The Economic Journal, **113** (*March*), C95–C102. © Royal Economic Society 2003. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

MEASURING VULNERABILITY*

Ethan Ligon and Laura Schechter

Traditional poverty measures neglect several important dimensions of household welfare. In this paper we construct a measure of 'vulnerability' which allows us to quantify the welfare loss associated with poverty as well as the loss associated with any of a variety of different sources of uncertainty. Applying our measure to a panel dataset from Bulgaria in 1994, we find that poverty and risk play roughly equal roles in reducing welfare. Aggregate shocks are more important than idiosyncratic sources of risk, but households headed by an employed, educated male are less vulnerable to aggregate shocks than are other households.

Economists have long used measures of poverty to summarise 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 recognised 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 in 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 they 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. Here we propose a simple definition of what we term *vulnerability*, and a simple technique for identifying vulnerable populations.

Our method may be contrasted with related efforts by several other authors (discussed at some length in Ligon and Schechter (2002)). Several papers have sought to address the issues of risk and poverty by estimating expected values of the poverty indices introduced by Foster *et al.* (1984). However, while useful for measuring poverty, these indices have several perverse features when trying to measure the welfare consequences of risk, and a policymaker who sought to allocate resources to minimise the expected value of one of these indices would tend to assign too much risk to poorer households.[1]

* We thank Emmanuel Skoufias for providing the data used in the application of this paper, and UN/WIDER for supporting, in part, research on this theme.

[1] Foster *et al.* (1984) define a family of poverty measures P_α, where α is a curvature parameter. Christiaensen and Boisvert (2000); Pritchett *et al.* (2000); Chaudhuri (2001) and Chaudhuri *et al.* (2001) all estimate EP_0. Importantly, a policymaker could reduce EP_0 simply by assigning as much risk to poor households as possible. Ravallion (1988) and Kamanou and Morduch (2001) seek to estimate EP_2; in this case, poor households are implicitly assumed to have increasing absolute risk aversion.

We proceed as follows. In Section 1, we define a utilitarian measure of vulnerability, and describe how it may be decomposed into distinct measures of poverty, aggregate risk, and idiosyncratic risk. In Section 2, we describe a method for estimating the vulnerability measure defined in Section 1, and show how our estimated measure is robust to measurement error in consumption expenditures, unlike measures proposed elsewhere.

We apply our techniques to a dataset from Bulgaria; these data are of particular interest because the data are collected at monthly intervals, and allow us, in Section 3, to explore the extent to which different groups suffered the effects of a major restructuring of the Bulgarian economy. Section 4 concludes.

1. Defining Vulnerability

We take a utilitarian approach to defining vulnerability in a risky environment. Suppose there to be a finite population of households indexed by $i = 1, 2, \ldots, n$, and let $\omega \in \Omega$ denote the state of the world. We focus on the distribution of household i's consumption expenditures, $c^i(\omega)$. To measure vulnerability, for each household we first choose some strictly increasing, weakly concave function $U^i : \mathbb{R} \rightarrow \mathbb{R}$ mapping consumption expenditures into the real line. Given the function U^i, we define the vulnerability of the household by the function

$$V^i = U^i(z) - \mathrm{E}U^i(c^i). \qquad (1)$$

Here z is some certainty-equivalent consumption such that *if* household i had certain consumption greater than or equal to this number, we would not regard the household as vulnerable. Thus, the choice of z is analogous to the choice of a 'poverty line' in the literature on poverty measurement.

One way to motivate a particular choice of z is explicitly to measure *relative* vulnerability within the population. In this case, consider an allocation in which every household receives the expected *per capita* consumption bundle with certainty. Since there is no inequality, there can be no relative poverty; since there is no uncertainty there can be no risk. Thus, for this allocation one would want our measure of vulnerability to be equal to 0. This is accomplished simply by setting z equal to expected *per capita* consumption expenditures.

Taking expectations of an increasing, concave function of consumption expenditures has the effect of making vulnerablity depend not only on the mean of a household's consumption, but also on variation in consumption. To understand the balance between poverty and risk in our measure of vulnerability better, note that we can decompose the measure into distinct components reflecting poverty and risk, respectively:

$$V^i = [U^i(z) - U^i(\mathrm{E}c^i)] + [U^i(\mathrm{E}c^i) - \mathrm{E}U^i(c^i)]. \qquad (2)$$

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 $\mathrm{E}c^i$ approaches the poverty line, an additional unit of expected consumption has diminishing marginal value in reducing poverty. For a suitable choice of $\{U^i\}$, it is easy to show that this poverty measure satisfies all the axiomatic requirements enumerated in Foster *et al.* (1984).

The second term of (2), which measures the risk faced by household *i*, is consistent with the ordinal measures of risk proposed by Rothschild and Stiglitz (1970) (though any monotone transformation would do as well). Further, this risk measure can usefully be further decomposed into two distinct measures of risk, one aggregate, the other idiosyncratic. Let $\mathrm{E}(c^i|\bar{\mathbf{x}})$ denote the expected value of consumption, c^i, conditional on knowledge of a vector of aggregate variables $\bar{\mathbf{x}}$. 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. Putting it all together, we have

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

Of course, the notation here is intentionally chosen to evoke comparisons with utility functions. Minimising vulnerability is similar to maximising the utilitarian social welfare function

$$\max_{\{c^i(\omega)\}} \sum_{i=1}^{n} \mathrm{E}U^i(c^i)$$

subject to some aggregate resource constraint.

Despite the notation, our proposed procedure need not be interpreted as a utilitarian social welfare function. One of several possible alternative interpretations would have a paternalistic donor or NGO choose some concave function for the $\{U^i\}$, with the shape of the function reflecting the *donor*'s preferences over the distribution and uncertainty of consumption expenditures. One happy consequence is that it is not necessary to be able to measure individual households' utility functions.

2. Estimating Vulnerability

Two additional steps are required before one can actually use data to compute a household's vulnerability. First, one must choose the functions $\{U^i\}$. Second, one must devise a way to estimate the conditional expectations which figure in our vulnerability measure. Each of these choices has an important impact on measured risk and poverty. Here, we assume that the $\{U^i\}$ take the simple form $U^i(c) = (c^{1-\gamma})/(1-\gamma)$ for some parameter $\gamma > 0$; as γ increases, the function U^i becomes increasingly sensitive to both risk and inequality. We normalise c so that the average of consumption over all households in all periods equals 1.

Despite the apparently static nature of the vulnerability function defined above, to estimate risk we rely on variation over time. Accordingly, we denote the time t realisation of household i's consumption expenditures as c^i_t, of household i's other idiosyncratic variables as $\mathbf{x}^i_t$, and of the vector of aggregate variables by $\bar{\mathbf{x}}_t$. We assume that $\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t, \mathbf{x}_t^i) = \alpha^i + \eta_t + \mathbf{x}_t^{i\prime}\boldsymbol{\beta}$, with $\boldsymbol{\theta} = (\alpha^i, \eta_t, \boldsymbol{\beta}')$ a vector of unknown parameters to be estimated.[2]

In practice, consumption expenditures are likely to be measured with error, and using observed consumption to measure vulnerability as in Section 1 would lead the analyst to conflate measurement error with idiosyncratic risk. To avoid this problem, we further decompose our measure of idiosyncratic risk into risk which can be attributed to variation in k observed time-varying household characteristics $\mathbf{x}_t^i = (x_{1t}^i, \ldots, x_{kt}^i)'$ and risk which can neither be explained by these characteristics nor by aggregate variables, but which is due instead to variation in unobservables and to measurement error in consumption. Let $\tilde{c}_t^i = c_t^i + \epsilon_t^i$, with $\{\epsilon_t^i\}$ a measurement error process having the property that $\mathrm{E}(\epsilon_t^i|\mathbf{x}_t^i, \bar{\mathbf{x}}_t) = \mathrm{E}(\epsilon_t^i c_t^i) = 0$. Then rewriting the expression for vulnerability yields

$$\begin{aligned} V^i = {} & [U^i(\mathrm{E}c) - U^i(\mathrm{E}c_t^i)] && \text{(Poverty)} \\ & + \{U^i(\mathrm{E}c_t^i) - \mathrm{E}U^i[\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t)]\} && \text{(Aggregate risk)} \\ & + \{\mathrm{E}U^i[\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t)] - \mathrm{E}U^i[\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t, \mathbf{x}_t^i)]\} && \text{(Idiosyncratic risk)} \\ & + \{\mathrm{E}U^i[\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t, \mathbf{x}_t^i)] - \mathrm{E}U^i(c_t^i)\}. && \text{(Unexplained risk \& measurement error)} \end{aligned} \tag{3}$$

We assume a stationary environment, and so we are led to estimate the unconditional expectation of household i's consumption by $\mathrm{E}c_t^i = (1/T)\sum_{t=1}^T c_t^i$. For the present application, we wish to choose $\boldsymbol{\theta}$ so as to predict c_t^i optimally in a least-squares sense. Measurement error in observed expenditures will not bias our estimates of $\boldsymbol{\theta}$, but will bias our estimate of total risk. However, given our assumptions on the measurement error process $\{\epsilon_t^i\}$, it follows that $\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t, \mathbf{x}_t^i) = \mathrm{E}(\tilde{c}_t^i|\bar{\mathbf{x}}_t, \mathbf{x}_t^i)$, so that measurement error in consumption expenditures will influence only our measure of *unexplained* risk. This last measure will be incorrect by the difference $\mathrm{E}U^i(\tilde{c}_t^i) - \mathrm{E}U^i(c_t^i)$, while our measures of aggregate and explained idiosyncratic risk will not be biased by this sort of measurement error.

Our parameterisation of $\mathrm{E}(c_t^i|\bar{\mathbf{x}}_t, \mathbf{x}_t^i)$ suggests the linear estimating equation

$$\tilde{c}_t^i = \alpha^i + \eta_t + \mathbf{x}_t^{i\prime}\boldsymbol{\beta} + v_t^i, \tag{4}$$

where the conditioning information $(\bar{\mathbf{x}}_t, \mathbf{x}_t^i)$ is understood to include the knowledge of the date and of the identity of the household,[3] where v_t^i is a disturbance term equal to the sum of both measurement error in consumption as well as prediction error, and where the household fixed effects α^i are restricted to sum to zero.

[2] Other models for these conditional expectations are possible, of course. Since the curvature of both the functions $\{U^i\}$ and of the prediction equation for consumption will matter in measuring and decomposing vulnerability, it seems desirable to base decisions regarding these objects on some kind of objective criterion. For example, if one were to assume that household preferences exhibited constant relative risk aversion, and thought that consumption allocations were close to a full insurance optimum (Wilson, 1968) then a log-linear specification for conditional expectations might be indicated.

[3] Thus, the $\{\eta_t\}$ capture the influence of changes in aggregates, and the $\{\alpha^i\}$ capture the influence of fixed household characteristics on predicted household consumption.

3. Vulnerability in Bulgaria

In this Section we estimate vulnerability in Bulgaria, and also decompose vulnerability into distinct components. We take $U^i(c) = (c^{1-\gamma})/(1-\gamma)$, the utility function most often used in the empirical literature concerned with behaviour under risk; note that the parameter γ can be interpreted as the households' coefficient of relative risk aversion. In rough keeping with estimates of this parameter found in the microeconometric literature, we take $\gamma = 2$.[4]

The data we will use in this study are from the Household Budget Survey (HBS) collected by the Central Statistical Office of Bulgaria, and previously described by Skoufias (2001). The dataset includes monthly information on 2,287 households over one year. Here we choose to use food expenditures as our measure of consumption; for similar calculations using broader measures of consumption see Ligon and Schechter (2002). We normalise consumption so that the food expenditures of the average household are equal to one; as a consequence, if resources were allocated in such a way that there was *no* vulnerability (i.e., no risk and no inequality), then each households' consumption and utility would be equal to one.

We take the vector $\mathbf{x}_t^i$ of time-varying household characteristics to include income, employment status, and the number of pensioners. We use restricted least squares to estimate the various conditional expectations which appear in (3), and then take simple averages of households' time series to estimate each of the components of (3). Table 1 shows these results on the estimated vulnerability of our sample (using food expenditures as our measure of consumption). Note first the figure of 0.1972 in the head of the second column; this is our estimate of the vulnerability of the population. Because of our normalisation of consumption, this figure has a simple interpretation: the utility of the average household in our sample is nearly 20% less than it would be if resources could be costlessly redistributed so as to eliminate all inequality and risk in consumption.

Table 1 also reports the decomposition of vulnerability into poverty, aggregate risk, idiosyncratic risk, and unexplained risk. Poverty is the largest single component of vulnerability, accounting for just over half of all of vulnerability. Various sorts of risk account for the remaining vulnerability of households, with unexplained risk (which may include measurement error) the most important source of vulnerability. Aggregate risk is much more important than idiosyncratic risk, although the considerably larger unexplained risk may be largely made up of unobserved sources of idiosyncratic risk.

We also look at the correlates of these components of vulnerability by estimating cross-sectional regressions of each component of vulnerability on a set of fixed household characteristics; results of these regressions are found in the columns of Table 1. For household characteristics which vary over the 12-month period, we

[4] While estimates of total vulnerability, poverty, and risk are all sensitive to to one's choice of γ (or, more generally, to the shape of the functions $\{U^i\}$), it is worth noting that the relative magnitudes of these different components are less sensitive as greater concavity reflects greater welfare losses associated with both risk and inequality. For example, while setting $\gamma = 3$ roughly doubles estimated vulnerability relative to $\gamma = 2$, the share of poverty in total vulnerability falls slightly, but remains close to one half in either case.

Table 1

Correlates and Breakdown of Vulnerablility in Food Consumption

Average Value (in utils)	*Vuln* 19.7156*** [18.9191, 20.5250]	= *Pov* 10.7900*** [10.1679, 11.4264]	+ *Agg Risk* 2.6430*** [2.4574, 2.8578]	+ *Idio Risk* 0.1472*** [0.0852, 0.2210]	+ *Unexp Risk* 6.1354*** [5.7690, 6.4939]
Variable	Coef	Coef	Coef	Coef	Coef
Primary Ed.	−4.3522	−5.0335**	−0.4169	−0.1722	1.2703
	(2.6999)	(2.3834)	(0.3009)	(0.2659)	(0.8333)
Secondary Ed.	−12.9181***	−12.1466***	−1.2685***	−0.6798**	1.1769
	(3.0069)	(2.6271)	(0.3225)	(0.2735)	(0.9419)
Post-Sec. Ed.	−16.4589***	−15.5674***	−1.3940***	−0.8678***	1.3703
	(3.4492)	(3.0623)	(0.3443)	(0.2642)	(0.9771)
Male	−3.3631	−2.0415	−0.7126**	−0.1312	−0.4778
	(2.4076)	(2.0156)	(0.2796)	(0.1092)	(0.5758)
Age	0.6755*	0.6554**	0.1454***	0.1295***	−0.2548
	(0.3872)	(0.3190)	(0.0362)	(0.0429)	(0.1576)
Age Squared	−0.0048	−0.0046	−0.0012***	−0.0011***	0.0021
	(0.0034)	(0.0029)	(0.0003)	(0.0003)	(0.0013)
Own Animal	−12.0656***	−12.0769***	−0.9161***	0.1859	0.7416
	(2.3172)	(1.9641)	(0.2524)	(0.1216)	(0.6960)
Land Cultivated	−0.0237	−0.0407	0.0088	0.0051	0.0032
	(0.1993)	(0.1801)	(0.0158)	(0.0063)	(0.0398)
Urban	8.6554***	7.3798***	0.5765**	−0.0694	0.7685
	(2.2916)	(1.8868)	(0.2292)	(0.1061)	(0.7627)
No. of Pens.	−7.3809***	−5.7357***	−0.9527***	−1.1472***	0.4548
	(1.8089)	(1.5443)	(0.2373)	(0.3958)	(0.7261)
No. of Emp.	−17.5215***	−15.3678***	−1.7694***	−2.0531***	1.6688*
	(1.9280)	(1.5413)	(0.2796)	(0.4834)	(0.9345)
Fam. Size	20.8202***	19.0939***	1.7719***	1.2708***	−1.3164**
	(1.1202)	(0.9319)	(0.1979)	(0.3609)	(0.6218)
R^2	0.3996	0.3509	0.4594	0.2694	0.2710

These regressions also include province dummies. Numbers in parenthesis are bootstrapped standard errors, and those in brackets are 90% confidence intervals. *** – significant at the 1% level, ** – significant at the 5% level, * – significant at the 10% level.

use the mean value of that characteristic as our right hand side variable. Note that the correlates of vulnerability are very similar to the correlates of poverty.

We find that households with more educated heads are less vulnerable, with college educated heads being on average 16% less vulnerable than households with uneducated heads. Nearly all of this reduction is due to educated households having higher expected consumption expenditures, but these highly educated households also face significantly less aggregate and idiosyncratic risk. Households which own animals or live in villages (as opposed to cities) are also less vulnerable, mostly because of their higher consumption. Despite agriculture's reputation for being risky, agricultural households bear no more risk than other households. 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. Those households with more pensioners or workers and smaller family size experience both higher levels of and lower idiosyncratic risk in food consumption. These households experience more unexplained risk. The gender of the household head has no significant effect on vulnerability, but female-headed households bear significantly greater aggregate risk; this contrasts with the results of Glewwe and Hall (1998), who find that female headed households are no more or less vulnerable to aggregate sources of risk than are male headed households. Households with older household heads are also more vulnerable than those with younger heads.

4. Conclusion

In this paper we propose a simple measure of vulnerability. This measure is simply the difference between the utility a household would derive from consuming some particular bundle with certainty and the household's expected utility of consumption. This measure can be naturally decomposed into distinct measures of poverty, exposure to aggregate risk, exposure to idiosyncratic risk, and unexplained risk plus measurement error.

Of particular note is that by adopting a utilitarian framework we correctly capture the effects of risk on household welfare. This contrasts with some other measures of vulnerability, which work with the expected value of one of the Foster-Greer-Thorbecke poverty measures. Use of such measures as a guide to policy would tend to underestimate the value of mechanisms for reducing risk, such as credit, savings, or insurance.

Using data from Bulgaria we estimate this vulnerability measure, and its components, and also look at the correlates of each of the components of vulnerability. Our estimates suggest that the welfare of the average Bulgarian household is 11% less than it would be if there was no inequality, and an additional 3% less than it would be in if there was no aggregate risk. Idiosyncratic risk stemming from observable sources (income shocks, unemployment incidence, changes in pensions), while significant, is unimportant in terms of magnitude.

We close by discussing two possible avenues for further research. The framework we have adopted here adds uncertainty in a satisfactory way, by thinking of the functions $\{u^i\}$ as *ex post* indirect utility functions. To permit a dynamic analysis, it is straightforward to think about instead regarding this function as the value function defined by Bellman's equation (with current consumption expenditures a function of asset holdings). However, we use time series variation in households' outcomes to identify the risk that they face. This makes it much more difficult to think simultaneously about extending our measure to permit any sort of dynamic analysis; doing so would require the structure of a proper dynamic model.

A second avenue is, perhaps, more immediately practical. As just noted, our present work requires panel data on households to estimate the vulnerability of those households. Such data are, of course, both expensive and time-consuming to collect. Faced with similar problems in poverty applications, other authors (Lanjouw and Lanjouw, 2001) have tried to match households from panel datasets with households in larger cross-sections, and to extend the inferences drawn accordingly. A similar procedure seems feasible here.

University of California, Berkeley

References

Chaudhuri, S. (2001). 'Empirical methods for assessing household vulnerability to poverty', unpublished manuscript.

Chaudhuri, S., Jalan, J. and Suryahadi, A. (2001). 'Assessing household vulnerability to poverty from cross-sectional data: a methodology and estimates from Indonesia', unpublished manuscript.

Christiaensen, L. J. and Boisvert, R. N. (2000). 'On measuring household food vulnerability: case evidence from Northern Mali', unpublished manuscript.

Foster, J., Greer, J. and Thorbecke, E. (1984). 'A class of decomposable poverty measures', *Econometrica*, vol. 52 (3), pp. 761–6.

Glewwe, P. and Hall, G. (1998). 'Are some groups more vulnerable to macroeconomic shocks than others? Hypothesis tests based on panel data from Peru', *Journal of Development Economics*, vol. 56 (1), pp. 181–206.

Kamanou, G. and Morduch, J. (2001). 'Measuring vulnerability to poverty', unpublished manuscript.

Lanjouw, J. O. and Lanjouw, P. (2001). 'Welfare in villages and towns: micro-level estimation of poverty and inequality', manuscript (March).

Ligon, E. and Schechter, L. (2002). 'Measuring vulnerability: the director's cut', Paper UN/WIDER Discussion DP 2002/86.

Pritchett, L., Suryahadi, A. and Sumarto, S. (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, vol. 98 (393), pp. 1171–82.

Rothschild, M. and Stiglitz, J. E. (1970). 'Increasing risk: I. A definition', *Journal of Economic Theory*, vol. 2, pp. 225–43.

Skoufias, E. (2001). 'Risk sharing in a transition economy: evidence from monthly data in Bulgaria', unpublished manuscript.

Wilson, R. (1968). 'The theory of syndicates', *Econometrica*, vol. 36 (January), pp. 119–32.