Mortality Comparisons and Age: a New Mortality Curve

John Creedy and S. Subramanian

WORKING PAPER 03/2022
February 2022

Working Papers in Public Finance

VICTORIA UNIVERSITY OF WELLINGTON
TE HERENGA WAKA

Chair in Public Finance

The Working Papers in Public Finance series is published by the Victoria Business School to disseminate initial research on public finance topics, from economists, accountants, finance, law and tax specialists, to a wider audience. Any opinions and views expressed in these papers are those of the author(s). They should not be attributed to Victoria University of Wellington or the sponsors of the Chair in Public Finance.

Further enquiries to:
The Administrator
Chair in Public Finance
Victoria University of Wellington
PO Box 600
Wellington 6041
New Zealand

Phone: +64-4-463-9656
Email: cpf-info@vuw.ac.nz

Papers in the series can be downloaded from the following website:
https://www.wgtn.ac.nz/sacl/centres-and-chairs/cpf/publications/working-papers

Mortality Comparisons and Age: a New Mortality Curve[1]

John Creedy

(Victoria University of Wellington)

and

S. Subramanian

(Independent Scholar; formerly, Madras Institute of Development Studies)

Abstract

This paper introduces a new mortality curve to illustrate and measure mortality and its relation to age. The curve draws on the 'Lorenz-Gini' framework of income-inequality measurement. The paper advances the cause of a 'mortality curve' analogous to the Lorenz curve, and a 'mortality-inefficiency' measure analogous to the Gini coefficient of inequality. The idea is to supplement the CDR with a mortality-inefficiency measure in a composite index of mortality which attends to both the mean and the dispersion of an age-distribution of deaths.

Keywords

Mortality Curve; Mortality-inefficiency measure; Crude Death Rate; Optimally distributed equivalent death rate; Lorenz Curve; Gini coefficient; Mean income; Equally distributed equivalent income

---

[1] Subramanian would like to thank, without implicating, Debraj Ray for many valuable discussions on the subject of this paper.

1. Introduction

> '…while the individual man is an insoluble puzzle, in the aggregate he becomes a mathematical certainty. You can, for example, never foretell what any one man will do, but you can say with precision what an average number will be up to. Individuals vary, but percentages remain constant. So says the statistician...' (Sherlock Holmes, in Arthur Conan Doyle's *The Sign of Four*).

The aim of this paper is to propose a new graphical device – a form of concentration curve – to describe and measure the combined influence of age-specific differences in mortality and the population age distribution. Of course, the central role of age is reflected in other measures and graphical approaches which are widely used by demographers and economists. These involve plots of the age distribution of deaths, profiles of life expectancy, and survival curves, in each case recognising gender and cohort differences. The use of an 'age-standardised' mortality rate (using a hypothetical standard population structure) is also often used to supplement the crude death rate, where the latter is expressed as a weighted sum over age groups of age-specific death rates, with weights equal to population shares. The age-adjusted death rate (ADR) is used instead of the crude death rate (CDR) because comparisons between countries or over time using the CDR conflate changes in age-specific mortality and the age distribution of the population. However, Curtis and Klein (1995, p. 3) suggest that 'the ADR is an artificial measure whose absolute value has no intrinsic meaning. The ADR is useful for comparison purposes only, not to measure absolute magnitude'.[2]

The curve proposed here takes its inspiration from the famous Lorenz curve, used to illustrate the inequality of incomes (or other non-negative variables), and defined, for a given distribution, as the 'first moment distribution' (the proportion of total 'income') plotted against the corresponding 'distribution function' (the proportion of people, ranked in ascending order). The mortality curve proposed here plots the proportion of total deaths against the corresponding proportion of people, where people are arranged from youngest to oldest. Since people are arranged in ascending order by age, rather than by their mortality rate, the curve is a type of concentration curve, and is referred to as an *M*-curve. Concentration curves are also familiar from other contexts, for example, poverty and income

[2] Thus, in exercises involving comparisons, the cardinal significance of any particular comparison is compromised by the fact that it is a variable function of the precise standard population employed. However, an alternative approach involves the use of decomposition methods: see Philip, Ray and Subramanian (2021).

mobility curves, and tax analysis where Lorenz-type curves are obtained for net income and individuals are ranked by pre-tax income.[3]

A further advantage of the curve – other than its ability to show 'at a glance' some important characteristics that are not evident from other profiles – is that a summary measure can be obtained, as with the Gini inequality measure and the Lorenz curve. The associated sub-optimality or inefficiency measure, I_M, is defined. Furthermore, the new mortality curve is shown to have some attractive properties in that it satisfies a number of basic consistency requirements. In addition, the curve allows for the explicit introduction of value judgements (regarding age at death) which may be adopted in policy debates. In particular, it is seen that the policy objective of reducing the CDR, or preventing a rise in the CDR, may not necessarily be consistent with other objectives relating to mortality, and revealed by the *M*-curve.

Section 2 defines the *M*-curve and its characteristics. It explains how value judgements regarding mortality can be explicitly introduced. Section 3 introduces an associated 'Generalised *M*-curve', which explicitly allows for a trade-off between the CDR and inefficiency. Section 4 provides more formal definitions of the curve and associated I_M measure. Section 5 shows the consistency of the approach with a number of basic criteria. Section 6 provides empirical examples. Brief conclusions are in Section 7.

2. The Mortality Concentration Curve

This section provides a basic description of the mortality curve. Subsection 2.1 defines the curve and explains the role of value judgements in making comparisons between curves for different countries or time periods. Subsection 2.2 introduces a measure of 'inefficiency' and discusses its relevance in policy evaluations.

2.1 The *M*-Curve Defined

The aim is to represent, in a single diagram, the essential combined characteristics of age-specific mortality rates and the population age distribution. Suppose, for observations over a specified time period (usually a year), individuals are ordered according to age, from youngest to oldest. The cumulative proportion of total deaths is then plotted against the corresponding cumulative proportion of the population. The resulting graph can be referred

[3] On poverty, see the TIP ('three Is of poverty') curves of Jenkins and Lambert (1997), and on mobility, see the TIM ('three Is of mobility') curves of Creedy and Gemmell (2019). On concentration curves in the measurement of tax progressivity, see Lambert (1993).

to as a ‘mortality concentration curve’, or ‘*M*-curve’, for short. It is contained within a box of unit ‘height’ and ‘width’, starting from coordinates (0,0) and ending at (1,1).

An example of a possible shape taken by the curve is given in Figure 1, showing the extent to which deaths are concentrated among the aged members of the population. If the age-specific mortality rate is identical for all ages, the curve follows the upward sloping diagonal line, and is independent of the age distribution.

Figure 1. A Hypothetical *M*-Curve

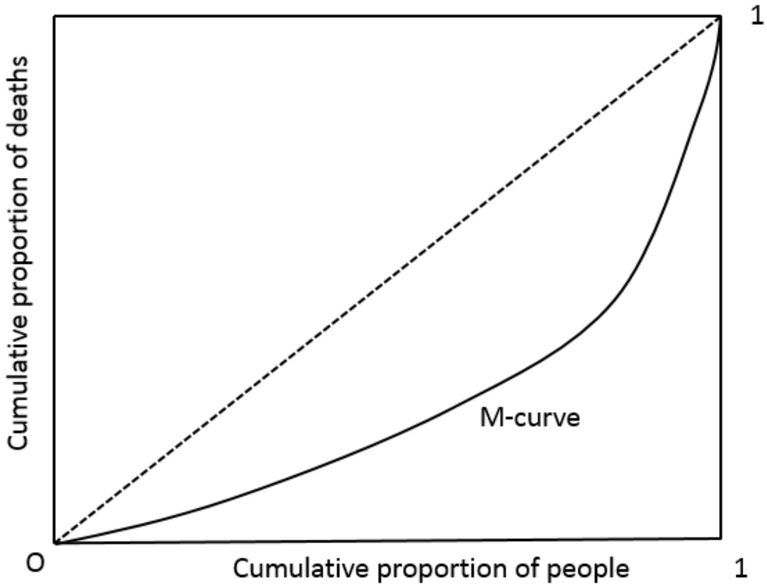

Unlike a Lorenz curve, and because individuals are first ranked by age rather than their age-specific mortality rate, there is no reason why the curve should be convex and lie always below the diagonal line from the left-hand lower corner to the top right-hand corner. In a society in which there is a relatively large degree of infant mortality, the curve might begin above the diagonal and have a concave shape, eventually moving below the diagonal and becoming convex.

This type of curve is illustrated in Figure 2. In practice, mortality and population frequency distribution data are usually available for a number of age groups, so that the curve consists of a number of piece-wise linear segments.

Figure 2. An *M*-Curve with High Infant Mortality

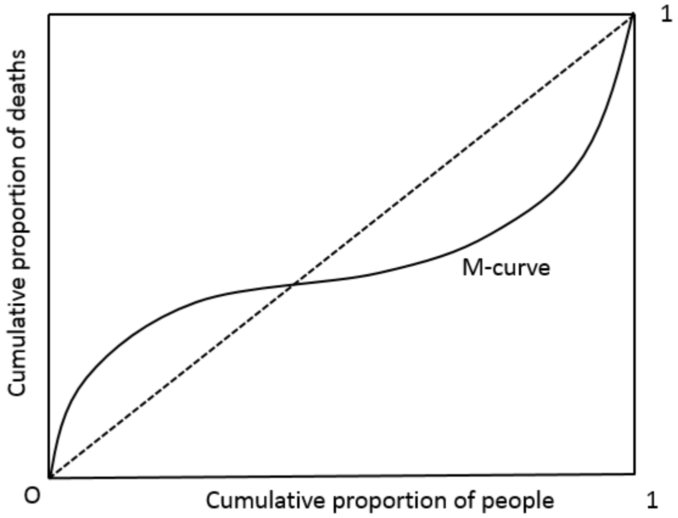

In the context of the Lorenz curve of income, it is relatively straightforward to define extremes of ‘equality’ and ‘inequality’. The former is represented by the upward-sloping diagonal line in the unit box, which arises when all incomes are the same. This presupposes that there are no relevant non-income differences between the individuals, or income units. Extreme inequality arises if one person only has all the income, while all other people have zero income: the Lorenz curve follows the lower and right-hand edges of the box. The statement that one Lorenz curve has less inequality than that of another distribution if it is everywhere closer to the ‘line of equality’ is founded on the value judgement summarised by the ‘Principle of Transfers’. This states that a transfer from a richer to a poorer person, which does not affect their ranks, is an improvement: such a transfer necessarily moves the Lorenz curve closer to the diagonal.

It must therefore be expected that any statements about a particular *M*-curve representing a preferred outcome, when compared with another curve, must also be based on explicitly-stated value judgements about mortality. Like the ‘principle of transfers’ it cannot be expected that such values receive universal support, only that they are shared sufficiently widely to warrant further analysis of the implications of holding them.

Suppose, then, that greater significance, or loss, is attached to a death, the lower the age at which it occurs. Again, this presupposes that there are no other relevant individual characteristics that may affect judgements. This implies that the ‘worst’ distribution would be one in which the deaths all take place in the youngest age-cohort: indeed, such a population could not be sustained. Correspondingly, the ‘best’ (or ‘least bad’) distribution is one in which all the deaths occur at the oldest age. In the former case the *M*-curve, denoted M_W, follows the left-hand side and top of the box, while in the latter case the *M*-curve, M_B, follows the base of the box and the right-hand side. This 'best' case could be said to represent a most 'efficient', or 'equitable' outcome. The term 'inequity' refers to differences in length of life, and the term 'inefficiency' may be thought to reflect in some sense a 'wastefulness of life-years' (before the biological maximum): the single term 'inefficiency' is used below.

Comparisons between countries or time periods can therefore be made using the respective *M*-curves. If one country has a curve that is everywhere closer to the M_B rectangle than the curve of another country, the former can be said unequivocally to display less inefficiency. Such comparisons are clearly independent of the overall death rate, CDR, since the values all range from 0 to 1. If the CDRs of two countries are identical, then the value judgements discussed above imply that the country with an *M*-curve closer to M_B than the curve of another country is judged to be preferred. The first is an 'improvement' on the second country: it involves a lower mortality-associated loss of welfare. Use of the terms 'preferred' and 'improvement', based on the adoption of quite basic value judgements, thus involves a step further than comparisons based on relative 'inefficiency'.

When evaluating *M*-curves, and debating public policy choices, a number of considerations may be thought to be relevant. For example, a country may have a very good healthcare system that avoids high child mortality, and also avoids extensive poverty, thereby ensuring healthy development. The country, further, has healthy nutritional standards and high standards of hygiene, with clean drinking water for everyone, and safe housing. Dangerous drugs are largely absent, as is conflict among groups. Safety standards in the workplace are high. In such a country most people may be expected to live a reasonably healthy life until death in old age. Some earlier deaths inevitably occur, associated with a range of causes not associated with poverty or poor social structure. But deaths are mainly concentrated in high age groups. The absence of these attributes contributes to the *M*-curve being closer to M_W.

An adverse change in the *M*-curve may arise at the same time as a reduction in the crude death rate. This also indicates possible social problems which need addressing. A contemporary case is perhaps connected with Covid. As a result of a strongly contagious virus, the CDR is of little use on its own as an indicator. For example, in many countries, Covid has mainly affected the aged and those with existing health problems. The *M*-curve may be little-changed by an attempt to reduce expected increases in the CDR. But this may not apply to those countries where there are seriously economic consequences, involving also access to health care, for people affected by lockdowns. The need to consider other dimensions of social policy and health care has been stressed by, for example, Ray and Subramanian (2020) and Gibson (2020).

2.2 A Measure of 'Inefficiency'

The previous discussion has involved *M*-curves which do not intersect, so that one is unequivocally closer to the 'most efficient' outcome. In practice the curves may well intersect, so that a quantitative measure of inefficiency is needed for comparisons. Corresponding to the Gini area-measure of income inequality, a measure of mortality 'inequity' or 'inefficiency' may be defined as the area between the observed *M*-curve and the M_B rectangle: this may be called the *I*-mortality measure, I_M. The Gini inequality measure 'normalises' the relevant area by dividing by the area contained within the extremes of inequality and equality. In the present context this area is unity.

The computation of I_M can be carried out in practice as follows. Generally, given the available data, *M*-curves are derived from grouped age distributions, so that the curves are actually formed by piece-wise linear segments. The area under the *M*-curve can therefore be computed by using the 'trapezoidal approximation method', which in the Lorenz curve context produces well-known expressions for the Gini measure. Suppose age is divided into *K* groups, 1,…,*j*,…,*K*. Let P_j be the cumulative proportion of the population whose age does not exceed the upper limit of the *j*th class, and Q_j the cumulative proportion of deaths at age not exceeding the upper limit of the *j*th class. Set $P_0 = Q_0 = 0$, and of course, $P_K = Q_K = 1$. The piece-wise linear *M*-curve is obtained by plotting the points $\{P_j,Q_j\}$ in the unit square, and connecting, with straight lines, the coordinates $(P_0,Q_0),(P_1,Q_1),\ldots,(P_K,Q_K)$. The inefficiency coefficient I_M is the sum of the areas of a number of trapeziums. It is given by:

$$I_M = (1/2)\sum_{j=1}^{K}(P_j - P_{j-1})(Q_j + Q_{j-1}). \tag{1}$$

Clearly, the larger the number of observations in (P_j,Q_j) space, the better is the approximation of the piece-wise linear curve to the actual *M*-curve. In particular, the fewer the observed coordinates of the *M*-curve, the smaller is the curvature of the piece-wise linear curve, and therefore the greater the under- (or over-) estimation of the true value of the inefficiency coefficient over the strictly concave (or strictly convex) segments of the *M*-curve

The M-curve and its associated inefficiency measure, I_M, can therefore be used to supplement information about the crude death rate. It has been seen that two societies can have identical CDRs but quite different M-curves. However, comparisons are more complex if societies have very different CDRs, even with agreement among judges about what values are attached to deaths at different ages. The following section shows how an explicit trade-off can be specified.

3 A Generalised *M*-Curve

This section examines the way in which comparisons can be made if the CDRs of two countries differ. Clearly some kind of trade-off is involved in making overall judgements. In the case of income distribution comparisons, a similar problem arises in making welfare comparisons using the Lorenz curve. In that context, Shorrocks (1983) showed that if the arithmetic means of the two distributions differ, the appropriate concept is that of the 'Generalised Lorenz (GL) curve', in which the values on the vertical axis of the Lorenz curve are multiplied by arithmetic mean income. A distribution with a Lorenz curve further from the line of equality may thus be preferred if the arithmetic mean income is sufficiently high that the Generalised Lorenz curve dominates that of the other distribution. The question therefore arises of whether a similar concept applies here.

In the income distribution context, the Gini measure was initially defined, as above, in terms of areas within the Lorenz curve. However, it is now well-known that the Gini, G, also arises from the adoption of a 'social welfare function' expressed as the Borda rank-order-weighted sum of incomes.[4] This is combined with the class of inequality measures defined in terms of the proportional difference between arithmetic mean income, μ, and an 'equally distributed equivalent' income (the equal income giving the same value of social welfare as the actual distribution).[5] This in turn gives rise to an 'abbreviated' welfare function, given by $W = \mu(1-G),$ which is itself equal to the equally distributed equivalent income.

Hence, the welfare index, W, is a function of total income per capita (as captured by μ) and the extent of equality in the distribution (as captured by $(1-G)$). Welfare is increasing in both μ and $(1-G)$: hence the trade-off between 'equity and efficiency' is explicit. To link this to the Generalised Lorenz curve, it is then necessary only to recognise that the area under

[4] It is a weighted sum of individuals' incomes, with weights equal to the 'inverse rank' of the individuals (with incomes arranged in ascending rank order). For further discussion, see, for example, Sen (1970).
[5] The famous Atkinson measure also belongs to this class, but has a different social welfare function than that giving rise to the Gini measure; see Atkinson (1970).

the Lorenz curve is $(1-G)/2$. Therefore, the area under the Generalized Lorenz curve is that of the Lorenz curve scaled by μ, and is thus $\mu(1-G)/2$: the above abbreviated welfare function is simply twice the area under the *GL* curve (Bishop *et al*., 2009).

The question therefore arises of whether equivalent results can be established in the present context. First, it is important to appreciate that a corresponding ‘welfare’ function is more appropriately described as an ‘ill-fare’ or loss function. Loss is captured by both the crude death rate, denoted *D*, and the ‘wastefulness’ or inefficiency of the distribution of deaths, as captured by I_M. An abbreviated loss function is thus $D^* = D(1+I_M)$: the loss is an increasing function of each of its arguments, *D* and $(1+I_M)$. The 1 is added to I_M to ensure that when there is zero inefficiency, the loss is *D*=D*. In the worst case of the *M*-curve, $I_M = 1$, so that *D** is twice the CDR.

These considerations suggest a Generalised Mortality curve, or *GM*-curve, which can be derived from the *M*-curve as follows. First, shift the *M*-curve up by the crude death rate, *D*. Then, scale the *M*-curve by multiplying by *D*. An example of a *GM* curve is shown in Figure 4. The area under this curve is a sum of two areas, A and B. Given the definition of I_M, area A is equal to DI_M, while Area B is equal to *D*. The sum of the two areas is thus $D(1+I_M)$, which is the abbreviated loss, *D**, as defined above. This also suggests, by analogy with the income distribution context, that the area under the Generalized *M*-curve reflects an ‘optimally-distributed equivalent death rate’, just as the area under the Generalised Lorenz curve is the equally-distributed equivalent income.

The Generalised Mortality curve can therefore be used to make ‘welfare loss’ comparisons between two countries or time periods. Figure 4 shows *M*-curves for two countries, 1 and 2. The curve for country 1 initially lies above that for country 2, before they intersect: it is clear from these curves that country 1 has the largest mortality inefficiency measure. Figure 5 shows the corresponding intersecting *GM* curves, reflecting the assumption that country 1 has a smaller CDR than country 2. Despite the lower death rate, it is also clear from Figure 5 that country 1 has a higher loss, *D**, than country 2, because the area under the *GM* curve is larger for country 1 than for 2. In judging country 2 to have a smaller ‘mortality loss’, the loss function trades off the higher crude death rate for a lower age-inefficiency.

Figure 3 A Generalised Mortality Curve

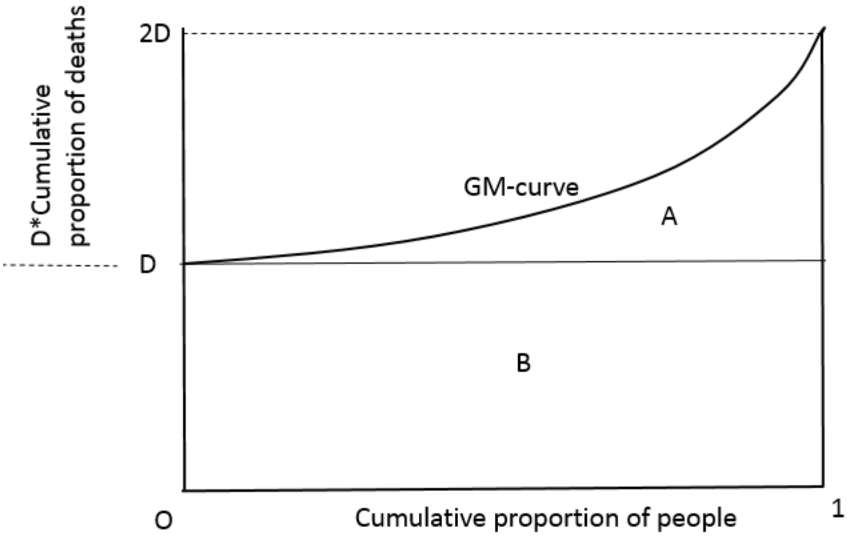

Figure 4 Two *M*-Curves

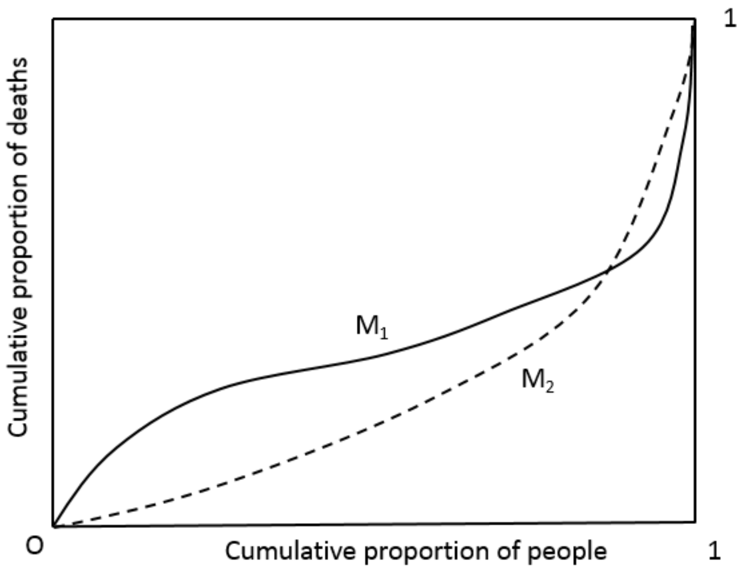

Figure 5 Two Generalised *M*-curves

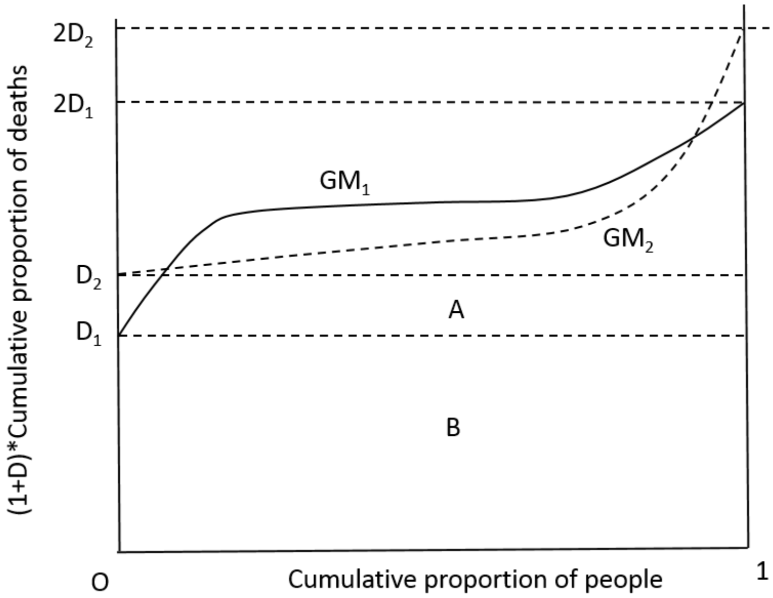

4. A More Formal Statement of The *M*-curve

This section provides a more formal statement of the *M*-curve. This is useful when examining the properties of the curve in the following section. Let a be a continuous random variable designating age, and distributed over the interval $A \equiv [0, \bar{a}]$, with $\bar{a} < \infty$. Let $n(a)(> 0)$ denote the size of the population aged a, and $m(a)(\leq n(a))$ is the number of deaths at age a, for all $a \in A$. The size of the total population is $n \equiv \int_A n(a)da,$ and the total number of deaths is $m \equiv \int_A m(a)da$. A population-mortality-distribution regime, or *regime* simply, is a pair of lists of age-specific population and age-specific death numbers. Regimes are designated by $R \equiv \langle n(a), m(a) \rangle_{a \in A}$, $\hat{R} \equiv \langle \hat{n}(a), \hat{m}(a) \rangle_{a \in A}$, and so on. The population density function is $\varphi(a) \equiv n(a)/n$ and the mortality density function is $\psi(a) \equiv m(a)/m$ for all $a \in A$. The corresponding cumulative distribution functions are $\Phi(a) \equiv \int_0^a \varphi(t)dt$ and $\Psi(a) \equiv \int_0^a \psi(t)dt$ for all $a \in A$. The Crude Death Rate, D, is the proportion of all deaths in the total population: $D = m/n$. In everything that follows, assume that $m \leq \min[n(0), n(\bar{a})]$: the reason for this assumption will be clarified shortly.

The M-curve is thus formally defined as the graph whose ordinates are given by the point (0,0) and the set of points $\{\Phi(a), \Psi(a)\}_{a \in A}$. For every regime R, the M-curve traces the function $M_a(R; \Phi(a)) = \Psi(a)$ for every $a \in A$. The curve commences at (0,0) in the unit square, is non-decreasing, and terminates at (1,1) of the unit square.

As explained above, the 'best-case' situation is defined as one in which all deaths occur at the oldest age $\overline{a}$: that this is feasible is guaranteed by the assumption that $n(\overline{a}) \geq m$. In this situation, the M-curve coincides with the horizontal axis for all values of $\Phi(a) < \Phi(\overline{a})$, and becomes unity at $\Phi(\overline{a})$: that is, the M_B-curve coincides with the curve describing the right angle formed by the lower and right sides of the unit square. In contrast, the 'worst'-case situation is the one in which all deaths are loaded on the youngest age 0 (the feasibility of which is again assured by the assumption that $n(0) \geq m$): that is, M_W coincides with the curve describing the right angle formed by the left and upper sides of the unit square.

The area under the M-curve, expressed as a proportion of the maximum such area (the difference between the areas under the M_W and M_B curves), is a natural normalised measure of how far away the regime is from the optimal one, and is interpreted as a measure of inefficiency. The difference between the areas under the M_W and M_B curves is the area enclosed by the unit square. Inefficiency, I_M, is then given simply for any regime, R, by:

$$I_M(R) = \int_A M_a(R; \Phi(a)) d\Phi(a) = \int_A \Psi(a) d\Phi(a) \qquad (2)$$

Integrating by parts, this is:

$$I_M(R) = \int_A (1 - \Phi(a)) \psi(a) da. \qquad (3)$$

Hence, I_M is essentially a weighted sum of age-specific shares of deaths, the typical weight at age a being the proportion of the population that is at least a years old. This is a distinct echo of the Borda rank-order weighting procedure that is employed in the derivation of the Gini coefficient of inequality, as discussed in the previous section.

5. Properties of *I*-Mortality and Unanimous Inefficiency Rankings

A 'mortality-inefficiency measure' T is a function which, for every regime R, specifies a real number which is intended to quantify the extent of inefficiency in the age-distribution of mortality associated with the regime. This section describes four properties which a mortality-inefficiency index might reasonably be expected to satisfy.[6] These are the axioms of *symmetry*, *aversion to young deaths*, *scale-invariance*, and *replication-invariance*. Symmetry requires that the measure of inefficiency should be invariant with respect to the precise personal identities of the individuals associated with any regime of population and mortality distributions. Aversion to Young Deaths demands that, other things equal, a 'transfer' of deaths from any age to a younger age causes the measure to rise; Scale-Invariance needs the inefficiency measure to be independent of the units of measurement (it does not matter whether age is measured in minutes or days or years); and Replication-Invariance is the property that the measure should be invariant with respect to age-specific replications of populations and deaths. To state these axioms more formally, first consider the following.

For all regimes $R, \hat{R}$:

$\hat{R}$ is said to be derived from R *through a permutation,* if the individuals associated with $\hat{R}$ are permutations, across ages, of the individuals associated with R;

$\hat{R}$ is said to be derived from R *through an old-to-young transfer of mortality*, if $\hat{n}(a) = n(a) \forall a \in A$, and [$\hat{m}(a) = m(a) \forall a \neq a', a''$ for some $a', a'' \in A$ satisfying $a' < a''$ and $0 < \hat{m}(a') - m(a') = m(a'') - \hat{m}(a'') \leq m(a'')$];

$\hat{R}$ is said to be derived from R *through a re-scaling*, when $R = \langle n(a), m(a) \rangle_{a \in A}$ and $\hat{R} = \langle n(\lambda a), m(\lambda a) \rangle_{a \in A}$, where λ is any positive scalar; and

$\hat{R}$ is said to be derived from R *through a k-replication,* where *k* is any positive integer, if $\hat{n}(a) = kn(a)$ and $\hat{m}(a) = km(a)$, $\forall a \in A$.

The Axioms, which apply to any mortality-inefficiency measure T defined on regimes of age-specific population and death distributions, can now be stated as follows.

Symmetry (Axiom S). For all regimes $R, \hat{R}$, if $\hat{R}$ has been derived from R through a permutation, then $T(\hat{R}) = T(R)$.

Aversion to Young Deaths (Axiom A). For all regimes $R, \hat{R}$, if $\hat{R}$ has been derived from R through an old-to-young transfer of mortality, then $T(\hat{R}) > T(R)$.

[6] This draws on Subramanian (2021).

Scale-Invariance (Axiom SI). For all regimes $R, \hat{R}$, if $\hat{R}$ has been derived from R through a rescaling, then $T(\hat{R}) = T(R)$.

Replication-Invariance (Axiom RI). For all regimes $R, \hat{R}$, if $\hat{R}$ has been derived from R through a *k*-replication, then $T(\hat{R}) = T(R)$.

It can be seen that the mortality-inefficiency measure, I_M, satisfies all four of these properties. In particular, that it satisfies Axiom A can be seen from the fact that I_M is a weighted sum of age-specific mortality shares, where the weight *declines* with age.

To consider whether there is a criterion by which unanimous rankings of regimes can be assured by a well-defined class of mortality-inefficiency measures, first define the notion of *M-dominance* analogously to the well-known concept of Lorenz-dominance in income inequality comparisons. Given any pair of regimes $R, \hat{R}$, R is said to *M-dominate* $\hat{R}$, written $R \succ_M \hat{R}$, if and only if the *M*-curve for *R* lies somewhere below and nowhere above the *M*-curve for $\hat{R}$; that is, if and only if $M_a(R; \Phi(a)) \leq M_a(\hat{R}; \hat{\Phi}(a)) \forall a \in A$ and the weak inequality is a strict inequality for some $a \in A$. If R *M*-dominates $\hat{R}$, the regime R is unambiguously less inefficient (in respect of the age-distribution of mortality) than the regime $\hat{R}$. Such an unambiguous judgement is not possible when *M*-curves intersect. The binary relation $\succ_M$ is a *strict partial ordering* (that is, a relation which is irreflexive, antisymmetric and transitive).

Given any class of inefficiency measures **T,** and any pair of regimes $R, \hat{R}$, R is said to **T**-*dominate* $\hat{R}$, written $R \succ_{\mathbf{T}} \hat{R}$, if and only if all measures belonging to the class $\mathbf{T}$ are united in ranking R as a less inefficient regime than $\hat{R}$; that is, if and only if $T(R) < T(\hat{R}) \forall T \in \mathbf{T}$. The binary relation $\succ_{\mathbf{T}}$ is a 'unanimity' partial ordering. Let $\mathbf{T}^*$ be the set of all mortality-inefficiency measures which satisfy Axioms S, A, SI and RI. Then, Foster's (1985) theorem on 'Lorenz-consistent inequality measures', adapted to the present context of mortality-inefficiency measures, suggests that, for any pair of regimes $R, \hat{R}$, if $R \succ_M \hat{R}$, then $R \succ_{\mathbf{T}^*} \hat{R}$. That is, the partial ordering $\succ_M$ is a sufficient condition for the partial ordering $\succ_{\mathbf{T}^*}$ to hold.

6. Empirical Illustrations

This section provides illustrations of alternative types of *M*-curve found in practice, showing how differences over time and between countries can be seen 'at a glance'. First, Figure 6 shows *M*-curves for India for 1971 and 2011. The data sources are provided in the Appendix. India's *M*-curves clearly reflect high child mortality, along with substantial improvement over the forty year period separating the two curves. The CDR halved from 0.0156 (or 15.6 per thousand) to 0.0071 (or 7.1 per thousand). Similarly, the inefficiency measure, I_M, dropped from 0.5432 to 0.2678, again approximately halving. Together these measures imply a drop in the value of the 'mortality loss function, *D**, from 0.024 to 0.009, a drop of 62.5 per cent.

Figure 6 Mortality Curves for India 1971, India 2011, and New Zealand 2019

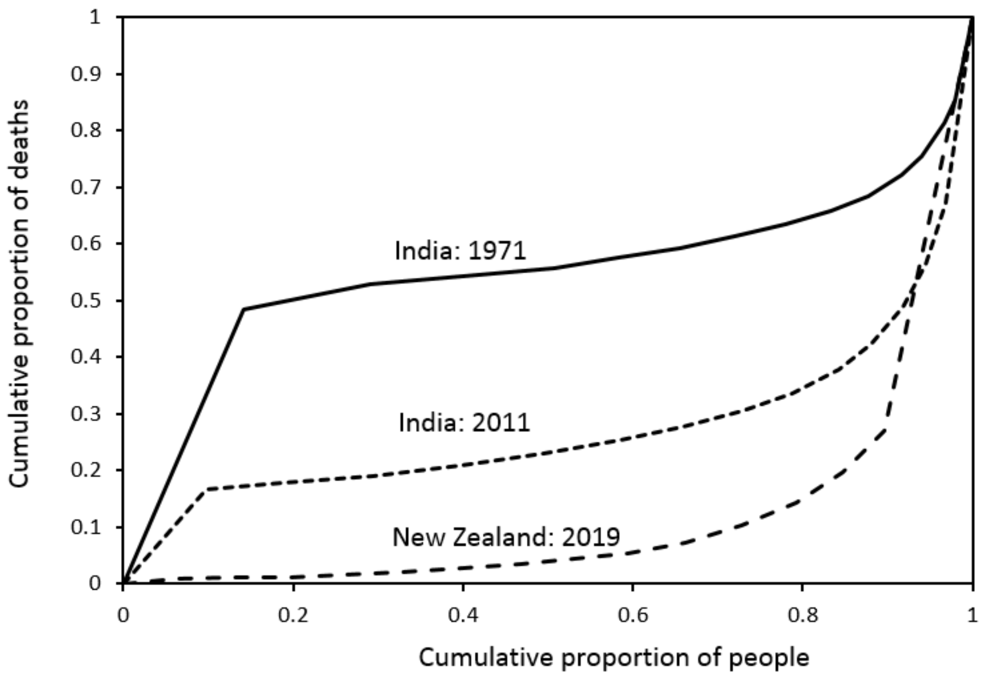

These results, for a 'developing country', may be contrasted with the mortality curve for a 'developed' economy, such as New Zealand. Figure 6 also includes the mortality curve for New Zealand in 2019. The greater 'efficiency' of NZ deaths is clearly shown by the convexity of the curve, compared with the sigmoid curves for India. The NZ curve is also considerably lower than the Indian curves, giving rise to an I_M value of 0.117. This is substantially lower than for India. Further, New Zealand's *D**, at 0.008, is lower than that of India in 2011 (0.009), although perhaps surprisingly the overall CDR for NZ in 2019 was similar to that of India in 2011. These examples demonstrate the value of supplementing CDR information with the *M*-curve and measures, I_M and *D**.

Finally, the observation above about the surprising near-equality of New Zealand's CDR in 2019 to India's CDR in 2011 requires some amplification. First, India's relatively 'low' CDR in 2011 might well have to do with inadequacies in the reporting of deaths. It is now well recognized that India's civil registration and vital statistics data could suffer from under-registration of deaths; see, for example, Basu and Adair (2021).

But apart from this, the observation in question is an invitation to see the need for supplementing information on a measure of central tendency with information on a measure of dispersion in order to obtain a more complete picture of mortality than is yielded only by the former. Table 1 presents information on age-group specific death rates for India in 2011 and New Zealand in 2019. It can be seen that for every age-group, the age-specific death rate for India exceeds that for New Zealand by a large factor, yet the average death rate for India is marginally *lower* than for New Zealand.

Table 1: Age-Specific Death Rates: India 2011 and New Zealand 2019

Age-Group	India, 2011: Age-Group Specific Death Rates (%)	New Zealand, 2019: Age-Group Specific Death Rates (%)	India's Age-Specific Death Rate/ NZ's Age-Specific Death Rate
0-4	1.22	0.11	11.09
5-9	0.10	0.01	10.00
10-14	0.07	0.02	3.50
15-19	0.13	0.04	3.25
20-24	0.16	0.06	2.67
25-29	0.18	0.06	3.00
30-34	0.23	0.08	2.88
35-39	0.27	0.09	3.00
40-44	0.40	0.13	3.08
45-49	0.55	0.20	2.75
50-54	0.83	0.32	2.59
55-59	1.22	0.47	2.60
60-64	2.01	0.70	2.87
65-69	3.32	1.08	3.07
70+	7.53	5.16	1.46
All Ages	0.71	0.73	0.97

Source: The relevant data sources can be found in the Appendix.

The reason is not hard to find, and is to be located in the age-distribution of the population. Specifically, the age-groups with the three highest death rates in both India and New Zealand are (unsurprisingly) the old-age groups of 60-64, 65-69, and 70+. The population in these age-groups as a proportion of total population is much lower for India, at 5.06 per cent, than for New Zealand, at 20.75 per cent. The New Zealand figure exceeds the India figure by a factor of more than 4. Briefly, the high-fatality age groups are much thinner on the ground in India than in New Zealand, which accounts for the fact that New Zealand's mortality statistics are lower than India's for every age group and yet end up showing a larger average.

This is one instance of the potential misleadingness of measures of central tendency when these are read without reference also to measures of dispersion. This fact has often been remarked, but is particularly well reflected in the pointed observation by Sherlock Holmes which is quoted at the beginning of this paper and serves to motivate a major strand of its concerns.

7. Conclusions

This paper has proposed a new graphical device, referred to as a mortality or *M*-curve, to describe and measure the combined influence of age-specific differences in mortality and the population age distribution. It can easily be constructed using information about age-specific mortality rates and the population age distribution (so long as the age grouping used is the same in each case). The need for such a curve arises partly because of the inadequacies of the crude death rate as an overall summary measure. The *M*-curve – a type of concentration curve – plots the proportion of total deaths against the corresponding proportion of people, where people are arranged from youngest to oldest. It has been shown that this device has a number of substantial advantages in addition to its ability to show 'at a glance' characteristics that are not evident from other profiles widely used to summarise mortality.

Taking inspiration from the income inequality literature on the famous Lorenz and Generalised Lorenz curves, an associated sub-optimality or inefficiency measure, I_M, was defined in terms of the area underneath the *M*-curve. This area can easily be computed using the well-known 'trapezoidal formula'. It is shown that this measure can be related directly to basic value judgements about the 'wastefulness' of early deaths and an explicit form of 'social welfare function', although in the present context the term 'loss function' is more appropriate. The loss function is expressed in terms of the crude death rate and (one plus) the inefficiency measure. The value judgement is simply that deaths occurring before some biological maximum are regarded as 'wasteful' of life-years, and hence the optimal outcome (assuming the absence of other characteristics of individuals) from the point of view of an independent judge is one in which no early deaths take place. The value judgement, by analogy with inequality aversion in the literature on income inequality, can be described as an 'aversion to young deaths'.

Establishment of a direct link from the value judgements to the inefficiency measure makes it possible to make 'welfare comparisons' between the mortality experience of different countries and time periods. A 'Generalised Mortality curve', or *GM*-curve, is defined in which the vertical axis of the *M*-curve is scaled by multiplying values by (1+CDR). A *GM*-curve that lies below that of another country over its whole length is then regarded as being preferred to the other. When such a 'dominance' result does not hold, explicit measures of the loss function can easily be used to make comparisons. Indeed, the value of the loss function is shown to be the area underneath the Generalised Mortality curve, and is obtained simply by multiplying the crude death rate by $1+I_M$.

The ability of the new curve to generate extra insights was illustrated by comparing *M*-curves and associated measures for India in 1971 and 2011, and New Zealand in 2019. A substantial improvement, in terms of both a reduction in I_M and the crude death rate, was found for India over the forty year period. The Indian *M*-curves were found to contrast with the New Zealand curve, which demonstrated considerably less 'inefficiency'. The results also highlighted the well-known inadequacy of the crude death rate in being influenced by the age distribution of the population. However, comparisons between countries on the basis of the overall loss function implied by basic value judgements can provide valuable information that does not rely on the use of an arbitrary 'standard' age distribution.

Appendix: Data Sources

India: Population Distribution Data, 1971: Derived from Statement 2.1C (Percentage Distribution in Five-Year Age-groups for Persons, Males and Females – India: 1961-2011), Registrar General & Census Commissioner, India: Census of India, 2011, Series-1: India, Report and Tables on Age (C-14, C-14SC & C-14ST), Volume-1. Available at:

http://lsi.gov.in:8081/jspui/bitstream/123456789/84/1/41036_2001_AGE.pdf

India: Deaths Distribution Data, 1971: Derived from Table 8 (Age-Specific Mortality Rate by Sex and Residence from 1971 to 1986 at Interval of 5 Years), Compendium of India's Fertility and Mortality Indicators, 1971-2013. Available at:

https://censusindia.gov.in/vital_statistics/Compendium/Srs_data.html

India: Population Distribution Data, 2011: Derived from Table 1 (Percent Distribution of Estimated Population by Age-Group, Sex and Residence, 2011) of Detailed Tables of SRS [Sample Registration System] Statistical Report, 2011. Available at:

https://censusindia.gov.in/vital_statistics/SRS_Report/12SRS%20Statistical%20Report%20Table%20-%2020111.pdf

India: Deaths Distribution Data, 2011: Derived from Table 8 **(**Table 8 Age-specific Death Rate by Sex and Residence, 2011) of Detailed Tables of SRS [Sample Registration System] Statistical Report, 2011. Available at:

https://censusindia.gov.in/vital_statistics/SRS_Report/12SRS%20Statistical%20Report%20Table%20-%2020111.pdf

New Zealand

For New Zealand it is necessary to go to each of the following web sites and use the on-line data selection facility to select and then download the required tables.

https://www.stats.govt.nz/topics/births-and-deaths

https://www.stats.govt.nz/topics/population

Table A1 presents the coordinates of the *M*-curves for India (1971), India (2011) and New Zealand (2019), derived from the data sources mentioned above. This might be of interest for those wishing to work further with the relevant distributional data.

Table A1 Coordinates of the M-Curves for India 1971 and 2011, and New Zealand 2019

Upper Limit of Age Group	India 1971 P_j	India 1971 Q_j	India 2011 P_j	India 2011 Q_j	New Zealand 2019 P_j	New Zealand 2019 Q_j
4	.1415	.4834	.0970	.1664	.0628	.0095
9	.2901	.5287	.1890	.1793	.1315	.0103
14	.4186	.5448	.2940	.1896	.1966	.0119
19	.5074	.5572	.3970	.2084	.2608	.0158
24	.5834	.5755	.4950	.2305	.3283	.0218
29	.6550	.5931	.5850	.2533	.4016	.0280
34	.7195	.6126	.6590	.2772	.4691	.0356
39	.7802	.6346	.7300	.3042	.5320	.0434
44	.8332	.6570	.7880	.3368	.5940	.0541
49	.8771	.6826	.8420	.3786	.6624	.0730
54	.9162	.7225	.8800	.4230	.7281	.1018
59	.9404	.7538	.9180	.4882	.7925	.1436
64	.9668	.8121	.9450	.5645	.8480	.1967
69	.9796	.8525	.9680	.6719	.8967	.2688
100*	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

Source: Data sources mentioned above.

Note: (1) P_j is the cumulative share in population, and Q_j the cumulative share in deaths, of those of age not exceeding the *j*th age-group's upper limit. (2) *100 years is a notional upper limit for the final, open-ended age group of 70-plus: the actual value is some number equalling or exceeding 70, and the precise value is irrelevant for generating the coordinates of the *M*-curve.

References

Atkinson, A.B. (1970) On the Measurement of Inequality. *Journal of Economic Theory,* 2(3): 244-263.

Basu, J. K. and T. Adair (2021) Have Inequalities in Completeness of Death Registration Between States in India Narrowed During Two Decades of Civil Registration System Strengthening? *International Journal of Equity in Health*, 20 (195). https://doi.org/10.1186/s12939-021-01534-y

Bishop, J.A., Chakraborti, S. and Thistle, P. (2009) An Asymptotically Distribution Free Test for Sen's Welfare Index, *Oxford Bulletin of Economics and Statistics*, 52(1): 105-113.

Creedy, J. and Gemmell, N. (2019) Illustrating Income Mobility: New Measures. *Oxford Economic Papers*, 71(3): 733-755.

Curtis, L.R. and Klein, R.J. (1995) Direct Standardization (Age-Adjusted Death Rates). *Healthy People 2000 Stat Notes*, March; (6): 1-10.

Foster, J.E. (1985) Inequality Measurement. In *Fair Allocation* (Ed. by H. Peyton Young), 38-61. Providence, Rhode Island: American Mathematical Society.

Gibson, J. (2020) Government Mandated Lockdowns Do Not Reduce Covid-19 Deaths: Implications for Evaluating the Stringent New Zealand Response, *University of Waikako Working Paper in Economics*, no. 6/20.

Jenkins, S.P. and Lambert, P.J. (1997) The Three "I"s of Poverty Curves, with an Analysis of U. K. Poverty Trends,' *Oxford Economic Papers*, 49(3): 317-327.

Lambert, P.J. (1993) *The Distribution and Redistribution of Income: A Mathematical Analysis*, Manchester: Manchester University Press.

Philip, M., Ray, D. and Subramanian, S. (2021) Decoding India's Low Covid-19 Case Fatality Rate. *Journal of Human Development and Capabilities*, 22(1): 27-51.

Ray, D. and Subramanian, S. (2020) India's Lockdown: An Interim Report. *Indian Economic Review*, 55: 31-79.

Sen, A. (1970) *On Economic Inequality*. Oxford: Clarendon Press.

Shorrocks, A. F. (1983) Ranking Income Distributions. *Economica*, 50: 3-17.

Subramanian, S. (2021) Age and Covid Fatality. *Development Studies Research*, 8(1): 236-243.

About the Author

John Creedy is Professor of Public Finance at Wellington School of Business and Government, Victoria University of Wellington, New Zealand.
Email: john.creedy@vuw.ac.nz

S. Subramanian is an Independent Scholar; formerly, Madras Institute of Development Studies.
Email: ssubramanianecon@gmail.com