



Social Protection Discussion Paper Series

Evaluating Different Approaches to Estimating Vulnerability

by Ethan Ligon and Laura Schechter

June 2004

Social Protection Unit
Human Development Network
The World Bank

Social Protection Discussion Papers are not formal publications of the World Bank. They present preliminary and unpolished results of analysis that are circulated to encourage discussion and comment; citation and the use of such a paper should take account of its provisional character. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author(s) and should not be attributed in any manner to the World Bank, to its affiliated organizations or to members of its Board of Executive Directors or the countries they represent.

For free copies of this paper, please contact the Social Protection Advisory Service, The World Bank, 1818 H Street, N.W., Washington, D.C. 20433 USA. Telephone: (202) 458-5267, Fax: (202) 614-0471, E-mail: socialprotection@worldbank.org. Or visit the Social Protection website at <http://www.worldbank.org/sp>.

Evaluating Different Approaches to Estimating Vulnerability

Ethan Ligon*
and
Laura Schechter†

Abstract

A number of researchers have recently proposed a variety of different ‘vulnerability’ measures designed to capture the welfare consequences of risk for poor households, and also proposed a variety of different approaches to estimating these various measures of household vulnerability. However, it’s possible to ‘mix-and-match’ estimators and measures. Here we conduct Monte Carlo experiments designed to explore the performance of different estimators with different measures, under different assumptions regarding the underlying economic environment.

We find that when the environment is stationary and consumption expenditures are measured without error that the best estimator is one proposed by Chaudhuri (2001), regardless of what measure of vulnerability is employed. If the vulnerability measure is risk-sensitive, but consumption is measured with error, a simple estimator proposed by Ligon and Schechter (2003) generally performs best. However, when the distribution of consumption is non-stationary, a modification of an estimator proposed by Pritchett et al. (2000) performs best.

Future research should focus on combining the efficiency of the Chaudhuri estimator with the good properties of the Ligon-Schechter (in environments with measurement error) and Pritchett (in non-stationary environments) estimators. However, even with present technology estimating vulnerability is simple and much more informative and useful than are static poverty measures, provided one has at least two rounds of panel data.

*Associate Professor, Dept. of Agricultural and Resource Economics, University of California, Berkeley.

†PhD. Candidate, Dept. of Agricultural and Resource Economics, University of California, Berkeley.

1. INTRODUCTION

Economists have long used measures of poverty to summarize the well-being of less fortunate households in a population. Typically either income or consumption expenditures are measured over some relatively short period of time (e.g., a year), and these are regarded as a proxy for the material well-being of the household. Policies are often explicitly crafted to reduce these poverty measures.

At the same time, economists have long recognized that a household’s sense of well-being depends not just on its average income or expenditures, but on the risk it faces as well, particularly for households with fewer resources. To consider an extreme case, a household with very low expected consumption expenditures but with no chance of starving may well be poor, but it still might not wish to trade places with a household having a higher expected consumption but greater consumption risk. It seems desirable to have a measure of household welfare which takes into account both average expenditures and the risk households bear.

Although numerous commonalities have been identified in the above approaches to estimating vulnerability, the current state of the literature is a collection of papers, each paper with its own vulnerability measure, its own approach to estimation, and its own dataset—we are in what Hoddinott and Quisumbing (2003) call the “let a hundred flowers bloom” phase of research on vulnerability. Though this has been useful for exploring the set of possible approaches to measuring vulnerability, it may be time to begin weeding the flower bed. The difficulty is that no one has yet systematically tried to compare approaches to estimating vulnerability on the same datasets, or tried using different estimators with different vulnerability measures.

Recent research on household ‘vulnerability’ has led to an increased appreciation of the welfare costs of risk. The key idea is simply that risk averse households will have lower levels of expected utility *ex ante* when those same households face greater variation in future consumption.

Hoddinott and Quisumbing (2003) describe several alternative definitions and approaches to estimating “vulnerability” using large scale household-level datasets. The survey provided by Hoddinott and Quisumbing draws on several different recent papers which seek to measure vulnerability (Chaudhuri et al., 2001; Ligon and Schechter, 2003; Pritchett et al., 2000). For each of these papers there are two steps involved in estimating vulnerability. First, one must estimate the distribution of future consumption expenditures for every household; second, one must construct a statistic (e.g., EP_α) from this estimated distribution, meant to capture the reduction in household welfare due to risk in household consumption expenditures. Although different authors use different methods at each of these steps, within limits it’s possible to mix and match. In particular, there’s no reason one couldn’t use, e.g., Chaudhuri’s methods for estimating the distribution of expenditures, but with Ligon-Schechter’s vulnerability measure. Another way of putting this: different approaches to *defining* vulnerability need not imply differences in methods for *estimation*, since the object which must be estimated is simply the distribution of consumption.

Unfortunately, because there’s considerable variety in the definition of vulnerability, in the estimators employed, and in the datasets used in example applications, the comparisons among approaches detailed in Hoddinott and Quisumbing (2003) are often comparisons between apples and oranges; one can’t easily evaluate the practical merits of various definitions

of vulnerability, the value of different estimators, or the power of different tests. Here, we address this problem by combining various definitions of vulnerability described in Hoddinott and Quisumbing (2003) with various of the different estimators also described therein. The various combinations of definitions and estimators will then be used to measure the vulnerability of households both in artificially generated data (where the actual *ex ante* welfare of households is known), and in two ‘real’ datasets, one from Viet Nam belonging to the LSMS family of household surveys, and another using data from the Bulgarian Household Budget Survey. Our aim is to provide the practitioner with—not a catalog of possible approaches (this is accomplished by Hoddinott and Quisumbing, 2003)—but with specific advice regarding what the *best* approach is to measuring vulnerability in any given dataset.

2. MEASURES OF VULNERABILITY

As noted above, one can think of there being two steps to estimating vulnerability. The first has to do with characterizing consumption expenditures; the second in summarizing the welfare consequences of this variation. We’ll call the first step ‘predicting consumption’; the second ‘measuring vulnerability’.

In this section we describe several measures of vulnerability. These are drawn from the existing literature, but for some we’ll make minor changes (mostly pertaining to changes in units and notation) so as to facilitate comparisons across measures.

2.1. An enumeration of vulnerability measures. A number of papers have, in recent years, sought to define and measure something called “vulnerability.” The earliest efforts (Amin et al., 1999; Glewwe and Hall, 1998) attempted to measure the *exposure* households have to shocks observed by the econometrician—the basic idea is, using household panel data, to regress some measure of changes in consumption expenditures on possible sources of shocks, and to look for a response; if variation in the shock can account for a significant proportion of the variation in expenditures, the household is said to be vulnerable to the shock. Later approaches (Chaudhuri et al., 2001; Ligon and Schechter, 2003; Pritchett et al., 2000) instead attempt to estimate *expected* welfare, variously defined. Households with estimated expected welfare below some threshold are said to be vulnerable. It’s worth noting that when panel data are available welfare realizations can be substituted for expenditures in a regression of possible sources of shocks, so that after describing some particular contributions to the ‘vulnerability-as-exposure’ literature in Section 2.2 we’ll focus entirely on the later ‘vulnerability-as-reduction-in-well-being’ contributions and extensions.

2.2. Exposure to observed risks. This approach to the measurement of vulnerability (Amin et al. (1999); Dercon and Krishnan (2000); Glewwe and Hall (1998)) focuses on the response of households’ consumption expenditures to various observable shocks, such as drought or idiosyncratic fluctuations in income. If household consumption expenditures covary with income shocks, then one may infer that a risk-averse household lacks the means to smooth or insure away these shocks to its expenditures. Note that this measure of vulnerability does not depend directly on a household’s level of consumption. Neither does it depend directly on the risk a household bears—a household with large variation in consumption which does not stem from variation in observables would have a low measured vulnerability. Of course, this latter feature could be regarded as a virtue; the method seems useful for identifying particular sources of risk, which may then be an appropriate focus of policy.

Amin, Rai, and Topa (1999). Amin, Rai, and Topa (1999) use panel data from Bangladesh to identify households whose consumption vary with income, after controlling for household fixed effects and aggregate variation in mean consumption. Accordingly, they estimate (in differences) a simple prediction equation with the idiosyncratic vector of variables x_t^i simply equal to household income,

$$(1) \quad E(c^i | \bar{x}, x^i) = \alpha^i + \eta_t + x_t^i \beta^i.$$

Note the use of a household-specific coefficient β^i , which the authors call the estimated “vulnerability” of household i . This is meant to capture the reduction in welfare associated with the additional risk a household bears if its consumption co-moves with risky (or time-varying) income. Thus, if two households have precisely the same consumption realizations in every state, but the second household has a more variable income stream, then Amin, Rai, and Topa (1999) would regard the second household as the less vulnerable.

Glewwe and Hall (1998). Glewwe and Hall (1998) and Glewwe and Hall (1995) measure something they call “vulnerability” in Peru, but in contrast to Amin, Rai, and Topa (1999), are particularly interested in the response of households’ consumptions to *aggregate* shocks. In particular, Glewwe and Hall (1995) estimate (in differences) a prediction equation with the idiosyncratic vector of time-invariant household characteristics x^i , but with time-varying coefficients,

$$(2) \quad E(c_t^i | \bar{x}, x^i) = \alpha^i + \eta_t + x^i \beta_t.$$

Contrast this with the prediction equation (1) of Amin, Rai, and Topa (1999). The key difference is that Glewwe and Hall focus on household level consumption responses to aggregate shocks (which they identify with β_t), while Amin, Rai, and Topa (1999) focus on household response to idiosyncratic shocks (which they identify with β^i).

Dercon and Krishnan (2000). Dercon and Krishnan (2000) take an approach similar to that of Amin, Rai, and Topa (1999) and Glewwe and Hall (1998), but estimate households’ exposure to both idiosyncratic and village level shocks. The authors work with an estimating equation of the form

$$c_t^i = \alpha^i + \gamma x_t^i + \beta x_t + e_t^i$$

where x_t^i contains aggregate, time-varying variables such as wages and prices, and where x_t contains observed idiosyncratic shocks faced by individuals and households (e.g. animal disease, personal illness). Thus β is a measure of households’ exposure to aggregate shocks, while γ is a measure of how vulnerable households are to assorted idiosyncratic shocks.

2.3. Expected poverty. Recall that Foster et al. (1984) define a family of decomposable poverty measures P_α .² A second approach to the measurement of vulnerability has been to adapt these standard poverty measures to a non-deterministic setting by estimating the *expected* value of P_α . Although methods for estimating these expected measures vary considerably, several papers share this approach to defining a measure of vulnerability. We divide these papers into two groups, depending on their favored choice of the parameter α , which governs the property of the poverty measure. Several authors have chosen to work with the headcount measure of poverty ($\alpha = 0$); others have chosen to work with the “squared

²Let z denote a “poverty line” level of consumption expenditures. For any household i , let $P_\alpha^i = \frac{\max\{z - c^i, 0\}^\alpha}{z^\alpha}$ be the poverty of household i ; then for any set of households I the poverty of the set is simply $\sum_{i \in I} P_\alpha^i$.

poverty gap” ($\alpha = 2$). Each of these alternatives has differing strengths and weaknesses, and so we discuss each in turn below.

Expected headcount ($\alpha = 0$). A number of authors (Chaudhuri (2001); Chaudhuri et al. (2001); Christiaensen and Boisvert (2000); Pritchett et al. (2000)) use a measure of household vulnerability which is simply the expected headcount measure of poverty. This measure is simple and comprehensive—it varies with households’ wealth and aggregate and idiosyncratic sources of risk. However, the measure suffers some of the same shortcomings of the headcount measure of poverty. Consider a household whose present consumption is somewhat above the poverty line, but which receives a very bad shock with small probability. Consistent with this story, we can imagine that expected consumption for the household might lie slightly below the poverty line, despite the fact that the probability of the household falling into poverty was small. If the household is risk averse, with von Neumann-Morgenstern utility, then it would *prefer* to consume its expected consumption with certainty. Thus, if offered the choice (say via the offer of actuarially fair insurance) the household would choose the consumption stream which would cause it to be ‘vulnerable’ according to this definition.

Expected poverty gap ($\alpha = 1$). We’re not aware of any author who has adopted the expected poverty gap as their preferred measure of vulnerability, but this choice seems at least as defensible as the adoption of EP_0 , and we include this measure in the experiments we conduct later in this paper. Like the expected headcount measure, the measured vulnerability of poor households will be smaller when these households face large risks, and won’t vary at all if these same poor households face only small risks (too small, that is, for there to be any chance for the household to vault over the poverty line).

Expected squared poverty gap ($\alpha = 2$). Finally, we turn our attention to Ravallion (1988). Ravallion doesn’t use the term “vulnerability,” preferring to think about “expected poverty.” Thus, Ravallion’s measure is intended to be the expected value of a concave function of household consumption.³

P_α still seems ill-suited to representing household risk attitudes. The first problem arises because these measures assign no weight to the welfare of households whose consumption is (perhaps only momentarily) greater than the poverty line. The second has to do with the nature of the risk preferences implicit in this measure—Foster, Greer and Thorbecke’s poverty measure P_α implies an absolute risk aversion of $(\alpha-1)z/(z-c)$, where z is the poverty line, c is household consumption, and α is a non-negative parameter. However, even if $\alpha > 1$ (so that households are risk-averse), this implies that households have increasing absolute risk aversion, which is sharply at odds with existing research on the risk preferences of poor households. Even more unreasonably, both absolute and relative risk aversion approach infinity as consumption approaches the poverty line (from below).

2.4. Expected Utility. Quite recently, a number of authors have proposed using some variant of “expected utility” as the basis of a vulnerability measure. The basic logic is simply that, since the class of von Neumann-Morgenstern utility functions was originally designed to capture risk preferences, and since these functions are widely estimated and used in actual applications, it seems sensible to adopt these functions to measure the welfare loss associated with risk. In particular, members of the HARA (“Hyperbolic Absolute Risk

³In practice, Ravallion uses time series variation in consumption to estimate dynamics in realized poverty, rather than attempting to capture the welfare consequences of risk; accordingly, though we consider the measure he employs we do not use his estimator.

Aversion”) family of utility functions have found favor in applications throughout economics, and so these seem suitable for use in evaluating vulnerability. The family can be described by the function

$$(3) \quad U(c) = \frac{(c - z)^{1-\gamma} - 1}{1 - \gamma},$$

$\gamma \geq 0$, $c \geq z$. Thus, the HARA bears a close resemblance to the P_α family, differing in that the parameter z is interpreted as a lower bound of consumption for all households, rather than as the lower bound only for wealthy households, and differing also in allowable values for the curvature parameter γ —in this case, the welfare costs of risk are decreasing as γ falls, while the opposite obtains for the P_α family.

Ligon and Schechter (2003). For their measure of vulnerability, Ligon and Schechter (2003) adopt a simple transformation of (3). The idea behind the normalization is simple: if there is no inequality and no uncertainty, then there can be no vulnerability.

To operationalize this idea, Ligon and Schechter (2003) let \bar{c} denote per capita expenditures. Then the vulnerability of household i is defined by

$$V^i = U^i(\bar{c}) - EU^i(c^i).$$

Ligon and Schechter choose units for consumption expenditures so that $\bar{c} = 1$; this normalization solves the problem of units, just as normalization by poverty line expenditures solves the problem for the P_α measure. Vulnerability for the population is computed by summing household vulnerability across all households. Note, then, that if every household consumes \bar{c} for sure, then there is *no vulnerability*; no household bears any risk, and there is no inequality (and hence no relative poverty).

Ligon and Schechter (2003) extend this idea by decomposing their measure of vulnerability into distinct (but not necessarily orthogonal) measures of poverty and risk, as follows: Taking expectations of an increasing, concave function of consumption expenditures has the effect of making vulnerability depend not only on the mean of a household’s consumption, but also on variation in consumption. To better understand the balance between poverty and risk in this measure of vulnerability, note that one can decompose the measure into distinct components reflecting poverty and risk, respectively:

$$(4) \quad V^i = [U^i(\bar{c}) - U^i(Ec^i)] + [U^i(Ec^i) - EU^i(c^i)].$$

Note that the first bracketed term, which measures poverty, involves no random variables—it is simply the difference between a concave function evaluated at the “poverty line” and at household i ’s expected consumption expenditure. The concavity of U^i implies that as Ec^i approaches the poverty line, an additional unit of expected consumption has diminishing marginal value in reducing poverty. For any $\{U^i\}$ in the HARA family it’s easy to show that this poverty measure satisfies all the axiomatic requirements enumerated in Foster et al. (1984).

The second term of (4), which measures the risk faced by household i , is consistent with the ordinal measures of risk proposed by Rothschild and Stiglitz (1970). Further, this risk measure can usefully be further decomposed into two distinct measures of risk, one aggregate, the other idiosyncratic. Let $E(c^i|x)$ denote the expected value of consumption, c^i , conditional on knowledge of a vector of aggregate variables x . Then we can decompose the risk household i faces into a term expressing the *aggregate risk* the household faces, and a term expressing

the *idiosyncratic risk* the household faces. Total vulnerability for household i is then given by

$$\begin{aligned} V^i = & [U^i(\bar{c}) - U^i(\mathbf{E}c^i)] && \text{(Poverty)} \\ & + [U^i(\mathbf{E}c^i) - EU^i(\mathbf{E}(c^i|x))] && \text{(Aggregate risk)} \\ & + [EU^i(\mathbf{E}(c^i|x)) - EU^i(c^i)]. && \text{(Idiosyncratic risk)} \end{aligned}$$

This decomposition is not peculiar to utility based measures of vulnerability; one can similarly decompose expected values of P_α .

Calvo and Dercon (2003). Calvo and Dercon (2003) follow Ligon and Schechter (2003) in adopting a utility-based measure of vulnerability, but claim that it’s desirable to add something to the measure which plays a role more like the poverty line employed by the P_α measure. Giving some credit to Clark et al. (1981) and Watts (1968) for the basic idea, they propose a measure of vulnerability

$$\frac{\mathbf{E}[\min\{z, c^i\}^{1-\gamma}] - z^{1-\gamma}}{z^{1-\gamma}}.$$

It’s easy to verify that this measure differs from the HARA version of the Ligon-Schechter measure only in that there’s now a “kink” in the function due to the min operator. Calvo and Dercon argue that this kink is desirable, giving as evidence an example in a household with expected consumption at or below the poverty line z . This household has a chance p of having its consumption exceed the poverty line by some amount y next period; Calvo and Dercon argue that our assessment of the vulnerability of this household ought not to depend on the magnitude y , despite the fact that clearly the household would prefer that y be larger rather than smaller.

2.5. Comparing Vulnerability Measures. In this subsection we use artificial data generated in accordance with Jalan and Ravallion (1999) (see Section 4.1), and compute household level vulnerability using each of the static measures described above, so as to facilitate comparisons.

We first calculate the correlations between various vulnerability measures and also the actual expected utility for the the households in our artificial sample. The Jalan-Ravallion data-generating process we exploit doesn’t hinge on the specification of a ‘true’ utility function, so we explore different possible specifications, including a collection of HARA utility functions with a range of possible coefficients of relative risk aversion ($\gamma \in \{1, 2, 3\}$).

In Table 1 we display correlations between the various measures of vulnerability described above, where the distribution of consumption expenditures is consistent with a stationary environment described by Jalan and Ravallion (1999), described in detail in Section 4.1, *infra*. The Ligon-Schechter measure is abbreviated “LS”, and the Calvo-Dercon “CD”; the various expected Foster-Greer-Thorbecke poverty measures are indicated by EP_α for $\alpha \in \{0, 1, 2\}$. Also shown in Table 1 are correlations between (minus) various HARA expected utilities, again given the distribution of consumptions described in Section 4.1, and the correlations between these HARA utilities and the vulnerability measures. Consensus estimates of HARA utility functions from household panel data suggest values for the parameter γ in the range of 1–3.

Not surprisingly, all of these measures and expected utilities are highly correlated. However, there’s considerable range in the degree of correlation. At the high end, in this instance

	LS	CD	EP_0	EP_1	EP_2	$-EU_1(c)$	$-EU_2(c)$	$-EU_3(c)$
LS	1.0000	0.9928	0.7852	0.9493	0.9827	0.9146	1.0000	0.9524
CD	1.0000	1.0000	0.7134	0.9203	0.9740	0.8610	0.9928	0.9731
EP_0	0.9980	0.9978	1.0000	0.9108	0.8175	0.9390	0.7852	0.5832
EP_1	0.9998	0.9998	0.9989	1.0000	0.9806	0.9544	0.9493	0.8150
EP_2	0.9999	1.0000	0.9983	0.9999	1.0000	0.9163	0.9827	0.8965
$-EU_1(c)$	0.9999	0.9997	0.9988	0.9999	0.9998	1.0000	0.9146	0.7675
$-EU_2(c)$	1.0000	1.0000	0.9980	0.9998	0.9999	0.9999	1.0000	0.9524
$-EU_3(c)$	0.9999	1.0000	0.9971	0.9995	0.9998	0.9995	0.9999	1.0000

TABLE 1. Correlations Between Vulnerabilities and Expected Utility. Correlations reported above the diagonal are simple Pearson correlation coefficients, while those below are Spearman rank correlation coefficients. Utility functions (denoted by $EU_\gamma(c)$) are HARA functions with $z = -0.1$.

the Ligon-Schechter measure happens to be constructed so as to be perfectly correlated with $-EU_2(c)$; at the low end the Pearson correlation of $-EU_3(c)$ and EP_0 is only 0.5832. Indeed, if any measure can be said “not to belong” it would be EP_0 —it would be more highly correlated with the various utility measures if households were taken to be risk-seeking, since while all these functions are decreasing in consumption expenditures only the function P_0 is concave. The Ligon-Schechter and Calvo-Dercon measures are very highly correlated, as one would expect—truncating the distribution of consumption, as the Calvo-Dercon measure implicitly does, only makes the welfare function slightly more convex at high levels of consumption, where it’s already nearly flat (for $\gamma = 2$, the HARA utility function rapidly asymptotes to zero). These two measures are also highly correlated with EP_2 —this is a little more surprising, until one recalls that these are the three measures we consider which are convex, so that increases in risk result in increases in vulnerability. We’ll call these the “risk-sensitive” vulnerability measures—we’ll see in our experiments below that this shared characteristic seems to be key to determining behavior. Among the expected poverty measures, the highest correlation is between EP_2 and $-EU_2(c)$, with a Pearson correlation coefficient of 0.98.

Correlation coefficients below the diagonal are Spearman rank correlation coefficients. These give some indication of the extent to which these different measures produce different *orderings* of vulnerability across households. For the data-generating process employed here cross-sectional variation in expected levels of consumption dominates time-series variation in any particular household’s consumption expenditures. All of these measures and utilities produce exactly the same ordering of non-random consumptions, and as a consequence the Spearman statistics all exceed 0.99.

3. ESTIMATING CONSUMPTION DISTRIBUTIONS

When measuring poverty, one can rely on surveys which measure household income or expenditures. While fraught with its own difficulties, these observed data are all one needs to compute static poverty measures. To actually measure the vulnerability of a household (or population), one cannot rely only on observable data—the whole point of adopting a notion of vulnerability is that household welfare is presumed to depend not just on what consumption expenditures are actually realized, but on what consumption expenditures *might* be.

Complicating matters, the notion of vulnerability is inherently forward-looking, since there’s only uncertainty about *future* consumption expenditures. Of course, we only have data on past realizations. Accordingly, the estimation problem we face involves using past realizations of consumption expenditures to estimate the probability of possible future consumption outcomes. When these probabilities remain the same across time, the environment is said to be *stationary*, and it’s possible to estimate the probability of different outcomes in a typical period, and regard this as vulnerability. However, in environments where there’s consumption growth over time or where institutions exist which permit households to save current windfalls against future shortfalls, the probabilities associated with different consumption realizations will vary over time. In such an environment it’s no longer possible to characterize the probability distribution of consumption expenditures in a typical period—change across periods may be the only constant. An alternative strategy in such an environment is to assume that the probabilities of *changes* in consumption expenditures is fixed across periods—such an environment is said to be *difference stationary*. Note that any stationary time series process is difference stationary, so that estimators which only assume difference stationarity will be correctly specified in a broader range of circumstances. However, there’s no free lunch: by only using information from changes over time, these estimators will ignore potentially useful information, and so will be less precise than estimators which exploit this stationarity.

Returning to the matter at hand, observing actual consumption expenditures cannot, by itself, provide enough information for us to compute vulnerability. For this, we need (an estimate of) the probability distribution of what consumption expenditures might have been, and to compute this estimate we need to make some identifying assumptions to allow us to map past outcomes into predictions about the future. To date, three different basic strategies have been adopted in the literature; we’ll consider each in turn, and discuss the particulars of their implementation by various authors.

3.1. Stationary Time Series. The basic idea behind the estimation strategy here is to assume that, for any particular household, the probability distribution of consumption in one period is identical to the probability distribution of consumption in any other period. As a consequence, if one observes consumption expenditures for each household for two or more periods, one assumes that each of these observations was just as likely to have happened to the same household in any other period.

Ligon and Schechter (2003) adopt this strategy to estimate vulnerability in Bulgaria. Using panel data with T periods of data on N households, they simply use the T observations on consumption expenditures for each household as evidence of ‘what might have been’ in the initial period. Accordingly, their main measure of vulnerability (“TS” for short) can be computed by

$$(5) \quad \hat{V}^i = U(\bar{c}) - \frac{1}{T} \sum_{t=1}^T U(c_t^i),$$

where c_t^i is observed consumption expenditures for household i in period t . The chief virtue of this approach is its transparency and simplicity—since U is assumed to be a known function, this is a calculation which could be trivially performed using a spreadsheet or simply a calculator. If the assumption of stationarity is correct, then as T grows large then a law of large numbers guarantees that the time series average computed here will converge to

the expected value we're trying to estimate. A second great virtue is that one doesn't need to assume anything about the variation of consumption expenditures *across* households—the method is appropriate when there's great heterogeneity across households, even when this heterogeneity is unobserved by the analyst.

There are two chief drawbacks to the Ligon-Schechter time-series estimator. First of all, the assumption of stationarity is crucial. If, for example, consumption expenditures are expected to grow over time, then observed consumption in one period will not have the same distribution as expenditures in a different period, and the time series average will not converge to the expected value. Second, the method obviously requires panel data, which is usually both expensive and time-consuming to collect.

Ligon and Schechter avoid a third problem, devising an estimator which we'll call the "lower bound" (LB) estimator. If measurement error contaminates the data on consumption expenditures, this will bias estimates of vulnerability, and if some (say wealthier) households have noisier data on consumption expenditures, this will lead to estimates which overstate their vulnerability relative to other (say poorer) households. Ligon and Schechter take an approach which has something of the flavor of instrumental variables to deal with classical measurement error. Let c_t^i denote the consumption expenditures of household i at time t , and let x_t^i denote a vector of idiosyncratic variables (which may help to predict consumption). Let $\tilde{c}_t^i = c_t^i e^{\epsilon_t^i}$, with $\{\epsilon_t^i\}$ a measurement error process having the property that $E(\epsilon_t^i | x_t^i) = E(\epsilon_t^i \log c_t^i) = 0$. In a sort of "first stage" they use a prediction equation of the form⁴

$$(6) \quad \log \tilde{c}_t^i = \alpha^i + \eta_t + \beta' x_t^i + u_t^i,$$

where the $\{\alpha^i\}$ are household fixed effects, the $\{\eta_t\}$ are time effects (constrained to sum to zero) which capture aggregate shocks to consumption, the x_t^i are deviations from the household mean, and the $\{u_t^i\}$ are residuals. Using restricted least squares to estimate this yields parameter estimates $(\hat{\alpha}^i, \hat{\eta}_t, \hat{\beta}, \hat{u}_t^i)$ which can be used to construct conditional expectations. Then using the decomposition trick employed in (5), one can write

$$\begin{aligned} \hat{V}^i &= [U(1) - U(\hat{E}c^i)] && \text{(Poverty)} \\ &+ \left[U(\hat{E}c^i) - \frac{1}{T} \sum_{t=1}^T U(\hat{E}(c_t^i | x_t)) \right] && \text{(Aggregate risk)} \\ &+ \left[\frac{1}{T} \sum_{t=1}^T U(\hat{E}(c_t^i | x_t)) - \frac{1}{T} \sum U(\hat{E}(c_t^i | x_t, x_t^i)) \right] && \text{(Idiosyncratic risk)} \\ &+ \left[\frac{1}{T} \sum_{t=1}^T U(\hat{E}(c_t^i | x_t, x_t^i)) - \frac{1}{T} \sum_{t=1}^T U(c_t^i) \right]. && \text{(Unexplained risk \& measurement error)} \end{aligned}$$

Thus, the contribution of measurement error to the estimate of household i 's vulnerability is confined chiefly to the final line of this expression. Unfortunately, the residual term u_t^i will typically depend not only on measurement error, but also on variables not observed by the analyst which nonetheless influence consumption. Accordingly, the LB estimate (the sum of

⁴The prediction equation actually employed in Ligon and Schechter (2003) uses levels of expenditures as the dependent variable. In principle this can cause problems, since with a linear prediction equation one may sometimes predict negative levels of consumption. Here we address this possible problem by taking logarithms.

the first three lines of this decomposition) may be regarded as a lower bound on household i 's vulnerability, while the quantity (5) may be regarded as an upper bound.

Chaudhuri (2001) also relies on an assumption of stationary time series to construct his estimator, and assume further that shocks to consumption are independent normal in their distribution. He makes a serious effort to accommodate heterogeneity, by permitting the variance of consumption changes to depend on observable fixed characteristics of the sample households, estimating \hat{c}^i and $\hat{\sigma}_\epsilon(x^i)$, which can then be employed to calculate household level vulnerability

$$\hat{V}^i = \int W(\hat{c}^i e^{\hat{\sigma}_\epsilon(x^i)\epsilon}) d\Phi(\epsilon),$$

where the function W is the welfare measure employed (e.g., $W(c) = \frac{1}{1-\gamma}[(\bar{c} - z)^{1-\gamma} - (c - z)^{1-\gamma}]$ for Ligon-Schechter, P_α for Pritchett or Chaudhuri). Thus, for each household Chaudhuri (2001) needs to construct two estimates of parameters, \hat{c}^i and $\hat{\sigma}_\epsilon(x^i)$. For \hat{c}^i he simply uses the time series geometric mean of c_t^i . The variance parameter he needs to compute is simply the time series variance of $\log c_t^i$; he constructs estimates of this by regressing the square of changes in $\log c_t^i$ on a vector of household characteristics x^i , and then dividing by two.^{5,6}

3.2. Difference Stationary Time Series. Pritchett et al. (2000) also rely on panel data and a stationarity assumption, but in their case the stationarity requirement is weaker—they require that *changes* in consumption expenditures be stationary and mean zero over time. This is consistent with a world in which consumption follows a random walk, as it might be if households had access to credit markets, but cannot accommodate a world with deterministic growth. As with Chaudhuri et al. (2001), Pritchett et al. (2000) rely heavily on the assumption that changes in consumption are distributed normally. This approach has two chief virtues. The first is that it's quite simple; one way to implement it is to use ordinary least squares to predict consumption expenditures in a first-stage regression, much as with the Ligon-Schechter estimator (given the assumptions adopted by the authors, OLS is the maximum likelihood estimator). The second is that this estimator is very efficient, permitting one to construct relatively precise estimates from relatively short panels. The chief defects of the method are related to these virtues—the distributional assumptions adopted are very strong, and permit no unobserved heterogeneity either across households or periods. The normality assumption adopted by these authors implies that consumption may sometimes be negative. We deal with this problem in two different ways. First, by setting any negative predicted consumption to zero (we call this the PSS₀ estimator). Second, by assuming that consumption is distributed log-normal (we call this the PSS₁ estimator). Neither of these estimators can be easily adapted to accommodate measurement error. As with Chaudhuri et al. (2001), the PSS estimators can accommodate heteroskedasticity across households, by permitting the variance of consumption changes to depend on observable fixed characteristics of the sample households.⁷ To calculate household level vulnerability, one first constructs an

⁵For our experiments we take x^i to include the level of first period income and its square and a constant.

⁶One problem with this estimator is that it may yield predictions of the standard deviation which are imaginary. To prevent this, we convert all negative predicted variances to zero.

⁷This also requires some kind of stationarity of the conditional distribution of consumption changes across households. Pritchett and coauthors address this by observing that if the analyst possessed a long enough

estimate of the variance of consumption changes $\hat{\sigma}_\epsilon(x_i)$; then one can calculate

$$\hat{V}^i = \int W(c_t^i + \hat{\sigma}_\epsilon(x^i)\epsilon)d\Phi(\epsilon),$$

Although the assumption of difference stationarity is weaker than the stationarity assumption adopted by Chaudhuri et al. (2001), attempts to use the method with panel datasets having more than two periods raises a difficult question—since consumption may now be distributed differently in different periods, just which period should we use for calculating vulnerability? If we have data on periods $t = 1, \dots, T$, should we be attempting to predict consumption (and vulnerability) for period T , which would be consistent with what Pritchett et al. (2000) do using a panel of two periods? What, then, should the household be assumed to know? As t progresses from 1 to T , the household learns more about what the distribution of consumption will be in T . If we compute vulnerability from the perspective of period 1, this suggests using

$$\hat{V}_{1,T}^i \equiv \int W(c_{i,1} + T\hat{\sigma}_\epsilon\epsilon)d\Phi(\epsilon),$$

while if computed from the perspective of period $T - 1$, we ought instead to employ

$$\hat{V}_{T-1,T}^i \equiv \int W(c_{i,T-1} + \hat{\sigma}_\epsilon\epsilon)d\Phi(\epsilon),$$

Either of these choices are defensible, but obviously these will yield very different estimates of vulnerability. Here we'll (somewhat arbitrarily) adopt a third choice; namely, to estimate vulnerability in period 2 from the perspective of period 1, yielding

$$\hat{V}_{1,2}^i \equiv \int W(c_{i,1} + \hat{\sigma}_\epsilon\epsilon)d\Phi(\epsilon).$$

3.3. Stationary Cross-Section. In addition to presenting a vulnerability estimator to be used with panel data, Chaudhuri (2001) suggests another option which can be used with a cross-sectional data set. Both estimators calculate vulnerability as

$$\hat{V}^i = \int W(\hat{c}^i e^{\hat{\sigma}_\epsilon(x^i)\epsilon})d\Phi(\epsilon).$$

If one has panel data, Chaudhuri suggests using the time series geometric mean of c_t^i to estimate \hat{c}^i . If, on the other hand, one only has cross-sectional data, he proposes using FGLS to estimate the equation $\log c^i = \beta x^i + \epsilon_i$ and then using $\hat{\beta}x^i$ as the estimate of \hat{c}^i .⁸ To estimate the variance of $\log c^i$ Chaudhuri regresses the square of the error in the above equation (ϵ_i^2) on a vector of household characteristics. We'll refer to this cross-sectional estimator as the "XS" estimator.

panel then one could rely solely on time series variation to estimate idiosyncratic variability, but make no effort to implement such an estimator.

⁸Note that as a consequence this estimator doesn't directly use any information on the actual level of consumption expenditures for household in estimating vulnerability; instead of using a household's consumption as the base level of consumption, the cross-sectional estimator uses only the part which is predicted by observed characteristics. For many data sets a large portion of consumption may not be predicted by observed characteristics—in such cases we would expect this estimator to perform less well than simply using observed household consumption, without making any effort to characterize the risk faced by the household.

4. THE DATA GENERATING PROCESSES

In this section we describe a variety of different data-generating processes (DGPs) which we'll use to conduct experiments later in this report.

4.1. Stationary Jalan-Ravallion. We have generated panel data on consumption, income, and “wealth group” using a data generating process (DGP) suggested by Jalan and Ravallion (1999). Using data from rural China, the authors posit that ability to smooth consumption may vary by wealth quantile. Their estimating equation is such that changes in household level consumption are affected by changes in household level income and changes in household size. They divide the data into five wealth quantiles and allow the coefficients to differ across groups.

Jalan and Ravallion (1999) estimate a regression of the form

$$(7) \quad \Delta \log c_t^i = \Delta \eta_t^{w_i} + \beta^{w_i'} x_t^i + \epsilon_t^i,$$

where i indexes the household, t the time period, and w_i indicates the wealth quantile of household i —thus, the parameters β^{w_i} are allowed to differ by wealth quantile. For the idiosyncratic variables x_t^i Jalan and Ravallion employ changes in log income and household size; for our purposes we'll drop the variation in household size.

For the stationary data generating process described here, we use Jalan and Ravallion's estimates, and generate data by first arbitrarily assigning each of N households to wealth quintiles. We then generate data on actual income y_t^i for household i at time t ; income is presumed to be the product of two factors, one observed by the analyst (θ_t^i), and the other (ξ_t^i) not. Each of these two factors is drawn from a log-normal distribution, with means and variances chosen to match quantile-specific moments reported by Jalan and Ravallion (1999).⁹ In addition, we also allow for measurement error in income. Thus, actual income will be made up of an observed portion θ and an unobserved portion ξ , so that

$$\log y_t^i = \theta_t^i + \xi_t^i,$$

while the income observed by the analyst is denoted \tilde{y}_t^i , which is equal to actual income times a measurement error e^{ν^i} . As with the components of actual income, the mean and variance of the normal distribution from which ν_t^i is drawn is allowed to vary by wealth quantile.

Once we've generated data on income, we draw data on period 0 consumption c_0^i for each household i . As with the components of income, this quantity is drawn from a log-normal distribution, with mean and variance set equal to wealth quantile-specific means of

⁹We first generate data on income in the initial period for households in the five wealth groups based on data given in Jalan and Ravallion (1999). The highest wealth group has a mean income of 912 Yuan with a standard deviation of 1136, and we have assumed the income of each wealth level to be lognormally distributed. The second wealth group has mean income of 479 and standard deviation of 414. The third wealth group has mean income of 380 and standard deviation of 323, the second to last wealth group has a mean income of 306 and a standard deviation of 254, while the lowest wealth group has a mean income of 227 and a standard deviation of 228. We generate income for the five wealth groups over five years, but then randomly reorder the years of data our dgp will use.

log consumption reported by Jalan and Ravallion (1999).¹⁰ We then compute consumption using

$$(8) \quad \log c_t^i = \log c_0^i + \eta_t^{w_i} + \beta^{w_i} \log y_t^i + u_t^i,$$

exploiting wealth-quantile specific values of β^w estimated by Jalan and Ravallion (2000).¹¹

Note that taking differences of (8) yields Jalan-Ravallion’s estimating equation (7)—despite the fact that Jalan-Ravallion’s estimating equation is specified in differences, it’s consistent with a stationary data-generating process (indeed, the null hypothesis of full insurance tested by Jalan-Ravallion implies that there exists a set of time effects such that log consumption minus these constants will be a stationary process).

4.2. Non-Stationary Jalan-Ravallion. Several of the estimators described above rely on time-series stationarity, so it behooves us to generate some non-stationary data to see how sensitive these estimators are to this sort of misspecification. The way in which we generate income and initial levels of consumption is identical—the only difference is that for our non-stationary Jalan-Ravallion data generating process we calculate consumption iteratively via

$$(9) \quad \log c_t^i = \log c_{t-1}^i + \eta_t^{w_i} + \beta^{w_i} \log y_t^i + u_t^i.$$

This data-generating process is also consistent with the Jalan-Ravallion estimator, but of course in this case the distribution of consumption at time $t + j$ depends on the sequence of shocks realized in earlier periods, and hence is non-stationary.

4.3. Vietnamese Household Survey. The two artificial data-generating processes described above are useful for conducting Monte Carlo experiments, but of course assume away much of the complexity of real-world data. Accordingly, we also work with a two-period panel dataset drawn from the 1993 and 1998 rounds of the Vietnamese Household Survey, a household survey belonging to the LSMS family. We’re able to link 2191 households across both periods.

4.4. Bulgarian Household Survey. Like many expenditure surveys, the Bulgarian Household Budget Survey (described in documents available from the Bulgarian National Statistical Institute, www.nsi.bg) collects data on household expenditures at monthly frequencies. We’ve used these monthly data over the course of 12 months for 2287 households in 1994-95.

5. EXPERIMENTS

By conducting Monte Carlo experiments we can compare the bias and efficiency of different of these estimators using datasets of various sizes. For example, the LB estimator of Ligon and Schechter (2003) employs a relatively weak restriction on conditional moments, while the *PSS* estimators use a much stronger parametric distributional assumption. We’ll use

¹⁰The highest wealth group has mean consumption expenditures of 606 Yuan in the baseline period, with a standard deviation of 763. The second wealth group has mean consumption of 365 and standard deviation of 260. The third wealth group has mean consumption of 322 and standard deviation of 224, the second to last wealth group has a mean consumption of 264 and a standard deviation of 190, while the lowest wealth group has a mean consumption of 209 and a standard deviation of 189.

¹¹These are (0.12, 0.13, 0.18, 0.24, 0.41), where the first coefficient is that on the wealthiest group and the last coefficient is that on the poorest group.

the results of our Monte Carlo experiments to ascertain which of these estimators is more precise, and to measure the bias of each estimator in small samples.

The experiments of Section 6 will focus on the effects of various kinds of mis-specification on various of the estimators we consider, whether these have to do with improperly specified distributions, failures of stationarity, etc. In contrast, the experiments of this section will not focus on various possible mis-specifications, but rather on relative performance when the data generating process is the stationary Jalan-Ravallion process of Section 4.1. Nonetheless, it's important to note that all of the estimators suffer from some degree of mis-specification relative to this DGP; in particular, while the Jalan-Ravallion data generating process allows different distributions of shocks for each of five wealth quintiles, none of the estimators we use will make use of this heterogeneity, nor of data on wealth. We permit this sort of mis-specification on the theory that when working with real world data we can never expect our estimator to accomodate all the complexity of the real-world data-generating process, so that in practice our efforts at estimation will always be mis-specified to some greater or lesser degree. By erring in making our estimator too simple, we make an effort to have our baseline experiments reflect to some degree the kinds of mis-specification which may be expected in practice.




FIGURE 1. Mean Squared Error in Estimating the Ligon-Schechter Measure of Vulnerability as a Function of \sqrt{N} .

5.1. Accuracy of Vulnerability Estimates as a Function of N . Here we consider each of several measures of vulnerability, and attempt to estimate household vulnerability using each of the estimators described in Section 3, using samples of varying size. The focus of our interest in this section will be on how the mean squared error (MSE) of our estimates changes as the number of households in the sample increases, holding the length of the panel fixed. We use this as the primary criterion for identifying a ‘preferred’ estimator for each of




FIGURE 2. Spearman Rank Correlation of Estimates of the Ligon-Schechter Measure of Vulnerability with Actual Vulnerability as a Function of \sqrt{N} .

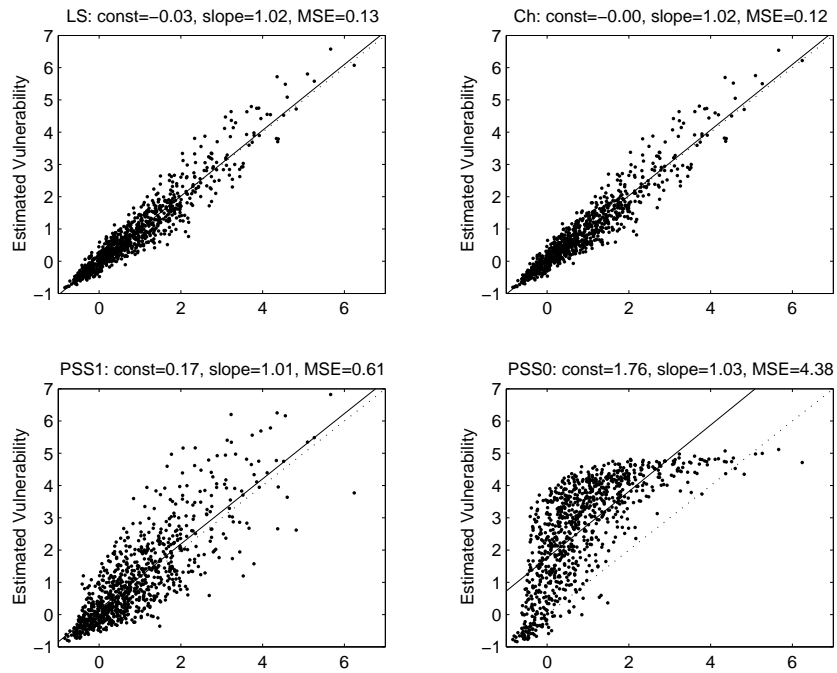


FIGURE 3. Estimates of Ligon-Schechter Vulnerability

the various measures of vulnerability we’ll describe. Note that for this experiment we hold the number of periods fixed, with $T = 5$.

We first turn our attention to Figure 1, which displays the logarithm of the MSE of each of the estimators we’ve described as the number of households (N) in the sample increases. Note that the MSE should not be expected to converge to zero as N increases; rather we expect it to converge to some positive constant.

Each estimator is marked with different symbols. For each household the Ligon-Schechter TS estimator is simply the time series average of a HARA utility function evaluated at realized observed consumption over the five periods;¹² the log of the MSE for this estimator is marked with circles. From the figure, one can see that this estimator converges to a mean square error of approximately 0.12 when N exceeds 2500, suggesting that the sample sizes characteristic of a typical Living Standards Measurement Survey are probably adequate for characterizing the vulnerability of the population as a whole.

The estimator of Chaudhuri (2001), marked with exes, performs just slightly better than does the TS estimator when estimating the Ligon-Schechter vulnerability measure. It appears to exhibit very similar behavior as the sample size increases. This may seem somewhat surprising, given the apparently greater complexity of the estimator he describes. His efforts to model heterogeneity in the variance of shocks pay off in this example, giving this estimator the smallest mean square error at every sample size—since the distributional assumption of log-normal shocks is satisfied for this particular data generating process Chaudhuri’s estimator is very nearly the maximum likelihood estimator—one could only do better by observing the actual wealth quintiles which determine the distribution of income and the variance of shocks, rather than simply income itself.

As mentioned above, we implement the Pritchett et al. (2000) estimator in two different ways. The first (PSS_0) involves working with levels of consumption, assuming with the original authors that consumption expenditures are normally distributed, but setting any negative predicted consumption to zero. The second way (PSS_1) is to assume that consumption is distributed *log*-normal. The estimator implemented in levels is indicated by a line with asterisks; the estimator implemented in logs is indicated with plusses. Note that each of these estimators performs abominably, with mean square errors orders of magnitude larger than the Chaudhuri or TS estimators. The cross-sectional “XS” estimator (marked with diamonds) is similarly poor.

Figure 2 is somewhat similar to Figure 1, but instead of using mean squared error as a criterion, it shows the Spearman rank correlation coefficient for each of the estimators with actual Ligon-Schechter vulnerability as a function of the square root of sample size; accordingly, in this figure, higher values are better. The ordering of estimators is just as it was using the MSE criterion, except that the performance of “XS” estimator is worse relative to the PSS_0 using this criterion. The best estimators (Chaudhuri and TS) have reassuringly high rank correlation coefficients—if one were to use either of these, this suggests that at least one would get the orderings of households’ vulnerabilities right in the overwhelming majority of cases. The two nonstationary estimators have remarkably similar performance, which is remarkably worse than the performance of the “TS” and Chaudhuri estimators. The “XS” estimator does worse still; as we’ll see later, if one has access only to cross-sectional

¹²A note on parameters—we set $z = -0.1$, and $\gamma = 2$ for these experiments. Further recall that we’ve chosen units so that the mean consumption of the average household is equal to one.

data it would be much better to simply use data on observed consumption expenditures, i.e., to use $W(c_t^i)$ as an estimator for V^i , and not to run any regressions at all.

The impression that the two Pritchett *et al* estimators perform poorly is confirmed by an examination of Figure 3. Here the estimator using levels is labelled PSS_0 in the lower right panel. We have plotted estimated Ligon-Schechter vulnerability for each of 900 households drawn from the Jalan-Ravallion world versus the actual vulnerability of these households. The level estimator PSS_0 is making an incorrect distributional assumption, and this is reflected in this panel. The estimator has a severe upward bias for the most well-to-do households, while it underestimates the vulnerability of those in the direst of straits.

The Pritchett *et al* estimator assuming a log-normal distribution (labelled PSS_1) makes the correct distributional assumption, and as a consequence there's no appreciable bias for this estimator (the dotted line in these figures is a 45 degree line, while the solid line is an OLS fit; when these two lines coincide there is no bias). However, by adopting the relatively weak hypothesis of difference stationarity the efficiency of this estimator is much lower than that of the TS and Chaudhuri estimators which adopt a stronger stationarity assumption.




FIGURE 4. Mean Squared Error in Estimating the EP_0 Measure of Vulnerability as a Function of \sqrt{N} .

We next turn our attention to the estimation of EP_0 . A glance at Figures 4 and 5 shows that, save for the very smallest sample sizes, the ordering of estimators is as it was for the estimation of the Ligon-Schechter measure. The Chaudhuri estimator performs better at all sample sizes, while the two Pritchett *et al* estimators perform much more poorly. However, relative to the earlier case the TS estimator (which simply averages outcomes for each household) performs considerably worse than does the Chaudhuri estimator. The reason for this relative drop in performance becomes clear in the top left panel of Figure 6. Since




FIGURE 5. Spearman Rank Correlation of Estimates of EP_0 (with Actual) as a Function of \sqrt{N} .

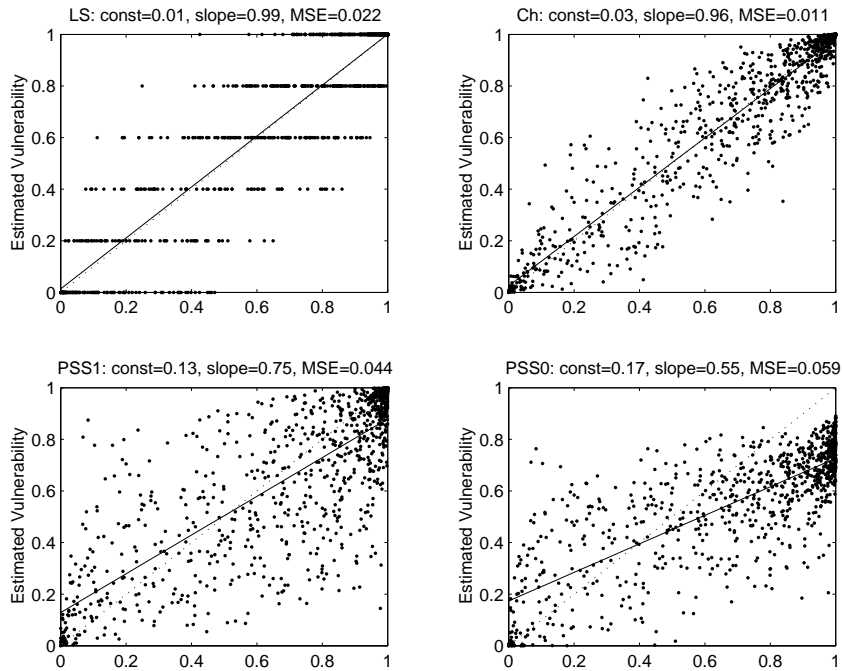


FIGURE 6. Estimates of EP_0 Vulnerability

in this case the TS estimator is simply averaging five different binary variables (indicating whether the household was above or below the poverty line), it's throwing out a great deal of information. Still, as one would expect the TS estimator is unbiased, just imprecise. In this particular experiment the Chaudhuri estimator displays a small bias.

In contrast to the nearly unbiased TS and Chaudhuri estimators, the two Pritchett *et al* estimators both display a quite substantial bias, overestimating the vulnerability of well-to-do households, and underestimating the vulnerability of the the most vulnerability. This underestimation of the vulnerability of the most vulnerability is particularly egregious for the estimator which assumes normally distributed consumption shocks, but both estimators do quite poorly both in terms of bias and imprecision. This is perhaps more surprising than elsewhere, since this is the measure the Pritchett *et al* estimator was designed to estimate.




FIGURE 7. Mean Squared Error in Estimating the EP_1 Measure of Vulnerability as a Function of \sqrt{N} .

Examining the estimation of EP_1 , we see in Figures 7 and 8 that once again there's a very clear division between the outcomes for the two estimators which assume stationarity (TS and Chaudhuri) and the two which do not (the two PSS estimators). Because the scale for EP_1 is the same as for EP_0 , it's possible to compare across the two poverty measures. The two PSS estimators do even worse in this setting than they did in estimating EP_0 . Once again (Figure 9) the PSS₀ estimator displays a very substantial bias, and makes light of the plight of the most miserable. The PSS₁ estimator's bias is much smaller than earlier, but is by far the least precise of the estimators, giving it an MSE in the same range as that of PSS₀. The TS and Chaudhuri estimators once again show a rather similar performance, but once again the TS estimator is edged out by the Chaudhuri estimator, which has a slightly smaller bias.




FIGURE 8. Spearman Rank Correlation of Estimates of EP_1 (with Actual) as a Function of \sqrt{N} .

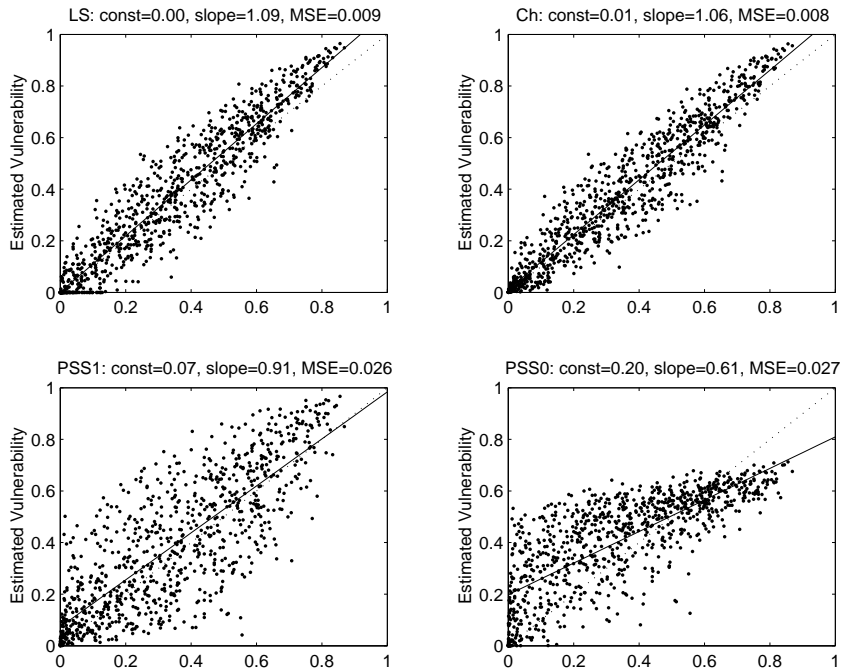


FIGURE 9. Estimates of EP_1 Vulnerability




FIGURE 10. Mean Squared Error in Estimating the EP_2 Measure of Vulnerability as a Function of \sqrt{N} .




FIGURE 11. Spearman Rank Correlation of Estimates of EP_2 (with Actual) as a Function of \sqrt{N} .

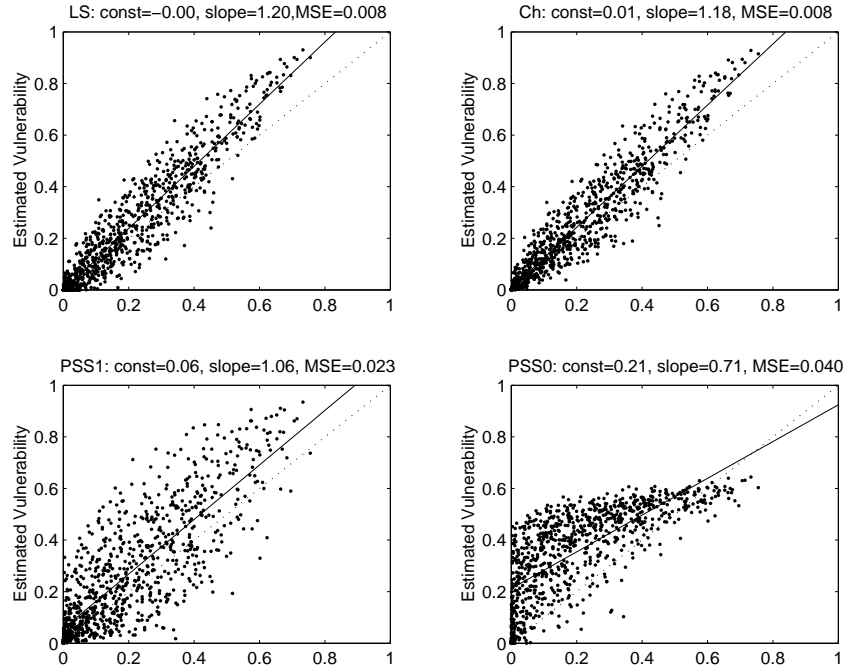


FIGURE 12. Estimates of EP_2 Vulnerability

Estimating the EP_2 measure gives results that are qualitatively similar to the case of estimating EP_1 . There are two chief differences. The first (Figures 10 and 11) is that the Chaudhuri estimator now joins the TS estimator in displaying a substantial bias; as Figure 12 indicates, both of these measures tend to overestimate the vulnerability of the least well-to-do, though both do a good job of estimating EP_2 at the other end of the distribution.

Finally, we turn our attention to the estimation of the Calvo-Dercon measure of vulnerability. We've already noted that the properties of this measure are very similar to that of the Ligon-Schechter measure—the basic difference in this environment is that households which might have had negative vulnerability using the Ligon-Schechter measure will now have their vulnerability set to zero, truncating the distribution.¹³ A comparison of Figure 13 with Figure 1 and of Figure 14 with Figure 2 reveals that the response of mean square error and rank correlations to changes in sample size is essentially identical to the case of the Ligon-Schechter measure. The scatterplots (Figure 15) are also very similar, the chief difference being that the distribution of vulnerability is truncated below at zero in the Calvo-Dercon case.

To summarize the results of this experiment:

- (1) Contrary to what one might have supposed, the best estimator turns out not to depend very much on what measure of vulnerability is chosen—for each of these measures using the stationary Jalan-Ravallion DGP the Chaudhuri estimator has the best performance (in terms of minimizing mean square error), and at all sample sizes.

¹³Strictly speaking, this isn't quite right; only in some states of the world will vulnerability be set to zero, and so the average across observed states need not be zero; still this serves as a rough intuition for the result.




FIGURE 13. Mean Squared Error in Estimating the Calvo-Dercon Measure of Vulnerability as a Function of \sqrt{N} .




FIGURE 14. Spearman Rank Correlation of Estimates of the Calvo-Dercon Measure of Vulnerability with Actual Vulnerability as a Function of \sqrt{N} .

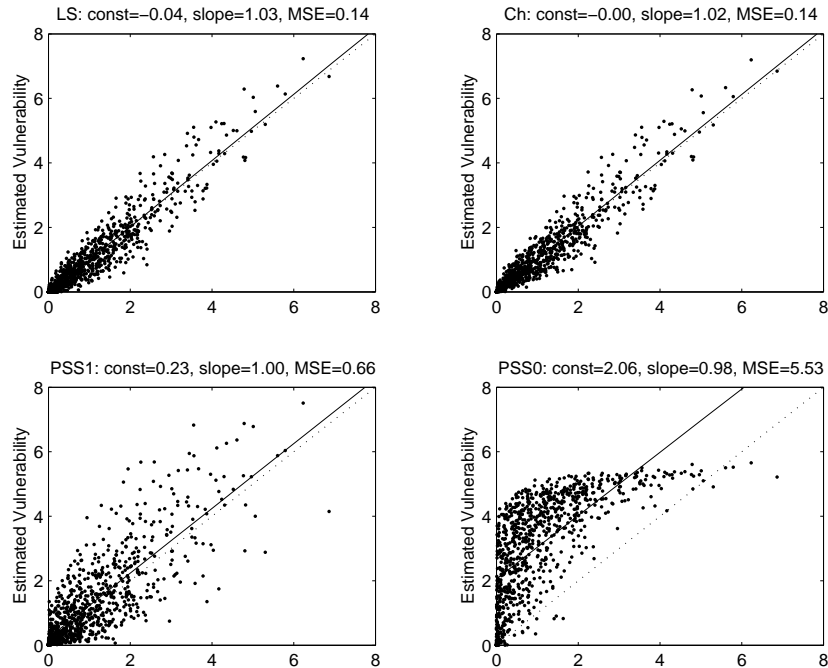


FIGURE 15. Estimates of Calvo-Dercon Vulnerability

- (2) For all measures except EP_0 , Ligon-Schechter’s TS estimator performs almost as well as the Chaudhuri. Because its implementation is much simpler, an analyst may sometimes prefer it for this reason.
- (3) Panel data is necessary for all of these estimators save for the “XS” estimator. However, the number of households represented need not be terribly large to get good performance—for any of the estimators, a panel of a couple thousand households is perfectly adequate.
- (4) If one has only a cross-section, one should think hard before trying to estimate vulnerability. Variation in consumption expenditures in the cross-section will depend both on inequality as well as risk, and it’s likely to be very difficult to disentangle these. In the experiments performed here, it would always be better just to use the static measures $W(c)$ rather than to use the “XS” estimator.

5.2. Accuracy of Vulnerability Estimates as a Function of T . In this subsection we’ll conduct an experiment very similar to that described in Section 5.1, with the difference that here we’re interested in the performance of the various estimators as the length of the panel (T) increases. Results on the mean squared error of the various estimators applied to the various measures are displayed in Figures 16, 18, 20, 22, and 24. Similarly, results on the Spearman rank correlation between estimators and actual vulnerability measures are shown in Figures 17, 19, 21, 23, and 25.

Qualitatively, results across these tables are quite similar, so we discuss them together. The most notable feature of these graphs has once again to do with the distinction between the stationary and non-stationary estimators. In particular, as the length of the panel increases, both the TS and Chaudhuri estimators do better, while the performance of the two Pritchett *et al* estimators is invariant to the length of the panel. In every case but one




FIGURE 16. Mean Squared Error in Estimating the Ligon-Schechter Measure of Vulnerability as a Function of T .




FIGURE 17. Spearman Rank Correlation of Estimates of the Ligon-Schechter Measure of Vulnerability with Actual Vulnerability as a Function of T .




FIGURE 18. Mean Squared Error in Estimating the EP_0 Measure of Vulnerability as a Function of T .




FIGURE 19. Spearman Rank Correlation of Estimates of EP_0 with Actual as a Function of T .




FIGURE 20. Mean Squared Error in Estimating the EP_1 Measure of Vulnerability as a Function of T .




FIGURE 21. Spearman Rank Correlation of Estimates of EP_1 with Actual as a Function of T .




FIGURE 22. Mean Squared Error in Estimating the EP_2 Measure of Vulnerability as a Function of T .




FIGURE 23. Spearman Rank Correlation of Estimates of EP_2 with Actual as a Function of T .




FIGURE 24. Mean Squared Error in Estimating the Calvo-Dercon Measure of Vulnerability as a Function of T .




FIGURE 25. Spearman Rank Correlation of Estimates of the Calvo-Dercon Measure of Vulnerability with Actual as a Function of T .

the ordering of the estimators in terms of their MSE is as in the previous experiment, with the two stationary estimators dominating the non-stationary, and the Chaudhuri estimator edging out the TS estimator.

		Stationary DGP						
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$	
Actual	1.0000	0.9522	0.9534	0.2305	0.8160	0.7252	0.8128	
TS	0.9498	1.0000	0.9986	0.2591	0.8566	0.7462	0.8546	
Chaud.	0.9518	0.9977	1.0000	0.2606	0.8567	0.7515	0.8541	
XS	0.2391	0.2624	0.2643	1.0000	0.3479	0.2080	0.3392	
PSS ₁	0.8096	0.8509	0.8539	0.3388	1.0000	0.8573	0.9984	
PSS ₀	0.8045	0.8440	0.8466	0.2126	0.9870	1.0000	0.8299	
$W(c)$	0.8100	0.8512	0.8542	0.3324	1.0000	0.9881	1.0000	
		Non-stationary DGP						
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$	
Actual	1.0000	0.8385	0.8456	0.3387	0.9995	0.8524	0.9985	
TS	0.8318	1.0000	0.9988	0.2907	0.8366	0.7328	0.8345	
Chaud.	0.8434	0.9985	1.0000	0.2929	0.8437	0.7377	0.8416	
XS	0.3295	0.2787	0.2820	1.0000	0.3338	0.1784	0.3289	
PSS ₁	0.9996	0.8309	0.8426	0.3269	1.0000	0.8451	0.9995	
PSS ₀	0.9796	0.8150	0.8264	0.1843	0.9804	1.0000	0.8310	
$W(c)$	0.9996	0.8309	0.8426	0.3229	1.0000	0.9813	1.0000	

TABLE 2. Correlations Between Estimates of Ligon-Schechter Vulnerability. Correlations reported above the diagonal are simple Pearson correlation coefficients, while those below are Spearman rank correlation coefficients.

		Stationary DGP						
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$	
Actual	1.0000	0.9064	0.9473	0.2296	0.7959	0.7883	0.6865	
TS	0.8635	1.0000	0.9571	0.2353	0.8087	0.7980	0.7614	
Chaud.	0.9512	0.9075	1.0000	0.2452	0.8390	0.8289	0.7291	
XS	0.2201	0.2244	0.2452	1.0000	0.3116	0.3610	0.2511	
PSS ₁	0.8101	0.7755	0.8518	0.3209	1.0000	0.9772	0.8741	
PSS ₀	0.7766	0.7591	0.8202	0.4620	0.9691	1.0000	0.8531	
$W(c)$	0.6470	0.7400	0.6810	0.2493	0.7981	0.7981	1.0000	
		Non-stationary DGP						
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$	
Actual	1.0000	0.8153	0.7891	0.2986	0.9967	0.9645	0.9071	
TS	0.7838	1.0000	0.9671	0.2391	0.8105	0.7874	0.7701	
Chaud.	0.8384	0.8730	1.0000	0.2328	0.7827	0.7711	0.6981	
XS	0.3154	0.2470	0.2624	1.0000	0.2990	0.3739	0.2504	
PSS ₁	0.9991	0.7837	0.8400	0.3120	1.0000	0.9600	0.9037	
PSS ₀	0.9536	0.7707	0.7987	0.4644	0.9523	1.0000	0.8425	
$W(c)$	0.7966	0.7948	0.6654	0.2505	0.7966	0.7966	1.0000	

TABLE 3. Correlations Among Estimators of EP_0 Vulnerability. Correlations reported above the diagonal are simple Pearson correlation coefficients, while those below are Spearman rank correlation coefficients.

5.3. Correlations Between Estimated Vulnerabilities. In this experiment, we provide some evidence bearing on the importance of one's choice of an estimator. In particular, we calculate the correlations between actual vulnerability and each of our five different estimators for each of our five measures of vulnerability, and also a simple poverty measure $W(c)$ (which varies according to which measure we're considering). Results for the stationary Jalan-Ravallion data generating process are presented in the top panels of Tables 2, 2, 3,

Stationary DGP							
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$
Actual	1.0000	0.9451	0.9516	0.2341	0.8100	0.7772	0.7883
TS	0.9442	1.0000	0.9938	0.2564	0.8500	0.8089	0.8384
Chaud.	0.9519	0.9924	1.0000	0.2615	0.8527	0.8140	0.8331
XS	0.2352	0.2548	0.2613	1.0000	0.3389	0.2787	0.3222
PSS ₁	0.8100	0.8500	0.8535	0.3343	1.0000	0.9405	0.9832
PSS ₀	0.8041	0.8448	0.8475	0.3181	0.9931	1.0000	0.9129
$W(c)$	0.7916	0.8395	0.8360	0.3219	0.9856	0.9801	1.0000
Non-stationary DGP							
Actual	1.0000	0.8478	0.8289	0.3294	0.9984	0.9455	0.9896
TS	0.8595	1.0000	0.9949	0.2788	0.8442	0.8139	0.8354
Chaud.	0.8429	0.9926	1.0000	0.2744	0.8241	0.8033	0.8118
XS	0.3272	0.2825	0.2809	1.0000	0.3259	0.2780	0.3153
PSS ₁	0.9998	0.8590	0.8425	0.3256	1.0000	0.9314	0.9912
PSS ₀	0.9893	0.8496	0.8328	0.3107	0.9891	1.0000	0.9107
$W(c)$	0.9857	0.8486	0.8242	0.3139	0.9859	0.9766	1.0000

TABLE 4. Correlations Among Estimators of EP_1 Vulnerability. Correlations reported above the diagonal are simple Pearson correlation coefficients, while those below are Spearman rank correlation coefficients.

Stationary DGP							
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$
Actual	1.0000	0.9453	0.9502	0.2274	0.8063	0.7397	0.7885
TS	0.9420	1.0000	0.9954	0.2542	0.8477	0.7635	0.8371
Chaud.	0.9517	0.9901	1.0000	0.2586	0.8492	0.7723	0.8335
XS	0.2394	0.2578	0.2661	1.0000	0.3433	0.2379	0.3272
PSS ₁	0.8096	0.8465	0.8540	0.3391	1.0000	0.8962	0.9891
PSS ₀	0.8053	0.8420	0.8483	0.2576	0.9911	1.0000	0.8485
$W(c)$	0.7915	0.8353	0.8366	0.3219	0.9855	0.9783	1.0000
Non-stationary DGP							
Actual	1.0000	0.8420	0.8306	0.3315	0.9985	0.8965	0.9922
TS	0.8438	1.0000	0.9965	0.2823	0.8381	0.7766	0.8294
Chaud.	0.8434	0.9930	1.0000	0.2812	0.8256	0.7747	0.8150
XS	0.3297	0.2798	0.2832	1.0000	0.3269	0.2234	0.3158
PSS ₁	0.9995	0.8430	0.8427	0.3296	1.0000	0.8786	0.9953
PSS ₀	0.9867	0.8326	0.8317	0.2407	0.9868	1.0000	0.8484
$W(c)$	0.9855	0.8312	0.8244	0.3139	0.9858	0.9745	1.0000

TABLE 5. Correlations Among Estimators of EP_2 Vulnerability. Correlations reported above the diagonal are simple Pearson correlation coefficients, while those below are Spearman rank correlation coefficients.

4, 5, and 6 (results for the non-stationary DGP are found in the lower panel, but we'll defer discussion of this to Section 6.2). Above the diagonal one finds Pearson correlation coefficients, while below the diagonal are Spearman rank correlation coefficients. These latter may be of particular interest, since they provide some indication of whether or not the various estimators deliver different *rankings* of household vulnerability. As one might expect, for the stationary DGP the two stationary estimators both are very highly correlated, both with each other and with actual vulnerability, while the non-stationary Pritchett *et al* estimators, while highly correlated with each other, have relatively low correlations with actual vulnerability.

There are two estimates in these tables which rely only on cross-sectional data: the XS estimator, which we've already discussed, and a measure labelled $W(c)$ in the tables. This

		Stationary DGP						
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$	
Actual	1.0000	0.9473	0.9498	0.2181	0.8057	0.6863	0.7986	
TS	0.9425	1.0000	0.9972	0.2477	0.8479	0.7012	0.8437	
Chaud.	0.9516	0.9904	1.0000	0.2505	0.8481	0.7100	0.8422	
XS	0.2418	0.2606	0.2683	1.0000	0.3430	0.1880	0.3275	
PSS ₁	0.8093	0.8447	0.8541	0.3429	1.0000	0.8165	0.9960	
PSS ₀	0.8014	0.8348	0.8432	0.1876	0.9818	1.0000	0.7697	
$W(c)$	0.7915	0.8335	0.8370	0.3219	0.9855	0.9694	1.0000	
		Non-stationary DGP						
	Actual	TS	Chaud.	XS	PSS ₁	PSS ₀	$W(c)$	
Actual	1.0000	0.8290	0.8369	0.3322	0.9991	0.8105	0.9966	
TS	0.8266	1.0000	0.9982	0.2855	0.8256	0.7004	0.8210	
Chaud.	0.8437	0.9935	1.0000	0.2882	0.8335	0.7073	0.8288	
XS	0.3313	0.2773	0.2845	1.0000	0.3262	0.1515	0.3175	
PSS ₁	0.9994	0.8255	0.8429	0.3312	1.0000	0.7950	0.9987	
PSS ₀	0.9726	0.8041	0.8210	0.1563	0.9730	1.0000	0.7701	
$W(c)$	0.9853	0.8127	0.8246	0.3139	0.9858	0.9606	1.0000	

TABLE 6. Correlations Between Estimated Calvo-Dercon Vulnerability and Expected Utility. Correlations reported above the diagonal are simple correlation coefficients, while those below are Spearman rank correlation coefficients.

measure corresponds to the welfare function introduced in Section 3, and is related to the vulnerability measure by

$$V^i = EW(c).$$

Thus, while the specification of $W(c)$ varies across the different vulnerability measures we consider, in no case does it incorporate any information about risk—it is best regarded as a sort of static poverty measure. Perhaps the most interesting feature of Tables 2–6 is how well these simple poverty measures perform relative to some of the other estimators. In particular, this simple measure always outperforms the XS estimator by a wide margin. In non-stationary environments it outperforms the stationary estimators, and generally behaves much as does PSS₁ does in both stationary and non-stationary measures for all but the EP_0 measure.

6. EXPERIMENTS WITH MISSPECIFICATION

In this section we conduct a variety of Monte Carlo experiments designed to help us understand the robustness of our various estimators to a selection of quite fundamental specification errors.

6.1. Measurement Error in Consumption Expenditures. We begin by investigating the effect of introducing a modest classical multiplicative measurement error into consumption. In particular, let observed consumption $\tilde{c}_t^i = c_t^i e^{-\sigma_\epsilon^2/2 + \epsilon_t^i}$, where σ_ϵ^2 is the variance of a log normal error. The adjustment $-\sigma_\epsilon^2/2$ makes it so that the expected value of the multiplicative error is equal to one regardless of the value of σ_ϵ . However, as the measurement error is uncorrelated with actual realized consumption its effect is to increase the apparent variation in observed consumption, leading to a bias in estimated vulnerability. This increase in bias is apparent in Figure 26. The variance of $e^{-\sigma_\epsilon^2/2 + \epsilon_t^i}$ is shown on the horizontal axis; the log of the MSE of the various estimators is shown on the vertical.




FIGURE 26. Mean Squared Error of Estimates of the Ligon-Schechter Measure of Vulnerability as a Function of σ_ϵ .




FIGURE 27. Spearman Rank Correlation of Estimates of the Ligon-Schechter Measure of Vulnerability with Actual as a Function of σ_ϵ .

Where in earlier experiments we've worked with five estimators, here we add a sixth—the Ligon-Schechter LB estimator of *explained* vulnerability, which is designed precisely to mitigate the problems associated with measurement error in consumption expenditures. This method proves its worth here—for estimating Ligon-Schechter vulnerability the LB estimator convincingly dominates all the other estimators except when the variance of the measurement error is very small. Results from Figure 26 and Figure 34 (for the Calvo-Dercon measure) are nearly identical, and this basic result is confirmed using evidence on the rank correlation between estimators and actual vulnerability, shown in Figures 27 and 35.




FIGURE 28. Mean Squared Error of Estimates of EP_0 as a Function of σ_ϵ .

As we saw earlier, the Ligon-Schechter TS estimator has problems with the EP_0 measure of vulnerability—too much information is thrown out by simply averaging poverty indicators. These problems are shared by the LB estimator, and extend to the case of measurement error—Figure 28 and Figure 29 indicate that the Chaudhuri estimator is the best estimator even in the presence of measurement error for this particular measure.

Figures 30 and 31 suggest that when estimating EP_1 any of the stationary estimators will do, though if measurement error is quite large the LB estimator ought to be preferred. Finally, evidence from Figures 32 and 33 indicate that the LB estimator really is to be strongly preferred for estimating EP_2 if there's any measurement error at all.

6.2. Non-stationary Consumption Expenditures. The reader will recall that in nearly all of our previous experiments the estimators described by Pritchett et al. (2000) perform very poorly. However, previous experiments have all employed the stationary Jalan-Ravallion DGP. The Ligon-Schechter and Chaudhuri estimators take advantage of this stationarity, while the Pritchett *et al* estimators rely on a weaker requirement of difference stationarity. Here we employ the non-stationary version of the Jalan-Ravallion DGP, and see how the various estimators we've described perform.




FIGURE 29. Spearman Rank Correlation of Estimates of EP_0 with Actual as a Function of σ_ϵ .




FIGURE 30. Mean Squared Error of Estimates of EP_1 as a Function of σ_ϵ .




FIGURE 31. Spearman Rank Correlation of Estimates of EP_1 with Actual as a Function of σ_ϵ .




FIGURE 32. Mean Squared Error of Estimates of EP_2 as a Function of σ_ϵ .




FIGURE 33. Spearman Rank Correlation of Estimates of EP_2 with Actual as a Function of σ_ϵ .




FIGURE 34. Mean Squared Error of Estimates of the Calvo-Dercon Measure of Vulnerability as a Function of σ_ϵ .




FIGURE 35. Spearman Rank Correlation of Estimates of the Calvo-Dercon Measure of Vulnerability with Actual as a Function of σ_ϵ .




FIGURE 36. Mean Squared Error in Estimating the Ligon-Schechter Measure of Vulnerability as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 37. Spearman Rank Correlation of Estimates of the Ligon-Schechter Measure of Vulnerability with Actual as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 38. Mean Squared Error in Estimating EP_0 as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 39. Spearman Rank Correlation of Estimates of EP_0 with Actual as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 40. Mean Squared Error in Estimating EP_1 as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 41. Spearman Rank Correlation of Estimates of EP_1 with Actual as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 42. Mean Squared Error in Estimating EP_2 as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 43. Spearman Rank Correlation of Estimates of EP_2 with Actual as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 44. Mean Squared Error in Estimating the Calvo-Dercon Measure of Vulnerability as a Function of \sqrt{N} in a non-stationary environment.




FIGURE 45. Spearman Rank Correlation of Estimates of the Calvo-Dercon Measure of Vulnerability with Actual as a Function of \sqrt{N} in a non-stationary environment.

Figures 36, 38, 40, 42, and 44 mirror our earlier experiment to understand the effects of sample size on the mean squared error of our estimators. The striking feature of *all* of these graphs is the nearly complete reversal of our earlier ordering—in this non-stationary environment the two stationary estimators do very poorly, while the PSS₁ estimator does rather well. However, unfortunately it’s not enough to use an estimator which relaxes the assumption of stationarity—the PSS₀ estimator also has lousy properties.

Figures 37, 39, 41, 43, and 45 mirror our earlier experiment to understand the effects of sample size on the (Spearman rank) correlations between the various estimators and actual vulnerability measure.

Matters only become worse when we turn our attention to the performance of these estimators as T increases. Figures 46, 48, 50, 52, and 54 report the logarithm of MSE for each estimator and each measure, while Figures 47, 49, 51, 53, and 55 report the Spearman rank correlation coefficient between estimates and the corresponding actual measure. As before, only the PSS₁ is reasonably well behaved—the MSE of the other estimators actually *increases* as the panel grows in length!

The fact that the PSS₁ estimator does well in this non-stationary environment while PSS₀ does poorly leads one to suspect that distributional assumptions are important beyond the issue of stationarity. Figures 56, 58, 60, 62, and 64 tend to confirm this suspicion. Here, we add measurement error (as in Section 6.1) to nonstationary consumption. For very low levels of measurement error the PSS₁ estimator provides the best performance, but the mean squared error of this estimator increases very rapidly for small values of σ_ϵ^2 . In our earlier experiments one estimator tended to dominate, regardless of measure, but that’s unfortunately not the case here. In three of five cases (the Ligon-Schechter measure, EP_2 ,




FIGURE 46. Mean Squared Error in Estimating the Ligon-Schechter Measure of Vulnerability as a Function of T in a non-stationary environment.




FIGURE 47. Spearman Rank Correlation of Estimates of the Ligon-Schechter Measure of Vulnerability with Actual as a Function of T in a non-stationary environment.




FIGURE 48. Mean Squared Error in Estimating EP_0 as a Function of T in a non-stationary environment.




FIGURE 49. Spearman Rank Correlation of Estimates of EP_0 with Actual as a Function of T in a non-stationary environment.




FIGURE 50. Mean Squared Error in Estimating EP_1 as a Function of T in a non-stationary environment.




FIGURE 51. Spearman Rank Correlation of Estimates of EP_1 with Actual as a Function of T in a non-stationary environment.




FIGURE 52. Mean Squared Error in Estimating EP_2 as a Function of T in a non-stationary environment.




FIGURE 53. Spearman Rank Correlation of Estimates of EP_2 with Actual as a Function of T in a non-stationary environment.

and the Calvo-Dercon measure) the ability of the LB estimator to correct for measurement error more than compensates for the failure of stationarity at moderate to large levels of σ_ϵ —it’s worth noting that these are the three cases in which increases in risk actually leads to increases in vulnerability. In the remaining two cases, the Chaudhuri estimator provides the best performance for the EP_0 estimator at moderate levels of σ_ϵ , while for EP_1 the MSE of all of the estimators actually seems to converge as measurement error grows large—presumably this is due to the linearity of this measure over consumption expenditures under the poverty line.




FIGURE 54. Mean Squared Error in Estimating the Calvo-Dercon Measure of Vulnerability as a Function of T in a non-stationary environment.

Figures 57, 59, 61, 63, and 65 show Spearman rank correlations between estimates and actual vulnerability measures.

It’s clear from the results above that one’s choice of an estimator ought to depend on several elements. First, the *measure* being estimated (EP_0 , EP_1 , or one of the risk-sensitive measures); second, whether or not the environment is non-stationary; and third, whether or not consumption expenditures are or are not measured with error. Table 7 provides a guide to which estimator ought to be employed in each circumstance.

To summarize the results of the experiments of this section, and highlight important elements of Table 7,

- (1) For the risk-sensitive vulnerability measures¹⁴ measurement error in consumption will lead to overestimates of vulnerability;

¹⁴That is, those measures which are strictly convex in consumption for poor households, so that those poor households facing more risk are more vulnerable. The risk-sensitive measures we consider here are the LS, CD, and EP_2 measures.




FIGURE 55. Spearman Rank Correlation of Estimates of the Calvo-Dercon Measure of Vulnerability with Actual as a Function of T in a non-stationary environment.




FIGURE 56. Mean Squared Error of Estimates of the Ligon-Schechter Measure of Vulnerability as a Function of σ_ϵ in a non-stationary environment.




FIGURE 57. Spearman Rank Correlation of Estimates of the Ligon-Schechter Measure of Vulnerability with Actual as a Function of σ_ϵ in a non-stationary environment.




FIGURE 58. Mean Squared Error of Estimates of EP_0 as a Function of σ_ϵ in a non-stationary environment.




FIGURE 59. Spearman Rank Correlation of Estimates of EP_0 with Actual as a Function of σ_ϵ in a non-stationary environment.




FIGURE 60. Mean Squared Error of Estimates of EP_1 as a Function of σ_ϵ in a non-stationary environment.




FIGURE 61. Spearman Rank Correlation of Estimates of EP_1 with Actual as a Function of σ_ϵ in a non-stationary environment.




FIGURE 62. Mean Squared Error of Estimates of EP_2 as a Function of σ_ϵ in a non-stationary environment.




FIGURE 63. Spearman Rank Correlation of Estimates of EP_2 with Actual as a Function of σ_ϵ in a non-stationary environment.




FIGURE 64. Mean Squared Error of Estimates of the Calvo-Dercon Measure of Vulnerability as a Function of σ_ϵ in a non-stationary environment.




FIGURE 65. Spearman Rank Correlation of Estimates of the Calvo-Dercon Measure of Vulnerability with Actual as a Function of σ_ϵ in a non-stationary environment.

Measure	Stationary?	Measurement Error?	Best estimator
EP_0	Yes	No	Chaud.
EP_0	Yes	Yes	Chaud.
EP_0	No	No	PSS ₁
EP_0	No	Yes	Chaud.
EP_1	Yes	No	Chaud.
EP_1	Yes	Yes	LB
EP_1	No	No	PSS ₁
EP_1	No	Yes	PSS ₁
Risk-sensitive	Yes	No	Chaud.
Risk-sensitive	Yes	Yes	LB
Risk-sensitive	No	No	PSS ₁
Risk-sensitive	No	Yes	LB

TABLE 7. Which estimator? The first three columns of the table enumerate different circumstances which effect rankings of estimators; the final column indicates the best estimator to use in each circumstance.

- (2) When the time-series distribution of consumption isn't stationary, this has a bad effect on the properties of the stationary estimators (TS, Chaudhuri, and LB) whether there's measurement error not; but

	TS	LB	Chaud.	PSS ₁	PSS ₀	$W(c)$
TS	1.0000	0.9834	0.9831	0.8531	0.7434	0.8559
LB	0.9854	1.0000	0.9985	0.7994	0.7069	0.7993
Chaud.	0.9843	0.9981	1.0000	0.7979	0.7092	0.7974
PSS ₁	0.8254	0.7836	0.7795	1.0000	0.9057	0.9981
PSS ₀	0.7339	0.6936	0.6951	0.9222	1.0000	0.8794
$W(c)$	0.8285	0.7871	0.7824	0.9983	0.8997	1.0000

TABLE 8. Correlation between different estimates of vulnerability in Viet Nam. Spearman rank correlation coefficients are found in the lower left part of the matrix, while Pearson’s coefficients are found in the upper right.

- (3) For any of the risk-sensitive measures of vulnerability, there’s some level of measurement error beyond which the LB estimator dominates all the others in terms of minimizing mean squared error.

7. ESTIMATING VULNERABILITY IN VIET NAM AND BULGARIA

We finally turn our attention to two real-world datasets, drawn from the Viet Nam Household Survey, and briefly described in Section 4.3, and Bulgarian Household Survey described at some length in Ligon and Schechter (2003).

Table 8 reports on the correlation between the various estimators of the Ligon-Schechter measure of vulnerability using data from the VHS. As with the Monte Carlo experiments conducted earlier in the paper, there’s a clear difference between the estimators which assume stationarity and the two PSS estimators, with each of these two groups displaying high within-group correlations, but much lower correlations across groups.

Table 9 has correlations between different vulnerability measures, each calculated with their respective ‘best estimator’¹⁵ and between different actual (not expected) poverty measures calculated using consumption in the first period. As one might expect, the table indicates that the various measures fall into two groups in such a way that within each group correlations are very high. In particular, the risk-sensitive vulnerability measures all give quite similar orderings, as do EP_0 and EP_1 . Notably, EP_2 is more highly correlated with the other risk-sensitive measures than it is with either of the other two expected poverty measures *or* with the poverty measure P_2 ; more generally rankings given by all vulnerability measures are much more similar to each other than they are to the rankings given by the actual poverty measures. This strongly suggests that even with only two periods of data, vulnerability measures already tell the researcher much more than simple poverty measures.

Table 10 reports average levels and decompositions of estimates of the Ligon-Schechter measure of vulnerability in the VHS. The first column of this table reports estimates of total vulnerability; subsequent columns use the techniques described in Section 2.4 to decompose these estimates of total vulnerability into estimates of poverty and risk; the LB estimator can be used to further decompose risk into several component parts. By summing the first row across the columns marked “Pov” (poverty), “Agg Risk” (aggregate risk), “Inc Risk” (income

¹⁵Remember that when one assumes non-stationarity and measurement error then one should use Chaudhuri’s estimator to calculate EP_0 , PSS₁ to calculate EP_1 , and LB to calculate any of the risk sensitive measures.

Viet Nam								
	LS	CD	EP_0	P_0	EP_1	P_1	EP_2	P_2
LS	1.0000	0.9265	0.9457	0.5984	0.7924	0.7493	0.9175	0.7223
CD	0.9352	1.0000	0.8008	0.5379	0.8159	0.8251	0.9552	0.8552
EP_0	0.9958	0.9286	1.0000	0.6591	0.7666	0.7070	0.7609	0.6126
P_0	0.6565	0.7045	0.6522	1.0000	0.7662	0.7094	0.4503	0.5062
EP_1	0.7843	0.8437	0.7764	0.8283	1.0000	0.9642	0.7148	0.9018
P_1	0.7582	0.8493	0.7499	0.8506	0.9754	1.0000	0.7179	0.9373
EP_2	0.9800	0.9316	0.9719	0.6589	0.7843	0.7683	1.0000	0.7575
P_2	0.7582	0.8493	0.7499	0.8506	0.9754	1.0000	0.7683	1.0000

Bulgaria								
	LS	CD	EP_0	P_0	EP_1	P_1	EP_2	P_2
LS	1.0000	0.9659	0.8882	0.5307	0.7060	0.6825	0.9309	0.6652
CD	0.9825	1.0000	0.7717	0.4442	0.6851	0.6622	0.9840	0.6835
EP_0	0.9995	0.9818	1.0000	0.5993	0.6196	0.6144	0.6847	0.5177
P_0	0.5906	0.5613	0.5909	1.0000	0.7040	0.7259	0.4055	0.5172
EP_1	0.6926	0.6680	0.6930	0.8483	1.0000	0.9821	0.6706	0.9610
P_1	0.6704	0.6548	0.6710	0.8783	0.9662	1.0000	0.6473	0.9376
EP_2	0.9900	0.9769	0.9891	0.5797	0.6800	0.6640	1.0000	0.6833
P_2	0.6704	0.6548	0.6710	0.8783	0.9662	1.0000	0.6640	1.0000

TABLE 9. Correlation between different measures of vulnerability in Viet Nam and Bulgaria, each measure estimated using its preferred estimator. Spearman rank correlation coefficients are found below the diagonal, while Pearson correlation coefficients lie above the diagonal.

	Vuln =	Pov +	Agg Risk +	Inc Risk +	HHSze Risk +	Unexp Risk
TS	0.1648 =	0.0803 +	0.0140 +	0.0004 +	0.0033 +	0.0667
Chaud.	0.2014 =	0.0361+			0.1653	
PSS ₁	0.4264 =	0.1420+			0.2844	
PSS ₀	1.5703 =	0.2000+			1.3703	

TABLE 10. Decomposition of Estimates of Vulnerability in Viet Nam. Note that the LB estimate is simply the sum of the first four terms on the right-hand side of the equality in the first row.

risk), and “HHSize risk” (risk associated with changes in household size), one can compute the LB estimate of vulnerability, equal to 0.0980 in this case. With our assumption that $\gamma = 2$, this quantity can be interpreted as the percentage increase in consumption expenditures necessary to compensate the average household for its loss in welfare due to poverty (8.03 per cent) and explained risk (1.77 per cent), relative to a utopia in which there was neither inequality nor risk. By adding unexplained risk we obtain the TS estimate of Ligon-Schechter vulnerability; this adds an additional 6.67 per cent to estimated vulnerability, but some or all of this amount may in fact be due to measurement error in consumption.

The remaining estimators don't decompose as naturally as do the Ligon-Schechter estimators, but since each estimator yields an estimate of expected consumption along with its variance, it's easy to use these estimates to construct separate measures of poverty and risk. Chaudhuri's estimator yields lower estimates of expected consumption than do the Ligon-Schechter estimators—using his estimator it appears that the compensating variation due to inequality in rural Viet Nam amounts to 3.61 per cent of average per household consumption, while the compensating variation for risk from all sources is 16.53 per cent. The remaining non-stationary estimators, not surprisingly, place an even greater emphasis on the importance of risk in reducing welfare, with PSS_1 yielding a compensating variation for poverty and risk of 14.2 and 28.44 per cent, respectively. Finally, the PSS_0 estimator yields estimates of these compensating variations equal to 20 per cent for poverty, and 137.03 per cent for risk (this last enormous figure because this estimator estimates a high probability of zero consumption for many households).

	Vuln =	Pov +	Agg Risk +	Idio Risk +	Unexp Risk
Vietnam					
Quintile 1	0.5465 =	0.3942 +	0.0173 +	0.0059 +	0.1292
Quintile 2	0.2813 =	0.1958 +	0.0154 +	0.0048 +	0.0653
Quintile 3	0.1356 =	0.0702 +	0.0141 +	0.0036 +	0.0477
Quintile 4	0.0129 =	-0.0425 +	0.0128 +	0.0027 +	0.0399
Quintile 5	-0.1518 =	-0.2156 +	0.0106 +	0.0017 +	0.0515
Bulgaria					
Quintile 1	0.5580 =	0.4536 +	0.0135 +	0.0018 +	0.0891
Quintile 2	0.2886 =	0.1998 +	0.0118 +	0.0016 +	0.0754
Quintile 3	0.1558 =	0.0729 +	0.0108 +	0.0016 +	0.0705
Quintile 4	0.0387 =	-0.0400 +	0.0098 +	0.0018 +	0.0672
Quintile 5	-0.1178 =	-0.1903 +	0.0083 +	0.0017 +	0.0625

TABLE 11. Ligon-Schechter Vulnerability by Consumption Quintiles.

So which of the various estimates of vulnerability presented in Table 10 should we use? In a real-world application we seldom have the luxury of knowing in advance whether, e.g., consumption expenditures are stationary or whether measurement error is important. Fortunately, the evidence presented in Section 6 also suggests ways in which the data itself can tell us which estimator ought to be employed.

We first consider the issue of whether or not measurement error is an important issue. The difference between the TS and LB estimators of vulnerability gives us an upper bound on how much measurement error there may be. Consulting Table 10, we see that in Viet Nam measurement error explains up to 40.5 per cent of total Ligon-Schechter vulnerability (6.67/16.48); via a simple bootstrap test, we find that this is significant (the test has a p -value indistinguishable from zero). Similarly, consulting Table 12 for Bulgaria, we see that measurement error explains as much as 39.5 per cent of total Ligon-Schechter vulnerability, and is similarly significant. Accordingly, we conclude either that the LB estimator ought to be preferred for both estimators, or that both estimators ought to be employed, giving upper and lower bounds on total vulnerability.

Consulting Table 7, we next seek to determine whether or not the consumption process is stationary or not. One crude but simple test of this is simply to add lagged log consumption expenditures to the basic LB regression (6); if this lagged variable turns out to help to explain variation in subsequent consumption, this is evidence against the null hypothesis of stationarity. Unfortunately, this is of no help in the VHS, as one needs at least three rounds of data to implement the test. However, in the Bulgarian data we have 12 rounds in total. Estimating the proposed regression, we find a coefficient on lagged log consumption equal to 0.20, with a standard error of 0.0063, and hence highly significant. As it happens, Table 7 still recommends the LB estimator in this instance.

Thus, Table 12 uses the LB estimator to estimate vulnerability for all the households in both samples. We decompose total vulnerability into several components for each household, as described in Section 3. It's then of some interest to compare vulnerability, poverty, and different sources of risk across quintiles (Table 11). The overall level of vulnerability in both countries is quite similar, though vulnerability is slightly higher in Bulgaria than in Vietnam.¹⁶ Interestingly, this greater vulnerability in the poorest quintile in Bulgaria seems to arise mainly due to the higher poverty level in Bulgaria. The poorest quintile in Bulgaria faces much less risk than the poorest quintile in Vietnam—accordingly, it's the poorest group across both samples that bears the greatest risk. However, all other quintiles face more risk in Bulgaria than in Vietnam.

In aggregate the Bulgarian population faces less explained risk, both aggregate and idiosyncratic than do the Vietnamese.¹⁷ On the other hand, all the Bulgarians except those in the poorest quintile experience more unexplained risk than do the Vietnamese. This either means that Bulgarians truly experience more unexplained idiosyncratic risk than do the Vietnamese, or that the Bulgarian data set has more measurement error.

In Table 12 we decompose vulnerability into poverty, aggregate risk, idiosyncratic risk, and unexplained risk, as in Ligon and Schechter (2003).¹⁸ Poverty is the largest single component of vulnerability. After that, unexplained risk is the second largest component, and aggregate risk is the third largest component. Explained idiosyncratic risk is quite small, although the considerably larger unexplained risk may be made up of much idiosyncratic risk due to unobservable shocks.

We also look at the correlates of these elements of vulnerability. To do this we regress each element of vulnerability on a set of fixed household characteristics. For household characteristics which vary over the 12 month period, we use the value of that characteristic in the first period as our right hand side variable. It is interesting to note that the correlates of vulnerability are extremely similar to the correlates of poverty. In addition, it is to be expected that the correlates of aggregate risk will be the same as the correlates of poverty.

¹⁶Because poverty is defined relative to per capita consumption in each country, this can't be due to differences in the average level of expected consumption; rather, it's an indication that there's more inequality in Bulgaria.

¹⁷In the Viet Nam data we calculate idiosyncratic risk from changes in income and changes in household size, while in Bulgaria we calculate idiosyncratic risk from changes in income, changes in the number of workers in the household, and changes in the number of unemployed in the household. Even though idiosyncratic risk is calculated using more variables in Bulgaria than in Vietnam, there is still more idiosyncratic risk experienced in Viet Nam.

¹⁸The following analysis draws heavily from our earlier paper; however, estimates provided here differ because the first stage regression (6) is specified in logs here, while it was specified in levels in our earlier paper.

Average Value (in utils)	Vuln	Pov	Agg Risk	Idio Risk	Unexp Risk
	0.1847*** [0.1789, 0.1914]	0.0992*** [0.0940, 0.1050]	0.0108*** [0.0101, 0.0117]	0.0017*** [0.0013, 0.0022]	0.0729*** [0.0707, 0.0751]
Variable	Coef	Coef	Coef	Coef	Coef
Primary Ed.	-0.0575** (0.0241)	-0.0551** (0.0226)	-0.0004*** (0.0002)	-0.0001 (0.0001)	-0.0018 (0.0024)
Secondary Ed.	-0.1818*** (0.0263)	-0.1717*** (0.0247)	-0.0013*** (0.0002)	-0.0001 (0.0001)	-0.0086*** (0.0027)
Post-Sec. Ed.	-0.2612*** (0.0293)	-0.2478*** (0.0275)	-0.0020*** (0.0002)	-0.0003* (0.0002)	-0.0111*** (0.0031)
Male	-0.0210 (0.0190)	-0.0164 (0.0178)	-0.0001 (0.0001)	0.0000 (0.0001)	-0.0046** (0.0018)
Age	0.0039 (0.0034)	0.0036 (0.0032)	0.0000 (0.0000)	0.0000 (0.0000)	0.0002 (0.0003)
Age Squared	-0.0000 (0.0000)	-0.0000 (0.0000)	-0.0000 (0.0000)	-0.0000 (0.0000)	-0.0000 (0.0000)
Own Animal	-0.0783*** (0.0189)	-0.0734*** (0.0178)	-0.0006*** (0.0001)	-0.0000 (0.0001)	-0.0042*** (0.0016)
Land Cultivated	-0.0014 (0.0022)	-0.0012 (0.0021)	-0.0000 (0.0000)	0.0000 (0.0000)	-0.0002 (0.0002)
Urban	0.0544*** (0.0182)	0.0529*** (0.0171)	0.0004*** (0.0001)	-0.0001 (0.0001)	0.0011 (0.0016)
# of Pens.	-0.0504*** (0.0144)	-0.0467*** (0.0136)	-0.0002** (0.0001)	-0.0000 (0.0001)	-0.0034** (0.0014)
# of Emp.	-0.1735*** (0.0142)	-0.1654*** (0.0134)	-0.0012*** (0.0001)	0.0003*** (0.0001)	-0.0072*** (0.0012)
Fam. Size	0.1488*** (0.0084)	0.1406*** (0.0080)	0.0011*** (0.0001)	-0.0001* (0.0001)	0.0072*** (0.0007)
R^2	0.2673	0.2712	0.2662	0.0247	0.0806

TABLE 12. Ligon-Schechter Vulnerability by Household Characteristics.

Aggregate shocks are, by definition, the same for all households, and so the poor households will experience greater impact on their utility from this component of risk.

We find that households with more educated heads are less vulnerable, with college educated head being on average 26% less vulnerable than households with uneducated heads. Most of this reduction is due to educated households having higher expected consumption expenditures, but these highly educated households also face significantly less aggregate and unexplained risk. Households which own animals or live in villages (as opposed to cities) are also less vulnerable, mostly because of their higher consumption. Given that one usually considers agriculture to be a more risky source of livelihood, it is interesting that these households experience no higher risk than other households. (On the other hand, decares of land cultivated has no significant effect on vulnerability.) Perhaps this is because of unobserved mutual insurance mechanisms which are at work.

Households which have many pensioners or workers but smaller family size are less vulnerable. This means that having a family which includes more income earning members (pensioners and workers) decreases vulnerability. Families with more unemployed and children, on the other hand, are more vulnerable. Perhaps this is because of the greater number of non-income generating members in these households. Female headed households are no more vulnerable overall than male headed households. Female headed households do experience lower unexplained risk than do male households. Households with older household heads are no more or less vulnerable than households with younger household heads.

8. CONCLUSION

In summarizing the results of our experiments it seems useful to draw three different sorts of distinctions. The first distinction is one between vulnerability measures. Consider any increase in the risk (in the sense of Rothschild and Stiglitz (1970)) borne by poor households. If a vulnerability measure increases when risk increases, then we call it “risk-sensitive” (the risk sensitive measures we consider are the Ligon-Schechter measure, the Calvo-Dercon measure, and EP_2).

The second and third distinctions have to do with the data-generating process. In particular, whether or not the DGP is stationary or not is important, and whether or not consumption expenditures are measured with error is important. With these distinctions drawn, we are ready to draw our main conclusion: In an environment with high levels of measurement error, the LB estimator should be used to estimate any risk-sensitive vulnerability measure.

Other conclusions may also be drawn from these experiments. If the distribution of consumption is stationary and measurement error is not an issue, then the estimator of Chaudhuri (2001) ought to be employed. Similarly, if the distribution of consumption is *non-stationary*, this distribution is known, and measurement error is unimportant, then the maximum-likelihood estimators described by Pritchett et al. (2000) will tend to perform best.

We conducted experiments using one estimator that uses only cross-sectional data. One suspects that it ought to be very difficult to draw reliable inferences about risk from cross-sectional data, since the variation in consumption across households will be due at least in part to inequality. Our experiments with the XS estimator strongly reinforce these suspicions; attempting to estimate risk from cross-sectional data led to estimates of vulnerability which were much inferior to simply looking at poverty measures. It’s hard to avoid the conclusion that if all one has is cross-sectional data, one shouldn’t bother even trying to estimate vulnerability, but rather stick with poverty measures instead.

Because all the remaining estimators considered here require panel data, it’s worth considering how the size of the dataset affects these estimators. One piece of good news is that with a number of households equal to a few thousand (roughly the size of a typical LSMS) estimates of vulnerability tend to be close to their limiting values. Another piece of good news is that it’s possible to get reasonable estimates of vulnerability even when the longitudinal dimension of the panel is as short as two years. However, the value of extending the length of the panel is very large, so long as the estimator is appropriately chosen, and this is likely to be even more true for yet-to-be-designed estimators which relax the assumption of stationarity while continuing to deal with measurement error.

Since poverty measures are the best measure of vulnerability when one has only cross-sectional data, is it worthwhile to estimate vulnerability at all? In general the answer to

this question will depend on the environment and dataset. However, we note that in our experiments vulnerability estimates gave much better measures of household welfare than did poverty measures. Furthermore, the information provided by vulnerability measures is potentially much richer—because one can decompose vulnerability into distinct measures of poverty and risk from various sources, one ought to be able to construct much better informed policy. It also seems likely that some sources of risk can be more easily and inexpensively addressed than poverty.

8.1. Directions for future research. We’ve shown that in some environments simple estimators can be used with panel datasets to yield reliable estimates of vulnerability, and that these measures of vulnerability yield much more information regarding household welfare than do poverty measures. We’ve also provided an illustration of how one might go about selecting the best estimator for a particular dataset; this suffices to make estimating vulnerability in a given panel dataset simple and practical.

However, there remain some clear directions for future research. We close by enumerating some of these, in no particular order.

Better Estimators: Of the estimators we’ve considered, three stand out: Chaudhuri’s panel estimator, Ligon and Schechter’s LB estimator, and the (modified) PSS_1 estimator of Pritchett et al. (2000). Each of these is designed to deal with different problems. Chaudhuri’s estimator takes advantage of information on the variance of residuals to improve efficiency, much as generalized least squares improves on OLS. PSS_1 is robust to nonstationarity in consumption. The LB estimator uses a method analogous to 2SLS to make it robust to measurement error in consumption.

Despite their various strengths, none of these three estimators deals with more than one problem at a time. Fortunately, this seems as though it may be fairly easy to change. It’s quite possible, for example, that estimating the first-stage of the LB estimator in differences (à la PSS_1) to deal with nonstationarity and then estimating the variance structure of the residuals (à la Chaudhuri) would give us the best of all three worlds—an efficient estimator robust to non-stationarity and measurement error.

Estimating Preferences: One claim advanced for (at least utility-based) vulnerability measures is that they avoid the paternalism inherent in poverty measures by reflecting the preferences of the households themselves. However, while there’s working consensus in the empirical literature on what functions usefully reflect household risk preferences (the HARA class), this still leaves at least one free parameter. Here we’ve chosen to simply fix this parameter at a value within a consensus range, but it would be better to be able to estimate it, and also to allow for various forms of heterogeneity in household preferences. Estimating preference parameters will require the analyst to observe the outcomes of household decisions (e.g., savings decisions, labor supply, etc.) which depend in part on risk attitudes.

Dynamics: Our experiments have revealed the possible importance of nonstationarity in consumption expenditures for estimation, and economic theory strongly suggests that consumption expenditures ought to be expected to be non-stationary whenever household bear idiosyncratic risk. But this means that vulnerability will change from period to period.

Elbers and Gunning (2003) take a sensible step forward by describing a utility-based measure of vulnerability which is naturally dynamic. This offers an elegant solution to the possible problem of nonstationarity, and incorporates into vulnerability measures not only information about the risks households face, but also about predictable variation in their consumption over time. Ligon (2003) offers a simple way to estimate vulnerability in a dynamic setting.

Making vulnerability measures fully dynamic sounds like a substantial complication, but actually resolves some of the most difficult conceptual and statistical issues which faced us in estimating the static measures of vulnerability in this paper. Further, much work currently undertaken by researchers interested in dynamic consumer behavior (e.g., Meghir and Pistaferri, 2004) could be exploited to estimate dynamic vulnerability without great difficulty.

Applications: Practitioners have not yet settled on one measure or estimator for estimating vulnerability in practice, something we hope this paper will help to change. We need to see the methods described here applied to different datasets in different settings, so as to get results which can be sensibly compared across populations. We've authored code for both Stata and Matlab,¹⁹ and look forward to helping others to use it.

REFERENCES

- Amin, S., A. S. Rai, and G. Topa (1999, October). Does microcredit reach the poor and vulnerable? Evidence from northern Bangladesh. Working Paper 28, Center for International Development at Harvard University.
- Calvo, C. and S. Dercon (2003). Vulnerability: An axiomatic approach (the v_γ family of vulnerability measures). Unpublished manuscript.
- Chaudhuri, S. (2001). Empirical methods for assessing household vulnerability to poverty. Unpublished Manuscript.
- Chaudhuri, S., J. Jalan, and A. Suryahadi (2001). Assessing household vulnerability to poverty from cross-sectional data: A methodology and estimates from Indonesia. Unpublished Manuscript.
- Christiaensen, L. J. and R. N. Boisvert (2000). On measuring household food vulnerability: Case evidence from Northern Mali. Unpublished Manuscript.
- Clark, S., R. Hamming, and D. Ulph (1981). On indices for the measurement of poverty. *The Economic Journal* 91, 515–526.
- Dercon, S. and P. Krishnan (2000). Vulnerability, seasonality and poverty in Ethiopia. *Journal of Development Studies* 36(6), 25–53.
- Elbers, C. and J. W. Gunning (2003). Vulnerability in a stochastic dynamic model. Discussion Paper TI 2003-070/2, Tinbergen Institute, Vrije Universiteit, Amsterdam.
- Foster, J., J. Greer, and E. Thorbecke (1984). A class of decomposable poverty measures. *Econometrica* 52(3), 761–766.

¹⁹Stata code designed to make it easy to estimate Ligon-Schechter measures of vulnerability is available at are.berkeley.edu/~ligon/papers.html; the matlab code actually used to conduct the experiments described in this paper is available on request.

- Glewwe, P. and G. Hall (1995). Who is most vulnerable to macroeconomic shocks?: Hypotheses tests using panel data from Peru. Living Standards Measurement Study Working Paper 117, The World Bank, Washington, D.C.
- 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* 56(1), 181–206.
- Hoddinott, J. and A. Quisumbing (2003). Methods for microeconometric risk and vulnerability assessments: A review with empirical examples. Unpublished Ms.
- Jalan, J. and M. Ravallion (1999). Are the poor less well insured? Evidence on vulnerability to income risk in rural China. *Journal of Development Economics* 58, 61–81.
- Jalan, J. and M. Ravallion (2000). Is transient poverty different? Evidence for rural China. *Journal of Development Studies* 36(6), 82.
- Ligon, E. (2003). Dynamic vulnerability. Unpublished manuscript.
- Ligon, E. and L. Schechter (2003). Measuring vulnerability. *Economic Journal* 113(486), C95–C102.
- Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica*. Forthcoming.
- Pritchett, L., A. Suryahadi, and S. Sumarto (2000). Quantifying vulnerability to poverty: A proposed measure, with application to Indonesia. SMERU Working paper.
- Ravallion, M. (1988). Expected poverty under risk-induced welfare variability. *Economic Journal* 98(393), 1171–1182.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2, 225–243.
- Watts, H. W. (1968). An economic definition of poverty. In D. P. Moynihan (Ed.), *On Understanding Poverty*. New York: Basic Books.