

Australian Government Australian Institute of Health and Welfare

Authoritative information and statistics to promote better health and wellbeing

Principles on the use of direct age-standardisation in administrative data collections

For measuring the gap between Indigenous and non-Indigenous Australians

September 2011

Australian Institute of Health and Welfare Canberra Cat. no. CSI 12 The Australian Institute of Health and Welfare is a major national agency which provides reliable, regular and relevant information and statistics on Australia's health and welfare. The Institute's mission is authoritative information and statistics to promote better health and wellbeing.

© Australian Institute of Health and Welfare 2011 (cc) BY

This product, excluding the AIHW logo, Commonwealth Coat of Arms and any material owned by a third party or protected by a trademark, has been released under a Creative Commons BY 3.0 (CC-BY 3.0) licence. Excluded material owned by third parties may include, for example, design and layout, images obtained under licence from third parties and signatures. We have made all reasonable efforts to identify and label material owned by third parties.

You may distribute, remix and build upon this work. However, you must attribute the AIHW as the copyright holder of the work in compliance with our attribution policy available at </www.aihw.gov.au/copyright/>. The full terms and conditions of this licence are available at </htp://creativecommons.org/licenses/by/3.0/au/>.

Enquiries relating to copyright should be addressed to the Head of the Communications, Media and Marketing Unit, Australian Institute of Health and Welfare, GPO Box 570, Canberra ACT 2601.

A complete list of the Institute's publications is available from the Institute's website <www.aihw.gov.au>.

ISBN 978-1-74249-214-8

Suggested citation

Australian Institute of Health and Welfare 2011. Principles on the use of direct age-standardisation in administrative data collections: for measuring the gap between Indigenous and non-Indigenous Australians. Cat. no. CSI 12. Canberra: AIHW.

Australian Institute of Health and Welfare

Board Chair Dr Andrew Refshauge

Director David Kalisch

Any enquiries about or comments on this publication should be directed to: Communications, Media and Marketing Unit Australian Institute of Health and Welfare GPO Box 570 Canberra ACT 2601 Tel: (02) 6244 1032 Email: info@aihw.gov.au

Published by the Australian Institute of Health and Welfare

Please note that there is the potential for minor revisions of data in this report. Please check the online version at <www.aihw.gov.au> for any amendments.

Contents

Acknowledgments	iv
Background to the development of these principles	. v
Summary	vi
1 Introduction	1
2 What is age-standardisation?	3
2.1 What is an age-standardised rate?	. 3
2.2 Methods of age-standardisation	. 3
2.3 Advantages and disadvantages of direct and indirect age standardisation	. 5
2.4 Choice of standard population	. 7
2.5 The preferred method in analysis related to the Indigenous population	. 8
3 Principles on the use of direct-age standardisation	9
3.1 Guiding principles	. 9
3.2 Other issues that could impact on the reliability of age-standardised rates	21
4 Conclusion	23
References	24
List of tables	25
List of figures	25

Acknowledgments

This report was prepared by Tetteh Dugbaza, Michelle Gourley, Indrani Pieris-Caldwell, Bernadette Kok and Fadwa Al-Yaman. Special thanks are due to Allan Nicholls, AIHW Statistical Adviser, who reviewed the report and provided valuable comments.

Contributions from Krys Sadkowsky and M. Shahidullah from the Australian Bureau of Statistics (ABS) are gratefully acknowledged.

Special thanks are due to the Working Group, made up of ABS, AIHW, Office of Aboriginal and Torres Strait Islander Health (OATSIH), COAG Reform Council (CRC), Productivity Commission (PC) and the Department of Families, Housing, Community Services and Indigenous Affairs (FaHCSIA), for stimulating interest in this topic and commissioning this work.

Background to the development of these principles

This report has been prepared following a workshop hosted by the Australian Bureau of Statistics (ABS) on 19 April 2011, attended by representatives from the ABS, Australian Institute of Health and Welfare (AIHW), Office of Aboriginal and Torres Strait Islander Health (OATSIH) at the Department of Health and Ageing (DOHA), COAG Reform Council (CRC), Productivity Commission (PC) and the Department of Families, Housing, Community Services and Indigenous Affairs (FaHCSIA) to discuss the best method of agestandardisation (direct or indirect) to be used when reporting Indigenous data.

Workshop participants agreed that the direct method is the best option for Indigenous data analyses, especially when the objective is to compare mortality rates between population groups, that is, between Indigenous and non-Indigenous Australians, and when monitoring trends over time. This will enable the measurement of the gap between Indigenous and non-Indigenous or other Australians, which is of importance to COAG reporting.

A number of participants at the workshop expressed interest in having a set of clear guidelines to address a number of concerns of analysts when undertaking agestandardisation. As a result, it was agreed that a working group should be formed to investigate the following issues:

- How a consistent approach across Indigenous analyses can be used when there are small cells (collapsing cells; minimum number of cells containing numbers required for analyses etc.)
- How to define a 'small' cell
- Whether collapsing cells when small numbers will change the resulting standardised measure
- How to manage the change in standard population following the 2011 Census
- What contextual information should be provided to the CRC or in other reports (such as age specific rates, ratios, numbers) to ensure that important information is not lost by producing a one-number summary measure, or in the absence of a summary measure

A sub-group led by the AIHW, working closely with the ABS, developed a paper and guidelines including a set of principles for using the direct standardisation method.

This report is the result of this effort.

Even though the original intention of undertaking this work was to produce guidelines specifically for the analysis and reporting of COAG "Closing the Gap" indicators, these guidelines have been drafted to be general in nature, and can be applied when undertaking analysis of all age-related events and for comparing event rates between populations or sub-population groups over time.

Summary

Disease and mortality rates, as well as other health and welfare indicators, are often used to evaluate the performance of government and community programs aimed at improving the health and welfare of the population. This involves comparing event rates for different populations over time or examining the trends in a particular population over time.

Populations whose event rates (e.g. death, disease, hospital separation or birth rates) are being compared may often have different age structures. A given population may also have different age distributions over time as a result of changes in death rates, fertility rates and migration. Because these event rates vary by age, comparisons are often confounded by differences in the age distributions of the populations being compared. For example, comparison of indicators of wellbeing between the Aboriginal and Torres Strait Islander population and non-Indigenous Australians is hampered by the different age distributions of the two populations.

Age-standardisation is a technique used to enhance the comparability of event rates from different populations or different sub-populations over time by making adjustments for the confounding effects of differences in age structure between the populations being compared.

There are two main methods of age-standardisation, the direct and indirect methods. Whether to use the direct or indirect method depends on the purpose of the analysis, what type of comparison is being carried out, including whether trend analysis is being carried out, as well as on a number of data quality considerations. Among the data quality considerations are the following:

- the availability of accurate and reliable event rates, classified by age, in the study populations
- the overall number of events in each of the study populations
- the distribution of events by age in the study populations
- the size of the overall population at risk and in each age group
- the number of populated cells
- the consistency of the relationship between age and the event of interest in each of the study populations.

After exploring these issues, this report recommends that the direct method of agestandardisation be used for purposes of comparing health and welfare outcome measures (e.g. mortality rates, life expectancy, hospital separation rates, disease incidence rates etc.) of the Aboriginal and Torres Strait Islander population and non-Indigenous Australians. In particular, the direct method of age-standardisation should be used to compare the gap in wellbeing between Aboriginal and Torres Strait Islander and non-Indigenous Australians. The direct method of age-standardisation is also recommended for use in investigating changes in health and welfare outcome measures for the Aboriginal and Torres Strait Islander population over time.

The report provides guidance on when and how to use the direct age-standardisation method and under what circumstances it should not be used.

1 Introduction

Demographers, epidemiologists and statisticians are often faced with critical questions when studying disease and mortality rates. They often have to determine if disease or mortality rates are increasing or decreasing, and whether they are higher in one population than in another. Disease and mortality rates are also used to evaluate the success or otherwise of government and community health programs. The ability to compare pre-program and post-program event rates is therefore also very important.

To be able to answer these and other questions, disease and death rates have to be compared. The rates to be compared can be rates for the same population at different periods in time, rates for sub-populations within the same population, or rates for different populations.

For example, as part of the Council of Australian Governments' (COAG) initiatives to close the gap in mortality between Indigenous and non-Indigenous Australians, data are needed to carry out time series analyses of Indigenous and non-Indigenous mortality to determine whether Indigenous mortality is declining over time and whether the gap between Indigenous and non-Indigenous mortality is narrowing (COAG 2008).

A summary measure such as the crude rate may be used for this comparison because it has minimal data requirements, is simple to calculate, and has a very simple interpretation. There is, however, a problem with the crude rate when it is used to measure the force of events, such as disease and mortality rates, that are heavily influenced by age. The populations being compared may have different age structures. For example, the Australian Indigenous population has a relatively young age structure, while the non-Indigenous Australian population has a much older age profile. Even for the same population, the age distribution may change over time, although this may occur slowly.

The risk of illness and death varies with age, sex, socio-economic and other environmental factors. Thus, populations with a large proportion of young children or older persons are likely to experience a higher number of deaths than populations with much lower proportions of older persons or young children (AbouZahr et al. 2010). The age structure of a population will therefore affect any estimated crude rates, making it difficult to determine whether morbidity and mortality are actually higher in one population than the other, or whether the observed crude rates are the result of the different age structures of the two populations. Figures 1a and 1b illustrate this problem.



Figure 1a shows a population with a relatively older age structure while Figure 1b shows a population with a much younger age profile. The age-specific rates are higher at each age in the older age population than in the younger population, yet the crude rate is lower in the older age population (15.8 per 1,000) than in the younger population (22.0 per thousand). The crude rates appear to be inconsistent with the age-specific rates.

Crude rates are however not inaccurate, but they are often misleading when used to make comparisons, as they fail to take into account other variables, particularly age, that may have an influence on the event being measured (Rose & Barker 1978). One way of controlling for the often confounding effects of age, as observed in Figures 1a and 1b, is to adjust for the effects of age on the rates, through a process called age-adjustment or age-standardisation.

2 What is age-standardisation?

Age-standardisation is a technique used to enhance the comparability of data from different populations by making adjustments for the confounding effects of compositional differences in age structure between the populations or sub-populations being compared (Earyes 2008).

2.1 What is an age-standardised rate?

Age-standardised rates are hypothetical rates that would have been observed if the populations being studied had the same age distribution as the standard population, while all other factors remained unchanged.

Although a crude rate is a weighted summary measure of the age-specific rates, it is heavily confounded by the age composition of the population, as has been demonstrated in Figures 1a and 1b. This makes any observed differences in crude rates difficult to interpret, and therefore not a useful measure for comparing between or across populations. Age-specific rates provide an accurate comparison between populations as they control for the effect of age structure. However, their use is very cumbersome as it involves comparing a large number of rates. It is more useful to have a single summary measure, such as an age-adjusted rate, for each population, that takes into account any differences in the age structure of the populations. When comparing rates adjusted for age, any remaining observed differences between the populations cannot be attributed to confounding by age.

2.2 Methods of age-standardisation

There are two main approaches to age-standardisation: the direct and indirect methods. Both consist of taking a weighted average of the age-specific rates. The difference between the two lies in the source of the weights and the rates. In the direct method, standardised rates are derived by applying the age-specific rates observed in the study population to a single standard population, while in the indirect method the age-specific rates from a standard population are applied to the age distribution of the study population.

In the case of mortality ratios, while the interpretation of a single standardised ratio is simple and straightforward, a problem arises when trying to compare a number of standardised ratios from different study populations with each other. 'To make such a comparison validly, the direct method, in which a common standard is used for the two exposed groups, is preferred' (Hennekens & Buring 1987).

It should be noted that age-standardised rates from both the direct and indirect methods are indexed numbers, or artificial summary measures that are only meaningful when they are used in comparison (Choi et al. 1999). They have no intrinsic value.

2.2.1 Direct age-standardisation

A directly age-standardised rate is defined as the weighted average of event rates, with the weights being equal to the proportion of people in each age group in a chosen standard population. At the AIHW, the current standard population is the total estimated resident population in Australia on 30 June 2001.

The directly age-standardised event rate $(ASR_{(dir)})$ for the study populations, or populations being compared, is obtained by applying the event rates (r_i) (e.g. death rates) for each age group of the study population to the standard population sizes for each age group (N_i) :

$$ASR_{(dir)} = \sum N_i r_i / \sum N_i \tag{1}$$

Thus, $ASR_{(dir)}$ may be regarded as a weighted mean of the r_i using the N_i as weights. The age-standardised rate is usually expressed per 1,000 or 100,000 population.

A simple summary of the incidence or mortality rate ratios between the study populations and standard population that accounts for the possible confounding effects of age is the comparative mortality figure (CMF). It is obtained by dividing the directly age-standardized rate for the study population by the standard population rate. It may be interpreted as the ratio of the number of deaths that would be expected in the study population if it had the same age structure as the standard population, divided by the number of deaths in the standard population (Breslow & Day 1987).

The ratio of CMFs calculated for two different study populations using the same standard rates and weights is simply the ratio of the two directly standardized rates. The CMFs are directly comparable because they were estimated in reference to the same standard population. For example, if the ratio of the CMF for Population A to that of Population B is 1.5, then it means that mortality is 50 per cent higher in Population A than in Population B.

2.2.2 Indirect age-standardisation

An indirectly age standardised rate is defined as the weighted average of the age-specific rate ratios, where the weights are the expected number of events in each age group of the population in comparison. Indirect standardisation applies the age-specific rates from the standard population to the age distribution of the study population. The indirect method calculates how many events would be expected in each group if the age-specific rates of the standard population were applicable.

More frequently, the ratio of observed events to expected events is presented. For mortality data, this ratio is called the Standardised Mortality Ratio (SMR). If incidences are used instead of deaths, then the ratio is called the Standardised Incidence Ratio (SIR). The formula for calculating the standardised mortality or incidence ratio is:

$$SMR = d / \sum R_i n_i$$

(2)

Indirectly age-standardised rates (ASR_(ind)) for each population can be calculated by multiplying the SMR by the crude rate of the standard population as shown in equation 3:

$$ASR_{(ind)} = (d/\sum R_i n_i) * R$$

(3)

where:

d = the total number of events in the study population

 R_i = age-specific rate in age group *i* in the standard population

 n_i = the population in age group *i* in the study population

R = crude rate for the standard population (usually expressed per 1,000 or 100,000 population).

2.3 Advantages and disadvantages of direct and indirect age standardisation

Indirect standardisation is less commonly used than direct standardisation, and is more conveniently thought of as a comparison of observed and expected events than in terms of standardised rates.

Table 1 presents a brief discussion of the advantages and disadvantages of the direct and indirect methods of age-standardisation.

8 8	0					
Direct age-standardisation	Indirect age-standardisation					
Advantages	Advantages					
 It preserves the consistency between the populations in comparison, i.e. if each age-specific rate in Population A is greater than each of the corresponding age- specific rates in Population B, then the directly 	 It has minimal data requirements. When calculating indirect standardisation rates, the age-specific numbers of event cases are not required; only the total number of observed events is required. 					
standardised rate for Population A will always be higher	 Indirect age-standardisation is considered useful when: 					
is preferred for comparing different populations against	 the age-specific rates for the population being studied are not known but the total number of events is known 					
each other (Earyes 2008). (see illustration of problem in Figures 1a & 1b)	 calculating rates for small populations where fluctuations in age-specific rates can affect the reliability of rates calculated using the direct method 					
 When using the same standard population, the directly ago-standardised rates can be readily compared over 	 comparing observed and expected events 					
time.	 It is more stable as it minimises the variance, giving a 					
 Because directly standardised rates can be readily compared, they can also be ranked as they are based 	smaller standard error and narrower confidence intervals than the direct method (Earyes 2008).					
on the same population weights. Indirectly age-	Disadvantages					
standardised rates cannot be ranked because each rate is based on a different population weight.	 In most cases, the indirectly standardised rates will not be strictly comparable, in particular, when the age 					
 The direct method is considered the best method to 	structures of populations are different.					
 use: when making multiple comparisons (e.g. by sex, age and state of usual residence) 	 The Ratios (or SMRs) from study populations can only be legitimately compared with the standard and not with each other because different weighting is used to generate each Ratio (the weighting depend on the age 					
 when undertaking time series analyses 	distribution of the study population).					
 for practical reasons, such as, to maintain consistency throughout a report. Disadvantages 	 Cannot be used for time series analysis as age-specific event rates in the standard population will vary over time. 					
 It is sensitive to small cell sizes. This can occur in the case of rare events, events that occur mostly in some age groups and not in others or where the 	 In most cases, indirectly standardised rates cannot be ranked because they measure performance relative to the standard. 					
 breakdown of the population into sub-groups (e.g. by sex, age, state of usual residence) leads to very small populations and events in some sub-groups. The direct method requires that the number of events be available, and be broken down by age. This information is not always available, and even when it is available, it may not always be reliable. 	 Although in most cases the indirect standardisation method can produce more stable results, this is not always the case as illustrated in Table 2 when the age distributions of disease events are substantially different. A problem with indirect standardisation is that the ratio of two SMRs determined by pooling observed and expected deaths across age groups may sometimes lie completely outside the range of the age- specific rate ratios (see Table 2). 					

Table 1: Advantages and disadvantages of direct and indirect methods of age standardisation

There are many similarities as well as differences between the two methods. However, the two methods will yield comparable results in most cases. Indeed, it has been argued that the choice of a standard population is more important than the choice of the direct or indirect method (Breslow & Day 1987).

It has been demonstrated however, that there are conditions where the two methods could produce substantially different results. The first is when there are non-negligible differences in the age distributions of the study population(s) and the standard population. In this situation, indirect standardisation produces biased results due to residual confounding by age, but direct standardisation is not affected. The second is when the rate ratio of the study population(s) compared to the standard population varies substantially with age. The third is when both of these factors occur together (Silcock 1959).

Table 2 presents misleading indirect age-standardised results due to substantially different age distributions and standardised mortality ratios (SMRs).

Denulation	Deaths/SMDa	Age group		Total	
Population	Deaths/SMRS	20-44	45-64	20-64	
А	Observed deaths	100	1,600	1,700	
	% of deaths	5.9	94.1	100	
	Expected deaths	200	800	1,000	
	SMR1 (%)	50	200	170	
В	Observed deaths	80	180	260	
	% of deaths	30.8	69.2	100	
	Expected deaths	120	60	180	
	SMR2 (%)	67	300	144	
	SMR1/SMR2 (%)	75	67	118	

Table 2: Misleading indirect age-standardisation results

Source: Re-adapted from Breslow and Day (1987:73), originally from Kilpatrick (1963).

For the above data, Population A and Population B have different age distribution of deaths. Consequently, in calculating the SMRs for Population A, more weight was given to the age group 45-64, while for Population B, much more emphasis was given to the age group 20-44.

On an age-specific basis, there are apparent excess deaths in Population B, but when data are pooled, apparent excess deaths are found in Population A. The CMF is not subject to this problem as the ratio of two CMFs is the ratio of directly standardised rates (which are the weighted average of the age-specific rate ratios).

One key advantage of the direct method over the indirect method is that rates calculated by the direct age-standardisation method can be ranked. For example, it can be said that the incidence rate for disease X is five times as high among the Indigenous population as among the non-Indigenous population. If a ratio is required, the directly standardised rates can be compared to each other, while the indirectly standardised rates can only be compared to the standard population rate.

Indirectly age-standardised rates cannot be compared legitimately with each other because different weightings are used to generate each SMR, with the weights depending on the age distribution of the study population. This inability to compare the mortality experience of

study populations is the major disadvantage of using indirect standardisation, and provides the motivation for using direct standardisation.

2.4 Choice of standard population

One of the key components in calculating an age-standardised rate is the choice of a standard population. Direct age-standardisation accounts for differences in the age structures of the populations being compared (study populations), by weighting their respective age-specific rates to the age distribution of a standard population. Indirect age-standardisation accounts for differences in the age distribution of events in the study populations by applying the age-specific event rates in the standard population to the age distribution in each of the study populations.

The standard population must be chosen carefully as the age profile of the standard population can affect not only the value of the age-standardised rates but also the relative standing of the populations being compared. Choosing a standard population with higher proportions of people at older ages tends to weight events at these ages disproportionately; choosing a younger age standard does the opposite (Ahmad et al. 2009).

Ideally the standard population should reflect an age distribution not greatly different from that of the study populations. If age-specific rates in the study populations have a roughly consistent relationship, then the choice of standard population should not substantially affect comparisons, but if the age-specific rates are not consistent, as for example, if rates are declining by age in one population and increasing by age in the other, then comparisons will depend on the choice of standard population (Anderson & Rosenberg 1998; Breslow & Day 1987). When the study populations are small, it can be difficult to decide whether there is a consistent relationship in age-specific rates, due to the large random variation associated with small numbers.

There are two basic types of standard populations, that is, internal or external. Internal standard populations are the total pooled population of the study groups to be compared. Internal standards are commonly used, but a limitation is that rates standardised to a specific study population are not as readily compared to age-standardised rates from other studies. External standards are standard populations drawn from sources outside the analysis.

Choice of an external standard is arbitrary, depending on the purposes of the study, but conventions apply. In most countries, and for most international comparisons, there are official standard populations that must be used for all age-standardisation procedures to ensure that the resulting age-standardised rates are comparable.

For studies that have an international focus, a standard population that is commonly used is 2000 World Standard Population (Ahmad et al. 2009). In Australia, the convention followed by the AIHW and the ABS is to use the most recent census for a year ending in one (1) as the standard population. Currently, age-standardisation is done using 2001 Census data, and this will continue until data become available from the 2011 Census. An advantage of choosing a commonly used standard population is that it allows comparisons of age-standardised rates with other published studies.

2.5 The preferred method in analysis related to the Indigenous population

The direct method of age-standardisation has been the dominant method used in health research of national and international organisations, such as the US Centers for Disease Control and Prevention (CDC) (Curtin & Klein 1995) and the World Health Organisation (WHO) (Ahmad et al. 2009). It is also the method of choice for the AIHW and the ABS because of its advantages over the indirect method when comparing Indigenous and non-Indigenous mortality rates, disease incidence and prevalence rates over time.

While the direct method of standardisation has been accepted as the best method for Indigenous mortality data, a number of problems associated with this method, in particular when dealing with small cell sizes, have also been identified. For example, small cell sizes is a major drawback of the direct method and while this can be dealt with by grouping ages into larger age categories, it is not always clear what an acceptable cell size should be. A rough guideline used by the CDC is that there should be at least 25 events over all age groups (Curtin & Klein 1995).

This report presents a set of guidelines that analysts can follow to address these issues.

3 Principles on the use of direct-age standardisation

There are a number of theoretical and practical concerns that should be taken into account when deciding to use the direct method to standardise event rates for purposes of comparing Indigenous and non-Indigenous measures over time. These considerations include:

- the age distributions of the study populations or the populations whose event rates are being compared, and whether the age-standardised rates will be different from their respective crude rates
- whether there are enough events to make the age-standardised rates meaningful and reliable
- the choice of a standard population, and how to manage the change in standard population following the 2011 Census
- how to define a 'small' cell in both the numerator and denominator, and how to adopt a consistent approach across Indigenous analyses when there are small cells
- what age groupings as well as lower and upper age-cut-offs are appropriate for calculating direct age-standardised rates
- the age distribution of the events of interest, that is, whether the events of interest are concentrated in only a few or particular age groups or whether they are spread throughout the age distribution of the population
- what contextual information should be provided (such as age-specific rates, ratios, numbers) to ensure that important information is not lost if there is no summary measure or if only a one-number summary measure is produced.

These considerations have been taken into account in developing a set of principles to guide the use of direct age-standardisation for preparing Indigenous morbidity and mortality rates over time for purposes of COAG 'Closing the Gap' reporting. However, they have wider application, and can also be used for comparative analysis, e.g. mortality analysis, over time involving multiple populations or sub-populations.

For developing the guidelines in this paper, the effect of cell size on the directly standardised rate was tested using different number of counts to determine the point at which the rate was stabilised.

3.1 Guiding principles

The guiding principles to be followed when undertaking the direct method of agestandardisation are summarised in Box 1. More detailed discussion on the analyses which support these principles follows.

Box 1: Guiding principles on the use of direct age-standardisation

Overarching principle: Before undertaking age-standardisation, analysts must investigate the data being used to understand the age-specific distributions of the populations being compared, and any limitations that may impact on the results.

Principle 1: The standard population used should be the Australian Estimated Resident Population as at 30 June 2001 from the 2001 Census until population estimates from the 2011 Census become available.

The population used as the denominator for the calculation of Indigenous age-standardised rates should be SERIES B of *Indigenous experimental estimates and projections 2006 to 2021* based on the 2006 Census until population estimates from the 2011 Census become available.

Principle 2: If the denominator is less than 30 in any one age group, then do not attempt to produce age-standardised rates.

Age groups may be collapsed to obtain a denominator of 30 or more (provided that this is in accordance with principles 3 and 4)

Principle 3: If the total number of events in one population, e.g. deaths in the Indigenous population, is less than 20, then do not attempt to produce age-standardised rates.

Combining several years of data, or aggregating jurisdictions should be considered to obtain a total of 20 or more events.

If this does not meet the purpose (i.e. data are required for time series or jurisdictional comparisons), or does not result in greater than 20 events, then other measures and contextual information should be reported instead of age-standardised rates which could include total number of events, crude rates, age-specific rates, age-specific rate ratios and median age at the time of the event.

Principle 4: Age-standardised rates should be calculated using the five year age groupings of 0-4 to 75+ (provided Principles 2 and 3 for denominator and numerator are met).

10 year age groups may be used to overcome small numbers (20 year age groups are too wide and should not be used)

Principle 5: Contextual information (most importantly age-specific rates and ratios) should be provided <u>in addition to</u> age-standardised rates when:

a) the age-standardised rates and rate ratios lie largely outside the range of the age-specific rates and rate ratios

b) the pattern of age-specific rates of the Indigenous and non-Indigenous populations differ substantially (e.g. deaths from a certain cause concentrate on younger ages for Indigenous population while for non-Indigenous they may occur at older ages)

c) age-specific death rates depart from the assumption of a uniform increase in death with age (e.g. injury which peaks in the young adult to middle ages and certain cancers amenable to treatment for some age groups).

Principle 6: For conditions restricted to a specific age group (e.g. conditions originating in the perinatal period, Sudden Infant Death Syndrome), it is recommended to report the age-specific rate for the age group of interest *instead of* the age-standardised rate.

Notes regarding application of these principles

Analysts should apply these principles keeping in mind that a consistent time series will need to be maintained and valid comparisons ensured. The guidelines should not result in a widely different methodology being used for data being compared across time or across jurisdictions. Analysts should make decisions regarding the application of these principles based on an assessment of the data from the outset of their analyses.

These principles may need to be reviewed at some point in the future in order to take into account any changes, for example, in Indigenous mortality and population estimates. A sensible time for such a review would be to coincide with a change of standard population (i.e. every 10 years).

The impact of a change in standard population on the resulting age-standardised rates will need to be assessed by the ABS and AIHW as soon as practical following the release of the 2011 Census population estimates.

Overarching Principle: Investigate data being used

Before undertaking age-standardisation, analysts should investigate the data being used, and understand the age-specific distributions of the populations being compared, and any limitations that may impact on the results.

The purpose of the study and whether age-standardised rates are the best measures to meet the objectives of the study are also important considerations.

There are many data issues that should inform the analyst as to whether or not to undertake age-standardisation, how to undertake the age standardisation and how reliable the estimates will be. These issues can be grouped into two broad categories relating to data quality and the structure of the data. These considerations are discussed further in this report.

Quality issues relate to data accuracy, completeness, and missing cells or "age not stated" cases. Poor quality data can yield misleading results and lead to poor policy decisions. Analysts need to decide if the quality of both the reported events and the population at risk are sufficient for the estimation of age-standardised rates. This includes whether there is a sufficient number of events to calculate age-standardised rates or whether other measures (e.g. age-specific rates and median age at death) and contextual information should be provided. The total number of events will affect the standard error of the estimates and the reliability of any comparisons based on the estimated age-standardised rates.

The size of any "age not stated" events should also be examined to determine whether they should be distributed pro-rata across the age groups or whether they should be ignored for the purposes of estimating age-standardised rates.

Other issues that could impact on the reliability and interpretation of any estimated agestandardised measures include the age distribution of the events of interest and the population at risk, as well as the consistency in the age-specific rates between the study populations.

In some cases, age-standardised rates may be unnecessary if the events of interest are restricted to only one or a couple of age groups (e.g. conditions originating in the perinatal period). In such a situation it may be better to report age-specific rates rather than age-standardised rates. There could also be other instances where it might be necessary to estimate age-standardised rates over a restricted age band, such as 15-64 years (e.g. for work-related injury to working age persons) or females in the childbearing ages (15-49 years) or young adults (age 15-34 years), if the events of interest are restricted to only particular age groups.

Age-standardised rates may also be unnecessary if the populations being compared have the same age structure. If the populations have the same age distribution, then age adjustment will not produce rates that are different from the crude death.

If one is comparing populations within a very narrow age range (e.g. 0-4 or 18-24 years), then it may not be necessary or meaningful to undertake age-standardisation.

Principle 1: Standard population

Use the Australian Estimated Resident Population, at 30 June 2001 from the 2001 Census, as the standard population, until population estimates from the 2011 Census become available.

The choice of standard population used in age-standardisation can have an impact on the resulting age-standardised rates and comparisons, particularly if the age-specific rates in the study populations have an inconsistent relationship with age.

The age structure of the standard population can affect not only the value of the agestandardised rates but also the relative standing of the populations being compared. The general consensus in the scientific literature is that unless the age-specific rates in the study populations have an inconsistent relationship with age, selection of the standard population is less likely to affect relative comparisons (e.g. rate ratios and rate differences), although it will affect the absolute values of the standardised rates. Ideally the standard population selected should reflect a distribution not greatly different from that of the study populations.

In Australia, the convention followed by the AIHW and ABS is to use the age distribution of the Australian population from a Census year that ends in '1' as the standard population which changes every 10 years. The current standard population used is the Estimated Resident Population (ERP) of Australia at 30 June 2001 from the 2001 Census. This will change to the Australian ERP at 30 June 2011 when data become available from the 2011 Census. An advantage of choosing a commonly used standard population is that it allows comparisons of age-standardised rates with other published studies.

The impact of a change in standard population on the resulting age-standardised rates will need to be carefully assessed when data from the 2011 Census become available (expected June 2013). It is anticipated that the change in standard population will have a significant impact on the age-standardised rates, particularly for Indigenous Australians, and, to a much lesser extent, on the rate ratios. This conclusion is based on previous experience, and supported by an examination of population data by the ABS, which found that the change in standard population from the Australian ERP as at 30 June 1991 to 30 June 2001 resulted in a significant increase in the age-standardised rates, and a small decrease in the rate ratios. Based on this, it is likely that all previous tables produced using the direct method of age-standardisation will need to be recalculated using the new standard population. As different relativities are likely to emerge, some description as to the way age standardisation works on populations with different age structures may need to be provided.

Recommendations

- 1. The Estimated Resident Population of Australia as at 30 June 2001 from the 2001 Census should continue to be used until population estimates from the 2011 Census become available.
- 2. When reporting age-adjusted rates, the standard population used should also be reported. It is important to document that the age-adjusted rates being compared have been adjusted to the same standard population.

Principle 2: Do not produce age-standardised rates if the denominator is less than 30 in any one age group

It is well established that as the denominator increases, the variance of the age-standardised rate decreases. A simulation study reported by the AIHW's National Injury Surveillance Unit has observed that the results from the direct standardisation method become unstable when the denominator is smaller than 30 (Berry & Harrison 2005:18). In this study, the precision of the directly age-standardised rate ratio was explored by modelling various combinations of the numerator and denominator for the 85+ age group to see what denominator size achieves only a minor change in the rate ratio. The authors tested the stability of the rate ratios using numerators ranging from zero to five and denominators ranging from zero to 75. The study found that when the denominator is around 30 or greater in each age group, the directly age standardised ratio is reasonably stable.

The AIHW has examined the size of the Indigenous population by jurisdiction, sex, year and five-year age group. These analyses show that of the five jurisdictions with adequate identification of Indigenous deaths (NSW, Qld, WA, SA and NT), there are no denominators which are less than 30 in a single age-group with the exception of the age group 85+ years for Indigenous males in South Australia (which is between 20 and 30).

Recommendations

- 1. Age-standardised rates should not be reported if the denominator is less than 30 in any one age group. Collapsing of age groups should be considered to obtain a population of 30, provided this is consistent with Principle 3 (the total number of deaths is 20 or more) and conclusions reached under Principle 4 (i.e. that 20 year age groupings appear too wide, and the upper age cut-off should be at least 75+).
- 2. Collapsing of cells to obtain a population of 30 or more need not be restricted to age groups only. Depending on the reporting objectives or purpose of the analysis, aggregations of jurisdictions or years of data could also yield a sufficiently large population for reporting of age-standardised rates.

Principle 3: Do not produce age-standardised rates if the total number of deaths is less than 20

Where possible, combine several years of data or aggregate jurisdictions to obtain 20 or more deaths. Where this is not possible, provide other measures and contextual information instead of age-standardised rates.

Age-standardised rates based on only a small number of events will be unstable and exhibit a large amount of random variation. A review of national and international literature suggests that directly age-standardised rates start to become unstable if there are less than 20 to 30 cases in the numerator. For example, the US National Center for Health Statistics recommends that there should be at least 20-25 total deaths over all age groups before attempting to calculate directly age-standardised death rates. Where fewer than 20 health events occur over a time period, they suggest considering combining years, or using indirect age-adjustment (Curtin & Klein 1995). This guideline was derived from an assessment of the Relative Standard Error (RSE), which relates to the stability of estimates to the number of cases. An RSE of greater than 23% roughly equates to fewer than 20 events (New York State Department of Health 2006; Klein et al. 2002). Analyses by the AIHW support literature findings that age-standardised death rates based on less than 20 deaths have high relative standard errors (25% or more), and consequently, wide confidence intervals (as large as 50% or more of the standardised rate)(see Figures 2a and 2b).



While the age-standardised rates themselves do not change significantly based on less or more deaths in total (see Figure 2b), the confidence limits depart substantially from the age-standardised rates when the number of cases is less than 20.

Recommendations

- 1. If the total number of events (e.g. deaths) is less than 20, then age-standardised rates should not be estimated or reported. Where possible, data pertaining to individual years or jurisdictions could be aggregated to obtain a total of 20 or more events.
- 2. If recommendation 1 does not meet the purpose of the estimates (i.e. if estimates are required for time series or jurisdictional comparisons), then it is recommended that additional contextual information be reported. This could include total number of deaths, crude rates, age-specific rates, age-specific rate ratios and median age at death.
 - It is important to note that the total number of deaths reported for certain conditions and in certain jurisdictions will change over time and may fall above or below the cut-off point of 20 deaths over a period of time. In order to maintain a consistent time series, the causes of death and jurisdictions for which age-standardised mortality rates are to be reported should be determined at the outset based on an examination of the data.

Principle 4: Use five-year age groups

Directly age-standardised rates should be calculated using five- year age groups from 0-4 to 75 years+.

Where numerator and denominator data are available for these age groups, ten-year age groups may also be used to overcome small numbers.

Having a greater number of age groups results in better control of the effect of any differences in the age distribution among the groups or time periods being compared, thus increasing the accuracy of the standardisation. However, with more age groups the data become sparse and there may not be enough events to populate all the strata, resulting in standardised rates with large variances.

The literature cautions that using broad age groups will produce a less precise adjustment and affect the comparability of age-specific death rates to other populations. A broad age group (e.g.15-44) will not be sensitive to any changes in age-specific rates within this category. The grouping of age groups with very different rates will result in a single rate that is insensitive to any changes in the age-specific rates within the group.

For Indigenous mortality analyses, for instance, five-year age bands, starting with the age group 0-4 years and finishing with the age group 75 years and over, may be appropriate. Where data are available, or where the event of interest varies substantially between infants (under 1 year) and children (1-4 years), it may be appropriate to break down the 0-4 age group into 'under 1'and '1-4' age groups.

For analyses of the Indigenous population, the age group '75 years and over' is commonly chosen as the oldest age group, instead of '85 years and over' or '100 years and over' which has recently become the WHO standard. This is due to the relatively small Indigenous population and deaths in these older age groups in certain states and territories. In time however, there could be changes in the size of the Indigenous population at older ages which would allow for an increase in the upper age cut-off to 85 years and over.

Effect of aggregating age-groups on resulting directly age-standardised rates

Rates based on 5, 10 and 20 year age groups

Examination of the effect of aggregating mortality data into ten and 20 year age-groups on the resulting age-standardised rates, rate ratios and rate differences suggests that there is very little difference between age-standardised rates using five-year or ten-year age groups. However, some differences are observed when 20-year age groups are used, particularly for certain causes of death. This can be seen in Figures 3 and 4 which present the rate ratios and rate differences for five jurisdictions combined.

This analysis suggests that 20-year age bands are too wide and conceal some of the specificities in death rates for particular causes of death at particular ages. A similar pattern was observed when estimates were based on individual jurisdictions and single years of data (e.g. 2007), for which smaller numbers of deaths are involved.





Rates based on upper age cut-offs of 55+, 65+, 75+ and 85+

Low upper age cut-offs (55 years+ or 65 years+) may be useful if there are not enough events at older ages to enable reliable age-specific rates to be estimated. However, low upper age cut-offs may distort the allocation of events to populations at risk as well as distort the age-standardised event rates if the pattern of events varies considerably at older age groups.

The effect of having different upper age cut-offs (55+, 65+, 75+, 85+) on age-standardised rate ratios and rate differences has been assessed (Figures 5a and 5b). The analysis shows that using low upper age cut-off of 55+ or 65+ has some impact on the resulting age-standardised rate ratios and rate differences (they are generally lower), however there is very little difference between rate ratios and rate differences based on upper age cut-offs of 75+ and those based on upper age cut-offs of 85+. A similar pattern is observed when estimates are based on different age bands (e.g. 10 and 20 year age-groups) and single year data (e.g. 2007) rather than aggregated years' data (e.g. 2003-2007).



Rates based on lowest age group of <1, 1-4 or 0-4 years

Infants (less than 1 year) and 1-4 year olds have quite different rates of mortality for most causes of death. The grouping together of age groups with very different demographic and health characteristics and event rates can result in a single rate that is insensitive to any changes in the age-specific rates within the group. In light of this, the AIHW has also looked at whether splitting the 0-4 age group into two age groups of 'less than 1 (<1)' and '1-4 years' has an effect on the resulting age-standardised rates, rate ratios and rate differences.

It can be seen from Figures 6 and 7 below that there is little difference in the rate ratios and rate differences when the age group 0-4 is used, compared to when it is split into <1 and 1-4 years, and used in the calculation of the age-standardised mortality rates. The exception is conditions originating in the perinatal period which are almost exclusively confined to infants (<1), with only a few deaths from this cause occurring in the age-group 1-4 years or the remaining age-groups (see principle 6 for recommendations on the reporting of conditions largely confined to a specific age-group).





Recommendations

- 1. 20 year age-groupings appear too wide and should not be used for the calculation of age-standardised mortality rates.
- 2. As there is little difference in the resulting rate ratios and rate differences using five or ten year age-groupings, it is recommended to follow the usual convention of using five year age-groupings in the calculation of directly age-standardised mortality rates. However if the distribution of the data across age-groups warrants collapsing of age-groups to overcome small numbers, then 10 year age-groupings may be used.
- 3. The upper age cut-off should be at least 75 years and over.
 - a. For estimates relating to the Indigenous population, it must be noted that Indigenous population estimates and projections from the 2006 Census are only available for age-groups of up to 65 years and over for Tasmania and ACT. This means that for these jurisdictions, upper age cut-offs for estimating age-standardised event rates will be 65+, and not 75+ as recommended above.
 - b. As some jurisdictions have very small Indigenous populations aged 85 years and over, it is recommended to follow the usual practice of using 75 years and over as the oldest age group in the calculation of age-standardised rates.
- 4. As there is little or no difference in rate ratios and rate differences produced using 0-4 compared to using <1 and 1-4 age groups in the estimation of age-standardised rates, it is recommended to follow the usual practice of using the 0-4 age group as the youngest age group in the calculation of age-standardised rates. This only applies to the calculation of age-standardised rates, and should not preclude presenting age-specific rates and distribution of events (e.g. deaths) for <1 and 1-4 age groups).

Principle 5: Consider whether additional contextual information is needed

Additional contextual information should be provided in addition to age-standardised rates when:

- the age standardised rates and rate ratios lie largely outside the range of the age-specific rates and rate ratios;
- the pattern of the age specific rates of the Indigenous and non-Indigenous populations differ substantially; and
- the age-specific rates depart from the assumption of a uniform increase in death with age.

a) the age standardised rates and rate ratios lie largely outside the range of the age-specific rates and rate ratios

The age-standardised rate and ratio may not always provide an accurate picture of the differences in the age-specific rates between the populations being compared. This may be particularly the case where the age-standardised rates and rate ratios lie largely outside the range of the age-specific rates and rate ratios.

An example of where the age-standardised rates can lie outside the range of most of the agespecific rates is notification rates for Chlamydia, shown in Table 3. In the table, the directly age-standardised rate ratio for Indigenous females compared to non-Indigenous females is 3.8, which is lower than the age-specific rate ratios (which lie between 5.3 and 69.8) with the exception of the 15-19 and 20-24 years age groups (around 3). The age-specific rate ratios also show that there are age-groups which are of particular interest in the Indigenous population as compared to the non-Indigenous population, for example 10–14 year-old females.

	0-4	5-9	10- 14	15- 19	20- 24	25- 29	30- 34	35- 39	40- 44	45- 49	50- 54	55- 59	60- 64	65- 69	70- 74	75+	Total
Indigenous	1.1	0.8	11.5	78.5	57.9	37.3	21.7	15.6	8.2	5.3	3.8	1.0	1.2	1.8	2.6	0.0	17.3
Non-Indig.	0.1	0.1	0.2	26.2	22.2	7.3	3.7	1.3	1.4	0.2	0.4	0.1	0.2	0.3	0.0	0.0	4.6
Ratio	16.5	5.3	69.8	3.0	2.6	5.1	5.9	12.1	6.0	24.8	9.6	9.0	6.6	5.7	-	-	3.8

Table 3: Age-specific notification rates for Chlamydia, by Indigenous status, females

Recommendation

Where the age-standardised rates and rate ratios lie largely outside the range of the agespecific rates and rate ratios, it is more informative to report the age-specific rates of the age groups of interest (which may have important policy implications), rather than attempting to convey the differences through one figure only.

b) the pattern of the age specific rates of the Indigenous and non-Indigenous populations differ substantially

The age-standardised rate can mask important differences in the pattern of age-specific rates for the Indigenous and non-Indigenous populations.

An example of where the pattern of age-specific rates is quite different for the Indigenous and non-Indigenous populations is external causes of death, including accidents, intentional self-harm and assault. As can be seen in Figure 8, mortality rates in the non-Indigenous population are fairly similar for age groups 15-19 to 70-74 years. However for the Indigenous population, a different pattern is evident, with rates steadily increasing to peak in the 35-39 age-group, decreasing until the 55-59 age-group and then increasing again in the older age-groups.



Recommendation

For conditions where the pattern of the age-specific rates of the Indigenous and non-Indigenous populations differ substantially, it is recommended to present the age-specific rates (e.g. in graph format) alongside the age-standardised rates.

c) the age-specific death rates depart from the assumption of a uniform increase in death with age

The age-standardised rate can mask trends that occur when the event of interest does not have a consistent relationship with age, exhibits drastically different patterns for different age groups over time, or changes in opposite directions over time for different age groups. The event of interest, such as overall mortality or deaths due to particular causes, may decline in some age groups, remain unchanged in others or even increase in other age groups. Consequently, there will be departures from the assumption of a uniform increase in death with age.

For example, in Figure 9 below, the death rate from all cancers for females aged 85 years and over has slowly increased, while the corresponding rate for females aged 0-4 has decreased substantially during the past 30 years. Simply reporting that the age-standardised rate for female cancer has remained constant at around 150 deaths per 100,000 persons over the 1977-2007 period will not identify that considerable improvement has been made in the survival of female infants (age 0-4) with cancer, particularly in the last 30 years.

Further, the age-standardised death rates do not show that for older age females (aged 85 years and over), the rate of death from cancer increased during the same period. As Figure 9 shows, such inconsistencies in the age-specific event rates over time will lead to a situation where the age-standardised death rates do not reflect the underlying age-specific death rates.



Recommendation

When reporting age-standardised rates it is essential to identify if and where age specific rates depart from the assumptions of consistency across age groups, and if they show different trends over time. The age ranges and results where the departures occur should be reported in addition to the age-standardised rate.

Principle 6: For conditions restricted to a specific age group, report only the age-specific rates for that age group

Some conditions are largely restricted to a specific age-group or age-range. Examples of this include conditions originating in the perinatal period and Sudden Infant Death Syndrome (SIDS), which are largely confined to infants; and cancers such as prostate cancer which is largely confined to the older age groups (e.g. 65 years and over).

Where the condition is restricted to only a specific age group or couple of age groups, then presenting an age-standardised rate is not particularly meaningful. For such conditions, it is preferable to report only the age-specific rate rather than an age-standardised rate.

Recommendation

For conditions restricted to a specific age group or couple of age groups, report only agespecific rates rather than the age-standardised rate.

3.2 Other issues that could impact on the reliability of age-standardised rates

Age-adjusted event rates are subject to various sources of error that could impact on the reliability and interpretation of the rates. These sources of error include random variation in the number of reported events as well as errors associated with the age distribution of the

reported events. As a result, it is useful to examine the standard errors and confidence intervals of the estimated age-standardised event rates before they are used or reported. It should be noted that where there are a number of cells with zero events, this will result in an underestimation of the true standard errors.

4 Conclusion

The decision to age-standardise mortality or any events data, as well as the choice of method, should be guided by a number of considerations, including the purpose of the analysis, the total number and age distribution of the events of interest, the age profile of the study populations, and how the analyst intends to use the resulting age-standardised rates.

Age-standardisation is recommended where the study populations have different age distributions, and using the crude rate to describe their mortality experiences might be misleading. The direct method of age-standardisation is recommended where the analyst intends to compare two or more populations over a period of time. For this comparison, it is important to use the same standard population, which in Australia, is the estimated resident population (ERP) of Australia as at 30 June 2001. This population is changed every 10 years.

Direct age-standardisation is recommended for reporting of COAG 'Closing the Gap' targets, because the aim of such reporting is to compare the gap in wellbeing between Indigenous and non-Indigenous Australians, and to determine trends in Indigenous wellbeing over time. Direct age-standardisation is also recommended for other national reporting where indicators of advantage and disadvantage are compared over time and between populations.

Age-standardised rates are relative indices, and should be used only for purposes of comparison. 'They are not actual measures of mortality risk, and they do not convey the magnitude of the problem' (New Mexico Department of Health 2010). It is important that analysts are informed by the principles guiding the use of the direct method of age-standardisation in order to avoid misapplication of the method and misinterpretation of the results of the analysis.

References

AbouZahr C, Mikkelsen L, Rampatige R & Lopez AD 2010. Mortality Statistics: a tool to enhance understanding and improve quality. Health Information Systems Knowledge Hub, Working Paper Series No. 13. School of Population Health, University of Queensland.

Ahmad OB, Boschi-pinto C, Lopez AD, Lozano R & Inoue M 2009. Age standardization of rates: a new WHO standard. Geneva: World Health Organisation.

AIHW 2010. GRIM (General Record of Incidence of Mortality) books. Canberra: Australian Institute of Health and Welfare.

Anderson RN & Rosenberg HM 1998. Age standardization of death rates: implementation of the year 2000 standard. National vital statistics reports 47:1-16.

Berry JG & Harrison JE 2005. A guide to statistical methods for injury surveillance. Canberra: Australian Institute of Health and Welfare.

Breslow N & Day N 1987. Statistical methods in cancer research. Volume II-The design and analysis of cohort studies. IARC Scientific Publication No 82. Lyon, France: International Agency for Research on Cancer.

Choi BCK, De Guia NA & Walsh P 1999. Look before you leap: stratify before you standardize. American journal of epidemiology 149:1087.

COAG 2008. National Indigenous Reform Agreement (Closing the Gap), Schedule F. Council of Australian Governments.

Curtin LR & Klein RJ 1995. Direct standardization (age-adjusted death rates), Statistical Notes, No. 6. March 1995. Hyattsville, Maryland: US Dept. of Health and Human Services, Public Health Service, Centers for Disease Control and Prevention, National Center for Health Statistics.

Earyes D 2008. Technical Briefing 3: Commonly used public health statistics and their confidence intervals [online]. Association of Public Health Observatories. York, UK.

Hennekens C & Buring J 1987. Epidemiology in Medicine. 37-45. Little, Brown & Co. Boston.

Kilpatrick S 1963. Mortality comparisons in socio-economic groups. Journal of the Royal Statistical Society. Series C (Applied Statistics) 12:65-86.

Klein RJ, Statistics NCfH, Control CfD & Prevention 2002. Healthy People 2010 criteria for data suppression. US Dept of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics.

New Mexico Department of Health 2010. Age-adjusted rates. 22 January 2010. Santa Fe: Indicator-based Information System for Public Health. Viewed 18 July 2011, <http://ibis.health.state.nm.us/resources/AARate.html>.

New York State Department of Health 2006. About Age Adjusted Rates, 95% Confidence Intervals and Unstable Rates. March 2006. New York: New York State Department of Health. Viewed 22 May 2011, http://www.health.state.ny.us/statistics/cancer/registry/age.htm.

Rose G & Barker D 1978. Epidemiology for the uninitiated: Comparing rates. The British Medical Journal 2:1282-3.

Silcock H 1959. The comparison of occupational mortality rates. Population Studies 13:183-92.

List of tables

Table 1:	Advantages and disadvantages of direct and indirect methods of age	
	standardisation	5
Table 2:	Misleading indirect age-standardisation results	6
Table 3:	Age-specific notification rates for Chlamydia, by Indigenous status,	
	females	19

List of figures

Figure 1:	The relationship between age and incidence rates of illness and death	1
Figure 2:	Relationship between number of cases, relative standard errors and	
C	confidence intervals	14
Figure 3:	Rate ratios, NSW, Qld, WA, SA NT Males, 2003-2007	16
Figure 4:	Rate differences, NSW, Qld, WA, SA NT Males, 2003-2007	16
Figure 5:	Cause-specific rate ratios and rate differences for various upper age cut- offs	17
Figure 6:	Cause-specific rate ratios: comparing 0-4 and under 1 and 1-4 age split, 2003-2007	17
Figure 7:	Cause-specific rate differences: comparing 0-4 and under 1 and 1-4 age split, 2003-2007	18
Figure 8:	Age-specific death rates for external causes	20
Figure 9:	Comparing female cancer deaths 0-4 and 85+ to age-standardised rate	21