DAVID E. BLOOM
AND MARK R. KILLINGSWORTH*

Pay Discrimination Research and Litigation: The Use of Regression

THE USE OF REGRESSION ANALYSIS in empirical studies of pay discrimination within organizations and in court cases is a relatively recent phenomenon. Until the late sixties and early seventies, researchers interested in race and sex discrimination in pay had little in the way of established analytical constructs for investigating these issues. Even if such analytical constructs had been fully developed, researchers interested in empirical questions of this kind would still have had to contend with another problem: the limited availability of computer software suited to the kind of large-size data set (e.g., Census data, data on individual employees at particular firms) that would normally be used in such analyses.

As regards litigation, for a considerable period after passage of the Civil Rights Act in 1964, employment discrimination cases were largely concerned with recruitment and hiring rather than with compensation, and frequently focused on persons in relatively unskilled jobs. Detailed empirical analysis was not generally undertaken (see *Harvard Law Review* [1971] for a discussion of developments in discrimination litigation between 1964 and 1971; and the review in Gwartney *et al.*, 1979). Whether or not regression analysis might have been useful in such contexts, it does not appear to have been used.

Times have changed. Many studies (especially Mincer's [1974] seminal work, which circulated widely in manuscript form prior to publication) have provided solid analytical foundations for empirical analyses of earnings; and development of powerful computer software—such as the Statistical Package for the Social Sciences (Nie *et al.*, 1970) and the Statistical Analysis System (Service, 1972)—now permits analysis of larger data sets at much lower cost and greater speed than was previously possible. At the same time, the kinds of problems and issues that arise in litigation under employment discrimination laws have become more complex, making the use of

*The authors are, respectively, Assistant Professor of Economics, Harvard University, and Assistant Professor of Economics, Rutgers—The State University of New Jersey.

INDUSTRIAL RELATIONS, Vol. 21, No. 3, (Fall, 1982). © 1982 by the Regents of the University of California.
0019/8676/82/1015/318/$1.00

Exhibit P-452

detailed statistical analysis seem more appropriate (e.g., see Judge W. H. Orrick's comments in *Boyd v. Bechtel Corporation,* 1979, pp. 612-613).

Our purpose in this article is to review the use of regression analysis in pay discrimination studies and in litigation.[1] We begin by describing the rise of regression analysis to its current prominence and then outline the basic notions underlying its use. Next, we consider studies of restricted samples, and note reasons why one might want to apply regression analysis to a restricted sample rather than to an organization's total workforce. We then examine conceptual and statistical problems that may arise in analyses of restricted samples. After showing how one may extend conventional regression methodology to derive appropriate (statistically consistent) measures of discrimination in such contexts, we conclude by discussing questions for future research.

Regression Analysis and the Courts: An Historical Overview

The earliest use of regression analysis in Federal employment discrimination litigation appears to be the 1973 court case, *United States v. U.S. Steel.*[2] Ironically, however, neither the district court's decision, nor a subsequent appellate court judgment (1975) actually discussed the regression studies presented; it was a discussion of the case in a professional journal by two of the "expert witnesses" involved (Haworth and Haworth, 1976) that made it clear that regression had been used.

The next relevant case (again, Federal) appears to be *Wade v. Mississippi Cooperative Extension Service* (1974). Here, as in *United States v. U.S. Steel,* the district court's decision did not mention regression analysis. However, the defendant (the Mississippi Cooperative Extension Service) appealed the district court's decision and, in its appeal, attacked the use of regression analysis by the plaintiffs. The appellate court's comments in its decision (1976, p. 517) are revealing:

Appellants essentially challenge the probative value of some of the statistical evidence introduced by plaintiffs, because of the sophisticated nature of the methods used to analyze the statistical data.... Although multi-variate regression analysis is

[1]Without implying that they necessarily share any of the opinions expressed here, we thank Orley Ashenfelter, George Duncan, Michael O. Finkelstein, Gilles Grenier, Charles C. Killingsworth, Vivienne Killingsworth, George Milkovich, Peter Nickerson, Cordelia W. Reimers, Elizabeth Scott, Finis R. Welch, and Donald Wise for helpful discussions, comments, and suggestions. We owe a special debt of gratitude to Paul Axel-Lute, who provided us with an extensive list of Federal court discrimination cases in which regression analysis was used or discussed; and to Floyd Hale and Alane Probst, who helped us track down copies of decisions in some of these cases.

[2]Tracing the history of the use of regression analysis in legal proceedings is hampered by the fact that not all court decisions explicitly discuss the regression studies that were presented by the contending parties. Also, not all cases culminate in a judicial decision; many are settled out of court. For further discussion of regression analysis in the context of legal proceedings involving charges of employment discrimination, see Baldus and Cole (1980); Fisher (1980); Finkelstein (1980); Feinberg (1982); and Feinberg and Straf (1982).

indeed a sophisticated and difficult method of proof in an employment discrimination case, there was additional evidence of specific instances of black and white workers with essentially similar experience and qualifications receiving disparate salaries. Thus, we find that while in some cases the statistical facts spoke for themselves, as in the absence of promotions of black professional workers, in other cases, there was evidence beyond the statistical facts and analysis that would support an inference of discrimination, as in the case of salaries.

While not exactly a ringing endorsement of regression analysis, the opinion seems to allow at least some scope for its use in discrimination cases. Indeed, the appellate court cited a lengthy note published in the *Harvard Law Review* (1975) the previous year, which discussed the potential uses of regression analysis in such cases.

Over time, the trickle of discrimination cases involving regression analysis has become a steady flow. Not infrequently, judges confronted with the complexities of regression have echoed or even emphasized some of the uncertainties and equivocal feelings that seem to have been just under the surface of the *Wade* decision. For example, in *Kyriazi v. Western Electric Company* (1973, p. 914), the judge remarked that the regression analysis presented in the testimony of the plaintiff's expert witness was "simply not comprehensible to the Court." Another case, *Wilkins v. University of Houston* (1979) followed the pattern of *Wade:* no mention of regression analysis in the decision of the district court, followed by fairly lengthy discussion in subsequent appellate decisions (1981a, 1981b), which contained a number of cautionary and even "rueful" (1981a, p. 1248) remarks about the complexity of regression analysis.

On the whole, however, the legal profession's growing familiarity with regression analysis seems to have bred acceptance rather than contempt and, as Finkelstein (1980, p. 737) has remarked, "The idea caught on with remarkable rapidity." In *Patterson v. Western Development Laboratories* (1976, p. 777), a judge criticized the plaintiffs for *failing* to undertake a regression analysis. Gradually, the lines of battle have shifted from arguments about whether regression analysis is at all appropriate in a legal setting to arguments about its proper implementation in the particular situation under consideration. For example, the appellate court decision in *James v. Stockham Valves* (1977), while apparently accepting the general notion of regression analysis, reversed the district court's decision (1975) in the case, and criticized regression studies presented by the defendant on the grounds that such studies used inappropriate variables. Other cases, such as *Presseisen v. Swarthmore* (1977), *Mecklenburg v. Montana Board of Regents of Higher Education* (1976), *Greenspan v. Auto Club of Michigan* (1980), and, most notably, the lengthy decision in *Vuyanich v. Republic National Bank* (1980),

have taken a similar approach: criticisms of regression studies have been concerned with the application of the method rather than with the method itself.

The Methodology of Regression: General Principles

In a regression analysis of pay discrimination, the purpose is to determine whether a sex (or race, etc.) difference in pay persists when one factors out the possible effects of other variables. Regression analysis of pay differences is complex for the same reason that it is powerful: its implementation entails the application of economic and statistical theory to the specification and estimation of economic relationships under appropriate rules of statistical inference.

The foundation of a regression analysis is the notion of a model which sets out the nature of the relationship under study. In the present context, this is called an *earnings function.* The earnings function used in studying questions about discrimination in pay at a particular organization typically takes the general form:

$$(1) \quad Y_i = b_{Y0} + b_{Y1}X_{1i} + \ldots + b_{Yk}X_{ki} + d_{Y1}D_{1i} + \ldots + d_{Yn}D_{ni} + u_{Yi}$$

where Y_i is a measure of individual i's pay (e.g., annual earnings or hourly wage rate, in dollars or units of natural logarithms); X_{ji} is individual i's amount of a measured productivity-related characteristic X_j (e.g., X_1 might represent the number of years of schooling individual i has completed, X_2 might represent individual i's years of service with the company, etc.); D_{ji} is individual i's value of some measured demographic characteristic D_j (e.g., D_1 might denote individual i's sex, D_2 might denote individual i's race, etc.); and u_{Yi} (the "error term") represents unmeasured factors that are associated with individual i's pay. Finally, b_{Y1} is the coefficient—or parameter—on the first X variable, X_1, while d_{Y1} is the coefficient on the first D variable, D_1 (and similarly for b_{Y2}, b_{Y3}, etc., and d_{Y2}, d_{Y3}, etc.). In effect, these coefficients or parameters are weights that indicate the amount of salary increase or decrease associated with changes in the variables they accompany, other things being equal. (The parameter b_{Y0} in effect represents the salary that is paid, on average, to persons whose values for all of the X and D variables are zero—a kind of baseline salary.)

Discrimination against persons possessing a particular demographic characteristic D_j is said to exist if the regression estimate of the coefficient d_{Yj} on that characteristic is nonzero (in the statistical sense) and large (in the ordinary language sense). For example, suppose that the variable D_1 denotes "sex is female" (and takes a value of unity for all individuals who are

female, and a value of zero for all individuals who are male); and suppose further that the regression estimate of the coefficient d_{Y1} in the model (1) is −1000. Then the regression results imply that being female rather than an otherwise-comparable male (that is, changing the value of D_1 from 0 to 1) is systematically associated with a reduction in salary of $1,000, on average. In this context, otherwise comparable means both (i) possessing the same values for all measured productivity related characteristics X_1 through X_k inclusive, and (ii) comparable in terms of all measured demographic characteristics D_2 through D_n other than sex (e.g., race, ethnicity, etc.).

In order for the regression estimate of a coefficient b_{Yj} (or d_{Yj}) not to be statistically biased, the measured characteristic X_j (or D_j) with which that coefficient is associated must be uncorrelated with the error term u_Y, other things being equal, where the "other things" in question are the *other* measured characteristics X and D (see Kmenta [1971], pp. 394-395, for elaboration). It is not hard to see why regression analysis will produce statistically biased results when this condition is not satisfied. For example, consider the use of regression analysis to determine whether the measured X variable, "years of school completed," has an effect on compensation, *ceteris paribus*. Suppose that there are no data available on employees' total years of work experience, so that this is one of the unmeasured characteristics that are comprehended within the error term u_Y. Suppose further that people who have more schooling also tend to have less work experience than do people with less schooling, other things being equal. Finally, suppose that both schooling and work experience do, in fact, have a positive effect on compensation.

Under these assumptions, the regression error term u_Y (which includes work experience) will be negatively correlated with years of schooling and positively correlated with compensation, *ceteris paribus*. It follows that a regression analysis may understate the impact of schooling per se on compensation: some of the apparent difference in pay between persons with more schooling and persons with less schooling is really attributable to differences in experience—but these experience-related differences in pay are erroneously attributed to differences in schooling rather than to differences in work experience because schooling is included explicitly in the earnings function while work experience is not.

This does *not* mean, however, that regression analysis will produce biased results simply because there are "unmeasured variables." The fact that a regression analysis does not include among the measured X and D variables all or even most of the factors that might have an important effect on pay is not sufficient to warrant the conclusion that the analysis will produce

biased results. (A number of court decisions seem to be in error in this regard; for example, see *Presseisen v. Swarthmore College* [1977]; *Eastland v. Freeman* [1981]; and *Valentino v. U.S. Postal Service* [1981].) In particular, the fact that an unmeasured variable that affects compensation is correlated with a measured variable (either a measured productivity related characteristic X or a measured demographic characteristic D) does not necessarily mean that failure to measure that unmeasured variable leads to biased results. Rather, bias arises only if the unmeasured variable is correlated both with compensation and with a measured variable *at the margin,* i.e., when all other measured variables are held constant. The only way to assess the magnitude of the bias is to measure the omitted variable, include it in a regression, and see what its inclusion does to the estimated pay differential (in this example, the sex differential) of interest. These or similar considerations seem to underlie a number of court decisions that have rejected claims that omission of variables necessarily invalidates regression studies, when such claims were not accompanied by an empirical demonstration of the magnitude of the alleged statistical bias involved (e.g., *Commonwealth of Pennsylvania v. Local Union 542, IUOE* [1978], and especially *Segar v. Civiletti* [1981] and *Trout v. Hidalgo* [1981]).

Most studies (and at least one court case: *Rosario et al. v. New York Times* [1979]) of discimination in pay at given organizations have implicitly or explicitly assumed that, in the organization's workforce as a whole, the error term u_Y in a regression model such as the one given by equation (1) is uncorrelated with the particular D variable of interest (e.g., race or sex), other things being equal. We discuss this assumption further on; here we temporarily adopt it as a maintained hypothesis and turn our attention to restricted samples. We first consider why one might want to study restricted samples and then focus on two issues, one conceptual and the other statistical. The conceptual issue is concerned with how regression results derived from restricted samples may be interpreted. The statistical issue is about the problem of bias: if one adopts the maintained hypothesis that u_Y is uncorrelated with the particular D variable of interest in an employer's total workforce, does this necessarily imply that the same is true of restricted samples or subgroups within that total workforce?

Restricted Samples

The literature on discrimination contains numerous examples of regression studies of "restricted samples"—subsets of employees selected

from an organization's total workforce on some basis other than simple random sampling.[3]

Often, restricted samples refer to sets of employees selected on the basis of either job type or job level. For example, studies of salary discrimination in universities not infrequently are concerned with analyzing samples that are restricted to persons who hold faculty (or, alternatively, administrative or staff) job titles; studies of for-profit organizations are sometimes restricted to white-collar or blue-collar employees; and at least one set of studies prepared for use in litigation under antidiscrimination laws has considered professional and nonprofessional employees separately (see *Vuyanich v. Republic National Bank* [1980]). In each case, the restricted sample is selected on the basis of job type. When formulating a sample on the basis of job level, the selection focuses on a particular position in a given organization's job hierarchy. For example, in studying a university, one might restrict one's analysis to persons holding the title of full professor (or associate professor, etc.); in studying a for-profit organization, one might restrict one's analysis to persons in a specific salary grade or job grade.

There are at least three reasons, not necessarily equally compelling, for analyzing restricted samples: a desire for homogeneity; a concern about noncompeting groups; and an interest in within-group as opposed to workplace-wide discrimination.

Homogeneity. The range of skills, attainments, and characteristics of the employees in a given organization frequently approximates that of the labor force as a whole. Restricting one's analysis to specific groups of employees (e.g., faculty vs. clerical workers) will often reduce substantially the range of variation in skill and training and, thus, result in a considerably more homogeneous set of individuals. The quest for greater homogeneity rests on two (not necessarily correct) premises: first, that regression analysis is not

[3]See Malkiel and Malkiel (1973); Gordon, Morton, and Braden (1974); Hoffman (1975); Ashenfelter and Pencavel (1976); Gordon and Morton (1976); Borjas (1978); Osterman (1979); the discussion of studies presented by plaintiffs in *Segar v. Civiletti* (1981) and *Trout v. Hidalgo* (1981); and the discussion of studies presented by defendants in *Mecklenburg v. Montana Board of Regents of Higher Education* (1976), *Valentino v. U.S. Postal Service* (1981; 1982), *Greenspan et al. v. Auto Club of Michigan* (1980), and *Vuyanich v. Republic National Bank* (1980). For studies that analyze "unrestricted" samples (of all employees, or samples randomly selected from all employees, in given organizations), see Oaxaca (1976); Smith (1977); and the discussion in *Greenspan et al. v. Auto Club of Michigan* (1980) of studies presented by the plaintiffs in that litigation.

Our remarks are specifically addressed to studies that are based on samples that are restricted in some way, particularly on the basis of job type or job level. However, the discussion applies almost in its entirety to a slightly different way of controlling for job type or job level, i.e., simply including one or more variables that indicate job type or job level in a study of an organization's entire workforce. (For an example of this kind of study, performed for purposes of litigation, see the discussion of regression analyses presented by the defendant in *James v. Stockham Valves*, 1975 and 1977.) Both sample restriction and inclusion of indicator variables are alternative ways of controlling for job type or job level; the latter procedure is simply somewhat less general than the former, since it constrains the coefficients on *other* variables (e.g., race, schooling, age, work experience), and their standard errors, to be the same regardless of job type or job level, while the former does not.

well-suited to consideration of a set of "heterogeneous" individuals; and, second, that homogeneity as such is a desideratum for regression analysis.

Noncompeting groups. A second reason for studying restricted samples has to do with the notion of noncompeting groups in the labor market and with the possibility that regression analysis of unrestricted samples may produce misleading results about the presence or absence of discrimination. The essential idea here is that differences among workers in preferences for nonpecuniary aspects of employment, and/or differences in workers' qualifications, may generate sizeable sex, race, or ethnic differentials in compensation, even if no discrimination exists.

For example, suppose that, because of past socialization, differences in tastes, etc., women tend to acquire skills in areas such as the humanities, nursing, library science, and education to a greater extent than men; and suppose further that market conditions in such fields are distinctly less favorable than they are in the natural and social sciences, which have attracted a disproportionate number of men. This will mean that women will tend to be paid less than men, on average, even if women possessing any given type of skill are paid on the same basis as men with the same type of skill.[4]

Similarly, differences in qualifications, prior experience, etc., may lead to substantial disparities in pay between men and women, whites and blacks, etc., even in the absence of discrimination. For example, an analysis of the compensation received by an airline company's personnel—pilots, flight attendants, ground crew, clerical workers—might give the appearance of sizeable discriminatory differentials in pay, even if the pay disparities were actually due to differences in qualifications.

In sum, not all pay differences are discriminatory; they may also arise from differences in preferences and/or in skills: individuals do not always *want* to compete with other people, and individuals are not always *able* to compete with other people. Analysis of restricted samples is sometimes regarded as a means of reducing the potential for ascribing salary differences between the sexes (for example) to discrimination when the differentials actually arise for other reasons (see Roberts, 1980, p. 193).[5]

[4]For example, see the appellate court's decision in *Valentino v. U.S. Postal Service* (1982), which puts particular emphasis on type of skill. The district court's decision in *Wilkins v. University of Houston* (1979) asserts "that *any* variations in salary among faculty men and women were due to market place values imposed upon the various schools of the University" (p. 1056; emphasis added)—which surely seems a rather extreme position.

[5]Of course, restricting analyses to particular job categories does not necessarily resolve all problems. For example, in the airline example given in the text, one might also want to analyze whether there are artificial restraints that prevent qualified women who want to work as pilots, ground crew, etc., from doing so. Furthermore, it is not inherently impossible to resolve the problem of noncompeting groups in an analysis of an unrestricted sample. For example, to address questions about heterogeneity of worker characteristics or noncompeting groups (e.g., heterogeneity of worker preferences and/or qualifications), one might simply use personnel files, application forms, and the like to derive a set of

Within-group vs. workplace-wide discrimination. A third reason for analyzing restricted samples is an interest in measuring within-group, as opposed to workplace-wide discrimination. (Here, with slight modification, we use the terminology of the appellate court in *Valentino v. U.S. Postal Service* [1982, p. 605, n. 207]; for similar terminology, see *Vuyanich v. Republic National Bank* [1980, p. 266]). In some cases, this is simply an objective of the researcher, who may want to determine whether discrimination by sex or race or ethnicity is greater in faculty than in nonfaculty jobs, in white-collar than in blue-collar jobs, etc. In other cases, this interest may be a direct result of legal considerations. For example, claims about unequal pay for equal work more or less by definition require analysis of restricted samples of employees who are all doing "equal work," meaning either the same job or similar jobs. Similarly, the courts have not infrequently restricted claims in class action litigation under antidiscrimination laws to specific jobs or job categories.[6]

Thus, researchers may desire to analyze restricted samples for a variety of reasons. But is it appropriate to do so by simply excluding employees in the job types or job levels that are not of direct concern, and then applying to the resulting restricted sample the same regression procedures that one would otherwise use to analyze an unrestricted sample? In brief, the answer to this question is, "Not necessarily." Rather, as we now explain, analyzing restricted samples raises both conceptual and statistical issues.

Restricted Samples: Conceptual and Statistical Issues

The conceptual issue raised by applying a regression model to a restricted sample is straightforward—though, unfortunately, not always fully appreciated. In the absence of any statistical problems of the kind discussed below, how can one interpret regression estimates of the coef-

variables, potentially quite extensive, that reflect such differences as manifested in workers' career histories prior to starting work with their present employer. (For discussion of one study that takes this approach, see *Greenspan et al. v. Auto Club of Michigan* [1980].) Unlike variables reflecting workers' careers at their present employer, these variables are, of course, not affected by practices of the present employer. Studies in which such prehire variables are included therefore are not subject to the kinds of difficulties that arise in the context of analyses of restricted samples (which in effect control for at-hire or post-hire variables denoting job level or job type at one's present employer). We suspect that, in some instances, judicial criticism of studies that do not control for *current* job type or job level is actually intended as criticism of failure to control for worker differences in the amount and type of skill possessed. If so, it is possible that such criticism might be satisfactorily addressed by inclusion of appropriate prehire variables of the kind described here.

[6]For example, see *Eastland v. Freeman* (1981), in which race discrimination claims were limited to persons in nonmanagerial white-collar jobs at the Tennessee Valley Authority; *Segar v. Civiletti* (1981), in which race discrimination claims were limited to persons in "special agent" positions of the Drug Enforcement Administration; *Stastny v. Southern Bell* (1978), in which sex discrimination claims were limited to persons holding (or denied employment in) management positions at a telephone company; and *Valentino v. U.S. Postal Service* (1981, 1982), in which sex discrimination claims were limited to persons holding positions at level PES-17 or higher on the U.S. Postal Service pay scale.

ficients d_Y on demographic variables D when these estimates have been computed using a restricted sample?

As in the unrestricted case, the regression estimate of the coefficient d_{Yj} measures the systematic association that holds, on average, between pay (Y) and the particular demographic characteristic in question D_j (e.g., sex or race), other things being equal, where the other things in question include the measured productivity related characteristics X_1 through X_k and all measured demographic characteristics except the one denoted by D_j. The main difference between the interpretation of such results for an unrestricted regression and for a restricted regression is that, in the latter, group membership (e.g., in a particular job type or job level) is *also* among the things that are held constant, while in the former it is not.

For example, suppose one analyzes monthly compensation for a restricted sample limited to those persons within an organization who hold clerical jobs; that the measured X and D variables are schooling, years of company service, sex, race, and ethnicity; and that the regression coefficient for the demographic variable D_1 that denotes "sex is female" turns out to equal −50. This means that, taking as given not only schooling, years of company service, race, and ethnicity, but also the fact of being employed in a clerical position, being female rather than male is systematically associated with a reduction in salary of $50 per month, on average.

Clearly, then, results derived from restricted samples may provide useful insights into the nature and sources of discrimination, even though they do not necessarily measure the full extent of discrimination.[7] It is important to note that within-group differentials of this kind can overstate as well as understate the workplace-wide extent of discrimination against a particular group (e.g., blacks or women). For example, suppose a university follows a systematic practice of paying male faculty more than female faculty holding the same academic rank, but also systematically favors females over males

[7]In particular, suppose one were to classify all of the employees in an organization into a set of mutually exclusive job types or job levels and compute the sex (or race, etc.) differential for each such group. Even in the absence of problems like heterogeneity and noncompeting groups, a weighted average of these differentials would generally not be identical to the workplace-wide differential that would be derived by applying the same regression model to all of the employees taken together and treated as a single group. In the absence of statistical problems of the kind discussed below, the difference between this workplace-wide differential and the weighted average of the within-group differentials would measure the portion of the workplace-wide differential attributable to differentials in assignment or access to the set of job groups. (For an empirical study that makes this point, though not in the context of discrimination, see Medoff and Abraham [1981].) One difficulty in interpreting legal discussions of such issues is that terms relating to salary discrimination do not always appear to be defined in the same way. For example, some court decisions appear to define wage discrimination as what we have called here the workplace-wide differential (see, e.g., *Mecklenburg v. Montana Board of Regents of Higher Education* and *Greenspan et al. v. Auto Club of Michigan).* However, some researchers and some court decisions appear to define wage discrimination as that portion of the workplace-wide differential that is *not* associated with differentials in assignment to different job types or job levels (see Smith [1977]; *Agarwal v. McKee* [1977]; and *Smith v. Union Oil Company* [1977]).

in making appointments or promotions to higher academic ranks. So long as pay increases with academic rank, the extent of *workplace-wide* salary discrimination against women might be negligible even if there are substantial salary differentials adverse to women *within* individual academic ranks.

However, all this ignores an important statistical issue: regression estimates of pay differentials may, in fact, provide a statistically biased estimate of the particular within-group differentials in pay they are intended to measure. This kind of bias is usually known as "selection bias," since it arises because of the particular way in which the persons in a restricted sample have been selected from a larger underlying population.

The nature of this statistical bias may be illustrated by reference to Pascal and Rapping's (1972) study of racial discrimination in major-league baseball. As Pascal and Rapping are careful to emphasize, the *measured* variables they used in their analysis (that is, the X variables in their version of equation (1)) do not necessarily capture all of the factors affecting major-league player salaries, and there certainly exist player characteristics (notably "star quality" and field leadership) that are difficult to quantify. Since such qualities and characteristics affect one's productivity as a baseball player, they presumably affect pay for major-leaguers; since they are unmeasured (at least in the sense that the researcher—though not necessarily coaches or team owners—will generally have a hard time measuring them), they are part of the error term, u_Y.

Consider next the factors that affect whether one will *be* a major-leaguer. Presumably, these include both measured characteristics (meaning productivity related and demographic factors) and unmeasured characteristics. Indeed, it is quite possible that some of the same unmeasured characteristics that may affect pay *as* a major-leaguer also may affect whether one will *be* a major-leaguer (and, thus, may affect whether one will be included in the restricted sample).

Finally, suppose there is a restraint or barrier against black entry into the major leagues, in the sense that, in order to get into the major leagues, blacks will have to have more "true" or total productivity (meaning the sum of measured and unmeasured productivity) than whites. (The limited evidence that Pascal and Rapping were able to obtain appears to support this view.) Then, in order to get into the major leagues, blacks who have the same measured productivity (that is, the same X) as whites will have to have more true productivity (including unmeasured factors such as star quality) than whites.

Obviously, then, black players who are in the major leagues will tend to have *unmeasured* productivity related characteristics to a greater extent than white major-leaguers who have the same *measured* productivity related

characteristics. A second implication is less obvious, but equally important: in these circumstances conventional regression analysis of salaries within the major leagues will generally suggest that there is no difference in salary between blacks and whites with the same *measured* productivity, on average; but the presence of barriers against black entry into the major leagues implies that blacks in the major leagues will have greater true productivity, on average, than will whites with the same measured productivity. If so, then, on average, black major-leaguers are the victims of salary discrimination.

Although in this case a regression analysis is likely to imply the absence of discrimination against blacks within the restricted sample when it is in fact present, it is important to note that under different conditions regression based on a restricted sample could have just the opposite result.[8] In either case, the fundamental statistical issue is the same: regression analysis based on a restricted sample may produce misleading estimates of the very within-group pay differentials that the analysis is designed to measure. A statistical bias of this kind will occur whenever (i) the unmeasured factors that affect whether one will be in the restricted sample are correlated with the unmeasured factors that affect pay for persons who are in the restricted sample; and (ii) inclusion in the restricted sample is affected by the demographic characteristic of interest. When both conditions obtain, the error term u_Y will generally be correlated with the demographic characteristic of interest, other things being equal. What this suggests is that a restricted sample may also be an outcome-based or endogenously selected sample, in the sense that the measured and unmeasured factors that affect whether one will be in the restricted sample are correlated with the unmeasured factors u_Y that affect pay for persons in the restricted sample. The essential point is that both pay and being in the restricted sample are outcomes that depend on decisions of the employer, i.e., they are "endogenous."

In formal terms, this analysis implies that, in studying restricted samples, one must consider not one but two relationships. The first—the relation of immediate interest—is the function for pay, Y, as given by equation (1) above. The second is a function for a variable R that determines inclusion in (or exclusion from) the restricted sample of persons whose compensation is to be analyzed. The function for R would take the following general form:

$$(2) \quad R_i = b_{R0} + b_{R1}X_{1i} + \ldots + b_{Rk}X_{ki} + d_{R1}D_{1i} + \ldots + d_{Rn}D_{ni} + u_{Ri}$$

[8]For example, suppose one uses regression to analyze a sample restricted to managers in a company; that the company systematically *favors* blacks in making assignments to managerial positions; and that unmeasured factors such as initiative, drive, and motivation that affect pay *as* a manager also affect whether one will *be* a manager. Then reasoning similar to that used in the baseball example shows that regression analysis of the restricted sample will appear to indicate systematic salary discrimination favoring white managers relative to black managers with comparable measured productivity related characteristics, even if no such discrimination within managerial jobs is in fact present.

where the X and D variables are defined as for equation (1); the b_R's and d_R's are parameters relating the measured productivity related and demographic characteristics to R; and u_R represents unmeasured factors that affect R (and, therefore, inclusion in or exclusion from the restricted sample).[9] Taken together, equations (1) and (2) constitute a two-equation model of inclusion in the restricted sample and pay if in the restricted sample. The crucial point is that, in order to avoid selection bias, estimation of the parameters b_Y, d_Y, b_R and d_R should allow for the possibility that the error terms u_Y and u_R may be correlated.

Two additional remarks on the nature of this kind of statistical bias seem appropriate. First, evidence to the effect that persons in the restricted sample are *either* very similar to *or* very different from persons not in the restricted sample in terms of such *measured* factors as educational attainment, age, sex, etc., is insufficient to establish either the presence or the absence of statistical bias. (Rather, such bias depends on relationships involving *unmeasured* factors.) Second, this kind of bias may arise whenever one analyzes a restricted sample using conventional regression methods, whether or not one's purpose is to investigate discrimination. For example, as Cain (1976, pp. 1245-1246) notes, a number of investigations of the effects of education on pay have used conventional regression methods to analyze compensation within samples restricted to persons holding low-level jobs; in general, these studies have found that education does not appear to be systematically related to pay, other things being equal. However, as Cain goes on to point out, analyses of this kind fail to take account of possible selection bias: they fail to consider whether the unmeasured variables that affect pay in low-level jobs might be related to the unmeasured variables that affect being (or not being) in those jobs. Thus, it is possible for conventional regression to understate the true effect of schooling on pay within lower-level occupations: the fact of the relatively small difference in pay within these jobs between people with more and people with less schooling is erroneously attributed to a small effect of schooling as such rather than to offsetting differences in unmeasured factors (such as motivation) because schooling is measured while these other factors are not.

[9]To understand the meaning of equation (2), note that one may say that, in deciding whether to place persons in positions (e.g., ranks or jobs) that are included within the restricted sample, the employer implicitly scores all employees on the basis of (i) characteristics not measured by the researchers, u_R, and (ii) productivity related and demographic characteristics for which the researcher does have measures, X and D, with the latter being given weights b_R and d_R. An employee's total score is R, as given by (2); and persons are assigned to positions (and, thus, are either included in or excluded from the restricted sample) depending on the value of their total score, R. (For example, one may say that, at a university, persons with a score or R value in excess of some threshhold are assigned to full professor positions; and so on.)

The Econometrics of Selection Bias in Restricted Samples

The problem of selection bias is potentially both vexing and serious. But it is not insurmountable. In this section we discuss the econometrics of restricted regressions and note some of the ways in which selection bias may be corrected. (The general reader can omit this section without loss of continuity.)

The two-equation model for (1) defining the pay individual i would receive if he or she were a member of a particular restricted sample of employees and (2) identifying the factors that determine whether individual i will, in fact, be in the restricted pay sample may be written as follows:

(3) $Y_i = X_i b_Y + D_i d_Y + u_{Yi}$

(4) $R_i = X_i b_R + D_i d_R + u_{Ri}$

where (3) and (4) are identical to equations (1) and (2) discussed earlier, except that here, for convenience, we write them in vector-matrix form. We assume that individual i will be in the relevant subgroup if and only if $R_i > 0$, where R_i is given by (4). Thus,

(5) $R_i > 0 \iff$ i is in the subgroup of interest

Finally, we assume that u_Y and u_R are mean-zero random variables, uncorrelated with X and D, that follow some unspecified joint distribution function with constant covariance matrix.

To answer questions about pay differentials, discriminatory or otherwise, within a subgroup of employees, researchers typically estimate the parameters of (3) by applying ordinary least squares to data on the persons actually in the restricted sample of interest. However, by ignoring the mechanism that determines membership in the restricted sample, (4)-(5), these researchers run the risk of inconsistently estimating the parameters of interest, b_Y and d_Y. To understand the implications of this problem, it is useful to first analyze the general nature of selection bias and then to examine the conditions under which the conventional regression estimate of a specific coefficient will suffer from selection bias.

Concerning the first question, begin by considering the regression function for compensation Y of the persons actually in the restricted sample of interest. By (3)-(5), this regression function is

(6) $$E[Y|X, D, R > 0] = Xb_Y + Dd_Y + E[u_Y|X, D, R > 0]$$
$$= Xb_Y + Dd_Y + E[u_Y|X, D, u_R > -Xb_R - Dd_R]$$

where, for any variable z, the symbol $E[z|.]$ means the expected value of z, conditional on the information specified after the | symbol. If u_Y and u_R are independent, then $E[u_Y|X, D, u_R > -Xb_R - Dd_R] = 0$. Under these

circumstances, conventional regression estimates will be consistent and efficient. However, if u_Y and u_R are not independent, then the conditional mean of u_Y (that is, the term $E[u_Y | X, D, u_R > -Xb_R - Dd_R]$) depends on u_R and, in particular, on the probability that an individual with characteristics X and D is observed in the restricted sample. In this case, using data restricted to persons actually in the subgroup to fit a regression of Y on X and D *only*, as in ordinary least squares regression, amount to ignoring the process (4)-(5) that determines whether an observation is found in the restricted sample, thereby omitting the second term on the right-hand side of (6) from the list of regressors. Standard specification-error arguments apply; in general, at least some of the estimates of the structural parameters b_Y and d_Y in (3) obtained in this way will be biased and inconsistent. Moreover, the estimated standard errors of all of the structural parameters will be biased and inconsistent. (See Heckman [1979] for elaboration of this point.)

However, in investigations of pay differentials, the primary concern is whether the estimated coefficient of the specific demographic variable (race, sex, etc.) of interest suffers from such bias. Without loss of generality, let the demographic variable of interest be D_1, and consider the conventional regression estimate of its coefficient in (3), d_{Y1}, when obtained using data for the restricted sample. To a first order approximation, this estimate is equal to

(7) $\delta E[Y | X, D, R > 0]/\delta D_1 = d_{Y1} + \delta E[u_Y | u_R > -Xb_R - Dd_R]/\delta D_1$
where the symbol σ indicates partial derivation.

Now, *if* $d_{R1} = 0$, or *if* u_Y and u_R are independent, then the second term on the right-hand side of (7) is necessarily zero. In either of these two cases, conventional regression analysis applied to data on the restricted sample of employees provides a consistent estimate of d_{Y1}.[10] Otherwise, however, the conventional regression estimate of the parameter d_{Y1} will be biased and inconsistent; the magnitude of the bias will be approximately equal to the (nonzero) magnitude of the second term on the right-hand side of (7).[11]

Econometric solutions for selection bias. Given the potential problem raised by applying ordinary least squares to data for a particular subgroup of employees, it is important to consider econometric procedures that can be used to test and correct for this problem. In so doing, it is useful to distinguish between methods suitable for censored data sets and methods suitable for truncated data sets. Censored data sets are those in which data on Y, X, and

[10]However, the least squares estimator of the standard error of d_{Y1} will be biased and inconsistent whenever condition (ii) is not met, regardless of whether condition (i) is satisfied.

[11]Note that, *mutatis mutandis*, exactly the same remarks apply to the consistency of conventional regression estimates of the b_Y parameters and of the partial derivatives of $E[Y | X, D, R, > 0]$ with respect to the X variables with which those b_Y parameters are associated.

D are available for employees in the restricted sample and in which data are available for employees not in the restricted sample. Truncated data sets are those in which data on Y, X, and D are available for employees in the restricted sample and in which data on X and D are not available for other employees.

The most widely used method for analyzing censored data sets is a two-stage, two-equation procedure developed by Heckman (1979).[12] This procedure is based on the assumption that the joint distribution of the error terms u_Y and u_R is bivariate normal. In the first stage of the procedure, one uses probit analysis to estimate the parameters of equation (4) using data on the X and D variables of the organization's entire employee complement (together with information about which of the employees are in the restricted sample). These estimates can then be used to construct a measure of the probability of being in the restricted sample (the inverse of Mills' ratio). In the second stage of Heckman's procedure, one adds this variable to the others in equation (3) to obtain the following "expanded" version of (3):

$$(3') \quad Y_i = X_i b_Y + D_i d_Y + \lambda_i c_Y + v_{Yi}$$

where v_{Yi} is an error term and λ_i is individual i's value of the inverse of Mills' ratio.[13] One estimates the parameters of this expanded version of equation (3) by applying conventional regression methods to data on the persons actually in the restricted sample. Heckman has shown that estimates of b_Y and d_Y derived from equation (3′) will be properly corrected for any correlation between unmeasured factors u_R in equation (4) and unmeasured factors u_Y in equation (3), and will therefore not suffer from selection bias. In addition, the t-ratio of the coefficient c_Y on the added variable λ in the expanded version of (3), (3′), can be used to test for the presence of such correlation.

One-stage procedures for estimating the parameters of equations (3) and (4) from a censored data set have also been developed; these use maximum-likelihood techniques. The efficiency of estimates obtained under such procedures is greater than that of estimates obtained under Heckman's method, but such one-stage procedures are also more time-consuming and difficult to implement. (See Griliches, Hall, and Hausman [1979] for an example of one such method.)

[12]It is important to emphasize that such a two-equation procedure does not assume anything in particular about whether the effects of the two sets of unmeasured factors are actually correlated; in particular, if no such correlation exists, none will be detected in the two-equation analysis and so no correction for selection bias will be—or need be—made. For a relatively nontechnical exposition of Heckman's procedure and an interesting empirical application, see Reimers (1982).

[13]Note the resemblance between (3′) and (6). Heckman (1979) shows that, under the assumption of bivariate normality of u_Y and u_R, the quantity $E[u_Y | u_R > -Xb_R - Dd_R]$ is in fact equal to $\lambda c_Y + v_Y$, where c_Y is the normalized covariance between u_Y and u_R, and where v_Y is uncorrelated with X, D, and λ.

In contrast with methods for censored data sets, methods for truncated data sets are more difficult to implement, since truncated data sets contain little or no information about employees not in the restricted sample. However, in Bloom and Killingsworth (1981) we develop a one-stage maximum-likelihood procedure that estimates equations (3) and (4) simultaneously, using only data on individuals in the restricted sample. Like the Griliches-Hall-Hausman and Heckman methods, this procedure provides a statistical test for possible correlation between u_Y and u_R, and corrects the estimates of the parameters of (3) for such correlation, to the extent that it is present.

Interpretation of results. Once obtained, how are the consistent parameter estimates of the model (3)-(4) to be interpreted? To address this issue, it is useful to focus on a particular demographic characteristic (race, sex, ethnicity, etc.), which we again denote by D_1, and to distinguish between two questions that might be investigated. First, on average, does the employer pay the black (female, Hispanic, etc.) persons actually in the subgroup more or less than the white (male, non-Hispanic, etc.) persons actually in the subgroup who are comparable in terms of observed productivity related characteristics (X) and other demographic characteristics (D)?

This question refers to what we will call the "conditional pay differential." One may write the conditional differential as

$$\Delta_{1c} = E(Y|X_1\ X_2, \ldots, D_1 = 1, D_2, \ldots, \text{is in restricted sample})$$
$$- E(Y|X_1, X_2, \ldots, D_1 = 0, D_2, \ldots, \text{is in restricted sample})$$

This measures the average difference in pay between persons who are similar in terms of all measured characteristics (all X variables and all D variables except D_1) other than the one of direct interest and who are in the subgroup in question. By (7),

$$\Delta_{1c} = d_{Y1} + \delta E[Y|X, D, R > 0]/\delta D_1$$
$$= d_{Y1} + \delta E[Y|X, D, u_R > -Xb_R - Dd_R]/\delta D_1$$

The second question is the following: on average, if any given set of persons were placed in the subgroup, would the employer pay those persons differently depending on whether they were black or white (male or female, Hispanic or non-Hispanic, etc.)? This question refers to what we will call the "structural pay differential." This may be defined as follows:

$$\Delta_{1s} = E(Y|X_1, X_2, \ldots, D_1 = 1, D_2, \ldots, u_Y = u_Y^*, \text{is in restricted sample})$$
$$- E(Y|X_1, X_2, \ldots, D_1 = 0, D_2, \ldots, u_Y = u_Y^*, \text{is in restricted sample})$$

where inclusion of the condition $u_Y = u_Y^*$ indicates that we are holding constant the influence of unmeasured factors that affect salary determination to the extent that they are also associated with unmeasured factors that affect inclusion in the subgroup. Thus, by (7),

$\Delta_{1s} = d_{Y1}$

As noted above, if the unmeasured characteristics u_Y that affect pay if in the restricted sample are correlated with the unmeasured characteristics u_R that affect being in the restricted sample, *and* if possessing demographic characteristic D_1 affects being in the restricted sample, other things being equal (i.e., if $d_{R1} \neq 0$), then Δ_{1c} and Δ_{1s} will not be equal. This is because, in this case, the conditional differential measures the *sum* of two effects. The first is the direct effect on pay d_{Y1} of the employer's compensation practices with respect to persons possessing demographic characteristic D_1 when in the restricted sample. The second is the indirect effect on pay $\delta E[Y|X, D, u_R > -Xb_R - Dd_R]/\delta D_1$ of the employer's selection practices with respect to persons possessing demographic characteristic D_1. On the other hand, the structural differential Δ_{1s} does not include the second of these two effects; rather, Δ_{1s} measures just the magnitude of the first or "direct" effect. However, note from our discussion of (7) above that Δ_{1c} and Δ_{1s} will be equal—that is, the indirect effect will be zero—when either (i) unobservable factors that affect being in the restricted sample, u_R, are uncorrelated with unobservable factors that affect compensation, u_Y; or (ii) D_1 does not affect inclusion in the restricted sample (that is, $d_{R1} = 0$).

In summary, conventional regression provides a first-order approximation to the conditional differential Δ_{1c}. However, unless either of the conditions (i)-(ii) just noted is satisfied, conventional regression will fail to provide an unbiased estimate of the structural differential Δ_{1s}. On the other hand, joint estimation of the model (3)-(4) based on techniques such as those described above provides statistically consistent estimates of all of the model's parameters (and their standard errors); and these estimates may be used to compute consistent measures of both differentials.

Future Research Needs

It is conventional to conclude with suggestions for future research. In the present case, it is imperative to do so.

One issue that deserves further scrutiny is the empirical importance of missing variables and the extent to which omission of such variables results in statistically biased measures of race or sex discrimination in pay and other aspects of employment. For example, does omission of variables measuring actual years of work experience bias regression measures of sex differentials in pay at industrial corporations? Does omission of variables measuring publications bias regression measures of sex differentials in pay in universities? Similarly—given claims raised in one recent court case *(Rosario et al. v. New York Times)*—one might ask if omission of variables measuring writing ability or motivation bias regression measures of race differentials

in pay. So long as these variables remain unmeasured, questions of this kind will remain unanswered. The most straightforward solution is to measure such previously unmeasured variables, include them in regression studies, and see what their inclusion does to race or sex differentials. This costs money (for data collection, data-coding, keypunching, and the like) and takes time, but it is clearly the most reliable way to answer questions about missing variables.

A second set of issues on which further research is needed concerns the distinction between within-group and workplace-wide discrimination. As regards the former, can one determine whether such discrimination is more or less pronounced at high-job levels than at low-job levels (e.g., at the full professor rank vs. lower academic ranks; at professional and managerial levels as opposed to clerical, service, etc., levels)? On the other hand, to what extent is workplace-wide discrimination a result of within-group discrimination (e.g., unequal pay for equal work), and to what extent is it a result of differential access (i.e., unequal work despite equal qualifications)? Malkiel and Malkiel (1973) concluded that almost none of the workplace-wide discrimination they found at the organization they studied was attributable to within-group discrimination. However, except for their study (which ignored selection-bias problems of the kind discussed here), there have been few attempts to decompose workplace-wide discrimination into components attributable to differential access and to within-group discrimination.

A third issue that deserves further attention is the empirical importance of selection bias in studies of restricted samples of employees in individual organizations. As noted above, the results of such studies are useful indicators of within-group differentials in pay only if they are not subject to selection bias. But is such bias present, and, if so, how important is it? The limited available evidence suggests that such bias *may* be important in at least some contexts, but further examination of this issue is long overdue. Studies of selection bias in contexts other than studies of individual organizations suggest that selection bias is potentially important. For example, Reimers (1982) finds that ignoring selection bias may result in misleading estimates of earnings functions for various groups, including whites, blacks, and Mexican-Americans. Similarly, Heckman, Killingsworth, and MaCurdy (1981) suggest that ignoring selection bias may yield badly biased estimates of functions for hours of work. Finally, Bloom and Killingsworth (1981) find that ignoring selection bias in studying faculty compensation at a large university substantially understates the male-female differential in faculty salaries. While suggestive, these studies are no substitute for analysis of selection bias problems in other kinds of restricted samples (e.g., persons in a particular occupation or salary range).

A fourth area involves generalizing the discussion of selection bias from questions of pay to other kinds of employment practices, such as promotions, layoffs, and transfers. For example, suppose one is interested in studying sex differences in promotions out of a particular job category. Suppose further that individuals are selected into that category on the basis of unmeasured as well as measured characteristics. Clearly, arguments similar to those developed above lead to the conclusion that analyses of promotion rates based on a sample restricted to persons in the category—even if such analyses make adjustments for differences in measured characteristics—may provide biased measures of the extent of discrimination, if any, in promotions out of the category.

Finally, although plausible and widely held, the hypothesis that unmeasured factors u_Y are uncorrelated with observed characteristics X and D within the organization as a whole is still only a hypothesis—and one that is not universally accepted.[14] Is it necessarily appropriate to treat an employer's current total workforce as the relevant "population"? One might instead argue that the total workforce itself ought to be regarded as a restricted sample taken from some larger underlying population—e.g., all employees past as well as present, all applicants for positions at the employer, or all potential applicants, etc. Indeed, under this alternative interpretation, our discussion above of potential selection bias problems applies with full force to studies of pay differentials within an employer's entire workforce. Brown (1982) treats all applicants for positions at a particular organization as the relevant population and treats the successful applicants (that is, those who were hired) as a restricted sample in a study of on-the-job performance. While Brown does not specifically address questions about race or sex discrimination (his data apparently refer only to white males), his results suggest that selection bias may be important in studies based on an employer's entire workforce.[15] Whether this is true in other contexts—particularly as regards estimation of race or sex differentials in pay and other rewards for employment—deserves investigation.

[14]For example, in discussing analyses of the *entire* faculty at Swarthmore College in testimony at the trial of *Presseisen v. Swarthmore* (1977), a statistical expert witness, Dr. Paul L. Meier, questioned the reliability of cross-section studies of Swarthmore faculty who were employed at Swarthmore as of a given date on the grounds that such studies necessarily omitted "inactives"—faculty who were employed as of some previous date but were no longer employed as of the date referenced by the study. The trial judge took note of this comment and seems to have agreed with it; see pp. 615 and 618. For similar comments by Meier in another case, see *Dickerson v. U.S. Steel* (1978, p. 1311).

[15]Unfortunately, Brown's evidence is not clear-cut; some aspects of his study raise *other* selection-bias issues. In particular, his data refer only to applicants (successful and unsuccessful) for first-line supervisory positons, not to all applicants.

References

Agarwal v. McKee, 15 *E.P.D.* 8301 (N.D. California, 1977).

Ashenfelter, Orley C. and John Pencavel. "Estimating the Effects on Cost and Price of the Elimination of Sex Discrimination: The Case of Telephone Rates." In Phyllis A. Wallace, ed., *Equal Employment Opportunity and the AT&T Case*. Cambridge, Mass.: MIT Press, 1976, pp. 111-124.

Baldus, David C. and James W. L. Cole. *Statistical Proof of Discrimination*, New York: McGraw-Hill, 1980.

Borjas, George. "Discrimination in HEW: Is the Doctor Sick or Are the Patients Healthy?", *Journal of Law and Economics*, XXI (April, 1978), 97-110.

Boyd v. Bechtel Corporation, 610 *F. Supp.* 610 (N. D. California, 1979).

Bloom, David E. and Mark R. Killingsworth. "Correcting for Selection Bias in Truncated Samples: Theory, with an Application to the Analysis of Sex Salary Differentials in Academe." Unpublished working paper, 1981.

Brown, Charles. "Estimating the Determinants of Employee Performance," *Journal of Human Resources*, XVII (Spring, 1982), 178-194.

Cain, Glen G. "The Challenge of Segmented Labor Market Theories to Orthodox Theory: A Survey," *Journal of Economic Literature*, XIV (December, 1976), 1215-1257.

Commonwealth of Pennsylvania v. Local 542, IUOE, 469 *F. Supp.* 329 (E. D. Pennsylvania, 1978).

Dickerson v. U.S. Steel, 472 *F. Supp.* 1304 (E. D. Pennsylvania, 1978).

Eastland v. Freeman, 528 *F. Supp.* 822 (N. D. Alabama, 1981).

Fienberg, Stephen E. "The Increasing Sophistication of Statistical Assessments as Evidence in Discrimination Litigation – A Comment on Kaye." Technical Report No. 216, Department of Statistics, Carnegie-Mellon University, 1982 (forthcoming, *Journal of the American Statistical Association*).

Fienberg, Stephen E. and Miron L. Straf "Statistical Assessments as Evidence." Technical Report No. 237, Department of Statistics, Carnegie-Mellon University, 1982.

Finkelstein, Michael O. "The Judicial Reception of Multiple Regression Studies in Race and Sex Discrimination Cases," *Columbia Law Review*, LXXX (September, 1980), 737-754.

Fisher, Franklin M. "Multiple Regression in Legal Proceedings," *Columbia Law Review*, LXXX (September, 1980), 702-736.

Gordon, Nancy M. and Thomas E. Morton. "The Staff Salary Structure of a Large Urban University," *Journal of Human Resources*, XI (Summer, 1976), 374-382.

Gordon, Nancy M., Thomas E. Morton, and Ina C. Braden. "Faculty Salaries: Is There Discrimination by Sex, Race and Discipline?", *American Economic Review*, LXIV (June, 1974), 419-427.

Greenspan et al. v. Automobile Club of Michigan, 22 *F. E. P. Cases* 184 (E. D. Michigan, 1980).

Griliches, Zvi, Bronwyn Hall, and Jerry A. Hausman. "Missing Data and Self-Selection in Large Panels," *Annales de l'INSEE*, XXX-XXXI (April-September, 1978), 137-176.

Gwartney, James, Ephraim Asher, Charles Haworth, and Joan Haworth. "Statistics, the Law and Title VII: An Economist's View," *Notre Dame Lawyer*, LIV (April, 1979), 633-660.

Harvard Law Review. "Employment Discrimination and Title VII of the Civil Rights Act of 1964," *Harvard Law Review*, LXXXIV (March, 1971), 1109-1316.

Harvard Law Review. "Beyond the Prima Facie Case in Employment Discrimination Law: Statistical Proof and Rebuttal," *Harvard Law Review*, LXXXIX (December, 1975), 387-422.

Haworth, Joan G. and Charles T. Haworth. "Statistical Inference in Civil Rights Decisions," *Proceedings* of the American Statistical Association, Social Sciences Section, 1976, Part I. Washington, D.C.: 1976, pp. 47-52.

Heckman, James J. "Sample Selection Bias as a Specification Error," *Econometrica*, XLVII (January, 1979), 153-162.

Heckman, James J., Mark R. Killingsworth, and Thomas MaCurdy. "Empirical Evidence on Static Labour Supply Models: A Survey of Recent Developments." In Zmira Hornstein, Joseph Grice, and Alfred Webb, eds., *The Economics of the Labour Market*. London: HMSO, 1981, pp. 73-122.

James v. Stockham Valves, 394 *F. Supp.* 434 (N. D. Alabama, 1975), 559 *F. 2nd* 310 (Fifth Circuit, 1977).

Kmenta, Jan. *Elements of Econometrics.* New York: Macmillan, 1971.

Malkiel, Burton G. and Judith A. Malkiel. "Male-Female Pay Differentials in Professional Employment," *American Economic Review,* LXIII (September, 1973), 693-705.

Mecklenburg v. Montana Board of Regents of Higher Education, 13 *E. P. D.* 11438 (N. D. Montana, 1976).

Medoff, James L. and Katherine G. Abraham. "Are Those Paid More Really More Productive? The Case of Experience," *Journal of Human Resources,* XVI (Spring, 1981), 186-216.

Mincer, Jacob. *Schooling, Experience and Earnings.* New York: Columbia University Press, 1974.

Nie, Norman H., C. Hadlai Hull, Jean G. Jenkins, Karin Steinbrenner, and Dale H. Bent. *Statistical Package for the Social Sciences.* New York: McGraw-Hill, 1970.

Oaxaca, Ronald. "Male-Female Wage Differentials in the Telephone Industry." In Phyllis A. Wallace, ed., *Equal Employment Opportunity and the AT&T Case.* Cambridge, Mass.: MIT Press, 1976, pp. 17-48.

Osterman, Paul. "Sex Discrimination in Professional Employment: A Case Study," *Industrial and Labor Relations Review,* XXXII (July, 1979), 451-464.

Pascal, Anthony and Leonard A. Rapping. "The Economics of Racial Discrimination in Organized Baseball." In Anthony Pascal, ed., *Racial Discrimination in Economic Life.* Lexington, Mass.: Lexington Books, 1974, pp. 119-156.

Patterson v. Western Development Laboratories, 13 *F. E. P. Cases* 772 (N. D. California, 1976).

Presseisen v. Swarthmore College, 442 *F. Supp.* 593 (E. D. Pennsylvania, 1977).

Reimers, Cordelia W. "Why Are Hispanic Americans' Incomes So Low? An Analysis of the Family Incomes and Earnings of the Major Hispanic Origin Groups in the U.S." *Final Report to U.S. Department of Labor,* Princeton University, 1982.

Roberts, Harry V. "Statistical Biases in the Measurement of Employment Discrimination." In E. Robert Livernash, ed., *Comparable Worth: Issues and Alternatives.* Washington, D.C.: Equal Employment Advisory Council, 1980, pp. 173-195.

Rosario et al. v. New York Times, 84 *F. R. D.* 626 (S. D. New York, 1979).

Segar v. Civiletti, 508 *F. Supp.* 690 (District of Columbia, 1981).

Service, Jolayne. *User's Guide to the Statistical Analysis System.* Raleigh, N.C.: North Carolina State University, 1972.

Smith, Sharon P. *Equal Pay in the Public Sector: Fact or Fantasy.* Princeton, N.J.: Industrial Relations Section, Princeton University, 1977.

Smith v. Union Oil Company, 17 *F. E. P. Cases* 960 (N. D. California, 1977).

Stastny v. Southern Bell Telephone Co., 458 *F. Supp.* 314 (W. D. North Carolina, 1978).

Trout v. Hidalgo, 517 *F. Supp.* 873 (District of Columbia, 1981).

United States v. U. S. Steel, 371 *F. Supp.* 1045 (N. D. Alabama, 1973), 520 *F. 2nd* 1043 (Fifth Circuit, 1975).

Valentino v. U.S. Postal Service, 511 *F. Supp.* 917 (District of Columbia, 1981), 28 *F. E. P. Cases* 592 (District of Columbia Circuit, 1982).

Vuyanich v. Republic National Bank, 505 *F. Supp.* 224 (N. D. Texas, 1980).

Wade v. Mississippi Cooperative Extension Service, 372 *F. Supp.* 126 (N. D. Mississippi, 1974), 528 *F. 2nd* 508 (Fifth Circuit, 1976).

Wilkins v. University of Houston, 471 *F. Supp.* 1054 (S. D. Texas, 1979), 26 *F. E. P. Cases* 1041 (Fifth Circuit, 1981a), 662 *F. 2nd* 1156 (Fifth Circuit, 1981b).