



IFS

NEW METHODS FOR COMPARING
LITERACY ACROSS POPULATIONS:
INSIGHTS FROM THE MEASUREMENT OF POVERTY

Kevin Denny

New methods for comparing literacy across populations: insights from the measurement of poverty.*

Kevin Denny

Department of Economics,

University College Dublin

Ireland

&

Institute for Fiscal Studies ,

London

24th February, 2000 Version 3.02

Summary:

This paper analyses levels of low literacy across twelve countries using the International Adult Literacy Survey. We go beyond existing work that only looks at the proportions below certain critical levels of literacy. Using methods developed for the measurement of poverty we calculate measures of literacy that are sensitive to the distribution of literacy within those defined as illiterate. This reveals a different pattern of the extent of literacy problems across countries and within some populations. These measures should be useful to policy makers who need to allocate resources to alleviate low literacy and numeracy.

Keywords: literacy, numeracy, poverty, education policy.

* I thank Statistics Canada for providing the IALS data, Cathal Guiomard, Colm Harmon, David Madden, Cormac Ó Gráda and Ian Preston for very helpful comments and Sandra Redmond for research assistance. The opinions expressed in the paper are those of the author and should not be attributed to Statistics Canada or the OECD who commissioned the data. Address: Dr. Kevin Denny, Department of Economics, University College Dublin, Belfield, Dublin 4, Ireland. Phone (+353 1) 7068399. Fax (+3531) 2830068. E-mail kevin.denny@ucd.ie .

1. Introduction

The importance of literacy and numeracy in society has become increasingly prominent as changes in information technology and globalisation put a greater premium on information processing skills¹. Until very recently attempts to compare literacy and numeracy was made difficult by the absence of data that was consistent across countries. The collection of the International Adult Literacy Survey (IALS) has allowed researchers to compare literacy and numeracy skills across countries on a systematic basis for the first time. Much of the popular discussion has focussed on the scores of particular countries and the proportion of the sample falling below particular levels. The extent of literacy problems in countries that are relatively wealthy and with well developed educational systems has surprised many. For example the World Bank reported that 23% of the adult population in Ireland is functionally illiterate². This low level of literacy is worrying when we consider the perceived increase in the demand for highly skilled individuals to deal with the advancement of technology in the workplace and the significant inverse relationship between formal education and skills and social and economic exclusion.

In this paper we present some new cross-country measures of illiteracy using ideas from the literature on poverty measurement. In the latter literature, there is often a distinction between *identification* (who is defined as poor) and *aggregation* (how the poverty levels of different individuals are aggregated to arrive at an overall figure). The same distinction arises here and in this paper we are solely concerned with the latter. We take as given the interpretation of the scores at certain thresholds as corresponding to a particular level of literacy. In the next section we outline the key issues and how analogous problems have

¹ See, for example, Borjas(1996) p273 for evidence on the changing premium to education in the United States.

² Those at level 1 on the prose scale of the IALS survey were defined as functionally illiterate.

been dealt with in the context of poverty. In section 3 we describe the dataset and the results of the analysis are outlined in section 4. Section 5 concludes.

2. Aggregate measures of poverty and illiteracy

Much of the research on measuring poverty relies on counting the number of proportion of people or households below a benchmark level of income or 'poverty line'. While using the proportions of a sample or a population below a given benchmark has certain advantages as a measure of deprivation it also has obvious drawbacks. In particular it is not sensitive to the distribution of incomes within the subset which are below the line. For example, it seems reasonable that a measure of poverty should record a rise if one of the poor's income falls, other things being equal. Secondly one would expect to record a rise in measured poverty if there is a transfer from a poor person to one who is better off (i.e. either less poor or not poor at all). These criteria, formulated by Sen (1976, 1979) are known as the *Monotonicity* and *Transfer Axioms* respectively. A third axiom due to Kakwani(1980), *Transfer Sensitivity*, is less obvious. It requires that if there is a transfer from a poor household (with income y) to a less poor one (with income $y+d>y$) then the magnitude of the increase in poverty must be lower the greater y is. A given transfer between the poor has less impact the less poor they are³.

In an influential article Foster, Greer and Thorbecke (1984) - hereafter FGT - propose the following class of poverty measures. Say $y=(y_1, y_2 \dots y_n)$ is a vector of incomes ranked from lowest to highest. A poverty line z is an income level such that by definition people whose incomes are lower than z are poor and not otherwise. The number of people who are poor is $q \leq n$.

³ See Ravallion(1994) for an accessible and lucid introduction to these issues. There is a fourth axiom, that of Focus, which stipulates the changes in the income of the non-poor should have no impact on measured poverty.

An individual's poverty shortfall or deficit is defined as:

$$\begin{aligned} g_i &= z - y_i & y_i &\leq z \\ &= 0 & y_i &> z \end{aligned} \quad (1)$$

They then define the following class of poverty indices:

$$P_\alpha = \frac{1}{n} \sum_{i=1}^q \left(\frac{g_i}{z} \right)^\alpha \quad \alpha \geq 0 \quad (2)$$

The parameter α measures the sensitivity of the index to the degree of 'poorness' of those classified by the benchmark as being poor. This index nests several special cases. If $\alpha=0$ equation (2) above corresponds to a simple "headcount"; the proportion of those below the threshold. Setting $\alpha=1$ amounts to aggregating the proportionate poverty gaps while $\alpha=2$ amounts to weighting each proportionate gap by itself so a poverty gap of 20% contributes *four* times as much as one with 10%. As α increases greater weight is being placed on the poorest in a society. It can be proven that (2) satisfies Monotonicity for $\alpha>0$, the Transfer Axiom for $\alpha >1$ and the Transfer Sensitivity Axiom for $\alpha >2$.

An important advantage of this class of measure is that it is additively decomposable; the overall score is a weighted sum of the scores for a set of mutually exclusive and exhaustive sub-groups where the weights are the corresponding proportions of the populations. Hence, one can measure the contribution of the poverty of a particular sub-group to overall poverty. This decomposability does not apply hold for measures that involve rankings, such as those that include Gini coefficients.

In this paper we take this class of index and apply it to the measurement of literacy by replacing the poverty line with an "illiteracy line". The merit of doing this depends on the extent to which the measurement of literacy is comparable to income of individuals or

households. Clearly there are differences. The hypothetical transfers often discussed in the context of poverty and inequality have no direct analogue in the context of measuring literacy. Income is "rivalrous", that is an increase in an individual's income of £1 necessarily implies £1 less to someone else and this is not the case with an individual's level of literacy or education. Moreover income is cardinal in a way that test scores are not⁴. Nonetheless the argument for an index to exhibit Monotonicity remains strong: it is reasonable that lower attainments on the tests *within* the group who are below the threshold should correspond to a higher measure of illiteracy.

The argument for the Transfer Axiom, the convexity of the underlying measure, may be less compelling. It amounts to putting a weight on a given shortfall or gap with weights increasing in the magnitude of the gap. It is not obvious that a given reduction in a test score (say 1 point) is more disadvantageous for an individual whose score is 100 than someone whose score is 150. However the same point could be said of income and the poor. If one is sufficiently poor, further declines in one's income may have a declining or even zero impact on one's standard of living.

One can adopt a "Veil of Ignorance" argument here: given that one can't know for sure it seems reasonable to adopt a more egalitarian approach and to attach a greater weight to those who are worse off. It is more difficult to give an intuitive argument for satisfying the Transfer Sensitivity Axiom and the estimates here do not in general satisfy it since we only compute the index for α equal to 0,1 and 2.

⁴ There is also a built-in upper bound to the distribution of scores, 500 in this case, though there is not for the distribution of income. This is hardly a major shortcoming since very few attain this score and we are interested in low scores only.

How does one choose α ? This requires judgement on the part of the investigator. Clearly it is arbitrary to some extent and may depend on the purpose of making the comparison. In the framework we have adopted it should be clear that the conventional "headcount" measure is no less arbitrary and that there are sensible *a priori* grounds for picking $\alpha > 0$. In applications of this approach it makes sense to see how sensitive are rankings to different values of α .

There is one method which may allow us to empirically determine the "appropriate" α . Say one is interested in literacy because of its impact on known variables such as earnings or employment status. One could use multivariate methods to fit a model which explains, say, the level of an individual's income in terms of a number of variables including their education socio-economic background and level of literacy. As a measure of literacy one could use the individual's literacy gap (g_i/z) raised to an arbitrary power and see what value fits the data best. In other words, in the construction of the index we use as a value of α , a measure of its impact on the variable of interest. In other work we have measured the impact of literacy scores on earnings in a subset of countries where it is shown that quantitative literacy (i.e. numeracy) is what matters (Denny *et al* (2000)) but the possible non-linearity has not been extensively explored.

In general however one is concerned about low literacy because it affects many aspects of a person's life some of which are not easily measurable if at all so this approach to determining the appropriate α may be limited in its usefulness.

3. The IALS Data

The data used in this paper come from the International Adult Literacy Survey (IALS) which was undertaken by twelve governments in association with the European Union, the OECD

and UNESCO between 1994 and 1996⁵. The purpose of the survey was to assess the literacy level of the adult population and to provide a common measure that would allow comparison of literacy proficiency across countries rather than a mere count of the number of 'illiterate' people in the population. For most countries, the survey consists of a sample of 2000 to 3000 from the adult civilian population aged between 16 and 65⁶. The language of interview is each country's respective national language. Sample design was the responsibility of each country.

The IALS is structured around three stages. Firstly, each individual was required to complete a background questionnaire, which provided information on age, sex, education, labour market experiences and literacy related activities. An individual was deemed to be an IALS respondent if they partly or fully completed the background questionnaire. Stage 2 involved the completion of six simple assignments; if the respondent answered incorrectly on more than two of these tasks the interview was terminated. This was in order to avoid assigning further tasks to those individuals whose literacy level is already known to be very low. Lastly, a main booklet of tasks was given to each respondent, which resulted in a score that measured their literacy level.

All assignments required the respondent to use materials from everyday life. For example, instructions from medicine bottles, the completion of order forms and reading a newspaper are listed amongst the tasks that were required in order to complete the test questionnaire.

⁵ The countries involved were Australia, Canada (French and English speaking), Belgium, Germany, Ireland, Netherlands, New Zealand, Sweden, Switzerland (French and German speaking), United Kingdom, United States and Poland. Unfortunately the data for Australia are not available to researchers outside of Australia. An additional seven countries are in the third round of IALS to be published soon.

⁶ All IALS countries were instructed to exclude residents in prisons, hospitals and psychiatric institutions.

The literacy level is measured on three scales: prose, document and quantitative⁷. Prose literacy is the knowledge required to understand and use information from texts such as newspapers, pamphlets and magazines. Document literacy is the knowledge and skill needed to use information from specific formats, for example, from maps, timetables and payroll forms. Quantitative literacy is defined as the ability to use mathematical operations, such as in calculating a tip or compound interest. In order to provide an actual measure of literacy, each individual was given a score for each task, which varied according to the difficulty of the assignment.

The scores for each scale range from 0 to 500 and are subsequently subdivided into five levels. Level 1 has a score range from 0-225 and would indicate very low literacy where, for example, instructions for a medicine prescription would not be understood. The interval 226 to 275 defines level 2, where individuals are limited to handling material that is not too complex and clearly defined. Level 3 ranges from 276 to 325 and is considered the minimum desirable threshold for most countries while level 4 (326-375) and level 5 (376-500) show increasingly higher skills which integrate several sources of information or solve complex problems⁸.

It is clear from the study design that the definition of literacy was not intended to be focused on literacy and numeracy as conventionally thought of - rather it was aimed at

⁷ The survey makes use of "plausible value" sampling methodology, which provides five measures of each of the three variables (prose, document and quantitative literacy) based on the fact that individuals will answer different parts of a given question. Thus, for each of the three literacy scales, each of the five values is equally plausible.

⁸ In constructing the scores each country was required to re-evaluate 20% of the tests to guarantee precision of results. In addition one statistical agency was required to re-evaluate 10% of another country's scores. The designers of the IALS were also very conscious of non-response bias. Interviewers were advised to return to households that did not give a response as many times as possible and the sample was carefully weighted to known population variables.

encompassing a broad range of skills used in the context of working, schooling and home duties which are much more cognitive in nature than the term ‘literacy’ at first suggests (OECD 1997). For example, respondents are asked to balance a chequebook, complete an order form, and maintain payroll data, tasks which require much more knowledge than simply the ability to read and write. Kirsch and Mosenthal (1992) applied regression analysis to determine the difficulty of each task involved in the US' National Adult Literacy Survey (NALS) in order to uncover what skills were being used in the completion of the tasks. Their findings suggest that the skills required to complete each task stretched beyond conventional ‘reading and writing’ abilities. This suggests that the definition of literacy in the NALS required more skill and ability than the traditional definition. Overall, they conclude that an ordered set of skills and knowledge is required to complete the different tasks. Since the IALS survey team adopted the NALS methodology and scaling procedures, this suggests that the various definitions of literacy in the IALS also require skills that are more cognitive in nature than the conventional understanding.

4. Results

Previous analyses of the IALS (OECD1995, 1997) have, in the main, used as comparisons measures of central tendency such as the mean or median or the proportions lying below certain values which are deemed to be critical. The threshold below which we define individuals as "illiterate", is the value that defines the upper bound of level 2 literacy: 275 .This corresponds to the value used by the World Bank (1999).

To calculate standard errors, one can use analytical formula for the asymptotic distribution of these indices. Kakwani (1990) shows that the sample estimates of the variance of (2) are given by:

$$\text{var}(\sqrt{n}.\hat{P}_\alpha) = \hat{P}_{2\alpha} - \hat{P}_\alpha^2 \quad (3)$$

In the related context of inequality measurement Mills and Zandvakli (1997) caution against the use of asymptotic approaches in small samples and recommend a distribution-free method such as bootstrapping. We provide bootstrap standard errors for each measure based on 80 replications. Experiments with more replications suggest that this number is easily sufficient⁹. Note that both bootstrapping and the method of Kakwani treat the actual critical value used as non-stochastic whereas in practice it is often estimated from the same sample as the statistics of interest. So, for example, one typically estimates the proportion of the sample below, say, the median of that sample. Hence random sampling generates two sources of uncertainty in the calculation of the threshold and in inferring the proportion below that threshold¹⁰. The measure used here is absolute rather than relative so this does not arise.

Table 1 presents three measures of prose literacy based on the FGT methodology, corresponding to values of α equal to 0, 1 and 2 (the corresponding estimates for document literacy and quantitative literacy provided in tables A1 and A2 in the Appendix. The first set of results ($\alpha=0$) corresponds to the "headcounts" reported in the existing literature and widely discussed elsewhere¹¹. We use the sample weights provided with the data. The rankings are reported in the adjacent column and tests for rank correlation are presented at the end. The new indices that we propose are those for α equal to 1 and 2. All three indices generate very similar rankings. Most of the countries displaying high levels of illiteracy by the headcount measure (such as Poland and Northern Ireland) continue to do so when we take account of the

⁹ This method of calculating standard errors differs from the "jackknife" procedure recommended by the data providers for calculating descriptive statistics, see Statistics Canada (1998), chapter 8. This is to deal with the "plausible value" sampling methodology. It is not clear how this method can be extended to our approach. As we just use one of the plausible values (see footnote 7) this should not be a problem and the bootstrap should provide relatively robust estimates in general, see Efron & Tibshirani (1993). Using any of the other "plausible values", such as "prose2", "prose5" and so on, give results that are not noticeably different. Details and a Stata file to replicate the results are available on request from the author.

¹⁰ Preston (1995) derives the distributional formulae allowing for both sources of sampling error which may or may not offset each other.

severity of the literacy shortfalls. Similarly most of the countries who are doing well continue to do so (such as Sweden and the Netherlands). There are important exceptions nonetheless. While Germany is middle ranking by the first index, increasing α improves Germany's position in the "league table" so that taking account of the severity of the literacy shortfalls implies that Germany is, by contrast, the country least troubled with high illiteracy. Interestingly, the same pattern occurs with the German-speaking Swiss sample. The populations that emerge worst in our new analysis are the US and the English speaking Canadians. For example the simple headcount ($\alpha=0$) gives the US and Belgium an almost identical incidence of low literacy, the US being marginally worse to an extent that is not statistically significant. But the severity of the literacy shortfall ($\alpha=1$) is lower in Belgium than the US by about 15% (0.101 - 0.086).

Taking account of the size of the literacy shortfalls does not always increase the difference implied by the headcount. For example, while there are substantially fewer English speaking Canadians in the low literacy category than there are German-speaking Swiss, the gap between the two is more or less eliminated with $\alpha=1$ or higher.

The results for quantitative and document literacy (in Tables A1 and A2 in the Appendix) show much less variation in rankings as one increases α . In the former, for example, no country's rank changes by more than two places. This suggests that the distribution of scores within those who are below the benchmark are not very different. However it is interesting that some of the patterns that emerged in looking at prose literacy also appear, albeit to a lesser degree, in the other literacy domains. As happened with the prose literacy results in table 1, the English-speaking Canadians have a higher ranking in Table A1 and A2 (that is

¹¹ To facilitate comparisons we have followed the convention of using the "first" of the plausible values for each test ("prose1" "doc1" "quant1"). Our results for $\alpha=0$ replicate those in Murray *et al* (1998) Table 12.2.

higher levels of illiteracy) as we take account of the size of the gap. This effect is more evident for document literacy than quantitative. This finding is also true for the United States. Equally the result in Table 1 that allowing for the severity of the literacy puts Germany in a better light, recurs when one examines quantitative and document literacy. As a broad generalisation there seems to be relatively little change between the second and third indices compared to the first and second so in applications it may be sufficient to use the more intuitive value of α equal to 1. There is, of course, no guarantee that this holds generally.

Accounting for the difference levels of literacy across countries and why different indices give a different picture is beyond the scope of this paper. The work by the OECD(1995,1997) and the authors (Denny *et al* (1999)) show that, as one expects, higher educational attainment is associated with higher levels of literacy 'though to an extent that varies by country. One may think of the different types of literacy as "joint products" of the educational system. This raises the question, why does a given distribution of educational attainment, generate a different distribution of scores across tests? For example the aggregate gap (that is when α equals 1) for French speaking Canadians is 0.149 for document literacy but 0.126 for quantitative literacy. So there is a proportionately greater discrepancy between the two (and indeed a different ordering) than is suggested by simply comparing the corresponding headcounts (58% and 60% respectively). One explanation may be that the impact of a given level of schooling is greater on document literacy than quantitative literacy. This requires further research. The point is that this question would probably not have been posed if one relied on indices that are not distributionally sensitive.

Table 1: Prose Results

	n	a=0	rank	a=1	rank	a=2	rank
Belgium	2246	0.466 <i>0.012</i>	6	0.086 <i>0.004</i>	5	0.028 <i>0.00</i>	5
Canada (English)	3951	0.452 <i>0.021</i>	3	0.097 <i>0.005</i>	7	0.037 <i>0.004</i>	9
Canada (French)	1709	0.534 <i>0.031</i>	11	0.129 <i>0.013</i>	13	0.052 <i>0.01</i>	13
Germany	2062	0.486 <i>0.013</i>	7	0.068 <i>0.003</i>	3	0.015 <i>0.00</i>	1
Great Britain	3811	0.520 <i>0.010</i>	9	0.104 <i>0.004</i>	11	0.038 <i>0.002</i>	10
Ireland	2423	0.524 <i>0.011</i>	10	0.102 <i>0.003</i>	10	0.033 <i>0.002</i>	7
Netherlands	3090	0.440 <i>0.011</i>	2	0.064 <i>0.002</i>	2	0.016 <i>0.000</i>	2
New Zealand	4223	0.458 <i>0.010</i>	4	0.085 <i>0.003</i>	4	0.027 <i>0.002</i>	4
Northern Ireland	2907	0.542 <i>0.011</i>	12	0.112 <i>0.003</i>	12	0.040 <i>0.002</i>	12
Poland	3000	0.772 <i>0.008</i>	14	0.191 <i>0.004</i>	14	0.074 <i>0.002</i>	14
Sweden	3038	0.367 <i>0.008</i>	1	0.063 <i>0.002</i>	1	0.020 <i>0.002</i>	3
Switzerland (French)	1440	0.515 <i>0.015</i>	8	0.095 <i>0.005</i>	6	0.034 <i>0.003</i>	8
Switzerland (German)	1398	0.550 <i>0.014</i>	13	0.099 <i>0.005</i>	8	0.033 <i>0.003</i>	6
USA	3045	0.466 <i>0.011</i>	5	0.101 <i>0.004</i>	9	0.039 <i>0.002</i>	11
Rank correlation tests (Spearman)							
$\alpha=0$ vs. $\alpha=1$		0.811					
P value		<i>0.004</i>					
$\alpha=0$ vs. $\alpha=1$		0.613					
P value		<i>0.020</i>					
$\alpha=1$ vs. $\alpha=2$		0.925					
P value		<i>0.000</i>					

NOTE: Standard errors below estimates.

While cross-country comparisons may make good headlines they are of limited interest from the point of view of public policy. However measures of illiteracy across groups, but within a country, are very useful since they provide valuable information about where resources might be best employed to alleviate the problem. This is where the decomposability property of the FGT indices is helpful. As an example, we decompose the results for the

United States across the four combinations of high/low education and men/women¹². The sizes of the sub-samples give the approximate weight of each sub-groups contribution to the overall measure because of the use of sample weights to correct for over/under sampling.

Table 2: Decomposing the US (prose literacy) results by education & gender

	n	a=0	a=1	a=2
Women low education	442	0.804	0.267	0.126
<i>s.e.</i>		<i>0.030</i>	<i>0.015</i>	<i>0.012</i>
Women high education	1137	0.352	0.057	0.016
<i>s.e.</i>		<i>0.017</i>	<i>0.004</i>	<i>0.002</i>
Men low education	413	0.832	0.266	0.123
<i>s.e.</i>		<i>0.026</i>	<i>0.015</i>	<i>0.012</i>
Men high education	997	0.417	0.068	0.020
<i>s.e.</i>		<i>0.017</i>	<i>0.004</i>	<i>0.002</i>

Note: Bootstrap standard errors based on 125 replications.

From Table 2 it is clear that the low incidence of literacy is strongly associated with low educational and this conclusion is not fundamentally dependent on which index is chosen. However while the simple incidence of low literacy ($\alpha=1$) is about 3% higher amongst low-educated males compared to females, taking into account the severity effectively eliminates this gap. That is the higher frequency of illiteracy amongst males is counteracted by a greater severity of illiteracy within the female respondents. Low literacy is moreover far from being absent amongst the relatively well educated. However the relative incidence across groups is not invariant to the choice of index. The first column of results would lead one to conclude that low literacy is twice as great a problem amongst low educated men compared to highly educated men. The second set of results would suggest that it is almost four times greater (.266/.068).

¹² "High" education in this context is equivalent to an ISCED level of at least 3, equivalent to higher secondary schooling or greater. That is they have at least graduated from High School.

As a second example we decompose the results for Great Britain by gender and age cohort with the younger cohort being 40 years old or less.

Table 3: Decomposing the British (prose literacy) results by age and gender

	n	a=0	a=1	a=2
Women : young	1123	0.469	0.080	0.027
<i>s.e.</i>		<i>0.016</i>	<i>0.006</i>	<i>0.003</i>
Women : old	958	0.613	0.133	0.049
<i>s.e.</i>		<i>0.020</i>	<i>0.007</i>	<i>0.006</i>
Men : young	863	0.461	0.090	0.033
<i>s.e.</i>		<i>0.023</i>	<i>0.007</i>	<i>0.005</i>
Men : old	867	0.564	0.121	0.046
<i>s.e.</i>		<i>0.020</i>	<i>0.008</i>	<i>0.006</i>

Note: Bootstrap standard errors based on 125 replications.

The choice of index does not change the broad picture with regard to age: the younger cohort have lower levels of illiteracy. This is probably a reflection of their higher educational attainment. In other work we have measured the impact on literacy of several factors including educational attainment of individuals and find that education is quite significant but to an extent which varies across countries (Denny *et al*(1999)). With respect to gender the picture is less clear: taking account of the severity of the illiteracy ($\alpha=1$ or 2) reverses the conclusion that young men have a (slightly) lower degree of illiteracy than young women.

5. Conclusions

This paper presents a new method of characterising the extent of low literacy in a particular population or subset thereof. Using a class of measures well known from the economic analysis of poverty we present two indicators of illiteracy in addition to a standard "headcount" of the number of people below a given threshold. We show that allowing for the severity of low literacy has some impact on rankings of different countries. For example, Germany improves its position relative to other countries while English-speaking Canadians and the United States fare much worse. Dis-aggregating our measures for the US by education level and gender, shows that when allowing for the severity of the problem makes the gap between men and women with low education disappear. In all these estimates we have accepted the threshold value for literacy chosen by the designers of the data as having certain intrinsic significance and widely used in the existing research on these data. In this regard it is worth quoting Deaton's observation about the arbitrariness of poverty lines¹³:

"...I see few advantages in trying to set a sharp line, below which people count and above which they do not. Poverty lines and poverty counts make good headlines, and are an inevitable part of the policy debate, but they should not be used in policy evaluation. Perhaps the best poverty line is an infinite one; everyone is poor, but some a good deal more so than others, and the poorer they are the greater the weight they should get in measuring welfare and in policy evaluation."

¹³ Deaton (1997) p 146.

Much the same point could be made about literacy and numeracy. A contrary view would be that there are distinct levels or thresholds in measured literacy which correspond to some concept of "fluency". If so, this is not clear from the data. One could indeed make a similar point about poverty lines: that income below a certain level may generate a distinct state of social exclusion.

Nonetheless we believe that the methods shown here make a useful contribution towards developing methods for comparing the functional literacy of nations and other groups. Due to innovative data sources such as the IALS there is now a basis for the systematic analysis of literacy problems in many countries. As a result, governments are better placed to tackle the problem of low literacy. An informed public policy for allocating resources to reduce the problem needs a method which takes into account not just the simple incidence of low literacy but the severity of it. The method outlined here is a step towards this end.

References

- Borjas George (1996) *Labor Economics* McGraw-Hill
- Deaton Angus (1997) *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy* Baltimore: Johns Hopkins University Press.
- Denny Kevin, Colm Harmon and Sandra Redmond (1999) Literacy and Education in Ireland *Economic and Social Review*, 30, 215-226.
- Denny Kevin, Colm Harmon and Sandra Redmond (2000) *Cognitive Skills, Educational Attainment and Earnings - Evidence from the International Adult Literacy Survey* University College Dublin, Economics Department Working Paper 2000/4.
- Efron Bradley, Robert J. Tibshirani (1993) *An Introduction to the Bootstrap*. London: Chapman & Hall.
- Foster James, Joel Greer, Erik Thorbecke (1984) A class of decomposable poverty measures. *Econometrica*, 52, 761-767.
- Kakwani Nanak (1980) On a class of poverty measures. *Econometrica*, 48, 437-446.
- Kakwani Nanak (1990) *Testing for Significance of Poverty measures; with application to Côte d'Ivoire*. Living Standards Measurement Study Working Paper no. 62, World Bank, Washington DC.
- Kirsch I.S. and P.B. Mosenthal (1990) Exploring Document Literacy: variables underlying the Performance of Young Adults. *Reading Research Quarterly*, 25, pp 5-30
- Mills, Jeffrey, Sourushe Zandvakli (1997) Statistical inference via bootstrapping for measures of inequality. *Journal of Applied Econometrics*, 12, 133-150.
- Murray, T. Scott, Irwin Kirsch & Lynn B. Jenkins (1998) *Adult Literacy in OECD countries: Technical Report on the First International Adult Literacy Survey*, National Center for Educational Statistics Washington DC.
- O.E.C.D. (1995) *Literacy, Economy and Society: Results of the First International Adult Literacy Survey* Paris: OECD .
- O.E.C.D. (1997) *Literacy Skills for the Knowledge Society*. Paris: O.E.C.D.
- Preston Ian (1995) *Sampling distributions of relative poverty statistics*, Applied Statistics (Journal of the Royal Statistical Society: Series C) 44, 91-99.
- Ravallion Martin (1994) *Poverty Comparisons* Chur: Harwood Academic Publishers (Fundamental of Pure and Applied Economics 56)
- Sen Amartya (1976) Poverty: an ordinal approach. *Econometrica* 44, 219-231
- Sen Amartya (1979) Issues in the Measurement of Poverty. *Scandinavian Journal of Economics*, 81, 285-307.
- Statistics Canada (1998) *International Adult Literacy Survey: Microdata Package: Guide*, Ottawa: Statistics Canada.
- World Bank (1999). *World Development Report*. Washington DC: World Bank.

Appendix Table A1 Quantitative Literacy

	n	a=0	rank	a=1	rank	a=2	rank
Belgium	2246	0.395 <i>0.013</i>	5	0.078 <i>0.004</i>	6	0.028 <i>0.002</i>	4
Canada (English)	3951	0.439 <i>0.021</i>	7	0.093 <i>0.007</i>	7	0.035 <i>0.003</i>	8
Canada (French)	1709	0.596 <i>0.038</i>	13	0.126 <i>0.013</i>	13	0.045 <i>0.007</i>	13
Germany	2062	0.333 <i>0.117</i>	2	0.038 <i>0.002</i>	1	0.008 <i>0.001</i>	1
Great Britain	3811	0.511 <i>0.010</i>	11	0.110 <i>0.004</i>	11	0.042 <i>0.003</i>	12
Ireland	2423	0.532 <i>0.105</i>	12	0.113 <i>0.004</i>	12	0.041 <i>0.002</i>	10
Netherlands	3090	0.385 <i>0.009</i>	4	0.060 <i>0.002</i>	3	0.017 <i>0.001</i>	2
New Zealand	4223	0.494 <i>0.008</i>	9	0.095 <i>0.003</i>	8	0.031 <i>0.001</i>	7
Northern Ireland	2907	0.498 <i>0.011</i>	10	0.110 <i>0.004</i>	10	0.041 <i>0.002</i>	11
Poland	3000	0.693 <i>0.010</i>	14	0.186 <i>0.004</i>	14	0.081 <i>0.002</i>	14
Sweden	3038	0.317 <i>0.009</i>	1	0.054 <i>0.003</i>	2	0.018 <i>0.001</i>	3
Switzerland (French)	1440	0.375 <i>0.014</i>	3	0.073 <i>0.005</i>	4	0.029 <i>0.003</i>	6
Switzerland (German)	1398	0.403 <i>0.015</i>	6	0.077 <i>0.003</i>	5	0.028 <i>0.002</i>	5
USA	3045	0.463 <i>0.013</i>	8	0.101 <i>0.003</i>	9	0.038 <i>0.002</i>	9

NOTE: Standard errors below estimates.

Rank correlation tests (Spearman)

$\alpha=0$ vs. $\alpha=1$ $\rho=0.9824$ p value=0.0000

$\alpha=0$ vs. $\alpha=2$ $\rho=0.9297$ p value=0.0000

$\alpha=1$ vs. $\alpha=2$ $\rho=0.9604$ p value=0.0000

Appendix Table A2 Document Literacy

	n	a=0		a =1		a =2	
Belgium	2246	0.394 <i>0.013</i>	2	0.073 <i>0.004</i>	4	0.025 <i>0.002</i>	4
Canada (English)	3951	0.444 <i>0.019</i>	5	0.107 <i>0.007</i>	8	0.047 <i>0.005</i>	12
Canada (French)	1709	0.578 <i>0.035</i>	13	0.149 <i>0.015</i>	13	0.064 <i>0.008</i>	13
Germany	2062	0.417 <i>0.013</i>	4	0.051 <i>0.002</i>	1	0.011 <i>0.001</i>	1
Great Britain	3811	0.503 <i>0.009</i>	9	0.110 <i>0.004</i>	9	0.043 <i>0.003</i>	9
Ireland	2423	0.570 <i>0.001</i>	12	0.116 <i>0.003</i>	11	0.040 <i>0.002</i>	7
Netherlands	3090	0.397 <i>0.008</i>	3	0.062 <i>0.002</i>	3	0.017 <i>0.001</i>	2
New Zealand	4223	0.506 <i>0.009</i>	10	0.099 <i>0.003</i>	6	0.033 <i>0.002</i>	6
Northern Ireland	2907	0.539 <i>0.011</i>	11	0.118 <i>0.003</i>	12	0.045 <i>0.002</i>	10
Poland	3000	0.761 <i>0.009</i>	14	0.22 <i>0.004</i>	14	0.101 <i>0.003</i>	14
Sweden	3038	0.328 <i>0.010</i>	1	0.056 <i>0.002</i>	2	0.018 <i>0.001</i>	3
Switzerland (French)	1440	0.452 <i>0.015</i>	6	0.083 <i>0.005</i>	5	0.028 <i>0.003</i>	5
Switzerland (German)	1398	0.472 <i>0.019</i>	7	0.101 <i>0.005</i>	7	0.042 <i>0.003</i>	8
USA	3045	0.497 <i>0.011</i>	8	0.115 <i>0.004</i>	10	0.046 <i>0.002</i>	11

NOTE: Standard errors below estimates.

Rank correlation tests (Spearman)

$\alpha=0$ vs. $\alpha=1$ $\rho=0.8989$ p value=0.0000

$\alpha=0$ vs. $\alpha=2$ $\rho=0.7363$ p value=0.0027

$\alpha=1$ vs. $\alpha=2$ $\rho=0.9121$ p value=0.0000