UCD CENTRE FOR ECONOMIC RESEARCH

WORKING PAPER SERIES

2021

Health Poverty

Bénédicte Apouey, Paris School of Economics
David Madden, University College Dublin

WP21/16

July 2021

UCD SCHOOL OF ECONOMICS
UNIVERSITY COLLEGE DUBLIN
BELFIELD DUBLIN 4

Health Poverty

Bénédicte Apouey and David Madden

Bénédicte Apouey
Paris School of Economics – CNRS, 48 boulevard Jourdan, 75014 Paris, France.
benedicte.apouey@psemail.eu

David Madden
School of Economics, University College Dublin, Belfield, Dublin 4, Ireland.
david.madden@ucd.ie

1. Introduction

Poverty is typically measured with respect to some measure of individual or household resources, such as income or expenditure. However, there is no formal reason why poverty can not also be measured with respect to other important dimensions of the human condition, such as health or education. In this chapter, we review the analysis of health poverty. We predominantly focus on research which has appeared in the economics or health economics literature, and so much of our review will concentrate on the degree to which the methodology of poverty analysis as it stands for income poverty can also be applied in a health setting. We will see that this depends greatly upon the nature of the data available to the analyst, and in many cases the nature of health data will limit the range of poverty measures available.

The chapter proceeds as follows: we first of all review the different formats in which health data may appear, given its importance in determining which poverty measures can be calculated. In the light of this, we then discuss the nature of the "health poverty line." We then move on to examine health poverty when data is available in ordered, categorical format, and then when it is available in continuous form. We also review poverty dominance in a health context and conclude with an analysis of multidimensional poverty when health is combined with other dimensions such as income or education.

2. The Nature of Health Variables

The choice of a poverty indicator depends on how the health variable in question is measured. For this reason, we briefly review the different types of health variables. Specifically, following Erreygers and Van Ourti (2011), we distinguish between different types of measurement scales: nominal, cardinal, ratio-scale, fixed, and ordinal scales.

Nominal Variables

Nominal variables are outcomes for which ordering individuals from poorer to better health is not possible. For these types of variables, measuring poverty is not possible.

Cardinal, Ratio, and Fixed Scale Variables

In cardinal data, the differences between individuals have meaning, while ratios do not (Erreygers and Van Ourti, 2011). The zero point is arbitrary; there is no true zero. For example, body temperature, the McMaster Health Utility Index (HUI, which captures individual functional health and varies between 0 and 1), the Short-Form Six-Dimension (SF-6D), and the EuroQol five-dimensional (EQ-5D) health measures are measured on a cardinal scale. The HUI, SF-6D, and EQ-5D have the advantage of incorporating many dimensions of health and accounting for "quality of life." However, they measure health in units which are not always easy to grasp.

For ratio-scale variables, the ratios between individuals make sense. The zero point has a meaning: this is when the health outcome is absent. Health care expenditures and body length (such as the body mass index (BMI)) are examples of ratio-scale variables.

For fixed scale variables, the scale is fixed, i.e. unique. As with ratio-scale data, the zero captures the case when the health outcome is absent. For instance, the number of chronic conditions or of doctor visits (per year for instance) are measured on a fixed scale.

Among these different types of variables, the empirical literature on health poverty has used cardinal and ratio-scale health variables, as far as we are aware. Note that these data are continuous or pseudo-continuous.

Ordinal Data

While many variables relating to specific dimensions of health are continuous, information on overall health is often only available in ordinal format. Indeed, individual general health status is generally assessed using the "self-assessed health" (SAH) measure, which is ordinal and categorical. For ordinal variables, individuals may be ranked according to their health outcome, but neither differences nor ratios between individuals make sense. SAH is derived from the following question: "How is your health in general?", with the following typical response categories: "very good," "good," "fair," "bad," and "very bad."[1] This variable has a number of advantages: in particular, SAH is an independent predictor of mortality, even in models that control for other health status indicators and covariates (Idler and Benyamini, 1997).

However, one of the limitations of SAH is reporting heterogeneity. In other words, individuals may use different thresholds to answer the SAH question, even if their "true" health is the same. The differences in reporting behavior may be related to background characteristics (such as gender, age, education, cultural group, etc.). Lindeboom and Van Doorslaer (2004) distinguish between two sources of reporting heterogeneity: index shift ("a parallel shift of the thresholds") and cut-point shift ("a change of the relative positions of the reporting thresholds").

[1] This wording of the question is recommended by the WHO Regional Office for Europe.

Importantly, the presence of reporting heterogeneity may bias the measurement of health poverty. Reporting heterogeneity may also be present for the same individual over time, if people's expectation as to what constitutes "very poor", "poor," or "fair" health, for instance, changes over survey waves.

In the case of ordinal health data, the use of traditional poverty measures, such as the FGT index introduced by Foster et al. (1984), is not possible.[2] In that case, either ordinal data are transformed into continuous ones and traditional poverty indices can then be employed, or the nature of ordinal data is preserved but specific tools, especially developed for ordinal data, should be used. We discuss both of these approaches below, reviewing first how ordinal data may be transformed into continuous data.

The simplest way to convert SAH into continuous data is to assign a value to each health category. In particular, some studies assign the values 1-5 to the categories. Other linear scales or concave scales may also be used. Assigning these values is arbitrary though, and health rankings will be sensitive to the specific scale chosen.

Another approach to convert SAH into a continuous variable is to assume that underlying the categorical empirical SAH, there exists a latent and continuous health variable which captures individual health status. For example, in some models, Wagstaff and Van Doorslaer (1994) assume that latent health has a standard lognormal distribution. Each SAH category can then be assigned a value which equals the mean of the latent variable in the category. The technique is employed by Van Doorslaer et al. (1997) in their study of health inequality in nine countries.

Regression-based techniques have also been developed to construct a continuous health index from ordinal SAH. In this case, each individual is assigned a health score derived from a regression. For instance, Jones and Van Doorslaer (2003) transform the ordinal SAH variable into a continuous variable by regressing it against a set of factors, using an ordered probit model, computing the linear prediction, and rescaling this prediction, in particular in the [0,1] interval. They also show that an interval regression approach can be used to derive an individual health score, using external information from a HUI to set the bounds of SAH. This approach has mainly been applied to the analysis of social health inequality, but has also been applied to health poverty (Madden, 2011).

As shown by Allison and Foster (2004), inequality measures and rankings are sensitive to the choice of the cardinalization. Similarly, we argue that different cardinalizations of ordinal variables (such as SAH) may lead to different poverty rankings.

Finally, we note that in some instances, such as the presence or absence of a specific disease or condition, the health variable will be binary. Binary variables can also arise when ordinal, cardinal, and ratio- and fixed-scale variables are converted into binary outcomes by dividing

[2] Similarly, the use of traditional tools to measure inequality with this type of data is not possible. See Apouey (2007) and Abul Naga and Yalcin (2008).

the sample into individuals with a "poorer" health outcome and individuals with a "better" health outcome, by choosing a cut-off point. Importantly, this cut-off point is arbitrary, and the conclusion regarding the evolution of health poverty over time, or the difference in poverty between countries, will depend on where this cut-off is set. This limitation has already been highlighted in the case of the measurement of health inequality (Wagstaff and Van Doorslaer, 1994; Van Doorslaer et al., 1997). For binary variables, the only possible poverty measure is the headcount ratio that captures the proportion of individuals in poorer health.

3. Choice of a Poverty Line

Poverty measurement has traditionally been divided into identification and aggregation. First, we need to identify those who are poor, and then, we need to aggregate this information into a scalar index which satisfies various desirable properties. The poor are typically identified as those who fall below a critical threshold, depending upon the dimension over which poverty is being measured. In the case of income poverty, this will be a critical level of income (usually after adjustment for family size and composition) and analysts must make the choice of whether this level should be fixed (the "absolute" approach) or allowed to vary, given changes in living standards over time (the "relative" approach).

The choice of poverty line in the case of health will be dictated by the nature of the health data and the issue of fixed versus absolute poverty lines also arises. Take, for example, the case of a continuous measure such as life expectancy (which is ratio scale). This has been gradually rising over time and so it seems reasonable that a poverty line based on life expectancy should also rise to reflect these improvements. However, in the case of another continuous variable, Body Mass Index (BMI), an adjustment over time seems less reasonable (not least as BMI has been rising over time!). Thus, for the case of a continuous health variable, the fixed-versus-relative question needs to be determined on a case-by-case basis, depending upon the health outcome.

In the case of an ordinal, categorical variable such as SAH, then the analyst must choose a specific category as the poverty line. The choice will typically depend upon the wording of the SAH question (which can differ from country to country and from survey to survey). However, in almost all cases, given that SAH typically has five categories, those who are identified as health poor will be those in the two lowest (or sometimes three lowest, depending upon wording) categories.

We now turn to discuss aggregation in health poverty, dealing first of all with ordinal data.

4. Health Poverty for Ordinal Data

This section focuses on health poverty measurement for ordinal data. The simplest poverty measure that can be used with ordinal data is the headcount ratio, i.e. the proportion of the population whose health is below a certain line. In the case of SAH, a headcount ratio can be defined as the share of individuals reporting fair or very poor health, for instance. However, while it takes into account poverty frequency, this ratio cannot be sensitive to poverty depth and to the distribution of health among the poor, by definition.

Hatzimasoura and Bennett (2011) were the first to develop a counterpart of the classical FGT index for ordinal data. Assuming that the health variable Y has S categories and c_k is the category that serves as the poverty line, their ordinal FGT index is:

$$\pi_\alpha(Y; c_k) = \sum_{j=1}^{k} p_j (\frac{k-j+1}{k})^\alpha$$

where p_j represents the share of individuals in health category c_j and α is a strictly positive parameter.

The index is a weighted sum of the probability of being in a category, the weights depending on the deprivation rank. Authors interpret their index as capturing weighted headcount ratios when setting the poverty lines at all categories from 1 to k. When $\alpha = 0$, the measure is a headcount ratio. The ordinal FGT measure does not have the limitation of the headcount ratio which does not satisfy the monotonicity and transfer axioms. Indeed, when $\alpha > 0$, the index satisfies a monotonicity property, which means that the index is sensitive to poverty depth (an individual in a poorer health category contributes more to poverty). Moreover, when $\alpha > 1$, the index satisfies the transfer property, i.e. the index is sensitive to distribution (i.e. inequality) among the poor. For these reasons, authors argue that choosing $\alpha \geq 1$ is preferable in empirical work. In addition to the monotonicity and transfer axioms, the index satisfies ordinal invariance, additive decomposability, and subgroup consistency.

The authors provide an empirical illustration using data on SAH from the Joint Canada/United States Survey of Health (JCUSH). When $\alpha = 1$, health poverty is greater in the US than in Canada, for any poverty line. A decomposition by income quintiles highlights that the difference between the two countries is larger in the first income quintile.

Brzezinski (2015) provides statistical inference for the Hatzimasoura and Bennett index, and employs the index to study poverty in SAH in the UK between 1991 and 2008, using three poverty lines (1st, 2nd, and 3rd category) and three values of α (0, 1, and 2). The headcount ratio indicates a statistically significant increase in health poverty over the period. Results for $\alpha = 1$ and 2 show either an increase (when the poverty line is the 3rd category) or stability (when the poverty line is the 2nd category) over time. When the poverty line is the 2nd category, the results for the different values of α taken together thus mean that the increase in the frequency of poverty is compensated by the decrease in poverty depth and in inequality among the poor.

While highlighting that the headcount ratio is not sensitive to poverty depth and arguing that the measures proposed by Hatzimasoura and Bennett (2011) assess depth-sensitivity in a restrictive way, Seth and Yalonetzky (2020) develop a class of poverty measures for ordinal data which takes poverty depth into account and is given by:

$$(\boldsymbol{p}, c_k) = \sum_{s=1}^{S} p_s\, \omega_s$$

where category c_k is the poverty threshold, $\omega_1 = 1$, $\omega_{s-1} > \omega_s > 0$ for $s = 2 \ldots k$ when $k \geq 2$, and $\omega_s = 0$ for all $s > k$.

This class is a "weighted sums of population proportions in deprivation categories." The weights (ω) are non-negative for all categories and strictly positive for deprived categories. The weight of the poorest category equals one. These weights increase for greater levels of deprivation, which explains why the authors refer to them as "ordering weights." The weights act like an "implicit scale" (Silber and Yalonetzky, forthcoming). The index is similar to the headcount ratio when the poverty line is the first category and when the ordinal variable is binary.

Moreover, this class of indices is normalized between 0 (no one is poor) and 1 (everyone is in the most-deprived state). Regarding properties, this class does not only satisfy anonymity and the population principle (duplicating the population in each category does not alter the value of the poverty index), but also ordinal monotonicity (poverty decreases when a poor individual moves to a better category), single-category deprivation (the poverty measure is the headcount when the poverty line is equal to the poorest category), focus (changes in the situation of non-poor people have no impact on the level of poverty, provided they remain non-poor), and subgroup decomposability.

Adopting a prioritarian approach, the authors then add the concept of precedence of the poorer among the poor (i.e. priority is given to the welfare improvements of the poorest), which leads to some restriction on ordering weights. An empirical illustration on sanitation deprivation in Bangladesh is given, exploiting data from Demographic Health Surveys in 2007, 2011, and 2014. The results show the usefulness of the indices, especially in cases when the headcount ratio decreases, but the situation of the poorest of the poor does not improve.

5. Health Poverty with a Continuous Variable

We now review work where health data is continuous. For such data, the scope for aggregation is widened. Since the gap between a person's health and the health poverty line now has a clear meaning, and different gaps can be compared, it is possible to use a simple (or weighted) sum of these gaps as a poverty measure. Thus, we can use the measures traditionally associated with income poverty in a health setting. However, it should be remembered that a property such as transfer, which is important in the literature on income inequality and poverty, does not have a direct application in health. While person-to-person transfers of income are feasible, this is not the case in a health setting. Thus, for health poverty, transfer is an important property conceptually, but cannot be practically applied.

Adapting the popular Foster, Greer, and Thorbecke (1984) measure to health poverty, given a poverty line, z, and a measure of health for individual i, h_i, then the FGT P^{α} measure is given by:

$$P^{\alpha} = \frac{1}{N}\sum_{i=1}^{H}\left(\frac{z - h_i}{z}\right)^{\alpha}$$

where N is the number of people in the population, H is the number of people who are health-poor, and α is a parameter chosen by the analyst. The usefulness of the P^{α} measure derives from its flexibility and the way it nests three different approaches to poverty measurement via the choice of α. Thus $\alpha=0$ gives the traditional headcount measure, $\alpha=1$ provides a measure which is the sum of the proportional poverty gaps, while $\alpha=2$ provides a weighted sum of proportional gaps, with higher weights on the larger gaps and hence sensitivity to the distribution of health among the health-poor.

In terms of specific applications of the P^{α} approach in health, given that BMI is one of the more readily available continuous health measures, it is hardly surprising that some of the earliest applications of the P^{α} measure were in the area of obesity (Joliffe, 2004; Madden, 2012).

Other health applications of the P^{α} approach examine more general health outcomes which are not necessarily clinically measured but are constructed from clinical and survey-based data. For example, Clark and Erreygers (2020) apply the P^{α} approach to three different health measures: the probability that a patient will have a cardiovascular (CVD) event (e.g. heart attack, stroke) in a specified time period, the SF-6D health scores obtained from a health survey and, finally, estimates of life-expectancy derived from a health survey and standard life tables.[3] For the latter two measures, a relative poverty line is used, and for the CVD event, probability the authors use a threshold of 20% which they claim has historically been the value above which treatment is initiated. Similarly, Simões et al. (2016) calculate P^{α} measures for Portugal using EQ-5D scores derived from a combination of Portuguese health data and UK health parameters.

6. Multidimensional Poverty

In this section, we examine health poverty as a component in a multidimensional (MD) poverty measure. Researchers now measure poverty across a number of different dimensions, not just income (Alkire et al., 2015). Health is clearly a dimension which it is desirable to include, but once again the precise details of its inclusion will depend upon the nature of the health data available.

There are important methodological and practical issues to consider in construction of MD poverty indices and these are covered elsewhere in this volume. In this chapter, we will merely

[3] The probability of a CVD event is calculated by applying the algorithm of the Framingham Coronary Heart Disease Risk Score to US data from the National Health and Nutrition Examination Survey (NHANES). See Clark and Erreygers (2020) for details.

review applications of MD poverty which include health, and these will incorporate some of the different approaches to analysing MD poverty.

The dashboard approach to MD poverty involves examining poverty across a number of dimensions and also including a summary measure of the degree to which poverty is correlated across the dimensions. Madden (2015) takes this approach and finds such correlation present but that it diminished over the course of the Great Recession as younger, health-rich families were drawn into income poverty. Madden (2011) also examines health and income poverty in Ireland following the approach of Bourguignon and Chakravarty (2003) which uses a constant elasticity of substitution weighted average of the individual FGT indices, finding that changes in poverty over the 2003-2006 period show very little sensitivity to either the degree of substitutability between dimensions or the weighting attached to poverty gaps of the very poor.

Probably the standard approach to measuring MD poverty is the dual cut-off approach of Alkire and Foster (AF, 2011). A poverty line is chosen for each dimension (this is the first cut-off), and being poor in that dimension counts as a deprivation. The second cut-off is then how many deprivations qualifies as MD poverty. The two limiting cases are the intersection approach (must be poor in all dimensions) or the union approach (poor in only one dimension). In most applications, an intermediate number is chosen.

Depending upon the nature of data available, the AF approach can incorporate measures accounting for the depth of poverty, and also the distribution of poverty gaps amongst the poor, as well as headcount-based measures. There are numerous applications of the AF methodology. One of the most comprehensive is the Global Multidimensional Poverty Index which applies the AF methodology to over 100 developing countries (Alkire et al., 2020). Three dimensions of poverty are employed: health, education, and living standards, and within the health dimension, two indicators are used (nutrition and child mortality), both measured on a binary basis.

7. Dominance Measures

Poverty rankings can be sensitive to the choice of (a) poverty line and (b) specific poverty index. These are *partial* orderings as rankings may change with different poverty lines/indices. Poverty dominance examines the conditions under which *complete* orderings are possible holding for all (or a very wide set) of poverty lines/indices.

This approach, which was originally developed for income in Foster and Shorrocks (1988a, 1988b), may also be applied in the context of health. Given two populations, P and Q, with cumulative distributions of health $F_P(h)$ and $F_Q(h)$ respectively, Foster and Shorrocks show that stochastic dominance of order α of F_P over F_Q is equivalent to poverty dominance in the sense that distribution P has less measured poverty than distribution Q for poverty measures of the FGT class for all values of the health poverty line, h^*. The order of stochastic dominance, α, determines the class of poverty measure over which dominance holds. Thus, first order stochastic dominance of P over Q implies that for all headcount measures, P will have lower poverty, i.e. no matter where we draw the health poverty line h^*, the fraction of population with health lower than h^* is less in P than in Q (see Davidson, 2008).

If first order stochastic dominance does not hold, but second order dominance does hold, then P will have lower poverty than Q for all gap-based measures of poverty, while if third order dominance holds, then P will have lower poverty for all weighted gap-based measures. These results apply in the case of a continuous measure such as BMI, and an application of this approach can be seen in Madden (2012). Of course, in many applications (e.g. BMI), poverty dominance over the complete range of h^* may not be desirable, but Foster and Shorrocks show that their results also hold over limited ranges of the poverty line.

The situation where the health measure is categorical and ordered is discussed in Seth and Yalonetzky (2020). In this case, the health poverty "line" is a category and anyone in or below this category is deemed poor. For first order dominance, the Foster-Shorrocks results follow through quite intuitively. Distribution P poverty dominates distribution Q if for all possible poverty line categories, the fraction in P (i.e. the headcount) is lower than in Q (for health applications, see Madden, 2009 and Allison and Foster, 2004). The Foster-Shorrocks results do not follow though directly for higher orders of dominance, however, since there is no obvious scaling to apply to the categories so that "gaps" can be calculated and aggregated. The approach taken by Seth and Yalonetzky is to accord higher priority (i.e. higher weights) to those poor who are in lower categories and they present dominance results for a range of plausible weights derived from what they term a "prioritarian view" which holds that "benefitting people matters more, the worse off those people are" (Parfitt, 1997). Of course, if dominance is not found for any order, then the analyst must choose a specific poverty line and poverty measure to rank the two distributions and of course these specific poverty lines and measures are open to challenge.

Finally, poverty dominance can also be applied in a multidimensional setting (e.g. Duclos et al. 2006, who analyse poverty over expenditure and nutrition for Vietnam).

8. Conclusion

This chapter reviewed the literature on health poverty, for different types of health data including ordinal data. In the health field, there is often no admitted poverty line, in contrast with the literature on income. Regarding ordinal data, Seth and Yalonetzky (2020) argue that poverty indices use implicit scales, i.e. a form of cardinalization, that emerge in the weights (except when the headcount ratio is used). In this context, the axiomatic derivation of indices for ordinal data is useful in that the properties the index should satisfy guide the choice of these weights. For continuous health data, the FGT poverty index has been used in several empirical studies. Moreover, the study of multidimensional poverty often includes the health dimension, though the precise multidimensional measure used will be influenced by the nature of the health data available.

Importantly, studies of unidimensional poverty in the health field are not plentiful, in contrast with the large literature on health inequality and social health inequality. Future research may be interested in studying the evolution of this poverty during crises, such as the Great Recession or the COVID pandemic.

References

Alkire, S., Foster, J.: Counting and multidimensional poverty measurement. Journal of Public Economics 95(7-8), 476-487 (2011).

Alkire, S., Foster J., Seth, S., Santos, M. E., Roche, J. M., Ballon, P.: Multidimensional poverty measurement and analysis. New York, Oxford University Press, USA (2015).

Alkire, S., Kanagaratnam, U., Suppa, N.: *The global multidimensional poverty index (MPI) 2020*. OPH I MPI Methodological Note 49, Oxford Poverty and Human Development Initiative, University of Oxford, (2020).

Allison, R.A., Foster, J. E.: Measuring health inequality using qualitative data. Journal of Health Economics 23, 505-524 (2004).

Abul Naga, R. H., Yalcin, T.: Inequality measurement for ordered response health data. Journal of Health Economics 27(6), 1614-1625 (2008).

Apouey, B.: Measuring health polarization with self-assessed health data. Health Economics 16(9), 875-894 (2007).

Bilger, M., Kruger, E.J., Finkelstein, E.A.: Measuring socioeconomic inequality in obesity: Looking beyond the obesity threshold. Health Economics 26(8), 1052-1066 (2017).

Bourguignon, F., Chakravarty, S. R.: The measurement of multidimensional poverty. Journal of Economic Inequality 1, 25-49 (2003).

Brzezinski, M.: Accounting for trends in health poverty: A decomposition analysis for Britain, 1991-2008. The European Journal of Health Economics 16, 153-159 (2015).

Clarke, P., Erreygers, G.: Defining and measuring health poverty. Social Science & Medicine 244, 112633 (2020).

Davidson, R.: "Stochastic dominance" in The New Palgrave Dictionary of Economics. Second Edition. Eds. Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan (2008).

Duclos, J. Y., Sahn, D. E., Younger, S. D.: Robust multidimensional poverty comparisons. Economic Journal 116 (514), 943-968 (2006).

Erreygers, G., Van Ourti, T.: Measuring socioeconomic inequality in health, health care and health financing by means of rank-dependent indices: A recipe for good practice. Journal of Health Economics 30, 685-694 (2011).

Ferreira, F. H. G., Lugo, M. A.: Multidimensional poverty analysis: Looking for a middle ground. The World Bank Research Observer 28(2), 220-235 (2013).

Foster, J., Greer, J., Thorbecke, E.: A class of decomposable poverty measures. Econometrica, 761-766 (1984).

Foster, J. E., Shorrocks, A. F.: Poverty orderings. Econometrica 56(1), 173-177 (1988a).

Foster, J. E., Shorrocks, A. F.: Poverty orderings and welfare dominance. Social Choice and Welfare, 179-198 (1988b).

Hatzimasoura, C., Bennett, C.: Poverty measurement with ordinal data. Institute for International Economic Policy, IIEP-WP-2011-14 (2011).

Idler, E. L., Benyamini, Y.: Self-rated health and mortality: A review of twenty-seven community studies. Journal of Health and Social Behavior 38(1), 21-37 (1997).

Jolliffe, D.: Continuous and robust measures of the overweight epidemic: 1971-2000. Demography 41(2), 303-314 (2004).

Lindeboom, M., Van Doorslaer, D.: Cut-point shift and index shift in self-reported health. Journal of Health Economics 23, 1083-1099 (2004).

Lustig, N.: Multidimensional indices of achievements and poverty: What do we gain and what do we lose? Journal of Economic Inequality 9(2) (2011).

Madden, D.: Health and income poverty in Ireland, 2003-2006. The Journal of Economic Inequality 9(1), 23-33 (2011).

Madden, D.: Health and wealth on the roller-coaster: Ireland, 2003-2011. Social Indicators Research 121(2), 387-412 (2015).

Madden, D.: Mental stress in Ireland, 1994-2000: A stochastic dominance approach. Health Economics 18(10), 1202-1217 (2009).

Madden, D.: A profile of obesity in Ireland, 2002-2007. Journal of the Royal Statistical Society. Series A (Statistics in Society), 175(4), 893-914 (2012).

Parfit, D.: Equality and priority. Ratio (new series) 10(3), 202-221 (1997).

Ravallion, M.: On multidimensional indices of poverty. The Journal of Economic Inequality 9(2), 235-248 (2011).

Seth, S., Yalonetzky, G.: Assessing deprivation with an ordinal variable: Theory and application to sanitation deprivation in Bangladesh. The World Bank Economic Review (2020).

Silber, J., Yalonetzky, G.: Measuring welfare, inequality and poverty with ordinal variables. Handbook of Labor, Human Resources and Population Economics, Klaus Zimmermann, editor. Springer Nature (forthcoming).

Simões, N., Crespo, N., Moreira, S. B., Varum, C. A.: Measurement and determinants of health poverty and richness: Evidence from Portugal. Empirical Economics 50(4), 1331-1358 (2016).

Van Doorslaer, E., Jones, A. M.: Inequalities in self-reported health: Validation of a new approach to measurement. Journal of Health Economics 22, 61-87 (2003).

Van Doorslaer, E., Wagstaff, A., Bleichrodt, H., Calonge, S., Gerdtham, U. G., Gerfin, M., Geurts, J., Gross, L., Häkkinen, U., Leu, R. E., O'Donnell, O., Propper, C., Puffer, F., Rodríguez, M., Sundberg, G., Winkelhake, O/. Income-related inequalities in health: Some international comparisons. Journal of Health Economics 16(1), 93-112 (1997).

Wagstaff, A., Van Doorslaer, E.: Measuring inequalities in health in the presence of multiple-category morbidity indicators. Health Economics 3, 281-291 (1994).

UCD CENTRE FOR ECONOMIC RESEARCH – RECENT WORKING PAPERS

WP20/22 Cormac Ó Gráda: 'On Plague and Ebola in a Time of COVID-19' August 2020
WP20/23 Morgan Kelly: 'Understanding Persistence' September 2020
WP20/24 Jeff Concannon and Benjamin Elsner: 'Immigration and Redistribution' September 2020
WP20/25 Diego Zambiasi: 'Drugs on the Web, Crime in the Streets - The Impact of Dark Web Marketplaces on Street Crime' September 2020
WP20/26 Paul Hayes, Séin Healy, Tensay Meles, Robert Mooney, Sanghamitra Mukherjee, Lisa Ryan and Lindsay Sharpe: 'Attitudes to Renewable Energy Technologies: Driving Change in Early Adopter Markets' September 2020
WP20/27 Sanghamitra Mukherjee: 'Boosting Renewable Energy Technology Uptake in Ireland: A Machine Learning Approach' September 2020
WP20/28 Ivan Petrov and Lisa Ryan: 'The Landlord-Tenant Problem and Energy Efficiency in the Residential Rental Market' October 2020
WP20/29 Neil Cummins and Cormac Ó Gráda: 'On the Structure of Wealth-holding in Pre-Famine Ireland' November 2020
WP20/30 Tensay Meles and Lisa Ryan: 'Adoption of Renewable Home Heating Systems: An Agent-Based Model of Heat Pump Systems in Ireland' December 2020
WP20/31 Bernardo Buarque, Ronald Davies, Ryan Hynes and Dieter Kogler: 'Hops, Skip & a Jump: The Regional Uniqueness of Beer Styles' December 2020
WP21/01 Kevin Devereux: 'Returns to Teamwork and Professional Networks: Evidence from Economic Research' January 2021
WP21/02 K Peren Arin, Kevin Devereux and Mieszko Mazur: 'Taxes and Firm Investment' January 2021
WP21/03 Judith M Delaney and Paul J Devereux: 'Gender and Educational Achievement: Stylized Facts and Causal Evidence' January 2021
WP21/04 Pierluigi Conzo, Laura K Taylor, Juan S Morales, Margaret Samahita and Andrea Gallice: 'Can ❤️s Change Minds? Social Media Endorsements and Policy Preferences' February 2021
WP21/05 Diane Pelly, Michael Daly, Liam Delaney and Orla Doyle: 'Worker Well-being Before and During the COVID-19 Restrictions: A Longitudinal Study in the UK' February 2021
WP21/06 Margaret Samahita and Leonhard K Lades: 'The Unintended Side Effects of Regulating Charities: Donors Penalise Administrative Burden Almost as Much as Overheads' February 2021
WP21/07 Ellen Ryan and Karl Whelan: 'A Model of QE, Reserve Demand and the Money Multiplier' February 2021
WP21/08 Cormac Ó Gráda and Kevin Hjortshøj O'Rourke: 'The Irish Economy During the Century After Partition' April 2021
WP21/09 Ronald B Davies, Dieter F Kogler and Ryan Hynes: 'Patent Boxes and the Success Rate of Applications' April 2021
WP21/10 Benjamin Elsner, Ingo E Isphording and Ulf Zölitz: 'Achievement Rank Affects Performance and Major Choices in College' April 2021
WP21/11 Vincent Hogan and Patrick Massey: 'Soccer Clubs and Diminishing Returns: The Case of Paris Saint-Germain' April 2021
WP21/12 Demid Getik, Marco Islam and Margaret Samahita: 'The Inelastic Demand for Affirmative Action' May 2021
WP21/13 Emmanuel P de Albuquerque: 'The Creation and Diffusion of Knowledge - an Agent Based Modelling Approach' May 2021
WP21/14 Tyler Anbinder, Dylan Connor, Cormac Ó Gráda and Simone Wegge: 'The Problem of False Positives in Automated Census Linking: Evidence from Nineteenth-Century New York's Irish Immigrants' June 2021
WP21/15 Morgan Kelly: 'Devotion or Deprivation: Did Catholicism Retard French Development?' June 2021

UCD Centre for Economic Research
Email economics@ucd.ie