WILEY

ON MEASURING LITERACY*

Kaushik Basu and James E. Foster

This paper presents a new approach to evaluating the aggregate literacy level in a country or region. The resulting 'effective literacy' measure takes into account the intrahousehold externality arising from the presence of a literate member. A simple and natural axiomatic characterisation is provided along with an empirical illustration using Indian data.

1. Proximate Illiteracy

A country's overall level of literacy is usually measured by taking the number of adults who are literate as a percentage of the total number of adults – the so-called *literacy rate.* Following Sen (1985) there has been increased use of the literacy rate and other social indicators to evaluate the overall standard of living in a country.[1] The present paper is concerned with a particular deficiency of the literacy rate as an indicator of the aggregate benefit from this important functioning. It draws attention to this inadequacy, develops a new measure, gives it a full axiomatic characterisation and offers an illustration of its use.

Suppose we learn that a certain country has a literacy rate of 40%. To be sure, this number is compatible with very different scenarios of the distribution of literate persons across households. In one scenario, the literate population could be highly concentrated and separate from the illiterate population so that, say, every household is either fully literate or fully illiterate; in another, the literate individuals might be 'evenly distributed' with, say, every household containing at least one literate member. In this paper we argue that, in a well-defined sense, a more even distribution of literacy across households leads to greater *effective* literacy. However, unlike in the similar task of measuring poverty (Sen, 1976; Clark *et al.*, 1981; Foster *et al.*, 1984), the concern for distribution need not reflect a concern about distribution *per se.* It is, of course, possible to argue that, as with all good things in life, a more even distribution of literacy is innately, ethically appealing. In the present paper, though, we contend that even if we ignore the inherent appeal of an equitable distribution, there are important *instrumental* reasons for being concerned about a 'better' distribution of literacy.

[1] See also Sen (1987) and UNDP (1990).

[1733]

To see this, consider the following examples that involve the use of literacy skills:[2]

A. A low-skilled job is available which requires the ability to read and write.

B. Agricultural extension workers come with information on how to plant and take care of high-yielding varieties. They leave behind brochures explaining these matters.

C. A medical facility is set up in a neighbouring village. The staff distributes pamphlets on methods of preventing disease and infection, as well as information on the various services offered by the facility.

Observe that while each of these opportunities – the low-skilled job, the agricultural information, and the health facility – is intimately connected to literacy, the connection is not of the same kind. In case A, the person has to be literate himself or herself to take advantage of the opportunity. In cases B and C this may not be necessary. All one really needs is access to a literate person who is willing to provide the requisite literacy services.

It is our contention that having a literate member *in the household* can make a substantial difference for each illiterate member in accessing information and accomplishing tasks that require literacy skills. In other words, literate household members generate a *positive externality* or a kind of *public good* for illiterate members. A wide variety of studies support this intuitive hypothesis. For example, Green *et al.* (1985) identified family literacy as a key variable in the propensity of Guatemalan peasant farmers to adopt modern farm practices and concluded that in this respect 'an illiterate farmer with a literate family is not at a disadvantage to a farmer who is literate himself.' More recently, Foster and Rosenzweig (1996a) found that the productivity of a household farm is linked to the education level of the most educated member of the household, and that these productivity gains are greatest at the lowest education levels.[3] It is therefore important to distinguish between two types of illiterate persons when assessing the distribution of literacy: a *proximate illiterate*, an illiterate person who lives in a household with at least one literate member, and hence has access to the public good; and an *isolated illiterate,* an illiterate person whose household has no literate members.

The key difference between the two motivating scenarios – each having the same 40% literacy rate, but a different pattern of literacy – should now be evident. In both cases, three out of every five persons are illiterate. But in the evenly-distributed scenario, all three are proximate illiterates who thus share in some of the benefits of literacy; whereas in the highly concentrated scenario all are isolated illiterates who do not. Our aim in this paper is to develop a measure of literacy which reflects this important distinction by taking account of the incidence of proximate illiteracy.

---

[2] For other examples of the use (and impact) of literacy, see Bhola (1984, 1994), Hayes and Valentine (1989), Margo (1990) and Wagner and Puchner (1992).

[3] See also the many references given in Green *et al.* (1985) and the case studies presented by Fingeret (1983). Dreze and Saran (1995) discuss a related type of externality in which the benefits of education accrue across individuals by 'one person taking decisions on behalf of another person'.

Before we proceed to measurement, some caveats are worth mentioning. There are surely cases where close proximity to a literate may be a handicap rather than a help. An illiterate woman married to a man scheming to extort more dowry may well be better off if the husband were illiterate and thereby less proficient in this effort. While we recognise that relative literacy levels can alter power relationships in a household, which in turn may affect the capabilities of household members to function, these issues are beyond the scope of the current paper and we do not address them here. Secondly, in certain cases there may be differential externalities to illiterate household members from, say, a father being literate, a mother being literate, or both being literate. Our formal presentation initially abstracts from this possibility. However, in Section 5, we extend our measure to account for differential effects that are gender-based, an alteration that may be especially important when significant inequalities between the sexes are evident. Thirdly, the household is not the only social unit that matters in questions of literacy. If no one is literate in the household, it could still help if one person in the village were literate or even one person in the *zilla* were literate. These more distant externalities are clearly missed by our household-based measure. However, our general approach easily accommodates larger social units, and this extension is discussed in the concluding section.

2. Effective Literacy

We propose in this paper a simple measure of a society's 'effective literacy' which takes account of the externalities mentioned above. This section provides the framework within which the measure can be defined.

Consider a country in which there are n adults and m households. Each household $h = 1, 2, \ldots, m$ has a *household literacy profile*, x^h, indicating each adult household member's level of literacy, where $x_j^h = 1$ is interpreted to mean that the jth member of household h is literate and $x_j^h = 0$ means that the member is illiterate.[4] We use the term *society* to refer to the vector of household literacy profiles $\mathbf{x} = (x^1, \ldots, x^m)$. So, for example, $\mathbf{x} = [(0, 1),$ $(1, 0, 0)]$ is a society of two households with two and three members, respectively, each of which has exactly one literate member. It is important to note that $\mathbf{x}$ conveys information on the household structure as well as the literacy levels in the country. The household structure can be ignored by concatenating the household vectors in $\mathbf{x}$ to obtain the *literacy profile x^o. For example, society* $\mathbf{x} = [(0, 1), (1, 0, 0)]$ *has the literacy profile $x^o = (0, 1, 1, 0, 0)$. Since we shall have occasion to vary the population and number of households in the country, we shall often use n_x and m_x to denote the (respective) numbers of adults and households*

---

[4] This makes the usual assumption that individual literacy is a $0 - 1$ variable. Of course, the underlying literacy variable is likely continuous and even multivariate (see for example Harris and Hodges, 1995, pp. 140-1) with arbitrary cutoff (see for example Peck and Kling, 1977). We abstract from these potentially important considerations.

in **x**. We denote the set of all possible societies (with arbitrary population size and number of households) by Δ.

A *measure of literacy* (*MOL*) is a mapping

$$\mathscr{L}: \Delta \to \mathscr{R},$$

from the set of all societies to the real numbers, where $\mathscr{L}(\mathbf{x})$ represents the overall level of literacy associated with society **x**. The traditional MOL is the *literacy rate R* defined by

$$R(\mathbf{x}) = \sum_i x_i^o / n_x$$

where $\sum_i x_i^o = \#\{i | x_i^o = 1\}$ is simply the number of literate persons in society **x**. By definition, $R(\mathbf{x})$ is the same for all societies having the same literacy profile x^o and, consequently, the household structure is ignored by R. We now turn to a new MOL whose definition crucially depends on the specific assignment of individuals to households.

We argued in the introduction that having a literate person in a household provides external benefits to illiterate members of the same household. For simplicity, let us assume that the magnitude of these benefits, say α (with $0 < \alpha < 1$), is independent of the characteristics of household members.[5] Then the *effective literacy profile for household h*, denoted $\tilde{x}^h$, can be defined as follows:

$$\tilde{x}_j^h = \begin{cases} 1 & \text{if } x_j^h = 1 \\ \alpha & \text{if } x_j^h = 0, \text{ and } x_k^h = 1 \text{ for some } k \neq j \\ 0 & \text{if } x_k^h = 0 \text{ for every } k. \end{cases}$$

The overall *effective literacy profile*, which we denote by x^*, is simply the literacy profile obtained from the resulting vector of effective household profiles, i.e., $x^* = (\tilde{x}^1, \ldots, \tilde{x}^m)^o$. So if the society is given by $\mathbf{x} = [(0, 1), (1, 0, 0), (0, 0)]$, then $x^* = (\alpha, 1, 1, \alpha, \alpha, 0, 0)$. This transformation leaves the literacy levels of the literates and isolated illiterates unchanged while assigning every proximate illiterate the 'effective' literacy level α. The magnitude of α reflects the extent to which having a literate member of household augments an illiterate's capabilities in this regard – a kind of effective literacy *equivalence scale*.[6] More generally, x_i^* is an indicator of i's *access* to literacy functionings.

Our overall *measure of effective literacy* $\mathscr{L}^*$ is defined by

$$\mathscr{L}^*(\mathbf{x}) = \sum_i x_i^* / n_x$$

or the original literacy rate measure applied to the effective literacy profile. Fig. 1 depicts the ordered profile of effective literacy from lowest to highest

---

[5] The framework can be generalised to allow α to depend on characteristics like gender and age of household members. One such extension is explored in Section 5.

[6] This analogy between 'equivalent income' and 'effective literacy' is quite apt. Both approaches use information on household structure to transform the raw data – which in the case of equivalence scales may include the economies of scale from living together in a household unit. See, for example, Deaton and Muellbauer (1980) and Coulter *et al.* (1992 *a*, *b*).
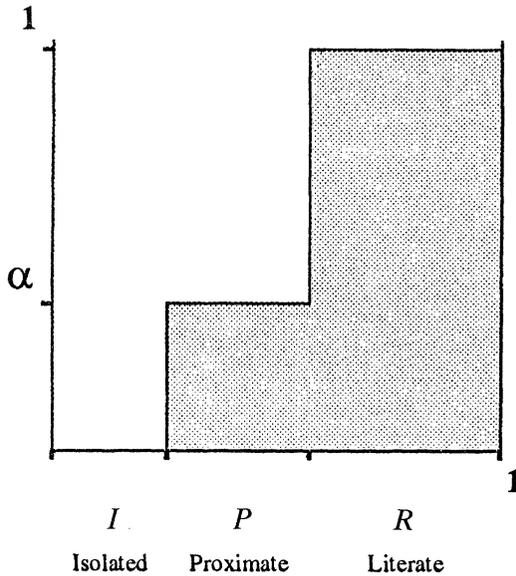
Fig. 1. *Effective Literacy Distribution*

levels, with isolated illiterates to the left, literates to the right and proximate illiterates in the middle. It is easy to see that $\mathscr{L}^*$ is the area below this graph or, equivalently,

$$\mathscr{L}^* = R + \alpha P,$$

the sum of the literacy rate R and α times P, the share of the population that is proximate illiterate.

Another related interpretation of $\mathscr{L}^*$ relies on the fraction I of isolated illiterates in the population:

$$\mathscr{L}^* = \alpha(1 - I) + (1 - \alpha)R.$$

The number $1 - I$ represents the fraction of the population with one or more literate persons in the household – which may be interpreted as an alternative indicator of literacy.[7] Our measure $\mathscr{L}^*$ is a convex combination of these two literacy indicators with α being the weight on $1 - I$. Note that if the externality indicator α were 0, the measure would reduce to the usual literacy rate R; if α were 1, so that a proximate illiterate were deemed equivalent to a literate, we would obtain the measure $1 - I$. In Section 4 we select $\alpha = 1/4$ and $\alpha = 1/2$ for reasons of illustration.

---

[7] See, for example, Rogers and Herzog (1966) or Sharma and Retherford, (1993, pp. 117-31), although both report the percentage of *households* that are isolated rather than the percentage of isolated illiterates.

© Royal Economic Society 1998

3. A Simple Characterisation

This section presents a number of intuitive properties for a measure of literacy to satisfy. It is then shown that these axioms completely characterise our effective literacy measure.

We say that society $\mathbf{x}$ is obtained from society $\mathbf{y}$ by a *permutation of households* if $(x^1, \ldots, x^m) = (y^{\pi(1)}, \ldots, y^{\pi(m)})$, where π is a permutation mapping.[8] When this definition applies, both societies are identical except that the household profiles are listed in different order. We say that $\mathbf{x}$ is obtained from $\mathbf{y}$ by a *permutation of individuals within a household* if x^h is a reordering of y^h for some household h, while the remaining household literacy profiles are unchanged. In this definition, only the members of household h are affected, and they are simply listed in a different order. Our first axiom says that the measure of literacy should not be affected by either type of permutation.

AXIOM A (*Anonymity*): If $\mathbf{x} \in \Delta$ is obtained from $\mathbf{y} \in \Delta$ by either a permutation of households or a permutation of individuals within a household, then $\mathscr{L}(\mathbf{x}) = \mathscr{L}(\mathbf{y})$.

This axiom requires the MOL to ignore the names of individuals or families in the evaluation of literacy. Moreover, it disallows any additional information on a person to be considered, such as the gender of the person or the gender of fellow household members.

We say that society $\mathbf{x}$ is obtained from society $\mathbf{y}$ by a *simple increment* if $x^h_j = 1$ and $y^h_j = 0$, while $x^{h'}_{j'} = y^{h'}_{j'}$ for all $(h', j') \neq (h, j)$. In other words, the only change in this case is that one person goes from being illiterate in $\mathbf{y}$ to being literate in $\mathbf{x}$. The next axiom requires the literacy measure to be consistent with this change.

AXIOM M (*Monotonicity*): If $\mathbf{x} \in \Delta$ is obtained from $\mathbf{y} \in \Delta$ by a simple increment, then $\mathscr{L}(\mathbf{x}) > \mathscr{L}(\mathbf{y})$.

If one person's literacy level rises, and the rest of the society stays unchanged, the Monotonicity axiom requires overall literacy to rise.

We say that society $\mathbf{x}$ is obtained from society $\mathbf{y}$ by a *household split* if $\mathbf{y}$ has m households, $\mathbf{x}$ has $m + 1$ households, y^m is the concatenation of x^m and x^{m+1}, and $x^h = y^h$ for all other households $h = 1, \ldots, m - 1$. In other words, the last household in $\mathbf{y}$ splits into the two final households in $\mathbf{x}$, while the rest of the households remain the same.[9] The split is called *externality-neutral* if either (i) both x^m and x^{m+1} contain a literate person, or (ii) neither of x^m or x^{m+1} contains a literate person. It is called *externality-reducing* if exactly one of x^m or x^{m+1} contains a literate person. The next axiom ensures that the MOL reflects the change in externality due to the changing household structure. If the split

---

[8] A permutation mapping π is a function on $\{1, 2, \ldots, m\}$ that is one-to-one and onto. Its effect is to change the order of the indexes.

[9] Given Axiom A, this definition is applicable to every household. A recent paper by Foster and Rosenzweig (1996*b*) develops and estimates a model of household splits to help understand how the Green Revolution affected the income distribution. Interestingly, their model posits a household education externality that influences an individual's decision to leave the household.

© Royal Economic Society 1998

does not affect anyone's access to a literate household member, then literacy is unchanged. If the split creates isolated illiterates, then literacy must decrease.

AXIOM E (*Externality*):   Suppose that $\mathbf{x} \in \Delta$ is obtained from $\mathbf{y} \in \Delta$ by a household split. If the split is externality-neutral, then $\mathscr{L}(\mathbf{x}) = \mathscr{L}(\mathbf{y})$; and if the split is externality-reducing, then $\mathscr{L}(\mathbf{x}) < \mathscr{L}(\mathbf{y})$.

Clearly this is the central axiom of our effective literacy proposal. In particular, if Axiom E is accepted, then the traditional literacy rate R is no longer viable as it regards all household splits with indifference.

We say that society $\mathbf{x}$ is *completely literate* if $x_j^h = 1$ for all h and all j; $\mathbf{x}$ is *completely illiterate* if $x_j^h = 0$ for all h and all j. The next axiom normalises the MOL.

AXIOM N (*Normalisation*):   If $\mathbf{x} \in \Delta$ is completely literate, then $\mathscr{L}(\mathbf{x}) = 1$; if $\mathbf{x} \in \Delta$ is completely illiterate, then $\mathscr{L}(\mathbf{x}) = 0$.

The final axiom allows the MOL to be calculated from regional data and aggregated to obtain the original level of literacy. We say that society $\mathbf{x}$ is *decomposed* into societies $\mathbf{y}$ and $\mathbf{z}$ if $y^h = x^h$ for all $h = 1, \ldots, n_y$ and $z^h = x^{n_y + h}$ for all $h = 1, \ldots, n_z$, where $n_x = n_y + n_z$. Note that this definition requires individual households to remain intact (i.e., no splitting) when society $\mathbf{x}$ is broken down into societies $\mathbf{y}$ and $\mathbf{z}$.

AXIOM D (*Decomposition*):   Suppose that $\mathbf{x} \in \Delta$ is decomposed into $\mathbf{y} \in \Delta$ and $\mathbf{z} \in \Delta$. Then $\mathscr{L}(\mathbf{x}) = (n_y / n_x) \mathscr{L}(\mathbf{y}) + (n_z / n_x) \mathscr{L}(\mathbf{z})$.

In other words, overall literacy in $\mathbf{x}$ is a weighted average of the levels of literacy in $\mathbf{y}$ and $\mathbf{z}$, with the weights being the population shares. Repeated application of Axiom D shows that the same relationship holds for any number of constituent societies. So, for example, if we know the effective literacy rates of Bihar, Kerala, West Bengal and the remaining states of India, then India's overall level of literacy is the weighted average of these component rates, with the weights being the states' (adult) population shares. Properties of this type have proved to be extremely useful in the measurement of poverty and inequality, particularly when it comes to analyses by region, ethnicity, or other population subgroups.[10]

Note that the traditional measure of literacy, R, satisfies each of Axioms A, M, N, D, and the first part of Axiom E. The second part of Axiom E is violated since R is unaffected by household splits. In addition, the other literacy measure $1 - I$ satisfies Axioms A, E, N, D, but not M, since a simple increment leaves $1 - I$ unchanged if it occurs in a household that already has a literate member. We now show that $\mathscr{L}^*$ is the unique measure satisfying all five of these axioms.

---

[10] See the careful analysis of Malaysian inequality and poverty by Anand (1983). A general discussion of decomposition postulates can be found in Foster and Sen (1997).

THEOREM 1: *A measure of literacy $\mathscr{L}$ satisfies axioms* A, M, E, N *and* D *if and only if $\mathscr{L} = \mathscr{L}^*$ for some α satisfying $0 < \alpha < 1$.*

Proof: Let $\mathscr{L}$ be a MOL which satisfies axioms A, M, E, N, and D. Consider any society $\mathbf{x} \in \Delta$. Applying Axiom D repeatedly yields $\mathscr{L}(\mathbf{x}) = \sum_h (n_h / n_x) \mathscr{L}(x^h)$, where n_h is the size of household h and $\mathscr{L}(x^h)$ is the level of literacy of the society containing the single household h. We need only determine the form of $\mathscr{L}$ over single household societies to derive the overall form of $\mathscr{L}$.

Let us define $r(x^h)$ and $s(x^h)$ to be the numbers of literate and illiterate members in x^h. Axiom A ensures that $\mathscr{L}(x^h) = \mathscr{L}(y^h)$ wherever $r(x^h) = r(y^h)$ and $s(x^h) = s(y^h)$. Consequently, if we define f by

$$f(r, s) = n_h \mathscr{L}(x^h),$$

where x^h is any household profile with $r(x^h) = r$ and $s(x^h) = s$, then f is a well-defined function. Moreover, from Axiom N we obtain $f(t, 0) = t$ and $f(0, t) = 0$, while Axioms D and E yield $f(t, t) = tf(1, 1)$, for all integers $t \geq 1$.

Now define $\alpha = f(1, 1) - 1$ and let $t \geq 1 + |r - s|$. Then by Axioms D and E,

$$f(r, s) + f(t, 0) = f(s, s) + f(r - s + t, 0)$$

for all integers r, s, t with $r, t \geq 1$. But then

$$f(r, s) = r + \alpha s \text{ for } r > 1$$

and hence

$$\mathscr{L}(x^h) = \frac{r(x^h) + \alpha s(x^h)}{n_h} \text{ for } x^h \text{ with } r(x^h) > 1.$$

In particular, $\alpha = 2\mathscr{L}(0, 1) - 1$, so that by Axioms E and N we have $\alpha > 2\mathscr{L}[(0), (1)] - 1 = 0$. Moreover, by Axioms M and N we have $\alpha < 2\mathscr{L}(1, 1) - 1 = 1$ and so $0 < \alpha < 1$ as desired. We will now show that the measure $\mathscr{L}$ must be the effective literacy measure $\mathscr{L}^*$ for this α.

Consider the effective literacy profile $\tilde{x}^h$ for household h. Clearly $\sum_j \tilde{x}_j^h = 0$ in the case where h in $\mathbf{x}$ has no literate members (i.e., $r(x^h) = 0$), while if h has at least one literate member (i.e., $r(x^h) \geq 1$), it follows that $\sum_j \tilde{x}_j^h = [r(x^h) + \alpha s(x^h)]$ by definition of $\tilde{x}^h$. In either case, then, $n_h \mathscr{L}(x^h) = \sum_j \tilde{x}_j^h$ by the above derivation. Substituting this into the decomposition $\mathscr{L}(\mathbf{x}) = \sum_h (n_h / n_x) \mathscr{L}(x^h)$ and recalling the definition of the effective literacy profile $x^* = (\tilde{x}^1, \ldots, \tilde{x}^m)^\circ$ yields $\mathscr{L}(\mathbf{x}) = \sum_i x_i^* / n_x = \mathscr{L}^*(\mathbf{x})$ as we originally set out to prove.

To complete the proof, one can easily check that $\mathscr{L}^*$ satisfies axioms A, M, E, N, and D. ∎

The above result offers support for our effective literacy measure $\mathscr{L}^*$ as the only MOL satisfying the five axioms. Of course, this is not a particularly noteworthy result if the axioms used to characterise the measure are not justified. The motivation for Axiom E has been rather extensively addressed;

but what about the remaining four axioms? Several arguments have been given on their behalf, including the fact that the traditional literacy rate R itself satisfies all four of them. But we can actually take this argument a step further. Consider the following alternative to Axiom E:

AXIOM I (*Indifference*): If $\mathbf{x} \in \Delta$ is obtained from $\mathbf{y} \in \Delta$ by a household split, then $\mathscr{L}(\mathbf{x}) = \mathscr{L}(\mathbf{y})$.

Axiom I is the embodiment of the observation that the literacy rate R is indifferent to household structure. It is easy to show the following result:

THEOREM 2: *A measure of literacy $\mathscr{L}$ satisfies axioms A, M, I, N and D if and only if $\mathscr{L} = R$, the traditional literacy rate.*

So the four basic axioms A, M, N, and D, provide a basic framework within which the measures R and $\mathscr{L}^*$ can be compared. If we add a requirement that literacy is indifferent to household structure, this leads to the above character-isation of R. But if Axiom I is rejected in favour of Axiom E, our effective literacy measure must be chosen. Whenever proximate illiterates become isolated illiterates as a result of a change in household structure, Axiom E requires acknowledgment of this fact; and this in turn justifies $\mathscr{L}^*$.

4. Literacy in India: An Illustration

Our previous discussion offered a conceptual basis for our measure of effective literacy. We now turn to an illustrative application of the measure to show how it may be used in practice.

The data we use are from the 1981 Indian Census (Registrar General of India, 1988, 1989) and the literacy tabulations of Sharma and Retherford (1993). Table 1 reports on several series drawn from the state-wise, cross-sectional data on adult literacy.[11] The first column is the traditional literacy rate R, where the states have been ordered from highest to lowest beginning with Kerala's 81.6% and finishing with Arunachal Pradesh's 25.6% level. The all-India rate of 43.3% has been calculated from the state-wise figures using the decomposition formula.

The Census also provides information on the number of literate members in households of various sizes, and this has been used to obtain the 'proximate illiteracy rate' P and the 'isolated illiteracy rate' I reported in the next two columns.[12] Our literacy measure $\mathscr{L}^*$ appears in Columns 4 and 5 for two values of α, namely, $\alpha = 1/4$ and $\alpha = 1/2$, while the non-isolated literacy rate $1 - I$ is reported in Column 6. The numbers in parentheses are the state's

---

[11] We exclude several states and union territories for ease of presentation. Note that India uses a cutoff of 7+ years of age in defining its 'adult' population for purposes of measuring literacy.

[12] Since the 'household size' variable reported in the Census is top-coded (at 6) and includes children, certain simplifying assumptions have been made – namely, (i) for households with six or more members, the mean household size for proximate and isolated households is the same; and (ii) the percentage of illiterates that are children is the same for proximate and isolated households. It turns out that our qualitative results are quite robust to alternative specifications.

© Royal Economic Society 1998

Table 1.

Literacy in Indian States: 1981

State	R	P	I	$\mathscr{L}^*$ $\alpha = 1/4$		$\mathscr{L}^*$ $\alpha = 1/2$		$1 - I$ $\alpha = 1$		G
Kerala	81.6	16.7	1.7	85.7	(1)	89.9	(1)	98.3	(1)	12
Mizoram	74.0	21.8	4.2	79.5	(2)	84.9	(2)	95.8	(2)	11
Goa	65.3	26.4	8.3	71.9	(3)	78.5	(3)	91.7	(3)	21
Maharashtra	55.8	31.4	12.7	63.7	(4)	71.6	(4)	87.3	(5)	29
Tamil Nadu	54.4	29.7	15.9	61.8	(5)	69.2	(6)	84.1	(7)	28
Gujarat	52.2	31.4	16.4	60.0	(7)	67.9	(8)	83.6	(8)	27
Himachal Pradesh	51.2	36.8	12.0	60.4	(6)	69.6	(5)	88.0	(4)	27
Nagaland	50.3	32.2	17.5	58.3	(9)	66.4	(9)	82.5	(11)	18
Tripura	50.1	29.9	20.0	57.6	(10)	65.1	(11)	80.0	(12)	23
Manipur	49.7	36.9	13.4	58.9	(8)	68.1	(7)	86.6	(6)	29
West Bengal	48.6	28.7	22.7	55.8	(12)	63.0	(13)	77.3	(15)	24
Punjab	48.2	34.7	17.1	56.9	(11)	65.5	(10)	82.9	(10)	16
Karnataka	46.2	32.4	21.4	54.3	(13)	62.4	(14)	78.6	(13)	26
Haryana	43.9	39.4	16.7	53.7	(14)	63.6	(12)	83.3	(9)	32
Meghalaya	42.0	28.7	29.3	49.2	(16)	56.4	(17)	70.7	(17)	9
Sikkim	42.0	35.4	22.6	50.8	(15)	59.7	(15)	77.4	(14)	26
Orissa	41.0	32.8	26.2	49.2	(17)	57.4	(16)	73.8	(16)	31
Andhra Pradesh	35.7	28.3	36.0	42.7	(18)	49.8	(20)	64.0	(22)	25
Madhya Pradesh	34.2	33.0	32.8	42.5	(19)	50.7	(19)	67.2	(19)	29
Uttar Pradesh	33.4	35.8	30.9	42.3	(20)	51.2	(18)	69.1	(18)	30
Bihar	32.1	32.0	35.9	40.1	(21)	48.1	(21)	64.1	(21)	30
Rajasthan	30.1	35.1	34.8	38.9	(22)	47.7	(22)	65.2	(20)	31
Arunachal Pradesh	25.6	28.1	46.3	32.6	(23)	39.6	(23)	53.7	(23)	21
India	43.3	31.7	25.0	51.3		59.2		75.0		

rankings according to each literacy measure; they indicate the extent to which the ordering by R corresponds to the orderings by $\mathscr{L}^*$ and $1 - I$. The 'gender gap' G between male and female literacy rates is listed in Column 7 for comparison purposes.

Let us first examine the overall picture of literacy as revealed by the all-India figures at the bottom of Table 1. The percentages of the population that are literate, proximate illiterate and isolated illiterate, are $R = 43.3\%$, $P = 31.7\%$ and $I = 25.0\%$, respectively. The proximate illiterate group comprises almost one-third of the entire population of India and more than half of the illiterates. Clearly a significant subset of the population has immediate access to some portion of the functionings typically associated with literacy – an empirical fact that is overlooked by studies that rely exclusively on the traditional rate R.

The quantitative impact on the assessment of literacy is captured by the extent to which our effective literacy rate $\mathscr{L}^*$ exceeds R. This in turn depends both on the observed P and the assumed magnitude α of the intrahousehold externality. When $\alpha = 1/4$, so that a proximate illiterate is 'equivalent' to one-fourth of a literate for purposes of evaluation, we obtain an effective literacy rate that is just above 51%. If α rises to $1/2$, the figure rises to about 59%. An upper bound on $\mathscr{L}^*$ is given by $1 - I = 75\%$; the remaining fourth of the population are isolated illiterates and make no contribution to $\mathscr{L}^*$.

We should emphasise that although our effective literacy measure $\mathscr{L}^*$ is

© Royal Economic Society 1998

higher than the traditional rate R, this in itself does imply greater incidence of literate persons nor increased satisfaction with the status quo. The approach simply offers a way of gauging actual achievement more accurately and differentiates between otherwise indistinguishable societies. Thus, $R = 43.3\%$ and $P = 31.7\%$ presents a rather different view of literacy achievements in India than $P = 0\%$ or for that matter $P = 56.7\%$. The measure $\mathscr{L}^*$ recognises and incorporates these distinctions.

Now moving to a state-wise view of literacy in India, we note a wide range of literacy rates R across the states, from 25.6% in Arunachal Pradesh to 81.6% in Kerala. There are also tremendous differences in the isolated illiteracy rate I, from 1.7% in Kerala to 46.3% in Arunachal Pradesh. It is interesting, therefore, to note that the proximate illiteracy rate P stays in a relatively narrow band – from 16.7% in Kerala to 39.4% in Haryana – about the all-India level of 31.7%. As we shall see, though, the state-wise variation in P is sufficient to lead to a reranking of states using $\mathscr{L}^*$ instead of R.

Before proceeding to this discussion, we should point out an interesting phenomenon: the three states with the highest levels of R have the lowest levels of P. The percentage of the illiterate population that is proximate illiterate (namely $P/(P + I)$) is certainly higher in the most literate states. But the lower absolute number of illiterates restricts the magnitude of P and hence the potential importance of the intrahousehold externality in evaluating literacy in high literacy regions. In this sense, our effective literacy approach may prove to be less important for assessments in developed countries – but is particularly well-suited for evaluating literacy in the developing world.[13]

Returning once again to Table 1, we see that apart from the especially low values of P for the top three states, there is no readily discernible relationship between P and R. High and low values of P can be found throughout the remaining distribution. The definition of $\mathscr{L}^*$ ensures that states with higher values tend to move up the distribution as α rises; while especially low values of P will have the opposite effect. For example, Manipur, with a value of $P = 36.9\%$ rises from the tenth spot to seventh as α goes from 0 to 1/2. West Bengal moves down the distribution from 11th to 13th over the same range of α, owing to its low value of $P = 28.7\%$. Table 2 focuses on the relative rankings of 10 'midrange' states as α takes on the values 0, 1/4, 1/2 and 1. In line with the above discussion, Himachal Pradesh, Manipur, Punjab and Haryana all rise at least two ranks as α increases from 0 to 1/2; Gujarat, Tripura, and West Bengal likewise fall due to their lower P values. When the intrahousehold externality is taken into account, the relative ranking of the states, and hence the regional picture of literacy in India, is clearly altered.[14]

This new way of viewing literacy leads us immediately to ask about the *causes* of the statewide variations in P and hence the rerankings of literacy according

---

[13] Indeed if all illiterates are proximate illiterates, then the rankings delivered up by $\mathscr{L}^*$ and R will be identical.

[14] The figure also depicts many cases of 'dominance' where the ranking of states is independent of α. For example, Nagaland is ranked lower than Tamil Nadu at both R *and* $1 - I$, and hence at all α between 0 and 1.

Table 2.

Literacy Rankings of 10 Indian States

R $\alpha = 0$	$\mathscr{L}^*$ $\alpha = 1/4$	$\mathscr{L}^*$ $\alpha = 1/2$	$1 - I$ $\alpha = 1$
Tamil Nadu	Tamil Nadu	Himachal Pradesh	Himachal Pradesh
Gujarat	Himachal Pradesh	Tamil Nadu	Manipur
Himachal Pradesh	Gujarat	Manipur	Tamil Nadu
Nagaland	Manipur	Gujarat	Gujarat
Tripura	Nagaland	Nagaland	Haryana
Manipur	Tripura	Punjab	Punjab
West Bengal	Punjab	Tripura	Nagaland
Punjab	West Bengal	Haryana	Tripura
Karnataka	Karnataka	West Bengal	Karnataka
Haryana	Haryana	Karnataka	West Bengal

to $\mathscr{L}^*$. Why, for example, is P so low in Tamil Nadu and so high in Haryana? While a complete answer is beyond the scope of this paper, we can at least speculate on the types of factors that may be pertinent. For example, to the extent that literacy is related to the status or caste of the individual, and assuming that most households stay within caste, this would tend to *lower P*. On the other hand, if literacy is associated with the gender of the individual then, assuming that most households have a mixture of genders, this would tend to *raise P*. Returning to the case of Haryana, note that both P and the 'gender gap' G between its male and female literacy rates (the last column of Table 1) are the highest in India. This suggests that the second of the two explanations might be applicable here. Tamil Nadu also has a high gender gap, and yet it has a modest level of $P = 29.7\%$. This points away from the second scenario and towards the first (or some other) alternative as a possible explanation of its lower ranking under $\mathscr{L}^*$ as α rises.[15] Formal testing of these and other potential hypotheses would, of course, require additional data on the gender, caste and other characteristics of household members.

5. Literacy and Gender

The above illustration has emphasised the potential differences between our new measure $\mathscr{L}^*$ and the traditional literacy rate R. In particular, the large variations in the proximate illiteracy rate P, can lead $\mathscr{L}^*$ to paint a rather different picture of literacy across states. But while the two measures may differ in their sensitivity to the distribution of literacy across households, they are quite similar in their treatment of gender, namely, both are *gender-neutral*. This section considers situations where a *gender-sensitive* literacy measure might be preferable.

First, following Anand and Sen (1995), one might specifically be 'concerned

---

[15] Other possible explanations for the variations in P include state-wise differences in the structure of families (prevalence of joint family households) or in the age structure of literacy (higher literacy in younger cohorts). On the latter possibility, see Margo's (1990, pp. 6–9) related discussion of literacy in the US South. There is a positive correlation between P and G, which suggests that a gender-based explanation may be important for several of the states.

© Royal Economic Society 1998

with inequalities in the opportunities and predicaments of women and men'
and consequently regard a higher gender gap – at a given average rate R – as
inherently undesirable. In response to this concern, Anand and Sen construct
a 'gender-equity-sensitive-indicator' for literacy based on the 'equally distribu-
ted equivalent (ede)' of Atkinson (1970). The resulting index is a measure of
overall achievement which, unlike R, is sensitive to the distribution of these
achievements across the two genders.[16]

The second motivation returns to our original, instrumental reasons for
caring about the distribution of literacy. Even when there is no particular
concern about gender inequality *per se*, if it so happens that females are more
effective than males in generating literacy externalities in the household, then
a *general* concern for greater effective literacy could translate into a *specific*
concern for greater female literacy. Indeed, there are many studies which
suggest the presence of a gender bias in the technology for creating external-
ities in the household, and hence that the set of functionings available to an
illiterate person by virtue of having access to a literate female is larger than the
set available when the only literate in the household is male.[17] The measure of
effective literacy can be modified to take this into account.

Distinguish between an *m-proximate illiterate*, who lives in a household with at
least one literate male and no literate females; and an *f-proximate illiterate*,
whose household contains at least one female literate. Compute an effective
literacy profile x^{**} by giving an effective literacy value of α_m to each *m*-
proximate illiterate and α_f to each *f*-proximate illiterate, while, as before
literates and isolated illiterates receive values of 1 and 0, respectively. Given the
above discussion, we take $0 < \alpha_m < \alpha_f < 1$. Then the *gender-adjusted measure of
effective literacy* is

$$\mathscr{L}^{**} = \sum_i x_i^{**} / n_x = R + \alpha_m P_m + \alpha_f P_f$$

where $P_m(P_f)$ is the share of $m(f)$-proximate illiterates in the population.

The measure $\mathscr{L}^{**}$ is clearly sensitive to the distribution of literacy between
the genders. But rather than preferring *equality* in literacy achievements, it
favours increases in female literacy because of the differential externality levels
(α_f versus α_m) conferred on illiterate household members.[18] Indeed, even if a
substantial gender gap in *favour* of female literacy were to exist, $\mathscr{L}^{**}$ would
recommend increasing the gap further by a 'transfer' of literacy from, say,

[16] Their index is $R_{ede} = (s_F F^{1-\varepsilon} + s_M M^{1-\varepsilon})^{1/(1-\varepsilon)}$, where s_F and F (respectively, s_M and M) are the
population shares and literacy rate for females (males), and $\varepsilon \geq 0$ is a parameter reflecting aversion to
inequality between the sexes. A similar methodology could be applied in the case of our effective
literacy measure by separating out the male and female components and aggregating according to the
ede.

[17] See, for example, the empirical findings of Murthi *et al.* (1995) or Nag (1983). This distinction
may be even more important if other indirect benefits of female literacy – in particular those received
by children – are taken into account. For example, infant mortality is well known to have a strong
negative correlation with maternal literacy (Caldwell, 1979).

[18] We are implicitly assuming that each is the first person to be educated in the household, and
hence the relevant comparison is between α_m and α_f. If most literates are males, however, the relevant
comparison may be between educating another isolated illiterate who is male or an *m*-proximate
illiterate who is female, in which case the magnitude of α_m versus $\alpha_f - \alpha_m$ becomes relevant.

male to female in a two-person household. Note also that while the gender-adjusted measure $\mathscr{L}^{**}$ satisfies Axioms A, M, N, and D (where the definition of a society is suitably altered to include gender information), it violates Axiom E since certain splits that were formerly 'externality-neutral' now alter the measured level of literacy.

6. Concluding Remarks

This paper presented a new approach to evaluating literacy and an effective literacy measure that take into account the intrahousehold externality arising from the presence of a literate member. A simple and natural characterisation of the new measure was given, along with an empirical illustration of its use. We also considered a 'gender-adjusted' specification in which the magnitude of the intrahousehold externality is sensitive to the gender of literate members.

Our new measures were constructed to include an important externality in the household – an externality missed by the traditional literacy rate. Consequently, they are likely to be superior to R in predicting or explaining other achievements that depend on literacy. For example, we can expect the rate of diffusion of a new technology for farming to be more closely linked to our effective literacy rate than to the usual literacy rate.[19] Similarly if the presence (or lack) of at least one literate household member influences the likelihood of children to become literate, the effective literacy rate should also be a better predictor of future generations' literacy levels. Of course, whether and to what extent our measures are better indicators of these and other attainments are empirical questions. But we have every reason to expect positive results in this regard.

In addition, changing the way that aggregate literacy is measured is likely to alter the perceived efficacy of actual literacy programmes, which in turn may influence their design. For example, R registers the same improvement when, say, (i) five persons from the same household become literate or (ii) five persons from different isolated illiterate households become literate. In contrast, $\mathscr{L}^*$ would see greater benefit from (ii) than from (i), while $\mathscr{L}^{**}$ will emphasise female literacy in its selection of the five persons in (ii). Consequently, we could expect some change in emphasis if these measures were used instead of the traditional literacy rate, namely, a shift towards ensuring that at least one person per household is literate (under $\mathscr{L}^*$) or focusing more fervently on the problem of female illiteracy (under $\mathscr{L}^{**}$).[20]

Finally, we should note that while our measure is based on the household as the unit of analysis, our general approach has much broader potential application. For example, suppose that each person in society has an extended 'social

[19] This hypothesis receives some support from the work of Green *et al.* (1985) and Foster and Rosenzweig (1996a).

[20] For example, a literacy programme that educates one female member in each completely illiterate, large household would raise $\mathscr{L}^{**}$ most rapidly. Practical considerations, though, can also affect the optimal design of programmes. For example, there may be significant economies of scale from teaching entire families, or entire villages, at the same time. Also, in many societies there are significant social barriers to female household members becoming literate ahead of male members. See, for example, Gustafsson (1991, p. 100).

© Royal Economic Society 1998

network' containing close friends and kin (and presumably the members of one's own household) who can be freely accessed for literacy (and other) needs.[21] Then our framework carries over immediately to this expanded environment. Simply define a proximate illiterate to be an illiterate person who has access to a literate person in his or her social network, and then let the effective literacy measure be the usual literacy rate R augmented by some share $\alpha > 0$ of the resulting proximate illiteracy rate P.

Of course, this kind of evaluation would require extensive information on relationships between people – information that is usually not available in conventional datasets. An alternative approach would be to partition the population into *observable* social units in such a way that most of these beneficial connections are likely to reside *within* a group and few are likely to extend *across* groups. Examples might include culturally defined units like ethnic groups, or geographically-defined units like neighbourhoods, villages and districts.[22] The underlying approach would still apply, but the definition of proximate literacy and the nature of the externality would have to be re-evaluated in this setting. For example, one could argue that as the average 'social distance' among members grows there will be a tendency for the literacy externality to become internalised in a 'market for literacy'. Literacy skills, like other personal services, can be withheld and therefore sold for a price; a literate person can charge a fee for writing a letter or for explaining the contents of a pamphlet. We have assumed that the groups are altruistic units with strong social ties, and hence that the full externality α is realised by illiterate members. But if literate and illiterate members are weakly linked, fees may be charged, which can result in lower net benefits from the externality. The exact form that such a market for literacy might take, and how its presence might alter α, are interesting topics for future exploration.

We should also note that there may be qualitative changes in the way α is configured as we shift among the various kinds of social units. For instance, in certain cases the extent of the externality could depend on the *percentage* of persons in the group who are literate, rather than on the mere presence of a literate member. The argument would be that a higher percentage of literate adults may ensure a more frequent access to literacy skills, and hence conveys a greater externality. Further work may shed some light on this important question.

Cornell University

Vanderbilt University

[21] Case studies of illiterate adults in the United States reveal intricate networks of social connections constructed by illiterates to help them function in a literate society. See, for example, Fingeret (1983) and Ziegahn (1991).

[22] See, for example, Benabou (1994), Borjas (1995) and Durlauf (1994) for discussions of external-ities in ethnic groups and neighbourhoods.

References

Anand, S. (1983), *Inequality and Poverty in Malaysia: Measurement and Decomposition.* New York: Oxford University Press.

Anand, S. and Sen, A. (1995), 'Gender inequality in human development: theories and measurement', Human Development Report Office Occasional Paper 19. New York: UNDP.

Atkinson, A. B. (1970), 'On the measurement of inequality', *Journal of Economic Theory*, vol. 2, pp. 244-63.

Benabou, R. (1984), 'Human capital, inequality, and growth: a local perspective', *European Economic Review*, vol. 38, pp. 817-26.

Bhola, H. S. (1984), *Campaigning for Literacy.* Paris: UNESCO.

Bhola, H. S. (1994), *A Source Book for Literacy Work.* London: UNESCO.

Borjas, G. J. (1995), 'Ethnicity, neighborhoods, and human-capital externalities', *American Economic Review*, vol. 7, pp. 365-90.

Caldwell, J. C. (1979), 'Education as factor in mortality decline', *Population Studies*, vol. 33, pp. 395-13.

Clark, S., Hemming, R. and Ulph, D. (1981), 'On indices for the measurement of poverty', Economic Journal, vol. 91. pp. 515-26.

Coulter, F. A. E., Cowell, F. A. and Jenkins, S. P. (1992a), 'Differences in needs and assessment of income distributions', *Bulletin of Economic Research*, vol. 44, pp. 77-124.

Coulter, F. A. E., Cowell, F. A. and Jenkins, S. P. (1992b), 'Equivalence scale relativities and the extent of inequality and poverty', Economic Journal, vol. 102, pp. 1067-82.

Deaton, A. and Muellbauer, J. (1980), *Economics and Consumer Behaviour*, Cambridge: Cambridge University Press.

Durlauf, S. N. (1994), 'Spillovers, stratification, and inequality', *European Economic Review*, vol. 38, pp. 836-45.

Dreze, J. and Saran, M. (1995), 'Primary education and economic development in China and India: overview and two case studies'. In *Choice, Welfare and Development* (K. Basu, P. K. Pattanaik and K. Suzumura eds.). Oxford: Clarendon Press.

Fingeret, A. (1983), 'Social network: a new perspective on independence and illiterate adults', *Adult Education Quarterly*, vol. 33, pp. 133-46.

Foster, A. D. and Rosenzweig, M. R. (1996a), 'Technical change and human-capital returns and investments: evidence from the Green Revolution', *American Economic Review*, vol. 86, pp. 931-53.

Foster, A. D. and Rosenzweig, M. R. (1996b), 'Household division, inequality and rural economic growth', mimeo. Philadelphia: University of Pennsylvania.

Foster, J. E. and Sen, A. (1997), '*On Economic Inequality* after a Quarter Century'. Annexe to the expanded edition of *On Economic Inequality* by A. Sen. Oxford: Clarendon Press.

Foster, J. E., Greer, J. and Thorbecke, E. (1984), 'A class of decomposable poverty measures,' *Econometrica*, vol. 52, pp. 761-76.

Green, S. E., Rich, T. A., and Nesman, E. G. (1985) 'Beyond individual literacy: the role of shared literacy for innovation in Guatemala', *Human Organization*, vol. 44, pp. 313-21.

Gustafsson, U. (1991), *Can Literacy Lead to Development? A Case Study in Literacy, Adult Education, and Economic Development in India.* Arlington: University of Texas.

Guthrie, J. T. (1983), 'Equilibrium of literacy', *Journal of Reading*, vol. 26, pp. 382-44.

Harris, T. L. and Hodges, R. E. (1995), (eds) *The Literacy Dictionary: The Vocabulary of Reading and Writing,* Reading Association.

Hayes, E. R. and Valentine, T. (1989), 'The functional literacy needs of low-literate adult basic education students', *Adult Education Quarterly*, vol. 40, pp. 1-14.

Margo, R. A. (1990), *Race and Schooling in the South 1890-1950: An Economic History.* Chicago: University of Chicago Press.

Murthi, M., Guio, A.-C., and Dreze, Jean (1995), 'Mortality, fertility and gender bias in India', *Population and Development Review*, vol. 21, pp. 745-82.

Nag, M. (1983), 'Impact of social and economic development on mortality: comparative study of Kerala and West Bengal', *Economic and Political Weekly*, vol. 18, pp. 877-900.

Peck, C. V. N. and Kling, M. (1977), 'Adult literacy in the Seventies: its definition and measurement', *Journal of Reading*, vol. 20, pp. 677-82.

Registrar General of India (1988), *Census of India, 1981: Household Tables.* New Delhi: Controller of Publications, Government of India.

Registrar General of India (1989), *Census of India, 1981: Social and Cultural Tables.* New Delhi: Controller of Publications, Government of India.

Rogers, E. M. and Herzog, W. (1966), 'Functional literacy among Colombian peasants', *Economic Development and Cultural Change*, vol. 14, pp. 190-203.

Sharma, O. P. and Retherford, R. D. (1993), 'Literacy trends in the 1980s in India', Office of the Registrar General and Census Commissioner of India, Occasional Paper No. 4.

Sen, A. (1976), 'Poverty: an ordinal approach to measurement', *Econometrica*, vol. 44, pp. 219-31.

Sen, A. (1985), *Commodities and Capability*. Amsterdam: North-Holland.
Sen, A. (1987), *The Standard of Living*. Cambridge: Cambridge University.
UNDP (1990), *Human Development Report*. New York: United Nations.
Wagner, D. A. and Puchner, L. D. (1992), 'World literacy in the year 2000', *Annals of the American Academy of Political and Social Science*, vol. 520.
Ziegahn, L. (1991), 'Beyond reciprocity: exchange around literacy', *Adult Basic Education*, vol. 1, pp. 79-97.