

MEASURING RISK BY LOOKING AT CHANGES IN INEQUALITY: VULNERABILITY IN ECUADOR

ETHAN LIGON

ABSTRACT. We describe a measure of welfare, “vulnerability”, which measures the difference between the highest feasible average level of utility in a population given aggregate resources, and the actual average level of utility. This measure can be decomposed into two components, related to inequality and to risk. We provide methods for computing vulnerability, inequality, and risk using only data on expenditures from repeated cross-sections of household data, and relate these to Atkinson’s family of inequality measures.

Using methods developed here and household-level Ecuadorean data from 1995 and 2006, we estimate the vulnerability and risk of different population groups. Taking the population altogether, we find that the crisis of the late nineties was not only a large shock for the country as a whole, but also greatly increased the risk faced by individual households in the Sierra, risk which was subsequently translated into greater inequality. After 1999, overall risk borne by the average household fell dramatically, with the consequence that inequality remained nearly constant from 1999–2006. Levels of rural risk are considerably greater than are urban; further, rural risks tend to be the consequence of spatial shocks, while urban risks are much more idiosyncratic in nature.

1. INTRODUCTION

We describe a measure of welfare meant to capture the social costs of both inequality and risk, which we term *vulnerability*. Previous efforts to measure vulnerability have generally relied on panel household datasets (Ligon and Schechter, 2003) to identify the risk faced by households. Here we describe an alternative simple method which allows one to estimate risk by relying only on repeated cross-sections. We apply these methods to Ecuador during a period in which that country experienced considerable political and macroeconomic instability.

Date: October 14, 2009

Printed October 15, 2009.

I thank Carla Calero, Melissa Hidrobo, and Steven Buck for help with the ECV data; Paul Glewwe and Laura Schechter for their comments; and Luis Tejerina for spurring me to tackle this project.

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

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

In recent years a number of researchers have sought to define and measure something called “vulnerability.”¹ These efforts fall into one of two groups. The first uses data on shocks (e.g., variation in income, illness, or employment status) to try and account for variation in household consumption expenditures—where a shock of one sort or another explains a significant proportion of the variation in consumption, the household is said to be ‘vulnerable’ (Amin et al., 2000; Glewwe and Hall, 1998; Dercon and Krishnan, 2000). The second group also begins (sometimes implicitly) with an attempt to account for variation in consumption, but then goes farther—by estimating not just the conditional mean of consumption but also its distribution, one can then estimate the expected value of nonlinear functions of household consumption meant to measure the welfare losses associated with variation in consumption [e.g., various poverty measures (Calvo and Dercon, 2003; Chaudhuri, 2001; Chaudhuri et al., 2001; Christiaensen and Boisvert, 2000; Pritchett et al., 2000; Kamanou and Morduch, 2001; Ravallion, 1988; Jalan and Ravallion, 1999, 2000), or household utility (Ligon and Schechter, 2003; Elbers and Gunning, 2003)]. This paper adopts the utilitarian approach.

1.2. How is Vulnerability Measured? Here we begin by specifying a simple model describing the problem facing a particular household. While the model itself is quite special, we'll use the model to illustrate

¹For more complete overviews of different means of quantifying vulnerability see Kamanou and Morduch (2001) and Hoddinott and Quisumbing (2003).

features of household behavior which would also obtain in a fairly wide class of models.

We 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. The household is forward looking: at time t , it makes forecasts about its future resources in all future periods, and based on these forecasts may engage in behavior to try and manage the risk associated with future shocks (e.g., saving, modifying its production practices, establishing relationships which may provide future insurance). Then at time t and conditional on the actions taken by the household to manage its risk, the households' expected future utility in period $t + j$ is given by

$$\mathbb{E}_t U(c_{t+j}) = \int U(c_{t+j}) dF_t(c_{t+j}),$$

where $U(c)$ is the household's momentary utility given a consumption realization c , and where $F_t(c_{t+j})$ describes the household's beliefs at time t regarding the distribution of time $t + j$ consumption. To repeat, this distribution may depend on actions taken by the household—in particular, savings decisions made in earlier period will help to determine F_t . The *vulnerability* of household i at t depends on risk faced in all future periods $t + j$, and can be expressed as

$$V_t^i = \frac{1}{T-t} \sum_{j=1}^{T-t} V_{t,t+j}^i = U(\bar{c}) - \frac{1}{T-t} \sum_{j=1}^{T-t} \mathbb{E}_t U(c_{it+j})$$

where \bar{c} is per capita consumption expenditures. Thus, $V_{t,t+j}^i$ is the difference between the utility household i would experience if it received a *per capita* level of consumption \bar{c} at every time $t + j$ with certainty minus the expected utility it actually receives in that period. Vulnerability over the subsequent $T - t$ periods is simply defined to be the average per-period vulnerability.

In any given period, the difference between these two levels of utility stem from inequality and risk, so it's worth noting that per-period vulnerability $V_{t,t+j}^i$ may be re-written as

$$(1) \quad V_{t,t+j}^i = [U(\bar{c}) - U(\mathbb{E}_t c_{it+j})] + [U(\mathbb{E}_t c_{it+j}) - \mathbb{E}_t U(c_{it+j})].$$

Here the first bracketed term is related to the relative wealth of the household, and the second to the *risk* borne by the household at time $t + j$, both viewed from the perspective of time t . As the distribution of future consumption 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 decrease future consumption. Rather, levels of vulnerability, wealth, and risk are what is experienced by the household *after* one takes into account whatever stratagems the household has employed to improve its welfare.

So far we've described a concept of vulnerability—vulnerability is the *feasible* average level of expected utility in the population minus the *actual* average level of expected utility. We're still some ways from actually being able to operationalize such a measure, however. To get from concept to operation, we need to be able to estimate objects such as the per-period vulnerability in (1). And for this, we need to (i) specify the utility function U ; (ii) estimate the highest feasible fixed per capita consumption, \bar{c} ; (iii) compute time t forecasts of time $t + j$ expected utility, $E c_{it+j}$, and (iv) compute time t forecasts of time $t + j$ utility, $E_t U(c_{it+j})$. More specifically, accomplishing steps (i)-(iii) will allow us to make statements about a component of our vulnerability measure we'll term *inequality*, while accomplishing step (iv) will allow us to characterize *risk*, the remaining component of vulnerability.

1.3. Inequality and Risk in Ecuador.

Add me!

2. DATA

For the application of this paper, we use data from five sources. Four of these sources are from cross-sectional household surveys, the second through fifth rounds of the *Encuesta de Condiciones de Vida* (ECV), conducted by the Ecuadorean Instituto Nacional de Estadística y Censos in 1995, 1998, 1999, and 2006. From these we obtain data on the distribution of fortnightly expenditures across households for a clustered, randomly selected sample of approximately five thousand households in rounds 2 through 4, along with sampling weights which are proportional to the reciprocal of the *ex ante* probability of a selected household being included in the sample. The fifth round of the ECV was conducted in 2005–06 (for brevity's sake, throughout this document we will refer to this as the 2006 ECV). This was a scaled-up version of the earlier rounds of the ECV, with 13,581 households providing data on fortnightly per capita expenditures.

A key variable for our purposes is a measure of total household expenditures. We report figures for per capita expenditures in Table 1. The units in this table are in thousands of current Sucres; the crisis and

Group	1995	1998	1999	2006
All	123.284	233.526	952.799	1721.632
Costa	121.722	213.323	884.137	1512.649
Sierra	127.256	262.992	1027.293	1990.091
Amazonia	94.751	170.091	—	1210.114
Urban	152.141	295.236	1161.392	2102.359
Rural	75.393	132.822	583.995	969.143

TABLE 1. Mean per capita expenditures, by year. Measured in thousands of current Sucres (in 1999 and after, there are 25,000 Sucres per US Dollar).

hyperinflation of 1998–1999 and subsequent ‘dollarization’ mean that one should interpret levels and changes in levels over time in this table only with great caution. However, one can see some meaningful patterns in averages across groups, and in changes in these patterns over time. For example, expenditures in the coastal region (the “Costa”) in 1995 were only four percent less than expenditures in the Sierra, while by 1998 the difference in expenditures between these two groups had increased to 23 per cent. The proportional difference between these groups fell to 13 percent in 1999, but then increased substantially to a difference of 31 percent in 2006.

Tracing the fortunes of the poorest region of the country, the Amazonian *Oriente* is made somewhat difficult by changes in the way this region was sampled across years. The biggest problem is that the 1999 round of the survey simply neglected this region altogether, but the treatment and sampling strategy used in this vast, relatively inaccessible region has consistently posed difficulties for surveyors. Nonetheless, a story can be told: per capita household expenditures in the *Oriente* were 34 percent less than expenditures in the Sierra and 28 percent less than the *Costa* in 1995. In 1998 the gap between the *Oriente* and the *Costa* narrowed to 18 percent, while the gap with the *Sierra* grew to 45 percent. By 2006 the Amazon region had fallen farther behind both of the other regions, with per capita household expenditures 23 percent below those in the *Costa* and fully 62 percent below those in the *Sierra*.

To some extent variation in household expenditures across regions may simply reflect differences in the rural-urban composition of the different regions, and differences in the economic shocks impacting rural and urban households. In both the *Sierra* and the *Costa*, roughly one quarter of all households were classified as urban, while in the *Oriente*

	1995	2006
Total observations	5812	13581
Unique Clusters	55	443
Common Clusters	49	49
Obs. in common clusters	5396	6331
Observations in common (Cluster,Zones)		
Unique	341	453
Common	119	119
Households in common	2308	1679
Observations in common (Cluster,Sector)		
Unique	281	229
Common	106	106
Households in common	3177	3924

TABLE 2. Matching Sub-populations Across the 1995 and 2006 Rounds of the ECV.

fewer than six percent were. However, the considerable variation in relative expenditures across the *Sierra* and the *Costa*, with their similarly composed rural-urban population means that the crisis of the late nineties and subsequent growth in this decade must have had important regional effects, even after controlling for differences in rural-urban composition of the three regions.

It's important for our purposes to keep careful track of the sampling scheme for the ECV, since we care about the distribution of our estimates of functions of the distribution of expenditures across the population. In brief, the census bureau divides the country into a collection of 20 dominions. Within each dominion, households are assigned to a particular census region, which we'll call the UPM. Any given UPM can be further disaggregated into provinces, cantons, parroquia, and zones.

The UPM is the primary sampling unit. The population of UPMs within each dominion is sampled, with a probability proportional to its population share within the dominion.² Each UPM is comprised of many sectors, each with between 50 and 200 households. Within each sampled UPM, a single sector is randomly chosen—thus, we can regard the sampling ‘cluster’ as either the UPM or the sector. Within

²But note that households selected from UPMs in the two principal cities of Quito and Guayaquil are ordered according to some socio-economic criteria, and then the UPMs are divided into three strata (“alto,” “medio,” “bajo”). The alto strata is made up of the best-off 30 percent of UPMs; the medio the next 40 percent, and the bajo the final 30 percent.

each sampled UPM-sector 18 households are randomly selected, and twelve of these are interviewed (the extra six households help to cover for refusals or other non-response). Whether by accident or design, some but not all of the UPM clusters sampled in 1995 are included in later samples; for example, of 55 clusters included in the 1995 sample, only 49 appear in the 2006 sample (see Table 2).

In addition to the household level data in the various rounds of the ECV, we use data from series collected by the International Monetary Fund in its *International Financial Statistics* series on the price level and on aggregate household expenditures in order to construct estimates of growth rates in real per capita expenditures.³

3. VULNERABILITY

In this section we describe a sequence of possible expenditure allocations, and compute the welfare of the average Ecuadorean household under each.

We move from more utopian allocations to less utopian. We begin imagining a society with no inequality, no growth,⁴ and no risk, and then add these imperfections sequentially.

Consider a population of n households, indexed by $i = P = \{1, 2, \dots, n\}$. Each household i consumes c_{it} at date t .

We assume that households possess von Neumann-Morgenstern time-separable preferences, and that each household has a common CES utility function of the form $U(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, where the parameter γ can be interpreted as the households' common coefficient of relative risk aversion. Future utility is discounted at a common rate $1/\beta - 1$.

We'll want to be able to compare inequality and risk across different sub-groups. Accordingly, we'll need to develop a notation to allow us to specify different groups. Let $G \subseteq P$ be a set of whole numbers indexing some set of households in the population. Related, let $\#G$ denote the number of elements of the set G so that, for example, we have $\#P = n$.

3.1. Utopian. Imagine a world in which all households had equal expenditures in every period—that is, a world with no inequality, no

³The main IFS series we use is that for household consumption expenditures over 1995–2004 (series 24896F.DZF), which is already measured in constant currency units.

⁴It may seem good for welfare to have growth in expenditures over time. However, while it's desirable to have *income* grow over time, if future income is high then society would be better off if it could borrow against this future income so as to make expenditures constant.

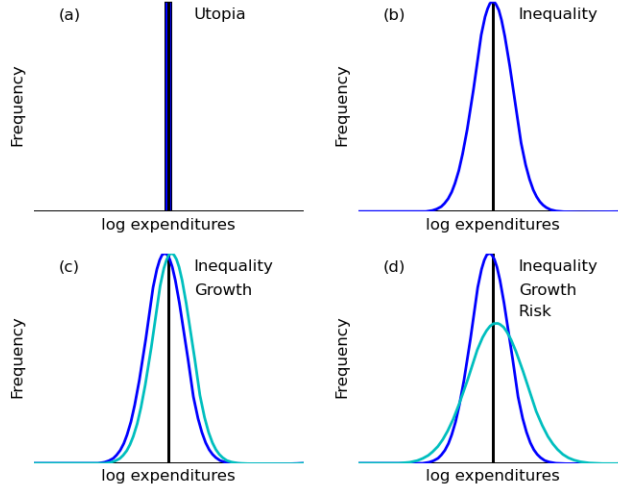


FIGURE 1. Decomposition of Vulnerability. Panel (a) provides a histogram of the logarithm of household expenditures in utopia (no inequality; no growth; no risk). Panel (b) provides a contrast in which there’s no growth and no risk (and hence no change over time), but there is inequality. Panel (c) adds growth; Panel (d) adds risk.

growth, and no risk. Since there is no growth and no risk, household i ’s consumption at any time t must be equal to its consumption at any other time $t + j$; accordingly, we can write household i ’s utopian consumption in *any* period simply as c_i . Since in addition each household is assumed to have identical preferences, as described above, then if all households are able to use credit markets, then

$$U'(c_i) = \beta^j R_{t,t+j} U'(c_i).$$

For a stationary allocation such as the one we’ve described to be optimal, such an equation must be satisfied for all t and j , and so it follows that the risk-free return between any two periods t and $t + 1$ must be a constant R , with $R = 1/\beta$.

One could imagine society implementing the utopian allocation we’ve described if three (implausible) conditions are met: first, that households are able to contract from behind a “veil of ignorance,” as in the thought experiment of Rawls (1971) or Harsanyi (1955) (thus eliminating inequality); that households are subsequently able to fully insure their consumption expenditures (Borch, 1962; Mace, 1991; Townsend,

1994); and that society is collectively able to borrow and lend at the (common) rate of time preference.

Since the optimal allocation in this setting involves equal expenditures across both households and time, one way to measure deviations from this optimum is to establish equality as a benchmark. Thus, with periods $t = 1, \dots, T$ and a population of households indexed by $i \in P = \{1, \dots, n\}$, the stationary, egalitarian allocation is the unique maximizing allocation of the expression

$$(2) \quad \frac{1}{T} \sum_{t=1}^T \frac{1}{n} \sum_{i=1}^n U(c_{it})$$

for any increasing, concave function U . Adopting the constant elasticity of substitution utility function described above, this becomes

$$(3) \quad W(P) = \frac{1}{1-\gamma} \left[\frac{1}{T} \sum_{t=1}^T \frac{1}{n} \sum_{i=1}^n c_{it}^{1-\gamma} - 1 \right].$$

We normalize the largest level of egalitarian, stationary expenditures to be one. Thus, at the optimum we have

$$W^0(P) = \frac{1}{1-\gamma} \left[\frac{1}{T} \sum_{t=1}^T \frac{1}{n} \sum_{i=1}^n (1)^{1-\gamma} - 1 \right] = 0,$$

which we adopt as our utopian benchmark level of welfare.

Figure 1 shows the utopian distribution of expenditures across time and the population in panel (a); all households have identical levels of expenditures in every period.

3.2. Inequality. In this section, we begin our slouch toward dystopia by describing an allocation of expenditures which permits inequality across households, while preserving our assumption that expenditures are constant across time. The resulting distribution of expenditures across the population is illustrated in panel (b); there's now variation across the population, with some households having lower expenditures and others having higher expenditures than the average in panel (a). This inequality is harmful to social welfare, since the value of an additional unit of expenditures is presumed to be greater for poorer households than it is to wealthier households. To imagine society implementing an allocation that involves inequality but no variation in household-level expenditures over time we would require 'only' that households could fully insure, and that society could borrow and lend at the rate of time preference.

Moving from illustration to measurement, let c_i be household i 's expenditures in every period, with a normalization so that mean expenditures across households at any period will be \bar{c} . In this case, our social welfare function (3) takes the value

$$(4) \quad W^I(P) = \frac{1}{1-\gamma} \left[\frac{1}{n} \sum_{i \in P} \left(\frac{c_i}{\bar{c}} \right)^{1-\gamma} - 1 \right].$$

Our measure of inequality can be thought of as the welfare loss associated with moving from the utopian allocation to an allocation with inequality. Thus,

$$\text{Inequality} = W^0 - W^I.$$

As it happens, our measure of inequality in this setting is essentially that of Atkinson (1970). This is a natural consequence of two facts: first, that we follow Atkinson in employing a utilitarian social welfare function; and second, that we also assume a common CES utility function. Atkinson's measure takes the form

$$(5) \quad A^\alpha(P) = 1 - \frac{1}{n\bar{c}} \left[\sum_{i \in P} c_i^{1-\alpha} \right]^{\frac{1}{1-\alpha}}.$$

Comparing this with (4), we have

$$(1-\gamma)W^I(P) + 1 = [1 - A^\gamma(P)]^{1-\gamma};$$

it follows that

$$(6) \quad W^I(P) = \frac{1}{1-\gamma} [(1 - A^\gamma(P))^{1-\gamma} - 1].$$

Our social welfare measure W^I has the interpretation of being the average momentary utility of a household. So long as γ is positive, then W will be an increasing and strictly concave function of each household's consumption. Note that this leaves open the issue of what *period* we're evaluating. But because at the moment we're interested in diagnosing the ills associated with inequality, we're free to assume that consumption expenditures, while unequal, are nonetheless perfectly insured and smoothed over time. Under these nearly utopian assumptions, realized consumption expenditures and the social welfare function will take the same value in every period, so that which period we actually evaluate doesn't matter. A further happy consequence of these assumptions is that if we have data for different households in different periods we can (with the appropriate normalization) use data from all of these in computing our measures of inequality.

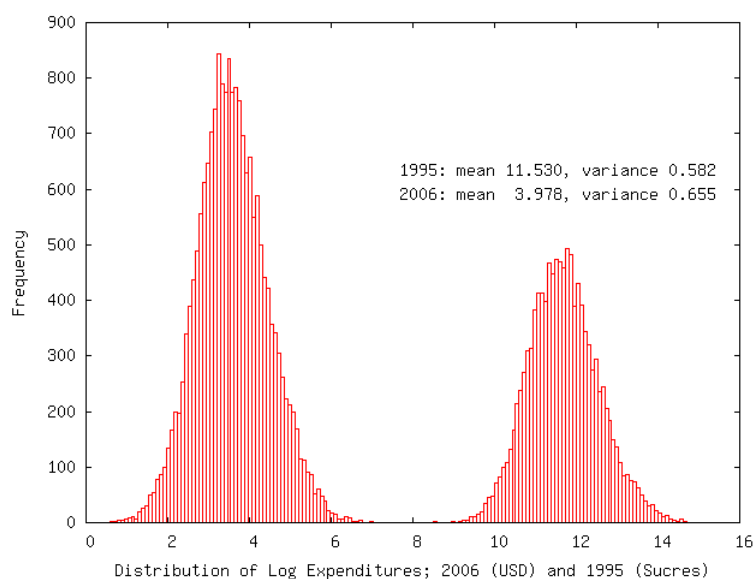


FIGURE 2. Histogram of fortnightly per capita expenditures.

Figure 2 shows the distribution of the logarithm of expenditures in current *Sucres*. The financial crisis and subsequent dollarization which separate the second and fifth rounds of the ECV lead to dramatic changes in the value of the *Sucre* (the official exchange rate at the time of the 2000 dollarization was 25,000 *Sucres* per US dollar, though as the hyperinflation and crisis happened in the late nineties, the value of the 2006 dollar in 1995 *Sucres* is approximately 1100). This dramatic difference in the value of the currency accounts for the two quite discrete distributions for expenditures in 1995 (the ‘hump’ to the right) and 2006 (the hump to the left). The height of the 2006 hump is much greater than the height of the 1995 hump; however, this is more a reflection of the larger sample size of the 2006 survey than a consequence of any change in the distribution of fortnightly expenditures. Though changes in the value of the *Sucre* and economic growth produce a large change in the mean of log expenditures (11.4 in 1995 to 4.05 in 2006), there is relatively little change in the variance of log expenditures across the two years.

So, to compute the level of inequality, we take data on household expenditures in each of the rounds of the ECV, and normalize these

Group	$\gamma = 2$			
	1995	1998	1999	2006
Pooled	0.454*** (0.007)	0.491*** (0.007)	0.495*** (0.008)	0.513*** (0.004)
Costa	0.383*** (0.009)	0.383*** (0.009)	0.370*** (0.010)	0.350*** (0.006)
Sierra	0.528*** (0.010)	0.599*** (0.010)	0.596*** (0.010)	0.616*** (0.006)
Amazonia	0.300*** (0.024)	0.182*** (0.018)	—	0.488*** (0.015)

TABLE 3. Atkinson inequality measures, by year and group. Figures in parentheses are standard errors. The notation “***” indicates that the reported figures are significantly different from zero at a 99 per cent confidence level.

by the mean expenditures in each of those rounds.⁵ With this normalization, if there was perfect equality then evaluating (3) would yield $W^0(P) = 0$ for any i . Taking this as a benchmark, our measure of inequality for the population becomes simply $-W^I(P)$. The normalized consumptions used to compute the value of this expression average one by construction, so any inequality in the distribution of expenditures (c_1, \dots, c_n) will be a mean-preserving spread of $(1, 1, \dots, 1)$. This fact combined with the concavity of W implies that inequality will be non-negative by construction. Considered simply as a measure of inequality, this measure both gives a complete ordering of possible expenditure distributions complete and is Lorenz-consistent.

Table 3 presents estimates of inequality across different regions and years, using the Atkinson inequality measure given in (5), with an inequality aversion parameter equal to two. Standard errors of these estimates are calculated using a method described by Biewen and Jenkins (2006) to take into account the effects of clustering and stratification. Taking all regions together, there’s an increase in inequality from 1995 to 1999, but inequality in the pooled population remains constant from 1999 to 2006. However, this increase in inequality over time is driven mainly by increases in inequality in a single region—the Sierra.

⁵Estimators of mean expenditures and of inequality statistics use sampling weights to make the estimates representative of the Ecuadorean population as a whole.

Inequality in the other economically important region, the Costa, actually falls from 1999 to 2006. Measured inequality in the Amazon region varies erratically over time (and data for this region is missing for 1999) but the relatively small part of the sample drawn from the Amazon (approximately 7 per cent) has only a small effect on the pooled estimates of inequality.

Inequality varies much more across regions than it does over time in our sample. In every year, inequality is greatest in the Sierra.

3.3. Growth. We'll think of aggregate economic growth as changes in mean expenditures which are known in advance. Reflecting a long tradition in growth economics, in considering the effects of growth on welfare we assume that changes in aggregate expenditures have no effect at all on individuals' shares of those aggregates—thus, that the distribution of resources is unaffected by changes in aggregate expenditures. This is illustrated in panel (c) of Figure 1; though there's inequality, if there's growth then all households' expenditures increase at a common rate, so that expenditures start at the left-most distribution, and move to the right. Though growth *per se* isn't harmful to welfare, every household would prefer the situation in panel (b); in panel (c) expenditures grow only because they're unable to borrow against future income. Though future expenditures are higher than they'd otherwise be, this comes at the expense of current expenditures being lower.

Though we have data on the distribution of household expenditures in four periods (1995, 1998, 1999 and 2006), these are not actually data we need for our immediate end. Since we're presently maintaining the hypothesis that the distribution of resources is unchanging over time, we can use the measure of Atkinson's measure of inequality $A^\alpha(P)$ given above, and simply compute welfare in any given period using an Atkinson social welfare function, which depends only on Atkinson's inequality measure and on estimates of aggregate expenditure growth.

To obtain such a measure, rather than normalizing by aggregate expenditures in every period, we instead normalize using expenditures in 1995. Figure 3 reports the time series. The same figure also reports our measure of welfare which varies across years when variation is perfectly predictable,

$$W_t(P) = \frac{1}{(1-\gamma)} \left(\frac{1}{n} \sum_{i=1}^n c_{it}^{1-\gamma} - 1 \right).$$

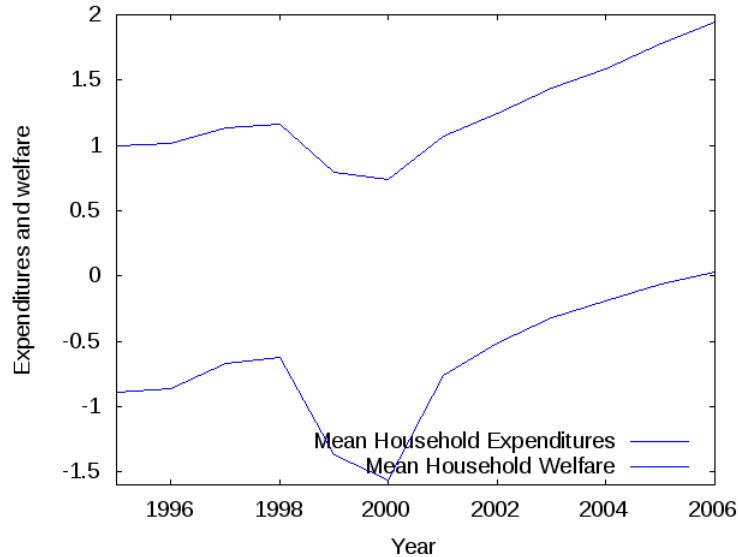


FIGURE 3. Top: Household expenditures by year, normalized by 1995 expenditures. Bottom: Welfare by year, allowing for inequality and changes in aggregate expenditures.

Denote aggregate expenditures at time t by \bar{c}_t . Given our maintained assumptions of full insurance and CES utility, individual i 's expenditures will be a constant share λ_i of the aggregate, so that inequality (measured *à la* Atkinson) will be unchanged across years. Indeed, if social *inequality* aversion α is set equal to individual *risk* aversion γ , then using our parametric assumption that $U(c) = (c^{1-\gamma} - 1)/(1 - \gamma)$, we have

$$(7) \quad W_t(P) = \frac{(1 - A_t^\gamma(P))^{1-\gamma} \bar{c}_t^{1-\gamma} - 1}{1 - \gamma}.$$

This measure allows us to express (momentary) social welfare at time t in terms of a combination of just three numbers: Aggregate expenditures \bar{c}_t , Atkinson's inequality $A_t^\gamma(P)$, and the (social) preference parameter γ .⁶ It is this expression which occupies the lower portion of Figure 3.

⁶A possible justification for setting the inequality aversion parameter equal to γ is given by an argument by Harsanyi (1955), anticipating Rawls (1971): Inequality aversion should be equal to risk aversion if individuals' expenditures in life are the outcome of a (fair) lottery entered into before birth.

We will find it useful to also characterize the rate of growth across periods. Let $\mu_{t,t+j} = \log\left(\frac{\bar{c}_{t+j}}{\bar{c}_t}\right)$. When households face no risk, this is equal to the average rate of growth of the reciprocal of marginal utility of consumption across the population. A complete collection of estimates of $\mu_{t,t+j}$ for our data is presented below the diagonal of the matrix in Table 4. For example, taking $t = 1995$, $t + j = 2006$, and $\gamma = 2$, data from the IFS on household expenditures implies that $\mu_{1995,2006} = 0.666$.

Year	1995	1998	1999	2006
1995	—	0.090	0.112	0.160
1998	0.153	—	0.022	0.070
1999	-0.226	-0.379	—	0.048
2006	0.666	0.512	0.891	—

TABLE 4. Growth and Variance. Figures reported below the diagonal are estimates of the growth rates $\mu_{t,t+j}$, while figures above the diagonal are estimates of the variances $\sigma_{t,t+j}^2$. For growth rates, column heads indicate the base year t , while row heads indicate the final year. For example $\mu_{1995,1999}$ can be found in the first column and third row. Conversely, for variance estimates, row heads indicate the base year while column heads indicate the final year.

The patterns of growth rates reported below the diagonal of Table 4 reveal the roller-coaster ride followed by the Ecuadorean economy during the late nineties. Real growth in expenditures from 1995–98 averaged a respectable 4.9 per cent, but the onset of the crisis resulted in a massive drop of 37.9 per cent in the year from 1998 to 1999. The story from 1999 to 2006 is one of steady recovery and growth, with average real growth in expenditures from 1999 to 2006 averaging an extraordinary 12.7 per cent per year, resulting in real expenditures in 2006 which were 67 per cent above the level of expenditures in 1995.

When credit markets are complete and common, then all growth is also common. This allows us to distinguish deterministic growth from the effects of shocks on growth. An intuition which may help to justify growth being common is simply that if one household knew that its future expenditures would grow at a faster rate than others' expenditures, then it would have an incentive to increase current expenditures by borrowing from those other households. If credit markets are perfect

and preferences are common, then differences in growth rates must be due to shocks.

3.4. Risk. To the extent that individual households do *not* have access to mechanisms which provide full insurance, then shocks to income will affect the distribution of expenditures since the effects of these shocks may not be shared with the rest of the population. It's only when some households don't have access to adequate insurance that we'd expect them to bear idiosyncratic risk, a situation illustrated in panel (d) of Figure 1. Here there's inequality, as in panel (b), and also growth, as in panel (c). But some households have positive idiosyncratic shocks, and others negative, and these risks aren't pooled. This leads to an increase in inequality over time—the distribution of expenditures moves to the right, with growth, but also 'spreads out'—a shock which affects expenditures today will also tend to affect expenditures tomorrow, per the permanent income hypothesis (Bewley, 1977; Deaton and Paxson, 1994).

In the absence of full insurance, households may still accomplish a considerable amount of smoothing by using credit markets to borrow and lend. Suppose that all households have access to a credit market, and can borrow or lend. A loan extended at time t which is repaid at time $t + j$ will return $R_{t,t+j}$. These returns can vary over time, but we assume for now that markets and returns are common across all households. Consumption satisfies the Euler equation

$$(8) \quad U'(c_{it}) = \beta^j R_{t,t+j} \mathbf{E}_t [U'(c_{it+j})],$$

where \mathbf{E}_t denotes the expectations operator conditioning on information available at time t . Only the variable c_{it+j} is unknown conditional on information available at time t .

3.4.1. Estimating a conditional moment of consumption growth from cross-sectional data. The term $\beta^j R_{t,t+j}$ in (8) is common to all households, so that

$$\sum_{i=1}^n U'(c_{it}) = \beta^j R_{t,t+j} \mathbf{E}_t \left[\sum_{i=1}^n U'(c_{it+j}) \right].$$

It follows that the expression $\beta^j R_{t,t+j}$ must be equal to a ratio involving the distribution of current and future consumption expenditures,

$$\beta^j R_{t,t+j} = \frac{\sum_{i=1}^n U'(c_{it})}{\mathbf{E}_t [\sum_{i=1}^n U'(c_{it+j})]}.$$

However, equation (8) also implies that

$$(9) \quad \frac{1}{\beta^j R_{t,t+j}} = \frac{1}{n} \mathbb{E}_t \left[\sum_{i=1}^n \frac{U'(c_{it+j})}{U'(c_{it})} \right],$$

so that

$$\frac{\mathbb{E}_t \sum_{i=1}^n U'(c_{it+j})}{\sum_{i=1}^n U'(c_{it})} = \frac{1}{n} \mathbb{E}_t \left[\sum_{i=1}^n \frac{U'(c_{it+j})}{U'(c_{it})} \right].$$

This last expression relates the expected growth in marginal utilities to a ratio which involves only cross-sectional moments. If we take the further step of invoking our parametric assumption that the utility equation takes the constant elasticity of substitution (CES) form (with coefficient of relative risk aversion γ), then we have

$$(10) \quad \frac{\mathbb{E}_t \sum_{i=1}^n c_{it+j}^{-\gamma}}{\sum_{i=1}^n c_{it}^{-\gamma}} = \mathbb{E}_t \sum_{i=1}^n \left(\frac{c_{it}}{c_{it+j}} \right)^\gamma.$$

Equation (9) and the expression (10) also allows us to relate inequality and growth rates for different subgroups. To see this, suppose that we observe consumption expenditures at t and $t+j$ for an a groups of households with an index set G_1 , and for a second group with an index set G_2 . Then we have

$$(11) \quad \frac{\mathbb{E}_t \sum_{i \in G_1} c_{it+j}^{-\gamma}}{\sum_{i \in G_1} c_{it}^{-\gamma}} = \frac{1}{\#G_2} \mathbb{E}_t \sum_{i \in G_2} \left(\frac{c_{it}}{c_{it+j}} \right)^\gamma.$$

The beauty of this expression is that in situations in which we don't have the panel data necessary to measure expenditure growth for some group G_2 at the household level (i.e., the data one would need to construct a sample analog to the right-hand side of (11)), we can still get at this moment using only repeated cross-sectional data on a (possibly) different group G_1 (i.e., the data one would need to construct a sample analog the the left-hand side of (11)), provided only that households in both groups had access to the same credit markets.

Let

$$(12) \quad M_{t,t+j}^b(G) = \frac{\sum_{i \in G} c_{it+j}^b}{\sum_{i \in G} c_{it}^b}.$$

This is simply the ratio of the b th cross-sectional sample moment at time $t+j$ to the b th cross-sectional sample moment at time t . Then $M_{t,t+j}^{-\gamma}(G_1)$ is the obvious sample analog to the left-hand side of (11), provided that the $t+j$ cross-sectional moment is known at time t . Note that this does *not* mean that any given individual knows the trajectory

of their future consumption growth—rather, it means that all individuals share similar and correct beliefs about the distribution from which future consumption realizations will be drawn (see Kocherlakota and Pistaferri, 2009, for a similar restriction). In any event, we have

$$(13) \quad M_{t,t+j}^{-\gamma}(G_1) = \frac{\frac{1}{\#G_1} \sum_{i \in G_1} c_{it+j}^{-\gamma}}{\frac{1}{\#G_1} \sum_{i \in G_1} c_{it}^{-\gamma}}$$

Note that no *panel* data is required here; only repeated cross-sections. From (11), this calculation then provides an estimate of the average intertemporal marginal rate of substitution for households in G_2 even in the complete absence of data on consumption growth rates for households in G_2 .

This trick of relating a ratio of cross-sectional moments for one group to the average intertemporal marginal rate of substitution for another is the key which allows us to draw inferences regarding risk, even in the absence of direct data on household-specific rates of expenditure growth.

3.4.2. Estimating the Distribution of Future Consumption. Calculating $M_{t,t+j}^{-\gamma}$ gives us an estimate of the average conditional moment of household-level consumption growth, even when we don't have a panel of households. However, to calculate the risk faced by the household we will typically need to know the probability distribution of future consumption, not just a single moment of this distribution. To estimate this probability distribution additional assumptions are required.

Accordingly, let $c_{it+j} = c_{it}e^{\epsilon_{it+j}}$, where ϵ_{it+j} is a continuously distributed random variable with a probability density function $f(\epsilon|\theta_{t,t+j})$, where $\theta_{t,t+j}$ is a vector of (possibly unknown) parameters which may vary across pairs of periods $(t, t+j)$. As this notation suggests, we also assume for any $(t, t+j)$ that ϵ_{it+j} is conditionally independently and identically distributed across the population.

We further assume that the support of ϵ_{it+j} is \mathbb{R} . This assumption regarding support along with our assumption that ϵ_{it+j} is independently and identically distributed implies a simple moment restriction relating aggregate growth in expenditures to individual growth,

$$(14) \quad \mathbb{E}_t \sum_{i \in G} c_{it+j} = \sum_{i \in G} c_{it} \mathbb{E}_t e^{\epsilon_{it+j}}$$

which yields the sample counterpart

$$(15) \quad M_{t,t+j}^1(G) = \int f(\log m|\theta_{t,t+j}) dm.$$

Next, from (11) one can see that we want to be able to compute time t expectations of growth rates of the marginal utility of consumption $m_{it+j} = \left(\frac{c_{it}}{c_{it+j}}\right)^\gamma = e^{-\gamma\epsilon_{it+j}}$. Using the usual “inverse Jacobian” approach to computing the probability density of a function of a continuous random variable, it follows that we have

$$E_t m_{it+j} = \frac{1}{\gamma} \int_0^\infty f\left(-\frac{1}{\gamma} \log m \mid \theta_{t,t+j}\right) dm.$$

Conditional on knowing the density f and the vector of parameters $\theta_{t,t+j}$ one could simply compute the conditional expectations which appear in (11). However, in the more usual case even if one assumes a particular density function f , one will still not know the parameters $\theta_{t,t+j}$. In this case, we use the moment restriction (11) and the sample counterpart of its left-hand side (13) to estimate the unknown parameter. The parameter vector $\theta_{t,t+j}$ satisfies the pair of moment conditions

$$(16) \quad E_t \frac{\sum_{i \in P} c_{it+j}}{\sum_{i \in P} c_{it}} = \int f(\log m \mid \theta_{t,t+j}) dm$$

and

$$(17) \quad \frac{E_t \sum_{i \in P} c_{it+j}^{-\gamma}}{\sum_{i \in P} c_{it}^{-\gamma}} = \frac{1}{\gamma} \int_0^\infty f\left(-\frac{1}{\gamma} \log m \mid \theta_{t,t+j}\right) dm.$$

Thus, if one knows the values of the left-hand side of these two moment conditions and $\theta_{t,t+j}$ has just two elements, one can simply compute these by solving the system (16) and (17) for $\theta_{t,t+j}$. More typically, of course, one won't know the value of the population moments that appear on the left-hand-side of these moment conditions, but these can be estimated using the obvious sample moment conditions

$$(18) \quad M_{t,t+j}^1(G) = \int_0^\infty f(\log m \mid \theta_{t,t+j}) dm$$

and

$$(19) \quad M_{t,t+j}^{-\gamma}(G) = \frac{1}{\gamma} \int_0^\infty f\left(-\frac{1}{\gamma} \log m \mid \theta_{t,t+j}\right) dm$$

for some sample G of the population.

At this point we must confront the awkward fact that the population changes over time. What we really want is a sample G drawn from the time t population, observed at both t and $t + j$. But this would typically require a longitudinal panel, and if we had such a panel we could take a more direct approach to estimating household-level consumption dynamics. The case which interests us here is the case in which we have

only repeated cross-sections at t and $t + j$. Sample households in such cross-sections will ordinarily be chosen so as to be representative of the current population, so that a sample of households at time t , G_t will represent a somewhat different population than a sample of households drawn at time $t + j$.

In this situation, we can resort to constructing a “pseudo-panel” of cohorts, along lines suggested by Deaton (1985). Our approach here will be to simply calculate the empirical distribution of the ages of household heads in year t , letting α_a denote the proportion of the heads of households in year t aged a years. Then to estimate the b th cross-sectional sample moment of household consumption for P_t in year $t + j$, we use

$$C_{t+j}^b(G_t) = \sum_a \alpha_a \frac{\sum_{i \in G_{t+j}} \mathbf{1}\{\text{age}_i = a + j\} c_{it+j}^b}{\sum_{i \in G_{t+j}} \mathbf{1}\{\text{age}_i = a + j\}},$$

where $\mathbf{1}$ is the Boolean indicator function, and age_i is the age of the head of the i th household.

Using this approach, our sample moment conditions become

$$(20) \quad \frac{C_{t+j}^1(G_t)}{C_t^1(G_t)} = \int_0^\infty f(\log m | \theta_{t,t+j}) dm$$

and

$$(21) \quad \frac{C_{t+j}^{-\gamma}(G_t)}{C_t^{-\gamma}(G_t)} = \frac{1}{\gamma} \int_0^\infty f\left(-\frac{1}{\gamma} \log m \middle| \theta_{t,t+j}\right) dm.$$

In the special case in which ϵ_{it+j} is normally distributed it’s possible to solve the moment conditions (18) and (21) on the back of an envelope—no integration is required. Let the mean of ϵ_{it+j} be given by $\mu_{t,t+j}$, and its variance by $\sigma_{t,t+j}^2$. We compute the mean directly, with $\mu_{t,t+j} = \log\left(\frac{\bar{c}_{t+j}}{\bar{c}_t}\right)$. This leaves the parameter σ to estimate. Exploiting our distributional assumptions,

$$\mu_{t,t+j} + \frac{\sigma_{t,t+j}^2}{2} = \frac{-1}{\gamma} \log M_{t,t+j}^{-\gamma}(G).$$

Since σ is a sufficient statistic for all the variation that a given household faces in expenditures, this is a very useful result, which allows us to relate the *risk* an individual household faces to *changes* in the cross-sectional moments of the consumption distribution.

The situation in which households faced *no* risk would imply that $\sigma = 0$. In the no-risk case, expected utility for household i at time $t + j$

can be written

$$E_t U(c_{it+j}) = U(E_t c_{it+j}) = \frac{\left[c_{it} e^{\mu_{t,t+j} + \sigma_{t,t+j}^2/2} \right]^{1-\gamma}}{1-\gamma} - \frac{1}{1-\gamma},$$

while expected utility for the same household *with* risk can be written

$$E_t U(c_{it+j}) = \frac{c_{it}^{1-\gamma}}{1-\gamma} e^{(1-\gamma)(\mu_{t,t+j} + \frac{1-\gamma}{2}\sigma_{t,t+j}^2)} - \frac{1}{1-\gamma}.$$

The welfare cost of risk is simply the first of these less the second, or

$$(22) \quad \text{Risk}_{t,t+j}^i = U(E_t c_{it+j}) - E_t U(c_{it+j}) \\ = \frac{c_{it}^{1-\gamma}}{1-\gamma} e^{(1-\gamma)(\mu_{t,t+j} + \frac{1-\gamma}{2}\sigma_{t,t+j}^2)} \left[e^{-\gamma \frac{\sigma^2}{2}} - e^{\gamma \frac{\sigma^2}{2}(\gamma-2)} \right].$$

3.4.3. *Risk from Changing Inequality.* Rearranging (13) gives us

$$(23) \quad M_{t,t+j}(G) = \frac{\frac{1}{\#G} \sum_{i \in G_{t+j}} c_{it+j}^{-\gamma}}{\frac{1}{\#G} \sum_{i \in G} c_{it}^{-\gamma}} \\ = \frac{\frac{1}{\#G} \sum_{i \in G} \left(\frac{c_{it+j}}{\bar{c}_{t+j}} \right)^{-\gamma}}{\frac{1}{\#G} \sum_{i \in G} \left(\frac{c_{it}}{\bar{c}_t} \right)^{-\gamma}} \left(\frac{\bar{c}_t}{\bar{c}_{t+j}} \right)^\gamma.$$

The expression $\sum c_{it}^{-\gamma}$ which appears in equation (10) and elsewhere is very closely related to the inequality measure of Atkinson (1970). To see this, recall that Atkinson's measure is defined up to a parameter α , meant to measure "inequality aversion," and is defined at time t by

$$(24) \quad A_t^\alpha(P) = 1 - \frac{1}{\bar{c}_t} \left[\frac{1}{n} \sum_{i \in P} c_{it}^{1-\alpha} \right]^{\frac{1}{1-\alpha}}.$$

Here the superscript α is a reminder that Atkinson's measure of inequality is a function of the preference parameter α , and the argument P is the index set for all the households in the population, so that $A_t^\alpha(P)$ denotes Atkinson's measure of inequality over the entire population at time t . Rearranging (24) gives us

$$\frac{1}{n} \sum_{i \in P} \left(\frac{c_{it}}{\bar{c}_t} \right)^{1-\alpha} = [1 - A_t^\alpha(P)]^{1-\alpha}.$$

Then, using the relationship between Atkinson’s inequality measure described above, this implies

$$M_{t,t+j}(G) = \left[\frac{1 - A_t^{1+\gamma}(G)}{1 - A_{t+j}^{1+\gamma}(G)} \right]^\gamma \left(\frac{\bar{c}_t}{\bar{c}_{t+j}} \right)^\gamma.$$

Using this, we don’t even need household-level data on consumption—if there are statistics on Atkinson inequality measures for the appropriate values of γ and data on aggregate consumption growth, then one may be able to calculate $M_{t,t+j}(G)$ on the back of an envelope.

If the idiosyncratic shocks ϵ_{it+j} are normally distributed, then the relationship between the variance parameters $\sigma_{t,t+j}^2$ and Atkinson’s measure of inequality is given by

$$(25) \quad \sigma_{t,t+j}^2/2 = \frac{1}{\gamma} \log \left[\frac{1 - A_t^{1+\gamma}(G)}{1 - A_{t+j}^{1+\gamma}(G)} \right].$$

It’s worth noticing at this point that for this estimate of the variance to be well-defined, it *must* be the case that inequality is increasing over time—if $A_t^{1+\gamma} > A_{t+j}^{1+\gamma}$ then this implies that the variance $\sigma_{t,t+j}^2$ is *negative*. It is not, of course, impossible that inequality should decrease over time. If it should seem to do so in a real-world dataset then one should examine the assumptions which lead to the expression (25), perhaps especially the assumption that all households have access to common credit markets, so that the permanent income hypothesis governs intertemporal variation in household-level expenditures (cf. Deaton and Paxson, 1994).

Table 5 presents estimates of $A_{1995}^{1+\gamma}(G_{1995})$ and $A_{2006}^{1+\gamma}(G_{2006})$; other than using a higher value of inequality aversion, the construction of this table mirrors that of Table 3 perfectly. Substituting these into our present expression then yields the result that $\sigma_{1995,2006}^2 = 0.160$.

With this estimate of $\sigma_{1995,2006}^2$ in hand, we’re ready to construct an estimate of the welfare costs of risk households faced in 1995 over the following decade. We can exploit this relationship between $\sigma_{t,t+j}^2$ and our Atkinson inequality measures to express the estimated risk for household i as a function of nothing more than Atkinson inequality measures at time t and $t + j$ and aggregate growth $\mu_{t,t+j}$:

$$\text{Risk}_{t,t+j}^i = \frac{c_{it}^{1-\gamma}}{1-\gamma} e^{(1-\gamma)\mu_{t,t+j}} \left[\left(\frac{1 - A_t^{1+\gamma}(P)}{1 - A_{t+j}^{1+\gamma}(P)} \right)^{\frac{1-\gamma}{\gamma}} - \left(\frac{1 - A_t^{1+\gamma}(P)}{1 - A_{t+j}^{1+\gamma}(P)} \right)^{\frac{(1-\gamma)^2}{\gamma}} \right].$$

Group	$\gamma = 3$			
	1995	1998	1999	2006
Pooled	0.583*** (0.009)	0.619*** (0.008)	0.628*** (0.008)	0.645*** (0.009)
Costa	0.496*** (0.010)	0.512*** (0.012)	0.512*** (0.014)	0.473*** (0.007)
Sierra	0.654*** (0.012)	0.718*** (0.010)	0.711*** (0.010)	0.730*** (0.008)
Amazonia	0.467*** (0.025)	0.356*** (0.019)	—	0.652*** (0.019)

TABLE 5. Atkinson inequality measures, by year and group. Figures in parentheses are standard errors. The notation “***” indicates that the reported figures are significantly different from zero at a 99 per cent confidence level.

4. RESULTS

Using the methods described in previous sections, we now turn our attention to using these methods to develop an understanding of the risk borne by Ecuadorean households over the period 1995–2006. Though it would be possible, of course, to consider other sub-periods, we wish to focus our attention here on the longest period of time available to us. We take $\gamma = 2$ throughout, for two reasons: first, this is a value of relative risk aversion in line with that chosen for many other empirical micro-econometric studies; and second, with $\gamma = 2$ our measure of welfare loss associated with vulnerability or risk turns out to have a very convenient interpretation as the amount (measured as a proportion of current expenditures) the household would be willing to sacrifice to eliminate the source of risk, inequality, or vulnerability. So, for example, a typical household’s measure of risk turns out to be exactly equal to $E\frac{1}{c} - \frac{1}{E_c}$. Since the units of expenditures are normalized so that per capita consumption expenditures in 1995 are equal to one, then when we average across a group of households, that tells us the amount that that group would have been willing to collectively sacrifice as a share of 1995 per capita expenditures.

With this interpretation in mind, we turn our attention to Table 6. The first row, labeled “Decomposition” reports estimates of total vulnerability, equal to 1.117, which is equal to the sum of the reported welfare costs associated with inequality (0.887) and total risk (0.168).

Variables	Vulnerability	Inequality	Risk
Decomposition	1.117*** (0.031)	0.887*** (0.015)	0.168*** (0.003)
Urban	0.357*** (0.024)	0.264*** (0.017)	0.116*** (0.002)
Rural	1.238*** (0.034)	0.860*** (0.024)	0.171*** (0.002)
Indigenous	4.586*** (0.069)	3.668*** (0.047)	0.337*** (0.004)

TABLE 6. Analysis of Vulnerability

Variables	Vulnerability	Inequality	Risk
Decomposition	1.117*** (0.031)	0.887*** (0.015)	0.168*** (0.003)
Quintile 1	2.683*** (0.052)	2.118*** (0.030)	0.286*** (0.003)
Quintile 2	1.452*** (0.035)	1.185*** (0.020)	0.200*** (0.002)
Quintile 3	0.639*** (0.036)	0.470*** (0.020)	0.135*** (0.002)
Quintile 4	0.075** (0.035)	-0.029 (0.020)	0.089*** (0.002)
Quintile 5	-0.470*** (0.035)	-0.511*** (0.020)	0.045*** (0.002)
Indigenous	2.974*** (0.074)	2.297*** (0.042)	0.211*** (0.004)

TABLE 7. Analysis of Vulnerability

Interpreting the inequality figure first, this suggests that a social planner maximizing average utility across households in Ecuador would be willing to sacrifice an average of nearly 89 per cent of average household per capita expenditures if she could instead face the problem of allocating resources beginning with a society in which there was no inequality in per capita household expenditures.

Of course, even in 1995 it would have been impossible to modify initial levels of inequality, unless the social planner had some sort of time machine. The best that one can possibly hope to do is to reduce *future* inequality, and the only way to do this is to try and modify the future probability distribution of expenditures for households, which is

tantamount to changing the risk that they face. Current risk breeds future inequality.

The total risk faced by the average household in 1995 was quite considerable, amounting to 16.8 per cent. Again, the interpretation of this is that the average household would have been willing in 1995 to sacrifice 16.8 per cent of their 1995 expenditures if they could have eliminated all idiosyncratic risk which would affect their expenditures in 2006 (their level of expenditures would still grow with the aggregate economy). This is really quite a large number; for comparison, Ligon and Schechter (2003) find figures for total risk in Bulgaria in the early nineties which are only about half of the Ecuadorean figure, and this was the period of the tumultuous transition from communism in Bulgaria.

The appropriate policy response to dealing with the risk documented in Table 6 depends on the source of the risk, and on what sub-populations are most affected. Table 6 also reports the vulnerability of various sub-groups. So, for example, the total vulnerability of urban households is only 0.357, compared with the Ecuador-wide average of 1.117. Urban households also bear less risk, 11.6 per cent compared with the economy-wide average of 16.8 per cent, but the biggest difference in vulnerability comes from the relatively low inequality for urban households, 26.4 per cent compared with 88.7 per cent. Note that this does not mean that there's less inequality *among* urban households than among rural; the differences here instead principally reflect differences between rural and urban households, and the fact that urban households' consumption expenditures are closer to the per capita figure than are rural households, simply because they're wealthier on average (see Table 1). Similarly, the enormous figure for average inequality for indigenous households (whether urban or rural) is primarily a reflection of their much lower levels of wealth, not a reflection of inequality within the group of indigenous households.

The figures for "Risk" are more usefully comparable across different groups. Here we see that rural households would, on average, be willing to pay 17 per cent of their 1995 expenditures in order to eliminate uncertainty regarding their 2006 expenditures, compared to an average of 16.8 per cent for the entire sample. Urban households face less risk, 11.6 per cent, while once again indigenous households bear the most, with 33.7 per cent.⁷

⁷Our approach in this paper uses changes in inequality to draw inferences regarding risk, and a look back at Table 3 will remind the reader that much the largest change in inequality we observe occurs in the Amazon region—the change is large

Are the very high levels of risk and inequality experienced by rural and indigenous households a consequence of the fact that they're rural or indigenous, or because they're simply poor? Table 7 documents our approach to answering this question. We've divided the sample population in the 1995 round of the survey into quintiles based on household expenditures in this initial round. Table 7 shows that levels of vulnerability fall monotonically by expenditure quintile: the poorest quintile has a vulnerability of 2.683, while the wealthiest has a vulnerability of -0.470 . Most of this variation in vulnerability is due to variation in inequality; this variation is present entirely by construction. However, it's still the case that the poorer the quintile, the greater the risk. Average risk borne by households in the poorest quintile is 0.286, compared with 0.168 for the sample as a whole. The wealthiest quintile's risk is only 0.045; however, it's worth noticing that this level of risk is still highly significant.

Variables	Risk	Regional Risk	Cantonal Risk	Parroquial Risk
Decomposition	0.168*** (0.003)	0.112*** (0.004)	-0.017 *** (0.006)	0.023*** (0.005)
Urban	0.116*** (0.002)	0.047*** (0.004)	-0.044 *** (0.007)	-0.027 *** (0.006)
Rural	0.171*** (0.002)	0.124*** (0.005)	0.011 (0.010)	0.072*** (0.008)
Indigenous	0.337*** (0.004)	0.384*** (0.010)	0.066*** (0.020)	0.132*** (0.017)

TABLE 8. Analysis of Risk

Table 8 breaks down the total risk faced by different groups to risk that can be explained by shocks experienced at the level of the three regions of the country and risk which can be explained by shocks experienced at the level of the Canton and the Parroquia (these do not sum to total risk because a residual idiosyncratic element remains). Overall, regional risk (that is, risks *across* regions, reflecting region-level shocks) account for a welfare cost of 0.112, almost exactly two-thirds

enough that one might be suspicious about the data from this region. What happens to our measure of vulnerability, inequality, and risk if data from the Amazon is elided? It turns out that our measures of pooled inequality and overall vulnerability aren't much changed, but that our estimates of risk are sharply reduced, by about 35 per cent (the figures corresponding to the first row in Table 8 when we leave out observations from the Amazon are (0.11, 0.09, -0.02 , 0.02)). Interestingly, decompositions by urban/rural or expenditure quintile (see below) are less dramatically affected.

Variables	Risk	Regional Risk	Cantonal Risk	Parroquial Risk
Decomposition	0.168*** (0.003)	0.112*** (0.004)	-0.017*** (0.006)	0.023*** (0.005)
Quintile 1	0.286*** (0.003)	0.200*** (0.009)	0.011 (0.019)	0.067*** (0.015)
Quintile 2	0.200*** (0.002)	0.103*** (0.006)	-0.051*** (0.013)	0.015 (0.010)
Quintile 3	0.135*** (0.002)	0.065*** (0.006)	-0.031** (0.013)	0.000 (0.011)
Quintile 4	0.089*** (0.002)	0.041*** (0.006)	-0.020 (0.012)	-0.006 (0.010)
Quintile 5	0.045*** (0.002)	0.026*** (0.006)	-0.018 (0.013)	-0.012 (0.010)
Indigenous	0.211*** (0.004)	0.293*** (0.013)	0.056** (0.026)	0.118*** (0.022)

TABLE 9. Analysis of Risk

of the total. This suggests quite a large degree of segregation among the economies of the *Sierra*, the *Costa*, and the *Oriente*. Conditional on regional risk, average cantonal risk is actually negative, indicating a negative correlation with other sources of shocks, so that which actually helps to provide some insurance to households. Risk at the lowest level of aggregation (the parroquia) adds something to the total; risk across parroquias within a given canton involves an additional welfare cost of 2 per cent on average.

A dissection of risk sources across rural and urban households is also instructive. Rural households in general bear much more risk related to geography, presumably because of the importance of agriculture to many of these households and the importance of spatial shocks to agricultural pursuits. Though total rural risk is 47 per cent greater than urban, rural risk can be entirely accounted for by spatial shocks at either the regional, the cantonal, or the parroquial level—remaining idiosyncratic risk is actually negative, a consequence of a negative correlation between idiosyncratic shocks and more aggregate sorts of risks. In contrast, urban risk is predominately idiosyncratic; though there's a significant positive regional component to risk, total risk associated with spatial shocks (those associated with regional, cantonal, or parroquial risks) are collectively negatively correlated with idiosyncratic shocks; these latter, if not partly insured against by regional shocks to

expenditures, would yield much larger levels of urban risk amounting to 14.0 per cent.

Looking at an alternative dissection of risk, Table 9 reports how different sources of risk affect different (initial) expenditure quintiles. This reveals a number of interesting patterns. First, the poorest quintile not only bears more total risk than any other; it actually bears significantly more risk in every category we consider. The poorest quintile would be willing to sacrifice 20.0 per cent of their expected expenditures in order to eliminate regional risk. They're the only quintile harmed by cantonal risk (for the other quintiles, cantonal variation is negatively correlated with other sources of risk, so that the canton functions as a source of insurance). And they're the only quintile with a statistically significant exposure to parroquial risk.

5. CONCLUSION

In this paper we've devised a set of methods for drawing inferences about the risk households face by using only data on inequality from multiple cross-sections. We apply these methods to the case of Ecuador, where a political and financial crisis in the late nineties seems likely to have had important impacts on household welfare.

We find that while the crisis of the nineties was important for the country as a whole, it had a particularly large impact on the risk faced by households in the country's *Sierra* region, and that this greater risk led to higher levels of inequality post-crisis. Between 1999 (the end of the crisis) and 2006 household-level risk fell dramatically.

Despite important temporal variation in the risk borne by the average Ecuadorean household, risk is not equally shared across households. Rural, indigenous, and poor households are particularly exposed to risk; further, the poorest quintile of households suffers disproportionately from shocks at several different levels of geographical aggregation: regional, cantonal, and parroquial.

There's a lesson for policymakers in these patterns of risk across quintiles. If one could eliminate all sources of risk for the poorest quintile, our measurements suggest that that would improve the welfare of the poorest roughly as much as a thirty per cent increase in their expected expenditures. Further, there's reason to be optimistic that this sort of insurance could be provided at relatively low cost: microfinance, social security, and similar reasonably well-understood mechanisms seem likely to have large effects on the risk borne by the poorest households at relatively low cost.

REFERENCES

- Amin, S., A. Rai, and G. Topa (2000). Does microcredit reach the poor and vulnerable? Evidence from northern Bangladesh. Unpublished Manuscript.
- Atkinson, A. B. (1970). On the measurement of inequality. *Journal of Economic Theory* 2, 244–263.
- Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory* 16, 252–92.
- Biewen, M. and S. P. Jenkins (2006). Variance estimation for generalized entropy and atkinson inequality indices: The complex survey data case. *Oxford Bulletin of Economics and Statistics* 68(3), 371–383.
- Borch, K. (1962). Equilibrium in a reinsurance market. *Econometrica* 30(3).
- Calvo, C. and S. Dercon (2003). Vulnerability: An axiomatic approach (the v_γ family of vulnerability measures). Unpublished manuscript.
- Chaudhuri, S. (2001). Empirical methods for assessing household vulnerability to poverty. Unpublished Manuscript.
- Chaudhuri, S., J. Jalan, and A. Suryahadi (2001). Assessing household vulnerability to poverty from cross-sectional data: A methodology and estimates from Indonesia. Unpublished Manuscript.
- Christiaensen, L. J. and R. N. Boisvert (2000). On measuring household food vulnerability: Case evidence from Northern Mali. Unpublished Manuscript.
- Deaton, A. (1985). Panel data from time series of cross-sections. *Journal of Econometrics* 30, 109–126.
- Deaton, A. and C. Paxson (1994). Intertemporal choice and inequality. *Journal of Political Economy* 102(3), 437–467.
- Dercon, S. and P. Krishnan (2000). Vulnerability, seasonality and poverty in Ethiopia. *Journal of Development Studies* 36(6), 25–53.
- Elbers, C. and J. W. Gunning (2003). Vulnerability in a stochastic dynamic model. Discussion Paper 2003–0702, Tinbergen Institute.
- Glewwe, P. and G. Hall (1998). Are some groups more vulnerable to macroeconomic shocks than others? hypothesis tests based on panel data from Peru. *Journal of Development Economics* 56(1), 181–206.
- Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics and interpersonal comparisons of utility. *Journal of Political Economy* 63.
- Hoddinott, J. and A. Quisumbing (2003). Methods for microeconomic risk and vulnerability assessments: A review with empirical examples. Unpublished Ms.

- Jalan, J. and M. Ravallion (1999). Are the poor less well insured? evidence on vulnerability to income risk in rural China. *Journal of Development Economics* 58, 61–81.
- Jalan, J. and M. Ravallion (2000). Is transient poverty different? Evidence for rural China. *Journal of Development Studies* 36(6), 82.
- Kamanou, G. and J. Morduch (2001). Measuring vulnerability to poverty. Unpublished Manuscript.
- Kocherlakota, N. and L. Pistaferri (2009). Asset pricing implications of pareto optimality with private information. *Journal of Political Economy* 117(3), 555–590.
- Ligon, E. and L. Schechter (2003). Measuring vulnerability. *Economic Journal* 113(486), C95–C102.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99, 928–956.
- Pritchett, L., A. Suryahadi, and S. Sumarto (2000). Quantifying vulnerability to poverty: A proposed measure, with application to Indonesia. SMERU Working paper.
- Ravallion, M. (1988). Expected poverty under risk-induced welfare variability. *Economic Journal* 98(393), 1171–1182.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press.
- Townsend, R. M. (1994, May). Risk and insurance in village India. *Econometrica* 62(3), 539–591.

UNIVERSITY OF CALIFORNIA, BERKELEY