SEX SEGREGATION IN U.S. MANUFACTURING

WILLIAM J. CARRINGTON and KENNETH R. TROSKE*

This study of interplant sex segregation in the U.S. manufacturing industry improves on previous work by using more detailed information on the characteristics of both workers and firms and adopting an improved measure of segregation. The data source is the Worker-Establishment Characteristics Database (a U.S. Census Bureau database) for 1990. There are three main findings. First, interplant sex segregation in the U.S. manufacturing industry is substantial, particularly among blue-collar workers. Second, even in analyses that control for a variety of plant characteristics, the authors find that female managers tend to work in the same plants as female supervisees. Finally, they find that interplant sex segregation can account for a substantial fraction of the male/female wage gap in the manufacturing industry, particularly among blue-collar workers.

This paper measures interplant sex segregation and explores the connection between segregation and the male/female wage gap. Early studies of cross-employer sex segregation (McNulty 1967; Buckley 1971; Blau 1977) found that women were typically segregated into the lowest-paying employers, even within occupations, and these results have been consistently replicated in more recent studies (Pfeffer and Davis-Blake 1987; Groshen 1991; Reilly and Wirjanto 1997; Carrington and Troske 1995; Griffin and Trejo 1995). However, there are at least two limitations of the existing literature. First, there is no publicly available database suited to the study of interfirm segregation that is truly representative of the national economy.[1] Thus, our knowledge of interfirm segregation is limited to snapshots from specialized samples. Second, the existing literature does not take proper account of the fact that the random assignment of workers to firms will typically generate some segregation, at least by conventional measures. Thus, previous studies may have overstated the extent to which men and women are systematically segregated.

This paper departs from previous work in three ways. First, we study data drawn from relatively large plants in the U.S. manufacturing industry, an important sector under-studied in previous work. Second,

*William J. Carrington is Senior Associate with Charles River Associates, Washington, D.C., and Kenneth R. Troske is Assistant Professor in the Department of Economics, University of Missouri–Columbia.

[1]The Equal Employment Opportunity Commission (EEOC) data studied by Griffin and Trejo (1995) come closest to being nationally representative, but they are not publicly available.

Industrial and Labor Relations Review, Vol. 51, No. 3 (April 1998). © by Cornell University.
0019-7939/98/5103 $01.00

our data include detailed information on workers' demographic characteristics that was missing in most previous studies. This information lets us better assess the extent to which human capital differences between men and women drive interplant segregation. Finally, we measure segregation in a different and better way. Standard segregation measures conflate the random allocation of workers to firms with systematic segregation that might be due to discrimination. In contrast, we use new measures that better distinguish between the systematic and random components of segregation.

The analysis is primarily motivated by policy considerations. Comparable worth policies equalize the wages paid to jobs of equal worth, where "worth" is defined in terms of each job's skill requirements, physical demands, and other attributes. Comparable worth typically increases the relative pay of jobs predominantly held by women, such as secretaries, clerks, and other "pink-collar" jobs. Implementation is usually on an employer-by-employer basis, so that disparities within employers are addressed by the policy, while disparities across employers are not. As Johnson and Solon (1986) pointed out, this implies that comparable worth can be effective only if men and women work in the same firms. Donohue and Siegelman (1991) made a similar point with respect to Title VII of the 1964 Civil Rights Act.[2] Thus, some of the most important policies concerning women in the labor market will be ineffective if men and women are heavily segregated in the workplace.

Our analysis is also motivated by the fact that several theories of the male/female wage gap predict sex segregation. For example, Becker's (1957) model predicts that discrimination by employers, employees, or customers leads to sex segregation and reduced wages for women. Lang (1986) generated similar predictions with a model based on intersex communications difficulties. Or if skill requirements vary widely among employers and if men and women have different skills, then employers will tend to hire men or women, but not both.[3] Following Macpherson and Hirsch (1995), we refer to this as the quality-sorting hypothesis. To the extent that we can isolate groups of workers with similar skills, this theory predicts that there should be no systematic within-skill-group segregation and no within-skill-group wage gap. Thus, several theories of the male/female wage gap predict that men and women will be segregated in the workplace, and our analysis provides some information on the validity of this class of models. And although the similar implications of all models within this broad class preclude any sharp distinctions, where possible we look for evidence that favors each particular model.

We begin by discussing the data set used in the analysis, how this data set is constructed, and the advantages these data offer for studying sex segregation. We discuss some weaknesses of existing measures of segregation and present our alternative measures. Next, we present our results on the interplant sex segregation of workers. We then turn to examining what systematic forces might account for the observed patterns of segregation. In particular, we focus on the gender of the supervisors in the plant, because both the Becker (1957) and Lang (1986) models predict there will be a positive correlation between supervisors' gender and the gender of the workers they supervise. Since most anti-discrimination policies are directed toward equalizing the wages of men and women, we next examine the relationship between the interplant sex

[2]Donohue and Siegelman's argument is that while the implementation of Title VII initially focused on all aspects of employers' personnel policy, Title VII lawsuits have become increasingly focused on wrongful discharge and on pay inequalities within firms.

[3]See Kremer and Maskin (1995), Doms, Dunne, and Troske (1997), and Troske (forthcoming) for evidence on cross-employer variation in the average skills of workers.

segregation of workers and the male/female wage gap. Finally, we offer some concluding remarks.

Data Sources

This study uses data from the Worker-Establishment Characteristics Database (WECD), a recently developed Census Bureau database that matches information on workers and employers.

The basic design of the WECD is as follows.[4] Manufacturing plants in the Census Bureau's Longitudinal Research Database (LRD) are associated with an industry and a geographic block code. Workers responding to the 1990 Census Long Form report the industry and street address of their employer, which has been linked to the block codes used in the LRD. This information can be used to match workers and firms on the basis of industry and block code. Workers are successfully matched only if in the LRD there is a unique plant in the appropriate industry and block code.[5] Thus, workers can fail to be matched if (a) there is no plant in the LRD in their reported industry and block (perhaps due to reporting error by the worker), or if (b) there is more than one plant in the LRD in their reported industry and block, in which case no unique assignment can be made. It is important to note that since the LRD samples only manufacturing establishments, the WECD only contains workers from manufacturing. Nevertheless, the combination of detailed employer and employee data is crucial to the measurement of interfirm segregation and is unmatched by other U.S. data sources.

The design of the WECD implies that we have information on a sample of manufacturing plants, rather than the entire population of plants in manufacturing. The incomplete sampling occurs because the LRD contains only a sample of all manufacturing plants, because plants that share an industry and block with another plant are excluded, and because for some small plants there were no workers who responded to the Census Long Form. We correct for any associated sample selection bias by weighting each plant by its sample weight.[6] Further, within any plant in our sample, we have an incomplete sample of workers. The incomplete sampling within establishments occurs because the Census Long Form is only administered to every sixth household, and because incomplete or inaccurate information on their employer's address or industry precludes matching some workers to their plant. Together, these factors imply that, among those plants represented in our sample, our sample includes roughly 1 out of 12 workers.[7]

Before moving on, let us emphasize the advantages of using the WECD to investigate interplant segregation. First, we have human capital indicators such as education, age, experience, and occupation of each worker.[8] In contrast, most previous segregation studies have used data with relatively crude information about workers' human capital attributes (for example, Groshen 1991; Griffin and Trejo 1995), or

[4]See Troske (forthcoming) for a more complete description of the WECD's development.

[5]This is strictly true only for plants and workers in urban areas. For rural areas, workers and plants are matched on the basis of industry and place code.

[6]Sample weights are inversely proportional to the predicted probability that the plant is in the sample. In particular, the weights are the product of fitted probabilities obtained from two probit regressions. The dependent variable in the first regression is the probability that a plant is unique in a location-industry cell, and, conditional on the plant being unique, the dependent variable in the second regression is the probability that the plant appears in the WECD. Results from these probit regressions are available from the authors upon request.

[7]While it would obviously be preferable to have complete information on each plant's work force, there is little reason to believe that there is much sample selection among workers *within* any given plant. Thus, we do not attempt to correct for sample selection bias arising from the incomplete sampling of workers within plants.

[8]This information comes from workers' responses to the Census Long Form.

data with no worker information at all (for example, Carrington and Troske 1995). Second, the WECD provides a useful set of plant characteristics, including sales, value-added, and the capital-labor ratio.[9] In conjunction with the relatively precise human capital measures, these plant characteristics permit an investigation of the characteristics that distinguish "male" and "female" establishments. Finally, the WECD covers the manufacturing sector, which has been relatively understudied in previous work.

Sample Selection Criteria

There are significant differences between men and women in the tendency to work part-time. In 1990, for example, roughly 25% of employed women worked fewer than 35 hours per week, while only 10% of employed men worked part-time by this definition (U.S. Statistical Abstract 1994). These gender differences in part-time status raise some important sample selection issues. If plants first choose their use of part-time workers and then hire without regard to sex, including part-time workers in the sample runs the risk of yielding a finding of "sex segregation" when what is really at work is the segregation of part-time and full-time workers. In contrast, if some plants' choice of part-time versus full-time workers is in part driven by preferences regarding the sex of their employees, then excluding part-time workers will lead us to miss an important dimension of inter-plant segregation.

Since we see no *a priori* reason to believe that either of these views is strictly correct, we experiment with different sample restrictions. For most of our work, we restricted the sample to workers who (a) usually worked more than 30 hours per week in the previous year and (b) worked more than 30 weeks in the previous year, which excludes most part-time workers. However, we examine the sensitivity of certain key results to the inclusion of part-time workers in the sample. Other than these restrictions on weeks and hours at work, our only restriction on workers is that a worker's wage not be too much of an outlier. In particular, we required that workers' actual log wage be within five standard deviations of their predicted log wage in a regression of log wages on a wide variety of individual and employer characteristics. This last restriction actually excludes very few people, so it has virtually no effect on our measures of segregation, but it leads to what we believe are slightly more representative measures of wage differences between men and women. Finally, in order to facilitate our study of segregation, we required that workers be in a plant where at least two other workers are also in the WECD. The application of these sample restrictions leads to a final sample of 123,183 men and 48,670 women spread across 8,308 manufacturing plants.

Table 1 presents selected summary statistics from our WECD sample of workers and firms and, for comparison purposes, a set of similar statistics from a random sample of all manufacturing workers in the 1990 Census. The first four rows of Table 1 indicate that WECD workers are somewhat better paid than average workers in the manufacturing industry, who are of course themselves well paid by economy-wide standards. This is largely because WECD plants are larger than the manufacturing industry average, and because wages are generally higher in large plants (Brown, Medoff, and Hamilton 1990). Rows 1 through 4 also show that there is a substantial wage gap between men and women in the WECD, and this gap is roughly consistent with that observed in manufacturing as a whole.

Rows 5–7 of Table 1 report the average age, percent ever married, and average potential experience of WECD workers, all of which are somewhat higher than the corresponding statistics for all manufacturing workers. Row 8 reports that the average WECD worker works at a very large establishment. This is again partly due to the generally large size of manufacturing establishments, but WECD establishments are

[9]This information comes from the LRD.

Table 1. Worker-Establishment Characteristics Database Sample Means.

Variable	Sample: 1990 Decennial Census Manufacturing (1) Men	(2) Women	WECD (3) Men	(4) Women
1. Hourly Wage	14.22	9.25	14.97	9.65
2. Log Hourly Wage	2.52	2.11	2.59	2.16
3. Annual Earnings	31,642	19,269	33,436	20,150
4. Log Annual Earnings	10.20	9.73	10.29	9.79
5. Age	39.3	39.0	40.5	39.7
6. Percent Ever Married	83.8	82.8	87.3	84.1
7. Potential Labor Market Experience	20.9	21.1	22.3	21.9
8. Total Employment at Plant	—	—	2,086	1,361
9. *Education*				
Less Than High School	18.4	21.7	16.7	19.6
High School Diploma	37.4	43.1	41.6	46.9
Some College	26.6	24.6	26.3	23.9
College Degree	13.2	8.6	11.7	7.8
Advanced Degree	4.6	2.1	3.7	1.8
10. *Region (%)*				
Northeast	20.5	21.3	26.4	29.3
Midwest	35.4	30.9	47.6	41.0
South	30.0	34.4	20.5	24.8
West	14.0	13.4	6.1	5.5
11. *Occupation (%)*				
Managers and Other Professionals	20.3	14.9	17.7	13.4
Clerical and Other Non-Production Workers	8.1	24.9	8.4	25.3
Sales Occupations	8.1	5.6	5.8	4.9
Operatives and Fabricators	33.5	39.5	36.6	39.8
Precision Production, Craft, and Repair Occupations	24.5	9.9	25.3	10.0
Laborers	5.3	5.2	6.2	6.6

large even relative to this industry baseline.[10] The rows below number 9 show that WECD workers have education similar to manufacturing workers in general, and that men in the WECD (as in manufacturing in general) have somewhat more average schooling than women.[11] Finally, the rows below the numbers 10 and 11 indicate that the restriction to the manufacturing sector has led our sample of WECD workers to be concentrated in the Northeast and Midwest and in those occupations with strong representation in manufacturing. Together, these facts imply that the WECD is a somewhat select group of workers. There are no obvious reasons why the pattern of segregation in the WECD should differ dramatically from that of the aggregate economy, but we must still recognize that these results may not be completely representative of the U.S. economy as a whole.[12]

[10]The large size of WECD plants arises from the WECD sampling frame. Workers are matched to plants based on industry and location, and unique matches can be made only to plants that are the only one in their industry and location. These unique plants tend to be larger than average (Troske, forthcoming).

[11]While the mean education level of WECD workers is similar to that of the Census manufacturing sample, WECD workers are somewhat less likely to be at either extreme of the educational distribution.

[12]Troske (forthcoming) shows that standard wage regressions estimated on the WECD are very similar to analogous regressions estimated on more clearly

Measuring Segregation

Economists typically summarize segregation patterns with indices whose range is the [0,1] interval. An index of zero corresponds to complete evenness, which occurs when groups are proportionately represented in each plant. For example, if the female population share were 50%, then the sample would be even if every plant employed equal numbers of men and women. In contrast, an index of one corresponds to complete unevenness, which occurs when every plant is either all male or all female. We use two particular indices in this study. The first is the popular dissimilarity index developed by Duncan and Duncan (1955). If we let w_i and m_i equal plant i's share of female and male employees in the sample, respectively, then the dissimilarity index is simply

$$(1) \qquad D = \sum_i \frac{1}{2} | w_i - m_i |.$$

The dissimilarity index may be interpreted as the share of men (or women) that would have to change plants in order to make the sample completely even. Thus, an index of one implies that the sample is completely uneven, while an index of zero implies that the sample is completely even.

Hutchens (1991) criticized the dissimilarity index because it is equally sensitive to mild and extreme departures from evenness, and recommended using instead the gini coefficient of segregation. If we first sort the plants on the basis of w_i/m_i, then the gini coefficient of segregation may be expressed as

$$(2) \qquad G = 1 - \sum_{i=1}^{T} w_i \left(m_i + 2 \sum_{j=i+1}^{T} m_j \right),$$

where T is the number of plants in the sample. As the name suggests, this index is analogous to the gini coefficient of variation widely used in income studies. The only difference is that this index measures interplant variation in gender work force shares (that is, w_i/m_i) rather than interpersonal variation in income.[13]

A weakness of the dissimilarity index and the gini coefficient is that they can both be positive when workers are allocated randomly to plants.[14] The problem is that unless the plant is very large, a random draw of employees will not typically reflect the population exactly. To see this clearly, consider a large sample of two-worker plants drawing randomly from a population with equal numbers of men and women. One-quarter of such plants would have two men, one-half would have one man and one woman, and one-quarter would have two women. Although such a distribution is completely random, it is far from even, and the gini coefficient and dissimilarity index would be .75 and .50, respectively. While this is an extreme case, it is easy to show that random allocation of workers to plants implies substantial unevenness for plants with as many as fifty or a hundred employees (Carrington and Troske 1997).

representative data, such as the Decennial Census. While this does not prove that segregation patterns in the WECD are representative of the wider economy, it does indicate that patterns of wage variation in the WECD are not anomalous.

[13]To see the difference between the two indices, consider a four-plant sample with the following distribution of men and women: plant 1 (50 women and 50 men), plant 2 (50 women and 50 men), plant 3 (75 women and 25 men), plant 4 (25 women and 75 men). Both the gini coefficient and the dissimilarity index would characterize this distribution as segregated (D = .25, G = .375). Now imagine either (a) having plants 1 and 2 trade a man for a woman, or (b) having plant 3 trade a man to plant 4 in return for a woman. Either move would result in a sample that was more segregated from the perspective of either index. However, the gini coefficient would view case b as causing a greater increase in segregation (.39 in case b vs. .38 in case a), whereas the dissimilarity index would treat the two cases symmetrically (.26 in either case). This is an example of the way in which the gini coefficient puts more weight on the tails of the gender share distribution.

[14]This critique also applies to other commonly used segregation measures such as Atkinson's index and Theil's Entropy Index.

This observation has three implications for the interpretation of segregation indices. First, it is often a mistake to conclude from positive and statistically significant segregation indices that there is any systematic sorting of men and women into different employers; such patterns are often equally explicable by chance. Second, random allocation generates far more unevenness among small plants than among large plants. In contrast with the example above, the random allocation of a 50/50 mix of men and women to a large sample of 1000-worker plants would lead to a gini coefficient of .04 and a dissimilarity index of .03. This implies that cross-sample comparisons of segregation are difficult to interpret unless the samples have plants of roughly equal size. Third, it is the number of workers per plant in the sample that dictates the importance of such random segregation. There is no reason to believe that small samples from large plants will be evenly distributed. Thus, since our sample has, on average, only one-twelfth of the workers in a plant, the fact that we look at relatively large plants does not mean we are free of this problem.

Carrington and Troske (1997) proposed the following modifications of the gini coefficient and the dissimilarity index as a means of distinguishing between systematic and random segregation. We couch the modification in terms of the gini coefficient, but the modified dissimilarity index is completely analogous. Let the *gini coefficient of random segregation* G^* be the gini coefficient that would occur if a very large number of workers were allocated randomly to employers, taking the sexes' population shares and the size distribution of plants as determined by the sample.[15] Put slightly differently, G^* is the average gini coefficient obtained if sample workers are assigned randomly to sample plants. Thus, if a sample contains mostly large plants and roughly even numbers of men and women, G^* will be close to zero. In contrast, if the sample contains mostly small plants and is either mostly male or mostly female, then G^* will be closer to one.

We use the gini coefficient of random segregation to adjust the standard gini coefficient so that it accounts for the role of random assignment in causing unevenness. In particular, we define the *gini coefficient of systematic segregation* as

$$(3) \qquad \hat{G} = \frac{G - G^*}{1 - G^*} \text{ if } G - G^* \geq 0$$

and

$$\hat{G} = \frac{G - G^*}{G^*} \text{ if } G - G^* < 0.$$

If the sample is more uneven than random allocation would imply, that is, if $G > G^*$, then $\hat{G} > 0$ is simply excess unevenness ($G - G^*$) expressed as a fraction of the maximum amount of such excess segregation that could possibly occur ($1 - G^*$). $\hat{G} = 1$ is analogous to complete unevenness, as with the standard gini coefficient, but $\hat{G} = 0$ implies that the sample is equivalent to random allocation. If, in contrast, there is excess evenness, that is, $G > G^*$, then $\hat{G}$ is negative and represents excess evenness in the sample ($G - G^*$), expressed as a fraction of the maximum amount of such excess evenness that could possibly occur (G^*).[16]

---

[15]See Carrington and Troske (1997) for a complete discussion of this issue. The basic idea behind the computation of G^* is as follows. The first step is to calculate the empirical number of firms in each size class s, $\hat{g}(s)$. Within any size class s, random allocation implies that the binomial function $B(m;s,p)$ is the fraction of firms that will have m women if p is the female population share. Thus, random allocation implies that the number of firms with s total employees and m women should be $N(m,s;p) = B(m;s,p)\hat{g}(s)$, where p and $\hat{g}(s)$ are determined by the sample. Note that $N(m,s;p) > 0$ for all nonnegative integer values of s and m such that $m \leq s$ and $\hat{g}(s) > 0$. $N(m,s;p)$ fully characterizes the distribution of employment that would be expected under random allocation, given p and $\hat{g}(s)$. The gini coefficient of segregation computed from this artificial distribution is what we call the gini coefficient of random segregation.

[16]The examples from the text illustrate the difference between G and $\hat{G}$. First, consider a large sample of two-person plants with equal aggregate numbers of men and women and a traditional gini coefficient of .75. In this case, $\hat{G} = 0$ because .75 is exactly what random assignment would imply. Second, consider a

To summarize, our $\hat{G}$ is different from the standard gini coefficient in two ways. First, we have set the baseline of 0 to correspond to random allocation rather than complete evenness. Second, we have remapped values of G that are greater than G^* into the [0,1] interval, and remapped values of G that are less than G^* into the [-1,0] interval. In the following analysis, we will examine a similarly modified dissimilarity index. We think that these modified indices provide more useful information than the traditional ones. However, we recognize that some readers will be more comfortable with the traditional indices. Thus, in the work that follows we report traditional indices of segregation, indices of random segregation, and the indices of systematic segregation that we developed here. Together, these indices provide a useful summary of segregation patterns.

Interplant Sex Segregation

Table 2 examines the extent to which men and women are segregated across manufacturing plants in the United States.[17] The six columns each report a different index. Columns (1)–(3) report results for the three versions of the gini coefficient, whereas columns (4)–(6) report analogous results for the dissimilarity index. While the columns differ by the index reported, the rows of Table 2 vary by the sample of workers considered. Row 1 reports results for all workers in the WECD, while other rows report results for samples stratified by broad industries (rows 2 and 3), by selected detailed industries (rows 4–8), by broad occupations (rows 9–14), by selected detailed occupations (rows 15–20), or by schooling (rows 21–24). The numbers without parentheses are the index values and the numbers in parentheses are bootstrap standard errors, which Boisso et al. (1994) have shown to be useful measures of sampling error in the present context.[18]

Row 1 shows that there is substantial sex segregation across U.S. manufacturing plants, as the traditional gini coefficient is .59 and the traditional dissimilarity index is .43. Thus, the distribution of men and women across plants is far from even. Some of the observed unevenness is attributable to random allocation of workers to plants, as the random gini coefficient is .24 and the random dissimilarity index is .16. Nevertheless, the random segregation indices of columns (2) and (5) are substantially less than what is actually observed in columns (1) and (4). This fact is reflected in the systematic indices of row 1, which indicate that excess unevenness is, depending on the particular index used, 45% or 33% of the maximum that could possibly be observed. Thus, men and women are systematically segregated in U.S. manufacturing.

These results are consistent with the broad class of models that predict systematic sex segregation. The rest of the table

large sample of 1,000 worker plants with equal aggregate numbers of men and women and a traditional gini coefficient of .30. In this case, $\hat{G}$ would be very close to .30 as well, as random assignment implies very little unevenness. As these examples illustrate, the extent to which $\hat{G}$ differs from G is entirely dependent on the likely role of random allocation in the particular sample at hand.

[17]It is important to remember that, like most previous authors (for example, McNulty 1967; Buckley 1971; Blau 1977; and Groshen 1991), we study segregation across plants rather than firms.

[18]See Efron and Tibsharani (1993) for a general introduction to bootstrapping. As the reader will of course know, standard errors are an estimate of how much the associated statistic would be expected to vary in repeated sampling from the population. The basic idea behind bootstrapping is to substitute the sample for the population. In particular, a series of 500 or so bootstrap samples are drawn randomly from the true sample *with replacement*, and the statistic (in this case the segregation index) is computed for each bootstrap sample. The "bootstrap standard error" is then the standard deviation of the distribution of statistics computed from the 500 bootstrap samples.

It may seem puzzling that the "Expected" indices of columns (2) and (5) have standard errors. However, recall that these indices are what would be expected *conditional* on the joint sample distribution of establishment sizes and gender. Since this distribution varies somewhat across bootstrap samples, the expected indices are themselves stochastic.

Table 2. Indices of Interfirm Sex Segregation.

Sample	Gini Index (1) Gini	(2) Random Gini	(3) Systematic Gini	Dissimilarity Index (4) Dis-similarity	(5) Random Dis-similarity	(6) Systematic Dis-similarity
1. All Workers	.59 (.01)	.24 (.01)	.45 (.01)	.43 (.01)	.16 (.00)	.33 (.01)
Within Broad Industry						
2. Nondurables	.59 (.01)	.26 (.00)	.45 (.01)	.43 (.01)	.17 (.00)	.31 (.01)
3. Durables	.55 (.01)	.23 (.01)	.41 (.01)	.40 (.01)	.15 (.01)	.30 (.01)
Within Selected Detailed Industry						
4. Meat Products	.43 (.05)	.21 (.02)	.28 (.06)	.30 (.05)	.14 (.01)	.19 (.06)
5. Apparel, Excluding Knit	.52 (.03)	.35 (.01)	.26 (.04)	.37 (.02)	.23 (.01)	.18 (.03)
6. Newspaper Publishing	.40 (.02)	.31 (.01)	.13 (.02)	.29 (.01)	.21 (.01)	.10 (.02)
7. Motor Vehicles and Equipm.	.41 (.03)	.15 (.02)	.31 (.03)	.32 (.02)	.10 (.01)	.24 (.02)
8. Household Appliances	.31 (.03)	.12 (.01)	.22 (.03)	.22 (.03)	.08 (.01)	.16 (.03)
Within Broad Occupation						
9. Prof./Managers	.63 (.02)	.54 (.02)	.19 (.01)	.44 (.01)	.36 (.02)	.13 (.01)
10. Sales/Service	.79 (.01)	.73 (.01)	.24 (.03)	.59 (.02)	.52 (.02)	.15 (.02)
11. Clerical	.68 (.01)	.59 (.01)	.21 (.02)	.49 (.01)	.40 (.01)	.14 (.01)
12. Craftsmen	.83 (.01)	.56 (.01)	.62 (.02)	.67 (.01)	.38 (.01)	.46 (.02)
13. Operatives	.80 (.01)	.38 (.01)	.68 (.01)	.62 (.01)	.24 (.01)	.50 (.01)
14. Laborers	.83 (.01)	.60 (.01)	.56 (.02)	.65 (.01)	.42 (.01)	.40 (.02)
Within Selected Detailed Occupation						
15. Engineers, Architects, and Surveyors	.75 (.05)	.71 (.05)	.13 (.03)	.54 (.06)	.50 (.06)	.08 (.03)
16. Mechanics and Repairers	.91 (.01)	.84 (.01)	.45 (.05)	.76 (.67)	.67 (.02)	.27 (.05)
17. Precision Production Occupations	.85 (.01)	.61 (.01)	.60 (.02)	.68 (.01)	.42 (.01)	.45 (.02)
18. Textile, Apparel, and Furnishings Machine Operators	.84 (.02)	.50 (.01)	.67 (.03)	.64 (.02)	.34 (.01)	.46 (.03)
19. Machine Operators, Assorted Materials	.83 (.01)	.56 (.01)	.62 (.01)	.66 (.01)	.38 (.01)	.45 (.01)
20. Fabricators, Assemblers, and Hand Working Occupations	.84 (.01)	.48 (.01)	.69 (.01)	.67 (.01)	.31 (.01)	.52 (.02)
Within Schooling Group						
21. < 12 Years	.80 (.01)	.51 (.01)	.59 (.01)	.62 (.01)	.34 (.01)	.42 (.01)
22. 12 Years	.69 (.01)	.35 (.01)	.53 (.01)	.52 (.01)	.23 (.00)	.38 (.01)
23. 13–15 Years	.63 (.01)	.45 (.01)	.32 (.01)	.45 (.01)	.29 (.01)	.23 (.01)
24. > 16 Years	.65 (.02)	.55 (.02)	.23 (.02)	.47 (.02)	.37 (.02)	.16 (.01)

Notes: Numbers in parentheses are bootstrap standard errors. See text for description of the samples and indices.

presents the results of crude attempts to empirically distinguish between the theories within this class. If segregation arises because men and women have different skills, as the quality-sorting hypothesis suggests, then levels of segregation should be much lower within groups of workers with similar skills. To examine this hypothesis, the rest of Table 2 examines sex segregation among relatively homogeneous groups of workers. Rows 2–8 of Table 2 examine the extent to which interindustry segregation can explain the aggregate segregation patterns of row 1. In particular, in rows 2 and 3 we take a crude attack by simply breaking industries out into durables and non-durables. The indices show that within-industry segregation is very similar to that of the aggregate, so that little aggregate segregation is due to interindustry segregation at this level.

Rows 4–8 of Table 2 examine segregation within a set of detailed industries that are well represented in the WECD.[19] The results show that segregation within detailed industries is generally substantially less than that within manufacturing as a whole. For example, the systematic gini coefficient for the meat products industry is only .28, whereas the corresponding index for nondurable manufacturing as a whole is .45. Similarly, the systematic dissimilarity index is .30 in durable manufacturing as a whole, but it is only .16 in the household appliance industry. These results suggest that much of the aggregate segregation in U.S. manufacturing is due to the sorting of men and women into different detailed industries. In addition, note that systematic segregation is particularly low in the newspaper publishing industry, which employs a relatively large number of white-collar workers. These results suggest a white-collar/blue-collar distinction that we will shortly explore further.

Rows 9–14 of Table 2 address this issue with measures of interfirm segregation within six crude occupational categories. Before we discuss the substance of these rows, note first that there is a mechanical reason why the random segregation indices of columns (2) and (5) are much higher for these "within-occupation" rows than they were for the previous analyses. This difference occurs because the restriction to particular occupations means that we have fewer sample workers per establishment. As the discussion of the previous section pointed out, the random allocation of fewer workers across the same number of employers leads to an increase in random unevenness, and hence to an increase in random segregation.[20]

More substantively, the move to a within-occupation analysis leads to reduced estimates of the systematic component of segregation for some occupations. For example, the systematic gini coefficient for sales and service occupations in our sample is only .21, and the systematic dissimilarity index for professionals, technicians, and managers is only .13. Thus, for these occupations, there is only a limited amount of systematic interplant segregation within U.S. manufacturing. However, the same is not true of more blue-collar occupations. For example, craftsmen, operatives, and laborers are all much more segregated than random allocation would predict. Thus, some but not all of the interplant segregation documented in row 1 is attributable to occupational segregation.

Much of the literature on occupational segregation is concerned with segregation across quite narrowly defined occupations. While a within-occupation analysis of all detailed occupations would be unwieldy, rows 15–20 of Table 2 analyze segregation within a few detailed occupations, which were chosen because they are well represented in our sample. Moving to samples defined by such narrow occupations leads to a small average number of workers per plant, which in turn leads to very high random indices of segregation. Thus, the

[19]We chose these industries because they had a large number of plants in the WECD. Interested readers may obtain results for other detailed industries from the authors.

[20]This illustrates the pitfalls of using traditional segregation indices to compare segregation across different samples. Without accounting for random unevenness, one would conclude that within-occupation segregation was more severe than it was in the entire population, which we believe would be the wrong conclusion.

high traditional segregation indices found in these samples (for example, .91 for the gini coefficient for mechanics and repairers) are largely due to random unevenness. Nevertheless, there is a substantial amount of systematic segregation within these occupations. As before, there is an important blue-collar/white-collar distinction between "engineers, architects, and surveyors" on one hand, and all the other occupations on the other.

As a final exercise, rows 21–24 of Table 2 analyze segregation within groups of workers stratified by educational attainment. The results again suggest that differences in educational attainment do not explain aggregate interplant gender segregation. Of course, this is not surprising, as the educational gap between men and women is not large. More interestingly, the results suggest that there is substantially less systematic segregation among workers with at least some post–high school education. For example, the systematic gini coefficient is .23 for workers with a college degree or more, but it is .59 for high school dropouts. These results are consistent with the results for occupation, where it was the blue-collar occupations that were the most segregated. These results may indicate that educated, white-collar men are less resistant to working with women than are their blue-collar counterparts, or it may be that affirmative action and civil rights pressures to integrate operate more strongly on white-collar workers. Alternatively, it may be that lower job exit rates of educated women lead their human capital to be more substitutable with that of men than is the corresponding human capital of blue-collar women.

We draw the following conclusions from Table 2. Like the authors of previous studies, we find that the distribution of men and women across plants is far from even. However, we show that much of this unevenness is potentially attributable to random allocation, particularly among white-collar workers, and thus that previous studies may have overstated the systematic component of interplant sex segregation. Nevertheless, we also find an important systematic component to observed segregation patterns, even among workers of similar skills. This finding provides support for theories of the male/female wage gap that emphasize segregation.

Differences Between Men's and Women's Employers

The systematic segregation evident in Table 2 indicates that there are systematic forces that drive some plants to hire mostly women while other plants hire mostly men. What distinguishes these two sets of plants from each another? Table 3 addresses this issue with regressions in which the unit of observation is a plant and the dependent variable is the female share of non-supervisory employees. Thus, using these regressions we try to explain why it is that some plants' supervisees are primarily women while other plants' supervisees are primarily men.[21] The right-hand-side variables in the regressions include controls for log plant employment, industry, region, the distribution of workers across broad occupations, the average age and education of supervisory and non-supervisory workers, and the percent of supervisors who are female.[22]

There are theoretical reasons for putting the female share of supervisors on the right side rather than the left side of the regression. Lang's theory clearly implies

[21]To be in the sample used to estimate these regressions, plants had to have at least one supervisor and one non-supervisor in our sample. This led to the exclusion of some of the smaller plants in the sample.

[22]This regression is heteroskedastic for two reasons. First, random allocation implies that the residual variance in this regression will be greater for smaller plants than for larger plants. Second, the dependent variable is bounded between 0 and 1. Since a relatively small fraction of plants are at either bound (16.0% of plants have 0% female nonsupervisors and 2.7% have 100% female nonsupervisors, and even less of the sample is at these bounds when we use sample weights), this does not induce much bias, but it might induce substantial heteroskedasticity. To correct for this problem, reported standard errors are based on White's (1980) method.

Table 3. Plant-Level Models of Employee Sex Composition.

Independent Variable	*(1)*	*(2)*	*(3)*	*(4)*
	Dependent Variable = Female Share of Non-Supervisory Employment			
1. Percent Female Supervisors	.120 (.019)	.102 (.017)	.113 (.019)	.100 (.017)
2. Log of Establishment Employment	.009 (.005)	.022 (.004)	.010 (.005)	.021 (.005)
3. Percent of Non-Supervisors with a College Degree	–.170 (.052)	–.060 (.050)	–.155 (.052)	–.052 (.052)
4. Percent of Establishment Workers in Managerial Occupations	–.017 (.096)	.268 (.090)	–.014 (.096)	.265 (.091)
5. Percent of Establishment Workers in Sales Occupations	.222 (.094)	.286 (.085)	.255 (.096)	.310 (.086)
6. Percent of Establishment Workers in Craft Occupations	–.115 (.084)	–.017 (.075)	–.100 (.084)	–.012 (.075)
7. Log of Average Hourly Wages of Non-Supervisory Employees	—	–.323 (.022)	—	–.319 (.023)
8. Log of Average Hourly Wages of Supervisory Employees	—	.035 (.018)	—	.038 (.018)
9. Log of Labor Productivity (× 100)	—	—	–.019 (.003)	–.008 (.002)
10. 2-Digit Industry Dummies	Yes	Yes	Yes	Yes
11. Region Dummies	Yes	Yes	Yes	Yes
12. Number of Plants in Sample	4,465	4,465	4,359	4,359
13. R-Square	.370	.465	.382	.469

Notes: All data drawn from the Worker-Establishment Characteristics Database. All regressions included controls for the average age and education of supervisors, as well as the average age of non-supervisory employees. To be in the sample for this table, establishments had to have both at least one supervisor and one non-supervisor in the WECD. The data are all weighted by the inverse of the probability that the plant appears in the WECD.

that communications difficulties between managers and workers lead female managers, who presumably speak the "language" of female workers, to employ more female workers. And if female managers are less likely than male managers to have a taste for discrimination against women, then Becker's theory predicts that female managers are more likely to supervise female employees.[23] In contrast, the quality-sorting hypothesis does not predict that otherwise identical male and female managers will systematically supervise workers of either sex. Thus, subject to some provisos, the correlation between the sex of supervisors and the sex of supervisees provides some clue as to the proper interpretation of our results on interplant segregation.

Row 1 indicates a strong correlation between the sex of supervisors and the sex of supervisees. For example, the regression of column (1) indicates that a 50% increase in women's share of a plant's supervisors is associated with roughly a 7.5% increase in women's share of the plant's non-supervisors. Thus, column (1) demonstrates that female managers tend to supervise female employees, although the relationship is far from completely systematic. The coeffi-

[23]While there is not much direct evidence on whether women are less likely than men to discriminate against women, such evidence as exists suggest that there probably is some same sex preference (Goldberg 1968; Ferber and Huber 1975). However, in her case study of a large corporation, Kanter (1977) argued that female managers may suffer from "tokenism" that causes them to be even tougher on female subordinates than are men.

cient in row 2 of column (1) indicates that women have a higher employment share among larger establishments. This result is consistent with the view that large employers are less likely to discriminate than are smaller employers, perhaps because they feel more pressure from federal antidiscrimination law (Carrington and Troske 1995). Rows 3–6 indicate that female non-supervisors are concentrated in plants with relatively few college-educated non-supervisors and in plants with many sales workers, both results that are consistent with differences between men and women in the WECD as a whole. While not reported here, the industry dummies are jointly significant, so that there is an important degree of industrial sorting.

Column (1) of Table 3 is consistent with the discrimination models of Becker and Lang, as female supervisors are more likely to hire and supervise female workers. However, these regressions can still be interpreted in the quality sorting framework without too much trouble. It is true that men and women are allocated somewhat differently within the broad class of supervisory occupations. The inclusion of industry dummies and occupation shares means that column (1) is comparing men and women within fairly narrowly defined industries and within broad occupations, but it remains possible that finer differences in the within-industry occupational structure of men and women are what drive the results. As an example, consider the printing and publishing industry in general, and the newspaper industry in particular. Female managers in this industry may be likely to work in editorial or advertising establishments, while men may be more likely to supervise the manufacture and distribution of the papers. To an extent, one might take this as evidence of discrimination of one sort or another, but it might also be attributed to differences in the type of human capital held by male and female managers and supervisees.

Columns (2)–(4) of Table 3 present three crude ways of trying to refine our results. To the extent that male/female differences in human capital or occupation can be reduced to a scalar measure, as opposed to simple qualitative differences, controlling for the wages or productivity of both supervisors and non-supervisors may provide a cleaner measure of the degree to which the correlation between sex of supervisor and sex of supervisees is driven by discrimination. Column (2) repeats the regression of column (1) with the inclusion of the average wage of non-supervisory employees and, separately, the average wage of supervisors. Row 13 shows that these variables add substantially to the explanatory power of the model, as the r-square increases by roughly a third. The added variables also attenuate the coefficient on percent female supervisors, although it remains moderately large and statistically significant. Column (3) omits the wage variables and instead adds the log of labor productivity, which is defined as the dollar value of plant shipments divided by the number of employees.[24] This variable has much less explanatory power, and it also has less of an effect on the coefficient on percent female supervisors. Finally, column (4) repeats the exercise when both labor productivity and the wage variables are included. The results are very similar to those of column 2. In sum, Table 3 shows that, holding the wages and productivity of their supervisees constant, men are more likely to supervise other men and women are more likely to supervise other women. This suggests that to the extent that human capital differences explain interplant gender segregation, it must be along qualitative rather than quantitative dimensions.

These results suggest that the systematic segregation documented in Table 2 is at least in part the result of supervisors' systematic choice of supervisees of their own sex. These choices may arise from discrimination due to prejudice, as in Becker's model, or from discrimination due to language differences between the sexes. Of course, given the relatively crude occupa-

[24]Similar results were obtained when we defined labor productivity to be value-added/employees.

tional classification system considered in the analysis, it remains possible that occupational segregation along finer dimensions is what drives the observed patterns of interplant segregation. However, together with findings of previous research by Bielby and Baron (1984), Groshen (1991), and others, the results we report here suggest that, even within narrow occupations, there is usually some interplant gender segregation.[25] Thus, we believe that there is an important systematic component of interplant gender segregation, particularly among less-educated and blue-collar workers, and that this segregation is probably not completely explained by sex differences in human capital acquisition.

Before moving on, let us note again that the preceding analysis was conducted on a sample of workers who worked at or near full-time. We focused on this sample because discrimination against full-time workers is particularly troubling, and because segregation in a broader sample might be partly due to segregation of part-time and full-time workers. However, since women are disproportionately part-time and because part-time/full-time segregation might itself reflect sex discrimination, it is worth comparing the above results to those obtained in a sample that includes part-time workers. In short, it makes very little difference in the analysis.[26] There are of course some minor changes in overall measures of segregation, and even some more significant changes in segregation within particular industries or occupations. However, the major results still hold: men and women are systematically segregated, particularly in blue-collar occupations and industries.

Interplant Segregation and the Male/Female Wage Gap

In this section we investigate two aspects of the relationship between interfirm segregation and the male/female wage gap. We first decompose the gender wage gap into between- and within-plant components. Our decompositions differ from those of earlier studies (for example, Groshen 1991) in that we control for human capital characteristics such as education and experience that have been unobservable to previous authors. The exercise is motivated by public policies such as Title VII and comparable worth that are largely devoted to eliminating within-plant pay differences between equally qualified men and women. As mentioned previously, these policies will not be effective unless within-plant pay differences are the primary source of women's low wages.

Our approach is to regress wages on a set of plant fixed effects, either before, after, or at the same time that we control for workers' personal characteristics, and to see how much of the gender wage gap can be explained by the location of men and women in different plants.[27] Let Y = log hourly wages, let X = a set of personal characteristics including terms in education, experience, and detailed occupation, and let Z = a set of plant fixed effects. Columns (1)–(3) of Table 4 then report results from a two-step procedure in which we first estimated $Y = X'\beta$, and then regressed the residuals of this first regression on the plant fixed effects Z. In essence, this procedure gives personal characteristics first crack at explaining the gender wage gap, and the plant effects are given the opportunity to explain the residual gender wage gap. Column (1) reports results for all workers, while columns (2) and (3) report separate results for blue-collar and white-collar workers. Row 1 shows that the unadjusted difference in log hourly wages

[25]One problem with these previous studies is that usually little effort was made to distinguish systematic from random segregation. Thus, some of the patterns of intraoccupational sex segregation found by previous authors could be due to random allocation of workers.

[26]Full results from the sample that includes part-time workers are available from the authors on request.

[27]By plant fixed effects we mean we have included a separate dummy variable for each plant in our data.

Table 4. Decomposing the Male/Female Wage Gap into Within- and Between-Plant Components.

Order of the Decomposition	Step 1: $Y = X'\beta + u$, Step 2: $Y - X'\hat{\beta} = Z'\gamma + u$ (1) Total	(2) White-Collar	(3) Blue-Collar	Step 1: $Y = Z'\gamma + u$, Step 2: $Y - Z'\hat{\gamma} = X'\beta + u$ (4) Total	(5) White-Collar	(6) Blue-Collar	Step 1: $Y = X'\beta + Z'\gamma + u$ (7) Total	(8) White-Collar	(9) Blue-Collar
1. Unadjusted Male/Female Wage Gap $\bar{Y}_m - \bar{Y}_w$	.429	.439	.446	.429	.439	.446	.429	.439	.446
2. Male/Female Wage Gap Adjusted for Personal Characteristics $(\bar{Y}_m - \bar{X}_m'\hat{\beta}) - (\bar{Y}_w - \bar{X}_w'\hat{\beta})$	.317	.280	.371	.366	.344	.416	.345	.302	.400
3. Male/Female Wage Gap Adjusted for Plant Fixed Effects $(\bar{Y}_m - \bar{Z}_m'\hat{\gamma}) - (\bar{Y}_w - \bar{Z}_w'\hat{\gamma})$	.342	.332	.194	.230	.285	.143	.275	.352	.197
4. Male/Female Wage Gap Adjusted for Personal Characteristics and Plant Fixed Effects $(\bar{Y}_m - \bar{X}_m'\hat{\beta} - Z_m'\hat{\gamma}) - (\bar{Y}_w - \bar{X}_w'\hat{\beta} - Z_w'\hat{\gamma})$	.230	.173	.119	.167	.190	.123	.191	.214	.151

Y = hourly wages

X = worker characteristics including flexible terms in experience and education, race, marital status, 10 occupation dummies, 4-digit industry dummies, MSA, region, MSA × region.

Z = a set of plant fixed effects.

between men and women is .429 for all workers, .439 for white-collar workers, and .446 for blue-collar workers in our sample.

Row 2 of Table 4 reports the residual wage gap left after we control for the effect of personal characteristics. This is approximately .280 for white-collar workers and approximately .371 for blue-collar workers. Row 3 reports the residual male/female wage gap after netting out the fixed effects estimated in the second step, without first adjusting for personal characteristics, and row 4 reports the residual wage gap after we net out both the personal characteristics and the plant fixed effects. Two facts stand out from these calculations. First, it is clear that plant fixed effects can statistically account for a substantial fraction of the male/female wage gap, both as a whole and separately for the blue- and white-collar samples, even after we control for a large array of other productive characteristics. Second, the role of the plant fixed effects is much more important for blue-collar workers than it is for white-collar workers. For white-collar workers, column (2) shows that human capital characteristics can account for 36% of the gender gap in log hourly wages. Subsequent controls for employer can account for an additional 24% of the gap so that, among white-collar workers, employer identity plays an important but secondary role in wage determination. In contrast, column (3) shows that employer identity plays a dominant role in statistically explaining the gender wage gap among blue-collar workers. This can be seen by the fact that while human capital characteristics can explain only 17% of the gender wage gap, plant fixed effects can account for almost 68% of the residual wage gap. Thus, among blue-collar workers, the male-female wage gap is largely accounted for by women's location in plants that generally pay low wages.

Columns (4)–(6) reverse the order of the decomposition by regressing wages on the plant fixed effects in the first step, and then regressing the residuals on the personal characteristics in the second step. This gives the plant effects first crack at explaining the male/female wage gap. Row 3 of column (4) shows that the male/female wage gap remaining after controlling for plant fixed effects is .230 for all workers, or about 52% of what it was without controlling for these effects. Column (6) shows that the role of plant fixed effects is particularly strong among blue-collar workers, where they explain over two-thirds of the male/female wage gap. Columns (7)–(9) regress *Y* on *X* and *Z* simultaneously, so that personal characteristics and plant fixed effects are given equal opportunity to explain the gender wage gap. The results are similar to those of the previous columns, as plant fixed effects explain a substantial fraction of the male/female wage gap.

The results of Table 4 show that there is an important, and in the case of blue-collar workers dominant, role played by interplant segregation in accounting for the wage gap between men and women. Women work in low-paying plants while men work in plants that pay relatively high wages. While broadly similar to the results of previous studies (for example, Groshen 1991), these results command more confidence because we have simultaneously controlled for human capital characteristics that were unobservable to previous analysts. In addition, we are unaware of any previous distinction between blue- and white-collar workers in this regard. Comparable worth policies are designed to even out interoccupational wage differences within firms, and by extension within plants. The important role of interplant pay differences suggests that such policies can have only a limited effect on the male-female wage gap, even if they are effective at evening pay differences within firms. In addition, the results are consistent with discrimination theories positing that women are crowded into a few non-discriminatory employers.

Our second exercise relates the wages of male and female workers to the gender makeup of their coworkers. This question is motivated by several observations. First, Becker's theory of coworker discrimination (although not his theory of employer discrimination) posits that workers in integrated plants may receive higher wages. For example, if male workers demand a higher wage in order to work with women, then men in integrated plants will receive higher wages than similar men working in all-male plants. A second motivation is the emphasis of the gender wage gap literature on the interoccupational relationship between average wages and gender composition. While this focus does not have a clear theoretical basis, our interfirm analysis of gender composition and wages is directly analogous.

Our approach here is to estimate an individual-level hourly wage regression with personal and plant characteristics on the right-hand side. The personal characteristics include terms in experience and education, sex, marital status, occupation, and race, and the plant characteristics include total employment and the female share of establishment employment. Table 5 presents the results of several such regressions.[28] Columns (1) and (2) report results for all workers. The columns differ in that column (1) includes only the variables listed above, while column (2) adds a measure of plant-level labor productivity.[29]

Both regressions indicate that, holding characteristics including sex constant, workers earn less if they work in plants with largely female workers. For example, row 2 of column (2) indicates that the representative man working in a 50% female plant

[28]The standard errors in these regressions have been corrected to account for heteroskedasticity and for the clustered sample design.

[29]Labor productivity is defined to be the dollar value of shipments divided by employment.

Table 5. Individual-Level Models of Hourly Wage Determination.

	Sample											
	By Industry						By Occupation					
	All Workers		Female-Intensive		Male-Intensive		Managerial		Sales		Laborers	
Independent Variable	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*	*(8)*	*(9)*	*(10)*	*(11)*	*(12)*
1. Female	−.083	−.086	−.090	−.091	−.083	−.087	−.072	−.074	−.111	−.109	−.050	−.054
	(.010)	(.010)	(.016)	(.016)	(.015)	(.015)	(.017)	(.017)	(.014)	(.014)	(.015)	(.015)
2. Female Share of Establishment Employment	−.214	−.190	−.254	−.218	−.188	−.169	.025	.031	−.178	−.159	−.297	−.268
	(.021)	(.021)	(.028)	(.028)	(.030)	(.029)	(.026)	(.025)	(.026)	(.026)	(.025)	(.025)
3. Female × Female Share of Establishment Employment	−.091	−.090	−.047	−.053	−.113	−.108	−.164	−.162	.071	.064	−.128	−.127
	(.022)	(.021)	(.030)	(.029)	(.039)	(.037)	(.036)	(.036)	(.027)	(.026)	(.029)	(.029)
4. Log of Establishment Employment	.068	.065	.062	.060	.071	.068	.035	.034	.056	.054	.083	.079
	(.004)	(.004)	(.006)	(.006)	(.005)	(.005)	(.003)	(.003)	(.003)	(.003)	(.005)	(.005)
5. Labor Productivity (× 1000)	—	.258	—	.312	—	.233	—	.106	—	.219	—	.301
		(.028)		(.040)		(.030)		(.028)		(.029)		(.033)
6. R-Square	.497	.504	.469	.475	.448	.457	.412	.414	.427	.433	.477	.488
7. Number of Observations	169,099		63,950		105,149		27,512		31,470		110,117	

Notes: The dependent variable is the log of the hourly wage, which is defined as annual earnings divided by annual hours of work. Each regression also included as regressors the following variables: a quartic in experience, five dummy variables for educational attainment, interactions of all education and experience terms, and dummy variables for race, marital status, 1-digit occupation, region, msa residence, and 2-digit industry. Standard errors have been corrected for heteroskedasticity and for the clustered sample design. All data are drawn from the Worker-Establishment Characteristics Database (WECD).

earns wages that are more than 10% lower than those of observationally equivalent men working in all-male plants. The combination of rows 2 and 3 shows that the relationship between wages and percent female in the establishment is even greater among women. The regression indicates that the representative woman in a 50% female plant earns wages that are over 15% lower than those of observationally equivalent women working in plants that are nearly all male. A comparison of columns (1) and (2) shows that these relationships are not much changed if we control for (admittedly crude) measures of labor productivity.

The regressions of columns (1) and (2) control for 1-digit occupation and 2-digit industry (within manufacturing, of course), so these results are not driven by the fact that men and women are sorted into broadly different occupations and industries. Nevertheless, it is possible that segregation along finer dimensions of occupation and industry is what drives these results. For example, the quality-sorting hypothesis suggests that, within any 2-digit industry, plants with many women tend to employ workers with relatively low skills. Alternatively, it could be that the mere fact of having female coworkers tends to drive down wages. There is no completely satisfactory way of sorting out these alternative interpretations, but the rest of Table 5 represents several crude attempts at doing so.

Columns (3)–(6) of Table 5 present wage regressions where the samples are divided along industrial lines. Columns (3) and (4) report regressions for industries that are relatively female-intensive (that is, the fraction of female employees in our sample exceeds 33%), while columns (5) and (6) report regressions for industries that are relatively male-intensive. The motivation for these regressions is that there might be differential scope for quality-sorting depending on whether women are an important component of the industry's overall work force. The results of these regressions may be summarized as being not too different from those of columns (1) and (2). Namely, workers with female coworkers earn lower wages than observationally similar workers with predominantly male coworkers.

Columns (7)–(12) take an analogous approach to samples stratified by broad occupation. The results display substantial variation across occupations. Columns (7) and (8) show that, among managers, there is no wage penalty for male managers who work with predominantly female underlings. However, female managers in predominantly female plants do earn substantially less than female managers in more integrated plants. Columns (9) and (10) show that, among sales and clerical occupations, both men and women are penalized for working in largely female plants, but the size of this penalty is substantially less for women. Despite this variation, there is still a common theme. Wages for both men and women tend to be lowest in those plants with a predominantly female workplace.

These results are somewhat hard to explain with any of the discrimination theories. In Becker's theory of employer discrimination and in Lang's language theory, competition forces the wages of workers within either sex to be the same across plants regardless of the plants' gender composition. In Becker's theory of coworker discrimination, prejudiced men should receive higher wages for working with predominantly female coworkers. None of these predictions are consistent with the data. In contrast, the quality sorting hypothesis suggests that highly female plants pay lower wages simply because their employees, both male and female, tend to have lower skills. This is roughly consistent with what we find here. The quality sorting hypothesis also suggests that there should be no gender wage gap once we control for labor productivity, but this is not what we find. However, our controls for labor productivity are sufficiently crude that this does not represent a strong test. Thus, of the theories considered here, quality-sorting is perhaps the most consistent with the data.[30]

[30]As with the results on segregation, we examined the sensitivity of these results to the inclusion of part-time workers in the sample. Not surprisingly, there

Conclusions

This paper has examined the extent and causes of interfirm gender segregation in U.S. manufacturing. A primary finding is that, consistent with earlier results, there is substantial interfirm sex segregation by conventional measures. However, some of this segregation is attributable to the random allocation of workers to plants, so previous studies may have overstated the systematic component of workplace gender segregation. This is particularly true among white-collar, highly educated workers. Nevertheless, among white-collar workers and particularly among blue-collar workers, we find that men and women are more segregated than random allocation would predict. A secondary finding is that managers and their subordinates tend to be of the same sex, even within industries and occupations. While this may be partly the result of qualitative differences in human capital acquisition, it is also consistent with theories of discrimination based on animus (Becker) or language (Lang).

We also examined the distribution of men's and women's wages across U.S. manufacturing plants. We found that differential pay across establishments can statistically account for a substantial fraction of the overall gender pay gap. Indeed, among blue-collar workers plant fixed effects explain more of the gender pay gap than do a full array of traditional human capital measures. Among other things, this finding suggests that comparable worth policies can have only limited success in reducing the gender pay gap. We also examined the relationship between wages and the fraction of female workers in a plant, holding the worker's sex and other characteristics constant. While there is some variation across occupations and industries, the basic finding is that both men and women earn less when they work in plants that are predominantly staffed by women. We argue that this finding is not explicable by theories of discrimination. The quality-sorting hypothesis of Macpherson and Hirsch (1995) offers one potential explanation of these facts, as perhaps low-skilled men and women are all concentrated in highly female plants. However, the fact that controlling for labor productivity leaves this relationship largely intact suggests that the quality-sorting hypothesis is not a complete explanation either.

are some minor changes in coefficients in the models that we estimate. However, the general results are quite unaffected by the inclusion of part-time workers. Full results from the sample including part-time workers are available from the authors upon request.

REFERENCES

Becker, Gary S. 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.

Bielby, William, and James Baron. 1984. "A Woman's Place Is with Other Women: Sex Segregation within Organizations." In Barbara Reskin, ed., *Sex Segregation in the Workplace: Trends, Explanations, Remedies*. Washington, D.C.: National Academy Press, pp. 27–55.

Blau, Francine D. 1977. *Equal Pay in the Office*. Lexington, Mass.: D.C. Heath.

Boisso, Dale, Kathy Hayes, Joseph Hirschberg, and Jacques Silber. 1994. "Occupational Segregation in the Multidimensional Case." *Journal of Econometrics*, Vol. 61, pp. 161–71.

Brown, Charles, James Hamilton, and James Medoff. 1990. *Employers Large and Small*. Cambridge, Mass.: Harvard University Press.

Buckley, John. 1971. "Pay Differences Between Men and Women in the Same Job." *Monthly Labor Review*, Vol. 94, No. 11, pp. 36–39.

Carrington, William J., and Kenneth R. Troske. 1995. "Gender Segregation in Small Firms." *Journal of Human Resources*, Vol. 30, No. 3, pp. 503–33.

____. 1997. "On Measuring Segregation in Small Unit Samples." *Journal of Business and Economic Statistics*, Vol. 15, No. 4, pp. 402–9.

Doms, Mark, Timothy Dunne, and Kenneth R. Troske. 1997. "Workers, Wages, and Technology." *Quarterly Journal of Economics*, Vol. 112, No. 1, pp. 253–90.

Donohue, John J. III, and Peter Siegelman. 1991. "The Changing Nature of Employment Discrimination Litigation." *Stanford Law Review*, Vol. 43, pp. 983–1033.

Duncan, Otis, and Beverly Duncan. 1955. "A Methodological Analysis of Segregation Indices." *American Sociological Review*, Vol. 20, No. 2, pp. 210–17.

Efron, Bradley, and Robert J. Tibshirani. 1993. *An Introduction to the Bootstrap.* New York: Chapman & Hall.

Ferber, Marianne Abeles, and Joan Althaus Huber. 1975. "Sex of Student and Instructor: A Study of Student Bias." *American Journal of Sociology,* Vol. 80, No. 4, pp. 949–63.

Goldberg, P. A. 1968. "Are Women Prejudiced Against Women?" *Trans-Action,* Vol. 5, pp. 28–30.

Griffin, Peter, and Stephen Trejo. 1995. "The Measurement of Gender Segregation Using Firm-Level Data." Long Beach: Working Paper, California State University.

Groshen, Erica L. 1991. "The Structure of the Female/Male Wage Differential." *Journal of Human Resources,* Vol. 26, No. 3, pp. 457–72.

Hutchens, Robert M. 1991. "Segregation Curves, Lorenz Curves, and Inequality in the Distribution of People Across Occupations." *Mathematical Social Sciences,* Vol. 21, pp. 31–51.

Johnson, George, and Gary Solon. 1986. "Estimates of the Direct Effects of Comparable Worth Policy." *American Economic Review,* Vol. 76, No. 5, pp. 1117–25.

Kanter, Rosabeth Moss. 1977. *Men and Women of the Corporation.* New York: Basic Books.

Kremer, Michael, and Eric Maskin. 1995. "Segregation by Skill and the Rise in Inequality." Cambridge, Mass.: Working Paper, Massachusetts Institute of Technology.

Lang, Kevin. 1986. "A Language Theory of Discrimination." *Quarterly Journal of Economics,* Vol. 101, No. 2, pp. 363–82.

Macpherson, David A., and Barry T. Hirsch. 1995. "Wages and Gender Composition: Why Do Women's Jobs Pay Less?" *Journal of Labor Economics,* Vol. 13, No. 3, pp. 426–71.

McNulty, Donald. 1967. "Differences in Pay Between Men and Women Workers." *Monthly Labor Review,* Vol. 90, No. 12, pp. 40–43.

Pfeffer, Jeffrey, and Alison Davis-Blake. 1987. "The Effect of the Proportion of Women on Salaries: The Case of College Administrators." *Administrative Science Quarterly,* Vol. 32, No. 1, pp. 1–24.

Reilly, Kevin T., and Tony S. Wirjanto. 1997. "Does More Mean Less? The Male/Female Wage Gap and the Proportion of Females at the Establishment Level." Working Paper, University of Waterloo.

Troske, Kenneth R. Forthcoming. "The Worker-Establishment Characteristics Database." In John Haltiwanger, Marilyn Manser, and Robert Topel, eds., *Labor Statistics Measurement Issues.* Chicago: NBER.

White, Halbert. 1980. "A Heteroscedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroscedasticity." *Econometrica,* Vol. 48, pp. 817–38.