



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

Papers downloaded from AgEcon Search may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

University of California, Berkeley
Department of Agricultural &
Resource Economics

CUDARE Working Papers

Year 2011

Paper 1112

Dynamics, risk, and vulnerability

Ethan Ligon

DYNAMICS, RISK, AND VULNERABILITY

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—simultaneous estimation of 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 households are situated seems too limited to permit general agreement on what the *right* dynamic model is. We show that simple restrictions on households’ intertemporal smoothing can be used to simultaneously estimate household risk preferences in a manner which is robust to a variety of different assumptions about the economic environment. Further, these simple restrictions and estimated preferences can then be used to robustly characterize the welfare costs of different sorts of variation in consumption expenditures.

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 utility *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 (for an evaluation of several approaches, see Ligon and Schechter, 2004).

Date: February 21, 2007.

This has led most researchers to focus on static measures of risk and 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. For example, 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 for any given environment. While research on the development, estimation, and testing of such models ought to be of the highest priority, we argue that because households do their best to smooth consumption over dates and states, this places Euler-type restrictions on the evolution of household consumption over time. One can exploit these restrictions to estimate *ex ante* measures of dynamic vulnerability even in the absence of a fully specified dynamic model. Perhaps surprisingly, the data requirements necessary for estimating dynamic measures need not be any greater than the data required to estimate static measures of vulnerability.

The rest of this paper proceeds as follows. In Section 2 we discuss different sources of variation which may be observed in consumption expenditures, and observe that in general identifying the importance of any one of these sources requires identifying all of them. However, even if one has a complete characterization of variation in consumption, to evaluate the welfare consequences of this variation one needs to take a stand on household preferences. Accordingly, in Section 3 we discuss the measurement of vulnerability when households have time-separable von Neumann-Morgenstern preferences, and relate the resulting measure of dynamic vulnerability with some of the static counterparts which have appeared in the literature. Section 4 develops empirical restrictions for a variety of possible economic environments, while Section 5 presents a class of GMM estimators which may be used to simultaneously estimate preference parameters and identify different sources of consumption variation. Section 6 describes the panel dataset we use for illustration, while Section 7 presents a characterization of dynamic vulnerability for the households in our data.

2. CHARACTERIZING VARIATION IN CONSUMPTION

In a typical household panel dataset we observe realizations of consumption for each household in the panel for a sequence of periods. We wish to distinguish three different kinds of variation in these data.

First is variation *across households*. This is the sort of variation that typically interests researchers who measure inequality or poverty—measures that attempt to capture differences in the welfare of different households, or in the distribution of wealth. Though a household’s position in the wealth distribution may change over time, this is an *ex post* measure of variation, and so is *not* meant to capture random variation in consumption expenditures. Whether there’s a social cost associated with this sort of non-random cross-sectional variation is ultimately an ethical question. For example, the usual Pareto criterion doesn’t support the notion that less cross-sectional variation is preferable to more. However, the cardinal utility criterion (perhaps most clearly expressed by Harsanyi (1955)) does permit one to argue for more equality in the distribution of consumption rather than less.

Second is *predictable* time-series variation in a given household’s consumption expenditures—this is just the sequence of future expected consumptions conditional on information available *ex ante*. Think of efforts to measure life-cycle variation in consumption, or of efforts to measure the persistence of innovations to consumption. As this last example suggests, this variation may be stochastic, but must be predictable using information revealed to the household over time. As the usual models of preferences suggest that households would prefer to smooth consumption over time, this sort of variation more clearly involves a real welfare cost, of sort which might be addressed via debt markets.

Third is *unpredictable* time-series variation in a given household’s consumption expenditures—this is simply equal to realized consumption expenditures minus the household’s forecast. It’s the welfare costs associated with this last source of variation which we’ll label “risk.”

It may be worth emphasizing the obvious: even for households with time-separable von Neumann-Morgenstern utility functions, the source of variation in consumption is *not* a matter of indifference. Consider a simple example economy of two households, in which realized consumptions in each of two periods are given by

Period	Household	
	1	2
1	1	2
2	2	3

Now consider two different economic environments which might have have generated these data. In both environments households' happiness depends on expected utility, averaged over time (no discounting). Each household has a logarithmic utility function, so that

$$W_i = E[\log(c_{i1}) + \log(c_{i2})],$$

where E denotes the expectation operator, c_{it} denotes household i 's consumption in period t , and W_i denotes household i 's ($i = 1, 2$) *ex ante* welfare. Let W_0 be the average utility of all households (alternatively, the expected utility of a household behind a Rawlsian “veil of ignorance”).

In the first environment (call it A), there's no uncertainty, so that $W_1 = \log(1) + \log(2) = \log(2)$, $W_2 = \log(2) + \log(3) = \log(6)$, and $W_0 = \log(12)/2$. In the second environment (B) consumption is a random variable, with $c_{it} \in \{1, 2, 3\}$, with probabilities (respectively) of $(0.25, 0.5, 0.25)$. Note that these probabilities match exactly the empirical distribution of consumption in the supposed data. However, in environment (B) expected utility $W_0 = W_1 = W_2 = \log(2)/2 + \log(3)/4 = \log(12)/4$, or only half the average welfare in the environment with no uncertainty.

3. MEASURING RISK AND VULNERABILITY

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

$$V_t = u(\bar{c}) - \int u(c_t) dF_t(c_t),$$

where \bar{c} is per capita consumption expenditures. It's worth noting that this expression may be re-written as

$$V_t = \left[u(\bar{c}) - u \left(\int c_t dF_t(c_t) \right) \right] + \left[u \left(\int c_t dF_t(c_t) \right) - \int u(c_t) dF_t(c_t) \right],$$

where the first bracketed term may be interpreted as a measure of the position of the household in the wealth distribution, and the second as the *risk* borne by the household. As F_t may be endogenous (for example, it will generally depend on both past and contemporaneous savings decisions), 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, levels 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 recursively. Let $x \in X$ denote a vector of state variables, and suppose that the household's problem of maximizing a discounted stream of expected utility can be represented as a dynamic program, with value function satisfying

$$(1) \quad W(x) = \max_{(c, x') \in \Gamma(x)} u(c) + \beta E[W(x')|x],$$

where $E[\cdot|x]$ denotes the expectations operator conditioning on the state variables x . The variable c is the consumption chosen by the household, subject to the requirement that $(c, x') \in \Gamma(x)$; note that $\Gamma(x)$ may not be observed by the researcher. However, so long as $\beta \in (0, 1)$, the set $\Gamma(x)$ is compact and convex for all x , u is strictly concave and bounded on the image of Γ , then the principle of the maximum implies that the consumption chosen by the household will be single-valued functions of x , which we write as $c(x)$ —thus, the realization of x simultaneously determines the momentary utility $u(c(x))$ of the household, and the distribution of next period's state variables x' . Accordingly, we define our new dynamic vulnerability measure by

$$(2) \quad V(x) = u(\bar{c}) - (1 - \beta)W(x),$$

or, recursively, by

$$(3) \quad V(x) = [u(\bar{c}) - u(c)] + \beta E[V(x')|x],$$

where now we can interpret the first bracketed term as a measure of poverty in the current period, and the second term as a combination of the welfare loss from intertemporal variation and risk. Fixing the current date to be t , recursive substitution into this expression allows us to decompose vulnerability,

$$(4) \quad \begin{aligned} V(x_t) &= [u(\bar{c}) - u(c_t)] \\ &+ (1 - \beta) \sum_{j=1}^{\infty} \beta^j [u(c_t) - u(E_t c_{t+j})] \\ &+ (1 - \beta) \sum_{j=1}^{\infty} \beta^j [u(E_t c_{t+j}) - E_t u(c_{t+j})], \end{aligned}$$

where the first summation is the welfare loss associated with intertemporal variation, and the second summation our new measure of risk. When averaged over a population, the three lines of (4) correspond to the different sources of variation identified in Section 2; in particular, the first line gives a measure of the welfare loss associated with inequality in the *ex post* allocation of consumption in period t (given Harsanyi's equally weighted utilitarian social welfare function, and supposing this inequality to persist forever). The second line measures the welfare loss associated with predictable variation in consumption expenditures—note that if future *expected* consumption is equal to c_t in every future period then the contribution of this term to total vulnerability is zero. The third line captures the welfare loss due to risk, in manner which follows from the treatment of risk in Rothschild and Stiglitz (1970). This measure of risk may, of course, be further decomposed as in Ligon and Schechter (2003) into, e.g., aggregate and idiosyncratic sources of risk.

4. EMPIRICAL RESTRICTIONS

If we knew the distribution of consumption expenditures in every period and knew household preferences, then we could simply integrate to compute the measures of vulnerability, inequality, welfare loss due to intertemporal variation, and risk. In actual application, computing measures of vulnerability ordinarily requires us 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.¹ 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 households indexed by $i = 1, \dots, n$, each with preferences over consumption given by $u_i(c) = A_i \frac{c^{1-\gamma}-1}{1-\gamma}$. Each household i may also take actions a_{it} which influence production at time t at a utility cost of $v_{it}(a)$, but we assume that utility from leisure is additively separable from utility from consumption. Households discount future utility at a rate $1/\beta - 1$. At date $t = 1, \dots, T$ some state $s_t \in S = \{1, 2, \dots, m\}$ is realized; the history of these states to t is denoted by s^t (with s^0 the null set), and the set of histories to date t is denoted S^t . Where no confusion results, we denote the realization of variables indexed by history s^t by using a t subscript; e.g., $c_{it} \equiv c_i(s^t)$. Thus, at date t each household i produces output $y_i(s^t)$. The population may collectively save or borrow at an interest rate $R(s^t) - 1$.

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

$$(5) \quad E_0 \sum_{i=1}^n \lambda_i \sum_{t=1}^T \sum_{s^t} \pi(s^t | s^{t-1}) \beta^{t-1} [u_i(c_i(s^t)) - v(a_{it})]$$

for some set of positive planning weights $\{\lambda_i\}$ subject to respecting aggregate resource constraints

$$(6) \quad \sum_{i=1}^n c_i(s^t) \leq \sum_{i=1}^n (y_i(s^t) + b_i(s^{t-1}) - b_i(s^t)/R(s^t))$$

for all histories $s^t \in S^t$, $t = 1, \dots, T$. 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

¹Chaudhuri et al. (2001) use a related strategy, but instead of using the time series of consumption expenditures to describe the distribution in any period, assume that this time series is drawn from a stationary log-normal distribution, and use each household's time series to estimate the mean and variance of this parametric distribution.

expenditures. By observing panel data on consumption expenditures, the analyst may be able to draw some inference regarding the nature of these constraints, and thus to estimate the distribution of consumption expenditures in future periods. We assume, however, that the analyst observes only an error-ridden measure of consumption, $\tilde{c}_{it} = c_{it}\nu_{it}$, with ν_{it} a measurement error process having finite second moments.

We consider some interesting special cases in turn, generally proceeding from more restrictive to less restrictive assumptions regarding the economic environment, bearing in mind that estimators based on less restrictive assumptions while possible more robust, will tend to yield less precise estimates and to allow only weaker inference.

Full insurance: 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 than it might otherwise appear. For example, if all households are able to pool their risks, then differences in the observed log marginal utility of consumption across households 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 may be valid. In terms of the model we've begun to lay out above, we have the first order condition to the planner's problem for $c_i(s^t)$,

$$(7) \quad \beta^t u'_i(c_i(s^t)) = \pi(s^t) \mu(s^t) / \lambda_i,$$

for all $i = 1, \dots, n$, and for all s^t , where $\mu(s^t)\pi(s^t)$ is the multiplier on the resource constraint after history s^t . The first order condition with respect to $b_i(s^t)$ yields a relationship between the price of consumption and returns $R(s^t)$,

$$1/R(s^t) = \sum_{\{s^{t+1}|s^t\}} \pi(s^{t+1}|s^t) \frac{\mu(s^{t+1})}{\mu(s^t)}.$$

Combining first order conditions then yields

$$(8) \quad \frac{u'_i(c_i(s^t))}{\beta u'_i(c_i(s^{t+1}))} = R(s^t)$$

for all $i = 1, 2, \dots, n$ and for all histories s^t . Note that the probabilities of different histories $\pi(s^t)$ have cancelled out of this expression, reflecting the absence of idiosyncratic risk also seen in (7). Accordingly, any variation in household i 's consumption depends on the aggregate quantity $\beta R(s^t)$.

Exploiting our assumption of CRRA utility, (7) can be rewritten as

$$\log c_i(s^t) = \alpha_i + \zeta(s^t),$$

where $\alpha_i = -\gamma^{-1} \log \lambda_i$ and $\zeta(s^t) = \gamma^{-1} \log \mu(s^t)$. Recalling our assumption there may be a multiplicative measurement error $\nu_i(s^t)$ associated with 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

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

where $\epsilon_{it} = -\log \nu_i(s_t)$, and 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 distribution of ϵ_{it} is stationary by assumption, it follows that the distribution of $\log c_i(h_t) - \zeta_t$ is also stationary over time, and assuming that the econometrician knows households' preferences, then the techniques of Ligon and Schechter (2004) may be employed to estimate vulnerability in any particular period.

To test the hypothesis that households' consumption is fully insured, we use an exclusion restriction, as in Townsend (1994). In particular, since the measurement error process is mean-independent of other variables in the contemporaneous information set, this suggests the moment restriction

$$(9) \quad E(\log c_{it} - \alpha_i - \zeta_t | s^t) = 0.$$

This suggests a test of full insurance, but doesn't identify the key preference parameter γ . Note, however, that we can exploit (8) to estimate γ , as below.

Credit markets: In this case, suppose that each household has access to credit markets with returns $R(s^t)$, but that insurance may be imperfect, so that households may face idiosyncratic risk. In this case, the Euler equation for each household's consumption expenditures is given by

$$(10) \quad u'_i(c_i(s^t)) = \beta R(s^t) E_t u'_i(c_i(s^{t+1})).$$

This is enough for us to see that the distribution of consumption in this case is non-stationary, even if $\beta R(s^t) = 1$, so long as there's any idiosyncratic risk at all. This implies that a strategy of simply using the empirical distribution of consumption over time isn't appropriate in this environment.

Making use of our assumption of CRRA utility it follows that households' one-period ahead forecast errors are related to

$$\xi_{it+1} = \left(\frac{c_{it+1}}{c_{it}} \right)^{-\gamma} - [\beta R_t]^{-1}.$$

Given knowledge of $\{\beta R_t\}$ and γ , this expression can be used to infer what the one-period ahead risk facing the household is, as well as the other components of vulnerability. The immediate difficulty we face is that these quantities may *not* be known. However, noting that if this Euler equation holds for every household in a sample of n households, then we have

$$[\beta R_t]^{-1} = \frac{1}{n} \sum_{j=1}^n E_t \left(\frac{c_{jt+1}}{c_{jt}} \right)^{-\gamma}.$$

Letting

$$M_{nt}(\gamma) = \frac{1}{n} \sum_{j=1}^n \left(\frac{c_{jt}}{c_{jt-1}} \right)^{-\gamma}$$

we have $[\beta R_t]^{-1} = E_t M_{nt+1}(\gamma)$, so that letting $x_{it} = c_{it}/c_{it-1}$,

$$(11) \quad E[x_{it+1}^{-\gamma} - M_{nt+1}(\gamma) | s^t] = 0.$$

This restriction and rational expectations then permits us to estimate the single parameter γ and thence to infer the risk and vulnerability facing households in the sample.

Credit markets and measurement error: The approach just taken to estimating γ using the Euler equation doesn't take into account the problem of measurement error in consumption, and as Runkle (1991) observes, this sort of measurement error can cause very serious problems for nonlinear estimators (such as the nonlinear GMM estimator suggested by (11)).

As a consequence here we'll develop a weaker Euler-type restriction which permits consistent estimation even in the presence of fairly general measurement error processes, an approach inspired by Chiolda (2004).

We begin by defining an error-ridden measure of consumption, \tilde{c}_{it} , which is related to actual consumption by $\tilde{c}_{it} = c_{it} e^{\nu_{it}}$, where $\{\nu_{it}\}$ is an iid measurement error process. Now, let $\tilde{x}_{it} = \tilde{c}_{it}/\tilde{c}_{it-1}$; then condition (8) implies

$$E \left[\tilde{x}_{it+1}^{-\gamma} \left(\frac{\nu_{it+1}}{\nu_{it}} \right)^\gamma - \tilde{M}_{t+1}(\gamma) \middle| s^t \right] = 0,$$

where $\tilde{M}_t(\gamma) = \text{plim}_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \left(\frac{\tilde{c}_{jt}}{\tilde{c}_{jt-1}} \right)^{-\gamma}$. Thus, if we knew the realizations of the measurement error process ν_{it} we could straightforwardly estimate the preference parameter γ , and the factors $\{\mu_i\}$ and $\{\beta R_t\}$. However, without observing the measurement error, we can construct a moment condition which can be used to identify γ and $\{\beta R_t\}$. In particular, we have

Proposition 1. *Let $\eta_{it} = \nu_{it+1}/\nu_{it}$, and let $\eta_t(b) = \text{plim}_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \eta_{jt}^b$. Now, if households' consumption processes satisfy the Euler equation (10) and if*

$$(12) \quad E \left[\tilde{M}_{t+1}(\gamma) \left(\frac{\eta_{it+1}^{-\gamma}}{\eta_{t+1}(-\gamma)} \right) \middle| s^{t-1} \right] = E \left[\tilde{M}_{t+1}(\gamma) \middle| s^{t-1} \right]$$

then

$$(13) \quad \text{plim}_{n \rightarrow \infty} E \left(\tilde{x}_{it+1}^{-\gamma} - \tilde{M}_{t+1}(\gamma) \middle| s^{t-1} \right) = 0.$$

Proof. Since every households' consumption profile satisfies (10), it will also satisfy (11). Let $\epsilon_{it+1} = x_{it+1}^{-\gamma} - M_{nt+1}(\gamma)$, and note from our definition of \tilde{x}_{it} and \tilde{M}_{nt} that this implies that

$$\begin{aligned} \tilde{x}_{it+1}^{-\gamma} - \frac{1}{n} \sum_{j=1}^n \tilde{x}_{jt+1}^{-\gamma} \frac{\eta_{it+1}^{-\gamma}}{\frac{1}{n} \sum_{j=1}^n \eta_{jt+1}^{-\gamma}} - \frac{1}{n} \sum_{j=1}^n \epsilon_{jt+1} \eta_{jt+1}^{-\gamma} \frac{\eta_{it+1}^{-\gamma}}{\frac{1}{n} \sum_{j=1}^n \eta_{jt+1}^{-\gamma}} \\ = \epsilon_{it+1} \eta_{it+1}^{-\gamma}. \end{aligned}$$

Now, $\{\epsilon_{it}\}$ is a martingale difference sequence by construction, so that $E(\epsilon_{it+1} \eta_{it+1}^{-\gamma} | s^t) = 0$, implying

$$E \left[\tilde{x}_{it+1}^{-\gamma} - \frac{1}{n} \sum_{j=1}^n \tilde{x}_{jt+1}^{-\gamma} \frac{\eta_{it+1}^{-\gamma}}{\frac{1}{n} \sum_{j=1}^n \eta_{jt+1}^{-\gamma}} \middle| s^t \right] = 0.$$

Taking the probability limit as the sample size n goes to infinity yields

$$E \left[\tilde{x}_{it+1}^{-\gamma} - M_{t+1}(\gamma) \frac{\eta_{it+1}^{-\gamma}}{\eta_{t+1}(-\gamma)} \middle| s^t \right] = 0;$$

subsequently exploiting the Law of iterated expectations and the condition (12) then yields the result. \square

The idea of the proposition is very simple. If all households have access to credit markets on the same terms, so that their consumption profiles satisfy (10), then with CRRA utility it

follows that all households expect the same rate of consumption growth. Measurement error in consumption will bias naive estimates of the rate of actual consumption growth, but if the expected growth rate of *measurement error* is the same across households, then the Euler equation with CRRA utility implies that with credit markets all households' expected *error-ridden* consumption growth rates will be equal. For it to be reasonable for expected growth rates in measurement error to be equal, it's important to take expectations conditioning on the s^{t-1} information set, prior to any knowledge regarding differing time t realizations of measurement error across households.

Hidden Actions, Limited access to credit markets: From the main result of Rogerson (1985), we know that the reciprocal of marginal utility will satisfy an Euler-type restriction of the form

$$\frac{1}{u'(c_i(s^t))} = \beta R(s^t) E_t \frac{1}{u'(c_i(s^{t+1}))}$$

(see also Kocherlakota (2005)). This can serve as an alternative to the usual “credit-market” Euler equation restriction in estimation, and implies a restriction very closely related to (11). In particular, by an argument precisely analogous to that above, efficient dynamic arrangements with hidden actions imply the empirical restriction

$$(14) \quad E[x_{it+1}^\gamma - M'_{t+1}(\gamma) | s^t] = 0,$$

where $M'_t(\gamma) = \frac{1}{n} \sum_{j=1}^n \left(\frac{c_{jt}}{c_{jt-1}} \right)^\gamma$, while with measurement error we have

$$(15) \quad E[\tilde{x}_{it+1}^\gamma - \tilde{M}'_{t+1}(\gamma) | s^{t-1}] = 0.$$

Note that these two restrictions can be nested with, respectively, (11) and (13) in a more general restriction which permits one to identify the environment as well as estimating the preference parameter γ , as in Ligon (1998).

Hidden Information, Limited access to credit markets: From Rogerson (1985) we know that when there's hidden information (independently distributed over time), a restriction similar to that which holds with hidden actions will be satisfied *except* that expectations must be conditioned only on commonly held information; i.e., one period earlier than in the case of hidden actions. Accordingly, we have

$$\frac{1}{u'(c_i(s^t))} = \beta E_{t-1} \frac{R(s^t)}{u'(c_i(s^{t+1}))}.$$

Since taking time $t - 1$ expectations also solves the problem of measurement error above, then the restriction (13) we originally devised to deal with measurement error when households all have equal access to credit markets *also* holds (modulo a sign on the estimated preference parameter γ) in environments *without* common access to credit markets and hidden information.

We've demonstrated that, even in the presence of measurement error, idiosyncratic time preferences, and uncertainty regarding the economic environment, panel data on household consumption expenditures can be exploited to estimate household risk preferences. In particular, *any* of these economic environments satisfy the restriction

$$(16) \quad E \left(\tilde{x}_{it+1}^{b_0} - \tilde{M}_{t+1}(b_0) \middle| s^{t-1} \right) = 0.$$

One can exploit (16) to derive estimates of the parameter b_0 . Provided that the risk aversion parameter γ is positive, then one can interpret the absolute value of an estimate \hat{b} of b_0 as an estimate of γ , so that we can write $\hat{\gamma} = |\hat{b}|$. The sign of \hat{b} provides information on the nature of the underlying economic environment—if significant and positive, then agents can be assumed to have access to credit markets, if significant and negative, then this suggests one of the private information environments, as in Ligon (1998) or Kocherlakota and Pistaferri (2007). Finally, if estimates of b_0 are not significantly different from zero, then one can't reject the hypothesis that markets are complete, and that agents face no idiosyncratic risk.

In sum, we've derived a restriction on households' growth rate of consumption (13) which can be used to estimate the preference parameter γ *even when the econometrician is ignorant* about the actual underlying economic environment—estimation of γ can proceed whether there's full insurance, self-insurance via credit markets, moral hazard, or private information. Further, estimation of γ is robust to measurement error in consumption so long as the measurement error process satisfies a rather weak requirement that expected changes in measurement error growth be the same across the sample.

5. ESTIMATION

In the previous section we've described empirical restrictions which might be useful for estimating the preference parameter γ , given beliefs about the prevailing economic environment (we'll return to the problem of inferring this below).

We now turn our attention to estimating risk and vulnerability, conditional on estimated preference parameters. With preferences in hand,

we only need to characterize the distribution of consumption for every household at every date, and we can then integrate to calculate vulnerability.

We begin by calculating the errors in our (not the households'!) forecasts, which can be inferred by ‘plugging’ our estimate \hat{b} into the restriction (11). In particular, we use

$$(17) \quad \xi_{it+1} = \left(\frac{\tilde{c}_{it+1}}{\tilde{c}_{it}} \right)^{\hat{b}} - M_t(\hat{b}),$$

to compute residuals from our estimation procedure, and then decompose these into orthogonal idiosyncratic and aggregate components, letting

$$\xi_{it} = \eta_t + \epsilon_{it}.$$

We next compile the realizations the aggregate components $\{\eta_t\}$ and the idiosyncratic component $\{\epsilon_{it}\}$ (the latter for every household in the sample, provided the panel is of reasonable length), and resample these empirical distributions so as to construct many possible future consumption paths for every household.

We use this resampling approach to construct estimates of vulnerability, but our measure of vulnerability depends not just on γ (for which we have an estimate) and the distribution of future paths of consumption, but also on households’ rate of discounting. The discount factor β isn’t identified by the restriction (16). It might be possible to use another restriction to estimate this parameter, but in practice our estimates of vulnerability don’t seem to depend on β very much—changing the value of β assumed more or less simply scales all of our welfare measures up and down. Accordingly, we can simply pick a reasonable value of this parameter, and proceed.

We then average over the values of discounted utility from the realized sample consumption streams to construct an estimate of vulnerability via monte carlo integration, in an approach similar to that of Kuhl (2003).

5.1. Decompositions of Vulnerability. While the methods described above give us a method to estimate total vulnerability, and Section 3 suggests a simple way to decompose household-level vulnerability into measures of poverty/inequality, welfare loss due to predictable intertemporal variation, and risk, it’s possible to go further and decompose the risk component, thus measuring the welfare cost of risk due to, e.g., aggregate shocks, idiosyncratic income shocks, shocks to household composition, etc.

The basic idea is to use data on realized shocks in the sample, and then to estimate how these shocks influence realized consumption. Let $\{z_k^t\}_{k=1}^K$ denote a collection of sets of information, which is increasing in the sense that $z_k^t \subset z_{k+1}^t \subseteq s^t$. Recall that risk R_{it} is given by

$$R_{it} = (1 - \beta) \sum_{j=1}^{\infty} \beta^j [u(\mathbb{E}(c_{it+j}|s^t)) - \mathbb{E}[u(c_{it+j})|s^t]].$$

However, any term of this expression can be decomposed, since

$$\begin{aligned} u(\mathbb{E}(c_{it+j}|s^t)) - \mathbb{E}[u(c_{it+j})|s^t] &= \{u(\mathbb{E}(c_{it+j}|s^t)) - \mathbb{E}[u(\mathbb{E}(c_{it+j}|z_1^{t+j}))|s^t]\} \\ &+ \sum_{k=1}^{K-1} \{\mathbb{E}[u(\mathbb{E}(c_{it+j}|z_k^{t+j}))|s^t] - \mathbb{E}[u(\mathbb{E}(c_{it+j}|z_{k+1}^{t+j}))|s^t]\} \\ &+ \{ \mathbb{E}[u(\mathbb{E}(c_{it+j}|z_K^{t+j}))|s^t] - \mathbb{E}[u(c_{it+j})|s^t] \}. \end{aligned}$$

A typical term in this contribution involves the quantity $\mathbb{E}[u(\mathbb{E}(c_{it+j}|z_k^{t+j}))|s^t]$. The innermost conditional expectation relates variation in consumption to variation in contemporaneous variables in z_k^{t+j} . Typically the latter won't be known at time t , so the outermost conditional expectation involves computing the expected utility which would result if variation in z_k^{t+j} was the only source of variation in consumption, using only information available at time t . Accordingly, the difference $\mathbb{E}[u(\mathbb{E}(c_{it+j}|z_k^{t+j}))|s^t] - \mathbb{E}[u(\mathbb{E}(c_{it+j}|z_{k+1}^{t+j}))|s^t]$ gives a measure of the welfare cost of variation in variables in $z_{k+1}^{t+j} \setminus z_k^{t+j}$, all from the perspective of time t .

To estimate these conditional expectations, we proceed in three steps. First, we estimate the innermost conditional expectations. Let $Z_{it}(k)$ be a matrix of variables pertaining to household i in the information set z_k^t . Then, using least squares, we estimate a sequence of parameters $\{\delta_k\}$ from

$$\log c_{it} = Z_{it}(k)\delta_k + v_{it}(k),$$

$k = 1, \dots, K$ where $v_{it}(k)$ are disturbances associated with the k th estimating equation; note that these will include the negative of logarithm of any multiplicative measurement error associated with c_{it} . Letting $\hat{\delta}_k$ denote the estimated parameters from the k th regression, and $\hat{v}_{it}(k)$ the residuals, it follows that

$$\mathbb{E}(c_{it}|Z_{it}(k)) = \exp(Z_{it}(k)\hat{\delta}_k)\mathbb{E}(e^{\hat{v}_{it}(k)}|Z_{it}(k)).$$

We estimate the factor $\mathbb{E}(e^{\hat{v}_{it}(k)}|Z_{it}(k))$ via a second regression, again using least squares,

$$(18) \quad e^{\hat{v}_{it}(k)} = Z_{it}(k)\phi_k + w_{it}(k),$$

thus constructing an estimate of the inner conditional expectations of c_{it}

$$(19) \quad \hat{c}_{it}^{(k)} = \exp(Z_{it}(k)\hat{\delta}_k)Z_{it}(k)\hat{\phi}_k,$$

where $\hat{\phi}_k$ denotes the least square estimates of ϕ_k in (18).

Naturally, we can estimate $\hat{c}_{it}^{(k)}$ only for periods in which we actually observe c_{it} , and in our present application (*contra* the static application of Ligon and Schechter (2003)) we wish to estimate $\hat{c}_{it+j}^{(k)}$ for any $j = 1, 2, \dots, \infty$. To do so, we pose a counterfactual supposition, to whit: Suppose that variation in household i 's consumption was determined entirely by variation in $Z_{it}(k)$, so that (13) would hold substituting $\{\hat{c}_{it}^{(k)}\}$ for $\{c_{it}\}$. We can then use the procedure described above for estimating the profile of all future consumption (using (13) to estimate forecast errors for each household, attributing these to aggregate and idiosyncratic components, resampling from the distributions of these two error components to construct an arbitrarily long realized sequence of forecast errors, computing the sequence of consumptions implied by this sequence of forecast errors, and repeating the resampling until our estimates of expected discounted utility converge), but to estimate the discounted expected utility of consuming $\hat{c}_{it}^{(k)}$.

It's important to note that, because the distribution of $\hat{c}_{it}^{(k)}$ will differ from that of c_{it} , this will affect the intertemporal behavior of the household, and hence the predicted path of consumption over time (for example, with access to credit markets and $\gamma > 1$, precautionary motives will tend to lead the household to save more when variation in consumption is greater, so that predicted consumption growth will typically be larger for c_{it} than for $\hat{c}_{it}^{(k)}$). We can correct for this by accounting for the difference in welfare due to differences in predicted consumption paths.

6. THE DATA

We provide a simple application of the measurement of risk and vulnerability. The data we will use in this study is from the Household Budget Survey (HBS) in Bulgaria, collected by the Central Statistical Office of Bulgaria, and previously described by Peters and Hassan (1995) and Skoufias (2001).² These data include information on 2287 households over 12 months. The sampling scheme employed by the HBS involved a clustered design, with clustering at the level of region, but was designed to be representative of the population of households

²This section is derived from Ligon and Schechter (2003).

residing in Bulgaria. However, households with only a single monthly observation or with per capita consumption in the bottom or top percentiles have been dropped. The survey includes variables such as age, gender, education, sector of the economy, and employment status. Most importantly this survey contains detailed information on household level income and non-durable consumption.

This survey was conducted during a very tumultuous period for Bulgaria. In 1991 price liberalization was undertaken and the share of administered prices in the Consumer Price Index went down from 70% to 24%, and by 1992 down even further to 16%. This price liberalization brought about severe output drops, perhaps caused by the disruption of productive links. The Gini Index between 1987 and 1989 was .23 and GDP per capita was 1730 Bulgarian Leva. Between 1993 and 1997 the Gini Index rose to 0.34 while per capita GDP fell to 1270 leva. In response to all of these changes, in 1994 the communists were re-elected to power and the government increased the share of controlled prices to 43 per cent. Using data from a period of such extreme shocks may make it possible to detect which households are insured against fluctuations.

A problem with most measures of consumption is that they do not reflect actual consumption when households consume out of their storage or their own production. This dataset avoids that problem. The HBS contains, for each food item, information on its stock at the beginning and end of each month, as well as flow quantities entering or leaving the household from production at home, gifts to or from friends, and quantities used as seed.³ Skoufias has created a food consumption variable for each food item which he calculates as

$$c_{it} = I_{it} + P_{it}Q_{it},$$

where I_{it} is defined as the value of purchases of that item and P_{it} as the national median unit value of that item. Q_{it} is the quantity in stock at the beginning of the month minus the quantity in stock at the end of the month, plus that obtained from reprocessing, from business organizations, from other sources, and produced at home. In addition he subtracts the quantity used for reprocessing or to feed animals, given out as presents or loans, sold, lost, wasted, or used for seed.

For non-food items it had not been possible to use the same approach.⁴ The HBS survey contains no information on their stock, only

³Food items include cereals, meats, milk, fish, eggs, dairy products, fruits, vegetables, sugar, fats, beverages, alcohol, and expenditures on eating out.

⁴Non-food items include tobacco, electricity, central heating and other energy, trash, water, telecommunications, education, gasoline, transportation, furniture,

on monthly expenditures and domestic production. Skoufias also created a measure of non-food expenditures from monthly purchases plus domestic production times the median unit value of that item. One can sum food and non-food consumption to find total consumption. This calculation of total consumption is, however, plainly imperfect, as seven of the 2287 households in the survey appear to experience negative levels of consumption during at least one month of the survey period. These households are excluded from the analysis.

The data set contains equally detailed income data. The measure of income we use includes salary, self-employment income, rent, interest, dividends, pension, unemployment benefits, disability payments, child allowances, maternal benefits, family benefits, other benefits, farm product sales minus farm product expenses, property sales, and other income. We also have data on, but do not include, transfers from friends and relatives and net loans, borrowings, and savings. All consumption and income variables are normalized by the national CPI with a base of June 1994. We have also expressed these in units of adult equivalent consumption.⁵

7. RESULTS

This section describes some simple results from applying the results of Section 5 to the panel dataset described in Section 6. We use the broader measure of total consumption in preference to the narrower measure of food consumption in these analyses. Units of total consumption are chosen so that the mean of consumption per adult-equivalent over all periods is equal to one. The distribution of these normalized consumptions is shown in Figure 1.

We use the restriction implied by (16) to estimate the key parameter b_0 (recall that $\gamma = |b_0|$, and that the sign of b_0 provides some key information on whether or not households have access to credit markets). We use a continuously-updated GMM estimator of the sort described by Hansen et al. (1996) and discussed in Imbens et al. (1998), which avoids the problem of a degeneracy in the moment conditions when evaluated at $b = 0$. For the set of instruments, we use appropriately lagged levels of adult-equivalent consumption, and lags in the growth of

health, clothing, entertainment and leisure, rent and home maintenance, insurance, cleaning, small appliances, domestic services, fees, and taxes.

⁵Our measure of adult equivalents assigns the consumption of adult males a weight of 1 and adult females a weight of 0.9 (adult means sixteen or older). Children ages 0 to 4 count as 0.32, ages 5 to 9 as 0.52, and ages 10 to 15 as 0.67. This is nearly the scheme used by Townsend (1994), save that our age brackets are slightly different.

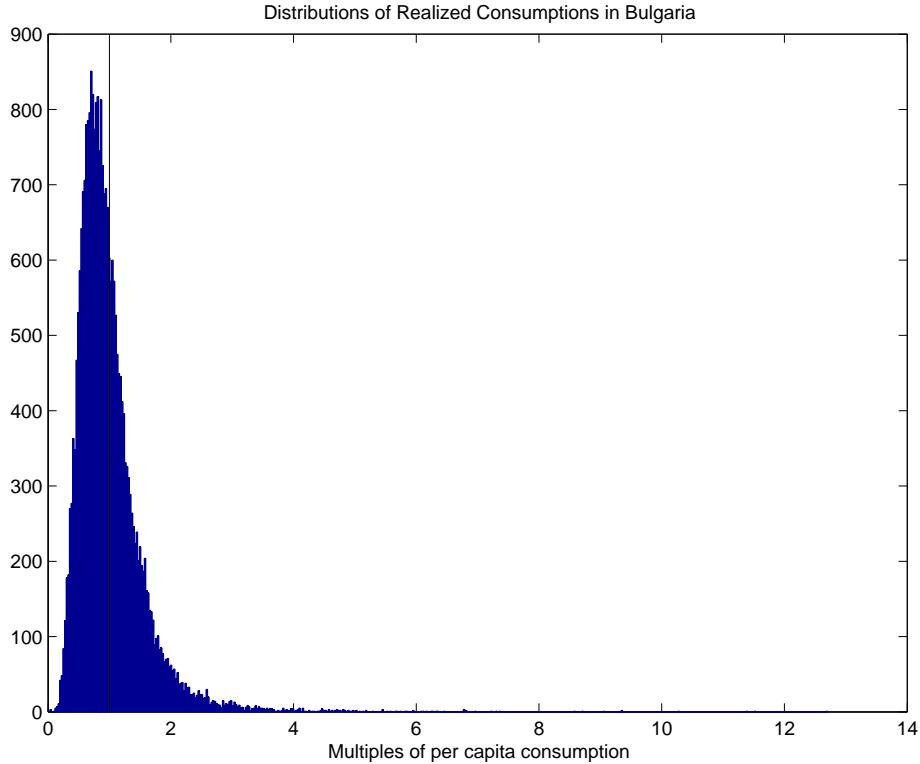
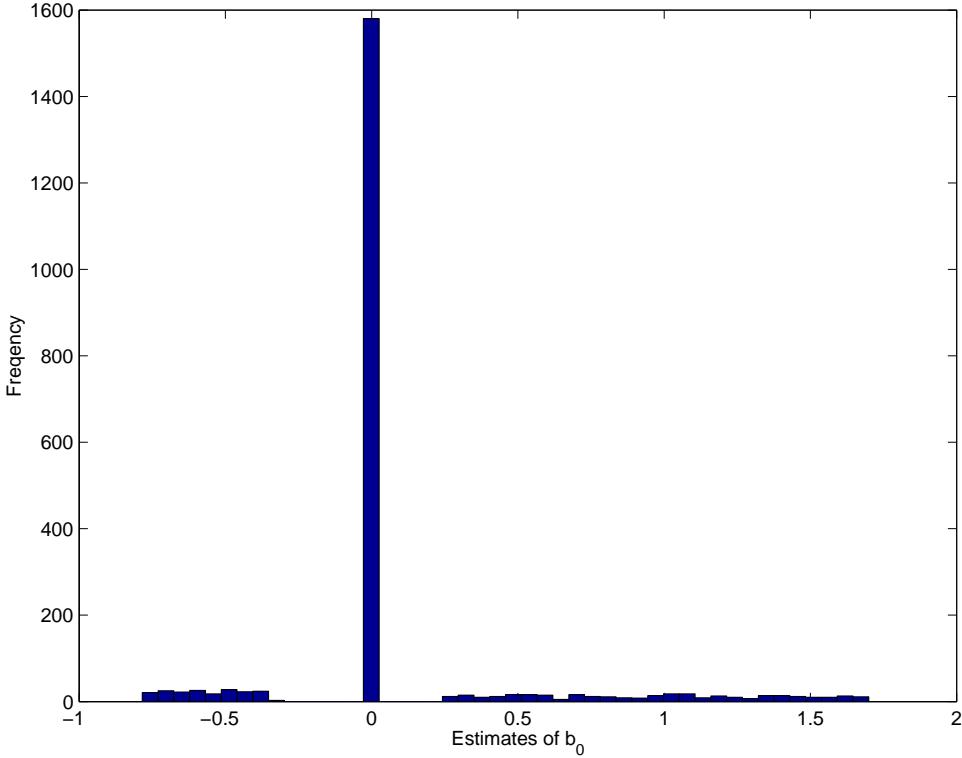


FIGURE 1. Distribution of monthly per adult-equivalent total consumption in Bulgarian Household Survey.

per capita household level income. This yields an estimate $\hat{b} = 0.9024$, but with an estimated standard error of 2.32. Thus, while the data seems to point to one of the private information environments, a conventional t -test wouldn't allow us to reject a null of full insurance or credit markets.

That said, a conventional t -test may not be a good test in this environment. By repeatedly resampling households and generating new estimates of b_0 it becomes apparent that the distribution of the estimates is decidedly non-normal; see Figure 2. In particular, the support of the distribution of \hat{b} has three distinct parts. A mass of estimates (75 per cent) is zero (within the limits of machine tolerance) around zero; for these bootstrap samples the data suggests that lagged levels of consumption and income changes can't predict consumption changes, so that the full insurance hypothesis can't be rejected. A second part (15.7 per cent) ranges between 0.26 and 1.70, and has a mean of 0.955 (close to our original estimate of 0.9024) and a standard deviation of 0.421. Thus, conditional on the estimate being positive, a conventional

FIGURE 2. Distribution of estimates of b_0

test would indicate that it was significantly positive. The remaining estimates (nine per cent) are negative, ranging from -0.784 to -0.335 with a mean of -0.560 and a standard deviation of 0.128 . Thus, these negative estimates are similarly conditionally significant.

This highly non-normal distribution of our estimates of b_0 seems to be a consequence of having quite weak instruments; Chioda (2004) seems to have encountered similar kinds of problems in trying to deal with measurement error as we have here. Her results, and those of Attanasio and Low (2000) suggest that we may need a considerably longer panel before we can expect to obtain reliable results.

Regardless of the (im)precision of our estimates of b_0 , and hence of γ , all is not lost if our aim is to estimate risk and vulnerability. Our original estimate of γ , 0.9024, isn't grossly implausible (though most micro-econometric estimates find a larger number), and so we adopt to proceed as outlined in Section 5.1 to estimate the risk and vulnerability for the average Bulgarian households, conditional on this estimate of γ . Using the same variables employed as instruments in

Variable	Stage 1		Stage 2	
	Estimate	Std. Err.	Estimate	Std. Err.
$\log(c_{it-1})$	0.5853	0.0054	-0.0014	0.0073
Time dummy 1	0.0129	0.0081	1.0167	0.0110
Time dummy 2	-0.0249	0.0081	1.0627	0.0110
Time dummy 3	0.0753	0.0081	1.0784	0.0110
Time dummy 4	-0.1069	0.0081	1.0720	0.0110
Time dummy 5	-0.1472	0.0081	1.0906	0.0110
Time dummy 6	-0.2468	0.0081	1.0770	0.0110
Time dummy 7	-0.2252	0.0081	1.1007	0.0110
Time dummy 8	-0.1843	0.0081	1.1180	0.0110
Time dummy 9	-0.2608	0.0081	1.1058	0.0110
Time dummy 10	-0.2796	0.0081	1.1067	0.0110
Time dummy 11	-0.2442	0.0081	1.0870	0.0110
Time dummy 12	-0.0185	0.0081	1.0912	0.0110
R^2	[0.3457]		[0.0210]	

TABLE 1. Estimates from Two Stage Prediction of Consumption

the GMM estimation of b_0 , we use the two-step procedure outlined above to construct predictions of c_{it} ; these predictions versus measured consumption can be seen in Figure 3, and point estimates for both stages in Table 1.

In the first stage, nearly all our estimates are significant. The coefficient associated with lagged consumption shows the “regression to the mean” that we’d expect in the presence of measurement error. Our GMM estimator of b_0 was, of course, designed to deal with the bias which results from this error, but for the current problem of simply trying to predict consumption this measurement error doesn’t cause serious problems. Both lagged consumption and the set of time dummies are all highly significant. The relatively large, negative values of the time dummies in the latter part of the sample period is strong evidence regarding the importance of aggregate shocks or seasonality affecting Bulgarian households during this period.

The second stage in principle allows us to take advantage of systematic variation in higher order moments of the log consumption distribution. However, in practice there is no evidence of heteroskedasticity or other violations of an assumption of identically distributed disturbance terms in the first stage regression. This is evidenced by the insignificant coefficient associated with lagged log consumption, and by the fact that

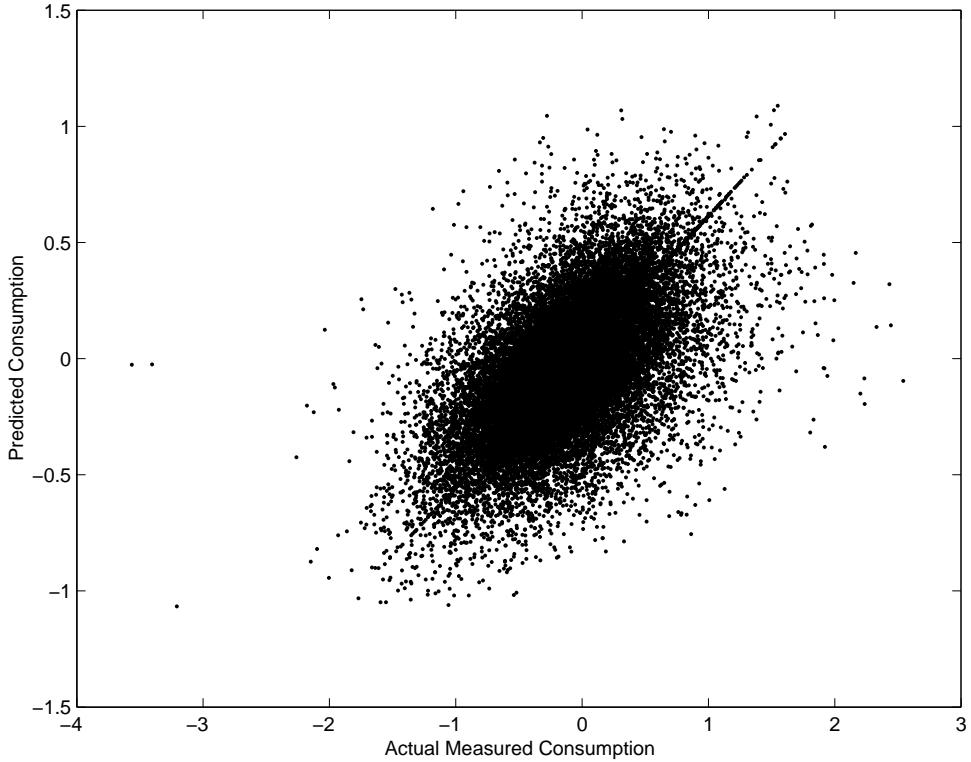


FIGURE 3. Predictions of log Household Adult-Equivalent Consumption versus log Measured Consumption.

the time dummies (though all significant) are not significantly different from one another.

Now, we wish to use the estimated distribution of future consumption paths to estimate risk and vulnerability. We fix $\beta = 0.97$ (as noted above, results aren't very sensitive to this choice), and use our earlier estimate of $\gamma = 0.9024$. We then compute forecast errors from (17), and decompose them into orthogonal aggregate and idiosyncratic components, with a common empirical distribution of aggregate shocks and a distinct empirical distribution of idiosyncratic shocks for each household. We draw sequences of shocks from each of these distributions, and use these to construct arbitrarily long sample paths for individual household forecast errors. These then are used to construct a distribution of possible future consumption paths, and realized utility along this path is computed, using the values of γ and β given above. We average many such paths for each household, thus obtaining a household-specific estimate of vulnerability; see the distribution

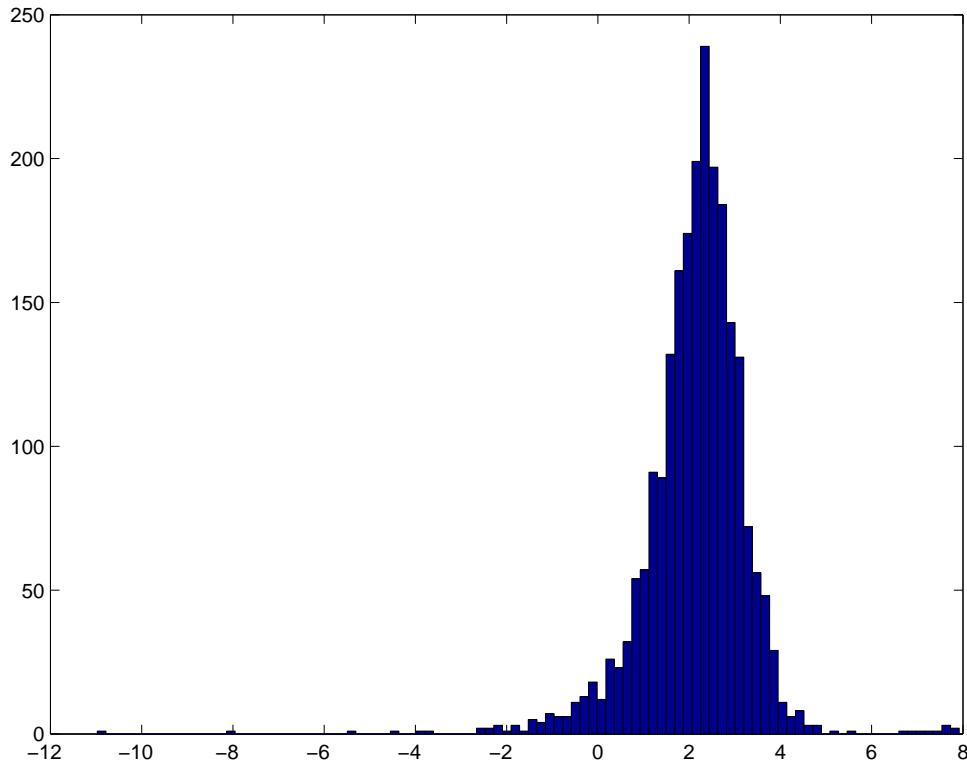


FIGURE 4. Estimates of Household Vulnerability.

of these estimates in Figure 4. The units of vulnerability are in utils; the mean household vulnerability is 2.135.

Welfare Component	Vulnerability	Poverty	Variation
Means	2.135	0.087	2.047
Std. Dev.	1.140	0.434	1.250
Correlations			
Vulnerability	1.000	-0.076	0.938
Poverty	-0.076	1.000	-0.416
Variation	0.938	-0.416	1.000

TABLE 2. Vulnerability, Poverty, and Welfare Loss due to Variation

From there we're interested in learning what proportion of this vulnerability is due to poverty/inequality, what part predictable time series variation, and what part risk. To estimate a measure of poverty

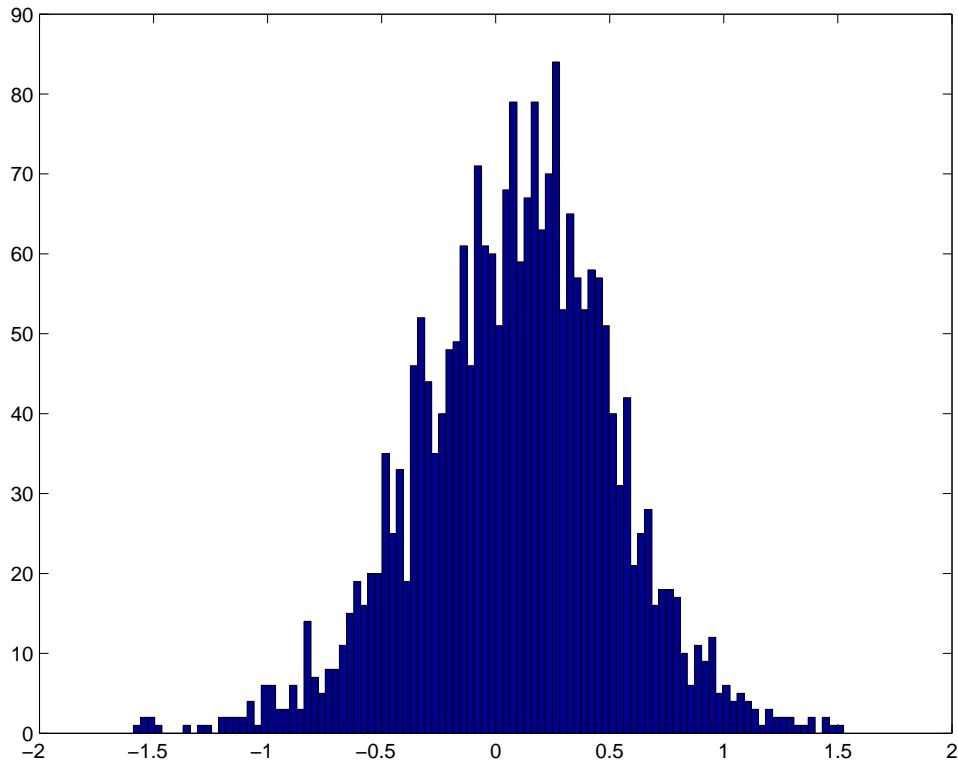


FIGURE 5. Estimates of Household Poverty.

we compare the utility that the household would derive if every household received the same adult-equivalent consumption in the first period as every other to the utility the household receives from their actual predicted period one consumption. The distribution of poverty across households is pictured in Figure 5. These are welfare costs measured in utils. The average poverty is 0.0873, so that ex ante inequality appears to play only a rather small role in reducing welfare (poverty of 0.0873 is only 4.1 per cent of total vulnerability).

Because total vulnerability is so much larger than poverty, this is an indication that future variation in consumption seems to weigh much larger in households' calculus than does initial inequality—if this society was presented with a choice between completely eliminating contemporary (but not future) inequality or fixing inequality but eliminating all future variation in consumption, fewer than eight per cent of all households would vote to eliminate inequality instead of variation, because risk and time series variation in future consumption are more than low levels of consumption for most households.

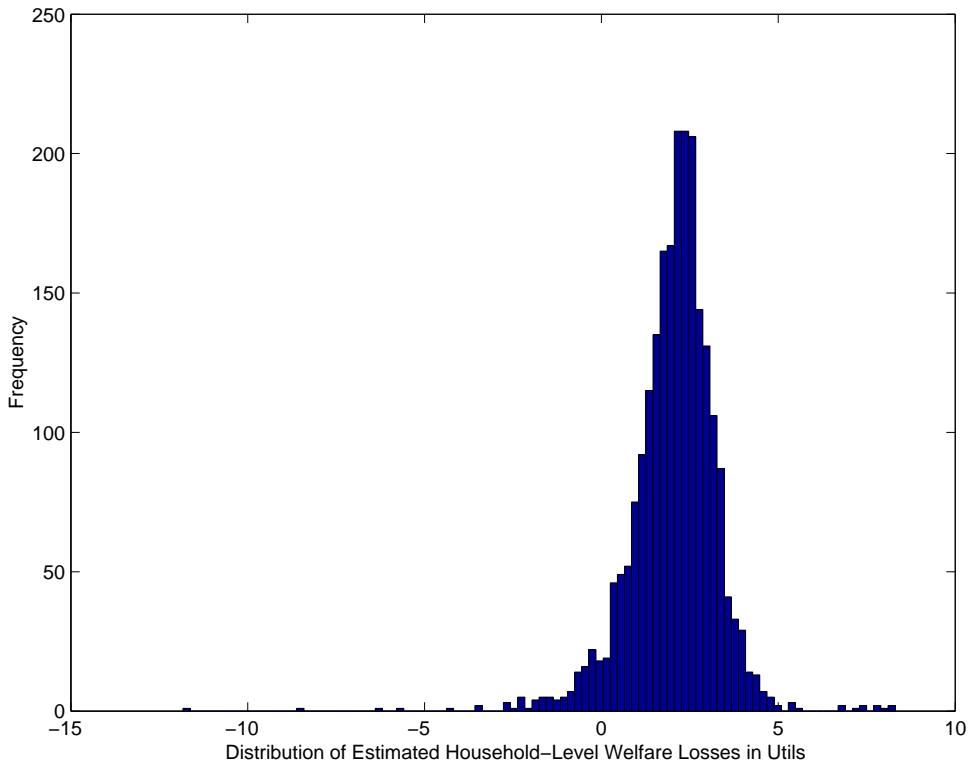


FIGURE 6. Estimates of household welfare loss due to risk and predictable variation.

This brings us to the decomposition of risk and predictable time series variation promised in Section 3. Unfortunately this decomposition turns out to be problematical for these data. The importance of the aggregate component of consumption growth and the relatively short time series means that it's very difficult to separately identify risk and predictable time series variation; though we can compute the decomposition, in practice we obtain very large estimates of both risk and of predictable time series variation, but of opposite sign. Since these two seem difficult to distinguish in these data, we've opted here to present only their sum, which is 2.047. This is reported in Table 2, along with data on the correlation between different components of vulnerability. In this connection, it's worth noting that poverty in these data is negatively correlated with the welfare costs of variation; there's no evidence in these data that the poor face more risk and other variation in consumption than do wealthier households, though of course this may be because of costly risk-coping strategies adopted by poor households.

8. CONCLUSION

There's a rapidly growing literature on ways to measure the welfare costs of risk and 'vulnerability', particularly for poorer households in developing countries. The usual approach to measurement in this literature involves using household-level panel datasets with data on consumption (or expenditures), and then exploiting time-series variation in consumption to draw inferences about the distribution of consumption at some future point in time (see, for example ?).

However, for this approach to be valid one must make some additional identifying assumptions. The literature to date has taken one of unsatisfactory approaches. In the first approach, typified by Ligon and Schechter (2003), the econometrician more or less avoids coming to grips with modeling the economic behavior of the household, and adopts some very strong and more or less *ad hoc* assumptions regarding the distribution of consumption and the nature of household risk aversion (in Ligon-Schechter that the former is independent and identically distributed in every period, and that the latter is known by the econometrician).

The second approach, typified by Elbers and Gunning (2003) or ?, takes a thorough-going structural approach. In this approach, understanding the economic behavior of the household is taken very seriously. Modeling this behavior involves specifying all the details of the economic environment and then solving the dynamic-programming problem facing the household. This is, in principle, much the more intellectually satisfying approach, and allows one to use the data to estimate key structural parameters such as preference parameters, rather than assuming that one knows them in advance. Unfortunately, for this approach to be computationally tractable the economic environment has to be drastically simplified. Accordingly, one might complain that the structural approach basically replaces an unpalatable, *ad hoc* statistical assumption with an heroic, implausible structural model. Further, the highly non-linear approach to estimation typically required by these dynamic programming approaches to estimation makes it very difficult to deal with the problem of measurement error in consumption (Runkle, 1991).

The present paper attempts to improve this state of affairs. A central result of the paper is to show that a rather wide variety of economic environments all imply a particular conditional moment restriction on households' consumption processes. Generalized method of moments estimators can be used to estimate a key risk-aversion parameter, to draw inferences regarding the nature of the economic environment, and

to generate an empirical distribution of the econometrician's forecast errors of household consumption. Further, drawing on work by ? and Kocherlakota and Pistaferri (2007), this approach to estimation can be made robust to multiplicative classical measurement error in consumption. By resampling this empirical distribution of forecast errors, one can construct possible future sample paths for household consumption, and estimate the welfare costs associated with variation in future consumption. These welfare costs, in turn, can be decomposed in to both predictable and unpredictable components—the last can be thought of as the welfare cost associated with risk.

By way of application, we use a panel dataset on Bulgarian households, collected a tumultuous period during that nation's transition from communism. We are able to use these data to estimate household risk aversion and to compute plausible estimates of the welfare costs associated with variation in future consumption. We are further able to measure the proportion of these welfare costs which are due to *aggregate* rather than idiosyncratic variation. In the case of these data, aggregate variation turns out to be extremely important, relative to individual variation.

This approach does have some serious shortcomings. The cost of using a very general moment restriction, which should hold in a wide variety of settings and which is robust to measurement error is that the precision of our estimates and the power of our tests is quite low. An immediate consequence in our application is that we are unable to separately identify the welfare costs of predictable time series variation in consumption (e.g., seasonal or life-cycle variation) from unpredictable sources of variation (risk). Consequently, the use of the methods described here might be better suited to an application in which one either has a longer panel, or in which observable sources of idiosyncratic shocks are more important relative to the aggregate than they seem to be in our Bulgarian data.

REFERENCES

Attanasio, O. P. and H. Low (2000, May). Estimating euler equations. NBER Technical Working Papers 0253, National Bureau of Economic Research, Inc. available at <http://ideas.repec.org/p/nbr/nberte/0253.html>.

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.

Chioda, L. (2004a). Estimating Euler equations with measurement error: a nonparametric approach. Unpublished Working Paper; Available at <http://www.princeton.edu/~lchioda/>.

Chioda, L. (2004b). Estimating Euler equations with measurement error: a nonparametric approach. Unpublished Working Paper; Available at <http://www.princeton.edu/~lchioda/>.

Deaton, A. (1992). Household saving in LDCs: Credit markets, insurance, and welfare. *Scandinavian Journal of Economics* 94, 253–273.

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

Hansen, L. P., J. Heaton, and A. Yaron (1996). Finite-sample properties of some alternative GMM estimators. *Journal of Business and Economic Statistics* 14(3), 262–280.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *Journal of Political Economy* 63.

Hoddinott, J. and A. Quisumbing (2003, December). Methods for microeconometric risk and vulnerability assessments. SP Discussion Paper 0324, The World Bank, Washington, DC.

Imbens, G. W., R. H. Spady, and P. Johnson (1998). Information theoretic approaches to inference in moment condition models. *Econometrica* 66(2), 333–358.

Kocherlakota, N. (2005). Zero expected wealth taxes: A Mirrlees approach to dynamic optimal taxation. *Econometrica* 73(5), 1587–1621.

Kocherlakota, N. and L. Pistaferri (2007, January). Asset pricing implications of pareto optimality with private information. Unpublished working paper; available at <http://www.stanford.edu/~pista/revision.pdf>.

Kuhl, J. J. (2003, May). Household poverty and vulnerability: A bootstrap approach. Unpublished; available at <http://www.econ.yale.edu/seminars/NEUDC03/Johannes.pdf>.

Kurosaki, T. (2006, January). Vulnerability and poverty: Concept, measurement, and implications to poverty reduction policies in Asia. Unpublished; available from <http://www.ier.hit-u.ac.jp/~kurosaki/vuln0601.pdf>.

Ligon, E. (1998). Risk-sharing and information in village economies. *Review of Economic Studies* 65, 847–864.

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

Ligon, E. and L. Schechter (2004, May). Evaluating different approaches to estimating vulnerability. Social Protection Discussion

Paper 0210, World Bank, Washington, DC.

Peters, R. K. and F. M. A. Hassan (1995). Social safety net and the poor during the transition: The case of Bulgaria. World Bank Policy Research Working Paper 1450.

Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica* 53, 69–76.

Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2, 225–243.

Runkle, D. E. (1991). Liquidity constraints and the permanent-income hypothesis: Evidence from panel data. *Journal of Monetary Economics* 27, 73–98.

Skoufias, E. (2001). Risk sharing in a transition economy: Evidence from monthly data in Bulgaria. Unpublished Manuscript.

Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.

Zimmerman, F. and M. Carter (2003). Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints. *Journal of Development Economics* 71(2), 233–260.

DEPARTMENT OF AGRICULTURAL AND RESOURCE ECONOMICS, UNIVERSITY OF CALIFORNIA, BERKELEY