DISCUSSION

Arlene Ash, Boston University Medical School

This paper makes the valid, although not entirely new, point that the relationship of salary to qualifications for employers within a company often differs strikingly from one job classification to another. The authors are concerned that a presumed non-discriminatory difference in the pricing of different jobs will be mistaken for discrimination against (say) women when women are disproportionately present as employees in the lower priced jobs.

If candidate pools for the different jobs have essentially no overlap (e.g., among university-level professors in distinct diciplines) discrimination studies commonly employ dummy variables to indicate jobs or homogeneous groupings of jobs (such as natural science departments and social science departments), adjusting for these differences in a straightforward way. The use of such an analysis implicitly accepts the employer's pricing differences between jobs as appropriate and non-discriminatory, although the actual level of pay differences by job are themselves legitimate subjects for study and open to challange from at least two directions.

One potential source of criticism is, of course, the issue of comparable worth (is it proper to pay truck drivers more than librarians?); another arises even if we in principle accept that pay differentials due to market forces (and not related to skill, training or job difficulty) are non-discriminatory. Actual job-pay differences at a particular place of employment are certainly influenced by market forces, but may not be determined by them. The dummy-variable-for-each-department methodology described above will not find any sex discrimination in a school for paying its English professors three times as much as its (comparably qualified) computer scientists, nor for doing just the reverse. When, as is usual, markedly different sex-ratios prevail in the different departments, the so-called "departmental difference" may be absorbing some of the difference due to sex discrimination. For example, anyone might become suspicious if the better-paid English professors in a school were mostly men while the less-well-paid computer scientsits were mostly women. But even when differences go in the expected direction are we willing to accept whatever differential exists (be it $100/year or $10,000/year) as due to "non-discriminatory" market forces? The usual dummy variable methodology is equivalent to doing just this.

The current paper, however, accepts as non-discriminatory whatever salary differences actually exist among jobs; for the purpose of future discussion we will do so as well. With this presumption, we agree with the authors that discrimination may take place in at least two separate ways, one in assignment of salary within job and the other in job placement. The bulk of their paper is devoted to developing methodology for detecting problems relating to unfair placement. I have problems both with the authors' statement of the problem and with each of their two proposed solutions. We will use their two-job hypothetical (as in Tables 1A & 1B) for discussion purposes.

In stating the problem, Conway and Roberts do not distinguish between settings in which the employment pools for the different jobs are disjoint and ones in which the same individuals are potentially candidates for either job. In the first situation, illustrated by, say, Job 1 is "English professor" and Job 2 "computer scientist", I would agree that Tables 1A & 1B do not provide any reason to suspect sex discrimination. However, suppose that Job 1 is "bank teller" and Job 2 is "management trainee at the bank". Then the two tables show clear evidence of a problem, in that among all employees with fourteen years of education, 3 out of 4 men, but only 1 in four women get the better job. A reasonable analysis of these date under the assumption of overlapping employment pools would ignore job classification and use a three-piece function to model the clearly non-linear jumps in salary that occur above and below fourteen years of education. Such an analysis would correctly identify a salary differential between men and women at the middle-level of education which requires further explaining, since in this example women with intermediate qualifications are far more likely than similarly situated men to be placed in a job "for which they are overqualified".

That Conway and Roberts' "reverse regression" methodology finds no suggestion of a problem in these data is to me, convincing evidence of the method's inappropriateness for detecting "shunting" when the pools are not disjoint. The paper should clearly state that the suggested methodology is not appropriate in any setting with overlapping candidate pools. It should explicitly warn people about its potential for abuse in a setting, such as university teaching, where the different professional ranks are essentially "honorifics", conferred or withheld by the institution, and are certainly not distinct job categories with disjoint labor pools. When job pools overlap, it is appropriate to omit job categories from the analysis, recovering as much explanatory power as possible from other factors. (At a university, for example, one might expect a salary discontinuity at six years, when the tenure decision usually occurs.)

So far we have only talked about differential treatment (with respect to either salary or job placement) of men and women among those actually hired. Another way in which employment discrimination may occur is through differential hiring standards for men and women. If we knew the sex and qualifications of the entire candidate pool for a particular job, Conway and Roberts reason that it should be possible to detect such "selection discrimination" into homogeneous job groups through logistic modelling. I see two major problems with this idea: the first has to do with the slippery concept of a "candidate pool", which is not easy to define. (Does it include every person with a Ph.D. in English? Those who were likely to have

seen notices informing them of an opening? Those currently living within an n-mile radius of the university who have published at least one book or article in the last year? Those who actually submitted applications?) Also, the pool of applicants can clearly be manipulated, for example, by inviting less qualified women to apply. Finally, not only are the data relating to non-hired individuals (in the words of this paper) "often absent", such data are virtually never present.

Even if acceptable candidate pool information were available, I wonder if the suggested logistic modelling has good power to detect selection discrimination when it is present. The paper provides no suggestions for how to test for the cumulative effects of sex differences in hiring, placement and salary within job when problems exist at more than one level. In general the authors seem far less concerned with the "type 2" error of failing to detect discrimination when it is present than with the problem of falsely rejecting the hypothesis of non-discrimination.

As a final point, I consider the authors' argument that, in the absence of information on those not hired (and, I would add, with disjoint candidate pools) selection discrimination is appropriately studied by comparing the average qualifications of men and women within jobs. The following hypothetical shows clearly that the untestable assumption of a similar distribution of qualifications among men and women in the original candidate pool is crucial, since the suggested methodology is prone to "see" discrimination against the group which is (on average) more highly qualified even when no such discrimination has occurred.

Suppose that a company has repeated job openings for a type of job for which it is willing to hire a person whose level of qualification, Q, is either 1, 2, 3 or 4. It adopts a "facially neutral" hiring policy in which one man and one woman, each with Q in the appropriate range, enter the candidate pool. The one with higher Q gets the job, ties being broken by the toss of a fair coin. If the distribution of Q in the population from which the candidate pool is formed is the same for men and women, a symmetry argument shows that the expected distributions of Q among those hired will be the same for both sexes. However, suppose that the distributions are as shown in Figure 1, below.

Figure 1: Percentages of people at each level of Q among prospective enrollees in a company's candidate pools for a given job, by sex.

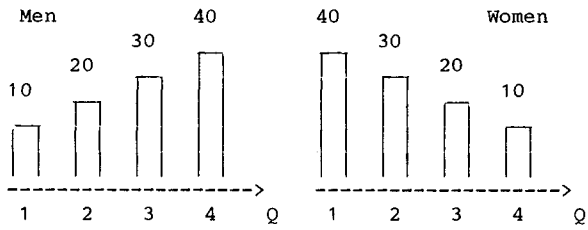

The policy of forming two-person candidate pools through a random selection of one man and one woman leads to the joint distribution indicated in Figure 2. The resulting hiring for men is summarized on the right. Similar calculations for the women would come from comparing the numbers below the diagonal to the column sums for each of the four levels of women's Q.

Figure 2: Joint distribution (probabilities x100) of the 4x4 possible compositions of one-man, one-woman candidate pools, with male hiring outcomes indicated

Man's Q	Woman's Q = 1	2	3	4	
4	16	12	8	4	->38 of 40 men w. Q=4 hired
3	12	9	6	3	->24 of 30 men w. Q=3 hired
2	8	6	4	2	->11 of 20 men w. Q=2 hired
1	4	3	2	1	-> 2 of 10 men w. Q=1 hired

From the joint distribution, it is easy to verify that the average Q of the men who are hired ((38x4+24x3+11x2+2x1)/75 = 3.31) is greater than the average Q for the hired women ((8x4+9x3+6x2+2x1)/25 = 2.92). When the rejected candidates' qualifications are absent, the Conway-Roberts methodology would mislead us into thinking that the men had been discriminated against. Interestingly enough, if the candidate pool information were available, their suggested methodology would show discrimination against women, despite the apparent neutrality of the hiring rules. See Figure 3.

Figure 3: Percentage hired, by level of Q, for men and women

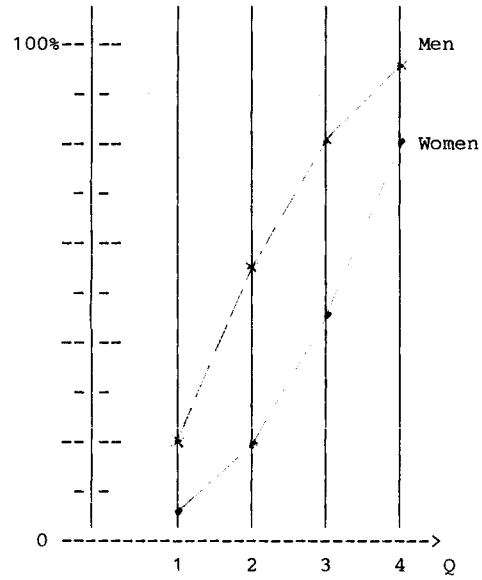

Which view is correct? Are the women discriminated against or are the men or neither? Our interpretation is the following: while the rule is "facially neutral", it is in fact discriminatory against women, which the direct regression (indicated in figure 3) correctly detects. The discrimination occurs because each woman competes against a candidate drawn from a more qualified pool than each man competes against. If we alter this example to let men and women enter the hiring pool in equal numbers on average, but not in sex-matched pairs, the inequality seen by direct regression would vanish; however, the reverse regression methodology would find an even higher level of discrimination against men than it detects in the original example (Among those hired, $\bar{Q}$ men = 3.4 while $\bar{Q}$ women = 2.7).

In conclusion, while the authors make some new points, especially with regard to the desirability of studying qualifications of those not hired as well as those hired, their paper has a much narrower range of applicability than it appears, being at most appropriate in settings in which 1) candidate pools for distinct jobs are essentially disjoint, 2) the size of pay differentials between jobs can plausibly be related to "acceptable" reasons (certainly "different levels of responsibility", possibly "market forces"), 3) a reasonably clear definition of the "candidate pool" can be agreed upon and 4) reliable data regarding the qualifications of those not hired are available. Even in this situation the authors have not convinced me that their proposed methodology has good power to detect employment discrimination when it is present.