Part III - Measures of Health Disparities

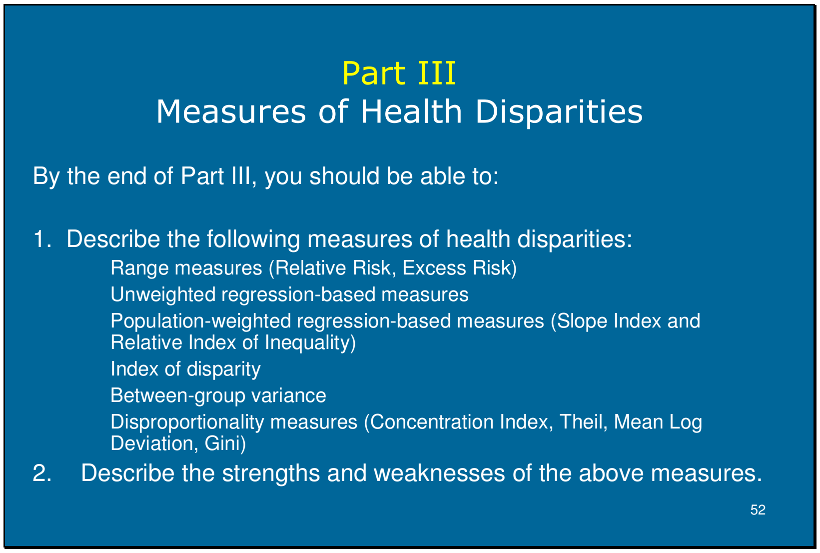

In Part III we review the most commonly used measures of health disparity.

By the end of Part III, you should be able to:

1. Describe the following measures of health disparities:

    Range measures (Relative Risk, Excess Risk)

    Un-weighted regression-based measures

    Population-weighted regression-based measures (Slope Index and Relative Index of Inequality)

    Index of disparity

    Between-group variance

    and Disproportionality measures (Concentration Index, Theil, Mean Log Deviation, Gini), and

2. Describe the strengths and weaknesses of the above measures.

Measures

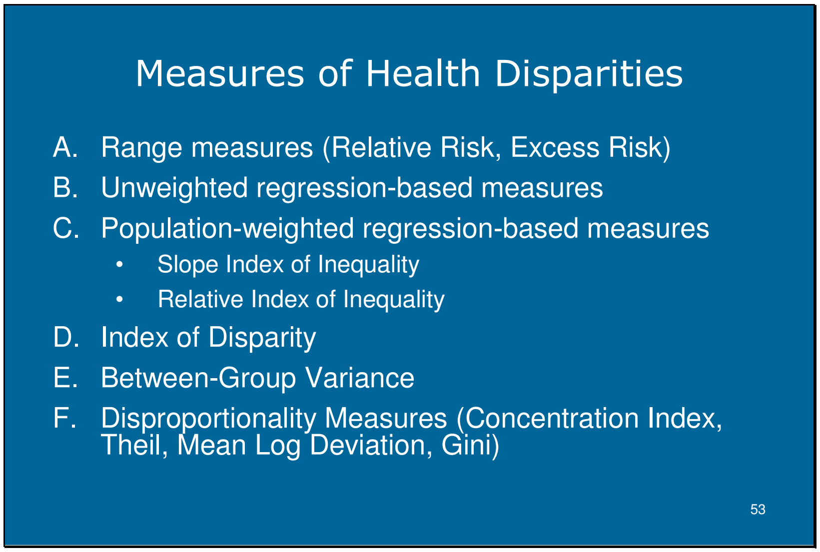

Part III will give you an idea of the general characteristics of each of these measures. For those of you who want more technical detail and a better understanding of how these measures are used in research and practice, we have provided references in the *Resources* section to key articles from the health disparity literature. When possible, we have also provided the text of the articles in a pdf file.

We will begin with the simplest measures:
Range measures
Un-weighted regression-based measures
Population-weighted regression-based measures
For many purposes, these will be all you need.

There may be situations, however, where you want to summarize disparities over time or across different groups, which can get technically more complicated. An overview of the following measures will provide you with a taste of what goes into these more complex calculations:
Index of disparity

Between-group variance

Disproportionality measures

Range measures

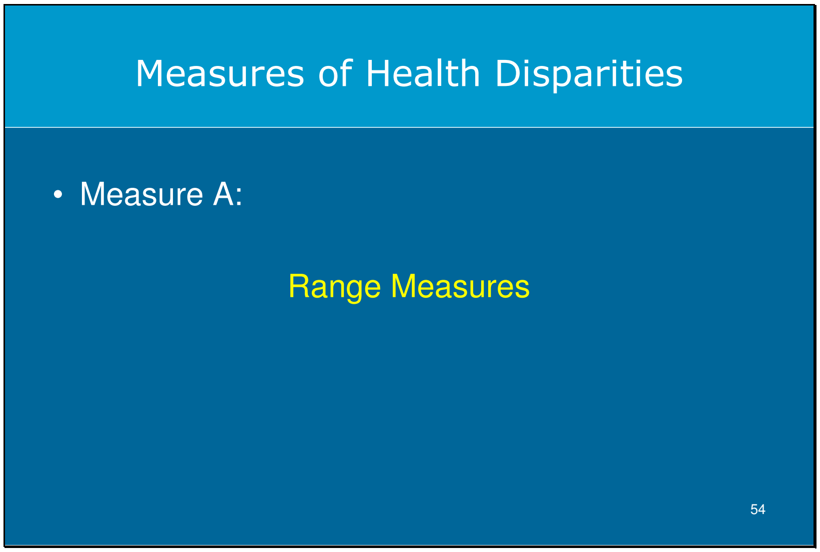

Range Measures typically compare two extreme categories.

Range measures

Range Measure Examples:
Relative Risk (RR), Excess Risk (ER)

Educational Disparities in BMI (1990 BRFSS)

Education Level	%	BMI	RR	ER
<8 years	5.66	26.6	1.09	2.2
Some High School	10.65	25.7	1.05	1.3
HS Grad / GED	34.10	25.1	1.03	0.7
Some College	25.95	24.6	1.01	0.2
College Grad	23.63	24.4	1.00	0.0

55

Using this table, let's examine Educational Disparity in Body Mass Index (BMI) according to the 1990 BRFSS. This is a typical data layout for examining disparities. Notice it contains a range of ordered educational groups, from less than eight years through college graduates.

In the first two columns, the table shows:

The percent of the population with less than 8 years of education (5.66%)

The percent of the population that has graduated from college (23.63%)

And so on.

The next column shows average levels of Body Mass Index within each educational group.

As you have seen before, we can easily calculate relative risks (RR in the chart). You can tell the reference group in this case is college graduates, since the relative risk value is equal to one (1) for that social group.

The disparity in terms of excess risk (ER in the chart), is displayed in the last column. Excess risk in this table has been calculated according to the absolute difference between BMI in the reference category, the college graduates, and in each of the education level categories. Relative measures of extreme groups are the ones typically used in epidemiology and public health.

Range measures typically compare the two extreme categories.
One of the extremes is used as the reference group, which is compared to the other extreme. In this case, the ratio of BMI among those with less than eight years (the least number of years of education) is compared to college graduates (the group with the most years of education).
The 26.6 BMI for those with less than eight years of education is divided by 24.4, which is the BMI for those in the reference group—college graduates—resulting in a relative risk of 1.09.

If we were to calculate excess risk as a measure of absolute disparity, we would subtract 24.4 from 26.6 and that absolute arithmetic difference is 2.2.

Notice is that we don't use *any* of the information about the groups in between. In other words, our measure of disparity, if we were to use a relative risk or an excess risk, is based only on information about the two extreme social groups. Notice also that in using these range measures we are not using any of the information in the first column on the relative size of the different educational groups.

Range measures

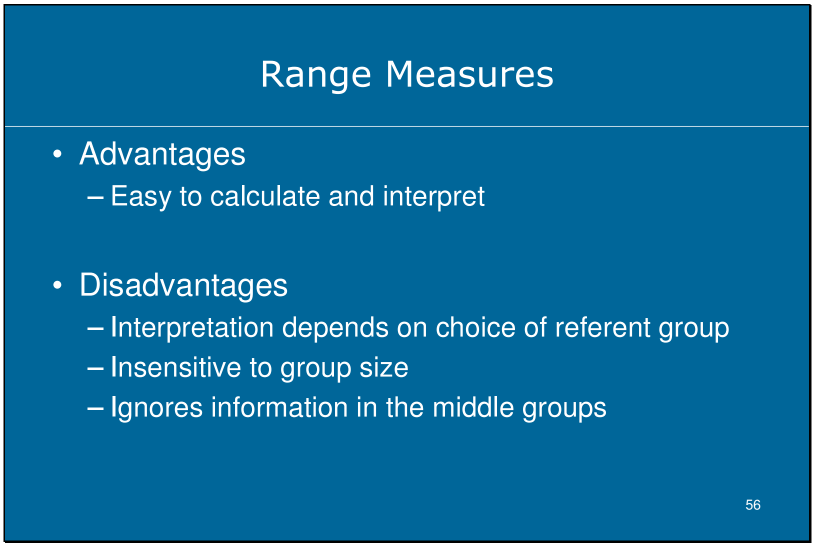

The advantage of range measures is that they are very easy to calculate and interpret since they are familiar to most people.

The disadvantages are several. The interpretation of range measures depends on the choice of the referent group. We discussed this in Part II.
When you change the reference category, the number you generate for the relative or the excess risk will differ.
These range measures are insensitive to the size of the groups. In the example of educational disparities in BMI, the measurement did not account in any way for the fact that only about 6% of the population has less than 8 years of education.
Range measures also ignore information on any group whose data falls in the middle range rather than the extreme.

Un-weighted regression measures

Un-weighted, Regression-Based Measures allow us to begin to incorporate information that exists in all groups, not just the two extremes, as in the range measures.

Un-weighted regression measures

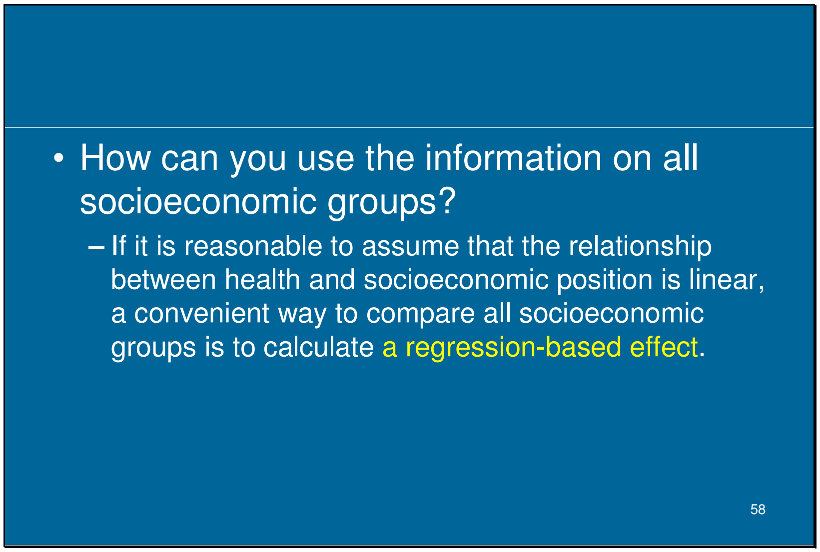

As we just saw, it does not seem intuitively right to ignore all the information that exists in middle groups, and rely exclusively on two groups for a comparison. If we can assume a linear relationship between the health indicator of interest and the indicator of socioeconomic position (such as education or income), then a convenient way of using all information for all socioeconomic groups is to calculate a regression-based effect measure.

Un-weighted regression measures

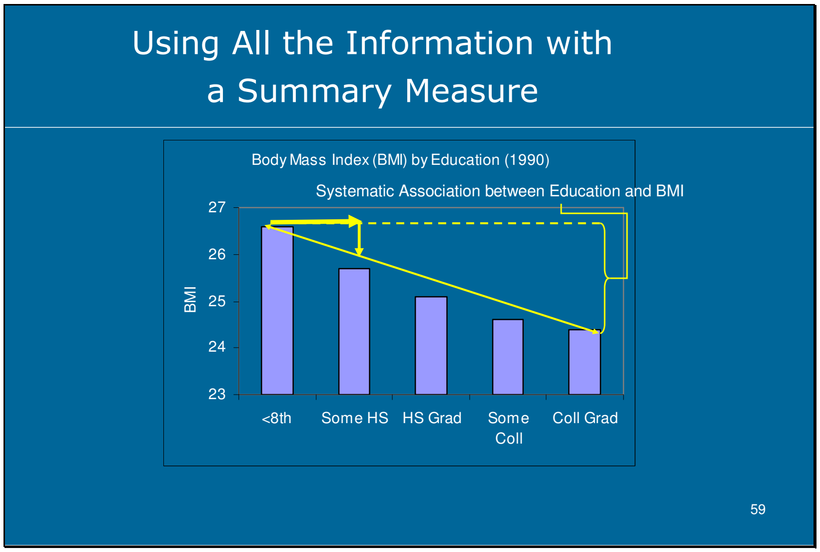

How is all the information used?

First, arraying the data allows you to ***regress*** (a statistical technique) the average BMI across the educational groups to calculate an average effect measure.

This difference between the college graduates and the less-than-8th-grade groups is expressed in a slope of the line, which represents the systematic association between education and BMI across all groups.

The interpretation of slope is that:

For an increase of one unit of education…

… the average decrease in BMI is a constant amount

In this case, a single number—the slope of a line—summarizes the data across the different groups rather than just using the information on the two extreme groups. How well this value summarizes a systematic association depends on various assumptions. The most important assumption is that the relationship between BMI and education is linear.

Un-weighted regression measures

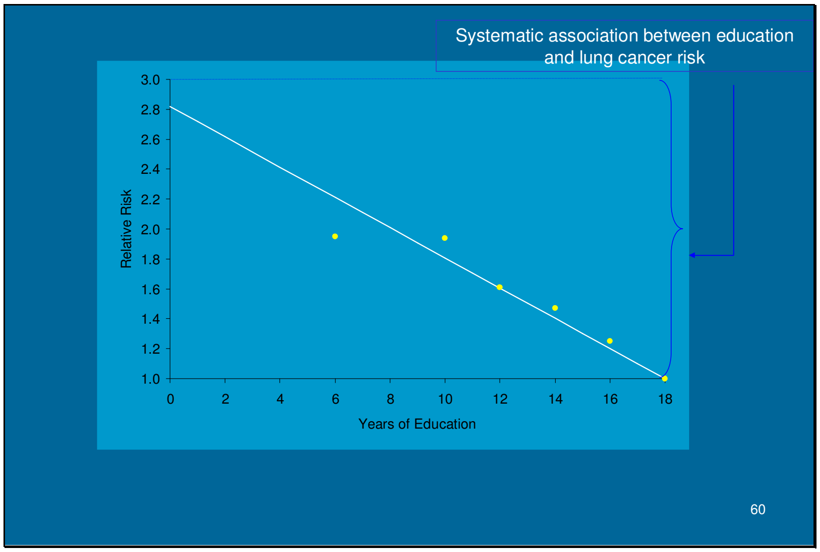

This example is from a paper by Steenland and colleagues that examines the systematic association between education and lung cancer risk.

In their study, the researchers calculated a set of relative risks using the highest education group (18 years) as their reference. In the graph, you can see the relative risk—or the association between education and lung cancer risk—for those with 16 years of education was about 1.3.

For those with only 6 years of education, there is approximately a twofold risk. If we want to summarize the information contained in the scatter plot, we could calculate and draw a regression line like the one shown. The slope of this line is the beta coefficient, described in discussion of the next measure, and the slope summarizes the information contained in all five of the data points into one number rather than five.

For more information about this particular study, refer to the *Resources* section.

Un-weighted regression measures

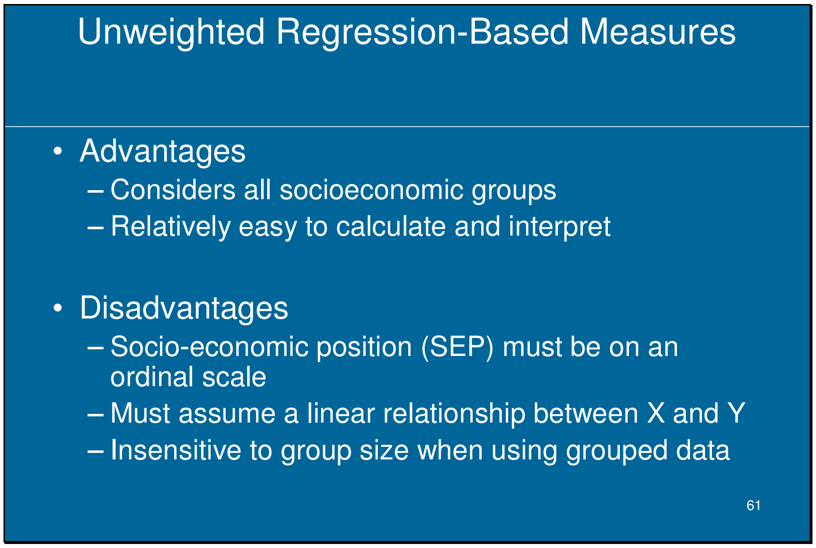

The advantages to un-weighted, regression-based measures are that they take into consideration information from all socioeconomic groups and they are relatively easy to calculate and interpret.

Like range measures, many people in public health are accustomed to seeing ***beta coefficients*** (that is, the slope of the line) that can be interpreted as a relative risk.

One of the disadvantages to un-weighted, regression-based measures is that our social grouping or socioeconomic position must be on an ordinal scale. In other words, the measures are valid only if you can order the groups. These measures also assume a linear relationship between the social group and the outcome. Lastly, they are insensitive to group size when using group data.

Population weighted regression measures

Population-Weighted, Regression-Based Measures allow us to incorporate information about the size of the social group by weighting.

Population weighted regression measures

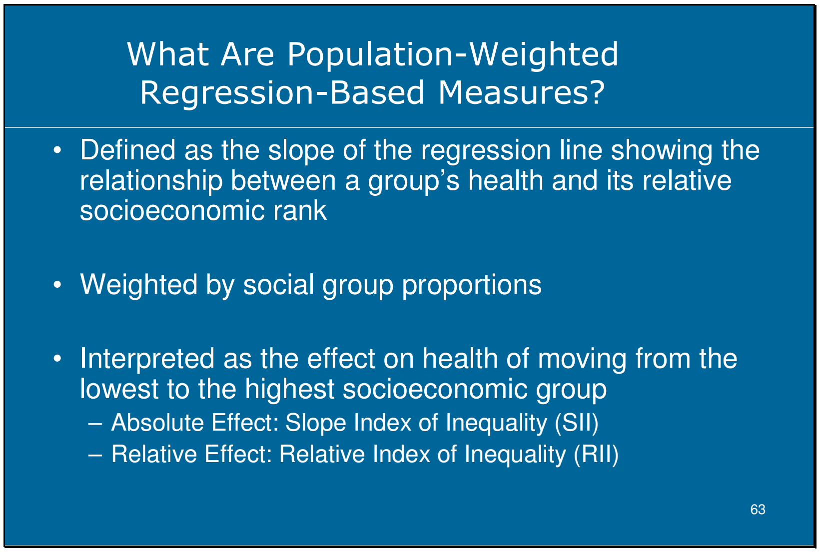

Population-weighted, regression-based methods are similar to the previous measures in that they involve finding the slope of a regression line, which measures the relationship between a group's health and its relative socioeconomic rank. Where population-weighted, regression-based methods differ from previous methods is that they enable us to incorporate information about the size of the social group by weighting.

These measures are interpreted as the effect on health of moving from the lowest to the highest socioeconomic group. In this section we look at two specific measures that account for the absolute and relative effects: the Slope Index of Inequality (the SII) and the Relative Index of Inequality (the RII).

Socioeconomic disparity as measured by the RII is becoming a more commonly used measure. The *Resources* section contains references to specific examples of how to use each of these measures in practice.

Population weighted regression measures

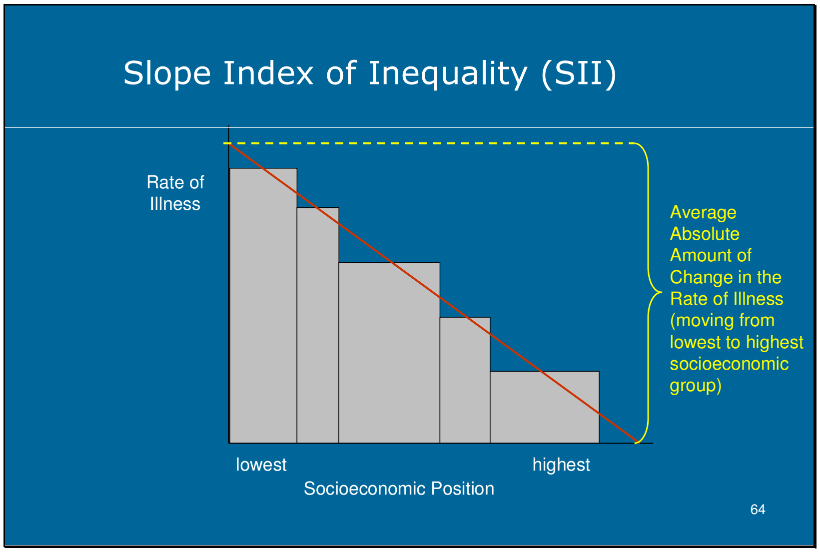

The approach to the Slope Index of Inequality (the SII) is similar to the one used for the un-weighted, regression-based measures.

We begin with a ranking of groups based on socioeconomic position, such as educational or income groups along the X-axis. We have also illustrated the size of the groups by adjusting the width of the bars as shown. (In the previous example, the width of the bars was all the same.) The X-axis depicts the relative rank of the socioeconomic group, with some indication of its size in the population, as expressed by the width of the intervals (bars).

Differing rates of illness are on the Y-axis.

If we use this data for regressing just like before, but weight the social groups by their population size, then the slope of the line indicates the average absolute amount of change in the rate of illness in moving from the lowest to the highest socioeconomic groups. It is the absolute amount because we are still using the same units we used in measuring the rate of illness. These units could have been infant mortality, heart disease, or any other rate of illness or health status indicator of interest. Note that this SII measure uses the information on all groups and information on the size of the groups.

Population weighted regression measures

SII Calculation

Distribution of Educational Position (1990)

Education Level	%	% Cumulative Population	Range % Cumulative Population	Midpoint
< 8 Years	5.66	5.66	0.0 – 5.66	2.83
Some High School	10.65	16.31	5.67 – 16.31	10.99
HS Grad / GED	34.10	50.42	16.32 – 50.42	33.37
Some College	25.95	76.37	50.43 – 76.37	63.40
College Grad	23.63	100.0	76.38 – 100.0	88.19

65

Let’s take a closer look at the basic data setup behind the calculation of the SII.

Again, we start with the categories of education and the proportion of the population in each of these groups. The next column is the cumulative percent. For example, 16.31 is the cumulative percent of those with less than eight years of education and those with some high school, which is simply the sum of 5.66 and 10.65. Notice that the cumulative percentage adds up to 100.

The range expresses the cumulative distribution of the population according to the socioeconomic position that each group occupies. For example, the group with some high school education occupies the range of 5.67 to 16.31% of the population. In the table, the third column shows the range in the cumulative distribution of education that each educational group occupies.

We need to know the range in order to calculate its midpoint for each socioeconomic group. The range midpoint is the value used in the regression to calculate the SII. Please refer to articles in the *Resources* section for more technical details.

Population weighted regression measures

SII Calculation

- Regress the health outcome (BMI) on the midpoint of socioeconomic categories, weighted by proportion in the population:

$$y = \beta_0 + \beta_1(\text{SEP midpoint}) + \varepsilon$$

  - Slope Index of Inequality (SII) = $- \beta_1$
  - Relative Index of Inequality (RII) = $(- \beta_1) / y$

66

Once we know the midpoints, we can regress the health outcome (the BMI in this case) on the midpoint of the socioeconomic position (SEP) categories.

A typical linear regression model is used where:
Y is the outcome, BMI
Beta-naught is the intercept of the regression line and the Y-axis
Beta-1 is the coefficient that relates BMI to the midpoint of the range of the distribution of Socioeconomic Position (SEP) and
An error term, Epsilon

Remember that:
Beta-1 is just the slope of the regression line, or the average change in the BMI per-unit increase in education category.

The Slope Index of Inequality is negative beta-1. The SII is interpreted as the absolute change in BMI involved in moving from the lowest to the highest socioeconomic group.

The Relative Index of Inequality is negative beta-1 (or the Slope Index of Inequality) divided by the population average for the health outcome (in this case BMI). The RII is an expression of the absolute disparity in the health outcome relative to the average level in the population.

Population weighted regression measures

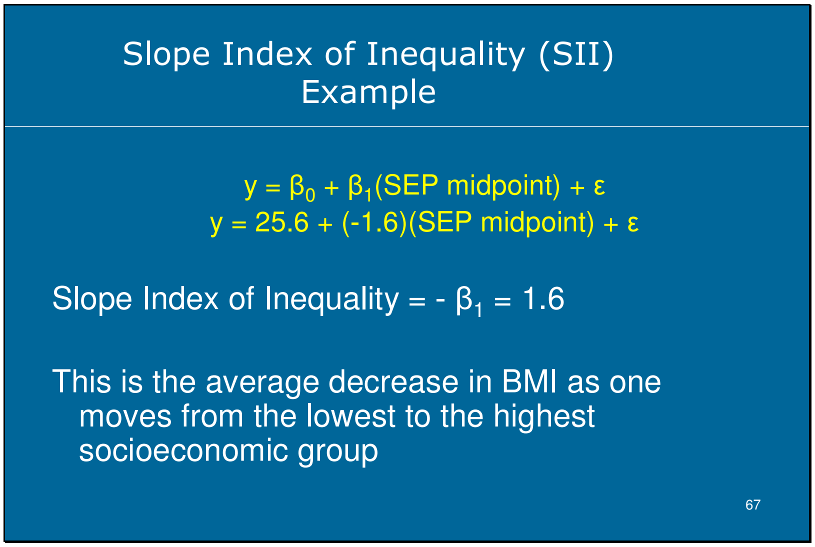

Let's see how this works with the data we have for BMI by years of education.

We'll start with the Slope Index of Inequality: $y = \beta_0 + \beta_1$ (SEP midpoint) + ε
This is the generic formula.

After performing the regression, we find that: y = 25.6 + (-1.6) (SEP midpoint) + the error term.

This suggests that there is a 1.6 unit decrease in BMI as you move from the lowest to the highest socioeconomic group. Therefore, beta-naught (or 25.6) is the BMI value of the hypothetically least-educated person. Beta-naught is the value of the BMI when the SEP midpoint equals zero and is the y-intercept of the regression line.

Population weighted regression measures

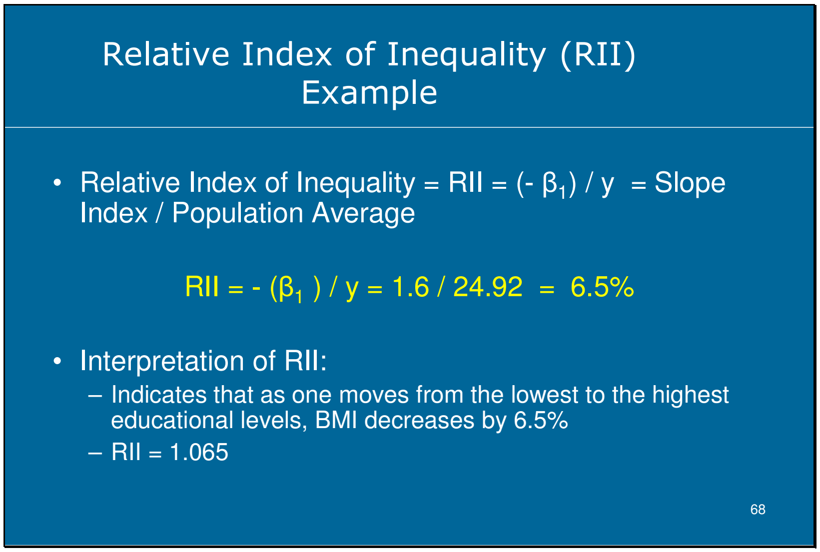

Once you know the Slope Index of Inequality, it is easy to find the Relative Index of Inequality (RII).

The RII is the SII divided by the mean BMI for the population. We can tell you, from calculations not shown, that the mean BMI value for the population is 24.92. Inserting these values into the formula gives you:
1.6 divided by 24.92 equals 6.5%

We can interpret this RII to mean that as one moves from the lowest to the highest educational group BMI decreases by 6.5%.

Applying the more commonly used rate ratio measures, an RII of 6.5% would be a rate ratio measure of 1.065.

Population weighted regression measures

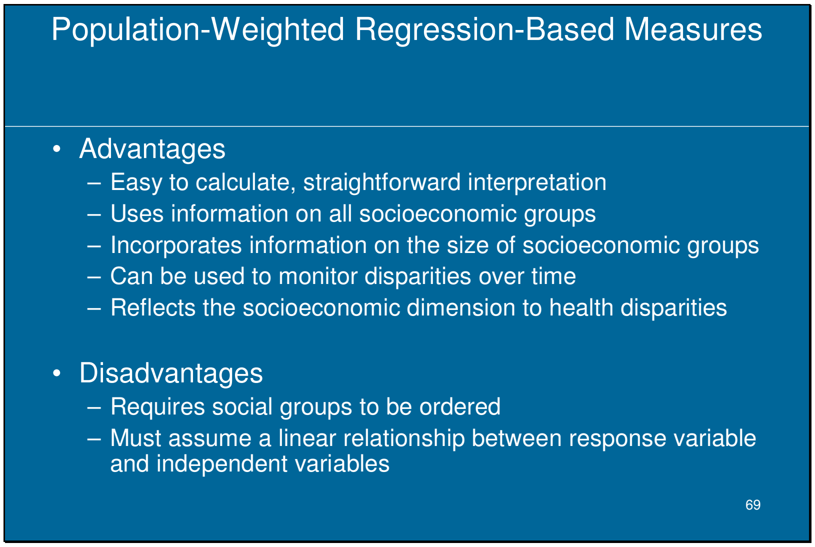

The advantages of the relative and slope indices of inequality include being fairly easy to calculate and having a reasonably straightforward interpretation, especially because they correspond to things that we're familiar with in the regression-modeling framework.

Most importantly, these indices use information on ***all*** the socioeconomic groups and incorporate information on the ***size*** of the socioeconomic groups. Also, you can use them to monitor disparities over time because they are sensitive to changes in the size of the socioeconomic groups, as well as changes in the rates of the health outcome. We think these are very important characteristics of a disparity measure.

Furthermore, these indices reflect the socioeconomic dimension to health disparities. The assumption is that we care more about a health disadvantage in a lower socioeconomic group than we do in a higher socioeconomic group. Some economists and philosophers argue that incorporating this concern is a desirable characteristic of a health inequality measure.

The major disadvantages to the SII and RII are that you can only use them when the social groups can be ordered. As we've seen before, many of the concerns of health disparities in the United States, as laid out in Healthy People 2010, do not involve ordered social groups.

Index of disparity

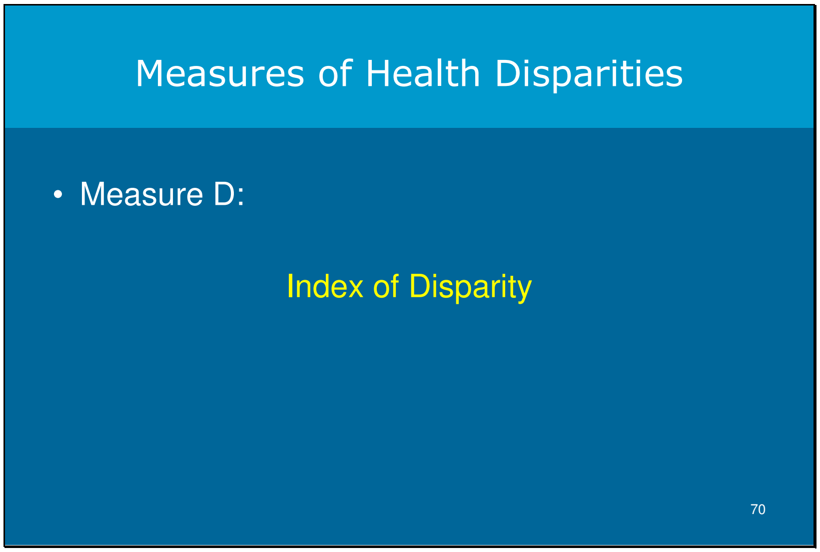

Measure D: Index of Disparity.

Keppel and colleagues from the National Center for Health Statistics have recently proposed the ***Index of Disparity*** as a recommended means for measuring health disparities. You may see also see it in the academic literature.

Index of disparity

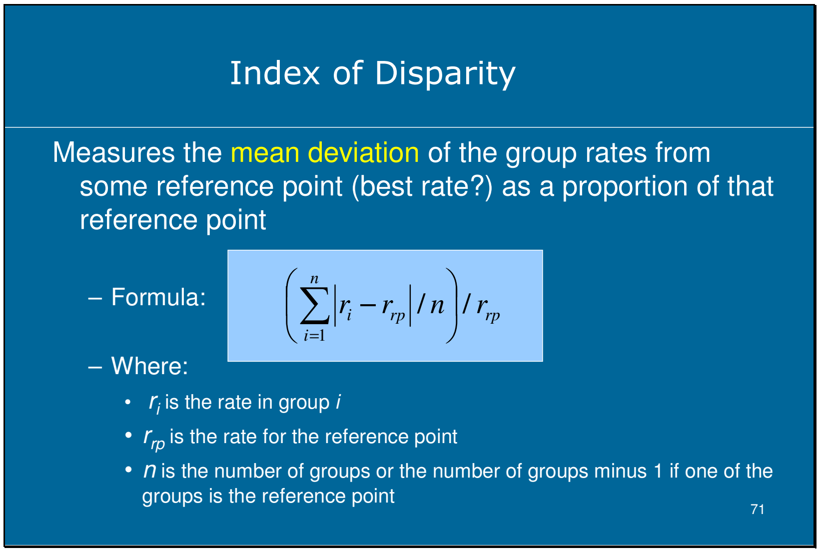

The index of disparity measures the mean deviation of several group rates from a given reference point (r_{rp}). The given reference point is usually the best group rate or total rate as a proportion of that reference point.

Keppel, et al., describe some of the more technical features of this measure in a paper cited in the *Resources* section. In essence, the calculation of the index of disparity simply involves the following process:

Subtracting each single group rate from the reference rate

Taking the absolute value of those differences

Summing all those differences, and

Expressing those differences as a proportion of the reference rate

Index of disparity

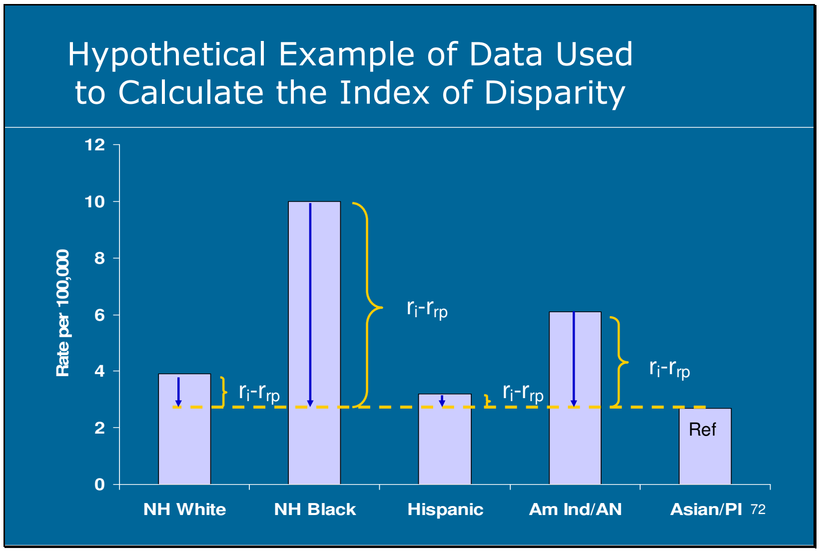

This is an example of what the index of disparity looks like in practice. Let's step through the process for determining this index.

1. Identify the reference rate. In this example, we want the best rate for this particular health outcome, which happens to occur among Asian/Pacific Islanders. The social groups deviate from this reference rate by different amounts. The largest deviation from this rate is among non-Hispanic blacks.

2. Sum up the deviations among all of the remaining social groups, as absolute values. In our example, that would mean summing up the deviations in rate from the reference group and the following:

Non-Hispanic whites

Non-Hispanic blacks

Hispanics

American Indians and Alaskan Natives

3. Average these deviations.

4. Divide the mean deviation we've just calculated by the reference rate, which is the rate among the Asian/Pacific Islanders.

Index of disparity

Index of Disparity: Calculation

How great is the mean deviation between race/ethnic-specific infant mortality rates and the total rate as a proportion of the total rate?

Mother's Race and Ethnicity	Infant Mortality Rate	$\lvert r_i - r_{rp} \rvert$
Non-Hispanic White, r_1	6.0	1.2
Non-Hispanic Black, r_2	13.9	6.7
Hispanic, r_3	5.8	1.4
Asian / Pacific Islander, r_4	5.5	1.7
American Indian / Alaska Native, r_5	9.3	2.1
Total Rate, r_{rp}	7.2	-
Sum of the Deviations $= \Sigma \lvert r_i - r_{rp} \rvert$		13.1
Mean Deviation $= \Sigma \lvert r_i - r_{rp} \rvert / n$		2.62
Index of Disparity = Mean Deviation / Reference Point $= (\Sigma \lvert r_i - r_{rp} \rvert / n) / r_{rp}$		0.36

This table provides a new example, and more detail for calculating the index of disparity. The best rate is the lowest infant mortality rate, which is 5.5 among the Asian/Pacific Islanders. The highest rate, 13.9, is indicated in the non-Hispanic black row.

In this example, the total rate is the reference point.

The deviation from the total rate, among non-Hispanic whites, is 1.2, which is the absolute value of the rate among non-Hispanic whites minus the total rate.

The deviation from the total among non-Hispanic blacks is 6.7, Hispanics 1.4, Asian / Pacific Islander 1.7, and American Indian/Alaskan Native 2.1.

If we sum all the deviations, we get 13.1.

The mean deviation, 2.62, is the sum divided by 5, the number of groups.

The index of disparity is .36, which is 2.62 (the mean deviation) divided by 7.2 (the total infant mortality rate) and is the mean deviation expressed in terms of the reference group rate.

Index of disparity

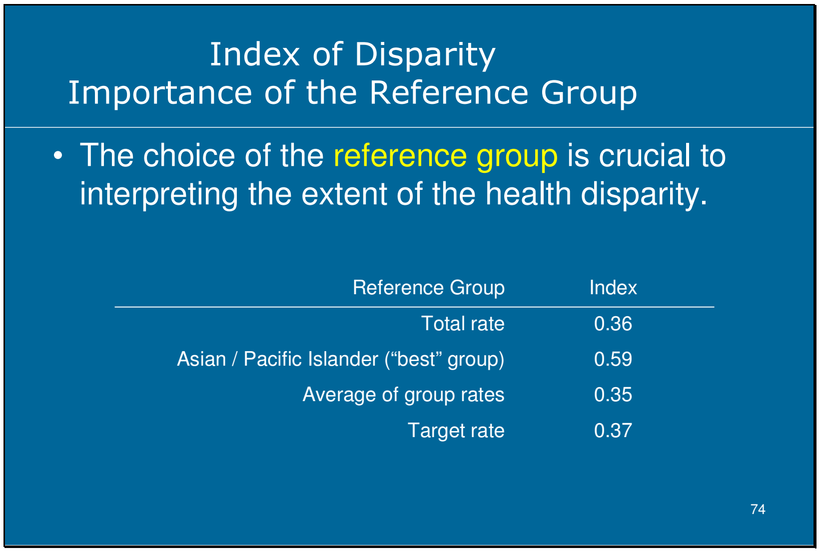

The size of the index of disparity depends on which reference group is chosen.

If we use the total rate as the reference group, as we did in the previous example, the index of disparity is 0.36.

If we use the best rate, that of the Asian Pacific Islander, the index of disparity would seem to be much larger, at 0.59.

If we use the average of all the group rates, the value would be 0.35.

If we use the target rate, as laid out in Healthy People 2010, the index of disparity would be 0.37.

As you can see, the choice of reference group is crucial to interpreting the extent of health disparity. The authors of the Index of Disparity recommend choosing the best group rate as the reference rate.

Index of disparity

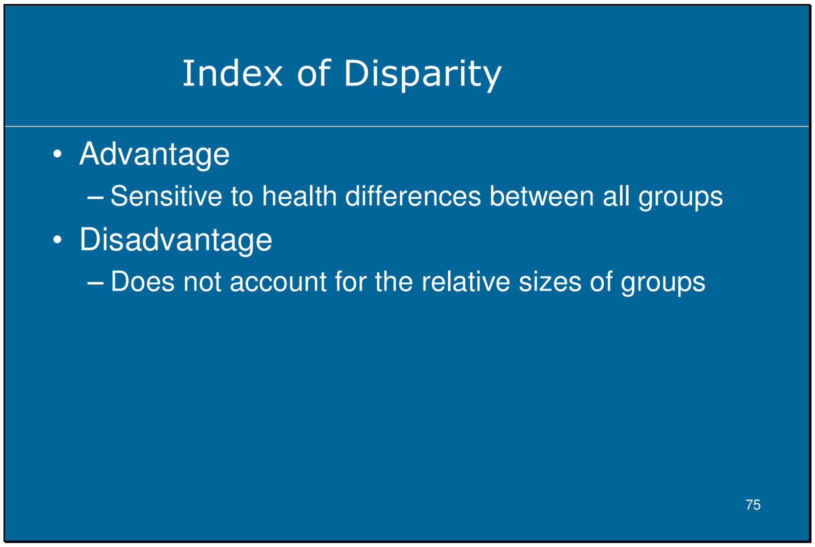

The index of disparity only compares the rate or the prevalence. It is sensitive to health differences only, *not* the size of the groups experiencing those rates or the prevalence of the different health states.

The advantage of the index of disparity is its sensitivity to health differences between all groups. The disadvantage is that it does not account for the size of the groups, and it only compares rates or prevalence of health status.

Between-group variance

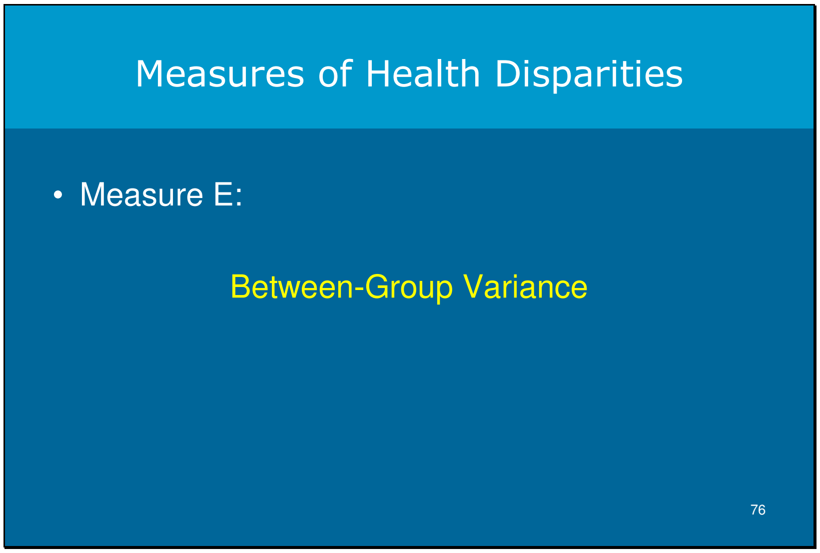

Measure E: Between-Group Variance

The Between-Group Variance measures the deviation of each group's rate from the population average and weights each group by its population size. This measure is similar to the index of disparity, except it has the desirable characteristic of including the size of the population.

Between-group variance

What is Between-Group Variance?

- Measures the deviation of each group's rate from the population average and weights each group by its population size
  - Formula: $$\sum_{j=1}^{J} p_j \left(y_j - \mu\right)^2$$
  - Where:
    - y_j is the rate in group j
    - μ is the population average rate
    - p_j is the group's share of the total population

77

Notice in the formula that we use the squared difference of each group's rate and the population average. This means rates that are further from the population average will actually have a greater influence when we calculate the summary index.

For example, if the disparity between Group A and Group B is 4, the squared difference is 16. On the other hand, if the difference is only 2, then the squared difference is 4.

Even though the difference between the two groups is double (2 vs. 4) their contribution to the disparity measure is much larger (4 vs. 16) because the values are squared. By *squaring* the difference, we are implicitly saying greater disparities should be weighted more than smaller disparities. This is an excellent example of how our values and ideas about disparities may or may not be reflected in the measure of disparity.

The index of disparity we discussed earlier does not use a squared term in its calculation. In that measure, all deviations from the reference have the same "weight." Between-Group Variance, which uses a squared term, implicitly reflects a belief that groups further away from the reference group should get higher weighting when calculating the size of disparity.

Between-group variance

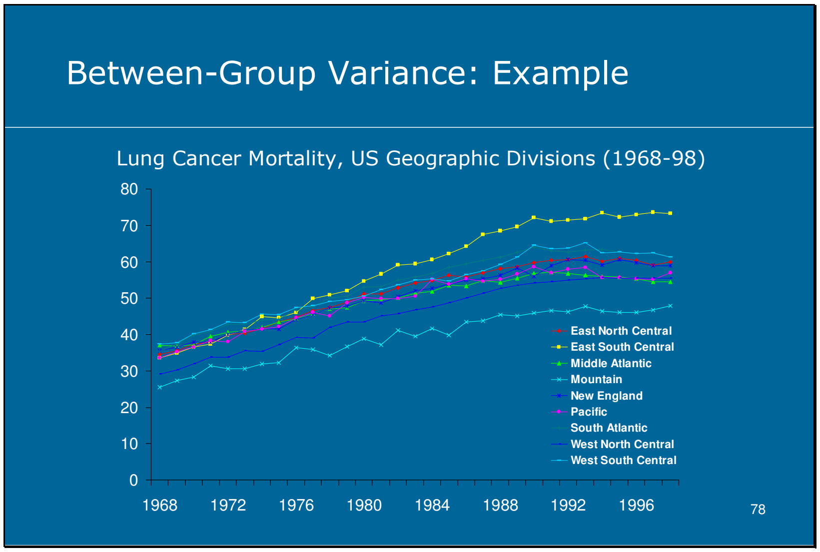

This is an example of data we might use if we wanted to answer the question, *"Have regional differences in lung cancer mortality increased over the last 35 years?"*

This is a typical question for health disparities investigators. But, where to begin?

In this example, nine different regional groups are represented. It is very hard to summarize the differences between all of them unless we use eight numbers to compare the mortality rates one-by-one, and group-by-group, and that does not take into consideration trying to analyze them over time.

A procedure like this would not be very efficient. In this type of situation, summary measures like the Between Group Variance are helpful.

Between-group variance

Between-Group Variance: Calculation

$$\sum_{j=1}^{J} p_j (y_j - \mu)^2$$

	1968			1998		
Region	% Pop	Rate	BGV	% Pop	Rate	BGV
New England	5.9	35.0	0.0	5.0	56.3	0.0
Middle Atlantic	18.5	37.0	0.9	14.2	54.0	1.3
East North Central	20.0	34.5	0.0	16.2	59.6	1.0
West North Central	8.1	29.2	2.5	6.9	54.8	0.4
South Atlantic	14.8	35.3	0.1	18.2	60.0	1.7
East South Central	6.3	32.6	0.3	6.1	70.1	10.4
West South Central	9.4	36.8	0.4	11.1	62.1	2.9
Mountain	3.9	27.5	2.1	6.3	46.2	7.4
Pacific	13.0	36.1	0.3	15.9	50.7	6.4
Total	**100.0**	**34.7**	**6.5**	**100.0**	**57.0**	**31.5**

79

In this example, we are using graphical data from a spreadsheet to help us calculate the Between-Group Variance.

Applying the formula for the Between-Group Variance to the information provided in the columns “Percent Population” and “Rate” gives us the Between-Group Variance (the “BGV”) for each group in 1968 and 1998. In 1968, the total Between-Group Variance was **6.5 deaths per 100,000**. By 1998, BGV increased to **31.5 deaths per 100,000**.

Compared to the average rate in the population, much larger differences existed among the regions in 1998; the size of the difference increased about fivefold to over 30. The regional disparity is increasing over time.

This conclusion is supported by what we see when we look at the graph again. We see the disparities spreading out across the regions over time. The advantage of measures like the Between-Group Variance is that it provides a quantifiable number for the change in disparity.

Between-group variance

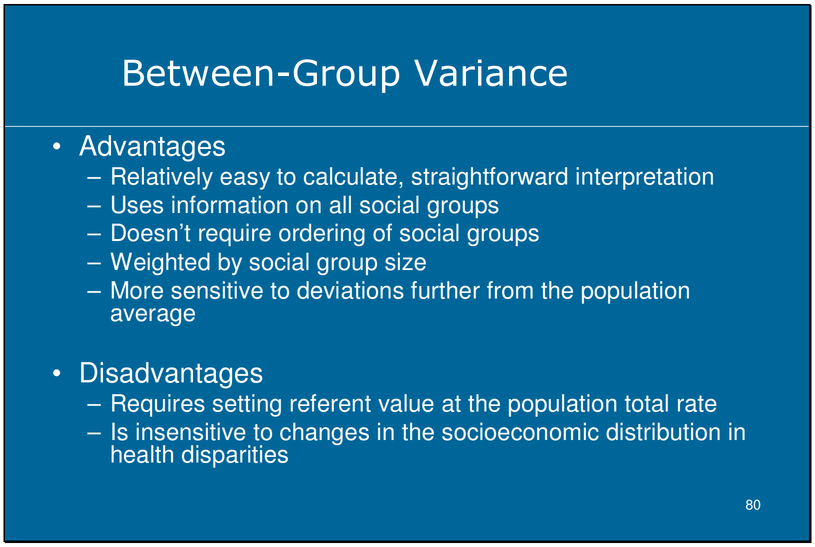

The advantages of the Between-Group Variance (BGV) include that it is relatively easy to calculate and is fairly straightforward in interpretation. It uses information on all social groups. It does not require ordering of social groups. (We just calculated Between-Group Variance for regions, which cannot be ranked). This measure is weighted by the group's size and is more sensitive to deviations further from the population average.

Disadvantages of the Between-Group Variance include that it requires setting a referent value at the total population rate. Also, BGV is insensitive to changes in the socioeconomic distribution in health disparities since it describes the change in the variation across social groups. It does not point to particular social groups that are experiencing improvements or declines.

The Between-Group Variance simply summarizes the amount of variation without regard to patterns of disparity between particular social groups.

Disproportionality measures

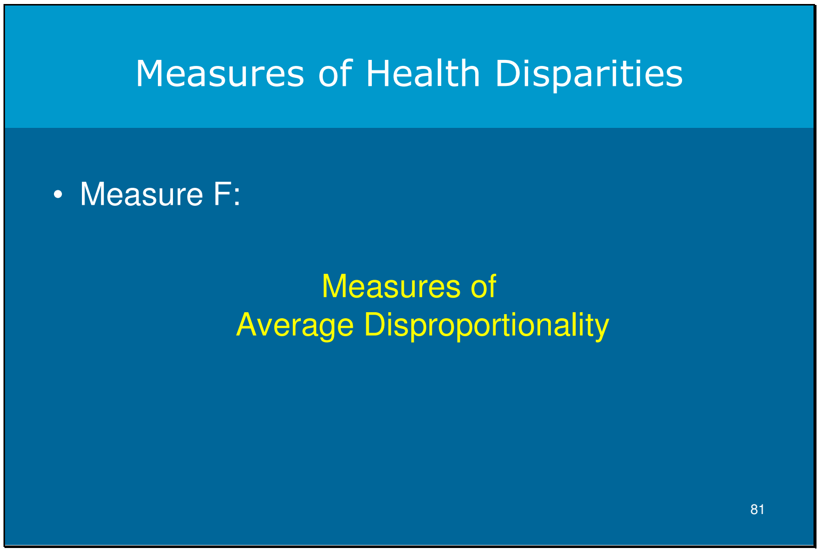

Measure F: Average Disproportionality

Like lots of things in life, there's no free lunch—the same is true in measuring health disparity.

The following measures are specifically designed to be like the other summary measures of health disparity, but they are somewhat more complicated in their calculation and interpretation.

These measures are more often used in disciplines like demography and economics. They are very rarely used in epidemiology and public health applications. However, they do have certain characteristics that make them attractive for the measurement of health disparities and they are more complicated to calculate.

To understand the application of these more complicated measures, which have some desirable characteristics but are not commonly used in public health, we'll begin with a discussion of disproportionality. After that, we'll work through examples of the Gini Index (or coefficient), Health Concentration Index, Theil Index, and Mean Logarithmic Deviation.

Disproportionality measures

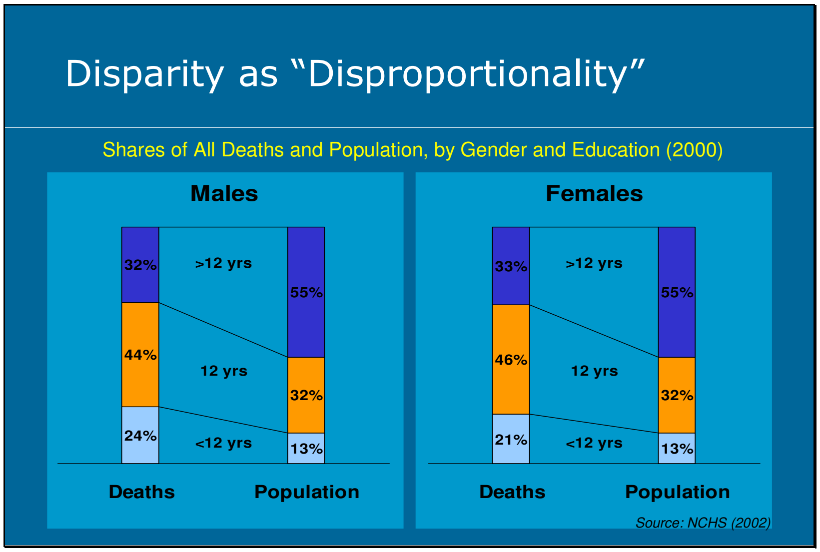

We frequently use the language of disproportionality in health disparities research, intervention, and policy.

For example, we often hear that certain social groups bear a *disproportionate* burden of ill health. How would this concept be incorporated into a specific measurement? Literally, the measure would show there is a disproportionate burden of ill health borne by a group, relative to its size in the population.

If a population subgroup comprises a given percentage of the population, then the disease burden of this population should be equivalent.

Let's start on the right-hand side of this graph, with those females having less than 12 years of education. In the population, these females comprise 13% of the population, and yet they comprise 21% of the deaths attributed to disease. This is disproportionate. If it was proportional, they would have 13% of the deaths.

Notice that women with more than 12 years of education comprise 55% of all females in the U.S. Yet, they experience only 33% of the deaths. This is also disproportionate.

Disproportionality measures

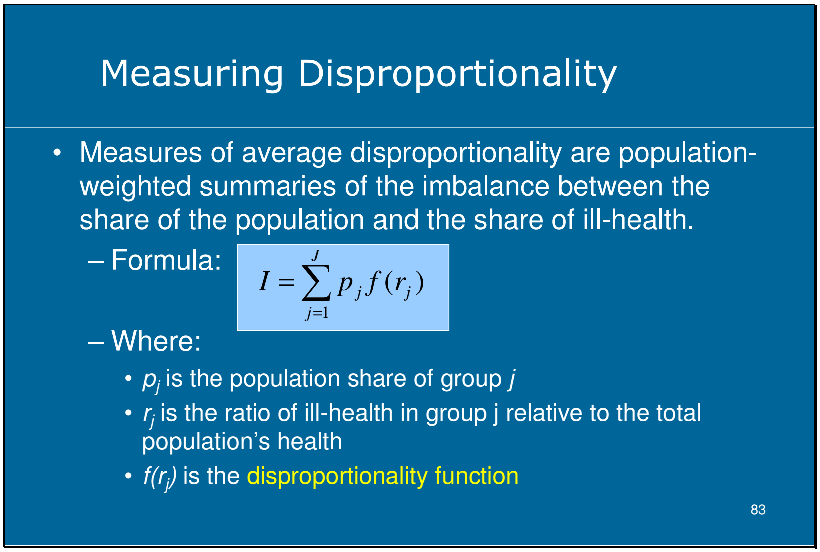

Measures of disproportionality are population-weighted summaries of the imbalance between the share of the population and the share of ill health. In other words, if a population group represents 10% in the population, it should experience 10% of the share of ill health for there to be *no* disproportionality.

These measures take a generic form, as shown in the formula. It is a summary measure of a function of the ratio of ill-health in each subgroup (r_j) relative to the total population's health, weighted by the population share of that subgroup, (p_j).

The key difference between the types of disproportionality measures is how they express the ***f***, the mathematical function.

Disproportionality measures

Commonly Used Disproportionality Functions

Index Name	*Disproportionality Function $f(r_j)$*
Gini Index or Coefficient (*G*)	Individual-level data: $\mid r_i - r_j \mid / 2$ Grouped data: $r_j(q_j - Q_j)$, where q_j is the proportion of the total population in groups less healthy than Group *j*, and Q_j is the proportion of the total population in groups healthier than Group *j* (i.e., $p_j + q_j + Q_j = 1$)
Health Concentration Index (*HCI*)	Same as for G, but groups are ranked by social group position instead of by health, so that q_j is the proportion of the total population in groups less advantaged than Group j, and Qj is the proportion of the total population in groups more advantaged than Group j (i.e., $p_j + q_j + Q_j = 1$)
Theil Index (T)	$r_j \ln(r_j)$
Mean Logarithmic Deviation (MLD)	$\ln(1/r_j) = -\ln(r_j)$

Adapted from Firebaugh, 2003.

84

Several commonly used measures use this general form, especially in economics, demography, sociology, and increasingly in epidemiology.

These measures include all of the following:

- The Gini Index
- The Health Concentration Index
- The Theil Index
- The Mean Logarithmic Deviation

Each differs in how it is constructed and each incorporates a particular view for how to express this function of disproportionality. Nevertheless, they all take the general form of trying to summarize the amount of disproportionality across population share and share of ill health.

We will provide an overview of these measures. Explaining the technical details of these measures is beyond the scope of this CD-ROM. However, you should be aware these measures of disproportionality exist.

For more details on how to calculate these measures, refer to the technical papers referenced in the *Resources* section.

Disproportionality measures – The Gini Coefficient

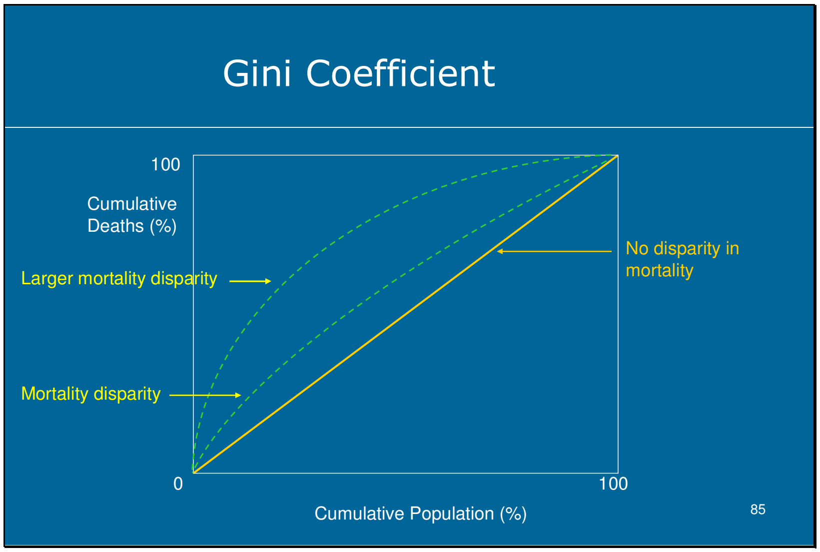

The Gini Coefficient can be depicted graphically. To start with, let's review the X and Y axes:
The cumulative proportion of the population, from 0 to 100%, is along the X-axis. The cumulative percentage of deaths (or another measure of disease burden) is on the Y-axis.

If no disproportionality in deaths exists, the Gini Coefficient equals 0 and 50% of the population would experience 50% of mortality, 10% of the population would experience 10 % of mortality, et cetera.

The diagonal line represents a population with no disproportionality between the cumulative proportion of the population and its cumulative experience of death.

When there is disproportionality, the ratio between cumulative proportion of the population and its cumulative experience of mortality is no longer 1 to 1. The Gini Coefficient, then, is represented as a curve and can range in value from -1 to 1, depending on which side of the diagonal it falls. As you can see, the depth of that curve indicates the depth of the disparity.

Frequently, the Gini Coefficient is used to measure income distributions, but it is not often applied to distributions of health in populations.

Disproportionality measures – The Gini Coefficient

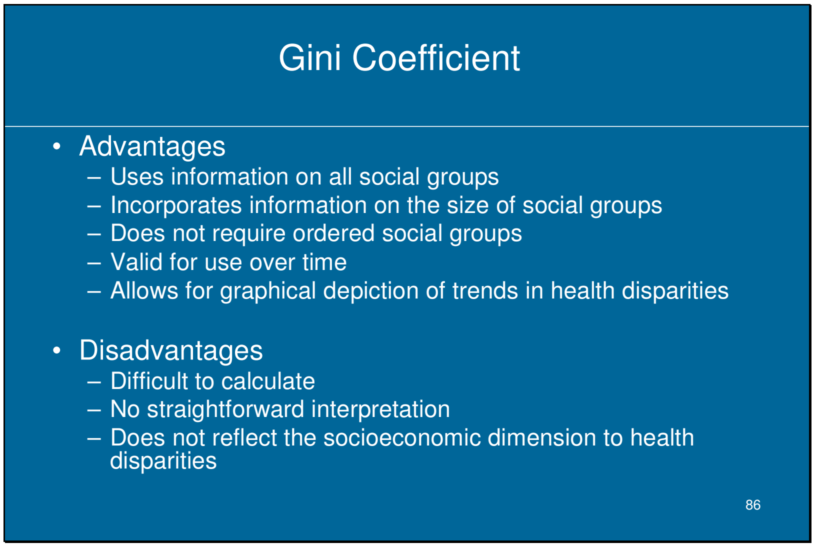

There are several advantages to using the Gini Coefficient as a measure of disparity.

First, it uses information on all social groups so everyone in the population is represented.

Second, the size of the social groups are represented in the measure.

Third, it does not require social groups to be ordered.

Fourth, it is valid for use over time.

And, finally, you can graphically depict this measure, which is often good for communicating with policymakers and the community.

There are disadvantages for using the Gini Coefficient as a measure of disparity. For example, it is somewhat difficult to calculate and its interpretation is not one to which we are commonly accustomed, especially as compared to relative risk. Another disadvantage is that it doesn't reflect the socioeconomic dimension of health. The basis for comparison is merely the cumulative proportion of the population against the cumulative proportion of the particular outcome of interest.

The Gini Coefficient is a measure of pure variation in health that does not explicitly include a consideration of social groups.

Disproportionality measures – The Health Concentration Index (HCI)

Health Concentration Index

- Ranks people by their socioeconomic position
- Plots the cumulative proportion of the population against the cumulative proportion of health/illness
  - Begins with the most disadvantaged population and ends with the least disadvantaged population
- If health is equally distributed, this plot is a 45° line and C = 0.
- If health is unequally distributed, the concentration index is equal to twice the area between the plotted curve and the 45° line.

87

Another disproportionality measure used increasingly in public health is the Health Concentration Index (HCI).

Think of this index as an extension of the Gini Coefficient, but instead of using the cumulative proportion of the population, the HCI also arrays the population according to rankings by socioeconomic position. In this sense, it is like the RII we previously discussed and, in fact, the HCI is mathematically related to the RII.

Like the Gini Coefficient, the HCI is usually depicted graphically:
Plot the cumulative proportion of the population, starting with the most disadvantaged group and ending with the least disadvantaged, against this cumulative proportion of illness along the X-axis.
Graph the cumulative percentage of disease burden along the Y-axis, as we did previously.

Like the Gini coefficient, if health is equally distributed, the diagonal at 45° shows the concentration index to be 0, and no social group disparity in health will be apparent.

Disproportionality measures – The Health Concentration Index (HCI)

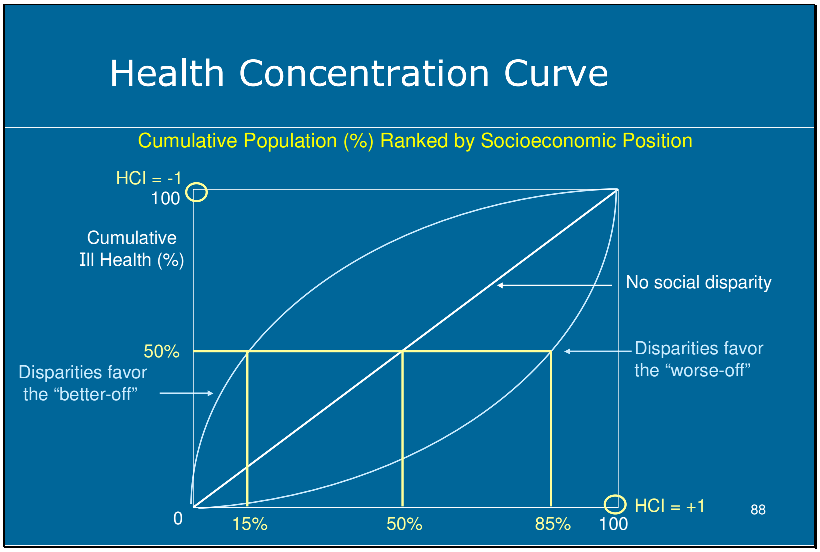

This is what a Health Concentration Index will look like.

The X-axis ranks the cumulative population by socioeconomic position, such as the cumulative proportion of the population by education, by income, or by some variable that can be rank-ordered.

The Y-axis plots the cumulative share of health.

The line along the diagonal represents the situation in which 50% of the population ranked by the socioeconomic indicator encounters 50% of the ill health. In other words, ill health is equally shared by each socioeconomic group. Along the diagonal, the concentration index is equal to zero and the interpretation is that there is no social disparity.

But what if we had a curve that looked like this?

In this case, the 15% of the population that is the least-well-off in terms of socioeconomic position accounts for half of all the ill health in the population.

This is typical of what we see in health disparity situations in the U.S.: The least-advantaged groups suffer a disproportionate burden of ill health and disparities tend to favor the better off. It is possible for you to see this kind of curve in other situations, since not all health outcomes involve worse health among the disadvantaged. Some health outcomes are experienced disproportionately among advantaged groups. If this were the case, we might see that the most disadvantaged 85% of the population have 50% of the cumulative burden of ill health. Disparities would favor the worse off. For example, we might expect this if looking at socioeconomic differences in breast cancer incidence or melanoma.

Disproportionality measures – The Health Concentration Index (HCI)

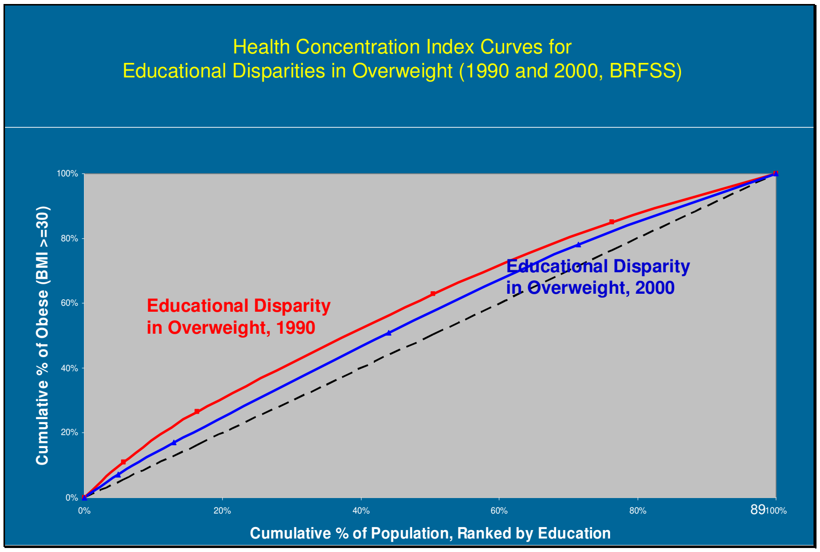

This is an example using the Health Concentration Index as a measure, based on data from the 1990 and 2000 Behavioral Risk Factor Surveillance Survey. Here we are interested in educational disparities in the proportion of the total population that is overweight.

On the X-axis is the cumulative percentage of the population as it ranked by education. On the Y-axis is the cumulative percent of obesity (a BMI greater than or equal to 30).

We can interpret that the educational disparity in obesity is smaller in 2000 as compared to 1990. In other words, we would say from this data that we have reduced the educational disparity.

Unfortunately, the reduction in educational disparity from 1990 to 2000 has occurred because all social groups are more overweight in 2000. This points to how important it is to understand that, while disparity is reduced, one still needs to understand *how* disparities are reduced to determine if the outcome is positive.

In this case, the disparity has lessened because the better educated are also becoming more obese, which is obviously not a desirable public health outcome.

Disproportionality measures – The Health Concentration Index (HCI)

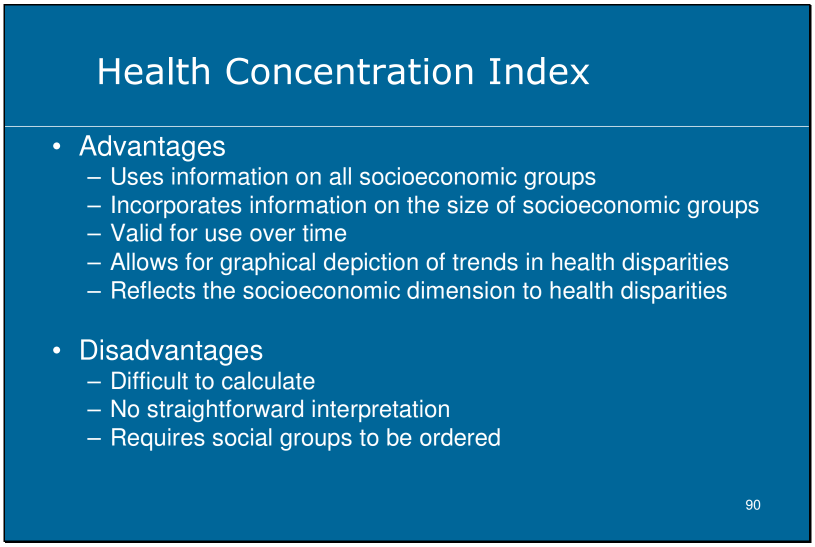

The advantages of the Health Concentration Index include the following.
Like the Gini Coefficient, it uses information on all groups and accounts for the size of the groups.

It is valid for use over time because it can account for both changes in the health measure and changes in the composition of the social groups.
It allows for graphical depiction of trends in health disparities.
Unlike the Gini Coefficient, the Health Concentration Index has the advantage of reflecting the socioeconomic dimension to health.

The HCI does have disadvantages. For example, it is somewhat more difficult to calculate and has no straightforward interpretation, as does a relative risk. Unlike the Gini Coefficient, it requires the social groups to be ordered. As a result, you cannot use a concentration index to examine geographic or race/ethnic differences where there is no natural ordering or ranking of the groups.

Disproportionality measures – Theil's Index and Mean Log Deviation

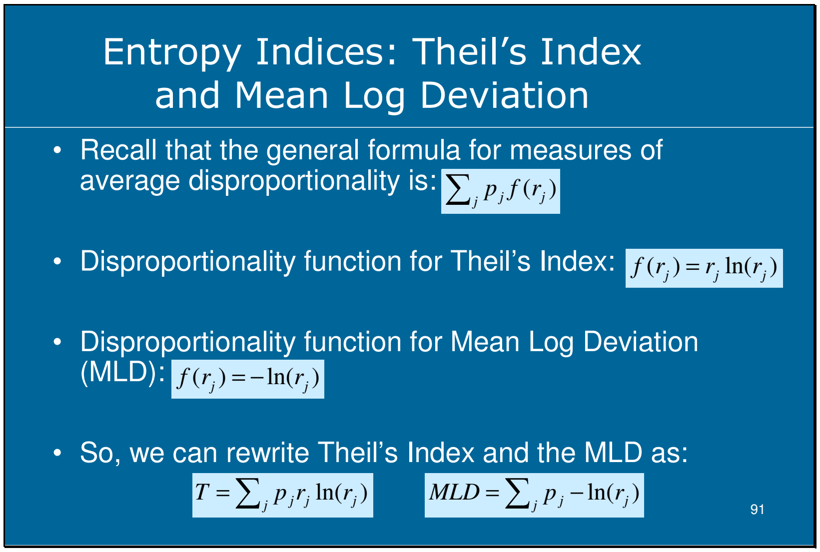

The entropy indices like Theil's Index and the Mean Log Deviation (MLD) are the most complicated measures we will discuss. However, we're not going to spend a great deal of time describing these. Examples of the more technical details of these indices are referenced in papers included in the *Resources* section.

We need measures like Theil's Index and Mean Log Deviation in disparities research so we can account for unordered groups.
These are some of the best measurement options we have when we want to have summary measures of race/ethnic disparity, for example.
These are measures that can summarize disparity over a large number of groups and do so over time in a reliable way.

Despite this, for a majority of people monitoring disparity in public health, this level of complexity may not be necessary. We present them to you for completeness.

Next, we will describe the sort of data that is used to derive Theil's Index.

Disproportionality measures – Theil's Index and Mean Log Deviation

Example of Theil's Index and the Mean Log Deviation
Applied to the Racial Disparity in Colorectal Cancer Mortality (2001)

	(a) Colorectal Cancer Mortality Rate	(b) Population (%)	(c) Rate Relative to the Total Rate	(d) Theil Index	(e) Mean Log Deviation
American Indian/ Alaska Native	10.4	0.009	0.541	–0.0029	0.0053
Asian / Pacific Islander	12.1	0.040	0.629	–0.0116	0.0184
Black	30.0	0.109	1.559	0.0751	–0.0482
White	18.3	0.843	0.950	–0.0410	0.0431
Total	19.2			***0.0198***	***0.0186***

92

This table shows rates of colorectal cancer mortality by race in the year 2001. Column (a) shows the colorectal cancer mortality rate in each race/ethnic group. You can see that 10.4 is the rate per 100,000 among American Indian and Alaskan Natives.

Column (b) shows the population proportion, which is 0.009 (or .9%) for American Indians / Alaska Natives.

Column (c) shows the colorectal cancer mortality rate in each group relative to the rate in the total population. The mortality rate for American Indian / Alaska Natives is 10.4; dividing that by the total rate, which is 19.2 yields 0.541.

Columns (d) and (e) show Theil's Index and the Mean Log Deviation respectively. These values are generated by applying the formula for average disproportionality using the disproportionality functions for Theil's Index and Mean Log Deviations. When summarized across all race groups, we get a value of .0198 for Theil's Index and .0186 for Mean Log Deviations. Note that the value is slightly higher for Theil's index; this is because it uses a slightly different disproportionality function that gives more weight to the rate differences in each group while the Mean Log Deviation accords more weight to the population size of each group.

Example. Educational disparity in mammography

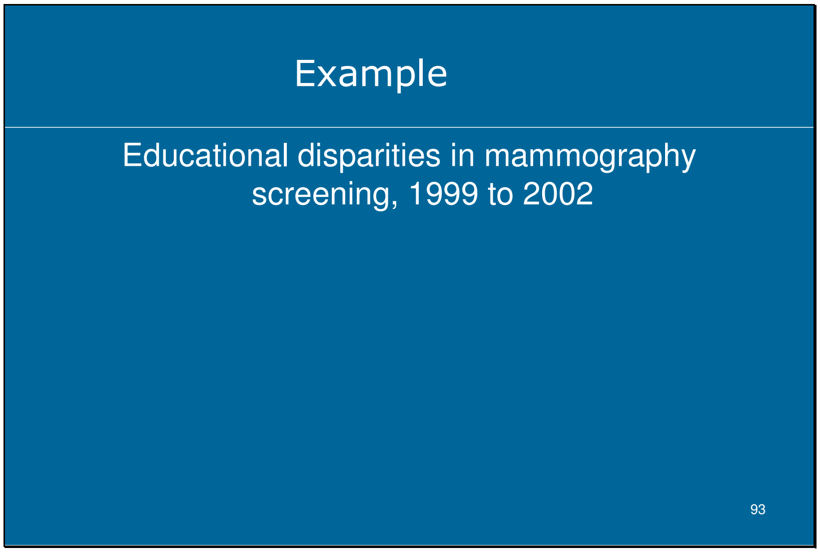

Let's look at an example using the Health Concentration Index to monitor the change in educational disparities in mammography screening from 1990 to 2002.

Example. Educational disparity in mammography

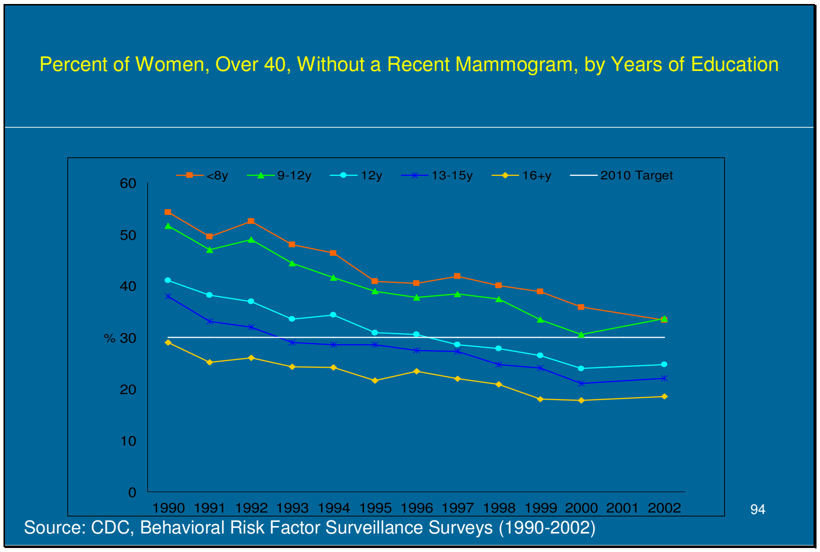

To begin, plot the underlying rates for different educational groups to get a sense of the pattern of disparity.

Here, we've plotted the percent of women over forty who haven't had a recent mammogram, grouped by years of education. The white line represents the Healthy People 2010 target rate. The way this underlying data is characterized—using group-by-group comparisons, using relative risks, using a total summary measure like concentration index, or using another summary measure—will depend on the purpose in analyzing the data. Whatever the choice, you should always plot the underlying data first to provide an idea of the problem you are investigating.

What can we conclude when we look at this data?

First, the slopes of the lines show us that the rates of lack of mammography screening are going down in all groups

The change in slopes indicates that rates are decreasing faster in recent years.

The rates seem to be going down a little faster among the least educated as compared to the more educated.

We could also conclude from the data that the absolute disparity between the highest- and the lowest-educated has reduced, as indicated by the smaller gap between the two groups in 1990 as compared to 2002.

How can we summarize this story?

Example. Educational disparity in mammography

Calculate the HCI

Education	Rate	Pop %	Cumulative Pop %	Midpoint	HCI
1990					
<8 years	54.2	0.09	0.09	0.05	4.8756
9-11 years	51.6	0.14	0.24	0.17	6.1429
12 years	41.1	0.36	0.60	0.42	8.6579
13-15 years	38.0	0.22	0.82	0.71	2.4226
16+ years	29.0	0.18	1.00	0.91	0.4762
Total	41.0				22.5752
			Health Concentration Index→		*-0.1025*
2002					
<8 years	33.3	0.06	0.06	0.03	1.8212
9-11 years	33.7	0.08	0.13	0.10	2.4003
12 years	24.7	0.34	0.47	0.30	5.7839
13-15 years	22.1	0.27	0.74	0.61	2.3366
16+ years	18.6	0.26	1.00	0.87	0.6289
Total	23.6				12.9709
			Health Concentration Index→		*-0.0998*

95

This is the data that would go into the calculation of the Health Concentration Index (HCI). It is evident from our discussion of HCI and by looking at the table data here, that education is arrayed by different groupings, mammography screening rates in each of the educational groups, the proportion of the educational groups in the population, the cumulative population proportion, the midpoint of that, and the actual calculation of the Health Concentration Index itself.

In 1990 the Health Concentration Index was -0.1025 and in 2002 it was -0.0998, suggesting that the educational disparity in mammography screening had reduced, as suggested by our initial graph.

Because the HCI is negative, we know that the disparities favor the better off. In other words, there is a greater burden of disparity among the less educated. If there was a need to come up with a number for how much the educational disparity in mammography had reduced from the 1990 levels, you could calculate the proportionate change in disparity by first subtracting 0.1025 and 0.0998, which equals 0.0027, and then dividing this by 0.1025 and multiplying by 100 which equals 2.6 %.

Example. Educational disparity in mammography

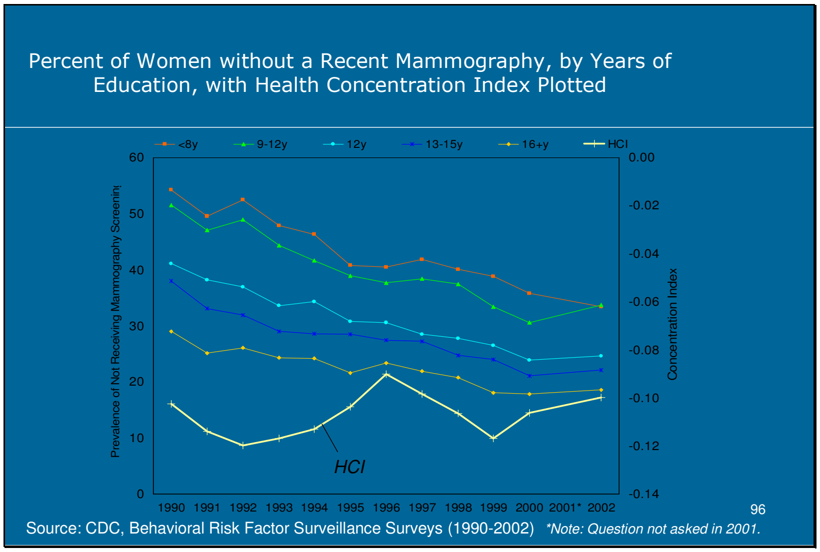

This graph again shows the underlying rates for not receiving mammography screening by the different educational groups, but it also plots the Health Concentration Index (HCI), a measure of relative disparity, over time.

You can see the increasing relative disparity from 1992 up to 1996 and then a decline to 1999.

Overall, however, there is a very small change, as indicated in that difference between -.10 (in 1990) and -.099 (in 2002) and the 2.6% reduction overall. Declines were seen in all groups such that the absolute disparity is reduced.

The combination of the graphical display of the underlying prevalence rates with some sort of summary measure like a Health Concentration Index allows for a more precise interpretation of the change in disparity.

Example. Educational disparity in mammography

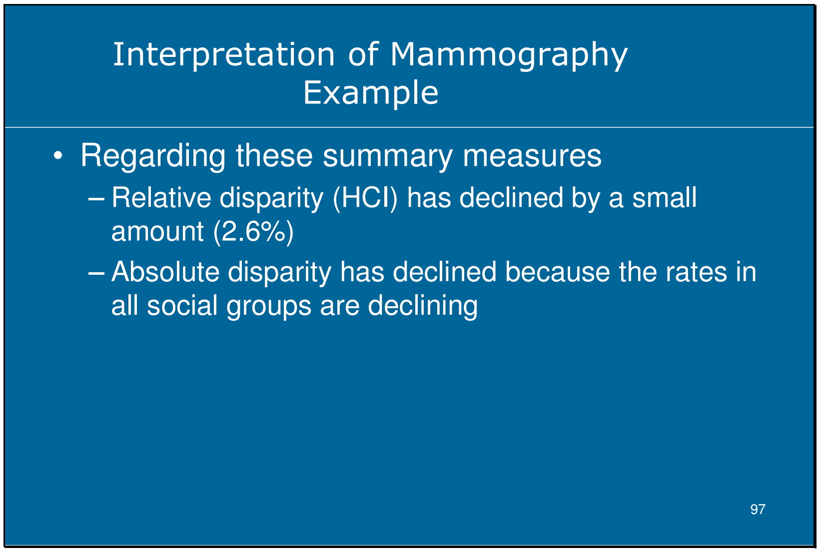

The interpretation of the disparity, then, would be that, “In regard to these summary measures, relative disparity (HCI) has remained about the same, but absolute disparity has declined because the rates in all social groups are declining.”