Estimating Vulnerability

Chris Elbers and Jan Willem Gunning[1]
Free University, Amsterdam

Paper to be presented at the conference

Staying Poor:
Chronic Poverty and Development Policy
University of Manchester, April 7-9, 2003.
word count: 3988
key words: vulnerability, expected poverty, risk, Ramsey model, consumption regressions

[1] We are grateful to Ravi Kanbur, Peter Lanjouw and Erik Thorbecke for very helpful discussions. Address for correspondence: Department of Economics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands; celbers@feweb.vu.nl and jgunning@feweb.vu.nl.

Abstract

Many existing measures of vulnerability lack a theoretical basis. In this paper we propose to measure vulnerability rigorously as the welfare of a household which solves an intertemporal optimisation model under risk.In such models, in essence a stochastic version of the Ramsey model, an important part of chronic poverty may be caused by the *ex ante* response of households to risks. Our simulation results indicate that whether or not a household is to be classified as vulnerable depends strongly on the time horizon considered. We use the model to assess the accuracy of existing regression-based vulnerability measures. We find that these methods can be vastly improved by including asset measures in the regression.

1 Introduction

It has long been recognised that a substantial part of poverty in developing countries is transient rather than chronic. Baulch and Hoddinott (2000) survey thirteen panel data studies of "poverty dynamics" - movements in and out of poverty - and conclude (p. 6) that "[i]n most of the studies, the category of 'sometimes poor' is larger, sometimes by a considerable amount, than the 'always poor'." Clearly, in such circumstances it may be very misleading to identify chronic poverty on the basis of one-off survey data. A household with a permanent income well above the poverty line might appear to be poor if it was observed just after experiencing an unfavourable shock. Conversely, a household which in most years experiences poverty could be misclassified as non-poor if its income was observed just after a positive shock. Establishing the extent of transient poverty is important since chronic and transient poverty have, obviously, very different policy implications. The problem is only slightly less serious if poverty measures are based on consumption rather than income: typically capital market imperfections severely constrain a household's scope for consumption smoothing (*e.g.* Deaton, 1990).

That a household's current poverty may be a bad guide to its future prospects explains the recent emphasis in the poverty literature on vulnerability (*e.g.* World Bank, 2001), a forward-looking poverty concept. Vulnerability is often understood as the *expected* poverty of an individual, household or group. (In this paper we will restrict the concept to households.) Vulnerability is then calculated as poverty at some future date for all possible realizations of income or consumption, weighted by the probability of these outcomes. For example, if poverty is measured by the headcount then the interpretation of vulnerability as expected poverty implies that vulnerability is measured as the probability that the household will find itself below the poverty line at the specified date.

Vulnerability is the net effect of three processes. It reflects, *first*, non-stochastic poverty determinants such as the soil quality of the holding or the education of the household's members, *secondly*, the household's exposure to shocks (*e.g.* unreliable rainfall) and, *thirdly*, its ability to cope with shocks (*e.g.* through insurance).

To apply the concept of vulnerability empirically one must estimate the distribution of the household's consumption at some future date. One approach in the literature is to assume that all households face the same, stationary distribution so that the distribution can be estimated from cross-section data. A second method is to allow for inter-household heterogeneity.

Maintaining the stationarity assumption one can then estimate household-specific distributions from time series data. These methods have the virtue of simplicity but, obviously, can lead to very misleading results if the stationarity assumption or, in the first case, the assumption of homogeneity are not satisfied. A third method is to regress a household's consumption on household characteristics and measures of realized shocks (*e.g.* an illness in the household). The estimated coefficients can then be used to predict the household's poverty for a particular (essentially arbitrary) realization of shocks. Households with high predicted poverty are then considered as vulnerable.[1]

In this paper we argue that these attempts are unsatisfactory: the expected poverty concept has unappealing characteristics (*e.g.* when household welfare increases expected poverty may rise) and the regression-based methods are likely to miss a large part of the impact of risk on household welfare. We propose an alternative methodology. The key step is to estimate a structural model of the household's consumption and (dis)saving, modelled as the outcome of intertemporal optimisation under uncertainty. This ensures that the household's responses (both *ex ante* and *ex post*) to shocks are explicitly taken into account. The estimated model can then be used to derive simulation-based estimates of vulnerability on the basis of a proper welfare concept (expected discounted utility). We illustrate this method with an example, using the parameter estimates of Elbers *et al.* (2002) who estimate a stochastic Ramsey model on panel data for smallholder households in Zimbabwe. We use this example to illustrate that vulnerability can change dramatically over time (both as a result of sustained growth and as a result of adjustment to shocks) so that outcomes are quite sensitive to the choice of time horizon. These results suggest that without a structural model vulnerability measures can be seriously misleading. We also show that much of the effect of risk on the mean of the ergodic distribution of consumption reflects the *ex ante* effect. The implication is that the usual identification of chronic poverty with structural determinants and transient poverty with risk breaks down: a household can be chronically poor because its response to risk lowers consumption permanently. Policies which are effective in reducing risk or improving households' ability to cope with risk will not only reduce transient poverty, they may also succeed in reducing chronic poverty.

[1] This method, used by Dercon and Krishnan (2000), does not take into account the probability of the realization considered and thereby fails to capture the extent to which the household is exposed to shocks. Clearly, the result cannot be interpreted as a measure of *expected* poverty since no information on the distribution of future consumption is used.

We also use the simulation estimates to assess the accuracy of regression-based vulnerability measures. Our key finding is that accuracy can be greatly improved if asset ownership (in our case: livestock) is included in the regression. Regressions which relate consumption to household characteristics (such as education or household size) and (*ex post*) shock measures but not to assets can be seriously misleading in identifying vulnerable households.

The structure of the paper is as follows. In the next section we review the methodology of vulnerability measures and we propose a definition which incorporates the time dimension. In section 3 we present the stochastic Ramsey model estimated by Elbers *et al.* (2002). In section 4 we compare the vulnerability estimates generated by this model and the measures derived from commonly-used regression specifications. Section 5 concludes.

2 Vulnerability Measures

When vulnerability is defined as expected poverty (*e.g.* Christiaensen and Subbarao, 2001) it may be measured as

$$V = \int_0^z p(c,z)dF(c) \tag{1}$$

where z is a poverty line, c consumption at a specified future date, $F(c)$ the distribution of consumption at that time and $p(c,z)$ a poverty measure, *e.g.* a member of the Foster-Greer-Thorbecke class

$$p(c,z) = \frac{(\max(z-c,0))^\alpha}{z}$$

where α is a non-negative parameter . The distribution F is taken as given and reflects both the household's exposure to shocks (idiosyncratic or covariant) and its ability to cope with them. In that sense F is a reduced form. (In the next section we will relax the assumption that F is given and assume instead that only the distribution of shocks is exogenous. The distribution F is then derived endogenously.) Note that for the headcount measure ($\alpha = 0$) $V = \int_0^z F(c)$: the vulnerability measure reduces to the probability that the household will experience poverty (in the sense that $c < z$). Hence probability measures of vulnerability (used, for example, by the World Bank) can be seen as special cases of (1).[2]

[2] The World Bank defines vulnerability as "he risk today to fall below the poverty line tomorrow" (Coudouel *et al.*, 2001, p. 37).

To apply (1) one needs an estimate of the distribution F. There are several approaches. First, one can use cross-section survey data to estimate the distribution of consumption (at a point in time) *across households* and use this (for each household) as F, *i.e.* as the distribution of consumption *across states of nature*. This would be valid if consumption had reached an ergodic distribution and this distribution was the same for all households. The homogeneity assumption that observed consumption represents draws from a single distribution can be relaxed by disaggregating (Kamanou and Morduch, 2001) by *e.g.* location or educational attainment, but this shifts the problem to a lower aggregation level.

Secondly, if panel data are available F can be estimated as the distribution of consumption *across time*, for a particular household. In this case the intertemporal average $\overline{c} = (c_1 + \cdots + c_T)/T$ is considered as the permanent component of consumption and all deviations from this mean as transient. Jalan and Ravallion (2000) use this method for China and McCulloch and Baulch (2000) do so for Pakistan. This method allows for inter-household heterogeneity but, as before, imposes the assumption that F is stationary. When in fact the deterministic part of consumption follows a negative trend this methodology will underestimate vulnerability by treating low consumption levels as unlikely deviations from the intertemporal mean. (This would affect the results of McCulloch and Baulch: their data have a negative trend.) Conversely, when there is a positive trend (as in the Zimbabwe data set analysed by Elbers *et al.* (2002) which exhibits very rapid growth or in Scott's (2000) analysis of Chilean data for the period 1968-86) then vulnerability would, of course, be *over*estimated.[3]

A third method is to regress changes in consumption on household characteristics using bootstrapping to generate a distribution of shocks from the regression residuals (Kamanou and Morduch, 2001). The estimated equation can the be used to predict future consumption and vulnerability can be calculated by using the distribution around this mean.

In our view these methods are unsatisfactory for four reasons. First, they rely on strong statistical assumptions, *e.g.* homogeneity or stationarity of the distribution of consumption.

Secondly, the expected poverty concept is unattractive. For example, in the case of the Foster-Greer-Thorbecke class of poverty measures an increase

[3]Ravallion (1988) does not measure vulnerability but he considers a closely related question: the welfare cost of variability. His money-metric for this is the amount by which income - when stabilised at the intertemporal mean - would have to be reduced for poverty to be equal to its intertemporal mean. Clearly, this procedure is very similar to the second method.

in risk (in the sense of a mean-preserving spread) will *increase* expected poverty (consistent with the reduction in welfare experienced by a risk-averse household) only for $\alpha > 1$ (Ligon and Schechter, 2002, cf. the earlier results of Ravallion 1988). This rules out the two most popular members of the class: the poverty gap measure ($\alpha = 1$) would record no change when risk increased whereas the headcount ($\alpha = 0$) would show an improvement, a reduction in expected poverty. Conversely, $\alpha > 1$ implies that the degree of absolute risk aversion increases with consumption, contrary to the evidence available.

Thirdly, the methods are essentially static: they focus on expected poverty at a particular moment. This makes sense only if the household has reached an ergodic state.

Finally, a household's vulnerability can be low either because it is not exposed to large shocks or because it is able to cope effectively with shocks. One would want to distinguish between the two cases. For example, a household may achieve consumption smoothing through means which are unnecessarily costly in terms of growth. There would then be a case for intervention (providing insurance to substitute for consumption smoothing through (dis)saving of liquid assets) but the case can be identified only if the household is classified as vulnerable in spite of its consumption smoothing. The distinction can only be made if both actual and counterfactual vulnerability can be estimated, the latter for the hypothetical case where the household faces no shocks. This requires a structural model so that behavioral responses to risk are taken into account.[4] With a structural model vulnerability can be assessed separately with and without risk so that the cost of the household's coping mechanism can be estimated.

3 Simulation-based Vulnerability Estimates

Our starting point is a Ramsey model: households optimize over an infinite horizon.[5] There is a single good, used for consumption, as a productive asset and as a store of value. Household h solves:

$$\max_{c_{ht},k_{ht}} \sum_{t=0}^{\infty} E_t \beta^t u(c_{ht})$$

[4] These responses are likely to induce persistence (Morduch, 1994, Elbers *et al.*, 2002).

[5] This section draws heavily on Elbers *et al.* (2002) where the model and estimation method are described in detail.

subject to

$$\begin{aligned} k_{h,t+1} &= w_{ht} - c_{ht} \\ w_{ht} &= s^y_{ht} a_{ht} f_h(k_{ht}) + s^k_{ht} \lambda (1-\delta) k_{ht} \\ k_{h0} &= \text{given} \end{aligned}$$

where c denotes consumption, k the capital stock, w wealth on hand, u the instantaneous utility function, β a discount factor, λ a parameter which converts assets to income and δ the depreciation rate. Time periods are identified by the subscript t. We assume that $0 < \beta < 1$, that $u(c)$ is a CRRA function: $u(c) = c^\gamma$, that the production function is of the CES type:

$$f(k) = (1 + \psi(k^{-\rho} - 1))^{-1/\rho}$$

and that total factor productivity is a function of the household's size (hhsize) and the highest educational attainment of its adult members (ed):

$$a_{ht} = (\alpha_0 + \alpha_1 \text{hhsize} + \alpha_2 \text{ed}) e^{\eta_h \zeta_{ht}},$$

where η and ζ are productivity shocks.

Unlike in the original Ramsey model, the household is exposed to risk: income $af(k)$ and assets $(1-\delta)k$ are both affected by shocks: s^y, s^k. These shocks have idiosyncratic and covariant components:

$$\begin{aligned} s^y_{ht} &= (\varepsilon^r_t)^{\pi_1} \varepsilon^y_{ht} \\ s^k_{ht} &= (\varepsilon^r_t)^{\pi_2} \varepsilon^k_{ht} \end{aligned}$$

We identify the covariant shocks with rainfall (denoted by the superscript r). The distribution of $\varepsilon_{ht} = (\varepsilon^y_{h,t}, \varepsilon^k_{h,t})$ and ε^r_t are lognormal, independent of each other and across time and $\ln \varepsilon_{ht}$ has correlation matrix

$$\begin{pmatrix} a_1^2 & a_1 b_1 \\ a_1 b_1 & b_1^2 + b_2^2 \end{pmatrix}.$$

When the household decides on c_t and k_{t+1} both k_t and the realizations (s^y_{ht}, s^k_{ht}) are known. Future shocks are, of course, unknown but the household does know the distributions of these shocks.

The program can be written in recursive form as the Bellman equation:

$$V(w) = \max_{\tilde{k}} \; u(w - \tilde{k}) + \beta E V(w(\tilde{k}, \tilde{s}^y, \tilde{s}^k))$$

with associated policy function

$$\varphi(w) = \arg\max_{\tilde{k}} u(w - \tilde{k}) + \beta EV(w(\tilde{k}, \tilde{s}^y, \tilde{s}^k))$$

where k and $\tilde{k}$ denote the capital stocks at the beginning and the end of the current period and similarly s and $\tilde{s}$ denote current and future shocks. In this form the model applies to every period so that time subscripts can be suppressed. Note that the policy function φ maps the current $w(k, s^y, s^k)$ into $\tilde{k}$, next period's k. Hence the policy function can be seen as an investment function, giving the optimal value of k_{t+1} as a function of w_t.

In this stochastic form of the Ramsey model risk affects household behaviour in two ways. First, if the household perceives a change in the distribution of shocks (*e.g.* an increase in rainfall risk in the form of a mean preserving spread of the covariant shock ε_t^r) then it will, in general, adjust its policy function φ so that (for the same values of the capital stock k_t and the shocks s_t^y, s_t^k) it will choose different values of k_{t+1} (and hence c_t). This effect of a change in risk on the household's policy function we term the *ex ante* effect. There also is an *ex post* effect: the change in risk will affect the size of the realised shocks so that the optimal values of k_{t+1} and c_t (controlling for k_t) are affected, even for an unchanged policy function.

V measures the household's perceived welfare. We interpret a low value of V as vulnerability. Note that this measure does not suffer from the problems identified in the previous section: there is no need to assume homogeneity or stationarity of the distribution of consumption (indeed we need not make any assumption: the distribution is determined endogenously); since V measures welfare the vulnerability measure cannot be inconsistent with household welfare; the measure is not static but is based on utility over an infinite horizon; and, finally, the impact on V of the *ex ante* and *ex post* effects of risk can easily be identified. This involves solving the model (a) under the assumption that there is no risk; (b) under the assumption that the household correctly perceives the distribution of the shock it faces but experiences no shocks ($s^y = s^k = 0$ throughout); and (c) under the assumption that the household experiences shocks drawn from the (correctly perceived) distributions.

It is typically *not* possible to solve this model analytically. We solve the Bellman equation by iteration on a finite grid of (k, s^y, s^k) values. The grid is, of course, a discrete approximation.

Elbers *et al.* (2002) estimate this model using a pseudo simulated maximum likelihood technique on an 18-year panel data set for 158 smallholder households in Zimbabwe. In this data set there are observations on livestock holdings and we identify this variable with the capital stock k. We use their parameters estimates, shown in Tables 1 and 2. These estimates imply a fairly high elasticity of factor substitution in agricultural production: since ρ is close to -0.5, the elasticity is close to 2. They also imply a close to unitary degree of relative risk aversion so that the utility function is approximately $u(c) = \ln c$. The Zimbabwean households were exposed to massive shocks in the period 1982-2000, including a very serious drought in 1991/2.

Table 1: Production Function Estimates

	estimate	t-score	
α_0	1429	4.64	
α_1	-9.842	-0.48	household size
α_2	54.038	1.75	education
ψ	0.5315	3.06	capital share
ρ	-0.5394	-0.53	

Table 2: Other Parameters

Parameter	Estimate	Standard error	
γ	0.0082	0.2598	close to log utility
β	0.7490	0.0007	discount rate 34%
λ	0.1969	0.0006	conversion parameter
δ	0.1330	0.0008	depreciation rate
π_1	0.0327	0.0361	rain elasticity
π_2	0.0421	0.0397	π_1, π_2 correlated
a_1	0.2561	0.0006	σ of $\ln \varepsilon^y$
b_1	0.2300	0.0013	
b_2	0.1314	0.0232	
$1+\tau$	1.0398	0.0086	τ rate of tech. progress

Figure 1 shows four 50-year paths of asset ownership (K/L_e: livestock per labour in efficiency units). The sample path represents a particular (randomly drawn) series of shocks for one of the households. Note that the shocks are very large: for much of the period asset ownerships changes by 50% in one or two years. The path denoted "average" represents the mean over 100,000 such paths. This shows that in this average sense the household

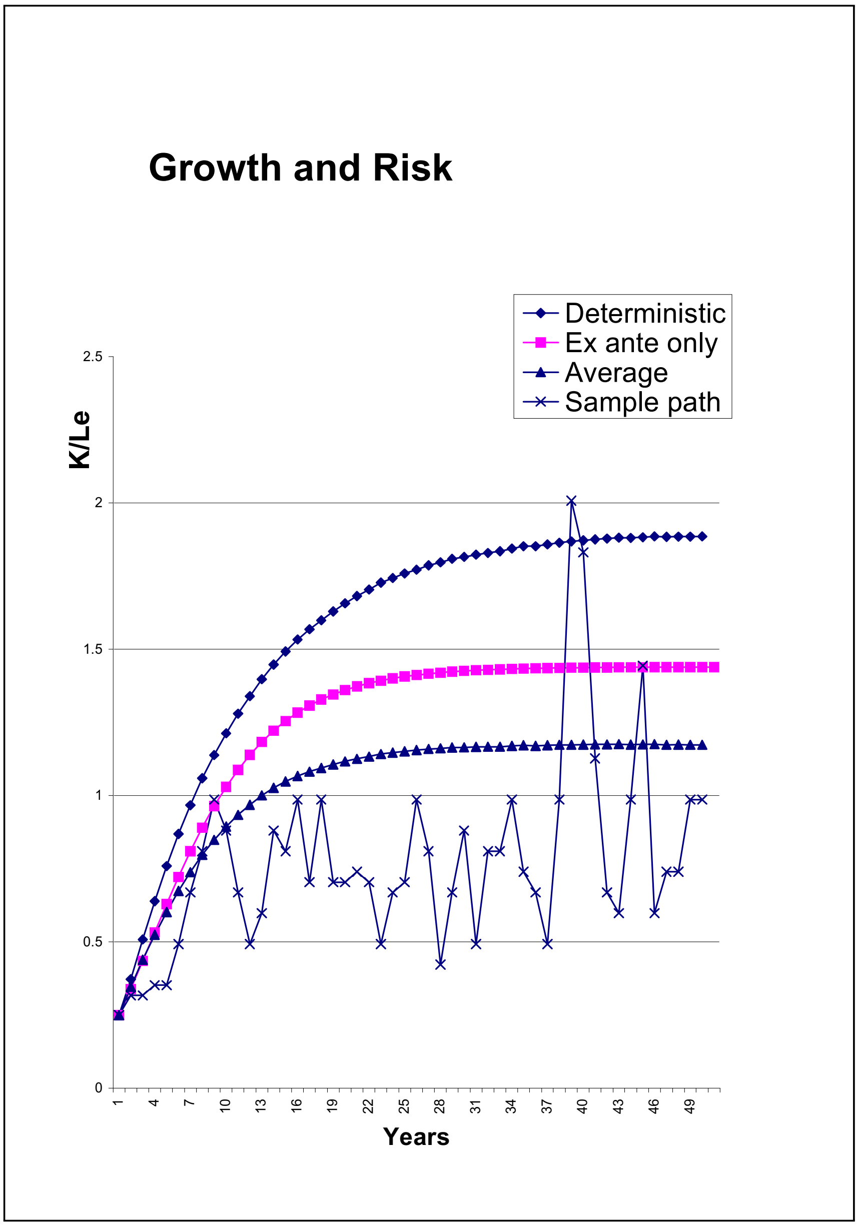

Figure 1: Growth and Risk. See text for explanation.

grows very rapidly, starting at 0.25 and reaching a level very close to the steady state value (1.2) after about 25 years. The remaining two paths show the effect of risk. This is massive: risk reduces the mean in the ergodic state from 1.9 (in the deterministic case) to 1.2. Two-thirds of this is the *ex ante* effect. Vulnerability methods which treat the mean over time of a household's consumption as the riskless counterfactual (*e.g.* Ravallion, 1988) would in this case miss most of the story: they would erroneously treat the "*ex ante* only" long run value (about 1.4) as the deterministic value (of 1.9). If measures of chronic poverty are based on mean consumption over time then a large part of chronic poverty could in fact reflect risk.

4 Comparing Vulnerability Measures

Theory - as exemplified by the stochastic Ramsey model of the previous section - implies a mapping from assets (k), shocks (s), productivity determinants (x), and characteristics of the distribution of shocks (σ) to consumption (c):

$$c_{it} = \xi(k_{it}, s_{it}, x_{it}, \sigma) \tag{2}$$

and from this mapping one can derive appropriate vulnerability measures, *e.g.* the household's expected discounted utility. However, in practice vulnerability measures are not based on equation (2) but typically on regressions of c_{it} on x_{it} and possibly on s_{it}.This approach is problematic in several ways.

First, functional form restrictions have to be imposed and these are often highly restrictive. For example, without interaction terms a linear specification makes the effect of s on c independent of x. There is no theoretical justification for such independence.

Secondly, by leaving out assets (k) the regression suffers from omitted variable bias. This is likely to be serious: if two households are observed after being hit by a negative shock and if they are identical in all respects except for the level of assets then their consumption decisions may well be very different: the household with the higher k can better afford to smooth consumption by using its assets.

Thirdly, omitted variable bias is also introduced by the exclusion from the regression of the shock characteristics σ. Recall that a change in risk affects household behaviour both *ex ante* (the policy function φ will be affected, *i.e.* the household will for the same values of (k, s^y, s^k) decide on a different level of investment and consumption) and *ex post (*since the shocks (s^y, s^k) are now drawn from different distributions). If s is included in the regression but σ is not then this *ex post* effect can in principle be identified but the *ex ante* effect will be missed. This is potentially serious: if one would like to estimate how much a policy-induced reduction in risk would contribute to welfare one needs both effects.

Finally, implicitly the regression treats the distribution of c across households (conditional on the regressors) as the distribution of shocks. This is appropriate only if the households are observed in an ergodic state, an assumption which is unlikely to be appropriate. (It certainly is not appropriate in our case. The rural households in our Zimbabwean sample were in 1983 - shortly after they were resettled - very far from the ergodic state. For example, for the sample household of Figure 1, livestock ownership was only about 20% of its ergodic mean.)

For all these reasons, vulnerability measures based on the usual consumption regressions may be wrong. However, the error need not be serious. We investigate this with a series of experiments. We take the estimated stochastic Ramsey model as the correct model and use it as the data generating process for a series of regressions (Table 3). In each case the data generated are for 1981.

The first regression relates consumption only to household-specific total factor productivity.[6] The performance of this regression is very poor. Clearly it cannot serve as a basis for identifying vulnerable households. The second regression includes initial-year livestock ownership as an additional regressor. The improvement of the fit is spectacular. Finally, in the last regression we also include the shocks experienced in that year.[7] This further improves the fit.

[6] In practice tfp is not observed and proxied by household characteristics. In our model tfp is determined by household size and education. To avoid artificial errors from choosing the wrong functional form in the regression we have used tfp as regressor.

[7] Clearly, in practice the researcher would have at best partial information on these shocks.

Table 3
Consumption regressions on simulated data

	TFP only		TFP and cattle		TFP, cattle, shocks	
	Coef.	s.e.	Coef.	s.e.	Coef.	s.e.
Constant	-19.25	17.08	-1.493	0.75	-4.494	0.47
TFP	8.325	7.04	0.5911	0.15	0.5072	0.09
TFP^2	-0.7471	0.72				
Cattle			3.635	0.12	3.433	0.07
s^y					3.163	0.37
s^k					0.1905	0.35
R^2	0.050		0.858		0.951	
Dependent variable is simulated after-shock 1981 consumption						

We now consider how such regressions can be used to identify vulnerable households. Figure 2 shows how the sample households are initially distributed over the (cattle,TFP) plane. The contour lines show combinations of productivity and cattle for which household welfare (V) is the same. The bold line separates the 50% worst off from from 50% best off. We arbitrarily consider the bottom 50% as "vulnerable". Note that the indifference curves are very steep: welfare is very sensitive to changes in initial asset ownership. This reflects the short horizon used in the present exercise. If we take a longer horizon (say 5 years) the indifference curves would become flatter. In the limit, when initial conditions are no longer relevant, the curves are horizontal.

Figure 3 shows the horizontal line which would be used to identify the bottom 50% on the basis of the first regression in Table 3. Clearly many vulnerable households would not be identified as such and vice versa. This is shown in Table 4 which shows that 40% of the households are misclassified by this regression.

Figure 4

In Figure 4 the separating line for the final regression is shown. Note that this line is very close to the 'true' 50% indifference curve. Indeed, very few households are misclassified when this regression is used. The fact that the indifference curve corresponding to the regression is steeper than the true indifference curves reflect the fact that the Ramsey program also takes into account the long-run effects of total factor productivity.

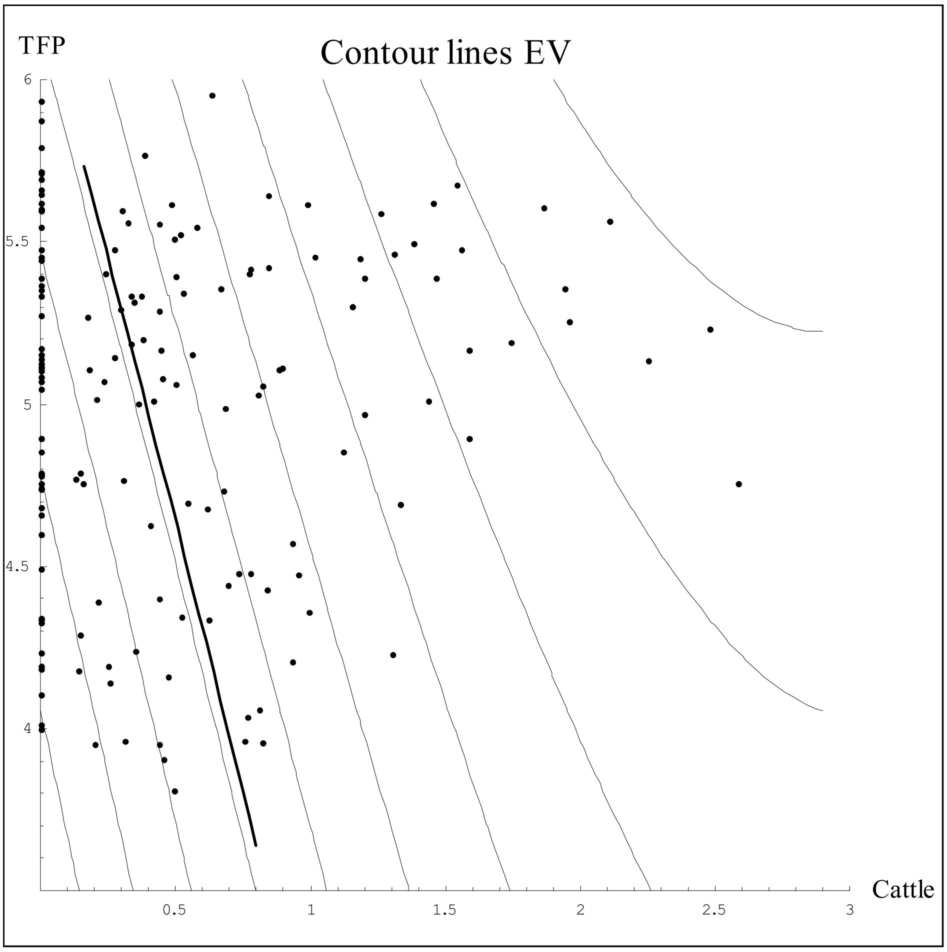

Figure 2: Expected one-year ahead program value (equation 1). Bold line separates 50% of sample. Dots indicate sample points.

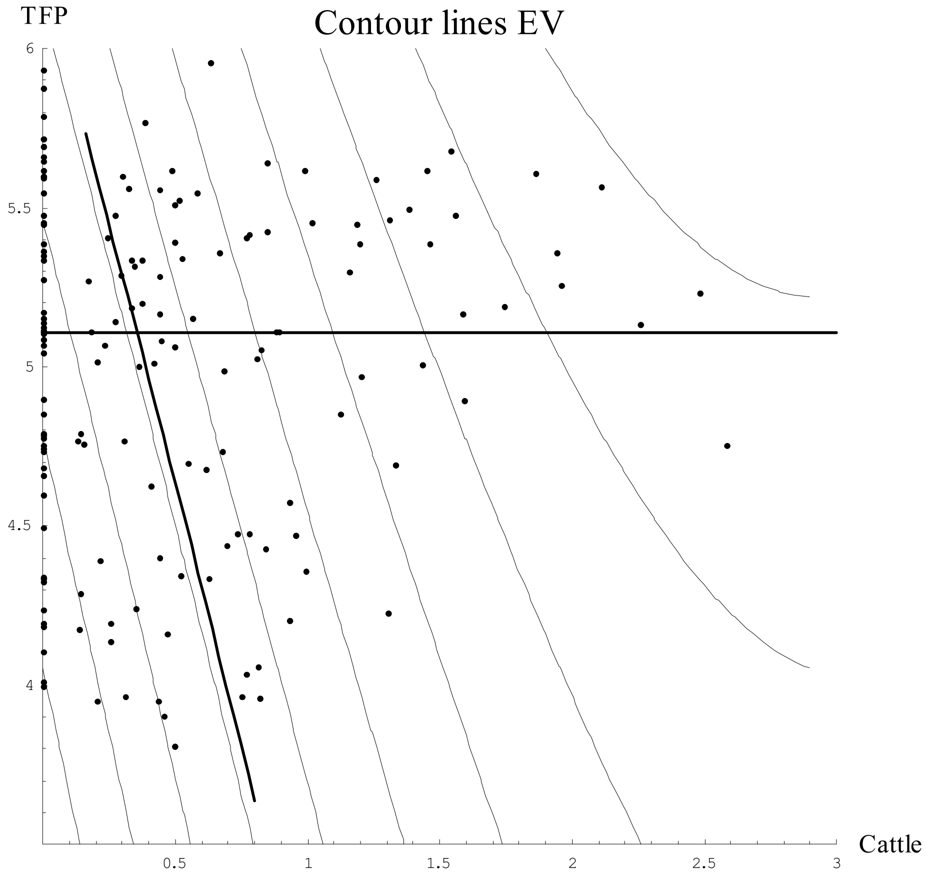

Figure 3: See also figure 2. The bold horizontal line separates 50% of the sample according to TFP differences only. (This corresponds to the first regression in Table 3.)

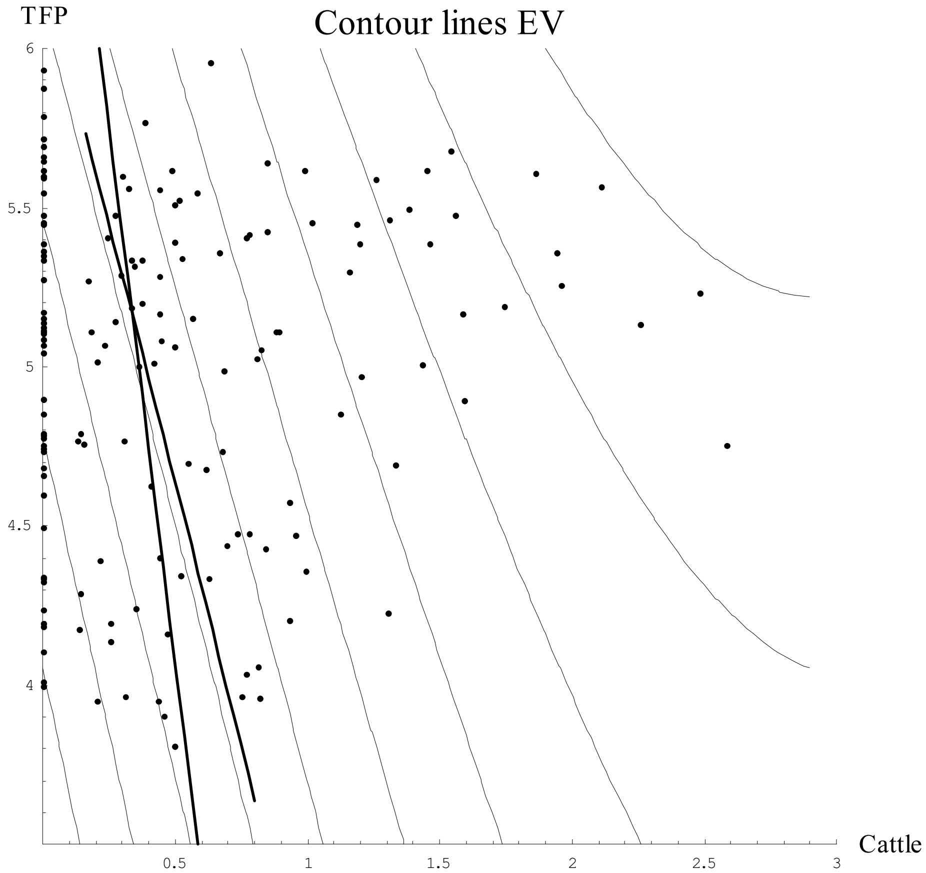

Figure 4: See also figure 2. The steeper bold line separates 50% of the sample according to TFP, cattle, and shock differences. (This corresponds to the third regression in Table 3.)

Table 4
Classifying the 50% poorest households.
Stochastic Ramsey vs. regression models

Model	TFP only	TFP and Cattle	TFP, Cattle and shocks
Correct poor	48	78	78
Correct non-poor	48	78	78
False poor	31	1	1
False non-poor	31	1	1
Total	158	158	158

'Poor' defined as below median according to the various models
'Correct' refers to classification according to the stochastic Ramsey model.

5 Conclusion

It has long been recognised that poverty measures based on cross-section data may be misleading indicators of household welfare if there is substantial unobserved heterogeneity in the sample, if households are not observed in an ergodic state and, perhaps most importantly, if they face risk. Existing vulnerability measures try to incorporate the effect of risk on welfare. While easy to apply they lack a theoretical basis. We have argued that vulnerability can be measured rigorously as the welfare of a household which solves an intertemporal optimisation model under risk. Using such a model (a stochastic Ramsey model estimated on panel data for smallholders in Zimbabwe) we showed that failing to distinguish between the *ex ante* and *ex post* effects of risk may lead to large errors in estimates of chronic and transient poverty.

Our analysis makes clear that vulnerability depends on the time horizon considered. In particular, if one takes a longer term perspective, vulnerability is less sensitive to initial conditions and, conversely, more sensitive to permanent productivity differentials.

We used the estimated model to assess the accuracy of existing methods in identifying vulnerable households. Regression-based methods using proxies for total factor productivity (*e.g.* education or soil quality) or measures of shocks experienced by the household as regressors can be very misleading. Our results show that a vast improvement can be obtained by including asset measures in the regression. With this amendment simple, regression-based methods can closely approximate the true model.

References

[1] Baulch, Bob and John Hoddinott (2000), 'Economic Mobility and Poverty Dynamics in Developing Countries', *Journal of Development Studies*, vol. 36, pp. 1-24.

[2] Christiaensen, Luc J. and Kalanidhi Subbarao (2001), 'Towards and Understanding of Vulnerability in Rural Kenya', mimeo, World Bank.

[3] Coudouel, A., J. Hentschel and Q. Wodon (2001), 'Well-being Measurement and Analysis', mimeo, World Bank.

[4] Deaton, Angus (1990), 'Saving in Deverloping Countries: Theory and Review', *World Bank Economic Review, Supplement: Proceedings of the 1989 ABCDE*, pp. 61-96.

[5] Dercon, Stefan and Pramila Krishnan (2000), 'Vulnerability, Seasonality and Poverty in Rural Ethiopia', *Journal of Development Studies*, vol. 36, pp. 25-51.

[6] Elbers, Chris, Jan Willem Gunning and Bill Kinsey (2002), 'Convergence, Growth and Poverty: Micro Evidence of Growth under Uncertainty', Tinbergen Institute Discussion Paper 02/52.

[7] Jalan, Jyotsna and Martin Ravallion (1998), 'Transient Poverty in Post-Reform Rural China', *Journal of Comparative Economics*, vol. 26, pp. 338-357.

[8] Jalan, Jyotsna and Martin Ravallion (2000), 'Is Transient Poverty Different? Evidence for Rural China', *Journal of Development Studies*, vol. 36, pp. 82-99.

[9] Kamanou, Gisele and Jonathan Morduch (2001), 'Measuring Vulnerability to Poverty', mimeo.

[10] Ligon, Ethan and Laura Schechter (2002), 'Measuring Vulnerability', mimeo.

[11] McCulloch, Neil and Bob Baulch (2000), 'Simulating the Impact of Policy upon Chronic and Transient Poverty in Rural Pakistan', *Journal of Development Studies*, vol. 36, pp. 100-130.

[12] Morduch, Jonathan (1994), 'Poverty and Vulnerablity', *American Economic Review, Papers and Proceedings*, vol. 84, pp. 221-225.

[13] Pritchett, Lant, Asep Suryahadi and Sudarno Sumarto (2000), 'Quantifying Vulnerability to Poverty: a Proposed Measure, Applied to Indonesia', mimeo, Social Monitoring and Early Response Unit (SMERU), Jakarta.

[14] Ravallion, Martin (1988), 'Expected Poverty under Risk-Induced Welfare Variability', *Economic Journal*, vol. 98, pp. 1171-1182.

[15] Scott, Christopher D. (2000), 'Mixed Fortunes: a Study of Poverty Mobility among Small Farm Households in Chile, 1968-86', *Journal of Development Studies*, vol. 36, pp. 155-180.

[16] World Bank (2001), *World Development Report 2000/2001: Attacking Poverty*, Washington, DC: World Bank.