Journal of Econometrics 61 (1994) 161–171. North-Holland

Occupational segregation in the multidimensional case

Decomposition and tests of significance

Dale Boisso, Kathy Hayes, and Joseph Hirschberg
Southern Methodist University, Dallas, TX 75275, USA

Jacques Silber
Bar-Ilan University, 52900 Ramat-Gan, Israel

A new multidimensional version of the *G*-segregation index is developed and applied to the study of occupational segregation. U.S. Current Population Survey data are used to measure the difference in occupational segregation between races as well as the change between time periods. Decomposition of the difference (change) into 'occupation mix' and 'gender composition' components indicates the contribution of each factor. Because these inequality measures are computed from sample data, distributional information required to test hypotheses is lacking. Two computer-intensive methods for estimating the distributional properties are demonstrated. The approximate randomization and bootstrap methodologies are used to test for statistically significant differences in segregation between races and for changes over time. In addition, the components of the decomposition are examined for statistical significance.

Key words: Gini; Randomization; Bootstrap

JEL classification: J16; J62

1. Introduction

Whereas economic inequality and its measurement have long been areas of interest to both economists and sociologists, the study of employment segregation seems to have been more popular among sociologists. In fact, the most frequently used index of segregation, the Dissimilarity Index, was proposed by sociologists [Duncan and Duncan (1955)]. It is only in recent years that economists have started to notice the conceptual similarities which exist between the measurement of income inequality and that of occupational segregation [cf. Butler (1987), Silber (1989b), Hutchens (1991)]. Taking the case of

Correspondence to: Kathy Hayes, Department of Economics, Southern Methodist University, Dallas, TX 75275, USA.

segregation by gender, for example, it is by now clear that there is segregation when there exists a significant dispersion, among the various occupations, of the ratio of female to male workers. Since income inequality is related to the dispersion of incomes, it seems likely that many of the indices measuring income inequality could also be used to determine the extent of occupational segregation by gender. The most popular of these indices in the income inequality literature is the Gini Concentration Ratio and, as suggested by Butler (1987) and Silber (1989b), if carefully used, this index should prove to be very useful in studying segregation, in particular when it is formulated in a simple way [cf. Lerman and Yitzhaki (1985), Silber (1989a)]. Thus, we intend to adapt the Gini inequality index to the study of occupational segregation.

A major problem with inequality measures computed from sample data is the absence of the distributional information required to test hypotheses concerning the relative values of these measures. In this paper we demonstrate two methods for estimating the distributional properties of these statistics. The first is based on the concept of Fisher's (1935) exact test which is approximated by a randomization test [see Edgington (1969)], and the second uses Efron's (1982) bootstrap estimate of the distribution of the statistic.

The purpose of the present paper is twofold. First, a new multidimensional version of the G-segregation index is developed. Secondly, computer-intensive techniques are used to conduct significance tests on the indices. Earlier work on the G-segregation index was restricted to two-dimensional decompositions. The new index can, in theory, allow for n comparisons within a decomposition. We provide two applications of this new index, using U.S. Current Population Survey (CPS) data to measure the difference in occupational segregation between races as well as the change between time periods. We also investigate a simple decomposition of each difference into 'occupational mix' and 'gender composition' components to ascertain the contribution of each factor to differences or changes in segregation.

Because the indices and the decompositions are merely point estimates based on a particular sample, it is interesting and instructive to test the null hypotheses implicit in the construction of the indices (and the decompositions), and provide the appropriate confidence intervals. An approximate randomization method is used to construct a sampling distribution of the statistic under the null hypothesis, thereby constructing a significance test of the point estimates. The bootstrap is used to construct a distribution of the point estimates from which we can calculate confidence intervals for the indices and the decompositions.

2. The multidimensional generalization of the G-segregation index

In a recent note Silber (1989b), following earlier work on the measurement of income inequality [cf. Silber (1989a)], suggested applying the properties of

Gini's Concentration Ratio to measure occupational segregation by gender. More precisely, an alternative to the Duncans' famous dissimilarity index [cf. Duncan and Duncan (1955)] was proposed, where occupational segregation by gender would be measured by an index G, defined as

$$G_s = m' G f, \tag{1}$$

where G is a mathematical operator called the G-matrix [cf. Silber (1989a)] in which the typical element g_{ij} is equal to zero if $i = j$, to -1 if $j > i$, and to $+1$ if $i > j$, and f and m' are, respectively, column and row vectors of the shares f_i and m_i of occupation i in total female and male employment, both sets of shares being ranked by decreasing value of the ratios f_i/m_i. In this way (1) can be thought of as a 'goodness of fit' test.

Using this idea as well as earlier work on various extensions of the Duncans' index [cf. Karmel and McLachlan (1988)], Silber (1992) recently suggested a multidimensional generalization of the Duncans' index which also amounts to comparing actual with expected shares. Let T_{ijk} be the number of individuals having the characteristics i, j, and k (for example, i could refer to the occupation, j to the gender, and k to the race of the individual), and define $T_{i..}$, $T_{.j.}$, $T_{..k}$, and T as

$$T_{i..} = \sum_j \sum_k T_{ijk}, \tag{2}$$

$$T_{.j.} = \sum_i \sum_k T_{ijk}, \tag{3}$$

$$T_{..k} = \sum_i \sum_j T_{ijk}, \tag{4}$$

$$T = \sum_i \sum_j \sum_k T_{ijk}. \tag{5}$$

The ratio (T_{ijk}/T) would therefore refer to the actual proportion of individuals having the characteristics i, j, and k, whereas the product $(T_{i..}/T)(T_{.j.}/T) \times (T_{..k}/T)$ would refer to the expected proportion of such individuals, assuming independence between the characteristics i, j, and k. Whereas Silber (1992) indicated that such a comparison of actual with expected shares allowed one to define a generalized Duncan index, it is also possible to apply the same idea to derive a generalized G_s^g of the G-segregation index. This new index G_s^g will be simply defined as

$$G_s^g = [(T_{i..}/T)(T_{.j.}/T)(T_{..k}/T)]\, G[T_{ijk}/T], \tag{6}$$

where G, as in (1), is the square G-matrix, whereas the shares defining the vectors to the right and left of the G-matrix in (6) are ranked by decreasing value of the ratios $[T_{ijk}/T]/[(T_{i..}/T)(T_{.j.}/T)(T_{..k}/T)]$.

To simplify the presentation consider the case where the individual is characterized only by his or her occupation i and gender j. Let A and B be two matrices giving, respectively, the number of individuals a_{ij} and b_{ij} having characteristics i and j (where $i = 1, \ldots, I$ and $j = 1,2$) in the groups A and B. A and B could refer to two different groups within a population at a given time or to the same population during two different time periods. The occupational segregation by gender in both groups will then be measured by the indices G_s^A and G_s^B defined as

$$G_s^A = [(a_{i.}/a)(a_{.j}/a)]\, G[a_{ij}], \tag{7}$$

$$G_s^B = [(b_{i.}/b)(b_{.j}/b)]\, G[b_{ij}], \tag{8}$$

where $a = \sum_i \sum_j a_{ij}$ and $b = \sum_i \sum_j b_{ij}$.

It will be shown that the difference $G_s^B - G_s^A$ may be expressed as the sum of two components, the first one related only to differences in the sex ratios in the various occupations, the second one being only the consequence of differences in the weights of the various occupations.

Define two hypothetical segregation indices G_s^{BA} and G_s^{AB} as

$$G_s^{BA} = \left[(b_{i.}/b) \left(\sum_i (a_{ij}/a_{i.})(b_{i.}/b) \right) \right] G[(a_{ij}/a_{i.})(b_{i.}/b)], \tag{9}$$

$$G_s^{AB} = \left[(a_{i.}/a) \left(\sum_i (b_{ij}/b_{i.})(a_{i.}/a) \right) \right] G[(b_{ij}/b_{i.})(a_{i.}/a)], \tag{10}$$

where G_s^{BA} is the value of the G index for a hypothetical group with the occupational distribution of group B and the gender distribution of group A, and G_s^{AB} is that for a hypothetical group with the occupational distribution of group A and the gender distribution of group B. The component related to differences in the gender ratio may then be written as

$$\Delta G_{\text{gender}} = (0.5)[(G_s^B - G_s^{BA}) + (G_s^{AB} - G_s^A)], \tag{11}$$

whereas the component related to differences in the shares of the various occupations will be equal to

$$\Delta G_{\text{occupation}} = (0.5)[(G_s^B - G_s^{AB}) + (G_s^{BA} - G_s^A)], \tag{12}$$

and as expected we have

$$G_s^B - G_s^A = \Delta G_{\text{gender}} + \Delta G_{\text{occupation}}. \tag{13}$$

Such a decomposition will be applied in section 4, firstly to a comparison of the occupational segregation by gender among white and nonwhite workers in the U.S. in 1990, secondly to the study of the changes which took place between 1985 and 1990 in occupational segregation by gender in the total U.S. labor force.

3. Hypothesis tests and confidence bounds

The Gini coefficient as well as the Gini-based G-segregation index are mere point estimates. Obviously, eq. (1) used to define the G-segregation index is a highly nonlinear function of the occupation, race, and gender specific employment levels. This fact makes an analytical solution for the variance of the point estimate intractable. The Gini-type indices are best evaluated by the use of a nondistribution-based resampling method because of the complex nature of the translation of the distribution of the data to the distribution of the index. Because of this, a number of authors have recommended the use of resampling methods to evaluate a Gini coefficient. Sanderstrom, Wretman, and Waldon (1988) demonstrate the ability of the jackknife to estimate variances. Dixon, Weiner, Mitchell-Olds, and Woodley (1987) discuss the ability of the bootstrap to estimate the confidence intervals. In this section we propose the use of two methods for making inferences from the estimated indices, one based on the bootstrap and the other on an approximate randomization test.

One method for testing the difference between point estimates of the indices and decomposition is the bootstrap as proposed by Efron (1979). The bootstrap provides a method for estimating the distributions of a statistic by resampling with replacement from the data set for each occupation to create multiple estimates of the G-segregation index and the decomposition. These multiple values of the same statistic generate a distribution of point estimates for each measure. These distributions can then be examined in order to establish a probability that the statistic's value will include the value implied under the null hypothesis (as we note below this may not be zero); in essence the bootstrap allows us to construct confidence intervals around the original point estimate. Alternatively, an approximate randomization technique can provide a distribution around the G-segregation index implied by the null hypothesis of no segregation. This distribution is used in the same manner as the traditional student's t-statistic to establish a probability that the measured G-segregation value was generated by a process that is centered at the null. This distribution can be examined for the inclusion of the point estimate. The remainder of this section provides further discussion of the two computer-intensive methods, the bootstrap and the randomization method, as they apply to the segregation analyses.

The bootstrap technique for the construction of confidence intervals has recently been the subject of a number of papers since Efron's original paper in 1979. A general overview is given in Efron and Gong (1983) and regression-based econometric applications are discussed in Freedman and Peters (1984a, b). The construction of the bootstrap distributions is done using a technique similar to that used in Schwartzbaum and Hirschberg (1991), where instead of resampling directly from the original data the gender and race proportions in each occupation are computed by constructing new pseudo samples of proportions in each occupation based on draws from a random number generator for a binomial distribution. The random numbers are calculated for each case using the number of persons observed and the proportion in each gender/race/occupation cell. The bootstrap is performed by reestimating the segregation index and the corresponding decompositions for each pseudo sample. These values are then used to construct an empirical statistical distribution of these statistics using Efron's (1982) percentage method.

Secondly we employ an approximate randomization test [see Noreen (1989) for an introduction to this technique]. This procedure is similar to the bootstrap, except that instead of obtaining the distribution of the statistic estimated from the observed sample it is estimated as if the sample had been generated by a process in which the null hypothesis (in the present case of no segregation) is true. This procedure also results in an estimated distribution. This method is related to the often employed test which is constructed by estimating a variance and assuming a distribution (student-t or normal) and establishing the 'p-value' of the estimated value under the null hypothesis that the value is zero. Theoretically this method for estimating the distribution under the null hypothesis is exact if we compute all the possible values of the statistic under the null hypothesis [this would be an application of Fisher's exact test – see Edgington (1969)]. However, in the present case this would involve far too many combinations to be practical. Sampling from the binomial distributions consistent with the null hypothesis thus generates a set of pseudo samples or possible real worlds that may be observed under the null hypothesis. The power and the validity of approximation randomization tests is demonstrated by Noreen (1989) and is shown generally by Hoeffding (1952).

4. An application

To conduct our analysis on segregation in the U.S. work force we used data from the Current Population Survey (U.S. Department of Commerce, Bureau of the Census) for March 1985 and March 1990. The observations within this data set are derived from monthly interviews with 57,000 households regarding, among other topics, employment, income, and demographic status. The observations were compiled into a classification of 51 two-digit occupations, where occupation was defined to be the longest-held job in the interviewee's past. For 1985, 48,793 observations were used in the analysis. The 1990 sample contained 53,746 observations.

The first application concerned occupational segregation between races based on the 1990 data. The focus of our attention was whether there existed evidence that race was a determining factor in the observed number of persons in each occupation. The segregation indices were calculated in accord with the procedure described in section 2, whereby the observations were first separated into two groups, one containing persons reporting their race as White and the second, Nonwhite, capturing all others.

The results appear in table 1 under heading True Mean. We found very little difference in the Gini segregation indices for the two categories White and Nonwhite, though this is not to say that there is no segregation, as the Gini's for White (0.3080) and for Nonwhite (0.3076) clearly indicate otherwise. The ΔG_{ii} of 0.00035 represents the difference between Gini indices for White and Nonwhite and can be separated into segregation resulting from a change in occupational mix (0.0025) and in the gender mix (− 0.0021). In words, it appears that the Nonwhite group had a higher level of occupational segregation but a lower level of gender segregation than the White group.

In our second application the 1985 and 1990 data were used to consider whether there was a change in the amount of occupational segregation over time. For computation of the Gini indices the observations were grouped

Table 1
Occupational and gender segregation between whites and nonwhites in 1990.

	Randomization[a]			Bootstrap		
	1%	Mean	99%	1%	True mean	99%
G_{22} (1990, nonwhite)	0.0343	0.0457	0.0585	0.3007	0.3080	0.3208
G_{11} (1990, white)	0.0126	0.0170	0.0214	0.3035	0.3076	0.3119
ΔG_{ii}	0.0174	0.0287	0.0421	− 0.0091	0.00035	0.0133
$\Delta G_{occupation}$	− 0.0075	− 0.0023	0.0001	− 0.0003	0.0025	0.0047
ΔG_{gender}	0.0179	0.0311	0.0461	− 0.0104	− 0.0021	0.0111

[a]Null is no segregation.

Table 2
Changes in occupational and gender segregation between 1985 and 1990.

	Randomization[a]			Bootstrap		
	1%	Mean	99%	1%	True mean	99%
G_{22} (1990)	0.0116	0.0160	0.0202	0.3049	0.3086	0.3125
G_{11} (1985)	0.0125	0.0166	0.0207	0.3217	0.3254	0.3298
ΔG_{ii}	− 0.0066	− 0.0006	0.0052	− 0.0224	− 0.0168	− 0.0116
$\Delta G_{\text{occupation}}$	− 0.0005	− 2.7E-5	0.0005	− 0.0087	− 0.0083	− 0.0079
ΔG_{gender}	− 0.0067	− 0.0006	0.0053	− 0.0142	− 0.0085	− 0.0034

[a]Null is no segregation.

according to their year, such that now time, not race, was the determining factor of interest.

The results are shown in table 2, again under the heading True Mean. In this case, the Gini coefficient for 1985 was 0.3254, whereas that for 1990 was 0.3086, indicating the change in the Gini measure of segregation, ΔG_{ii}, over time was − 0.0168. That is, there was a decrease in segregation between 1985 and 1990, and this adjustment resulted from changes in both the occupational mix and gender composition. Segregation due to occupational structure decreased by 0.0083, while that explained by gender composition decreased by 0.0085.

However, these results represent point estimates of the level of segregation and, as such, each is just one bit of information needed to make statements about whether segregation exists or not. We must conduct tests of significance and/or construct confidence intervals for these Gini coefficients. As discussed above, two approaches are taken to test for statistical significance of the indices and the decompositions.

Recalling the results of table 1 in which the segregation of Whites and Nonwhites is analysed for 1990, there does appear to be segregation in employment for each race, but the difference between races is miniscule. The decomposition results suggest that occupational segregation may be more important in determining segregation for nonwhite workers, while gender segregation is more important in determining segregation for white workers. The question of course is to what extent these conclusions are correct based on our randomization and bootstrapped results. One might guess that the segregation results are indeed significant, i.e., G_{11} and G_{22} are different from the null. However, the difference is so small that one might guess that segregation differences between races are not statistically different from each other.

The bootstrap and randomization results are presented in tables 1 and 2. Included in the tables are the 1st and 99th percentiles of the randomized and bootstrapped distributions, as well as the mean of the approximate randomization. Using this information the True Mean can be compared to the significance bounds created by the randomization procedure to see if the original point estimate of each measure is statistically different from the null hypothesis. Then, the mean of the randomization procedure, i.e., the null hypothesis, will be compared to the distribution of the bootstrap to determine if the null lies within the confidence bounds created by the bootstrap procedure.

Theoretically one would expect the null hypothesis to be zero as the Gini index, in theory, varies from zero to one. An examination of the results in tables 1 and 2 suggests that the mean of the approximate randomization in fact is not zero. The data used to construct this distribution were created under the assumption of no segregation; however, because the number of individuals in each occupation varies, the 'true' null is different from zero. This is particularly important when conducting the bootstrap analysis. Had the randomization distributions not been completed, we would have assumed incorrectly that the appropriate null was zero.

First, in comparing the point estimates to the distribution created by the randomization technique we note that the True Mean of G_{22} and G_{11} lies outside the randomization distribution. The same is true for the difference in the G_{ii}, ΔG_{ii}, and for each of the decompositions. Second, the randomization mean is compared to the bootstrapped distribution, to determine if the null lies within the confidence bands of the bootstrapped values. For both G_{11} and G_{22} the randomized mean is clearly less than the 99 percent confidence band. Similarly, the randomized means for ΔG_{ii} and the decompositions do not lie within the confidence bound created by the bootstrapped procedure.

Our first conclusion, that there is segregation by gender in the 1990 distribution, is verified by the two methodologies. Somewhat surprisingly, we can conclude that there is a difference in degree of segregation between the races and that occupational segregation is more important in generating segregation for Nonwhites and gender segregation is more important in generating segregation for Whites.

Turning to table 2, we can examine the change in segregation between 1985 and 1990. For both years the point estimate of segregation is quite large and similar in magnitude to the results in table 1. The change in the segregation indices is small relative to the size of index, but suggests the segregation may have fallen over the time period. Based on the decompositions we would attribute the decline about equally to a reduction in occupational segregation and gender segregation.

Comparing these results to the distribution created by the randomization procedure indicates that there was statistically significant segregation in both 1985 and 1990. The point estimates of both $\Delta G_{\text{occupation}}$ and ΔG_{gender} lie outside the 1–99 percentile range of the randomized distribution. Although the decline in segregation appears small, it is statistically different from the null hypothesis as it lies outside the distribution created by the randomization. Furthermore, we can conclude that declines in both occupational and gender segregation are responsible for the statistically significant decline in overall segregation. Both $\Delta G_{\text{occupation}}$ and ΔG_{gender} point estimates lie outside the randomized distribution.

5. Conclusions

In the present study an extension of the use of the Gini Concentration Ratio has been proposed which allows one to measure occupational segregation by gender. This generalization was then applied to analyse differences between races and changes over time in the level of occupational segregation and to derive a decomposition of these variations into two components, which reflect differences in the occupational mix and in the gender composition.

Noreen's computer-intensive techniques were employed to generate sampling distributions via approximate randomization to test hypotheses about the indices and using bootstrap sampling to derive confidence intervals for the G-segregation measures. The results of this statistical analysis indicated that the differences between races and betweeen periods in the level of occupational segregation by gender were small, but statistically significant. Even the two components of the decomposition of the difference in segregation levels were found to be significant in a statistical sense.

References

Butler, R.J., 1987, New indices of segregation, Economics Letters 24, 359–362.

Dixon, Philip M., Jacob Weiner, Thomas Mitchell-Olds, and Robert Woodley, 1987, Bootstrapping the Gini coefficient of inequality, Ecology 68, 1548–1551.

Duncan, O.D. and B. Duncan, 1955, A methodological analysis of segregation indices, American Sociological Review 20, 210–217.

Edgington, Eugene S., 1969, Approximate randomization tests, Journal of Psychology 72, 143–149.

Efron, Bradley, 1979, Bootstrap methods: Another look at the jackknife, Annals of Statistics 7, 1–26.

Efron, Bradley, 1982, The jackknife, the bootstrap and other resampling plans (Society for Industry and Mathematics, Philadelphia, PA).

Efron, Bradley and Gail Gong, 1983, A leisurely look at the bootstrap, the jackknife, and crossvalidation, American Statistician 37, 36–48.

Fisher, R., 1935, The fiducial argument in statistical inference, Annals of Eugenics 6, 391–398.

Freedman, David A. and Stephen C. Peters, 1984a, Bootstrapping a regression equation: Some empirical results, Journal of the American Statistical Association 79, 97–106.

Freedman, David A. and Stephen C. Peters, 1984b, Bootstrapping an econometric model: Some empirical results, Journal of Business and Economic Statistics 2, 150–158.

Hoeffding, W., 1952, The large-sample power of tests based on permutations of observations, Annals of Mathematical Statistics 23, 169–192.

Hutchens, R.M., 1991, Segregation curves, Lorenz curves and inequality in the distribution of people across occupations, Mathematical Social Sciences, 31–51.

Karmel, T. and M. McLachlan, 1988, Occupational segregation – Increasing or decreasing?, Economic Record 64, 187–198.

Lerman, R.I. and S. Yitzahki, 1985, Income inequality effects by income source: A new approach and applications to the United States, Review of Economics and Statistics 67, 151–156.

Noreen, Eric W., 1989, Computer intensive methods for testing hypotheses (Wiley, New York, NY).

Sandstrom, A., J.H. Wretman, and Bertil Walden, 1988, Variance estimators of the Gini coefficient – Probability sampling, Journal of Business and Economic Statistics 6, 113–119.

Schwartzbaum, J.A. and J.G. Hirschberg, 1991, Exploring the order of odds ratios using the bootstrap, Statistics in Medicine 10, 1273–1281.

Silber, J., 1989a, Factor components, population subgroups and the computation of the Gini index of inequality, Review of Economics and Statistics 71, 107–115.

Silber, J., 1989b, On the measurement of employment segregation, Economics Letters 30, 237–243.

Silber, J., 1992, Occupational segregation indices in the multidimensional case: A note, Economic Record 68, 276-277.