ESTIMATING VULNERABILITY: STATIC WELFARE MEASURES IN A DYNAMIC WORLD

ETHAN LIGON

Abstract. Recent research on household 'vulnerability' has led to an increased appreciation of the welfare costs of risk. Measuring the risk borne by a particular household has generally involved the use of panel data, and in particular the use of time series variation in household expenditures to estimate the risk borne by the household in any given period. This has led researchers to focus on static measures of vulnerability, since once used to identify the distribution of consumption expenditures in a single period the time series variation can no longer be used to describe the intertemporal profile of the distribution of consumption expenditures—simultaneously estimating inequality, risk, and time series variation in household vulnerability requires the additional structure of a dynamic model. Unfortunately, our present understanding of the economic circumstances in which most poor households are situated seems too limited to permit general agreement on what the *right* dynamic model is. While research on the development, estimation, and testing of such models ought to be of the highest priority, here we argue that because there is considerable evidence that households do their best to smooth consumption over both dates and states. Accordingly, a static vulnerability measure which correctly estimates the distribution of consumption expenditures at a point in time can, in many circumstances, also capture the welfare costs of intertemporal variation.

1. Introduction

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 utilility *ex ante* when those same households face greater variation in future consumption (for a recent survey see, e.g., Hoddinott and Quisumbing, 2003).

Measuring the risk borne by a particular household has typically involved the use of panel data. In particular most approaches to estimating the risk borne by the household in any given period have relied on the use of time series variation in household expenditures.

Date: October 8, 2003.

This has led most researchers to focus on static measures of vulnerability, since once used to identify the distribution of consumption expenditures in a single period the time series variation can no longer be used to describe the intertemporal profile of the distribution of consumption expenditures—simultaneously estimating inequality, risk, and time series variation in household vulnerability requires the additional structure of a dynamic model. In a recent paper, Elbers and Gunning (2003) specify a stochastic dynamic model precisely in order to be able to describe the trajectories of vulnerability for sample households in Zimbabwe.

Unfortunately, our present understanding of the economic circumstances in which most poor households are situated seems too limited to permit general agreement on what the *right* dynamic model is. While research on the development, estimation, and testing of such models ought to be of the highest priority, here we argue that because households do their best to smooth consumption over dates and states, one can exploit this fact to estimate *ex ante* measures of vulnerability even in the absence of a fully specified dynamic model.

Accordingly, a static vulnerability measure which correctly estimates the distribution of consumption expenditures at a point in time can, in many circumstances, also capture the welfare costs of intertemporal variation.

2. Measuring Vulnerability

Here we begin by specifying a simple model, describing the problem facing the household. While the model itself is quite special, we'll seek to use the model to illustrate features of household behavior which would also obtain in a fairly wide class of models.

We'll begin the process of modeling household behavior by supposing that a particular household has von Neumann-Morgenstern preferences defined over a single consumption good in each of many periods, so that the households' expected utility in period t is given by

$$\int u(c_t) dF_t(c_t),$$

where $u(c)$ is the household's momentary utility given a consumption realization c, and where F_t is the distribution of consumption for the household at time t. This distribution, of course, may depend on actions taken by the household—in particular, savings decisions made in earlier period will help to determine F_t. Following Ligon and Schechter

(2003) we define the *vulnerability* of the household at t by

$$V_t(z) = u(z) - \int u(c_t)dF_t(c_t),$$

where z may be regarded as analogous to a poverty line level of consumption expenditures. It's worth noting that this expression may be re-written as

$$V_t(z) = \left[u(z) - u(\int c_t dF_t(c_t))\right] + \left[u(\int c_t dF_t(c_t)) - \int u(c_t)dF_t(c_t)\right],$$

where the first bracketed term may be interpreted as the *poverty* of the household, and the second as the *risk* borne by the household. As F_t is endogenous, the second term *should not* be interpreted as the welfare improvement to be had from eliminating all risk, since this sort of change in the environment will generally lead to differences in household behavior—for example, elimination of future risk would eliminate precautionary motives for saving, and so might increase future poverty. Rather, levles of vulnerability, poverty, and risk are what is borne by the household *after* one takes into account whatever strategems the household has employed to improve its welfare.

Now suppose that we are faced with the problem of taking explicit account of the fact that forward looking households will care not only about their vulnerability in period t, but at all future dates. How ought we to calculate the welfare consequences of future risk, and of time series variation in levels of consumption? Let us suppose that we can represent the household's problem as a dynamic program. Let x denote a vector of state variables, and suppose that the economic environment is stationary, so that *conditional* on this set of state variables the household solves

$$W(x_1) = \max_{(c_1,x_2)\in\Gamma(x_1)} u(c_1) + \beta \mathrm{E}_1 W(x_2 + \epsilon). \tag{1}$$

Here E_1 denotes the expectations operator, conditioning on x_1, $\Gamma(x_1)$ denotes the set of feasible choices of consumption and future state variables (c_1, x_2) given the current state x_1, and W is the household's value function. Note that future values of the state depend on the realization of the shock ϵ; it is this uncertainty which necessitates taking expectations. The distribution of ϵ may depend on both x_1 and on the choices made by the household.

A few things are worth noting about this problem, before specializing it further. The first is simply that provided that $\Gamma(x)$ is compact and u increasing, concave, and bounded, then a single-valued solution to the households problem exists, and we can write the household's consumption as a function of the current state, $c_1 = g(x_1)$. Closely related, the state variable $x_2 = h(x_1)$ where ϵ is a shock realized after. If x_1 is a scalar (e.g., household wealth), then under additional standard conditions g and h will be monotone, so that for the econometrician observing the household's behavior, observing x_1 provides no information beyond that provided by observing c_1; indeed, we can define $\tilde{W}(c) = W(g^{-1}(c))$, so that the household's behavior satisfies the recursion

$$\tilde{W}(c_1) = u(c_1) + \beta \mathrm{E}_1 \tilde{W}(c_2).$$

Accordingly, if one knows enough about the household's problem, then observing consumption suffices to predict future trajectories of consumption, utility, and vulnerability.

It should also be clear that to capture the welfare of the household, any measure of vulnerability ought to be forward-looking. In particular, letting $V(z)$ denote our forward looking measure of vulnerability, we may write

$$V(z) = \mathrm{E}_0[u(z) - (1-\beta)\tilde{W}(c_1)] = (1-\beta)\mathrm{E}_0 \sum_{t=1}^{\infty} \beta^{t-1} V_t(z).$$

Note that this measure, being *ex ante*, should not depend on realizations of V_t, but only on the household's time 0 *expectations* of these realizations. The contribution of variation in consumption at time t to *ex ante* vulnerability is *not* $V_t(z) = u(z) - \int u(c_t) dF_t(c_t)$, but the time zero expectation of this quantity, $\mathrm{E}_0 V_t(z) = u(z) - \mathrm{E}_0 \int u(c_t) dF_t(c_t)$. Thus, once one knows the solution to the household's problem, observing a sequence of realizations of state variables such as assets, cattle, or productivity shocks will *not* aid one in computing the welfare loss associated with risk.

3. Stationarity

Computing measures of vulnerability requires one to estimate the distribution of consumption. One approach to doing so is illustrated by Ligon and Schechter (2003), who simply use the empirical distribution of consumption expenditures over time for a given household as a proxy for the variation the household might expect in a single period. This is reasonable if the distribution of consumption expenditures is stationary, so that realizations of consumption expenditures over time may be regarded as draws from the same time-invariant distribution. Accordingly, it behooves us to consider the circumstances under which the distribution F_t will indeed not vary over time; and further, to describe alternative ways of estimating the distribution of consumption at t when the distribution of consumption is non-stationary.

Let us begin with the description of a simple example environment. There's a population of household indexed by $i = 1, \dots, n$, each with preferences over consumption given by $u^i(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, so that utility from leisure is additively separable from utility from consumption. At date $t = 1, \dots, T$ some state $s_t \in \mathcal{S} = \{1, 2, \dots, S\}$ is realized; the history of these states is denoted by h_t (with h_0 the null set). Thus, at date t each household i produces output $y_i(h_t)$. The population may collectively save or borrow, leaving a total quantity of $x(h_t)$ of the consumption good to be allocated across households at date t.

To describe the efficient allocation of resources across these households we'll find it convenient to compute allocations that a central planner would choose to implement; the planner will seek to maximize the objective function

$$\sum_{i=1}^{n} \lambda_i \sum_{t=1}^{T} \beta^{t-1} \sum_{h_t} \pi(h_t|h_{t-1}) u_i(c_i(h_t) \tag{2}$$

subject to respecting aggregate resource constraints

$$\sum_{i=1}^{n} c_i(h_t) \leq x(h_t) \tag{3}$$

for all histories h_t. This resource constraint will invariably constrain the planner, though since sometimes (as when the planner has access to outside credit markets, or when there's some other means of moving the consumption good across periods or states) $x(h_t)$ will be an endogenous variable. Of course the planner may also face other constraints, and these other constraints will help to determine the evolution of the distribution of household consumption exenditures. We consider some interesting special cases in turn.

Full insurance with no aggregate variation: Since risk averse households will actively try to smooth their consumption over both dates and states, the requirement that consumption should be stationary is more reasonable that it might otherwise appear. For example, if all households have access to adequate insurance and there's no source of aggregate risk, then the distribution of consumption expenditures will be stationary even when the distribution of, e.g., individual income is not. Under these circumstances using time series variation in levels of consumption to estimate risk at a point in time will be valid. In terms of the model we've begun to lay out above, our assumption that there's no aggregate variation means that $x(h_t) = x$ for all h_t; then from the first order condition to the planner's problem for $c_i(h_t)$ we have

$$u_i'(c_i(h_t)) = \mu/\lambda_i$$

where mu is the multiplier on the resource constraint (in this case, though there are many constraints, they're all identical, so we need only a single multiplier) for all $i = 1, \ldots, n$, and for all h_t—accordingly, household i's consumption will be a constant, which doesn't depend on either the state of the world or the date. As a consequence, the distribution of future consumption expenditures is (trivially) stationary. If we imagine that consumption is measured with some error, and that the distribution of this error is also stationary, then simply using the empirical distribution of household consumption over time to stand in for variation at a point in time is perfectly sound.

Full insurance; stationary aggregate shocks: Let us now extend the case above to make it somewhat more realistic. Here we imagine that $x(h_t)$ is random, but has a stationary distribution. This might be the case, for example, for an isolated village or closed economy, having a stationary distribution of income. In this case, the first order conditions to planner's problem becomes

$$u_i'(c_i(h_t)) = \mu(h_t)/\lambda_i.$$

Here $\mu(h_t)$ is the multiplier associated with the resource constraint for history h_t; since this depends on aggregate resources, which are now variable, household consumption will vary with shocks to the aggregate endowment. Exploiting our assumption of CRRA preferences, we can re-write this first order condition as

$$\log c_i(h_t) = \alpha_i + \eta(h_t),$$

where $\alpha_i = -\gamma^{-1} \log \lambda_i$ and $\eta(h_t) = \gamma^{-1} \log \mu(h_t)$. From here, if we add a multiplicative measurement error $e^{-\epsilon_{it}}$ to household i's time t consumption (assuming, as before, that the distribution of this measurement error is stationary), then it's a very short step to a regression of the form

$$\log c_{it} = \alpha_i + \eta_t + \epsilon_{it},$$

where the t subscripts now denote time t realizations of the variables they adorn; this is essentially the regression employed by Deaton (1992) to characterize full insurance allocations. Since the distributions of $x(h_t)$ and $\epsilon_i(h_t)$ are stationary by assumption, it follows that the distribution of $c_i(h_t)$ is also stationary.

Limited insurance; access to credit markets: In this case, suppose that each household has access to credit markets with returns $R = 1/\beta$. In this case, the Euler equation for each household's consumption expenditures is given by

$$u'(c_i(h_t)) = \mathrm{E}_t u'(c_i(h_{t+j}).$$

This is enough for us to see that the distribution of consumption in this case is non-stationary so long as there's any variation at all, and suggests that our earlier strategy of simply using the empirical distribution of consumption over time isn't appropriate in this environment. This also implies, of course, that (from the standpoint of earlier periods) the distribution of i's marginal utility at time t is the same as the distribution of time t *predictions* of marginal utility at all dates subsequent to t. We also have

$$u'(c_i(h_{t+j})) - u'(c_i(h_t)) = u'(c_i(h_{t+j})) - \mathrm{E}_t u'(c_i(h_{t+j});$$

This just says that the change in marginal utility between t and $t+j$ is equal to the time t forecast error of marginal utility at $t+j$. There is, unfortunately, no obvious, easy way to guarantee that these forecast errors are generally stationary, but at least the theory isn't in obvious conflict with this assumption. Accordingly, in this case we might proceed by using panel data and calculating the empirical distributions of 1-step ahead forecast errors for each household, and then interpreting this as the distribution of possible innovations at time t. As an important bonus, if we don't have strong priors regarding the value of γ, we can use the moment restriction $\mathrm{E}_t c_{it+1}^{-\gamma} = c_{it}^{-\gamma}$ to simultaneously estimate this preference parameter. Once we have estimates of γ and of the distribution of one-step ahead forecast errors, we can calculate our measure of vulnerability without further difficulty.

Note an important feature of this case. If a household has a rapid decrease in marginal utility over time (increase in consumption), we would interpret this as evidence that the household faced a great deal of risk; this would have the consequence of *increasing* our estimate of the household's vulnerability.

References

Deaton, A. (1992). *Understanding Consumption*. Oxford: Clarendon Press.

Elbers, C. and J. W. Gunning (2003). Vulnerability in a stochastic dynamic model. Discussion Paper TI 2003-070/2, Tinbergen Institute, Vrije Universiteit, Amsterdam.

Hoddinott, J. and A. Quisumbing (2003). Methods for microeconometric risk and vulnerability assessments: A review with empirical examples. Unpublished Ms.

Ligon, E. and L. Schechter (2003). Measuring vulnerability. *Economic Journal 113*(486), C95–C102.

Stanford University and the Department of Agricultural and Resource Economics, University of California, Berkeley