For example, in 1990, using the squared poverty gap measure, the severity of poverty in Guizhou was estimated to be 3.26% compared to less than 0.15% in Guangdong.

Consumption expenditure per capita is the individual welfare measure. The consumption measure is comprehensive, in that it includes imputed values for consumption from own production valued at local market prices, and it also includes an imputed value of the consumption streams from the inventory of consumer durables.

The poverty lines are those constructed by Chen and Ravallion (1996). These are based on a normative food bundle set by SSB, which assures that average

[6] Inspite of the care which goes into collecting the household data, there may still arise some measurement error in consumption expenditures on account of imputed values of consumption from own production. However, the consequences of measurement error if any, is not an issue in this paper since we are comparing the cross-sectional dimension to the panel aspect of the same data. So if the cross-section estimates are contaminated by measurement error so will the panel data estimates be.

nutritional requirements are met with a diet which is consistent with Chinese tastes; this is valued at province-specific prices. The food component of the poverty line is augmented with an allowance for non-food goods, consistent with the non-food spending of those households whose food spending is no more than adequate to afford the food component of the poverty line.

5.2. Econometric strategy adopted in the three country studies

Philippines and Indonesia

For the Philippines and Indonesia, the data available are primarily from a single cross-section and are roughly similar in terms of coverage. Hence, except for some minor details, the econometric strategy we employ is identical in the two cases.

We begin by assuming that the stochastic process generating the consumption of a household h is given by:

$$\ln c_h = X_h \beta + e_h \tag{5.1}$$

where c_h is per capita consumption expenditure, X_h represents a bundle of observable household characteristics, characteristics such as household size, location, educational attainment of the household head, etc., β is a vector of parameters, and e_h is a mean-zero disturbance term that captures idiosyncratic factors (shocks) that contribute to different per capita consumption levels for households that are otherwise observationally equivalent.

The set of covariates for the Philippines included: linear and quadratic terms in family size, household composition variables including number of adults (ages 15–60) and the dependency ratio, characteristics of the household head such as indicator variables for female headship, marital status, educational attainment, age and age squared, occupational characteristics including dummy variables for sector of employment where ten different sectors are included, indicator variables for ownership of land, use of electricity, membership in cooperatives, and a range of barangay (community) characteristics.

To allow for spatial heterogeneity in the returns to these characteristics, we estimated (**??**) separately for each of 11 regional domains, starting with the National Capital Region (Metro Manila), which is all urban, and followed by rural and urban domains for each of the following five clusters of provinces: North Luzon, South Luzon, Western Visayas, Eastern Visayas and Mindanao.

In the case of Indonesia the covariates we included in X_i were: household size (level and its square), proportion of household members in the age-groups 6-12 years, 13-15 years, 16-18 years, proportion of adults in the household, whether the head of the household is single, married, divorced, age and age-squared of the head of household, and a series of dummies for whether the household head is illiterate, has attended primary school, attended junior-high school, attended senior high school, has some tertiary eduction; whether the head of the household is male, whether the household head is self-employed with no assistance, self-employed with some assistance from family and temporary workers, self-employed with permanent employees, and salaried workers in either the government or private sector.

We estimated (**??**) separately for each of 13 geographical domains-the province of Jakarta (which is completely urban) and the rural and urban areas of the following six clusters of provinces: Sumatra, West Java, Central Java and Yogyakarta, East Java and Bali, Kalimantan and Sulawesi, and the rest of Indonesia.

We assume that the variance of e_h is given by:

$$\sigma^2_{e,h} = X_h \theta \tag{5.2}$$

We estimate β and θ using a three-step feasible generalized least squares (FGLS) procedure suggested by Amemiya(1977).

First we estimate equation (5.1) using an ordinary least squares (OLS) procedure. We use the estimated residuals from equation (5.1) to estimate:

$$\widehat{e}^2_{OLS,h} = X_h \theta + \eta_h \tag{5.3}$$

using OLS. The predictions from this equation are used to to transform the equation as follows:

$$\frac{\widehat{e}^2_{OLS,h}}{X_h \widehat{\theta}_{OLS}} = \left(\frac{X_h}{X_h \widehat{\theta}_{OLS}} \right) \theta + \frac{\eta_h}{X_h \widehat{\theta}_{OLS}} \tag{5.4}$$

This transformed equation is estimated using OLS to obtain an asymptotically efficient FGLS estimate, $\widehat{\theta}_{FGLS}$. Note that $X_h \widehat{\theta}_{FGLS}$ is a consistent estimate of $\sigma^2_{e,h}$, the variance of the idiosyncratic component of household consumption.

The estimates:

$$\widehat{\sigma}_{e,h} = \sqrt{X_h \widehat{\theta}_{FGLS}} \tag{5.5}$$

are then used to transform equation (5.1) as follows:

$$\frac{\ln c_h}{\widehat{\sigma}_{e,h}} = \left(\frac{X_h}{\widehat{\sigma}_{e,h}} \right) \beta + \frac{e_h}{\widehat{\sigma}_{e,h}} \tag{5.6}$$

OLS estimation of equation (5.6) yields a consistent and asymptotically efficient estimate of β. The standard error of the estimated coefficient, $\widehat{\beta}_{FGLS}$, can be obtained by dividing the reported standard error by the standard error of the regression.

Using the estimates $\widehat{\beta}$ and $\widehat{\theta}$ that we obtain we are able to directly estimate expected log consumption:

$$\widehat{E}\left[\ln c_h \mid X_h\right] = X_h\widehat{\beta} \tag{5.7}$$

and the variance of log consumption:

$$\widehat{V}\left[\ln c_h \mid X_h\right] = \widehat{\sigma}^2_{e,h} = X_h\widehat{\theta} \tag{5.8}$$

for each household h. By assuming that consumption is log-normally distributed, we are then able to use these estimates to form an estimate of the probability that a household with the characteristics, X_h, will be poor, i.e, to estimate the household's vulnerability level. Letting $\Phi(.)$ denote the cumulative density of the standard normal, this estimated probability will be given by:

$$\widehat{v}_h = \widehat{\Pr}\left(\ln c_h < \ln z \mid X_h\right) = \Phi\left(\frac{\ln z - X_h\widehat{\beta}}{\sqrt{X_h\widehat{\theta}}}\right) \tag{5.9}$$

Two substantive issues arise in the implementation of the procedure outlined above, both having to do with the estimation of the variance of consumption. The first has to do with the possibility of measurement error in the observed data on consumption expenditures. Measurement error is a major concern in most consumption (and income) measures drawn from household surveys. The presence of such errors can lead to significant overestimates of the variance of log consumption from (5.3) and (5.4). Why? Because the mean of the squared residuals from (5.1) will be biased upwards by the variance of the measurement error and that bias will be transmitted to the estimate of the intercept in equations (5.3) and (5.4). And if that were the case, we would overestimate predicted mean consumption *levels* (which, given log-normality of consumption, is an increasing function of the variance of log consumption). To control for this, we make a multiplicative adjustment to the estimated variances such that the predicted mean consumption equals the actual mean consumption for each of the geographic domains for which we estimate a separate set of regressions.

This adjustment also corrects for overestimates of variance because of unobserved, but deterministic components of consumption. For instance, suppose two households look identical in terms of the observables *we include* in the consumption equation (5.1). Nevertheless, they have different consumption prospects because of some unobserved but deterministic factor-e.g. rural cultivating households who live in areas with more fertile soil may have better consumption prospects though they appear to be identical to households in areas with less fertile soil. This will bias upwards the mean of the squared residuals from (5.1).

A second, somewhat different complication stems from the possibility of *unobserved local shocks* that are common to households in particular areas. For instance, suppose a particular area is subject to a localized shock, which is reflected in the consumption data from that area. Households from that area will, depending on whether the shock was positive or negative, have higher or lower consumption levels than otherwise observationally equivalent households from other areas. If we include a set of area dummies in log consumption equation (5.1) to capture the effects of such localized common shocks and include the estimated dummies in estimating the mean of log consumption we would bias (either upwards or downwards) the latter estimate. If we instead include a set of area dummies in the variance-estimating equations, we risk overestimating the variance of log consumption for households in areas that experience large relative shocks. A reason for including area dummies would be to control for unobserved deterministic components of consumption. But since we address that issue through the adjustment we describe above, we chose ultimately not to include any area dummies in any of the regressions we estimated.

A third more minor issue is the fact that, given the simple linear specification we have adopted, there is no guarantee that the estimate of $\sigma^2_{e,h}$, $X_h\widehat{\theta}$, will be positive. In practice we did not find this to be a problem except for a few observations, so we simply dropped them from the sample. An alternative would have been to choose a different specification for the variance-estimating equation (5.3), such as a logistic specification (as in Elbers et al.(2001)). That would force the estimate to always be positive, though the estimate would then have to be constructed from a Taylor approximation.

Rural southern China

With six years of annual data available to us, we adopt the following fairly general specification of the consumption process:

$$\ln C_{ht} = \alpha_j + X_{ht}\beta_j + Z_h\gamma_j + \Gamma_j t + \eta_h + e_{ht} \tag{5.10}$$

Here, $\ln C_{jht}$ is the log per-capita consumption of household h in province j in year t, X_{ht} is a vector of time-varying household characteristics, Z_h is a vector of time-invariant observable household characteristics, Γ_j is a trend growth rate common to all households in province j, η_h is a time-invariant unobservable household-specific effect and e_{ht} is a disturbance term capturing period-specific shocks (both idiosyncratic and covariate) to household consumption as well as measurement error.

As the subscript j on the parameters indicate, we estimate (5.10) separately for each of the four provinces. We adopt this disaggregated estimation strategy because we wished to allow for some heterogeneity in the structural parameters underlying the consumption processes of households in these different regions. Given the differences in the structures of local economies in different provinces, it is likely that key structural parameters—for instance, the returns to education or experience—may differ across regions.

The most commonly identified causes of poverty in rural China are living in remote and mountainous areas, limited transport, power and other rural infrastructure, being a minority ethnic group, being illiterate, and being prone to ill-health (see, for example, World Bank, 1992). The question here is whether these factors are equally important in estimating mean consumption and inter-temporal variance. We therefore include as explanatory variables household-specific human assets, and community effects, the latter measured by a set of county specific variables.

The household variables include: schooling variables (the proportion of adult household members with different levels of schooling (left out category is those with high school or higher levels of education), proportion of children (defined as household members under 15 years of age) with primary and secondary school education (left-out category is the proportion staying at home or are declared to be illiterate); a wide range of demographic variables to (age and age2 of the household head to capture any life-cycle effects, the proportion of kids in the household at various ages with the proportion of children under infants 5 years as the left out category, whether the household is an exclusively farm household and household size. Rural labor markets appear to be thin in this setting, so demographic characteristics of the household can matter to productivity; these variables may also pick up differences in consumption behavior.

The geographic variables are dummies for the topographical features of the county of residence (plains, coast, and mountains), dummy variables for whether the county falls into a specific need-based category (revolutionary base area, minority area or border area), availability of medical facilities, and rural nfrastructure variables like density of roads, cultivated area that is irrigated etc.The most commonly identified causes of poverty in rural China are living in remote and mountainous areas, limited transport, power and other rural infrastructure, being a minority ethnic group, being illiterate, and being prone to ill-health (see, for example, World Bank, 1992). The question here is whether these factors are equally important in estimating mean consumption and inter-temporal variance. We therefore include as explanatory variables household-specific human assets, and community effects, the latter measured by a set of county specific variables.

The household variables include: schooling variables (the proportion of adult household members with different levels of schooling (left out category is those with high school or higher levels of education), proportion of children (defined as household members under 15 years of age) with primary and secondary school education (left-out category is the proportion staying at home or are declared to be illiterate); a wide range of demographic variables to (age and age2 of the household head to capture any life-cycle effects, the proportion of kids in the household at various ages with the proportion of children under infants 5 years as the left out category, whether the household is an exclusively farm household and household size. Rural labor markets appear to be thin in this setting, so demographic characteristics of the household can matter to productivity; these variables may also pick up differences in consumption behavior.

The geographic variables are dummies for the topographical features of the county of residence (plains, coast, and mountains), dummy variables for whether the county falls into a specific need-based category (revolutionary base area, minority area or border area), availability of medical facilities, and rural nfrastructure variables like density of roads, cultivated area that is irrigated etc.

In estimating (5.10), we allow for the potential correlation of the time-invariant household effects with the observable household characteristics, both time-varying and time-invariant, included in the vectors, X_{ht} and Z_h respectively, and treat the η_h as fixed effects. Further, we treat the time-varying household characteristics as *predetermined*, rather than *strictly exogenous* variables from an econometric standpoint. In other words, while we assume that the elements of X_{ht} are *contemporaneously uncorrelated* with e_{ht}, we allow for the possibility that the levels of these variables may in part be determined by past consumption shocks realized by the household. Formally, we assume that:

$$E[X_{ht}e_{ht}] = 0 \tag{5.11}$$

but:

$$E[X_{ht}e_{h,t-k}] \neq 0 \qquad \text{for } k > 0 \tag{5.12}$$

For some of the household characteristics such as the level of financial wealth at the beginning of the year, and the grain stocks held by the household entering the year, (5.12) is clearly the right assumption to make. There is considerable evidence from a variety of settings that households respond to positive and negative income shocks (and hence consumption shocks) by, respectively, accumulating and drawing down assets. And that would imply that a household's asset holdings entering into any given year reflect (and are correlated with) shocks experienced by the household in previous years. For other characteristics, among them household size, levels of educational attainment of various household members, and dependency ratio, a case can be made for strict exogeneity except in the face of large shocks. Nevertheless, we adopt a conservative approach and assume that (5.12) applies for these variables as well.

With fixed effects and predetermined but not strictly exogenous covariates, the standard within-estimator cannot be used to estimate (5.10). We therefore estimate the equation in first differences and instrument the changes in the predetermined variables using lagged changes and levels of the same variables.

References

[1] Alderman, Harold and Christina Paxson.1992. "Do the Poor Insure? A Synthesis of the Literature on Risk and Consumption in Developing Countries," Policy Research Working Papers, Agricultural Policies, WPS 1008, The World Bank, (October).

[2] Alderman, H., J. Hoddinott and B. Kinsey (2002). 'Long term consequences of early childhood malnutrition'. Mimeo, FCND, International Food Policy Research Institute, Washington D.C.

[3] Alderman, H., Behrman, J., Lavy, V. and Menon, R. (2001). 'Child Health and School Enrollment: A Longitudinal Analysis'. Journal of Human Resources, Vol. 36, pp. 185-205.

[4] Alderman, H., P.A. Chiappori, L. Haddad, J. Hoddinott and R. Kanbur, 1995, 'Unitary versus collective models of the household: Is it time to shift the burden of proof?' World Bank Research Observer , vol. 10 no. 1, pp. 1-19.

[5] Amemiya, T. 1977. The maximum likelihood estimator and the non-linear three stage least squares estimator in the general nonlinear simultaneous equation model. Econometrica, vol. 45, pp. 955-968.

[6] Amin, Sajeda, Ashok Rai, and Giorgio Topa .2003. "Does Microcredit Reach the Poor and Vulnerable? Evidence from Northern Bangladesh." Journal of Development Economics 70: 59-82.

[7] Bell, C. 1988. Credit markets and interlinked transactions, In H. Chenery and T.N. Srinivasan (eds.), Handbook of Development Economics vol. 1, Amsterdam: North Holland.

[8] Blundell, R. and T. MaCurdy (1999). "Labor Supply: a Review of Alternative Approaches," in O. Ashenfelter and D., eds., Handbook of Labor Economics, Vol. 3. Amsterdam: North-Holland Publishing Co.

[9] Chaudhuri, S., J. Jalan, and A. Suryahadi. 2002. "Assessing Household Vulnerability to Poverty: A Methodology and Estimates for Indonesia," Columbia University Department of Economics Discussion Paper No. 0102-52, New York: Columbia University.

[10] Chaudhuri, S. and L. Christiaensen. 2002. "Assessing Household Vulnerability to Poverty: Illustrative Examples and Methodological Issues" Presentation at the IFPRI-World Bank Conference on "Risk and Vulnerability: Estimation and policy applications," September 23-24, 2002, Washington DC.

[11] Christansen, Luc and Kalanidhi Subbarao.2001. "Towards an Understanding of Vulnerability in Rural Kenya," Mimeo, The World Bank: Washington DC.

[12] da Corta, L. and D. Venkateshwarlu, 1992. Field methods for economic mobility, in S. Devereux and J. Hoddinott (eds) Fieldwork in developing countries, London: Harvester Wheatsheaf.

[13] Deaton, A., 1992, Understanding consumption, Oxford: Clarendon Press.

[14] Deaton, A., 1997, The analysis of household surveys, Baltimore: ohns Hopkins University Press.

[15] Deaton, Angus and Christina Paxson (1994) "Intertemporal Choice and Inequality," Journal of Political Economy, Vol. 102, No. 3, pp. 437-467.

[16] Deolalikar, A. (1996). 'Child nutritional status and child growth in Kenya: Socioeconomic determinants'. Journal of International Development, Vol. 8, pp. 375-394.

[17] Dercon, S. 1996. Risk, crop choice, and savings: Evidence from Tanzania. Economic Development and Cultural Change, vol. 44, pp. 485-513.

[18] Dercon, Stefan (2002) "Risk-Sharing and Vulnerability" Comments prepared for a workshop on measuring vulnerability, June 27, 2002.

[19] Dercon, S. and P. Krishnan. 1996. Income portfolios in rural Ethiopia and Tanzania: Choices and constraints. Journal of Development Studies, vol. 32, pp. 850-875.

[20] Dercon, S. and P. Krishnan, (2000a). 'In sickness and in health: Risk-sharing within households in rural Ethiopia.' Journal of Political Economy, Vol. 108, pp. 688-727.

[21] Dercon, S., and P. Krishnan. 2000b. Vulnerability, seasonality, and poverty in Ethiopia. Journal of Development Studies 36 (6): 25-5

[22] Fafchamps, M., 1993, 'Sequential labor decisions under uncertainty: An estimable household model of West-African farmers', Econometrica, vol. 61 no. 5, pp. 1173-1198.

[23] Fafchamps, M., C. Udry, and K. Czukas. 1998. Drought and saving in West Africa: Are livestock a buffer stock? Journal of Development Economics 55 (2): 273-305.

[24] Foster, A., 1995, 'Prices, credit markets and child growth in low-income rural areas', Economic Journal, vol. 105 no. 430, pp. 551-570.

[25] Foster, J., J. Greere, and E. Thorbecke. 1984. "A class of decomposable poverty measures" Econometrica 52 (3): 761-766.

[26] Gertler, P., E. Rose and P. Glewwe, 2000. Health in. M. Grosh and P. Glewwe (eds), Designing household survey questionnaires for developing countries, vol 1. World Bank: Washington D.C.

[27] Glewwe, P. and G. Hall, 1998, 'Are some groups more vulnerable to macroeconomic shocks than others? Hypothesis tests based on panel data from Peru', Journal of Development Economics, vol. 56 no. 1, pp. 181-206.

[28] Haddad, L. J. Hoddinott and H. Alderman (eds). 1997. Intrahousehold resource allocation in developing countries: Models, methods and policy. Johns Hopkins University Press: Baltimore.

[29] Harrower, S. and J. Hoddinott (2002) "Consumption Smoothing and Vulnerability in the Zone Lacustre, Mali," Draft. International Food Policy Research Institute (IFPRI): Washington DC.

[30] Heitzmann, K. R. S. Canagarajah, and P. B. Siegel. 2002. Guidelines for Assessing the Sources of Risk and Vulnerability. Social Protection Discussion Paper 0218. Washington DC: World Bank.

[31] Hoddinott, J. and Kinsey, B. (2001). 'Child health in the time of drought.' Oxford Bulletin of Economics and Statistics, 2001, vol. 63, pp. 409-436.

[32] Hoddinott, J. and Kinsey, B. (1999). 'Adult health in the time of drought'. mimeo, International Food Policy Research Institute, Washington D.C.

[33] Hoddinott, J. and A. R. Quisumbing 2003. Data Sources for Microeconometric Risk and Vulnerability Assessments. Manuscript, International Food Policy Research Institute, Washington DC.

[34] Holzmann, R. 2001. Risk and Vulnerability: The forward looking role of social protection in a globalizing world. Paper prepared for "The Asia and Pacific Forum on Poverty-Policy and Institutional Reforms for Poverty Reduction," Asian Development Bank, Manila, February 5-9, 2001.

[35] Holzmann, R. and S. Jorgensen. 2000. "Social Risk Management: Anew Conceptual Framework for Social Protection and Beyond" Social Protection Discussion Paper No. 0006, The World Bank.

[36] Jacoby, H. and E. Skoufias, 1997, 'Risk, financial markets, and human capital in a developing country', Review of Economic Studies, vol. 64 no. 3, pp. 311-336.

[37] Jacoby, H. and E. Skoufias. 1998. "Testing Theories of Consumption Behavior using Information on Aggregate Shocks: Income Seasonality and Rainfall in Rural India," American Journal of Agricultural Economics, Vol. 80, No. 1, (February), pp. 1-14.

[38] Kamanou, Gisele and Jonathan Morduch. 2002. "Measuring Vulnerability to Poverty," NYU Wagner Working Paper No. WP1012. New York: New York University.

[39] Kochar, A., 1999, 'Smoothing consumption by smoothing income: Hours-of-work responses to idiosyncratic agricultural shocks in rural India', Review of Economics and Statistics, vol. 81 no. 1, pp. 50-61.

[40] Ligon Ethan (2001) " Targeting and Informal Insurance" UC Berkeley Working Paper.

[41] Ligon, E. and L. Schechter. 2002. Measuring Vulnerability. Manuscript, University of California, Berkeley.

[42] Mace, Barbara (1991): "Full Insurance in the Presence of Aggregate Uncertainty," Journal of Political Economy, Vol. 99, pp.928-56.

[43] Martorell, R., J. Behrman, H. Delgado, R. Flores, J. Hoddinott, J. Maluccio, A. Quisumbing, U. Ramakrishnan, A. Stein and B. Torun, 2000. Early nutrition, human capital and economic productivity. Research proposal, Emory University.

[44] McCloskey, D., 1976, 'English open fields as behavior towards risk', Research in Economic History, vol. 1 no. 2, pp. 124-171.

[45] McElroy, M. 1990. The empirical content of Nash-bargained household behavior. Journal of Human Resources, vol. 25, pp. 559-583.

[46] Morduch, J., 1990, 'Risk, production and saving: Theory and evidence from Indian households', mimeo, Harvard University.

[47] Morduch, J., 1995, 'Income smoothing and consumption smoothing', Journal of Economic Perspectives, vol. 9 no. 2, pp. 103-114.

[48] Morduch, J., 1999, 'Between the market and state: Can informal insurance patch the safety net?', World Bank Research Observer, vol. 14 no. 2, pp. 187-208.

[49] Pritchett, Lant, Asep Suryahadi and Sudarno Sumarto (2000) "Quantifying Vulnerability to Poverty: A Proposed measure, Applied to Indonesia," Policy Research Working Paper no. 2437. The World Bank: Washington DC.

[50] Quisumbing A. 2002a. "Consumption Smoothing, Vulnerability, and Poverty in Rural Bangladesh," Draft. International Food Policy Research Institute (IFPRI): Washington DC.

[51] Quisumbing, A. 2002b. "Consumption Smoothing, Vulnerability, and Poverty in Rural Ethiopia," Draft. International Food Policy Research Institute (IFPRI): Washington DC.

[52] Quisumbing, A. R. and J. A. Maluccio. 2003. "Resources at Marriage and Intrahousehold Allocation: Evidence from Bangladesh, Ethiopia, Indonesia, and South Africa" Oxford Bulletin of Economics and Statistics, forthcoming.

[53] Rose, E. 1999. Consumption smoothing and excess female mortality in rural India. Review of Economics and Statistics 81 (1): 41-49.

[54] Rosenzweig, Mark R. and Kenneth I. Wolpin (1993): "Credit Market Constraints, Consumption Smoothing and the Accumulation of Durable Production Assets in Low-Income Countries," Journal of Political Economy, Vol. 101.

[55] Schechter Laura.2001. "Identifying Vulnerable Households" Draft. Dept of Agricultural Economics , University of California at Berkeley.

[56] Sen, A., 1981, Poverty and famines. Oxford: Oxford University Press.

[57] Sen, A., 1988, 'The concept of development', In H. Chenery and T.N. Srinivasan (eds.), Handbook of Development Economics vol. 1, Amsterdam: North Holland.

[58] Skoufias, Emmanuel. (1996). "Intertemporal Substitution in Labor Supply: Micro Evidence from Rural India," Journal of Development Economics 51 (December), 217-237.

[59] Skoufias, E. 2002a. "Measuring Household Vulnerability to Risk: Estimates from Russia," Draft. International Food Policy Research Institute (IFPRI): Washington DC.

[60] Skoufias, E. 2002b. "Household Vulnerability to Risk: Estimates from the PROGRESA Evaluation Sample in Rural Mexico," Draft. International Food Policy Research Institute (IFPRI): Washington DC.

[61] Skoufias, E. and A.R. Quisumbing (2002) "Consumption Insurance and Vulnerability to Poverty: A Synthesis of the Evidence from Bangladesh, Ethiopia, Mali, Mexico and Russia," Draft. International Food Policy Research Institute (IFPRI): Washington DC.

[62] Stein, A., J. Behrman, R. Grajeda, J. Hoddinott, R. Martorell and M. Ramirez, 2003. Education and health over the life course in Guatemala. Research proposal, Emory University.

[63] Tesliuc, E. and K. Lindert. 2002. Vulnerability: A quantitative and qualitative assessment. Guatemala Poverty Assessment Program. Washington DC: The World Bank.

[64] Townsend, Robert (1994): "Risk and Insurance in Village India," Econometrica, Vol. 62, pp. 539-91.

[65] Townsend, R., 1995, "Consumption insurance: An evaluation of risk-bearing systems in low-income countries", Journal of Economic Perspectives, vol. 9 no. 3, pp. 83-102.

[66] Walker, T. and J. Ryan, 1990, Village and household economies in India's semi-arid tropics, Baltimore: Johns Hopkins University Press.

Figure 1
The multiple interlocking paths to poverty

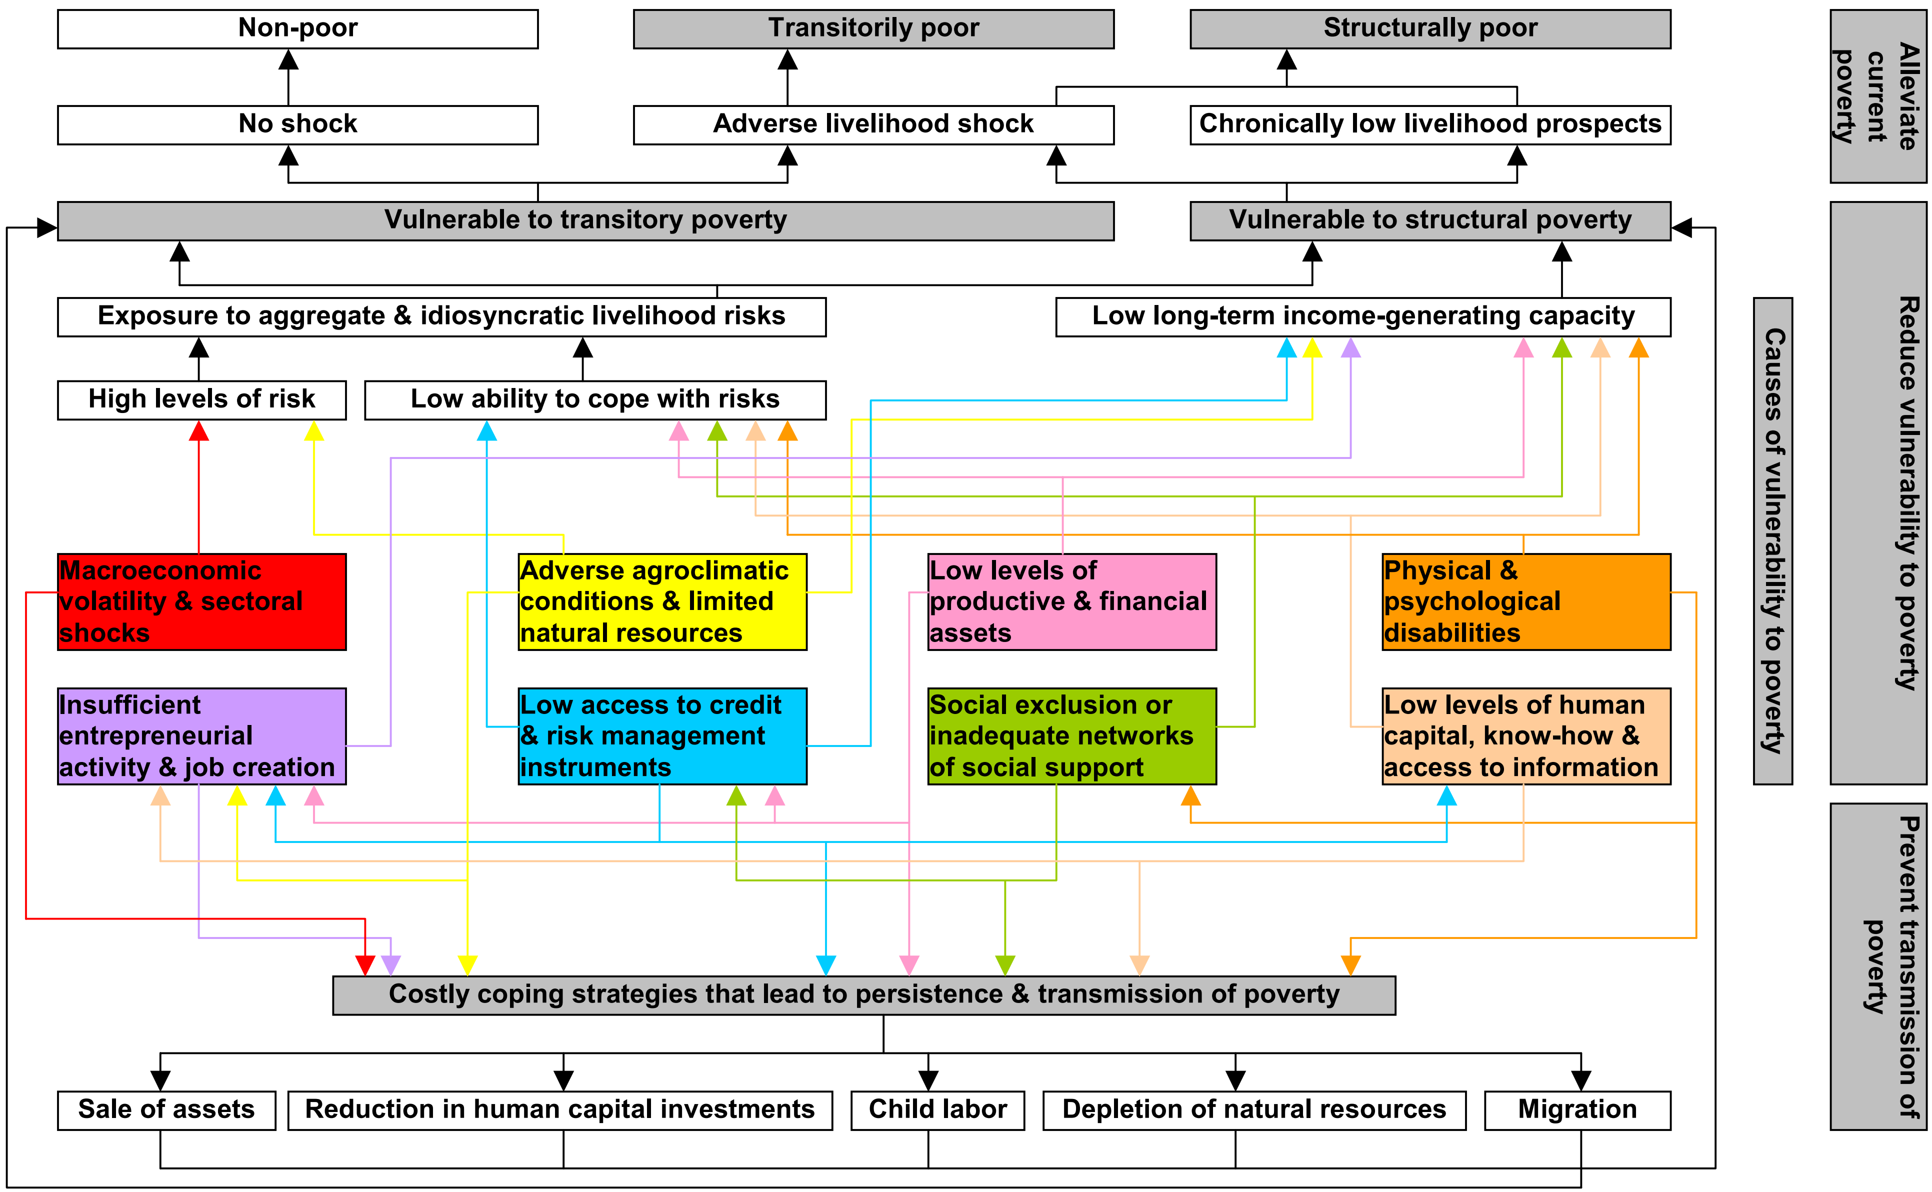

Figure 2
Estimated distribution of vulnerability in southern China in 1985
(Rural Household Survey, 1985-1990)

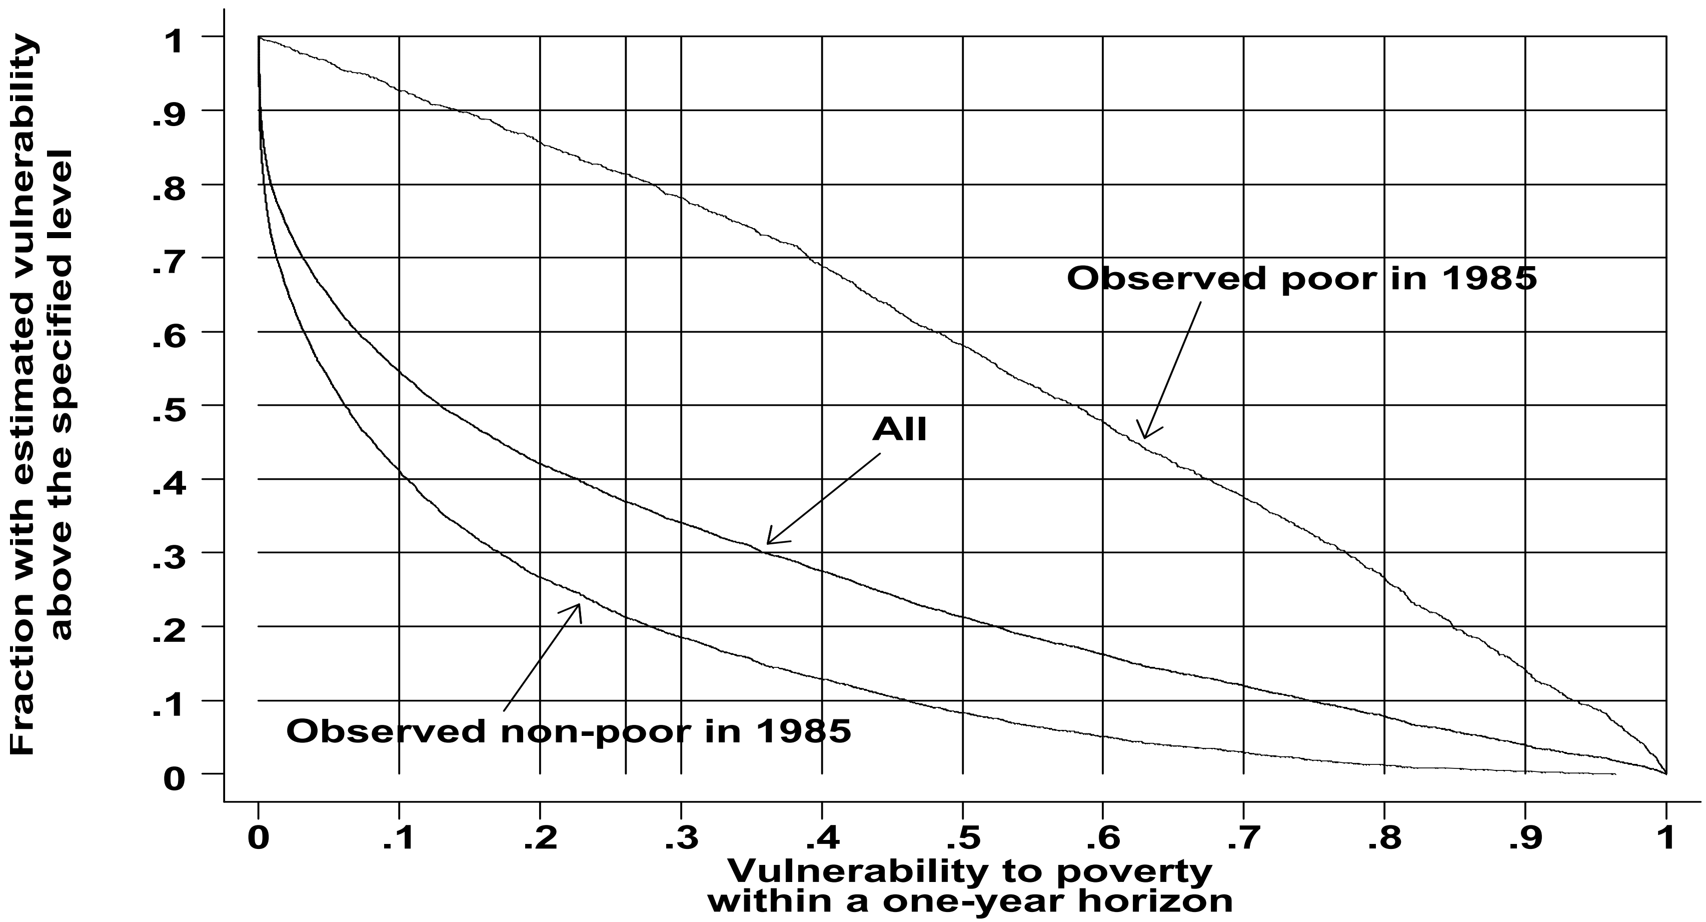

Figure 3

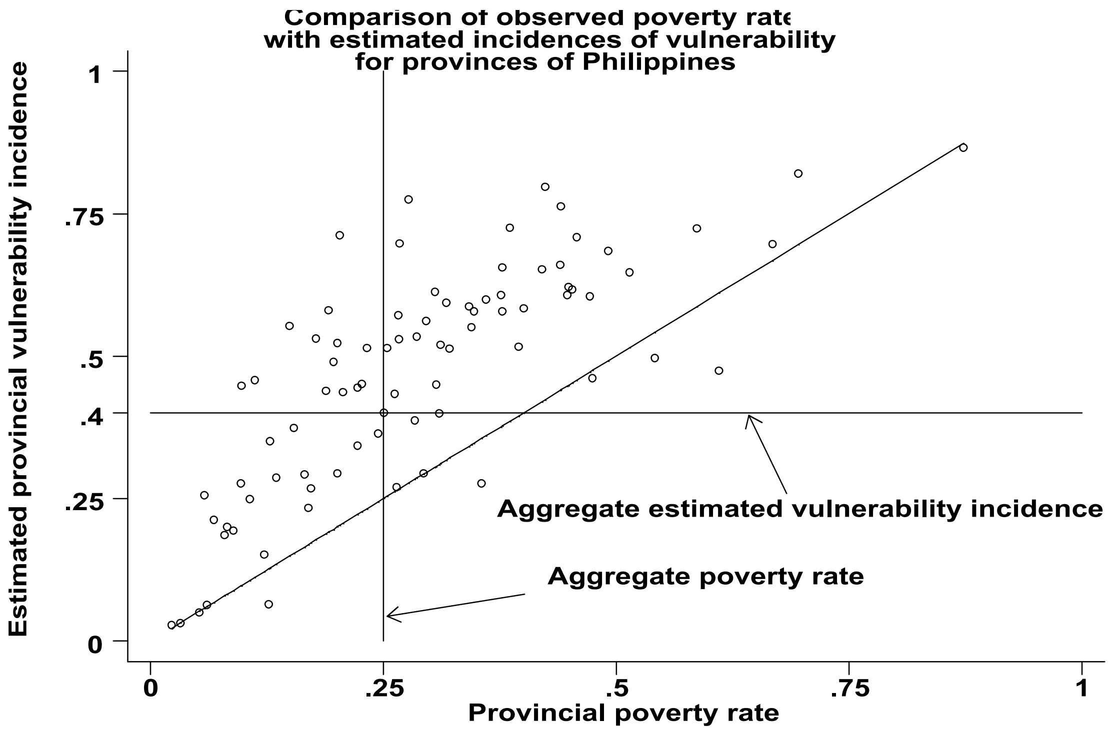

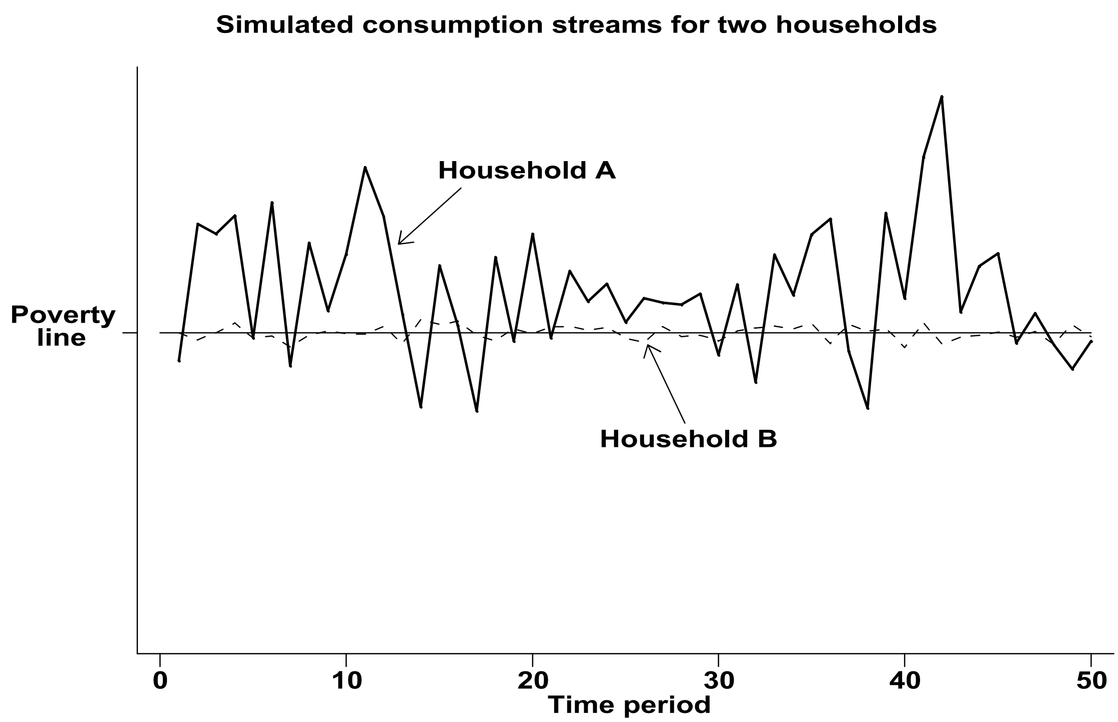

Figure 4

Figure 5
Do households with higher mean consumption also face greater consumption volatility?
Evidence from southern China (Rural Household Survey, 1985-1990)

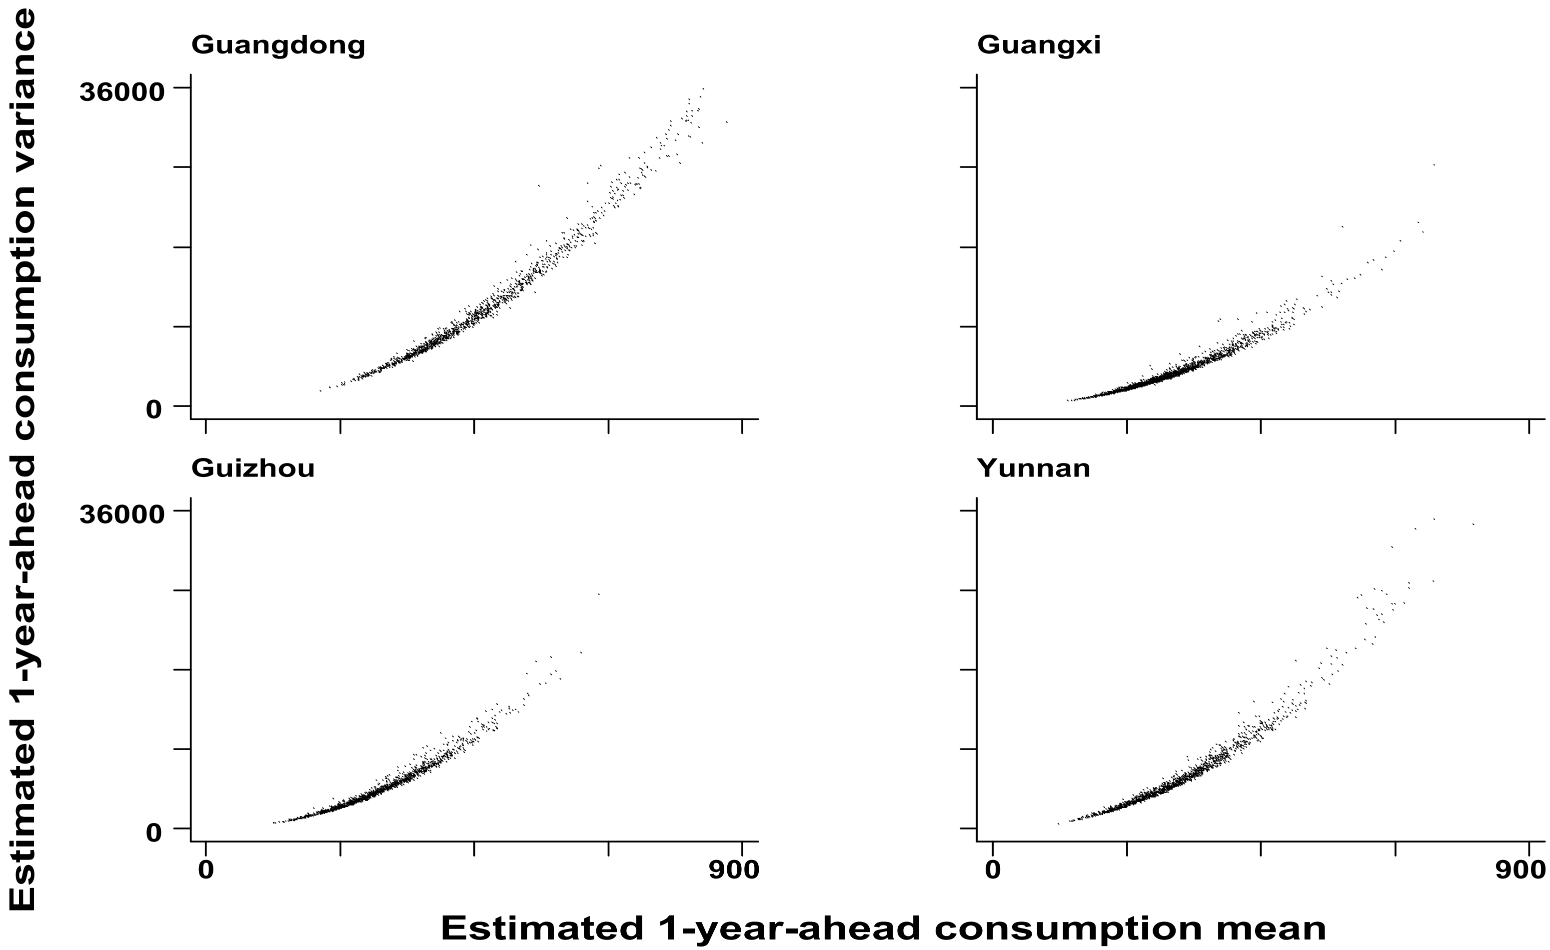

Figure 6
Geographic incidence of structural poverty and vulnerability in southern China
(Rural Household Survey, 1985-1990)

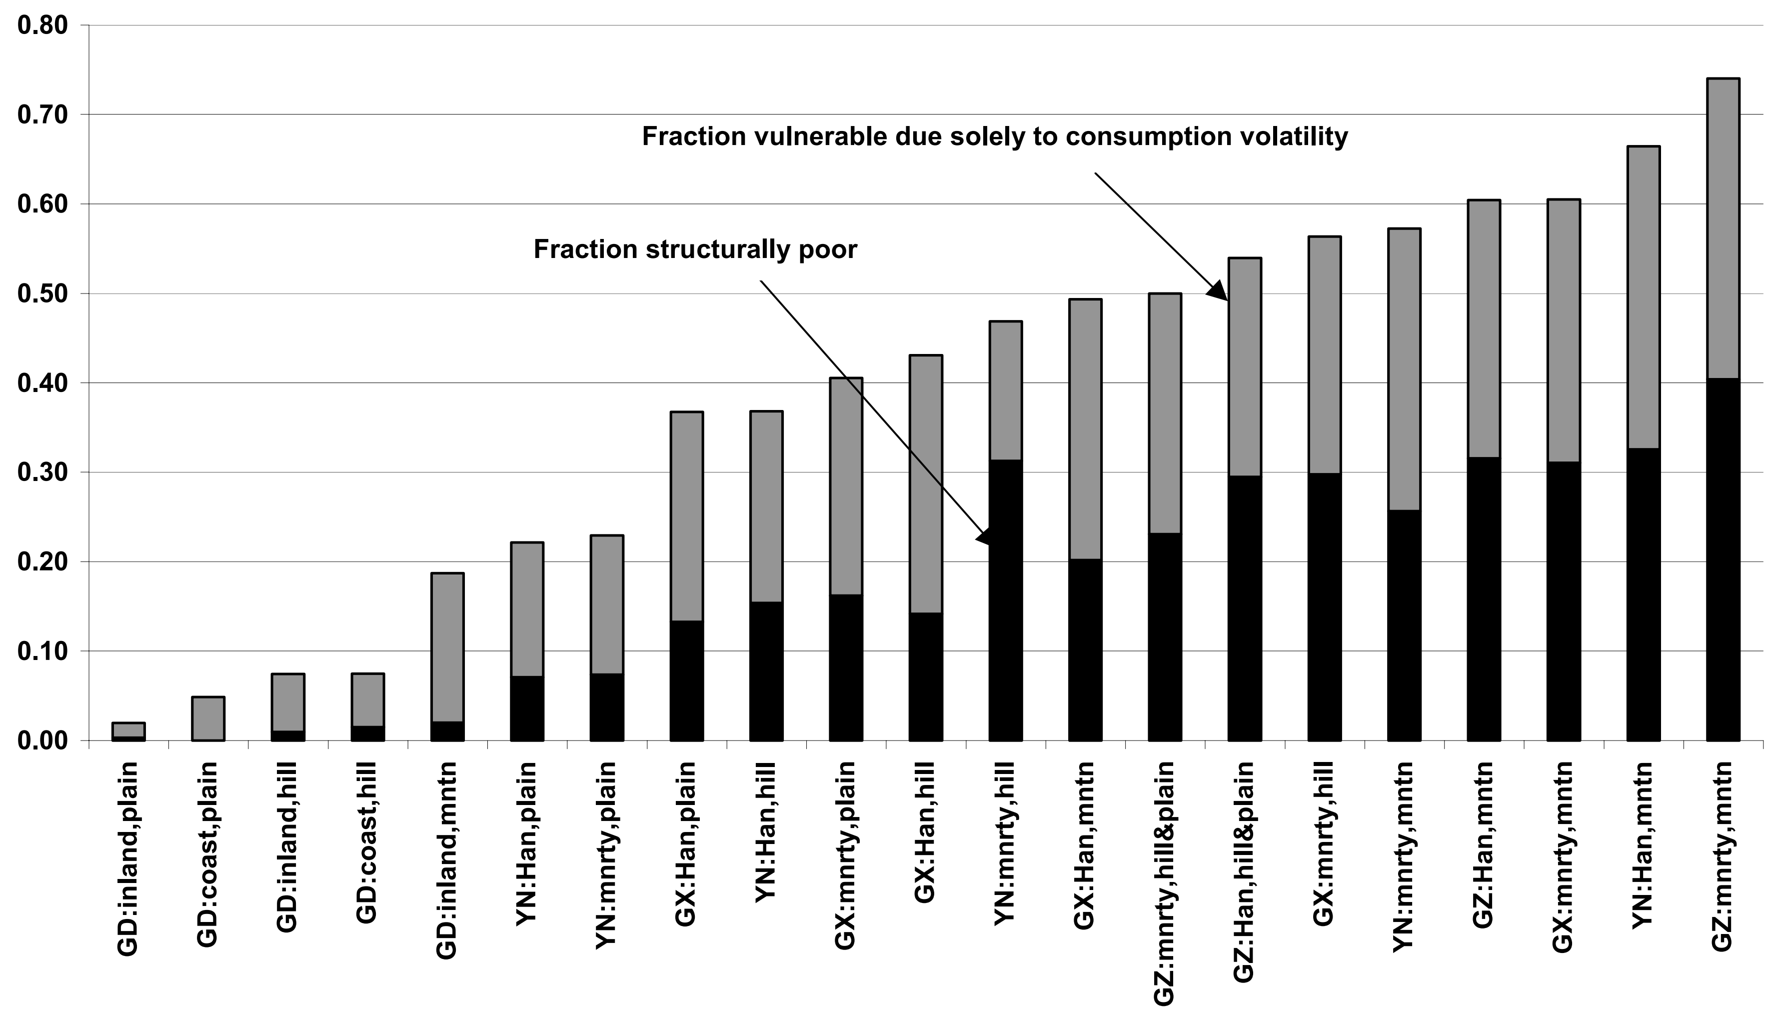

Indonesia (mini-SUSENAS, 1998-1999)

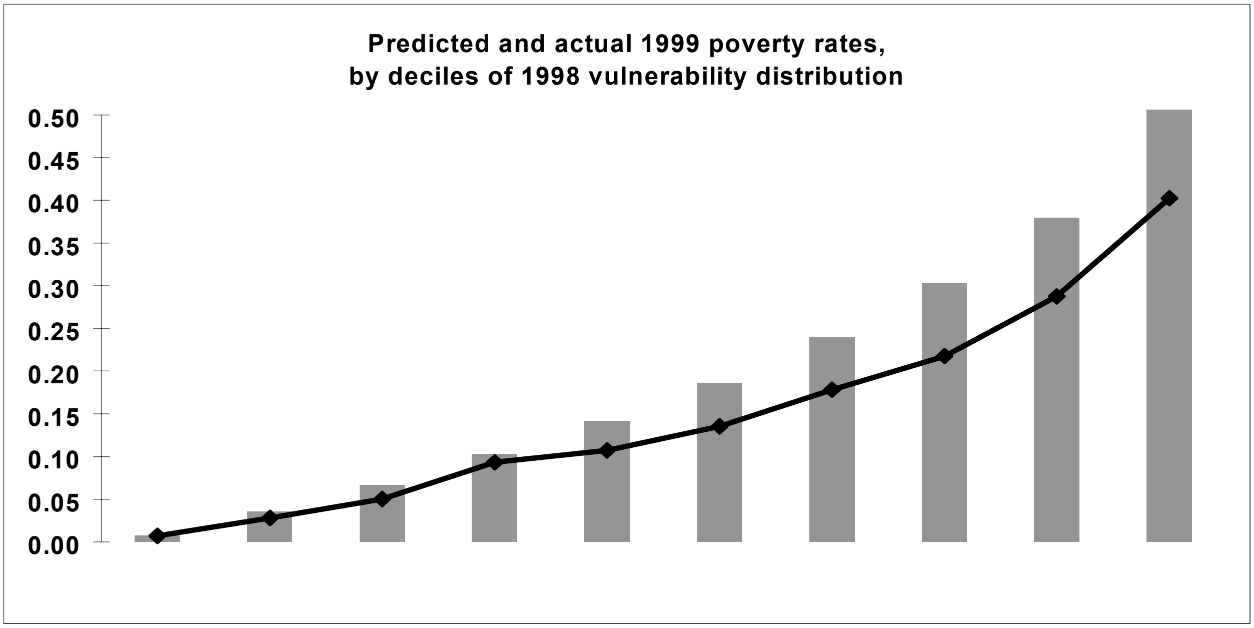

Philippines (FIES, 1997, APIS, 1998)

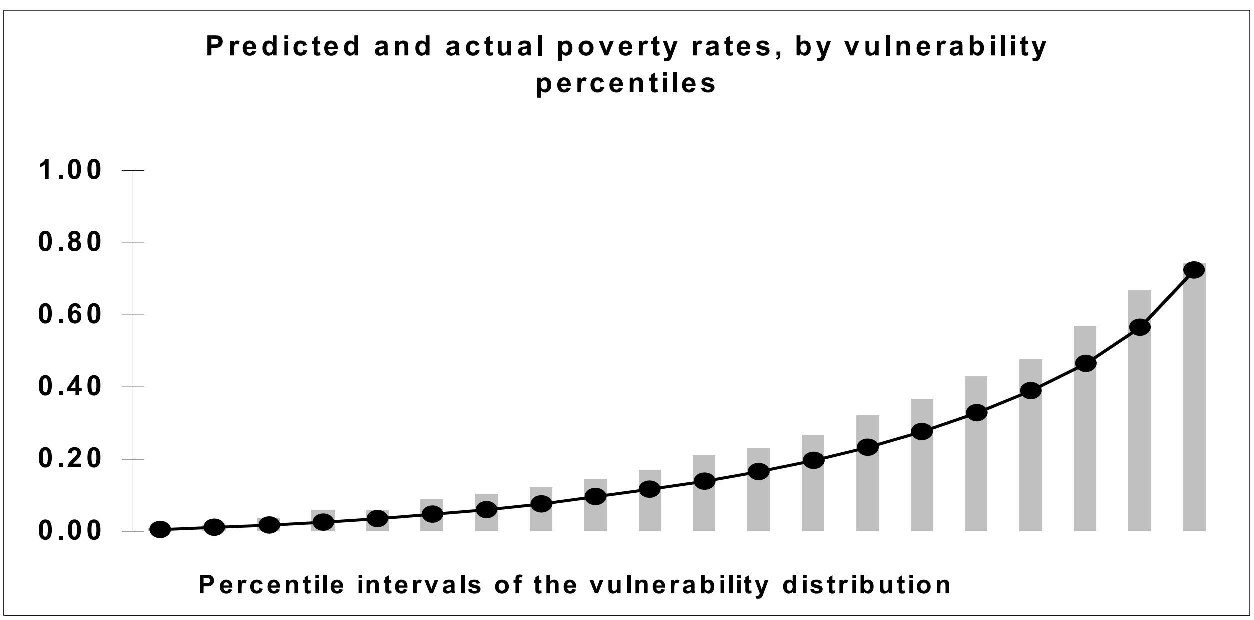

Figure 7

Table 1
Aggregate poverty and vulnerability in rural south and southwest China, 1985

	Overall	Amongst the non-poor	Amongst the poor	Amongst the non-vulnerable	Amongst the vulnerable	Amongst the highly vulnerable
Mean per-capita expenditure (yuan/year)	324.36	367.92	200.99	374.42	237.31	220.13
Fraction poor	0.26	0.00	1.00	0.08	0.57	0.69
Mean vulnerability	0.25	0.13	0.58	0.04	0.62	0.79
	0.04	0.02	0.12	0.00	0.11	0.16
	0.01	0.00	0.03	0.00	0.03	0.05
	-0.05	-0.06	0.01	-0.08	0.01	0.03
Fraction vulnerable	0.37	0.21	0.80	0.00	1.00	1.00
Fraction relatively vulnerable	0.22	0.09	0.60	0.00	0.61	1.00
Fraction highly vulnerable	0.21	0.08	0.57	0.00	0.57	0.93

Table 2
Poverty and vulnerability within different segments of the population, Indonesia, December 1998

	Population share	Share of poor	Share of vulnerable	Share of highly vulnerable	Fraction poor	Mean vulnerability	Fraction vulnerable	Vulnerability to poverty ratio	Fraction highly vulnerable
Overall					0.23	0.23	0.44	1.92	0.09
By location:									
Rural	0.61	0.80	0.82	0.91	0.30	0.30	0.60	1.99	0.13
Urban	0.39	0.20	0.18	0.09	0.12	0.13	0.20	1.66	0.02
Sumatra: urban	0.06	0.03	0.02	0.02	0.08	0.10	0.10	1.27	0.00
Jakarta: urban	0.05	0.00	0.00	0.00	0.02	0.03	0.02	1.00	0.00
West Java: urban	0.10	0.05	0.04	0.00	0.12	0.13	0.23	1.95	0.00
Central Java & Yogyakarta: urban	0.07	0.06	0.07	0.06	0.22	0.23	0.48	2.16	0.08
East Java & Bali: urban	0.06	0.02	0.01	0.00	0.08	0.11	0.12	1.52	0.00
Kalimantan & Sulawesi: urban	0.04	0.03	0.03	0.01	0.18	0.16	0.30	1.66	0.02
Rest of Indonesa: urban	0.01	0.00	0.00	0.00	0.08	0.13	0.21	2.59	0.00
Sumatra: rural	0.13	0.12	0.09	0.10	0.16	0.17	0.27	1.74	0.01
West Java: rural	0.12	0.16	0.16	0.20	0.31	0.30	0.62	1.98	0.16
Central Java & Yogyakarta: rural	0.12	0.17	0.20	0.17	0.34	0.35	0.78	2.30	0.14
East Java & Bali: rural	0.14	0.17	0.20	0.17	0.28	0.30	0.65	2.34	0.11
Kalimantan & Sulawesi: rural	0.09	0.12	0.13	0.20	0.31	0.34	0.69	2.22	0.21
Rest of Indonesa: rural	0.03	0.07	0.04	0.06	0.56	0.35	0.74	1.31	0.21
By education of household head									
No schooling	0.12	0.17	0.19	0.36	0.34	0.37	0.74	2.16	0.28
Primary	0.57	0.71	0.76	0.60	0.29	0.28	0.61	2.16	0.10
Junior	0.11	0.06	0.03	0.01	0.12	0.13	0.14	1.35	0.01
Secondary	0.16	0.05	0.02	0.02	0.07	0.08	0.03	0.55	0.01
More than secondary	0.05	0.01	0.00	0.01	0.01	0.03	0.00	0.00	0.00
By employment status of household head									
Unemployed/unpaid	0.13	0.11	0.12	0.12	0.19	0.22	0.43	2.23	0.08
Self-employed: no help	0.24	0.25	0.25	0.13	0.22	0.22	0.46	2.10	0.03
Self-employed: some help	0.31	0.37	0.39	0.43	0.27	0.28	0.57	2.11	0.12
Salaried (private & public)	0.32	0.27	0.24	0.31	0.19	0.19	0.33	1.78	0.08

Table 2 (continued)
Poverty and vulnerability within different segments of the population, Indonesia, December 1998

		Population share	Share of poor	Share of vulnerable	Share of highly vulnerable	Fraction poor	Mean vulnerability	Fraction vulnerable	Vulnerability to poverty ratio	Fraction highly vulnerable
By demographic categories										
Household head less than 60		0.86	0.86	0.85	0.83	0.22	0.22	0.45	2.00	0.08
Household head greater than 60		0.14	0.14	0.15	0.17	0.22	0.25	0.49	2.20	0.10
Female household head		0.08	0.08	0.09	0.13	0.22	0.24	0.46	2.07	0.13
Male household head		0.92	0.92	0.91	0.87	0.22	0.23	0.45	2.03	0.08
Household head not currently married		0.10	0.10	0.10	0.15	0.20	0.23	0.44	2.17	0.12
Married household head		0.90	0.90	0.90	0.85	0.22	0.23	0.45	2.02	0.08
Dependency ratio less than 0.25		0.79	0.81	0.79	0.81	0.23	0.23	0.45	1.99	0.08
Dependency ratio greater than 0.25		0.21	0.19	0.21	0.19	0.21	0.23	0.46	2.22	0.07
By community characteristics										
Transport facilities:	No	0.09	0.15	0.11	0.11	0.41	0.28	0.61	1.48	0.12
	Yes	0.91	0.85	0.89	0.89	0.21	0.22	0.44	2.13	0.08
Industry:	No	0.20	0.25	0.21	0.25	0.29	0.25	0.48	1.63	0.10
	Yes	0.80	0.75	0.79	0.75	0.22	0.23	0.44	2.02	0.07
Bank:	No	0.79	0.83	0.82	0.88	0.24	0.24	0.47	1.90	0.10
	Yes	0.21	0.17	0.18	0.12	0.18	0.20	0.37	2.04	0.05
Cooperative:	No	0.48	0.58	0.57	0.61	0.28	0.27	0.53	1.88	0.11
	Yes	0.52	0.42	0.43	0.39	0.18	0.20	0.37	1.99	0.07
Access to clean water	No	0.74	0.92	0.87	0.91	0.29	0.27	0.52	1.83	0.11
	Yes	0.26	0.08	0.13	0.09	0.07	0.14	0.22	3.01	0.03

Table 3
Exploring the proximate cause of vulnerability in rural south and southwest China, 1985

Population segment		Fraction of population in segment: Vulnerable to poverty	Vulnerable to transitory poverty	Vulnerable to structural poverty	Ratio of transitory vulnerable to structurally poor
Overall		0.37	0.16	0.21	1.31
Per-capita financial wealth quintiles	Lowest	0.52	0.19	0.33	1.74
	Middle	0.36	0.16	0.20	1.25
	Highest	0.08	0.05	0.04	0.80
Per-capita housing wealth quintiles	Lowest	0.62	0.21	0.41	1.95
	Middle	0.34	0.18	0.16	0.89
	Highest	0.06	0.04	0.02	0.50
Per-capita grain stocks quintiles	Lowest	0.46	0.14	0.32	2.29
	Middle	0.45	0.19	0.26	1.37
	Highest	0.25	0.13	0.12	0.92
Per-capita cultivated area quintiles	Lowest	0.37	0.15	0.22	1.47
	Middle	0.41	0.18	0.23	1.28
	Highest	0.30	0.14	0.16	1.14
Farming main livelihood?	Yes	0.38	0.16	0.22	1.38
	No	0.23	0.12	0.11	0.92
State employee in household?	No	0.38	0.16	0.22	1.38
	Yes	0.14	0.07	0.07	1.00
Household size	Large	0.50	0.19	0.31	1.63
	Small	0.30	0.14	0.16	1.14
Dependency ratio	High	0.51	0.19	0.32	1.68
	Low	0.31	0.14	0.17	1.21
Any illiterate adult household members?	Some	0.42	0.16	0.26	1.63
	None	0.28	0.15	0.14	0.93
Any adult household members with post-primary education?	None	0.45	0.17	0.28	1.65
	Some	0.30	0.15	0.16	1.07
Guangdong: coastal plains		0.05	0.02	0.02	1.00
Guangdong: coastal hilly areas		0.10	0.09	0.01	0.11
Guangdong: inland plains		0.03	0.02	0.01	0.50
Guangdong: inland hilly areas		0.13	0.09	0.03	0.33
Guangxi: plains		0.36	0.21	0.16	0.76
Guangxi: hilly areas		0.41	0.22	0.19	0.86
Guangxi: minority areas in plains		0.37	0.17	0.20	1.18
Guangxi: minority hilly areas		0.55	0.24	0.31	1.29
Guizhou: hilly areas		0.50	0.15	0.35	2.33
Guizhou: minority hilly areas		0.57	0.18	0.39	2.17
Yunnan: plains		0.14	0.07	0.07	1.00
Yunnan: hilly areas		0.46	0.18	0.27	1.50
Yunnan: minority areas in plains		0.18	0.11	0.07	0.64
Yunnan: minority hilly areas		0.45	0.19	0.26	1.37
Proportion of cultivated area in county irrigated	Low	0.46	0.18	0.29	1.61
	High	0.25	0.13	0.11	0.85
Road density in county	Low	0.40	0.16	0.24	1.50
	High	0.33	0.15	0.17	1.13
Density of medical personnel in county	Low	0.45	0.18	0.27	1.50
	High	0.27	0.13	0.14	1.08

Table 4
Estimating the contribution of risk to vulnerability in rural south and southwest China, 1985

Population segment		% of vulnerability in terms of expected poverty gap ratio attributable to risk		
		All households in segment	Those vulnerable to poverty	Those vulnerable to structural poverty
Overall				
Per-capita financial wealth quintiles	Lowest	71.7	48.8	19.0
	Middle	83.2	58.8	24.7
	Highest	96.5	74.9	41.7
Per-capita housing wealth quintiles	Lowest	65.3	45.9	18.3
	Middle	87.6	67.1	29.5
	Highest	98.2	81.9	32.7
Per-capita grain stocks quintiles	Lowest	69.2	43.0	18.0
	Middle	77.8	54.7	22.1
	Highest	89.6	66.6	28.9
Per-capita cultivated area quintiles	Lowest	79.9	53.5	22.1
	Middle	80.8	57.4	24.7
	Highest	83.9	57.4	21.9
Farming main livelihood?	Yes	89.2	63.6	26.3
	No	80.4	55.2	22.6
State employee in household?	No	80.9	55.7	22.8
	Yes	93.6	67.2	32.4
Household size	Large	74.1	51.3	20.8
	Small	85.2	59.5	24.8
Dependency ratio	High	72.8	50.4	21.2
	Low	85.0	59.4	24.1
Illiterate adult household members?	Some	76.9	50.9	20.0
	None	88.4	66.5	30.7
Any adults with post-primary education?	None	74.9	50.5	20.8
	Some	86.2	61.6	25.6
Guangdong: coastal plains		97.2	68.8	37.6
Guangdong: coastal hilly areas		99.0	92.1	36.6
Guangdong: inland plains		99.4	85.6	42.4
Guangdong: inland hilly areas		97.4	83.9	41.2
Guangxi: plains		90.0	75.0	41.9
Guangxi: hilly areas		86.3	68.3	31.0
Guangxi: minority areas in plains		85.9	66.3	38.7
Guangxi: minority hilly areas		76.5	58.3	25.9
Guizhou: hilly areas		66.9	40.4	15.8
Guizhou: minority hilly areas		62.7	40.5	13.0
Yunnan: plains		93.8	66.2	30.8
Yunnan: hilly areas		76.5	54.6	24.1
Yunnan: minority areas in plains		94.1	74.3	37.1
Yunnan: minority hilly areas		77.6	53.6	20.1
Proportion of cultivated area in county irrigated	Low	75.0	50.9	20.8
	High	90.1	67.3	29.2
Road density in county	Low	77.8	52.3	20.7
	High	85.0	60.4	26.0
Density of medical personnel in county	Low	76.7	53.1	21.6
	High	86.8	60.7	25.5