

RACE AND GENDER DIFFERENCES IN WAGES: The Role of Occupational Sorting at the Point of Hire

Andrew M. Penner*

University of California, Irvine

Previous work suggests that occupational sorting and differences in starting salary play a large role in race and gender wage gaps. This study uses unique data from the human resources department of a financial company to examine the role of occupational sorting in race and gender differences in initial salary offers. While this company exhibits large race and gender differences in salary offers, controlling for occupational differences accounts for all of the race effects and reduces the gender effect to substantive insignificance. These findings underscore the importance of occupational sorting mechanisms in creating race and gender differences in wages.

INTRODUCTION

The existence of race and gender differences in labor market outcomes is well established; women are typically found to earn 25 percent less than men, and blacks and Hispanics earn, respectively, 22 and 32 percent less than whites (Bielby 2000). These differences are important not only in and of themselves, but also because of the role they play in larger systems of inequality. There is thus a great deal of interest in assigning responsibility and blame to the source of these differences, and, in particular, to understanding the role that employers might play in creating these inequalities.

The impetus to understand the role that employers play in creating inequality is partially driven by the desire to gain a better understanding of inequality writ large, but also arises from concerns over discrimination. Given the importance of the issue, it is not surprising that there is a substantial literature on how firms are implicated in producing inequality. Studies often examine questions surrounding hiring, promotions, and terminations in an attempt to understand how firms create inequality in these arenas and how this impacts wage differences. This literature suggests that wage differences found at later stages in careers can be traced back to differences at the point of hire (e.g., Gerhart 1990). Hiring then becomes central to our understanding of inequality in the labor market and particularly to understanding the role that employers play in its creation.

Typically, the hiring process at a large firm can be broken down into a number of stages.¹ The first stage consists of the creation of an applicant pool and involves advertising the position and recruiting potential applicants. Next, the applications received are screened, and human resource personnel decide whom to interview. On the basis of this

*Direct all correspondence to Andrew M. Penner, Department of Sociology, University of California, Irvine, Irvine, CA 92697-5100; e-mail: andrew.penner@ucl.edu

interview with human resource personnel, applicants are invited for a second interview, this time by their potential supervisor and/or colleagues. Upon the successful completion of this process, applicants are extended job and wage offers. Given the importance of differences at the point of hire, it is important to understand how it is that this hiring process is implicated in the creation of wage inequality.

Surprisingly little is known about how wage inequality at the point of hire is created, and even less is known about the role that employers might play in this process. This is, in part, because finding data to address these issues is exceedingly difficult. While there are several studies that examine the hiring process in a small number of occupations at a single firm, understanding how wage inequality is created involves understanding how it is that individuals are sorted into different occupations. This sorting into different occupations is important for wage inequality because within any given occupation there is typically little variation in pay, and most pay differences are between occupations. Understanding the role of employers in creating inequality thus becomes an issue of understanding what role employers play in matching people to jobs, particularly at the point of hire.

Approaching the problem from the angle of causal factors reveals two sources besides employers: individual job seekers and society. Although individual job seekers, employers, and society undoubtedly interact in complex ways, for heuristic purposes, it is helpful to think of them each independently.

Individual Job Seekers

At the micro level, it is possible to argue that labor market differences are a function of individuals' choices. It is plausible that individual men and women, or blacks, whites, Asians, and Hispanics, have different tastes and preferences influencing the jobs they desire and pursue. The strength of this argument is that it is almost certainly true that at some level, individual choices regarding education, occupation, and effort play a role in creating labor market differences. However, these choices are not made in a vacuum, so that even if inequality is entirely a mere aggregation of job seekers' choices, factors like organizational policies and social context are still important. Thus, where individual job seekers are concerned, even where it is the aggregation of their preferences that results in labor market differences, research often focuses on the factors shaping and constraining their choices and preferences (e.g., Kanter 1977; Correll 2001, 2004).

Employers

Substantially more attention has been paid to the ways that firms can create differences in wages, with much of the focus surrounding issues of discrimination (see, e.g., Rees and Schultz 1970; Becker 1971; Kirschenman and Neckerman 1991; Petersen and Saporta 2004; Pager and Quillian 2005). There are four different areas in which we can think of wage differences arising from employer discrimination. First, employers could pay different wages to employees of different genders and races who do the same work. Second, employers could create wage differences by hiring different applicants for different types of jobs that pay different wages. Third, like hiring discrimination,

discriminatory promotional practices could also affect the jobs available to women and minorities. Finally, recruiting and advertising strategies could be used to yield certain types of applicants for certain jobs, thus producing the same effect as hiring or promotion differences. While discrimination could potentially play a substantial role in creating inequality, it is difficult to establish that discrimination exists, as the data required to differentiate between applicant and firm actions are substantial.

Society

There are three main ways in which society is viewed as contributing to gender and race differences in wages. First, society shapes individual preferences and affinities. Second, society and social structures provide the context for the decisions that individuals make, so that biases in the education system and policies like parental leave affect career choices and trajectories. While these two societal mechanisms impacting labor market differences are closely related, the distinction between them is that the first acts through a process of internalization and changes people, whereas the second merely changes the external constraints in which people make their choices. Finally, there is what is referred to by Petersen and Saporta (2004) as *valuative discrimination*, where work that women and minorities do is systematically undervalued. Theories of *valuative discrimination* link the wages of an occupation to the composition of the workers employed in that occupation, arguing that shifts in workforce composition are accompanied by shifts in the wage. This is the most macro level of the three societal considerations, and is also the most explicitly studied (e.g., England 1992).

Given that the number of mechanisms contributing to differences intertwined in any particular situation makes disentangling them exceedingly difficult, research examining race and gender differences in labor market outcomes has often focused on where differences arise to gain insight into what axes are important. While this does not ultimately answer questions of causality, it does provide an indication of what factors are likely to play a vital role. For example, Petersen and Morgan's (1995) finding that gender differences are due primarily to occupational sorting suggests that at the individual level, occupational preferences and preparation are important, at the firm level, hiring and promotional practices of firms are important, and at the societal level, stereotypes about masculine and feminine occupations are important.

What Is Known and Unknown

Previous research has shown that women and minorities are typically at a disadvantage in the labor market. Minorities are less likely to be hired (Kirschenman and Neckerman 1991; Pager 2003), earn less (Tomaskovic-Devey 1993; Bielby 2000), and are more likely to be laid off (Sicherman 1996; Elvira and Zatzick 2002).² Further, while studies disagree over whether minorities are promoted at lower rates than whites (Powell and Butterfield 1997; Pergamit and Veum 1999), even when they are promoted at similar rates, minorities enter at lower positions and are thus underrepresented in higher-level positions (DiPrete 1989). Findings on gender differences are somewhat more complicated. The considerable male advantage in earnings is due almost entirely to the occupational

sorting processes that result in men and women working in different jobs so that when men and women do the same work for the same employer, men earn only slightly more than women (Petersen and Morgan 1995). Women are found less likely to be hired in some arenas (e.g., Neumark 1996), but more likely to be hired in others (e.g., Fernandez and Weinberg 1997). While women are underrepresented in top-level jobs, some studies find that once occupational level and dead-end career ladders are accounted for, women are actually promoted at a higher rate than men (DiPrete 1989; Spilerman and Petersen 1999; Petersen and Saporta 2004). Finally, studies find that women are not more likely to be fired (Elvira and Zatzick 2002; Petersen and Saporta 2004), although they are more likely than men to quit for nonmarket reasons (Sicherman 1996).

In addition to establishing the dimensions of race and gender labor market inequality, much of the current literature is concerned with understanding how firm-level processes create and maintain these differences. Studies examining career trajectories within firms suggest that wage differences can largely be traced back to differences at the point of hire (DiPrete 1989; Gerhart 1990; Petersen and Saporta 2004). While research finds that race and gender wage differences in the overall labor market are largely attributable to occupational sorting (Tomaskovic-Devey 1993; Petersen and Morgan 1995), the effect of occupational sorting on wage differences at the point of hire has yet to be examined. Given the centrality of hiring in questions of wage inequality, this constitutes a key omission, and it is this omission that serves as the focus of the current study. In looking at occupational sorting at the point of hire, this study allows us to see whether wage differences in this context arise from differences in the jobs people are hired for or from differences in pay to people in the same jobs. While a number of recent firm case studies consider the impact of race and gender in the hiring process and provide insight into *how* people are hired into certain jobs (e.g., Fernandez and Sosa 2005; Petersen, Saporta, and Seidel 2005), little is known about the *effects* of occupational segregation on wage differences at the point of hire.

Although understanding the effects of occupational sorting at the point of hire is important, a common weakness of studies examining the hiring process in firms is that they examine only a small number of occupations. As a result, while these studies are informative about the job-matching processes that impact occupational sorting at the point of hire, they cannot provide an indication of the effect of occupational sorting on wages at the point of hire. For example, Petersen et al. (2005) analyze hiring into five occupations, while Fernandez and Sosa (2005) look at hiring into one occupation. By contrast, this study analyzes differences in the salary offers made to successful applicants from 866 different detailed occupational codes. In analyzing differences across so many occupations, I am able to examine how differences in occupation impact differences in salary offers. However, as I only have data on successful applicants, I am unable to determine whether sorting is due to the actions of applicants or employers. This study thus addresses the questions of whether and how much occupation sorting matters at the point of hire, and not the question that previous studies have addressed, which is what processes might be involved in occupational sorting at the point of hire.

Framework

While it is important to think about causal agents, in practice it has proved quite difficult to assign causality to a specific actor, as potential causes are plentiful and adjudicating between them is difficult. Thus, although this study focuses on a particular company and examines the role of occupational sorting in creating inequality, I focus more on where inequality comes from than on who creates it. This study, then, seeks to identify more precisely where differences arise to gain insight into their nature.

To do this, I identify differences at the point of hire as a key factor in wage inequality, and identify to what degree these differences are accounted for by occupational sorting as opposed to within occupation differences in salary offers. In addition to looking at differences in salary offers, I am also able to observe the race and gender differences in market salary rates, which indicate how much the firm expects to pay employees in different jobs. Because the market salary rate of the job is established before the job opening is announced and is thus not affected by characteristics of the applicants, examining race and gender differences in market salary rates provides insight into the importance of the job-matching process. This study is, as far as I know, the first study to analyze differences in a firm's market salary rate.

DATA AND METHODS

The Company

This study uses human resource data on job openings in multiple occupations from one firm to examine wage differences at the point of hire. The data come from a regionally based Midwestern financial company that has offices throughout the United States. The company had on the order of 20,000 employees during the period of this study, January 1999 to June 2001, and their recruitment and hiring practices are standard for the industry. As the data are from the central human resources department at the firm, they are extremely detailed and contain information on the entire pool of job openings and offers during this period.

These data are comparable in scope to data used in other case studies like Fernandez and Weinberg (1997) and Petersen et al. (2005). Like those case studies, these data provide detailed information from the human resource records of a large company. In the context of this analysis, the data have two shortcomings. First, while the data include information on all job offers, they are missing information about rejected applications. Thus, this study cannot examine how occupational sorting occurs, only its consequences. That is, while ideally an examination of occupational sorting would observe how a pool of potential applicants chooses what jobs to apply for, and then how employers choose which applicants to hire, these data can only address how the outcome of this process affects wage differences. Second, these data do not include any human capital variables, such as education or experience.

For this analysis, only cases where data from the file on applications matched data from the file on job openings were kept, providing complete job data. Data were additionally restricted to include only cases with complete information on race, gender, and

salary offer amount. Further, as internal applicants are no longer at the point of hire and are often seeking promotions, this analysis is restricted to external applicants. The final data set includes 11,537 job openings in 866 occupations,³ and 11,643 offers to 11,459 individuals. It is worth noting that the unit of analysis is the offer, so that some individuals are counted multiple times.⁴

Dependent Variables

The analyses in this article focus on two dependent variables, the annual salary offered and the market salary rate. While the salary offer is fairly straightforward, the market salary rate is worth explaining further.

The market salary rate can be thought of as the firm's assessment of the going annual salary rate for that occupation. There are 80 market salary rates in the data, ranging from \$12,000 to \$200,000. The average salary offer is 87 percent of its market rate, and only 20 percent of offers are above their market rate. Thus, the market salary rate is neither an average of the offers nor an upper bound for salary offers, but falls somewhere in between. As the market salary rate is set before applicants apply, differences in market rates can only result from the job-matching process.⁵ This study not only analyzes race and gender differences in market salary rates, but also introduces the market salary rate as an independent variable in models of salary offers.

Both market salary rate and salary offer amounts are transformed by taking their natural logarithms. Modeling the natural logarithms of the amounts allows for the coefficients to be interpreted as approximating proportional changes. That is, a coefficient for female of $-.05$ would indicate that women do approximately .05 (5 percent) worse than men.⁶ In the models where the log of the market salary rate is predicting the log of the salary offer, the coefficient P for the logged market salary rate can be interpreted as an elasticity. That is, a 1-percent increase in market salary rate predicts a P -percent increase in salary offered.

Independent Variables

The independent variables in this study include basic demographic characteristics such as race and gender, the applicant's source, and the Equal Employment Opportunity (EEO) category of the occupation for which they received an offer. While the demographic variables are self-explanatory, the other two merit explanation.

Applicant source categorizes applicants according to how they heard about the job and contains seven categories: referral, walk-in, advertisement, Internet, temporary agency, search firm, and other, which includes methods like college recruiting fairs and phone inquiries. This variable is introduced primarily to serve as a control variable, but is also of interest as it relates to a substantial literature that is concerned with the use of referral networks in negotiating the job-finding process (Granovetter 1995; Fernandez and Weinberg 1997; Fernandez, Castilla, and Moore 2000; Petersen et al. 2000). Previous research on networks and hiring suggests that in some settings referral networks enable applicants to better navigate the hiring process, and thus contribute to race differences in initial salary (Petersen et al. 2000). However, it is unclear how these networks impact

occupational sorting, as Fernandez and Sosa (2005) find that networks exacerbate the male disadvantage in hiring rates for a typically feminine job at a call center, while Fernandez and Fernandez-Mateo (2006) show that networks are all but orthogonal to race differences in hiring rates at a different firm.

EEO codes represent nine broad occupational categories for which companies are required to keep records. Although not very detailed, they allow for a presentation of descriptive statistics broken down across rough occupational categories. Of the nine EEO categories, this company had job openings in six: officials and managers, professionals, technicians, sales workers, office and clerical, and service workers. Office and clerical workers account for over half of the job offers, and are thus used as the baseline category in the models. Service worker jobs, on the other hand, account for only 33 offers, and so while they are included in the analyses, I do not report findings separately for them.

Limits of the Firm Case Study Approach

A drawback of firm case studies approach is that the ability to make generalizations is sacrificed for detail. While it is possible to have a high level of detail in a larger context (see, e.g., DiPrete 1989), this rarely occurs because such highly detailed data are difficult to obtain, and finding comparable datasets with the same variables is difficult. Generalizing from human resource data is further problematic, because the firms that allow researchers access to their data could differ substantially from firms that do not. On the positive side, human resource data provide great depth in examining labor market differences, as they provide the researcher a record of what the employer saw and how they acted. In contrast to audit data, which have similar capabilities but are explicitly constructed, a strength of human resource data is that they arise naturally. This allows human resource data to more readily provide information on a greater variety of jobs, including jobs higher up in the corporate hierarchy, and to provide information about actual hires and salary offers.

Models

The models presented in this article are estimated using ordinary least squares regression, and dependent variables are transformed by taking their natural logarithm so that the results approximate proportion differences. The models examine occupational sorting in three ways. First, I introduce dummy variables for EEO codes. This introduces controls for broad occupational sorting, so that the race and gender effects estimated report the differences within EEO codes. The second model seeks to understand what it is about occupational sorting that is important by controlling for the market salary rate of the occupation. This shows to what degree occupational sorting is a function of differences in the pay rate anticipated by the firm. Finally, I estimate models with fixed effects for occupations (866 groups) to obtain the race and gender differences comparing only offers within the same detailed occupation. The results from these three analyses are compared with the coefficients from the baseline model.

TABLE 1. Descriptive Statistics

	Percentage	Market salary rate (\$)		Salary offer (\$)	
		Mean	SD	Mean	SD
Male	29	42,782	24,277	40,121	27,501
Female	71	28,603	12,833	24,387	13,457
White	72	34,997	19,348	31,187	21,521
Black	19	25,236	10,398	21,468	10,393
Hispanic	4	29,506	14,481	25,982	16,291
Asian	5	31,435	18,633	28,269	20,769
Overall		32,774	18,202	29,015	20,042
N	11,643	11,643	11,643	11,643	11,643

SD, standard deviation.

RESULTS

Descriptive Statistics

Table 1 reports basic descriptive statistics, showing that of those receiving offers roughly 70 percent are female, 20 percent are black, 4 percent are Hispanic, and 5 percent are Asian. Differences in salary offers and market salary rates are in the expected direction, with men and whites doing better than females and minorities. Comparing the market salary rate with salary offers reveals that they are similar and that the average salary offer is smaller than the market rate for all of the groups.

The results presented in Table 2 show that, overall, women's market salary rates are 67 percent of men's (column 8, bottom row), and that blacks, Hispanics, and Asians have market salary rates that are, respectively, 72, 84, and 90 percent of whites (columns 9–11, bottom row). Table 2 also reports the female and minority market salary rates as a percent of the male and white market salary rates within each of the EEO codes. It is noteworthy that even the largest gender and race differences within specific EEO codes are still smaller than the overall race and gender differences. This suggests that the overall race and gender differences are mostly accounted for by sorting on who receives offers in the different EEO codes. Columns 3 and 4 confirm the plausibility of sorting on EEO codes, as they show that the different EEO codes contain different percentages of men and whites, and that EEO codes with larger average market salary rates (column 7) have a higher percentage of men and whites.

Panel A of Table 3 compares the average overall gender and race differences in market salary rate with the average gender and race differences within EEO codes. This allows us to see race and gender differences net of EEO codes, and to what degree differences in salary offers are accounted for by sorting on EEO codes. The contrast is quite stark, as differences within EEO categories are substantially smaller than those overall. Controlling just for these five broad occupational categories accounts for around 70 to 80 percent of the overall gender and race differences in market salary rates. Panel B shows that differences in salary offers follow a similar pattern.

TABLE 2. Market Salary Rate by Gender, Race, and Equal Employment Opportunity Code

	1	2	3	4	5	6	7	8	9	10	11
	Percent		Percent white		Mean market salary rate (\$)		Percent of male/white market salary rate				
N	Percent	Percent male	Percent white	Men	White	Overall	Female	Black	Hispanic	Asian	
Office and clerical	8,381	72	20	67	25,178	24,863	24,273	96	91	97	96
Professionals	1,007	9	42	83	47,942	43,733	43,463	84	90	99	105
Officials/managers	1,623	14	55	87	63,203	58,952	58,592	84	90	97	103
Technicians	203	2	69	87	61,776	61,006	61,059	96	96		104
Sales workers	396	4	66	93	69,629	65,995	65,742	84	92		
Overall	11,643	1	29	72	42,782	34,997	32,774	67	72	84	90

Notes: The overall row statistics also include 33 service workers. Column 1 reports the number of applicants receiving offers, 2 the percentage distribution, 3 the percent male, 4 the percent white, 5 the mean market salary rate for males in dollars, 6 the mean market salary rate for whites in dollars, 7 the overall mean market salary rate in dollars, 8 the mean female market salary rate as a percentage of the male, 9 the mean black market salary rate as a percentage of white, 10 the mean Hispanic market salary rate as a percentage of white, and 11 the mean Asian market salary rate as a percentage of white. Empty cells have less than 10 observations.

TABLE 3. Gender and Race Differences in Market Salary Rate and Salary Offer

	Average overall	Average net of EEO	Percent of gap explained
Panel A: Market salary rate			
Female as percent of male	67	93	80
Black as percent of white	72	91	69
Hispanic as percent of white	84	97	82
Asian as percent of white	90	98	78
Panel B: Salary offer			
Female as percent of male	61	90	74
Black as percent of white	69	92	74
Hispanic as percent of white	83	96	76
Asian as percent of white	91	99	89

Note: The average overall column reports the raw gender and race differences, while the average net of EEO reports the weighted average of the gender and race differences within Equal Employment Opportunity (EEO) codes.

Market Salary Rate

Table 4 presents results from models of the market salary rate. Model 1 reports the baseline differences for women and minorities. As was seen in Tables 2 and 3, these differences are substantial, ranging from $-.194$ for Hispanic men to $-.558$ among black women. The R^2 for model 1 is also noteworthy, revealing that 17 percent of the variation in the market salary rate is accounted for just by race and gender. Model 2 adds dummy variables for the applicant source. While the introduction of these variables reduces the size of the gender and race gaps by around 3 percentage points,⁷ the remaining differences are still quite substantial. Comparing the R^2 coefficients for the two models shows that controlling for applicant source accounts for another 9 percent of the variance in the market salary rates. Thus, while the results indicate that employee referrals do better than advertisement respondents and that applicant source has implications for the market salary rates, race and gender differences in market salary rate are largely unaccounted for by differences in applicant source in this firm.

Model 3 adds EEO codes into the regression and finds that controlling for EEO codes substantially reduces the differences in market salary rate. Net of EEO codes the gaps are closer to $-.1$, whereas previously they were on the order of $-.3$. As would be expected, the R^2 value is also substantially higher for this model; this model accounts for nearly 70 percent of the variation in market salary rate. Occupational sorting thus plays an important role in gender and race differences in market salary rates, as is visible even with a crude proxy such as EEO codes. It is worth noting that separate analyses show that once EEO codes are included, applicant source does not account for any additional variation in market salary rate or reduce the effects of race and gender. Thus, any effects of applicant source on either market salary rates or race and gender differences in

TABLE 4. Effects of Gender and Race on Logged Market Salary Rate

Variables	Model 1	Model 2	Model 3
Demographic			
White women	-.359***	-.315***	-.080***
Black men	-.401***	-.356***	-.112***
Black women	-.558***	-.500***	-.168***
Hispanic men	-.194***	-.174***	-.049*
Hispanic women	-.454***	-.403***	-.102***
Asian men	-.247***	-.218***	-.070***
Asian women	-.402***	-.351***	-.087***
Applicant source (Advertisement omitted)			
Employee referral		.129***	.070***
Temporary agency		-.004	.052***
Walk-in		-.036*	.026**
Search firm		.619***	.262***
Internet		.312***	.171***
Other		.220***	.107***
EEO group (Office and clerical omitted)			
Officials and managers			.774***
Professional			.508***
Sales workers			.878***
Technicians			.824***
Constant	10.596***	10.435***	10.099***
R ²	.172	.258	.684
N	11,643	11,643	11,643

* $p < .05$, ** $p < .01$, *** $p < .001$.

EEO, Equal Employment Opportunity.

market salary rates appear to operate primarily through the sorting of applicants into broad occupational categories.

Salary Offers

Table 5 reports results from the models of salary offers. Model 1 reports the baseline race and gender differences, and finds differences similar to those found in the market salary rate. The amount of variation in the two dependent variables accounted for by just race and gender is likewise similar at 18 percent. In Model 2, controls for applicant source are introduced, showing that as with market salary rate, little is gained by controlling for applicant source. The gender and race gaps are slightly smaller once applicant source is taken into account, and the R² indicates that including applicant source in the model accounts for 29 percent of the variation in salary offers. Model 3 introduces dummy variables for EEO codes, revealing again that even rough occupational codes make a considerable difference. As with the analysis of market salary rates presented in Table 4, salary offers exhibit much smaller race and gender differences and much more of the variation is accounted for when we control for these basic occupational categories.

TABLE 5. Effects of Gender and Race on Logged Initial Salary Offer

Variables	Model 1	Model 2	Model 3	Model 4	Model 5
Demographic					
White women	-.442***	-.385***	-.126***	-.079***	-.017***
Black men	-.455***	-.400***	-.130***	-.050***	-.002
Black women	-.617***	-.546***	-.179***	-.053***	-.022***
Hispanic men	-.198***	-.174***	-.035	-.002	.007
Hispanic women	-.522***	-.459***	-.125***	-.063***	.002
Asian men	-.247***	-.210***	-.049*	.002	.000
Asian women	-.465***	-.400***	-.109***	-.059***	.029***
Applicant source (Advertisements omitted)					
Search firm		.770***	.376***		
Temporary agency		-.015	.046***		
Walk-in		-.113***	-.044***		
Employee referral		.106***	.042***		
Other		.242***	.117***		
Internet		.356***	.197***		
EEO groups (Office and clerical omitted)					
Officials and managers			.850***		
Professional			.584***		
Sales workers			.946***		
Technicians			.943***		
Log market salary rate				1.010***	
Occupation fixed effects					✓
Constant	10.490***	10.322***	9.950***	-.215***	10.136***
R ²	.180	.285	.694	.846	.941
N	11,643	11,643	11,643	11,643	11,643

* $p < .05$, ** $p < .01$, *** $p < .001$.

EEO, Equal Employment Opportunity.

Models 4 and 5 build on the previous models by exploring whether more detailed occupational controls account for more of the race and gender differences. In Model 4, the market salary rate of the job is controlled for. This allows us to see the effect of market salary rate on the salary offer, and whether there are any race and gender differences net of this. Predictably, the effect of market salary rate is such that a 1-percent increase in the market salary rate results in a near 1-percent (1.01) increase in the salary offer.⁸ Overall, we see that market salary rate, race, and gender account for 85 percent of the variation in initial salary offers, 15 percentage points more than when EEO codes are used. The gender and race effects are also smaller net of market salary rate. Comparing Model 4 with Model 1 shows that controlling for the job's market salary rate accounts for upwards of three-quarters of the gender and race effects. Thus, sorting on market salary rates accounts for much of the race and gender inequality in salary offers. However, as the coefficients from this model show that gender and race impact

differences between the market salary rate for the job and the salary offer, they are also noteworthy for the gender and race effects that remain, even though these effects are smaller than in previous models.

Model 5 provides information about race and gender differences in salary offers within occupations by introducing occupation fixed effects for the 866 detailed occupation codes. The results can be interpreted as comparing the offers of men and women in the same occupation, and allow us to see if anything is gained by looking at detailed occupational sorting once the market salary rate of the job is known (as in Model 4). Comparing the R^2 values from Models 4 and 5 reveals that Model 5 accounts for more of the variance in salary offers. We also see that the effects of race and gender are smaller, with the largest disadvantage among black women, who receive offers that are on average 2.2 percent less than the average offer to white men for the same occupation. Thus, while controlling for market salary rate accounts for a considerable portion of race and gender differences, detailed occupations do even better. This suggests that the role of occupational sorting is not completely reducible to the market salary rate of the occupation.⁹ Further, the negligible race and gender differences found by Model 5 indicate that the baseline differences are entirely the result of sorting on occupation.

DISCUSSION

This study shows that occupational sorting accounts for the substantial race and gender differences in salary at the point of hire in this firm. In doing so, it highlights the similarities between the mechanisms producing race and gender wage differences in this context. While the role of occupational sorting in creating gender differences has received much attention, in questions of racial inequality, the role of occupational sorting is less prominent. The centrality of occupational sorting to both race and gender wage gaps suggests that to understand labor market inequality more broadly, it is important to understand how it is that occupational sorting operates. One of the interesting findings in this regard is that a large portion of the wage differences are accounted for by sorting into very broad groups. White women, for example, receive salary offers on average 36 percent ($\exp[-0.442] - 1$) less than white men in this firm. Controlling for detailed occupational codes in the fixed effects analysis reveals a 2 percent gender difference, or a reduction of 94 percent in the salary differences. However, detailed occupational codes are not always available, and this study shows that controlling for just EEO codes reduces the gap from 36 percent to 8 percent, a reduction of 78 percent. Although the prevailing wisdom suggests that EEO codes are not particularly useful in accounting for occupational sorting (e.g., Bielby and Baron 1986), this finding suggests that, at least in terms of wage inequality at the point of hire, much of the effect of occupational sorting occurs at a broad level. Thus, while detailed occupational codes are preferable, it is noteworthy that even quite broad codes can account for a sizable portion of differences in initial salary.

The salience of occupational sorting is also visible in the results for market salary rates. Race and gender differences in market salary rates mirror those in the salary

offers. This is of interest, because the market salary rates were established before applications were received and therefore cannot be adjusted based on the demographic characteristics of successful applicants. Differences in market salary rates are thus entirely the result of occupational sorting processes. Further, the role of occupational sorting is visible in the departure of salary offers from the market salary rate. While controlling for the market salary rate of the job accounts for the majority of race and gender differences, it does not completely account for them, and examining detailed occupational codes accounts for more of the gap than controlling only for the market salary rate. This suggests that even race and gender differences in discrepancies between market salary rates and salary offers are accounted for by differences in occupational sorting.

This study also sheds light on how occupational sorting works from the firm's perspective. It is possible that human resource personnel could offer higher salaries for certain types of jobs because they receive certain types of applicants. Human resource personnel could, for example, see that successful applicants for an engineering job are overwhelmingly white men and adjust their offers accordingly. However, the findings here suggest that this does not occur at this firm. That is, the bulk of wage differences arise not from adjustments made to individual offers, but rather from the process by which applicants and jobs are matched. This is not to say that occupational wage setting is not important, but that it seems unlikely that it is a factor when human resource personnel are making decisions about specific offers.¹⁰

It is worth returning to the question of causal factors and mechanisms. As previous firm-level studies have shown that most of the race and gender differences in earnings and promotions are attributable to differences at initial hire, subsiding over time and with rising seniority (DiPrete 1989; Gerhart 1990; Petersen and Saporta 2004), this study focuses on the origins of those initial differences in one firm. The sizable race and gender differences in salary offers that exist in this setting are found to be almost entirely attributable to occupational sorting, a finding that underscores the importance of studies examining how the hiring process interacts with occupational sorting. In thinking about causal agents, the central issue becomes how it is that occupational sorting occurs in the hiring process, suggesting that an understanding of how job seekers and employers interact to impact the job-matching processes that drive occupational sorting is needed. Here, previous research is of some help, suggesting that both applicants and employers are implicated. For example, previous work on gender differences finds that not only do applicants have gender-typed preferences for certain jobs, but the gender stereotypes of human resource personnel play a role in steering applicants toward jobs for which they believe applicants would be well suited (Fernandez and Sosa 2005). There is also evidence that, in some contexts, homophilous networks cut through supply and demand side considerations to exacerbate occupational sex segregation through referral processes (Fernandez and Sosa 2005; but see Fernandez and Fernandez-Mateo 2006). While there is no evidence that referral effects impact race and gender differences in this setting, it seems reasonable to expect that similar applicant and employer factors would be at play in this context.

In thinking more concretely about how firms might impact occupational sorting at the point of hire, there are three arenas to consider: recruitment, human resource steering, and hiring decisions. Recruitment differences occur when firms encourage certain people to apply for certain jobs. This produces different applicants to different jobs, and could play quite a powerful role in occupational sorting by creating different applicant pools for different occupations. There are a few studies that discuss the implications of recruitment based on gender and race (e.g., Rees and Schultz 1970; Holzer 1987), but little in the way of detailed information on job seekers who considered applying or might have considered applying and did not.¹¹ Anecdotal evidence suggests that recruitment differences could play a powerful role in occupational sorting, as even slight changes in the wording of newspaper advertisements can have large effects on the applications received (Petersen and Togstad 2006). Human resource steering takes place at the next phase of the hiring process once applicants have contacted the firm, but before their application has been completed. While applicants can sometimes apply directly to specific jobs, in others cases, the application process is mediated by human resource personnel. By encouraging and discouraging certain people to apply for certain jobs, human resource personnel can play an active role in the occupational sorting process (see, e.g., Fernandez and Mors 2008). This is potentially an important source of race and gender differences about which very little is known. Finally, even if all jobs receive similar applicants, by extending certain people offers to certain jobs, firms' hiring rates could play a substantial role in occupational sorting. Evidence from other contexts suggests that differential hiring rates do not play a substantial role in occupational sorting for gender, but might for race (Petersen et al. 2005).

While this study cannot identify where in the hiring process occupational sorting occurs, it does demonstrate that race and gender differences in wages at the point of hire are overwhelmingly a function of occupational sorting. This is significant because it rules out within-job pay discrimination as a factor in creating wage differences at the point of hire, and underscores the importance of studies examining the job-matching process. The importance of the job-matching process is further highlighted by the results for applicant source. Applicant source is found to have some effect on salary offers, but is largely orthogonal to race and gender differences. To the degree that applicant source does impact race and gender differences, this effect is mediated by occupational sorting. Thus, while applicant source (including referral networks) plays a role in occupational sorting, it does not begin to account for the substantial impact that occupational sorting has on race and gender differences in this firm.

In conclusion, this study shows where wage differences at the point of hire arise in this firm, and what axes are important. As differences in initial salary are found to be entirely the result of occupational sorting processes, within occupation wage discrimination can be eliminated as a factor in the creation of both race and gender wage differences in this context. Thus, in demonstrating the inequality-generating potential of occupational sorting processes, this study highlights the centrality of the job-matching process in gender and race wage inequality. Future research would do well to explore job-matching processes in a broad range of occupations, and in particular,

differences in applicant preferences and firm recruitment practices seem important and understudied.

ACKNOWLEDGMENTS

I would like to thank Trond Petersen and Roberto Fernandez for their generous comments and numerous insights. This article also benefited from the feedback of Henry Brady, Tonya Dewey, Neil Fligstein, Mike Hout, Lars Jarkko, Peter Kivisto, Julie Kmec, and several anonymous *Sociological Quarterly* reviewers. Previous versions of this article were presented at the 2004 RC28 summer meeting in Rio de Janeiro and the 2004 ASA annual meeting in San Francisco. Research for this article was conducted while supported by a National Science Foundation Graduate Research Fellowship.

NOTES

¹While this archetypical process is the model for the company studied here, it is likely that there is variation in the hiring processes used in different branches and occupations, particularly in the number of interviews.

²While generalization is difficult in all social science, this is especially true of studies surrounding labor market inequalities, as the case study data necessary for fine-grained analyses do not lend themselves readily to generalization. For the sake of brevity, in reviewing the findings of other studies, I do not detail their limitations; however, it is important to note that all of these studies occur in specific empirical contexts and should not be viewed as context-free.

³Examples of the 866 detailed occupation codes include portfolio monitor analyst II, mail clerk II, client relations rep I, HR help desk specialist, and telemarketing specialist I.

⁴This occurs when an applicant successfully applies to more than one job opening. There were 120 applicants who received 2 job offers, and 27 applicants who received 3 job offers.

⁵That is, given that the market salary rate is established before the race and gender of the successful applicant are known, differences in the market rates cannot be the result of the decision to pay a successful applicant more on the basis of their race or gender. However, although occupational sorting processes are the proximate cause of race and gender differences in the market salary rates, at the macro level, the process of setting market rates is likely responsive to valuative discrimination. Thus, while in this analysis I take the market salary rates as given (because from the perspective of applicants and hiring managers they are), it is important to remember that they are not set in a vacuum, and that at the macro level, they are likely responsive to the demographics of the people with that particular occupation.

⁶Strictly speaking, the coefficients reported pertain to log unit differences; proportions can be obtained by exponentiating the coefficients and subtracting one. When coefficients are small in absolute value (between .1 and -1), they approximate the relative change in the dependent variable.

⁷For example, model 1 shows that black women receive salary offers that are 42.8 percent ($\exp[0.558] - 1$) less than white men, while model 2 shows that, controlling for applicant source, they receive offers that are 39.3 percent ($\exp[-0.5] - 1$) less than white men.

⁸It is worth noting that the constant term in Model 4 is quite different from those in the other models. The constant should be interpreted in conjunction with the effect of log market salary

rate to indicate that people on average do not earn their market rate salary. That the constant is negative while the effect of log market rate salary is greater than one reflects the fact that people receiving offers for jobs with higher market rates receive salary offers that are a higher percentage of their market rates.

⁹The difference between Models 4 and 5 is because of the larger number of men and whites receiving offers for jobs that have larger market salary rates, as jobs with larger market salary rates tend to pay more as a percentage of the market rate. This is confirmed by the fact that results from a model (not reported) including fixed effects for the 80 pay groups yielded coefficients similar to the 866 occupation fixed effects. This suggests that there is something associated with being in a market salary rate group that is not captured by the group's market salary rate.

¹⁰At the macro level, it is still possible that there is some kind of valuative discrimination at play. It is even possible that the firm's market salary rates are influenced by their knowledge of what kind of groups are likely to apply, rather than what other companies are paying. It seems more likely that the market and society as a whole devalues these occupations. An analysis of how it is that firms establish their market salary rates could be very informative about this process.

¹¹It is unclear whether it is even possible to obtain both the employer-side data and detailed information about why potential applicants chose not to apply necessary to explore this research agenda. The study that comes closest is Fernandez and Sosa (2005), which matches firm hiring data to census data from the area surrounding the firm.

REFERENCES

- Becker, Gary. 1971. *The Economics of Discrimination*. Chicago, IL: University of Chicago Press.
- Bielby, William T. 2000. "Minimizing Workplace Gender and Racial Bias." *Contemporary Sociology* 29:120–29.
- Bielby, William T. and James N. Baron. 1986. "Men and Women at Work: Sex Segregation and Statistical Discrimination." *American Journal of Sociology* 91:759–99.
- Correll, Shelley J. 2001. "Constraints into Preferences: Gender, Status, and Emerging Career Aspirations." *American Sociological Review* 69:93–113.
- . 2004. "Gender and the Career Choice Process: The Role of Biased Self-Assessments." *American Journal of Sociology* 106:1691–730.
- DiPrete, Thomas A. 1989. *The Bureaucratic Labor Market: The Case of the Federal Civil Service*. New York: Plenum.
- Elvira, Marta M. and Christopher D. Zatzick. 2002. "Who's Displaced First? The Role of Race in Layoff Decisions." *Industrial Relations* 41:329–61.
- England, Paula. 1992. *Comparable Worth: Theories and Evidence*. Hawthorne, NY: Aldine de Gruyter.
- Fernandez, Roberto M., Emilio J. Castilla, and Paul Moore. 2000. "Social Capital at Work: Networks and Employment at a Phone Center." *American Journal of Sociology* 105:1288–356.
- Fernandez, Roberto M. and Isabel Fernandez-Mateo. 2006. "Networks, Race, and Hiring." *American Sociological Review* 71:42–71.
- Fernandez, Roberto M. and Marie Louise Mors. 2008. "Competing for Jobs: Labor Queues and Gender Sorting in the Hiring Process." *Social Science Research* (forthcoming).
- Fernandez, Roberto M. and Lourdes Sosa. 2005. "Gendering the Job: Networks and Recruitment at a Call Center." *American Journal of Sociology* 111:859–904.
- Fernandez, Roberto M. and Nancy Weinberg. 1997. "Sifting and Sorting: Personal Contacts and Hiring in a Retail Bank." *American Sociological Review* 62:883–902.

- Gerhart, Barry. 1990. "Gender Differences in Current and Starting Salaries: The Role of Performance, College Major, and Job Title." *Industrial and Labor Relations Review* 43:418–33.
- Granovetter, Mark. 1995. *Getting a Job* (2d ed.). Chicago, IL: University of Chicago Press.
- Holzer, Harry J. 1987. "Hiring Procedures in the Firm." Pp. 243–74 in *Human Resources and the Performance of the Firm*, edited by Morris Kleiner, Richard Block, Myron Roomkin, and Sidney Salsburg. Madison, WI: Industrial Relations Research Association.
- Kanter, Rosabeth M. 1977. *Men and Women of the Corporation*. New York: Basic Books.
- Kirschenman, Joleen and Kathryn M. Neckerman. 1991. "'We'd Love to Hire Them But . . .': The Meaning of Race for Employers." Pp. 203–32 in *The Urban Underclass*, edited by Christopher Jencks and Paul E. Peterson. Washington, DC: The Brookings Institute.
- Neumark, David. 1996. "Sex Discrimination in Restaurant Hiring: An Audit Study." *The Quarterly Journal of Economics* 111:915–41.
- Pager, Devah. 2003. "The Mark of a Criminal Record." *American Journal of Sociology* 108:937–75.
- Pager, Devah and Lincoln Quillian. 2005. "What Employers Say versus What They Do." *American Sociological Review* 70:355–80.
- Pergamit, Michael R. and Jonathan R. Veum. 1999. "What Is a Promotion?" *Industrial and Labor Relations Review* 52:581–601.
- Petersen, Trond and Laurie A. Morgan. 1995. "Separate and Unequal: Occupation-Establishment Sex Segregation and the Gender Wage Gap." *American Journal of Sociology* 101:329–65.
- Petersen, Trond and Ishak Saporta. 2004. "The Opportunity Structure for Discrimination." *American Journal of Sociology* 109:852–901.
- Petersen, Trond, Ishak Saporta, and Marc-David L. Seidel. 2000. "Offering a Job: Meritocracy and Social Networks." *American Journal of Sociology* 106:763–816.
- . 2005. "Getting Hired: Sex and Race." *Industrial Relations* 44:416–43.
- Petersen, Trond and Thea Togstad. 2006. "Getting the Offer: Sex Discrimination in Hiring." *Research in Social Stratification and Mobility* 24:239–57.
- Powell, Gary N. and D. Anthony Butterfield. 1997. "Effect of Race on Promotions to Top Management in a Federal Department." *Academy of Management Journal* 40:112–28.
- Rees, Albert and George P. Schultz. 1970. *Workers and Wages in an Urban Labor Market*. Chicago, IL: University of Chicago Press.
- Sicherman, Nachum. 1996. "Gender Differences in Departures from a Large Firm." *Industrial and Labor Relations* 49:484–505.
- Spilerman, Seymour and Trond Petersen. 1999. "Organizational Structure, Determinants of Promotion, and Gender Differences in Attainment." *Social Science Research* 28:203–27.
- Tomaskovic-Devey, Donald. 1993. *Gender and Race Inequality at Work*. Ithaca, NY: ILR Press.