International Conference

The many dimensions of poverty

Brasilia, Brazil – 29-31 August 2005
Carlton Hotel

Multidimensional Poverty in a Risky Environment Peru, 1998-2002

Conference paper

Cesar Calvo
Department of Economics/University of Oxford,
United Kingdom

Multidimensional Poverty in a Risky Environment

Peru, 1998-2002

Cesar Calvo

Department of Economics

University of Oxford

SECOND DRAFT

PLEASE DO NOT QUOTE

August 2005

1. Introduction

Human deprivation is not confined to consumption shortfalls. Recent literature is crafting both analytical and empirical frameworks that will reshape our understanding of poverty as a multidimensional reality. Low consumption, below some minimal standard, is but one of the faces of poverty, and indeed the predicament of the poor may be often due to some other form of deprivation, such as disease, illiteracy, malnutrition, and also a sense of insecurity and defencelessness as they endeavour to survive in an environment characterised by uncertainty about the future. In this paper we focus on the latter, and argue that such lack of 'peace of mind' is a relevant form of deprivation. We explore the implications of this view on the measurement of multidimensional poverty, as we study the living conditions of the poor in Peru, between 1998 and 2002.

The paper rests on the assumption that a backward-looking approach to multidimensional poverty may not suffice. While most of the literature follows the mainstream tradition and measures ex-post poverty, we advocate an ex-ante approach here. The distinction goes beyond the practicality of the point in time we choose to analyse – it will not be only about finding the best method to predict (multidimensional) poverty in an uncertain future. Uncertainty plays a much more crucial role here. Forward-looking agents surely worry about the threats of future poverty, and this welfare effect should be accounted for in some way.

Several empirical studies à la 'Voices of the Poor' (Narayan, Patel et al. 2000) witness to the desire for security as a relevant dimension of wellbeing. Psychological stress implies a welfare loss. Indeed, exposure to shocks threatening to deplete their assets, or severely restrict their consumption possibilities, or more generally affect any other dimension of wellbeing (health, nutrition, etc) is a form of predicament which the poor are all too aware of. It is in this sense that we claim that poverty is also about being deprived of 'peace of mind'. More precisely, we will use the term '*vulnerability*' to refer to the *threat of suffering any form of poverty*. This threat is tantamount for intense stress.

Since the 2000/1 World Development Report, more and more frequent use has been made of the term ‘vulnerability’, not least by practitioners. At least to some extent, the frequency of references to ‘vulnerable groups’ or ‘vulnerability assessments’ is explained by a rather loose and informal use of the term. While recent works inter alia by Ligon and Schechter (2003), Basu and Nolen (2004) and Calvo and Dercon (2005) are contributions towards a clearer and more precise definition of the notion of vulnerability, consensus is coming about only slowly. Given our interest here, we build on Calvo and Dercon, who develop a measure where vulnerability is allowed to relate both to predicted shortfalls in any particular outcome, and to the effect of uncertainty and insecurity on ex-ante well-being.

To put it differently, we intend to focus on ‘downside risks’, and on the stress it imposes on the poor. In practice, this will imply we are concerned with uncertainty only insofar as it makes it impossible to rule out critically low levels of wellbeing in the future. We do not argue here that general uncertainty is a burden and a source of stress. Our point is narrower, and so to speak, more tightly committed to the concept of poverty. We focus on the fear of future shortfalls in any one dimension of wellbeing, even if no current deficit exists. This lack of security is what we consider as a relevant form of deprivation.

Secondarily, our point is further strengthened by mounting empirical evidence of the perverse impact of risk exposure on poverty (e.g. Rosenzweig and Binswanger 1993, Dercon 1996, Elbers and Gunning 2003). If opportunities to escape poverty come at the cost of greater uncertainty, households may well shy away from them. Imagine a farmer refusing to increase land productivity by investing in fertilisers, simply because she would feel exposed to greater losses (fertiliser costs included) if rain turned out to be insufficient. The feeling of insecurity can thus reinforce other forms of deprivation, and we would miss a sizeable part of the big picture if we failed to pay attention to it.

Peru provides a suitable case study to test this approach. Between 1998 and 2002, macroeconomic performance was disappointing. GDP had grown at a 7.1% yearly rate in the five preceding years (1993-1997), only to allow internal political turmoil, along with region-wide instability, to fetter growth down to 1.6%. In terms of per capita GDP, there was no growth at all between 1998 and 2002. Recession times are typically periods of stress, when current hardship is compounded by the fear of further decline in the future – nobody knows for sure how much more severe a recession might become, nor when it will come to an end.

The paper is structured as follows. Section 2 provides a brief overview of the Peruvian economy during the period of our study. We lay down our methodology in section 3, which primarily draws on Calvo and Dercon. The dataset is described in section 4, and section 5 presents our main results. Section 6 concludes.

2. Peru, 1998-2002

Along with many other Latin American countries, Peru committed most of the nineties to bold macroeconomic reforms, with the aim to remove state interventions and liberalise all markets. Until 1997, results were impressive in terms of both overall growth (on average GDP rose 3.4% every year) and poverty reduction – the headcount fell from 54.2% in 1991 to 46.4% in 1997. Unfortunately, the party was over by 1998, when the period of our analysis starts.

A number of factors conflagrated to put an end to the economic boom. Most evidently, South America as a whole was rocked by the currency crises in Brazil first, and in Argentina later on. The Asian crisis in 1997 had changed the minds of international investors previously willing to risk their funds in developing markets, and fears to invest in Peru became all the more acute when contagion effects reached Brazil and more severely, Argentina.

Furthermore, the future performance of the economy turned more uncertain as the government progressively lost its grip of political developments in the country. Even though the precise time was unknown, increasing rumours of corruption and cronyism made clear that the end of the regime would come sooner rather than later, and also that it would be abrupt and probably violent, with unpredictable consequences on the path the country would choose in the aftermath. This political environment undoubtedly hampered the growth process.

Episodes of political instability pervade Peruvian history and have been often blamed for the failure of development projects. There is however at least one other major structural factor determining the country's performance, namely its geography. Escobal and Torero (2000) provide an excellent study of the relation between geographic conditions and economic growth in Peru. We draw heavily on their work here, as much as on Loayza and Polastri (2005).

The unpopulated Amazonia makes up for 58% of the territory, while the Andes mountains occupy another 31% and cross the country from north to south, with the rain forest and the coastal region on either side. Petty farming is the most common activity in the Andes, as it has been for many centuries. Large-scale mining projects are also important, even though far smaller in terms of the labour force they absorb.

Not surprisingly, 49% of the 27 million inhabitants live in the coast, which happens to be one of the driest areas in the planet. Living conditions are however enhanced by the Humboldt current of cold water flowing along the coast and preventing excessively hot temperatures. Most major cities are located in this region, including the capital Lima, with its seven million inhabitants.

Deep cultural differences between these regions have often been highlighted. However, it will suffice here to turn our attention to one other consequence of this geographic variety – namely, the fact that the presence of the Andes mountains impose *very high transport and communication costs*, both within the Andes region and between it and the rest of the country. Isolation is not an uncommon characteristic of life in the Andes, with obvious consequences on trade possibilities and access to public services and anti-poverty programmes.

Given our purposes, this brief and much simplified description implies that the analysis must take account of some form of geographical decomposition. Given our sample size, a distinction between urban and rural areas is probably as far as our descriptive statistics can go.

A few results from our LSMS data sets will enrich our country description, even though a detailed presentation of these sets is postponed to section 4. Table 1 uses full cross sections to ascertain the behaviour of consumption, schooling and health figures between 1998 and 2002, along with access to some basic services. Consumption is measured in terms of the local poverty line (thus a figure below 1 reveals a consumption-poor household). Schooling is proxied by the percentage of adults who completed secondary schooling, whereas our health indicator is the percentage of individuals with no reported illness in the quarter prior to the survey interview. Several caveats apply to this last indicator, as we shall discuss further on.

Both in urban and rural areas, consumption fell between 1998 and 2002, which reveals the magnitude of the crisis. Self-reported health also worsened, probably due however to wording changes in the survey questionnaires (details will be discussed shortly). Finally, schooling did improve in rural areas, but this result is surely due to changes in enrolment patterns unfolding beyond this five-year period.

Table 1
Consumption, Schooling, Health and Basic Services, 1998-2002

Household Characteristics	1998		2002	
	Urban	Rural	Urban	Rural
Consumption	2.17	1.01	1.86	0.88
Schooling	0.62	0.14	0.63	0.18
Health	0.69	0.74	0.47	0.44
Water supply	0.71	0.28	0.83	0.35
Sewage system	0.66	0.17	0.75	0.12
Electric power	0.79	0.21	0.93	0.24

Own computation.

Our main point here is the contrast between urban and rural areas. As compared with their rural counterparts, urban households spend more on consumption and achieve better schooling. In terms of health, however, fewer illnesses are reported in rural areas (except by the end of the period). Every year, differences in each of these dimensions is statistically significant (again, except for self-reported health in 2002, when the gap cannot be rejected to be zero, at 1% significance level).

The last three rows show equally striking differences with regards to access to public services. In 1998, only 28% of rural households were supplied with safe water, while only 17% had access to a hygienic sewage system. Five years later, access to water and electric power had improved (which did not happen in the case of sewage systems), but the urban-rural gap remained.

3. Methodology

Our methodology will be more easily understood by considering uni-dimensional poverty first. In that case, the vulnerability to future poverty of any i-th individual (V_i) can be measured as proposed by Calvo and Dercon (2005):

$$V_i = 1 - E\left(\tilde{x}_i^{\alpha}\right), \text{ where } 0 < \alpha < 1 \qquad (1)$$

where x stands for some particular outcome y (e.g. consumption, health), expressed as a fraction of the corresponding poverty line z. In other words, $x = \frac{y}{z}$. The ˜ symbol signifies that the outcome has been censored at the threshold line, i.e. $\tilde{x} \leq 1$.

This measure is bounded between 0 and 1, and proves to satisfy a number of desirable properties – we shall highlight only two here. Crucially, it abides by a 'focus axiom', ensuring that the individual is not made any less vulnerable to poverty by an increase in 'non-poor outcomes'. For instance, imagine only two scenarios will be possible in the future – either it will rain, or it will not. In the former case, a farmer would be rich. In the latter, she would be poor. The focus axiom implies that her vulnerability level should remain unchanged if her harvest in case of rain were doubled, say by some technological development. Both the likelihood and the severity of the threat of poverty remain unaltered by this new technology, and hence the farmer must be considered as vulnerable to poverty as she was before. The censoring of y is evidently related to this focus axiom.

Secondly, this measure is sensitive to an increase in risk (as defined by a transfer of probability weight from the centre to the tails, with no change in expected outcomes). In other words, it is faithful to the intuition of insecurity and uncertainty as a cause of stress *on its own right*, i.e. greater exposure to risk matters and impinges negatively on wellbeing, even if all in all expected outcomes remain unaltered.

It is important to remark that the measure is not sensitive to any risk in general, but to downside risk in particular – more precisely, it is concerned with the possibility of outcomes falling below some critical threshold, and with the uncertainty about how far below this threshold an outcome may be realised.

Of course, the informational requirements of this measure are very demanding. It feeds on the probability distribution of future outcome levels, and moreover, on a specific distribution for each particular individual. The expected value of y does not suffice. We need density weights for all possible outcome levels, as it becomes clear if we rewrite (1) as

$$V_i = 1 - \left(\int_{-\infty}^{1} x^{\alpha} f_i(x) dx + \int_{1}^{\infty} f_i(x) dx \right) \tag{2}$$

At this point, some studies have relied on cross-sectional information, in the hope that they will in some way describe the outcome distribution faced by any particular individual (e.g. Pritchett, Suryahadi et al. 2000, Chauduri, Jalan et al. 2002). However, this is a rather strong assumption. Data permitting, the time dimension of a panel should be the natural informational source for this estimation. For instance, one could take the following model as a starting point:

$$y_{i,t} = \beta_1 y_{i,t-1} + \beta_2 w_{i,t-1} + \beta_3 t + \mu_i + \eta_t + \varepsilon_{i,t} \tag{3}$$

where $w_{i,t-1}$ are exogenous variables, and μ_i and η_t are individual- and time-specific effects. Ideally, equation (3) would be estimated for each individual separately, and the resulting error term $\varepsilon_{i,t}$ would be construed as the main source of uncertainty – the probability distribution of future outcomes $y_{i,t+1}$ (conditional of current observables $y_{i,t}$ and $w_{i,t}$) must rely on this interpretation. Needless to say, the existence of measurement error would greatly undermine this approach.

As we shall see, the short time dimension of our data confines the scope for choice in our estimation procedures. With only five points in time, we will later have to content ourselves with *only one* variable on the right-hand side of equation (3).

We then test whether the *deviations from the resulting predictions* (which we are forced to construe as unexpected shocks) are normally distributed. If accepted, then we estimate (2) under this assumption.[1]

The extension of this framework to the multidimensional case remains to be discussed. Of course, the main intuition still holds as we generalise the analysis, namely, individuals will be wary of the threat of poverty in any one particular dimension. However, an additional source of concern arises.

In fact, the threat now arises that the individual may be hit by a form of 'double hardship', i.e. some states of the world may exist, where shortfalls occur in both wellbeing dimensions simultaneously. Put it in operative terms, the probabilistic correlation between these dimensions matters. The points discussed by Atkinson and Bourguignon (1982) apply here, even though of course their work refers to the correlation of outcomes *across individuals* of a population, as opposed our concern here for correlations *across states of the world* for each particular individual.[2] In their terms, the analysis will be driven by our decision to take these dimensions as either substitutes or complements.

We take a flexible approach, and allow consumption and health to substitute for each other, but this substitutability diminishes as either of them falls to extremely low levels. As an illustration, imagine a worker who may well be willing to sacrifice his health for some overtime payment – understandably, such trade-off cannot be pushed beyond all limits. When close to death, no person would risk any further loss in his health status.[3] Likewise, if consumption possibilities are severely constrained, it will not be sensible for this worker to refuse some overtime work. This approach is reflected by the following measure, evidently analogous to (1):

---

[1] We resort to Mathematica for these calculations.

[2] Likewise, Tsui (2002) and Bourguignon and Chakravarty (2003) refer to this issue as they discuss their 'poverty-non-decreasing rearrangement' and 'non-decreasing poverty under correlation increasing switch' properties, respectively.

[3] On the other hand, the view of consumption and health as complements is not devoid of arguments. For instance, the literature around the Human Development Index (HDI) has provided a number of arguments for health to be considered as a condition *sine qua non* of consumption-derived wellbeing. Going to the extreme, it is clear that only alive individuals can consume – in this sense, life expectancy and health conditions come first.

$$V_i = 1 - \mathrm{E}\left[\left(\delta_1 \tilde{x}_{1i}^{\ \alpha} + \delta_2 \tilde{x}_{2i}^{\ \alpha}\right)^{\frac{1}{\alpha}}\right], \text{ where } 0 < \alpha < 1 \text{ and } \delta_1 + \delta_2 = 1 \qquad (4)$$

This measure is clearly reminiscent of Bourguignon and Chakravarty (2003) and Deutsch and Silber (2005), and more precisely of equations (18) and (38), respectively. Needless to say, the expected-value operator appears in this paper for the first time. As opposed to the simple measure in (1), we lack here a full body of axioms providing the theoretical foundations of (4). However, it is not difficult to see that it retains the same spirit as (1). Increases in risk matter, and of course the focus axiom applies. Note that this implies that poverty episodes in one dimension are not allowed to be relieved by high performance (that is, above the poverty line) in the other.

The behaviour of (4) in the face of risks is determined by its second- and cross-derivatives in any particular state of the world (i.e., momentarily ignoring the expectation operator):

$$\frac{\partial^2 V_i}{\partial x_{1i}^{\ 2}} = (1-\alpha)\delta_1\delta_2\left(\delta_1 \tilde{x}_{1i}^{\ \alpha} + \delta_2 \tilde{x}_{2i}^{\ \alpha}\right)^{\frac{1}{\alpha}-2} \tilde{x}_{1i}^{\ \alpha-2} \tilde{x}_{2i}^{\ \alpha} > 0, \text{ and}$$

$$\frac{\partial^2 V_i}{\partial x_{1i}\partial x_{2i}} = (1-\alpha)\delta_1\delta_2\left(\delta_1 \tilde{x}_{1i}^{\ \alpha} + \delta_2 \tilde{x}_{2i}^{\ \alpha}\right)^{\frac{1}{\alpha}-2} \tilde{x}_{1i}^{\ \alpha-1} \tilde{x}_{2i}^{\ \alpha-1} > 0$$

where the positive values are secured by the condition $\alpha < 1$. Normalisation in turn warrants $\alpha > 0$.

As announced, the positive cross-derivative commits our measure to a view of health and consumption as substitutes. For instance, if a very healthy individual is compared to another in poor health conditions, our assumption implies that an increase in consumption will have a greater impact on the multidimensional poverty of the latter.

Finally, we turn to equation (3). If anything, the presence of more than one well-being dimensions suggests the use of seemingly-unrelated-regressions techniques, whereby potential correlations in the distributions of health and consumption could be accounted for. Yet again our sample size bars such exercise. In practice we estimate the following two simple OLS regressions:

$$c_{i,t} = \beta_3 t + \mu_i + \varepsilon_{i,t} \tag{5}$$

$$h_{i,t} = \beta_2^a gr_{i,t}^a + \beta_2^b gr_{i,t}^b + \ldots + \eta_t + \varepsilon_{i,t} \tag{6}$$

where $gr^a_{i,t}$, $gr^b_{i,t}$, ... stand for the number of members of household i in each age-sex group, at time t. All variables were expressed in logarithmic form.[4]

A crucial difference between the two equations above is that (5) is estimated for each individual separately, and thus is based on five observations (for each individual). A regression with one right-hand side variable is already an abuse, and hence we dare not add a second one. Our resulting errors are *deviations from the household-specific time trend.*

In the case of health, we resort to the cross-sections in order to estimate (6). The choice is prompted by the fact that household composition, much more than time, seems to be a sensible determinant of illness episodes at the household level. Since this implies that more than one explanatory variable will be needed (and furthermore, each of them has little variation over the five-year period), we are forced to switch to the cross-sectional dimension. Residuals are unexplained variations in illness reports. In practice, the tests and computations below use their negatives, which we construe as unexpected (positive) health shocks.

4. The data

We use data from the National Household Surveys (ENAHO, for their name in Spanish) between 1998 and 2002. These data were collected in the last quarter of each year, and make up a panel of 272 households. The sample size is admittedly limited, especially when compared to the yearly cross-sections, which peak in 2002 with

[4] If zero was a possible value for the variable at hand, we first added one unit before taking the logarithm. This was the case of health and the demographic variables.

19,673 households.[5] While the sample size can be noticeably increased by accepting households with only four observations, we prefer to preserve the time dimension of the panel as large as possible. As explained above, household-specific variance estimates will be based on the variation of household outcomes *over time*.

As most LSMS, the ENAHO surveys contain abundant information on demographic and socio-economic variables. In particular, they are especially careful and accurate with regards to social anti-poverty programmes as a source of (in-kind) income for poor families. We benefit from this feature, as consumption is one of the two dimensions which our analysis will focus on. Consumption per capita is measured following standard procedures, as well as the poverty line. As price levels vary across regions, we express all consumption figures in terms of their region-specific poverty lines. Thus we also pave the way for the estimation of vulnerability as specified by equations (1) and (4).

Our second wellbeing dimension is family health, as measured by the number of household members with no illness report in the three months preceding the survey interview. Enough studies have cast doubts on the reliability of self-reported conditions as an objective measure of true health status (e.g. Butler, Burkhauser et al. 1987). We do not intend to rebut these studies here – we are aware of the limitations of our indicator, yet we stick to it as a useful component of our illustrative exercise of the concept of vulnerability to multidimensional poverty.

Moreover, health reports are further marred by two changes in the interview questions. Firstly, the query about "disease" was replaced in 2001 by a question about "disease or any symptom of illness", which explains why our illness reports are higher in the last two years (as in the case of Table 1). Secondly, the question in 2002 asks for information about "the last four weeks" (and not "three months", as in all previous years). We attempt to make up for this discrepancy by combining these reports with those of "chronic illnesses".[6] While both issues can severely distort our results, we

[5] The other cross-section sample sizes are 7938, 4016, 4963, and 18179, in 1998, 1999, 2000, and 2001, respectively.

[6] In fact, the resulting figures are similar to the three-month reports of illness (and symptoms) in 2001.

expect our residuals-based treatment (as depicted above) to mitigate them considerably.

Our health variable has no obvious ‘poverty line’ at the household level, i.e. how many ill members are necessary for the family to qualify as health-wise poor. For illustrative purposes, we simply set this threshold at two-thirds of the household size. One shortcoming of this choice might be a form of bias acting against small families, in the sense that one ill individual is enough to label a two-member household as poor, while a four-member household would still be classified as non-poor. However, this does not need to be counterintuitive – indeed, a larger family should be able to cope more efficiently with the illness of one of its members.

Finally, before turning to our results, a crucial caveat refers to the restricted power of our normality tests, given our short time dimension. In fact, the tests happen to require five observations at the very least. We use the Shapiro-Francia test in the case of univariate normality, and the Shapiro-Wilk test in the bivariate case.[7]

5. Main results

The size of our panel set discourages decompositions into more than two subgroups. Thus, in spite of the geographical diversity of the country, our analysis will be confined to the comparison of urban and rural areas. Table 2 tests for differences in the means of consumption and health in the entire cross-sections vis-à-vis the panel set. To simplify computations, sample weights were ignored here, as they are hereafter.

While our panel set seems to produce a representative picture of health conditions, it tends to underestimate consumption figures in urban areas, and overestimate them in rural areas. Consequently, we should expect the gap between urban and rural households to be underestimated in our results henceforth.

---

[7] A STATA do-file performing the Shapiro-Wilk test is available from the author, who failed to find any add-on command to carry out this test when only five observations exist. The references can be found in Sarhan and Greenberg (1956) and Royston (1982, 1983).

Table 2
Health and Consumption in the Cross-Sections and the Panel Set

Year	Health				Consumption			
	Urban		Rural		Urban		Rural	
1998	2.84	*	1.30		3.67	*	-4.79	*
1999	1.73		1.33		2.85	*	-4.61	*
2000	-0.37		1.90		5.24	*	-1.31	
2001	-1.95		-5.36	*	4.85	*	-7.18	*
2002	-1.90		-3.80	*	3.09	*	-8.02	*

* denotes statistical significance (at 5% sign. level). Own computation.

Table 3 shows reports consumption poverty in 1998 and 2002, both with cross-sectional and panel data. On all accounts, poverty rose. In the case of urban households, the headcount increased from 25% to 36%, while the poverty gap rose from 7% to 11%. Of course, indices are much higher in rural areas. In 2002, 72% of rural households were poor. Also, just as expected, the urban-rural gap is underestimated by the panel results.

Table 3
Consumption-Poverty Indices, 1998 and 2002

Poverty Index	Cross-Section		Panel	
	Urban	Rural	Urban	Rural
1998				
FGT0	0.25	0.64	0.29	0.53
FGT1	0.07	0.26	0.09	0.18
2002				
FGT0	0.36	0.72	0.39	0.56
FGT1	0.11	0.31	0.12	0.21

Own computation.

Following the outline in section 4, we estimate OLS regressions as in (5) and (6), and focus hereafter on deviations from their predictions. Relevant results are reported in Table 4, where the variance of the model is estimated by the quadratic deviations.

Urban and rural areas turn out exhibit similar patterns with regards to these deviations. In the case of consumption, around 94% of households failed to reject the null hypothesis of normality in either area, and likewise standard deviations are very similar. We may recall here that we are presenting sample averages of *household-*

specific standard deviations. In other words, our results suggest that urban and rural households face similar degrees of uncertainty.

Results for self-reported health are also similar. If we ignore a small percentage of cases where no variation was observed (3.7% altogether, with no illness reports in any year), the remaining observations can be thought to be normally distributed in 94% and 91% of households, in urban and rural areas respectively.

Table 4
Deviations from time-trends of consumption and health

Area	Consumption Std Dev	Consumption % Normal	Health % No Var	Health Std Dev	Health % Normal	Both % J. Normal
Urban	0.25	94.7%	5.3%	0.41	94.4%	100.0%
Rural	0.24	93.4%	1.6%	0.43	90.8%	100.0%
Total	*0.25*	*94.1%*	*3.7%*	*0.42*	*92.7%*	*100.0%*

Own computation.

The fact that standard deviations are similar is unexpected. Note this might announce that provided trends do not differ drastically, no major discrepancy will arise between poverty and vulnerability figures as we compare urban and rural households, since both groups face *similar degrees of uncertainty*. On the other hand, a number of studies have stressed the existence of area-specific risks – for instance, Moser (1998) describes urban 'commoditisation hazards', such as inflation or unemployment, which indeed qualitative surveys have found to be a relevant concern of the urban poor (Narayan, Patel et al. 1999). Our results here are suggesting that either such area-specific risks are not dominant, or their magnitude is similar in urban and rural areas.

Finally, the last column of Table 4 shows that no observation (from among those where some variation in health was reported) rejected the hypothesis of bivariate normality. While the restricted power of our test must be borne in mind, these results certainly strengthen the normality assumption, which is crucial as we turn to calculate an expression similar to (2), with the obvious extension implied in (4). For simplicity, we assume normality applies in all cases, except when self-reported health is unchanged throughout, and the household is consequently assumed to face no health risks.

Consumption and health receive the same weight ($\delta_1 = \delta_2 = 0.5$), and the risk sensitivity parameter is assumed $\alpha = \frac{1}{3}$. Summary results are reported in Table 5. If we are willing to accept the illness of the one-third of the household size as the health-poverty line, then both urban and rural areas seem to be dominated by the threat of poor health outcomes, which doubles the threat of consumption hardship. Vulnerability to health deprivation is the greater concern in our data. Of course, a relaxation of the health-poverty line would easily change this conclusion – our result is admittedly arbitrary.

Table 5
Vulnerability Indices, by Area, 2002

Area	Consumption Vulnerability	Health Vulnerability	Multidimensional Vulnerability
Urban	0.08	0.15	0.26
Rural	0.09	0.17	0.28
Total	*0.08*	*0.16*	*0.27*

Own computation.

The comparison between urban and rural areas provides some further insights. Rural households are more vulnerable both to consumption poverty and to health poverty, but the difference with respect to urban indices is only marginal. The similarity of health and consumption vulnerability results is striking because we had seen that the urban households consume typically much more than their rural counterparts, whereas health differences are not as stark (Table 1). In our panel data, figures were 1.38 (urban) and 0.75 (rural) for consumption in 2002, and 0.48 for health in *both* areas.

Thus, one explanation for the similarity of urban and rural consumption vulnerability levels can be sought in the similarity of the standard deviations of consumption in those areas, i.e. in their similar degrees of uncertainty. Loosely speaking, we find here that uncertainty dominates (consumption) levels as a driving force of our vulnerability results. To put it differently, a mere backward-looking view of urban households would overestimate their advantage over rural households – in fact, they face just as much risk, and this threat bears heavily on their well-being.

Unsurprisingly, the last column in Table 5 shows that when both dimensions are considered simultaneously – as specified by (4) –, no major difference is observed.

In the setting of vulnerability to multidimensional poverty, correlations matter. For instance, imagine that health and consumption risks are negatively correlated across states of the world (that is, for on individual). In such case, both risks may counterbalance. In other words, vulnerability to health poverty and vulnerability to consumption poverty can compensate each other. Consequently, we could find households with *limited multidimensional vulnerability, and yet high uni-dimensional vulnerability levels.*

In fact, we do find that on average, health and consumption shocks (i.e. deviations from their predicted values) are negatively correlated both in urban and in rural areas, albeit only slightly.[8] The figures are -0.11 and -0.14, respectively.[9] As a rough assessment of the relevance of this effect, we calculate the quantile position of each household according to its consumption-, health- and multidimensional-vulnerability, and note that in 10% of cases, the position resulting from the latter is below (by at least 30%) the lowest position among the two uni-dimensional measures.

We come to a clearer picture of the links between vulnerability in one or another dimension by describing '*profiles of the vulnerable*'. Table 6 reports Tobit regressions predicting vulnerability levels in 2002, i.e. *the threat of uni- or multi-dimensional poverty in 2003*. The right-hand side variables are all measured in 2002, in the way of a prediction model, and include household composition variables, household- and cluster-level schooling, as well as twenty region-specific dummies (with Lima as the omitted category). Again, we appeal here to our previous discussion on the geographical diversity of the country.

Rather than looking for causal relations, we are simply finding the correlates of vulnerability. Some insights come forth promptly. For instance, the relative

[8] This negative correlation implies that potential episodes of low consumption are not necessarily compounded by poor health – if anything, the worsening of health conditions coincides with positive shocks on consumption expenditures. As an example, it may help to think of a period of little rain, causing no respiratory illnesses, but producing meagre harvests. Needless to say, the self-reported nature of our health indicator can also underlie this result.

[9] Recall that the signs of the errors in the prediction of the number of ill members were changed.

abundance of significant region-specific effects is in keeping with the geographic variety we discussed earlier. More interestingly, we note that none (except for one) of our variables has a significantly explanatory power on health vulnerability (column [3]). We will turn to this result shortly.

For the sake of comparison, column [1] predicts the household probability of being poor in 2002, in terms of consumption, i.e. a simple probit on actual consumption falling below the poverty line. All coefficients exhibit expected signs – poverty is more likely among households with more children (or girls younger than 17), or a female head, or an uneducated head, or among those located in the Andes. This is a pretty standard description of the poor in Peru. It is thus interesting to find that *it does not quite overlap with the profile of the consumption-vulnerable*.

In column [2], again all the expected signs obtain, except for the proportion of household members with complete secondary schooling, which is positively correlated with consumption vulnerability. Interestingly, some new variables surface, while others lose significance. For instance, household composition variables has no relevant role here, whereas the northern coast turns out to be a vulnerable region. This is but a natural consequence of our conceptual distinction between (backward-looking) poverty and (forward-looking) vulnerability – either can be large with no need of the other being significant, as in the case of the northern coast, or household with many children.

Of course, *the reason for the discrepancy can lie either in the change in expected consumption between 2002 and 2003, or in the degree of uncertainty*. For instance, we could guess that life conditions in the northern coast are more volatile, e.g. due to climatic reason, such as the well-known Niño phenomenon.

In fact, the correlation of consumption and consumption-vulnerability across households (as opposed to across states of the world, for a given household) is negative, but not large: -0.18. All in all, predicted consumption is however still more strongly correlated with final vulnerability than the household-specific standard deviation – the statistics in this case are -0.88 and 0.18, respectively.

Table 6
Vulnerability profiles, 2002
(Tobit regressions)

Variables	Consumption Poverty† [1]	Consumption Vulnerability [2]	Health Vulnerability [3]	Multidim Vulnerability [4]
Males [0-12]	0.666 ***	-0.011	-0.024	-0.032 *
	0.22	0.01	0.02	0.02
Males [13-17]	0.064	-0.012	0.024	0.010
	0.28	0.02	0.02	0.02
Males [over 65]	0.600	-0.025	-0.007	-0.026
	0.43	0.03	0.03	0.04
Males [other]	0.488	-0.030	-0.015	-0.041
	0.30	0.02	0.02	0.03
Females [0-12]	0.730 ***	0.001	0.012	0.019
	0.23	0.01	0.02	0.02
Females [13-17]	0.641 **	0.003	0.003	0.007
	0.27	0.02	0.02	0.02
Females [over 65]	-1.020 **	-0.040	0.005	-0.044
	0.51	0.03	0.04	0.04
Females [other]	-0.281	-0.031	0.018	-0.027
	0.36	0.02	0.03	0.03
HH: Female	0.505 *	-0.032 *	0.018	-0.013
	0.28	0.02	0.02	0.02
HH: Sec. Schooling	-0.664 **	-0.032 *	-0.020	-0.058 **
	0.28	0.02	0.02	0.02
Sec. Schooling (% in household)	-0.104	0.013 *	0.013	0.029 ***
	0.13	0.01	0.01	0.01
Sec. Schooling (% in cluster)	-1.205	0.007	-0.145 **	-0.126 *
	0.88	0.05	0.06	0.07
Northern Coast, Rural	0.626	0.096 ***	-0.005	0.098 ***
	0.42	0.03	0.03	0.04
Northern Coast, Urban	-	0.113 *	-0.094	0.045
		0.06	0.07	0.08
Central Coast, Rural	-0.356	0.030	0.000	0.036
	0.46	0.03	0.03	0.04
Central Coast, Urban	-0.321	0.053	0.032	0.088 *
	0.61	0.03	0.04	0.05
Southern Coast, Rural	0.085	0.000	-0.016	-0.011
	0.54	0.03	0.04	0.05
Southern Coast, Urban	-	0.145 **	0.096	0.238 **
		0.07	0.08	0.10
Northern Andes, Urban	1.369 ***	0.071 ***	-0.020	0.060 *
	0.42	0.02	0.03	0.03
Central Andes, Urban	0.875 **	0.056 **	0.043	0.092 ***
	0.43	0.03	0.03	0.03
Southern Andes, Rural	-	0.024	-0.067	-0.029
		0.04	0.05	0.06
Southern Andes, Urban	0.993 **	0.074 **	0.026	0.105 ***
	0.49	0.03	0.03	0.04
Rain-forest, Rural	-0.291	0.020	0.029	0.043
	0.52	0.03	0.04	0.05
Rain-forest, Urban	0.884 **	0.083 ***	-0.046	0.045
	0.42	0.03	0.03	0.04
Constant term	-0.989 *	0.085 **	0.181 ***	0.287 ***
	0.58	0.03	0.04	0.05
Sample size	260	265	265	265
Chi-squared	114.5 ***	46.0 ***	30.4	47.0 ***

†: Probit regression. *, **, and *** denote statistical significance (at 10%, 5%, and 1% sign levels). Own computation.

In the case of vulnerability to health poverty, these correlations describe a different situation. Predicted health and standard deviations correlate with household vulnerability, with almost identically strong force: -0.31 and 0.33, respectively. Indeed, this contributes to the explanation of why no variable in column [3] has predictive power (except for the percentage of educated adults in the cluster, which makes the household less vulnerable).[10] The role of uncertainty is more important here. If health risks pervade the lives of most people, regardless of how old or educated they are, where they live, and so on, then the insignificance result above is not as surprising.

As we lastly turn to multidimensional vulnerability (on column [4]), we find that by and large, results are dominated by the patterns we observed in the case of consumption. Note that this obtains in spite of the equal weights we attach to either outcome. Multidimensional poverty is more of a threat for households with little boys, or with an uneducated head, or in an uneducated cluster, or for those living in some specific regions.[11] While this description has an interest of its own, let us also pay attention to the fact that discrepancies with respect to the profile of consumption- and health vulnerability do exist.

In particular, take the case of the rural rain-forest. While households in that region are greatly exposed to consumption poverty, this does not translate (as in most other cases) into a particularly high vulnerability to *simultaneous* health and consumption deprivation. To some extent, this is due to their lower health vulnerability – yet this effect is far from statistically significant in column [3]. As we had seen, household-specific correlation of consumption and health *across states of the world* can also underlie this kind of result. And indeed, this region exhibits the greatest correlation (-0.23) – however, we hasten to admit that the our sample size renders regions with typically few observations (25, in the case at hand). Nevertheless, the fact remains that other variables did retain their significance as we move from consumption- on to multidimensional vulnerability, in spite of the effect on health vulnerability being equally negligible.

[10] The fact that demographic variables have no bearing is surprising, since indeed the vulnerability measure builds on illness report predictions which in turn, are based on household composition variables.

[11] We also find here the odd result of secondary schooling among household members increasing vulnerability, for which we have no good excuse.

Finally, the presence of little boys and the urban central coast exhibit the opposite pattern – their effects on multidimensional vulnerability are significant, whereas they are not particularly vulnerable to uni-dimensional poverty. Correlation can also contribute here to an explanation. In the urban central coast, the coefficient on column [4] is on the verge of a 5% significance level (t=1.93), and in fact, the correlation between consumption and health across states of the world is the second greatest in that region (-0.20). However, the intuition is less clear…, and indeed it opens new questions.

For instance, if spells of consumption poverty do not coincide with those of health poverty, then episodes of severe, combined deprivation would be unlikely. So far (and this is the viewpoint we have taken above), one could argue there is no reason to expect high multidimensional vulnerability. However, we must not neglect that the negative correlation also implies that a state with no consumption poverty will be likely to suffer health poverty instead. Hence we need to compare two competing effects, and this is in fact done by the functional form of our multidimensional vulnerability measure – however, the precise manner of this comparison remains unclear. We have to content ourselves here with pointing out this issue, which could be a matter of further research.

6. Concluding remarks

We have presented an empirical approach to the vulnerability of Peruvian households to multidimensional poverty. We have thus intended to pay attention to the fact that wellbeing can also hit critically low levels due to the stress caused by insecurity and powerlessness. We have argued that this is a relevant dimension where deprivation can take place, which can be argued to be as important as disease, illiteracy, or malnutrition.

Urban and rural areas are clearly differentiated in Peru. While no clear gap appears as we compare urban and rural (self-reported) health conditions, rural households are indeed much poorer in terms of consumption. Thus, as we find that both areas are equally vulnerable to future consumption poverty, we must seek for an explanation in the similarity of the standard deviations of consumption in both areas. Uncertainty can thus dominate (consumption) levels as a source of vulnerability, at least as measured by an index à la Calvo and Dercon (2005). This result strengthens our argument in favour of addressing vulnerability as a relevant wellbeing dimension (and policy target) on its own right.

On the other hand, our approach also suggested that household-specific correlations in wellbeing dimensions (in our case, health and consumption) can play a crucial role, which to date remains unexplored and requires further research. For instance, this could lead to a high degree of vulnerability to multidimensional poverty, even if uni-dimensional vulnerability were negligible. In fact, an effect along these lines seems to be in operation in the urban rain-forest of the country. All in all, correlations in our data happen however to be weak in magnitude, and unlikely to have major consequences.

Apart from region-specific effects, consumption vulnerability seems to be correlated with the characteristics of the household head, in particular whether the head is a female, and whether she has completed secondary schooling. In the case of vulnerability to health poverty, this second characteristic is the only significant determinant we were able to find.

References

Atkinson, A. and F. Bourguignon (1982), 'The Comparison of Multi-Dimensional Distributions of Economic Status', *Review of Economic Studies*, 49: 183-201.

Basu, K. and P. Nolen (2004), *Vulnerability, Unemployment and Poverty: A New Class od Measures, its Axiomatic Properties and Applications*, BREAD Working Paper 69.

Bourguigon, F. and S. Chakravarty (2003), 'The Measurement of Multidimensional Poverty', *Journal of Economic Inequality*, 1: 25-49.

Butler, J. S., R. Burkhauser et al. (1987), 'Measurement Error in Self-Reported Health Variables', *Review of Economics and Statistics*, 69(4): 644-650.

Calvo, C. and S. Dercon (2005), *Measuring Individual Vulnerability*, Oxford Economics Discussion Paper 229.

Chauduri, S., J. Jalan et al. (2002), *Assessing Household Vulnerability to Poverty from Cross-Sectional Data: A Methodology and Estimates from Indonesia*, Columbia Economics Discussion Paper 0102-52.

Dercon, S. (1996), 'Risk, Crop Choice and Savings: Evidence from Tanzania', *Economic Development and Cultural Change*, 44(3): 385–414.

Deutsch J. and J. Silber (2005), 'Measuring Multidimensional Poverty: An Empirical Comparison of Various Approaches', *Review of Income and Wealth*, 51(1): 145-174.

Elbers, C. and J. Gunning (2003), *Growth and Risk: Methodology and Micro-evidence*, Tinbergen Institute Discussion Papers 03-068/2.

Escobal, J. and M. Torero (2000), *Does Geography Explain Differences in Economic Growth in Peru?*, IADB Research Network Working Paper R-404.
Ligon, E. and L. Schechter (2003), 'Measuring Vulnerability', *The Economic Journal*, 113(486): C95-C102.
Loayza, N. and R. Polastri (2005), *Poverty and Growth in Peru*, World Bank Background Report for Peru's Poverty Assessment, unpublished.
Moser, C. (1998), 'The Asset Vulnerability Framework: Reassessing Urban Poverty Reduction Strategies', *World Development* 26(1): 1-19.
Narayan, D., R. Patel et al. (2000), *Voices of the Poor. Can Anyone Hear Us?*, New York, The World Bank and Oxford University Press.
Pritchett, L., A. Suryahadi et al. (2000), *Quantifying Vulnerability to Poverty: A Proposed Measure, Applied to Indonesia*, World Bank Social Monitoring and Early Response Unit, unpublished.
Rosenzweig, M. and H. Binswanger (1993), 'Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments', *The Economic Journal*, 103: 56–78.
Royston, J. P. (1982), 'An Extension of Shapiro-Wilk's W Test for Normality to Large Samples', *Applied Statistics*, 31(2): 115-124.
Royston, J. P. (1983), 'Some Techniques for Assessing Multivariate Normality Based on Shapiro-Wilk W', *Applied Statistics*, 32(2): 121-133.
Sarhan, A. E. and B. G. Greenberg (1956), 'Estimation and Location of Scale Parameters by Order Statistics from Singly and Doubly Censored Samples', *The Annals of Mathematical Statistics*, 27(2): 427-451.
The World Bank (2001), *World Development Report 2000/2001. Attacking Poverty*, New York, The World Bank and Oxford University Press.
Tsui, K. (2002), 'Multidimensional Poverty Indices', *Social Choice and Welfare*, 19: 69-93.

International Poverty Centre
SBS – Ed. BNDES, 10º andar
70076-900 Brasilia DF
Brazil

povertycentre@undp-povertycentre.org
www.undp.org/povertycentre
Telephone +55 61 2105 5000