MEASURING INEQUALITIES IN HEALTH IN THE PRESENCE OF MULTIPLE-CATEGORY MORBIDITY INDICATORS

ADAM WAGSTAFF[1] AND EDDY VAN DOORSLAER[2]

[1]School of Social Sciences, University of Sussex, Brighton UK [2]Department of Health Policy and Management, Erasmus University, Rotterdam, The Netherlands

SUMMARY

This paper considers the problems which arise in seeking to measure socioeconomic inequalities in health when the health indicator is a categorical variable, such as self-assessed health. It shows that the standard approach—which involves dichotomizing the categorical variable—is unreliable. The degree of measured inequality is found to depend on the cut-off point chosen and the choice of cut-off point to affect the conclusions one can reach about trends in or differences in health inequality. The paper goes on to propose an alternative approach which involves constructing a latent health variable and then measuring inequalities in this latent variable by means of a variant of the health concentration curve.

KEY WORDS—Inequalities in health, self-reported morbidity, latent variables.

A major problem confronting researchers seeking to measure socioeconomic inequalities in morbidity is that many—if not most—morbidity indicators are categorical rather than continuous. Obvious examples include the presence or absence of chronic and acute conditions, and self-assessed health. In contrast to the former two indicators, which are dichotomous, self-assessed health is typically measured on a scale comprising at least three categories. Our concern in this paper lies specifically with such multiple-category health indicators.

One approach to comparing the degree of inequality in self-assessed health in two years or two countries is simply to compare the proportions in, say, each income group falling into each self-assessed health category.[1] Such a comparison does not lend itself, however, to the quantification of inequality, since the appropriate inequality measures—notably the slope index of inequality[2,3,4,5] and the health concentration index[6,7]—require information on health in the form of either a continuous variable or a dichotomous variable. The approach adopted to date by researchers seeking to measure inequalities in health using multiple-category health indicators has been to dichotomize the self-assessed health indicator by choosing a cut-off point. The measurement of socioeconomic inequality then proceeds by analysing the variation across, say, income groups in the percentages of persons reporting their health as, say, 'good' or worse than 'good'.[8,9]

This approach suffers from two obvious shortcomings. First, it entails a loss of information, since some of the cross-sample variation is lost as a result of the dichotomization. Second, the choice of the cut-off point is inevitably arbitrary. More importantly, the choice of cut-off point can affect the degree of measured inequality. In the empirical results reported below, the lower the cut-off point, the greater is the degree of

measured inequality. This can result in different conclusions being drawn concerning changes in inequality over time, depending on the cut-off point chosen. We suspect, though we do not have any empirical results confirming this, that the same may well be true in cross-country comparisons. This problem is compounded by the fact that the number of categories used in questions on self-assessed health varies across surveys. The self-assessed health question in the UK *General Household Survey*, for example, contains three categories, while similar questions in the US *National Medical Expenditure Survey* and the Dutch *Health Interview Survey* contain four and five categories respectively.

In this paper we propose an alternative approach to the measurement of inequalities in health in the presence of multiple-category morbidity indicators, involving the construction of a continuous standardized latent health variable. Since the standard health concentration curve cannot always be used to analyse inequalities in a standardized latent health variable, we also develop a new concentration curve which *is* suited to such a task. We illustrate our approach—and the deficiencies of the traditional approach—using data from the Dutch *Health Interview Survey*.

THE HEALTH CONCENTRATION CURVE

Our concern is with the measurement of relative socioeconomic inequality in health. By 'relative' inequality we mean that our measure of inequality should not depend on the mean level of health, so that if everyone's health doubled, measured inequality would remain unchanged. The issue of relative inequality is, of course, an issue of interest in its own right. But in the context of self-assessed health there is an additional motive for undertaking such an analysis, namely that the scope for an analysis of absolute inequality—at least in the context of international comparisons—is rather limited, since differences in overall proportions falling into each category almost certainly reflects differences in the wording of questions in the relevant questionnaires. By 'socioeconomic' inequality we mean that measured inequality ought to reflect the socioeconomic dimension to inequality. It ought to be able to distinguish, for example, between a situation where the least healthy in society are

millionaires and one in which the least healthy are unskilled manual workers.

Constructing the health concentration curve

Suppose we rank individuals by their socioeconomic status (as measured by their income, or social class, or whatever), beginning with the most disadvantaged. Suppose too that we have available a continuous measure of health— an assumption which we relax in a moment. Finally, assume that the data are grouped according to individuals' socioeconomic status. This latter assumption simply eases the exposition. It is not essential but is also not unrealistic, since data frequently *are* grouped in this context, not least because typically they are first age-sex standardized.

The health concentration curve—labelled g_{hlth} in Figure 1—plots the cumulative percentage of the sample (beginning with the most disadvantaged and ending with the least disadvantaged) against the cumulative percentage of health. Thus if f_t is the fraction of the sample in group t (T groups in total), x_t is the group's mean health, these cumulative percentages are given by

$$p_t = \sum_{\gamma=1}^{t} f_\gamma \quad \text{and} \quad c_t = \frac{1}{x} \sum_{\gamma=1}^{t} f_\gamma x_\gamma, \qquad (1)$$

where $x = \sum_{t=1}^{T} f_t x_t$ is the mean health of the sample. If health is equally distributed, g_{hlth} will coincide with the diagonal. If good health is concentrated amongst the least disadvantaged,
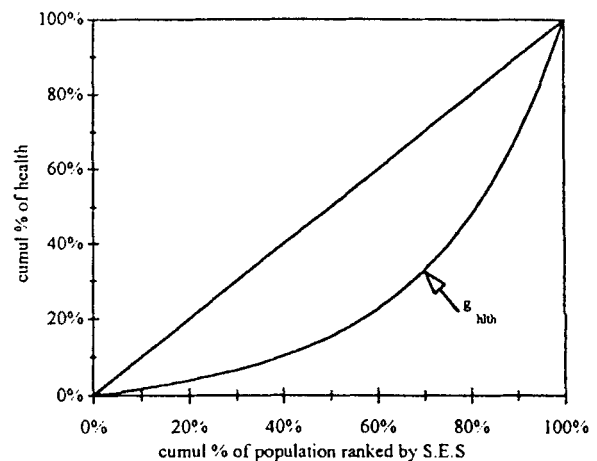

Fig. 1. Health concentration curve.

g_{hlth} will lie beneath the diagonal as in Figure 1, while the opposite will be true if good health is concentrated amongst the most disadvantaged.

The health concentration index, C, is defined as twice the area between g_{hlth} and the diagonal. If the concentration curve is assumed to be piecewise linear—other assumptions are possible[10]—the index can be computed simply using the formula[11]

$$C = \sum_{t=1}^{T} p_t c_{t+1} - \sum_{t=1}^{T} p_{t+1} c_t \qquad (2)$$

The index ranges from -1 (when only the most disadvantaged person is healthy) to $+1$ (when only the least disadvantaged person is healthy). It takes a value of zero when there are no inequalities in health and therefore g_{hlth} coincides with the diagonal. It is worth noting that the index also takes a value of zero when inequalities favouring (or disfavouring) the least disadvantaged are exactly offset by inequalities favouring (or disfavouring) the most disadvantaged. In this case g_{hlth} crosses the diagonal and the two areas between g_{hlth} and the diagonal before and after the cross-over point are the same size and hence cancel one another out.

The concentration curve can easily be adapted to handle situations where it is ill-health—rather than health—whose distribution is being investigated. In this case the concentration curve lies above the diagonal when ill-health is concentrated amongst the most disadvantaged. The concentration index is computed as twice the area between the illness concentration curve and the diagonal, and is negative when ill-health is concentrated amongst the most disadvantaged.

Relationship with slope index of inequality

As we have shown elsewhere,[7] the concentration index is closely related to the slope index of inequality (*SII*)—an index widely used in the literature on inequalities in health. The *SII* is equal to the slope of a regression of health on the relative ranks of the socioeconomic groups, a group's relative rank, R_t, being defined as

$$R_t = \sum_{y=1}^{t-1} f_y + \tfrac{1}{2} f_t, \qquad (3)$$

and therefore indicating the cumulative proportion of the population up to the midpoint of each interval. Since grouped data are being used, the appropriate estimation method is weighted least squares (WLS) rather than ordinary least squares.[7] If $\hat{\beta}_{WLS}$ is the WLS estimate of the slope from the regression of health on R_t, the concentration index can be calculated as

$$C = 2 \cdot \mathrm{var}(R_t) \, \frac{\hat{\beta}_{WLS}}{\bar{x}} \qquad (4)$$

where $\mathrm{var}(R_t)$ is the variance of the relative rank and $\bar{x}$ is the mean of x.

THE DANGERS ASSOCIATED WITH CHOOSING ARBITRARY CUT-OFF POINTS

Suppose that health—or ill-health—is measured by a categorical variable rather than a continuous variable. Where the variable is dichotomous and grouped data are being used, the obvious strategy is to calculate for each socioeconomic group the percentage of persons falling into the bottom category (e.g. reporting a chronic condition) and then plot the cumulative percentage of the sample (ranked by socioeconomic status) against the cumulative percentage of 'sick' persons.

This procedure can also be followed when the morbidity indicator has more than two categories, once the indicator has been dichotomized. However, this practice, as we indicated in the introduction, is not entirely satisfactory. To illustrate the problems we report some results obtained using data for adults from the Dutch *Health Interview Survey*. Socioeconomic status is measured by net income of the head of household plus any (net) income of the partner. The self-assessed health question reads 'How is your health in general?' and the five possible answers are 'very good', 'good', 'fair', 'sometimes good, sometimes bad' and 'bad'.[12]

Table 1 and Figure 2 show the implications of altering the cut-off point for the degree of

Table 1. Illness concentration indices for different cut-off points

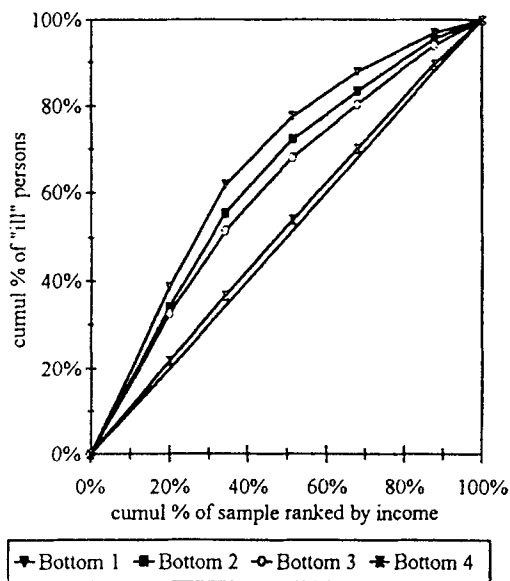|	1989/90	1991/92
Bottom 1 category	-0.3235	-0.3402
Bottom 2 categories	-0.2848	-0.2659
Bottom 3 categories	-0.2056	-0.2176
Bottom 4 categories	-0.0340	-0.0354

Fig. 2. Effects of health inequality of raising cut-off point.

measured inequality for 1989/90 and 1991/92. (The data have not been age-sex standardized, since we wish simply to focus on the implications of choosing different cut-off points and standardization is unlikely to make much difference to these conclusions.) As is apparent from Table 1 and Figure 2, as the cut-off point is raised, the illness concentration curve moves closer to the diagonal, coinciding with the diagonal as the cut-off point is raised to the top category. Thus the lower the cut-off point, the greater is the degree of measured inequality—a result which is consistent with the findings of other researchers. [8] The concentration indices in Table 1 also reveal that the ordering of the two pairs of years depends on where the cut-offpoint is set: measured inequality is smaller in 1991/92 than in 1989/90 if the cut-off point is set so as to include only the two bottom categories of the multiple-category variable in the bottom category of the new dichotomous variable, but higher in 1991/92 if any other cut-off point is chosen. These results clearly cast doubt on the reliability of the dichotomization approach.

AN ALTERNATIVE APPROACH

An alternative is to assume that underlying the categorical self-assessed health variable is a con-tinuous latent variable representing the individual's self-assessed health. Suppose that the categorical self-assessed health variable—which we denote by y—has J categories, with category 1 corresponding to the worst health and J corresponding to the best. If we denote the latent health variable by y^*, we can assume that y and y^* are related as follows:

$$
\begin{aligned}
y &= 1 \quad \text{if} \; -\infty < y^* \leqslant \alpha_1 \\
y &= 2 \quad \text{if} \; \alpha_1 < y^* \leqslant \alpha_2 \\
y &= 3 \quad \text{if} \; \alpha_2 < y^* \leqslant \alpha_3 \\
&\vdots \\
y &= J \quad \text{if} \; \alpha_{J-1} < y^* \leqslant +\infty
\end{aligned}
\tag{5}
$$

where the α_j are thresholds. Evidently it is important to ensure that the scale of y^* is independent of the number of categories J. Moreover, since our concern lies with relative inequality, we wish to abstract from any differences in mean health across surveys.

A standard normal latent health variable

A simple strategy—we consider an alternative in a moment—is to assume that y^* has a standard normal distribution. Values of y^* can then be computed for each respondent as follows. [13] The $J - 1$ thresholds are estimated as

$$
\hat{\alpha}_j = \Phi^{-1}\left(\sum_{i=1}^{j} n_i/N\right), \quad j = 1, 2, \ldots, J-1
\tag{6}
$$

where Φ^{-1} is the inverse standard normal cumulative density function, n_j is the number of cases in category j and N is the total number of cases. In effect, then, the thresholds are set so that the area under the standard normal distribution is divided up in proportion to the numbers falling into each of the J categories. Once the thresholds have been estimated, the mean values of y^* in each of the J intervals $\alpha_{j-1} < y^* \leqslant \alpha_j$ can be estimated as normal scores using the formula

$$
\hat{z}_j = (N/n_j)\,[\phi(\hat{\alpha}_{j-1}) - \phi(\hat{\alpha}_j)]
\tag{7}
$$

where the $\hat{z}_j$ are the normal scores in question and $\phi(\cdot)$ is the standard normal density function.

Using the method outlined above, each individual can be assigned one of the J values of y^*. The analysis of inequalities in health can then proceed by analysing inequalities in the continuous latent health variable y^* rather than analysing inequalities in the ordinal health indicator y. A moment's reflection will reveal, however, that the concentration curve device in Figure 1

cannot be used for such an analysis, since it requires that the mean of the variable whose distribution is being analysed be positive. (Otherwise one cannot calculate the percentage shares.) In our case the (weighted) mean of y^* is zero and hence the standard concentration curve cannot be used.

The literature on generalized Lorenz curves[14] suggests an alternative, namely a generalized concentration curve. Suppose that instead of plotting p_t against c_t as in the concentration curve diagram above, we plot p_t against c_t', where

$$c_t' = \sum_{\gamma=1}^{t} f_\gamma y_\gamma^*, \qquad (8)$$

which is simply the cumulative amount of latent health, $N \sum_{y=1}^{t} f_y y_y^*$, averaged across the sample. Evidently

$$c_0' = 0$$

and, since y^* is a standard normal variable, we also have

$$c_t' = 0.$$

If health does not vary across the sample,

$$c_t' = 0 \forall t$$

since $y_t^* = 0 \forall t$, and the latent health concentration curve simply runs along the horizontal axis. If, by contrast, ill-health is concentrated amongst the lower socioeconomic groups, these groups will tend to have negative values of y_t^* while the upper groups will tend to have positive values of y_t^*. If y_t^* rises monotonically as one moves up the socioeconomic distribution, the result will be a teardrop-shaped latent health generalized concentration curve along the lines shown in Figure 3. Clearly the further this curve is away from the horizontal, the greater is the degree of socioeconomic inequality in health.

It is worth noting that although the shape of our concentration curve is similar to that of the absolute Lorenz curve discussed by Moyes[15] and Jenkins,[16] the comparison is potentially misleading, since the zero mean of our latent health variable ensures that our concentration curve is relevant for the analysis of *relative* inequality rather than absolute inequality (the analysis of which clearly needs to allow for different means).

The above suggests that an appropriate index for measuring the degree of inequality in latent health would be (twice) the area between the hori-
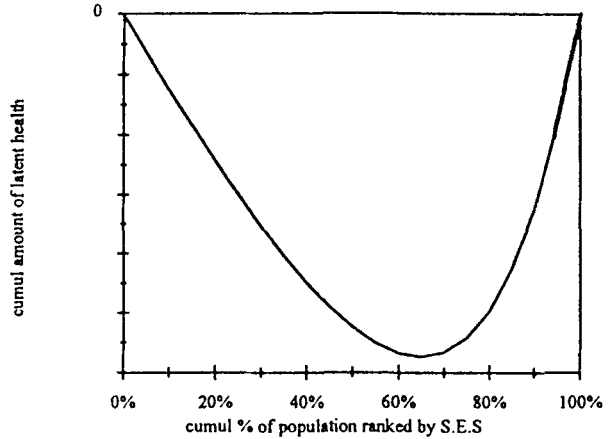

Fig. 3. Generalized latent health concentration curve.

zontal axis and our generalized concentration curve. If the concentration curve is assumed to be piecewise linear, it can be verified that this concentration index, C', can be calculated using the formula

$$C' = \sum_{t=1}^{T} p_t c_{t+1}' - \sum_{t=1}^{T} p_{t+1} c_t', \qquad (9)$$

which is positive when the concentration curve drops below the horizontal axis, zero when the concentration curve runs along the horizontal axis and negative when the concentration curve lies above the horizontal axis. C' is also zero when the concentration curve crosses the horizontal axis but the area to the left of the cross-over exactly equals the area to the right, the implication being that inequalities favouring, say, the better-off are exactly offset by inequalities favouring the worse-off

As is to be expected, there is a relationship between the proposed index of inequality and the slope index of inequality. It is well known in the income inequality literature that the Gini coefficient—defined as twice the area between the Lorenz curve and the diagonal—can be calculated by dividing the slope coefficient from a regression of income on the recipient's relative rank by mean income and multiplying by twice the variance of the relative rank variable.[17] (It is this result that was used in the discussion above to derive the relationship between the *SII* and the concentration index.) Since the generalized Lorenz curve is simply the standard Lorenz curve scaled up by mean income, it follows that twice the area between the generalized Lorenz curve and the

relevant diagonal is equal to the same regression coefficient multiplied by twice the variance of the relative rank variable. By analogy, our new concentration index, C', can be calculated by using WLS to run a regression of latent health on the group's relative rank, and then multiplying the slope coefficient from this regression by twice the variance of R_t. It can be verified that one indeed obtains precisely the same result using this regression method as one does using the formula for C' above.

A standard lognormal latent health variable

A weakness of the above approach is that it rules out any skewness in the distribution of the latent health variable. Since many measures of health are known to have skewed distributions (typically persons suffering from serious ill-health are in a minority and a large proportion of any general population sample report good health), it seems desirable to allow for skewness.

One possible strategy is to use the standard lognormal distribution rather than the standard normal distribution. To ensure that the latent health variable is skewed in the appropriate direction, we assume that $z = \ln[-y^*]$ has a standard normal distribution. We assume, in other words, that $y^* = -e^z$ has a standard lognormal distribution. Evidently, in contrast to the case where y^* was assumed itself to have a standard normal distribution, in this case y^* is decreasing in health. The mean and variance of y^* are respectively[18]

$$E[y^*] = E[-e^z] = -e^{0+0.5 \cdot 1} = 1.649$$

$$\mathrm{var}(y^*) = \mathrm{var}(e^z) = e^{2.0+1}(e^1 - 1) = 4.671,$$

so that again in contrast to the previous case, the mean of y^* is now positive. Calculation of the values of y^* is straightforward:[19] one simply divides up the area under the standard normal distribution and works out the normal scores as before, but now these are to be interpreted as the logarithms of the negative of the corresponding value of y^*. The corresponding values of y^* can therefore be obtained simply by exponentiating the normal scores and taking their negatives. Since y^* is always positive, the standard illness concentration curve can be used to assess inequalities in health.

SOME EMPIRICAL ILLUSTRATIONS

In this section we report the results of some empirical illustrations of the proposed new approach to measuring inequalities in health with multiple-category morbidity indicators.

Inequalities in health in the Netherlands: 1989/90 and 1991/92

The sensitivity of the ordering of the two pairs of years to the choice of cut-off point in the traditional dichotomization approach prompts the question: which pair of years has the more equal health distribution if the new approach is used?

Table 2 shows the frequencies and the estimated normal scores—the $\hat{z}_j$,—for 1989/90 and 1991/92 on the assumption that y^* has a standard normal distribution. Table 3 shows the sample frequencies and the mean values of latent health for each income group for the two pairs of years. The latter are graphed as concentration curves in Figure 4; the 1991/92 concentration curve cuts the 1989/90 curve three times, but for the most part lies further from the horizontal axis than the 1989/90 curve. This is reflected, naturally, in a larger concentration index for 1991/92. Thus although the concentration curves do not give an unambiguous result, the picture seems to be one of rising inequality in self-assessed health in the Netherlands over the period in question, though the increase is small and may not be statistically significant. Nonetheless, the result contrasts with the picture which emerged when the dichotomization approach was used with the cut-off point set at the second lowest category.

Tables 4 and 5 and Figure 5 show the corresponding results when y^* is assumed to have a

Table 2. Calculation of latent health variable assuming y^* has a standard normal distribution

| Category | 1989/90 | | 1991/92 | |
	f	y^*	f	y^*
1	2.4%	−2.3543	2.5%	−2.3404
2	5.3%	−1.6609	5.1%	−1.6595
3	12.6%	−1.0960	13.3%	−1.0859
4	54.2%	−0.0712	53.0%	−0.0699
5	25.5%	1.2590	26.1%	1.2467

Table 3. Inequalities in latent health—The Netherlands, 1989/90 and 1991/92 assuming y^* has a standard normal distribution

	1989/90		1991/92	
Income group	% sample	y^*	% sample	y^*
1	24.6%	−0.2502	19.9%	−0.2765
2	16.7%	−0.0394	14.4%	−0.1312
3	17.3%	−0.0047	17.2%	0.0076
4	16.1%	0.0987	16.5%	0.0974
5	16.1%	0.1961	19.8%	0.1183
6	9.1%	0.2290	12.2%	0.2737
Total	100.0%	0.0000	100.0%	0.0000
C'		0.0942		0.0978

Table 4. Calculation of latent health variable assuming y^* has a standard lognormal distribution

	1989/90		1991/92	
Category	f	y^*	f	y^*
5	25.5%	0.2839	26.1%	0.2875
4	54.2%	1.0738	53.0%	1.0724
3	12.6%	2.9922	13.3%	2.9620
2	5.3%	5.2638	5.1%	5.2564
1	2.4%	10.5309	2.5%	10.3851

standard log normal distribution. Note that in Table 4 the ordering of the self-assessed health categories has been reversed, reflecting the fact that y^* is now increasing in ill-health. The negative values of the illness concentration index indicate that inequalities in health favour the better-off and the indices' respective values are in accordance with the result obtained when y^* was assumed to have a standard normal distribution—inequality was apparently slightly higher in 1991/92 than in 1989/90. However, as is clear from Figure 5, the concentration curves cross (and therefore the result is not unambiguous) and, moreover, the curves are extremely close to one another.

Standardization for age and sex differences

The new approach proposed above can easily be adapted to allow the data to be age-sex standardized. Since standardization is usually undertaken using dichotomous variables, it may be useful to set out the procedure in the case where multiple-category variables are involved.

Table 5. Inequalities in latent health—The Netherlands, 1989/90 and 1991/92 assuming y^* has a standard lognormal distribution

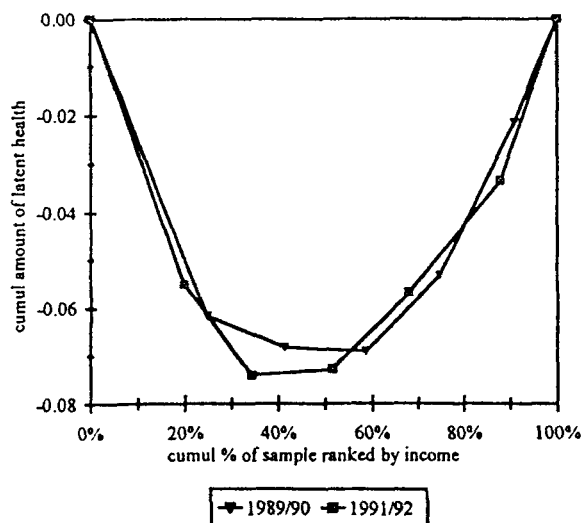|	1989/90		1991/92	
Income group	% sample	y^*	% sample	y^*
1	24.6%	2.0782	19.9%	2.1157
2	16.7%	1.6450	14.4%	1.8714
3	17.3%	1.5672	17.2%	1.5404
4	16.1%	1.3231	16.5%	1.3415
5	16.1%	1.1618	19.8%	1.2894
6	9.1%	1.1417	12.2%	1.0873
Total	100.0%	1.5624	100.0%	1.5646
C		−0.1220		−0.1254

Fig. 4. Generalized latent health concentration curves—The Netherlands, 1989/90 and 1991/92.

Recall first how the standardization works for dichotomous variables. Let m_{it} be the morbidity rate for demographic group i in socioeconomic class t and f_{it} be the fraction of the sample in the same group. Then the standardized morbidity rate for socioeconomic class t is

$$m_t^* = \sum_i f_i m_{it} \tag{10}$$

and the standardized number of persons reporting illness is

$$M_t^* = \sum_i f_i N_t m_{it}$$

$$= \sum_i \frac{f_i}{f_{it}} M_{it} \tag{11}$$

where N_t is the number of persons in socioeconomic group t and M_{it} is the number of persons reporting illness in demographic group i and socioeconomic group t. By analogy the standardized number of persons falling into morbidity category j can be calculated as

$$M_{jt}^* = \sum_i \frac{f_i}{f_{it}} M_{ijt} \tag{12}$$

where M_{ijt} is the number of persons in demographic group i in socioeconomic group t falling into morbidity category j.

Table 6 shows the unstandardized and standardized frequencies for each income group for the pooled sample of *Health Interview Surveys* for 1989, 1990, 1991 and 1992. Figure 6 shows the corresponding concentration curves on the assumption that y^* has a standard normal distribution. It is evident that the effect of the age-sex standardization is to reduce the degree of ine-
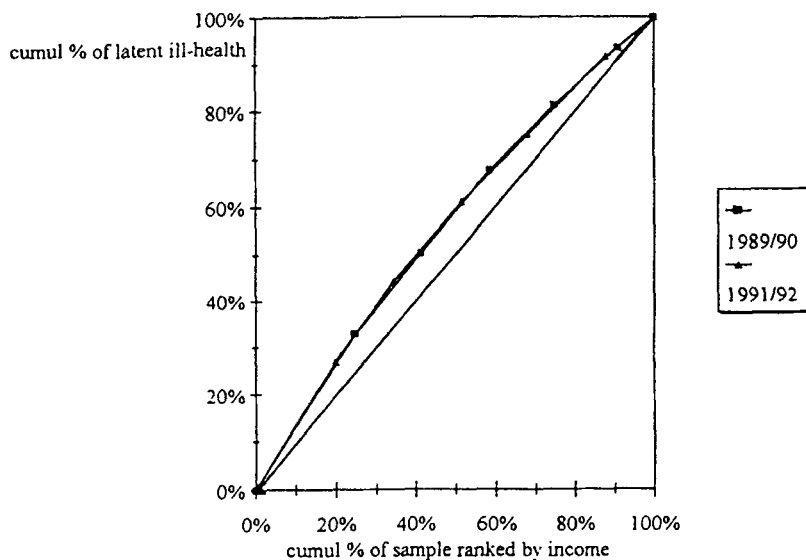

Fig. 5. Latent ill-health concentration curve—The Netherlands, 1989/90 and 1991/92.

Table 6. Standardized frequencies for self-assessed health variable

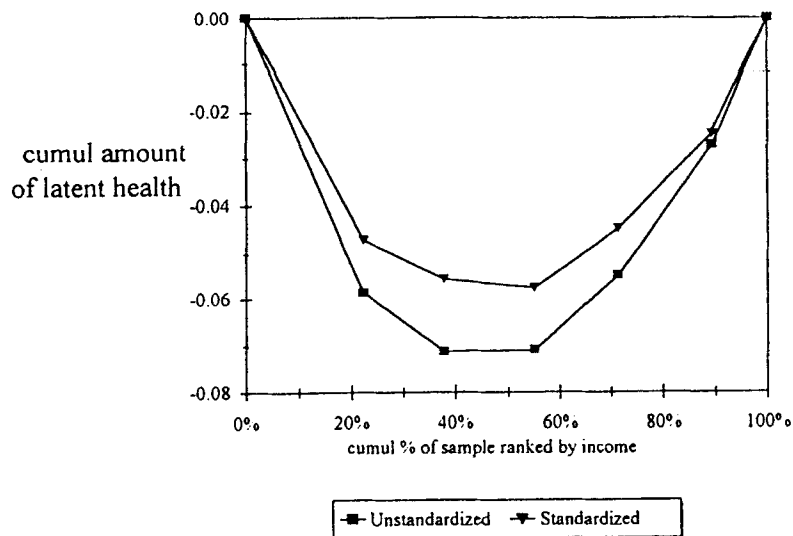| Income group	Unstandardized					Standardized				
	1	2	3	4	5	1	2	3	4	5
1	196	357	803	2023	781	197	347	709	2062	846
2	94	194	431	1480	703	87	187	415	1496	717
3	74	176	404	1744	826	77	183	422	1732	810
4	38	109	339	1719	847	41	112	364	1718	817
5	34	101	293	1948	975	43	116	346	1914	930
6	17	38	154	1093	682	16	39	178	1091	659
Total	453	975	2424	10007	4814	460	985	2434	10012	4778

Fig. 6. Effect of age-sex standardization on inequality in latent health.

quality in latent health—a result which is not altogether unexpected.

CONCLUSIONS AND DISCUSSION

It has been seen in this paper that dichotomizing categorical health indicators when analysing inequalities in health is potentially unreliable. We showed that the degree of measured inequality appears to depend on the cut-off point chosen—in our case, at least, the higher the cut-off point, the lower was the degree of measured inequality—and that one's conclusions concerning trends in (or

differences in) health inequality can be sensitive to the choice of cut-off point. We found in the case of the Netherlands, for example, that inequality in self-assessed health could have increased or fallen at the end of the 1980s depending on where the cut-off point is set in the categorical question concerning self-assessed health.

The paper went on to propose an alternative approach, which involves postulating the existence of a continuous latent self-assessed health variable underlying the categorical variable used in the health interview survey. In this paper we assumed that this latent health variable has either a standard normal distribution or a standard log-

normal distribution, the latter assumption having the attraction of allowing for the fact that most measures of health have a skewed distribution. An attraction of these assumptions is that they force the latent health variable to have the same mean (zero in the case of the standard normal assumption and 1.649 in the case of the standard lognormal assumption) and the same variance (one in the case of the standard normal assumption and 4.671 in the case of the standard lognormal assumption), irrespective of the number of categories in the original question and irrespective of the distribution of cases across these categories. The latter allows us to abstract from any differences in mean self-assessed health across surveys—something that is desirable given our interest in *relative* inequality. Having obtained values of the latent health variable for each case, we then measured inequalities in health with reference to a health concentration curve, appropriately modified in the case where the latent health variable is assumed to have a standard normal distribution. Applying this approach to the Dutch data suggests that—irrespective of whether latent health is assumed to have a standard normal distribution or a standard lognormal distribution—inequalities in health rose slightly at the end of the 1980s.

The approach outlined in this paper could, of course, be applied to any ordinal health indicator, though given that it is the sample proportions that determine the thresholds, it is only in the case of multiple-category indicators that the approach seems likely to be more useful than the dichotomization approach. Many health interview surveys ask whether the respondent suffers from specific chronic conditions and if so whether they limit his or her daily activities. The responses for each condition could be combined into a three-category ordinal indicator—(i) the person does not have the chronic condition in question, (ii) he or she has it but it is non-limiting, and (iii) he or she has it and it is limiting—and then the approach outlined in this paper could be used to assess the extent of inequalities in the chronic condition. Other possible uses of the method clearly exist.

There is one issue which merits discussion, namely the possibility that different income groups may have different expectations concerning their health. There is some evidence that the worse-off tend to underestimate their ill-health.[1] The implication is that we would expect people in,

say, the lowest income group to report their health as, say, fair when a better-off person with the same conditions would report it as being poor. This means that persons in the lower income group will tend to be in higher self-assessed health categories than is warranted given their objective health and *vice versa* for people in the higher income groups. Our value of y^* will therefore be too high for the lower income groups and too low for the higher income groups. The degree of 'true' inequality in health will therefore be understated by our approach. But the same is true, of course, of the dichotomization approach. Ideally one would like to allow for different mean values of y^* for different income groups (which we do) *and* for different thresholds (which we don't). We *could* allow for different thresholds in our approach but only by forcing each income group to have the same mean value of y^*. The only way to allow for both would seem to combine information on self-assessed health with 'objective' measures of health. This is a task we leave for future research.

REFERENCES

1. Blaxter, M. A comparison of measures of inequality in morbidity. In: Fox, J., (ed), *Health inequality in European Countries*. Gower: Aldershot, 1989.
2. Preston, S. H., Haines, M. R. and Pamuk, E. Effects of industrialization and urbanization on mortality in developed countries. In: *Solicited papers Vol. 2, IUSSP 19th International Population Conference, Manila*. Liege: IUSSP, 1981.
3. Pamuk, E. Social class inequality in mortality from 1921–1972 in England and Wales. *Population Studies*. 1985; **39**: 17–31.

4. Pamuk, E. Social-class inequality in infant mortality in England and Wales from 1921 to 1980. *European Journal of Population*. 1988; **4**: 1–21.

5. Wilkinson, R. G. Class mortality differentials, income distribution and trends in poverty 1921–1981. *Journal of Social Policy*. 1989; **18**: 307–335.

6. Wagstaff, A., van Doorslaer, E. and Paci, P. Equity in the finance and delivery of health care: some tentative cross-country companions. *Oxford Review of Economic Policy*. 1989; **5**: 89–112.

7. Wagstaff, A., Paci, P. and van Doorslaer, E. On the measurement of inequalities in health. *Social Science and Medicine*. 1991; **33**: 17–31.

8. Netherlands Central Bureau of Statistics. *International variations in socioeconomic inequalities in self-reported health*. The Hague: SDU Publishers, 1992.

9. Van Doorslaer, E., Wagstaff, A. and Rutten, F. (eds). *Equity in the finance and delivery of health care: An international perspective*. Oxford: Oxford University Press, 1993.

10. Kakwani, N. C. *Income inequality and poverty: Methods of estimation and policy applications*. Oxford: Oxford University Press, 1980.

11. Kakwani, N.C. and Podder, N. On the estimation of Lorenz curves from grouped observations. *International Economic Review*. 1973; **14**: 278–292.

12. Netherlands Central Bureau of Statistics. *Netherlands health interview survey 1981–1991*. The Hague, SDU Publications, 1992.

13. Jöreskog, K. G. and Sörbom, D. *PRELIS: A program for multivariate data screening and data summarization. A preprocessor for LISREL*. Mooresville IA: Scientific Software Inc, 1988.

14. Shorrocks, A. F. Ranking income distributions. *Economica*. 1983; **50**: 1–17.

15. Moyes, P. A new concept of Lorenz dominance. *Economics Letters*. 1987; **23**: 203–7.

16. Jenkins, S. P. Income inequality and living standards: changes in the 1970s and 1980s. *Fiscal Studies*. 1991; **12**: 1–28.

17. Lerman, R. I. and Yitzhaki, S. A note on the calculation and interpretation of the Gini index. *Economics Letters*. 1984; **15**: 363–368.

18. Maddala, G. S. *Econometrics*. McGraw-Hill: London, 1977.

19. Cowell, F. A. *Measuring inequality*. Philip Allan: Oxford, 1977.