Dickinson, David L., and Ronald L. Oaxaca. (2009) Statistical Discrimination in Labor Markets: An Experimental Analysis. **Southern Economic Journal**, 76(1): 16-31.(July 2009) (ISSN: 0038-4038) Published by the Southern Economic Association (July 2009). Archived in NCDOCKS with permission of the managing editor.

Statistical discrimination in labor markets: An experimental analysis

David L. Dickinson and Ronald L. Oaxaca

ABSTRACT

This article reports results from controlled laboratory experiments designed to study second-moment (that is, risk-based) statistical discrimination in a labor market setting. Since decision makers may not view risk in the same way as economists or statisticians (that is, risk = variance of distribution), we also examine alternative measures of risk: the support of the distribution and the probability of earning less than the expected (maximum) profits for the employer. Our results indicate that employers made statistically discriminatory wage offers consistent with loss aversion in our full sample (though the result is driven by the male employer subsample). If one can transfer these results outside of the laboratory, they indicate that discrimination estimates based only on first-moment (mean-based) discrimination are biased. The public policy implication is that efforts and legislation aimed at reducing discrimination of various sorts face an additional challenge in trying to identify and limit relatively hidden, but significant, forms of statistical discrimination.

ARTICLE

1. Introduction

When membership in a particular group conveys valuable information about an individual's skills, productivity, or other characteristics, a nonprejudiced agent may still find it rational to statistically discriminate. Examples of statistical discrimination include wage or hiring decisions in labor markets, racial profiling in law enforcement, determinants of loan approval rates, voting the party ticket in elections, or differential premiums for insurance, among others. In some settings statistical discrimination is legal and acceptable (for example, insurance rates); whereas, in others it is controversial and/or illegal (for example, racial profiling and employment discrimination). Existing research has focused on first-moment statistical discrimination: that is, discriminatory wage offers to females or lower loan approval rates for minority applicants are based on average productivity and default rates, respectively. Agents attribute average group characteristics to each individual from that group when it is costly to gather information.

In this article we explore the possibility that statistical discrimination extends beyond differential treatment based on average group characteristics. Specifically, discrimination may also exist if agents base decisions on productivity distribution risk (or default rates, accident rates, etc.). Using labor markets as an example, risk-averse employers may make lower wage offers to females if their productivity variance is believed to be higher, even though average productivity may be identical to that of males. If such variance-based statistical discrimination is empirically documented, then existing measures of statistical discrimination are biased, and measures of prejudiced-based discrimination may be overstated. Some have found field evidence of statistical discrimination based on higher-order moments of a distribution (Ayers and Siegelman 1995; Goldberg 1996; List 2004), but we also recognize that statistical variance may not be the only behaviorally important measure of distributional risk. Thus, we contribute additionally by examining the support of the productivity distribution (Tversky and Kahneman 1973; Curley and Yates 1985; Griffin and Tversky 1992; Babcock et al. 1995) and the possibility for loss (Kahneman and Tversky 1979) as two other potentially important cognitive assessments of risk. In short, some discrimination labeled as personal prejudice or taste based may really be just a different form of statistical discrimination than what is typically examined.

We report results from a controlled laboratory experiment in which subjects are engaged as employers and workers in a laboratory double-auction labor market. We choose a labor market context for our experiments for several reasons. First, we believed that because payoffs to employer-subjects are determined by an outcome variable minus a contract payment to another subject, the labor market context would aid subjects in understanding payoffs in our experiment. Second, statistical discrimination is very relevant to labor markets, highlighted by the many existing empirical studies of statistical discrimination that examine labor markets. Finally, the use of labor market context in a competitive double-auction market environment is a logical context with precedence in the experimental economics literature (Fehr et al. 1998). That said, the insights we gain from our data extend to other contexts, and the implication of our results is that statistical discrimination may be more pervasive than previously thought. Our results show that subjectemployers make significantly lower wage offers when the probability of loss is greater, and this measure of risk mattered more than the statistical variance or support distribution.

Statistical theories of discrimination have been advanced by Arrow (1972), Phelps (1972), Aigner and Cain (1977), and Lundberg and Startz (1983). Some studies base statistical discrimination on noisier productivity signals for certain worker groups, while others base it on imperfect or incomplete information. (1) Most researchers advance theories that depend on differences in average productivity characteristics; although, others note that statistical discrimination need not be based on differences in average productivity (e.g., Aigner and Cain 1977; Curley and Yates 1985). For risk-averse individuals, it seems clear that a less risky outcome distribution would be preferred to a more risky distribution; although, "risk" may be defined in ways other than just a statistical variance, as we note. Empirical evidence alluding to statistical discrimination can be found in a variety of settings; although, it is often difficult to identify taste-based versus statistical discrimination (see discussion in Arrow 1998). Probably the only easily observable forms of statistical discrimination are the legal forms, such as those found in the insurance industry. In labor markets there is some direct evidence from employer interviews that race is used as a proxy in employment decisions (Wilson 1996). Neumark (1999) uses field data to uncover discrimination not based on productivity characteristics, but Altonji and Pierret (2001) find little evidence for statistical discrimination based on race. (2)

Given identification and causation issues inherent in field data examinations of discrimination, some have used controlled experiments to study statistically based discrimination resulting from imperfect information (Anderson and Haupert 1999), asymmetric information (Davis 1987), or ethnic stereotypes (Fershtman and Gneezy 2001). We employ a full information environment to examine higher-order statistical discrimination and to explore which of several risk measures is more behaviorally important. Our design is such that causation can go in only one direction (that is, exogenous wage distributions imply that wage contracts cannot affect future worker productivity), and the market institution for determining wage contracts is one that produces strong convergence to the competitive equilibrium prediction. Nevertheless, we find evidence for statistical discrimination based on one important measure of worker risk.

2. Experimental Design

We implement a two-sided auction market design to simulate a labor market. Specifically both employers (buyers) and workers (sellers) negotiate in an open-pit fashion, with no central auctioneer. Workers are more plentiful than employers, and so there is an equilibrium level of "unemployment" in this design. Both supply and demand for labor are induced upon the experimental subjects using standard experimental techniques, discussed in Smith (1982). (3)

The baseline design we use is simple in that it generates clear equilibrium predictions. Specifically the demand side of the experimental market consists of five employers, each capable of hiring one unit of labor in each experimental market round. The productivity of a unit of labor in the baseline (treatment 1) is certain and fixed at three units of output (each unit of output sells for SI experimental), and so the demand for labor is perfectly elastic at $3.00 up to five units of labor. The supply side of the market consists of 10 workers, each with a reservation wage of S0.40, and each is able to sell at most one unit of labor services in each experimental market round. As such, the supply curve is perfectly elastic at SO.40 up until 10 units of labor. The predicted market wage is S0.40, and the predicted market quantity of labor traded is five units. We used the labels "worker," "employer," and "wages" to facilitate the subjects' understanding of the connection between productivity and final payoff, but it was clear to all subjects that no labor task would be completed in the experiment. In this way we maintain strict control over productivity in the experiment. Figure 1 shows the experimental design graphically.

The baseline experimental design is quite similar to that used in Smith (1965), though Smith does not use a labor market context. That is, at the predicted equilibrium the entire market surplus is allocated to one side of the market (the buyers of labor). In our design the employers are not given information on worker reservation wages, and workers are not informed as to the value (to employers) of a unit of output. Payoff information is therefore private to each subject, as in Smith (1965), who shows that, even when market surplus at equilibrium is designed to be extremely imbalanced, this trading institution produces strong convergence of equilibrium prices to the competitive equilibrium prediction. Any evidence of statistical discrimination in the uncertain productivity treatments would then be significant, given the strong competitive tendencies inherent in our baseline design.

The stochastic or uncertain productivity treatments are labeled treatments 2, 3, and 4. The difference across these uncertain productivity treatments lies in the particular (known) productivity distribution for the labor pool. After hiring a unit of labor in an uncertain productivity treatment the employer discovers the realized productivity of that unit of labor by means of an ex post random draw. Specifically, in treatment 2, productivity of the labor pool is either one, two, three, four, or five units of output with probability 10%, 10%, 60%, 10%, and 10%, respectively. Productivity is determined by a random draw from a Bingo cage, and an independent draw is conducted for each employer who hires a unit of labor. Though wage contracts are made with a specific experimental subject in any given trading round, it is made clear that productivity draws are independent of the actual worker-subject (that is, you cannot contract in the next round with John Doe to ensure productivity of five just because it happened to turn out that way in the current or past rounds when contracting with John Doe). The independence of the productivity draw from the specific worker-subject controls for differences that employers in naturally occurring work environments would have in sorting and selecting workers from a given labor pool. We simply assume that employers are equal on this dimension, and so hiring any worker from a given pool of workers with a specific productivity distribution is similar to taking a random draw from the productivity distribution.

Treatments 3 and 4 also involve uncertain productivity distributions of the labor pool, but they differ from treatment 2 in terms of the specific distribution. In treatment 3 productivity of the labor pool is either one, two, three, four, or five units of output with probability 20% for each possible outcome. In treatment 4 productivity of the labor pool is either two or four units of output with probability 50% for each.

The expected competitive employer profit is $2.60 experimental dollars because the expected revenue is $3.00 and the competitive wage is $0.40. There were a total of seven experimental sessions in which the order of the treatments was randomized. Each of the four treatments in an experimental session lasted four periods. There were a total of 35 employers in our experiment, and we observe wage contracts for each employer a total of 16 times. Hence, we have a panel with 560 observations.

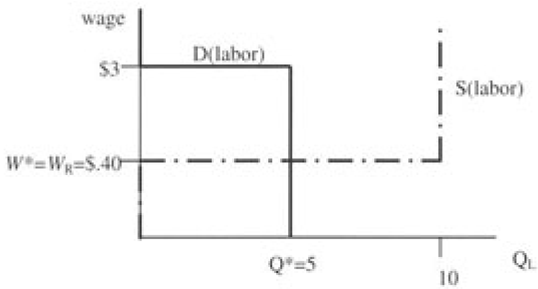

Figure 1. Experimental Design

Table 1 describes the experimental design in terms of how each of the treatments varies with respect to distinct measures of productivity distribution risk. This design allows us to examine several candidate variables for statistical discrimination: discrimination based on the variance of labor productivity, based on the support of the productivity distribution, or based on the probability of less-than-expected competitive profits for the employer. A comparison of wage contracts in treatment 1 to treatments 2, 3, and 4 allows us to test these different hypotheses of statistical discrimination. Binary comparisons among treatments 2, 3, and 4 allow us to look at the joint effects of varying combinations of variance, support, and probability of lessthan-expected competitive profits for the employer. The difference between treatment 3 and treatment 2 reflects the joint effects of a higher variance and greater probability of less-thanexpected profits in treatment 3. The difference between treatment 4 and treatment 2 reflects the joint effects of a smaller support and a greater probability of less-than-expected profits in treatment 4. Finally, the difference between treatment 4 and treatment 3 reflects the joint effects of a smaller variance, a smaller support, and a greater probability of less-than-expected profits in treatment 4. For the statistical analysis discussed next, we also create independent variables that isolate the effects of changes in each distinct measure of distributional risk.

Table 1. Experiment Treatment Design

Treatment	Description Productivity (Probability)	Productivity Mean	Productivity Variance	Productivity Distribution Support	Likelihood of Productivity < Mean Productivity
1	3 (1.00)	3	0	3	0
2	1, 2, 3, 4, 5 (0.1, 0.1, 0.6, 0.1, 0.1)	3	1	1–5	0.20
3	1, 2, 3, 4, 5 (0.2, 0.2, 0.2, 0.2, 0.2)	3	2	1–5	0.40
4	2, 4 (0.5, 0.5)	3	1	2–4	0.50

3. Results

We report results on a total of seven experiment sessions, using 105 unique college-aged subjects (35 employers and 70 workers). Our sample has 53% female subjects overall: 57% female employers (N = 20 out of 35 employers) and 51% female workers (N = 36 of 70). Six of the seven treatments had either four, five, or six female workers (out of 10 total worker-subjects in the session). Session 7 abnormally had eight of 10 workers that were female. Each session lasted about 1.5 hours, and average earnings were $I8.13.

Our results are summarized in Tables 2, 3, and 4. In Table 2 we use dummy variables to control for the uncertainty productivity treatments 2, 3, and 4 (T2, T3, T4, respectively), to control for rounds 2, 3, and 4, and to control for treatment order within a four-treatment experiment (for example, T03 = 1 if a treatment occurs third in a particular session). Because our data consist of repeated observations on employers, panel data methods seem appropriate. Fixed effects and random effects estimators account for differences in wage contracts across employers and possible correlation in the error terms across rounds for an individual employer's wage contracts. Given our particular orthogonal design, the estimated parameters of the wage contract equations are identical for fixed effects, random effects, and ordinary least squares (OLS) with a single constant term. The estimated standard errors are also identical for fixed effects and random effects but differ from those obtained from OLS (see Oaxaca and Dickinson 2005 for details). Because we are able to reject the classic OLS model in favor of both fixed effects and random effects, we interpret this as support for using the fixed/random effects estimated standard errors in our analysis.

The Table 2 results show that, for the full sample, treatment 4 significantly lowers wage contracts offered to workers, but the results from the gender-specific samples show that this is due entirely to the behavior of the male employers. (4) The estimated wage contract effect of treatment 4 relative to treatment 1 captures the combined effects of all three of our risk measures. Wage contracts offered by male employers are about 12 cents lower in treatment 4 compared to the certain worker productivity treatment 1. This represents an average wage offer decrease of about 20% given the average wage contract level of about 60 cents for male employers. Female employers, on the other hand, did not offer significantly different wages across treatments. While this is consistent with female employers being risk neutral, the literature suggests otherwise. We return to this point later. Across rounds, the estimated coefficients indicate that wage contracts converge toward equilibrium in later rounds of each treatment, and wages also converge downward for a given treatment the later it occurs in the experiment.

Table 2. Wage Contracts (Random Effects)

Variable	Full Employer Sample Coeff.	Full Employer Sample Std Error	Male Employer Sample Coeff.	Male Employer Sample Std Error	Female Employer Sample Coeff.	Female Employer Sample Std Error
Constant	0.861	0.038***	0.906	0.061***	0.812	0.050***
T2	−0.003	0.029	−0.063	0.042	0.048	0.042
T3	−0.028	0.036	−0.069	0.049	0.012	0.052
T4	−0.062	0.029**	−0.116	0.044***	−0.027	0.039
Round 2	−0.119	0.028***	−0.116	0.039***	−0.121	0.039***
Round 3	−0.143	0.028***	−0.153	0.039***	−0.135	0.039***
Round 4	−0.149	0.028***	−0.167	0.039***	−0.135	0.039***
TO2	−0.181	0.034***	−0.154	0.048***	−0.193	0.048***
TO3	−0.183	0.029***	−0.193	0.042***	−0.153	0.041***
TO4	−0.186	0.029***	−0.177	0.044***	−0.172	0.040***
R^2	0.137		0.168		0.123	
Nobs	560		240		320	

** and *** indicate significance at the 0.05 and 0.01 levels, respectively, for the two-tailed test.

Table 3 presents treatment effects comparisons (that is, coefficient comparisons) from within the uncertain productivity treatments. Treatment 4 compared to treatment 2 (T4-T2) reflects the combined effect of the smaller support but higher probability of less-than-expected profits in treatment 4. In all samples the combined effect is negative but statistically significant only in the full sample and female employer subsample. We note that the depressed wage contract effects of treatment 4 relative to any of the other treatments are larger in absolute value than any other binary treatment comparisons. These results reflect the dominance of the loss aversion motive.

Table 3. Binary Comparisons among the Uncertain Productivity Treatments

Comparison	Full Employer Sample Coeff. Diff.	Full Employer Sample Std Error	Male Employer Sample Coeff. Diff.	Male Employer Sample Std Error	Female Employer Sample Coeff. Diff.	Female Employer Sample Std Error
T3-T2	−0.025	0.031	−0.006	0.041	−0.036	0.046
T4-T2	−0.059	0.030**	−0.053	0.049	−0.075	0.043*
T4-T3	−0.034	0.036	−0.047	0.051	−0.039	0.053
Nobs	560		240		320	

Coefficient comparisons from Table 2 results.
* and ** indicate significance at the 0.10 and 0.05 levels, respectively, for the two-tailed test.

Though these results presented thus far offer some initial evidence of statistical discrimination based on distributional risk, it is also the case that the treatment effects specification does not directly control for differences in the productivity distribution's variance, support, or probability of below average profits. This follows from the fact that certain treatments vary more than one of these distributional characteristics (see Table 1). In formulating our statistical design, we had not originally considered the loss aversion factor associated with the variation in the probability of less-than expected profits (that is, the probability of a less than average productivity draw). We therefore also estimate a model using explicit controls for individual changes in each of these distributional characteristics in Table 4.

In Table 4 wage contracts are regressed on variables for variance, support, and loss probability, where variance and loss probability are coded as defined in Table 1 (that is, loss probability is measured relative to expected [competitive] profits). The variable for productivity distribution support is defined as the width of support: 0, 0.4, 0.4, and 0.2 for treatments 1, 2, 3, and 4, respectively. Given our particular orthogonal design, the statistical model in Table 4 is the same as that in Table 2 but offers an alternative way to view the results pertaining to the treatment effects. Consequently, as in Table 2, the Table 4 results are from a fixed/random effects specification, and estimates are presented for the entire employer sample as well as the gender-based employer subsamples. (5) For comparison purposes, we estimate two versions of the model for each sample. The first model controls only for distributional variance as the risk measure of interest; whereas, the second model controls for variance, distribution support, and loss probability. Among the risk measure variables in Table 4, we can see in the overall sample that the only significant predictor of wage contract differences is the probability of loss, due entirely to the male employer subsample. Specifically, male employers offer significantly lower wage contracts when faced with higher probability of profits less than average from a worker. Interestingly, if we control only for productivity distribution variance in the first model of the full and employer samples, as might typically be done, this significant effect would (incorrectly) be attributed to traditional risk aversion. Our design is therefore able to discriminate between what looks like a risk aversion effect to show that it is really a Joss aversion effect.

Table 4. Wage Contracts with Alternative Measures of Risk (Random Effects)

	Full Employer Sample				Male Employer Sample				Female Employer Sample			
Variable	Coeff.	Std Error	Coeff.	Std Error	Coeff.	Std Error	Coeff.	Std Error	Coeff.	Std Error	Coeff.	Std Error
Constant	0.863	0.037***	0.861	0.038***	0.898	0.060***	0.906	0.061***	0.826	0.048***	0.812	0.050***
Variance	−0.023	−0.017	0.006	0.042	−0.042	0.024*	0.048	0.058	−0.001	0.025	−0.014	0.062
Support			0.056	0.123			−0.141	0.17			0.209	0.176
Loss Prob.			−0.159	0.092*			−0.273	0.136**			−0.108	0.129
Round 2	−0.119	0.028***	−0.119	0.028***	−0.116	0.038***	−0.116	0.039***	−0.121	0.039***	−0.121	0.039***
Round 3	−0.143	0.028***	−0.143	0.028***	−0.153	0.038***	−0.153	0.039***	−0.135	0.039***	−0.135	0.039***
Round 4	−0.149	0.028***	−0.149	0.028***	−0.167	0.038***	−0.167	0.039***	−0.135	0.039***	−0.135	0.039***
TO2	−0.165	0.033***	−0.181	0.034***	−0.143	0.047***	−0.154	0.048***	−0.185	0.045***	−0.193	0.048***
TO3	−0.197	0.028***	−0.183	0.029***	−0.213	0.040***	−0.193	0.042***	−0.176	0.039***	−0.153	0.041***
TO4	−0.199	0.028***	−0.186	0.034***	−0.218	0.039***	−0.177	0.044***	−0.177	0.039***	−0.172	0.040***
R^2	0.132		0.137		0.159		0.168		0.116		0.123	
Nobs	560		560		240		240		320		320	

*, **, and *** indicate significance at the 0.10, 0.05, and 0.01 levels, respectively, for the two-tailed test.

On the other hand, female employers did not significantly alter wage contracts in response to changes in any of the distributional risk variables. The result is consistent with females being risk or loss neutral. It may also be the case that female employers possess similar risk attitudes as male employers but with different subjective beliefs regarding the productivity distributions (for example, optimism as to the likely productivity draw from a distribution with larger support). If beliefs as well as risk preferences are important determinants of wage contracts, there may be systematic differences in both of these across genders (for example, males being either less optimistic, or having risk aversion that dominates any optimism toward the likely productivity draw). Though we do not generate data on beliefs, we do not consider optimism to be a likely explanation for our results. The reason is that subjects were given very explicit details on the exact productivity distribution.

Existing research on gender differences has shown that females are generally less driven by competition and more averse to negotiations than males (Niederle and Vesterlund 2007): That is, female employers might negotiate worse outcomes (higher wage contracts) in general, independent of the riskiness of the worker productivity distribution. In our experiment employer payoffs are partly determined by one's ability to compete with other employers while negotiating with workers in the double-sided auction institution. Babcock and Laschever (2003) document that females are generally more averse to negotiations than males. If this aversion to competitive negotiations is most prominent when worker productivity is uncertain, then this "negotiations aversion" may interact with female risk attitudes toward worker productivity. Across the seven treatments, the percentage of workers that were female were 60%, 50%, 40%, 40%, 50%, 40%, and 80%; whereas, the percentage of those female workers unemployed across all 80 wage contracts (five employers times 16 rounds each) in the respective sessions were 66%, 50%, 43%, 40%, 48%, 45%, and 76%. We do not therefore find evidence to indicate that female workers are, on average, more likely to be unemployed relative to male workers.

Suppose that female subjects in our sample are averse to negotiations risk and expected payoffs are a function of the productivity distribution risk as well as negotiations risk. If employer-worker matching is essentially random, then we would expect male employers and workers to have better contract outcomes than females. Female employers would offer higher wages, on average, and this would counteract any tendency to lower wages in response to worker productivity distribution risk. We do not, however, find such evidence that female employers offer higher wage contracts, ceteris paribus, or that females do worse in mixed-gender negotiations. (6) Aversion to negotiations may also manifest itself in gender-matching patterns, with female employers more likely to contract with a female worker. That is, women may be more averse to negotiating with men than with other women. In our sample single-gender contracts--male-male or female-female agreements--are statistically significantly more likely than mixed-gender contracts (306 to 254 individual wage contracts: $p = 0.01$ for the one-sided binomial test). However, female employers contracting with female workers (56% of the time) are not much more likely than male employers contracting with male workers (53% of the time). Overall our evidence with respect to female aversion to negotiations is weak. The gender difference in those unemployed or the propensity for same sex contracts is not great, and any such aversion is not displayed in the wage contracts themselves.

There is thus only weak evidence that females may be more averse to mixed-gender negotiations than males, but we do not find evidence that women fare worse in mixed-gender pairs or that they offer generally higher wage contracts. In short, the fact that females may be more averse to competitive negotiations does not explain the wage results from our gender-specific samples. Our experimental data indicate that males react more significantly to the probability of loss than females. Though we cannot fully explain the nature of this gender difference result, the overall significance is that we find evidence for statistical discrimination not based on average group differences. Considering the labor market context, our full data sample results indicate that this variable--a higher potential for less-than-average payoffs--can significantly decrease the wage that an employer would pay to workers from that more risky labor pool. (7)

The context of the statistical discrimination may be important for this result, but it implies that individuals respond significantly to increased distributional risk. If subjects feel somehow entitled to earn expected profits, then, for the entire sample, we find evidence consistent with statistical discrimination resulting from loss aversion. We are careful to note that our experiments examine only a particular range of variance and support of the productivity distribution. We do not take these results to imply that distributional variance and support are not behaviorally important measures of risk. Rather, we have shown that a dominant influence on behavior may be loss probabilities. And, because loss probability is not naturally independent of distribution variance and support, it is important that our results identify this more significant independent effect of loss probability on wage contracts. It is indicative of the fact that the reference point of average or expected profits is an important determinant of wage negotiations outcomes.

4. Concluding Remarks

This article has examined a very simple framework for studying second-moment statistical discrimination. In a general sense, this type of statistical discrimination is really about how aversion to various measures of risk might manifest themselves in a market setting. Despite the strong competitive equilibrium convergence properties of the double-auction institution, we were able to uncover indications of statistical discrimination, mainly among male subjects. The robust result from this study is that we find evidence of loss aversion among employers; although, this is again only significant among male subjects. Results from our female employer subsample indicate that females did not alter wage contracts to workers from more risky productivity distributions. This gender difference cannot be explained by a hypothesis of female aversion to negotiations/competition in the double-auction experiment environment. The only hypothesis consistent with this result would be female risk- or loss-neutrality or a combined effect of risk attitude and belief differences across gender. At this point we have no explanation as to why there should be a gender difference, though perhaps the labor market context we use may play a role. We do not report the results here, but we also examined whether or not the gender composition of the contract pair had any effect. The results showed that gender composition of the contract pair had no effect on wage contracts (these results are available on request).

There is an important message that emerges from these data. Statistical discrimination can exist in many forms, and only the most obvious forms of statistical discrimination--based on differences in average productivity among worker-groups--are likely to be measured in field studies. Even studies that examine distributional variance may not be capturing all the statistical discrimination in the data. Productivity risk from distinct worker-groups should be a concern, and our results indicate that current measures of statistical discrimination are predictably biased when this is not taken into account. Specifically, statistical discrimination will be underestimated when one ignores more hidden forms of this type of discrimination. (8) Furthermore, measures of prejudice-based discrimination may be overestimated if one fails to account for the likelihood that a certain component of unexplained wage differentials is due to a form of statistical discrimination not usually considered. Policy prescriptions aimed at reducing discrimination in various markets may require reassessment if the reason behind the discrimination has a different motive than typically thought.

Appendix A: Instructions--EMPLOYERS

This is an experiment in economic decision making. Please read and follow the instructions carefully. Your decisions, as well as the decisions of others, will help determine your total cash payment for participation in this experiment.

In this experiment, you are an Employer. Other individuals in the experiment will be workers. As an employer, you will have the ability to hire one unit of labor (at most) in each decision round from a pool of workers. You may wish to do this because a unit of labor will be assumed to produce a certain amount of output for you for that round. To keep things simple, whatever output a unit of labor produces, we will assume that you will sell each unit of that output for a market price of $1 (one experimental dollar). You will have the ability to hire a unit of labor in each round for a series of decision-making rounds. In each decision round, your experimental earnings will be determined by your employer "profits." Profits are calculated as total revenues minus total costs. Your employer profits in each round are then simple to calculate--your total revenues are given by the quantity of output that the unit of labor will produce for you (multiplied by the $1 that you receive for each unit of output), and your total costs are just given by whatever you agree to pay for the worker for his/her unit of labor.

You will receive specific and more detailed instructions on labor productivity shortly.

You are not required to purchase a unit of labor in each round. Rather, if you do not purchase a unit of labor in a given round, your profits for that round are zero (since total revenue and total cost are zero). If you do hire a unit of labor in a given round, your profits for that round will depend on both the productivity of labor (i.e., how much output the unit of labor produces for you) and the wage that you pay for that unit of labor. For example, if a worker produces three units of output for you, and if you agree to pay that worker $2, then your profits for that decision round would be $1 (remember, three units of output are assumed to be sold by you for $1 each, and so total revenues are $3). If, on the other hand, you agree to pay that worker S4, then your profits for that round would be -$1. In other words, one dollar would be subtracted from your total experimental earnings in that case. As such, your experimental earnings would be higher if you did not hire a unit of labor in a given round, as opposed to hiring a unit of labor and earning negative profits. The way in which you earn money in this experiment (through your profits) is private information to you and should not be discussed with other employers or with the workers.

In this experiment, there are a total of 5 employers and 10 workers. Each worker in the experiment has the ability to sell one unit of his labor to only one employer in each decision round, and each employer can hire only one unit of labor per decision round. As an employer, you will be allowed to freely "shop" around within the pool of workers in your attempt to hire one unit of labor for the round. Similarly, each worker will be allowed to freely shop among the employers in order to sell his/her unit of labor. Each round will last for a maximum of 2.5 minutes. The wages you and a worker mutually agree to and your per-round experimental profits will be calculated on the decision sheet that you have also been given. If you and a worker agree on a wage for a given round, the decision sheet also includes a space for you to document the identification number of the worker you purchased your unit of labor from for that round.

FOR TODAY'S EXPERIMENT, YOUR CASH EARNINGS ARE RELATED TO YOUR EXPERIMENTAL EARNINGS BY THE FOLLOWING EXCHANGE RATE: $1 EXPERIMENTAL = $_J_U.S.

Specific (Treatment) Instructions for--EMPLOYER

TREATMENT 1

For the next few rounds, each of the workers in the worker pool will be equally productive, and a unit of labor from any worker will produce 3 units of output. As such, if you mutually agree with any worker on hiring his/her unit of labor in a particular round, you know that the productivity of the worker will be 3 units of output.

TREATMENT 2-4 (combined for exposition only)

For the next few rounds, different workers may have different productivities, and you will not know the productivity of any given worker until after you have hired a unit of labor from that worker. You will, however, be given some general information on the entire group of workers.

The pool of workers for the following rounds has these characteristics (productivity refers to how many units of output a worker's unit of labor will produce for you):

Treatment 2

10% chance that a worker has productivity of 1

10% chance that a worker has productivity of 2

60% chance that a worker has productivity of 3

10% chance that a worker has productivity of 4

10% chance that a worker has productivity of 5

Treatment 3

20% chance that a worker has productivity of 1

20% chance that a worker has productivity of 2

20% chance that a worker has productivity of 3

20% chance that a worker has productivity of 4

20% chance that a worker has productivity of 5

Treatment 4

50% chance that a worker has productivity of 2

50% chance that a worker has productivity of 4

Neither you nor the workers know exactly how productive a worker will be until after the unit of labor is hired. You may seek to mutually agree upon a wage with any worker, but you will not know his/her productivity until after you have made your wage agreement with the worker. The workers do not know how productive their labor will be for an employer either. Workers see the same general worker characteristics that you see above.

Once the round is over, for all employers who hired a unit of labor, a random draw will be made from a Bingo cage to determine the productivity of the unit of labor. A separate draw will be made for each employer. Profits for each employer can then be calculated using the random draw of productivity to determine the total revenue that is generated by that unit of output. Your total costs are still just the agreed-upon wage for the unit of labor that you hired.

Finally, it is important for you to realize that each new round under this set of instructions will be conducted similarly. You may have made a wage agreement with a particular individual in a previous round which resulted in a productivity of 1,2, 3, 4, or 5. However, that does not affect in any way the probabilities for productivity for a future round, even if you re-hire the same person. In other words, if you make an agreement with Jane Doe in round one, and the random productivity draw says that the productivity for that unit of labor is 3, that does not imply that you can make an agreement with the same Jane Doe in the next round and be guaranteed a productivity of 3. The productivity that Jane Doe's unit of labor provides for you or any other employer in any round will always be determined by a new draw from the Bingo cage. Each round should be treated as independent from any other round in terms of determining worker productivity after agreements have been made--even though the pool of workers is still physically composed of the same individuals. Please raise your hand if this is confusing in any way!

ALL TREATMENTS

Each decision round is 2.5 minutes long, and the experiment will continue in this fashion until you are given different instructions. If you and a worker agree on a wage for a given round, the decision sheet also includes a space for you to document the identification number of the worker you purchased your unit of labor from for that round.

Your decision sheet for these rounds is attached to these instructions. Please raise your hand if at any point you have questions about how each round will proceed and/or how to correctly fill out your decision sheet.

TOTAL PROFITS FOR THIS DECISION SHEET—

Appendix B: Instructions--WORKERS

This is an experiment in economic decision making. Please read and follow the instructions carefully. Your decisions, as well as the decisions of others, will help determine your total cash payment for participation in this experiment.

In this experiment, you are a Worker. Other individuals in the experiment will be employers. As a worker, you will have the ability to sell one unit of labor (at most) in each decision round to only one employer. You may wish to do this because selling a unit of labor will provide you with a wage for that round. You will have the ability to sell a unit of labor in each round for a series of decision-making rounds. In each decision round, your experimental earnings will be determined by the wage you can obtain from selling your unit of labor. Employers may be interested in paying you a wage for your unit of labor because your labor produces output for the employer, which we will assume the employer can sell for profit.

You will receive specific and more detailed instructions on labor productivity shortly.

You are not required to sell a unit of labor in each round. Rather, if you do not sell a unit of labor in a given round, you will still earn a minimal $.40 for that round. If you do sell your one unit of labor in a given round, then your experimental earnings for that round will be the wage you mutually agree upon with the employer. For example, if you agree with an employer to sell your unit of labor for $1.00, then your earnings for that round would be $1.00 (one experimental dollar). If you agree with an employer to sell your labor for $.25, then your earnings for that round would be $.25. If you do not sell your unit of labor to any employer, then your earnings for that round are $.40. As such, your experimental earnings would be higher if you did not sell your unit of labor in a given round, as opposed to selling it for less than $.40. The way in which you earn money in this experiment (through wages) is private information to you and should not be discussed with other workers or with the employers.

In this experiment, there are a total of 5 employers and 10 workers. Each worker in the experiment has the ability to sell one unit of his labor to only one employer in each decision round, and each employer can hire only one unit of labor per decision round. As a worker, you will be allowed to freely "shop" around among the employers in your attempt to sell one unit of labor for the round. Similarly, each employer will be allowed to freely shop among the pool of workers in order to hire his/her unit of labor. Each round will last for a maximum of 2.5 minutes. The wages you and an employer mutually agree to and your per-round experimental profits will be calculated on the decision sheet that you have also been given. If you and an employer agree upon a wage for a given round, the decision sheet also includes a space for you to document the identification number of the employer you sold your unit of labor to for that round.

FOR TODAY'S EXPERIMENT, YOUR CASH EARNINGS ARE RELATED TO YOUR EXPERIMENTAL EARNINGS BY THE FOLLOWING EXCHANGE RATE: $1 EXPERIMENTAL = $_1_U.S.

Specific (Treatment) Instructions for WORKER

TREATMENT 1

For the next few rounds, each of the workers in the worker pool will be equally productive, and a unit of labor from any worker will produce 3 units of output. As such, if you mutually agree with any employer on selling your unit of labor in a particular round, the employer will know that the productivity of your unit of labor will be 3 units of output.

TREATMENT 2-4 (combined for exposition only)

For the next few rounds, different workers may have different productivities, and employers will not know the productivity of any given worker until after the employer has hired (and you have sold) the unit of labor. As a worker, you will not know either what your own productivity will be for that employer until after your labor unit is sold. You will, however, be given some general information on the entire group of workers. The employers are given this general information as well, and productivity refers to how many units of output a worker will produce for the employer who purchases his/her unit of labor.

The pool of workers for the following rounds has these characteristics:

Treatment 2

10% chance that a worker has productivity of 1 10% chance that a worker has productivity of 2 60% chance that a worker has productivity of 3 10% chance that a worker has productivity of 4 10% chance that a worker has productivity of 5

Employer Decision Sheet					
Round #	Productivity of Worker	Output price	Mutually agreed-upon wage	Worker ID#	Profits =(productivity times output price, minus the wage)
1		$1			
2		$1			
3		$1			
4		$1			

Treatment 3

20% chance that a worker has productivity of 1 20% chance that a worker has productivity of 2 20% chance that a worker has productivity of 3 20% chance that a worker has productivity of 4 20% chance that a worker has productivity of 5

Treatment 4

50% chance that a worker has productivity of 2 50% chance that a worker has productivity of 4

Neither you nor the employers know exactly how productive a worker will be until after the unit of labor is hired. You may seek to mutually agree upon a wage with any employer, but the employer will not know your productivity for that round until after you have made your wage agreement with the employer.

Once the round is over, for all employers who hired a unit of labor, a random draw will be made from a Bingo cage to determine the productivity of the unit of labor (for the purposes of the employer's calculation of profits). A separate draw will be made for each employer. As a worker, your experimental earnings for each round are still determined by the wage agreed upon with the employer (or $.40 in a round when you do not sell your unit of labor to any employer).

Finally, it is important for you to realize that each new round under this set of instructions will be conducted similarly. An employer may have made a wage agreement with you in a previous round which resulted in a productivity of 1, 2, 3, 4, or 5. However, that does not affect in any way the probabilities for your productivity for a future round. In other words, if you make an agreement with an employer in round one, and the random productivity draw says that the productivity for your unit of labor is 3, that does not imply that your productivity is guaranteed to be 3 in the next round. The productivity that your unit of labor provides to any employer (even the same one) in any round will always be determined by a new draw from the Bingo cage. Each round should be treated as independent from any other round in terms of determining worker productivity after agreements have been made---even though the pool of workers is still physically made of the same individuals. Please raise your hand if this is confusing in any way!

Worker Decision Sheet			
Round #	Mutually agreed-upon wage	Employer ID#	Earnings =(agreed-upon wage *or* $.40 if your unit of labor was not sold)
1			
2			
3			
4			

ALL TREATMENTS

Each decision round is 2.5 minutes long, and the experiment will continue in this fashion until you are given different instructions. If you and an employer agree upon a wage for given round, the decision sheet also includes a space for you to document the identification number of the employer you sold your unit of labor to for that round.

Your decision sheet for these rounds is attached to these instructions. Please raise your hand if at any point you have questions about how each round will proceed and/or how to correctly fill out your decision sheet.

TOTAL PROFITS FOR THIS DECISION SHEET--

TABLES

Table 1. Experiment Treatment Design

Treatment	Description Productivity (Probability)	Productivity Mean
1	3 (1.00)	3
21	1,2,3,4,5 (0.1, 0.1, 0.6, 0.1, 0.1)	3
31	1,2,3,4,5 (0.2, 0.2, 0.2, 0.2, 0.2)	3
4	2,4 (0.5, 0.5)	3

Treatment	Productivity Variance	Productivity Distribution Support	Likelihood of Productivity < Mean Productivity
1	0	3	0
21	1	1-5	0.20
31	2	1-5	0.40
4	1	2-4	0.50

Table 2. Wage Contracts (Random Effects)

	Full Employer Sample		Male Employer Sample	
Variable	Coeff.	Std Error	Coeff.	Std Error
Constant	0.861	0.038 ***	0.906	0.061 ***
T2	-0.003	0.029	-0.063	0.042
T3	-0.028	0.036	-0.069	0.049
T4	-0.062	0.029 **	-0.116	0.044 ***
Round 2	-0.119	0.028 ***	-0.116	0.039 ***
Round 3	-0.143	0.028 ***	-0.153	0.039 ***
Round 4	-0.149	0.028 ***	-0.167	0.039 ***
T02	-0.181	0.034 ***	-0.154	0.048 ***
T03	-0.183	0.029 ***	-0.193	0.042 ***
T04	-0.186	0.029 ***	-0.177	0.044 ***
[R.sup.2]	0.137		0.168	
Nobs	560		240	

	Female Employer Sample	
Variable	Coeff.	Std Error
Constant	0.812	0.050 ***
T2	0.048	0.042
T3	0.012	0.052
T4	-0.027	0.039
Round 2	-0.121	0.039 ***
Round 3	-0.135	0.039 ***
Round 4	-0.135	0.039 ***
T02	-0.193	0.048 ***
T03	-0.153	0.041 ***
T04	-0.172	0.040 ***
[R.sup.2]	0.123	
Nobs	320	

** and *** indicate significance at the 0.05 and 0.01 levels, respectively, for the two-tailed test.

Table 3. Binary Comparisons among the Uncertain Productivity Treatments

	Full Employer Sample		Male Employer Sample	
Comparison	Coeff. Diff.	Std Error	Coeff. Diff.	Std Error
T3-T2	-0.025	0.031	-0.006	0.041
T4-T2	-0.059	0.030 **	-0.053	0.049
T4-T3	-0.034	0.036	-0.047	0.051
Nobs	560		240	

	Female Employer Sample	
Comparison	Coeff. Diff.	Std Error
T3-T2	-0.036	0.046
T4-T2	-0.075	0.043 *
T4-T3	-0.039	0.053
Nobs	320	

Coefficient comparisons from Table 2 results.

* and ** indicate significance at the 0.10 and 0.05 levels, respectively, for the two-tailed test.

Table 4. Wage Contracts with Alternative Measures of Risk (Random Effects)

		Full Employer Sample		
Variable	Coeff.	Std Error	Coeff.	Std Error
Constant	0.863	0.037 ***	0.861	0.038 ***
Variance	-0.023	-0.017	0.006	0.042
Support			0.056	0.123
Loss Prob.			-0.159	0.092 *
Round 2	-0.119	0.028 ***	-0.119	0.028 ***
Round 3	-0.143	0.028 ***	-0.143	0.028 ***
Round 4	-0.149	0.028 ***	-0.149	0.028 ***
T02	-0.165	0.033 ***	-0.181	0.034 ***
T03	-0.197	0.028 ***	-0.183	0.029 ***
T04	-0.199	0.028 ***	-0.186	0.034 ***
[R.sup.2]	0.132		0.137	
Nobs	560		560	

Male Employer Sample

Variable	Coeff.	Std Error		Coeff.	Std Error	
Constant	0.898	0.060	***	0.906	0.061	***
Variance	-0.042	0.024	*	0.048	0.058	
Support				-0.141	0.17	
Loss Prob.				-0.273	0.136	**
Round 2	-0.116	0.038	***	-0.116	0.039	***
Round 3	-0.153	0.038	***	-0.153	0.039	***
Round 4	-0.167	0.038	***	-0.167	0.039	***
T02	-0.143	0.047	***	-0.154	0.048	***
T03	-0.213	0.040	***	-0.193	0.042	***
T04	-0.218	0.039	***	-0.177	0.044	***
[R.sup.2]	0.159			0.168		
Nobs	240			240		

Female Employer Sample

Variable	Coeff.	Std Error		Coeff.	Std Error	
Constant	0.826	0.048	***	0.812	0.050	***
Variance	-0.001	0.025		-0.014	0.062	
Support				0.209	0.176	
Loss Prob.				-0.108	0.129	
Round 2	-0.121	0.039	***	-0.121	0.039	***
Round 3	-0.135	0.039	***	-0.135	0.039	***
Round 4	-0.135	0.039	***	-0.135	0.039	***
T02	-0.185	0.045	***	-0.193	0.048	***
T03	-0.176	0.039	***	-0.153	0.041	***
T04	-0.177	0.039	***	-0.172	0.040	***
[R.sup.2]	0.116			0.123		
Nobs	320			320		

*, **, and *** indicate significance at the 0.10, 0.05, and 0.01 levels, respectively, for the two-tailed test.

ENDNOTES

(1) Cornell and Welch (1996) consider that it is less costly to assess workers with similar backgrounds; thus a "screening" discrimination results. Lang (1986) also considers discrimination from differential communication costs across groups.

(2) Additional evidence of statistical discrimination is found in mortgage lending (Ladd 1998). car price negotiations (Ayers and Siegelman 1995; Goldberg 1996; Harless and Hoffer 2002), sports card price negotiations (List 2004), vehicle repair estimates (Gneezy and List 2006). and law enforcement decisions (Applebaum 1996). See also the discussion in Loury (1998).

(3) That is, workers are assigned cost values and paid the difference between the negotiated wage and the assigned cost value. Employer demand values depend on the productivity of the worker hired, and employers are paid the difference between the marginal revenue product of the worker and the negotiated wage.

(4) Average wage contracts for each of treatments 1, 2, 3, and 4 were 0.65 ([sigma] = 0.36), 0.62 ([sigma] = 0.26), 0.57 ([sigma] = 0.26), and 0.54 ([sigma] = 0.25), respectively.

(5) As before, the random effects estimates are identical to those from fixed effects or OLS specifications due to our particular design, though the estimated standard errors in OLS will differ from those in fixed or random effects (see Oaxaca and Dickinson 2005).

(6) We conduct a wage regression identical to the full employer sample in Table 4, while including a dummy variable for female employer. The coefficient on this variable is statistically no different from zero (p = 0.84). We also find statistically insignificant effects of gender-composition dummy variables. These results are available from the authors on request.

(7) This result is due to the single-period framework we utilize. In a multiperiod framework where market participants can have repeated interactions, this result may not hold.

(8) This assumes that groups with lower average productivity are the same groups that have riskier distributions. Otherwise, these two forms of statistical discrimination would have opposing effects in the data.

REFERENCES

Aigner, Dennis J., and Glen G. Cain. 1977. Statistical theories of discrimination in labor markets. Industrial and Labor Relations Review 30:175-87.

Altonji, Joseph G., and Charles R. Pierret. 2001. Employer learning and statistical discrimination. Quarterly Journal of Economics 116:313-50.

Anderson, Donna M., and Michael J. Haupert. 1999. Employment and statistical discrimination: A hands-on experiment. Journal of Economics 25:85-102.

Applebaum, Arthur Isak. 1996. Response: Racial generalization, police discretion and Bayesian contractualism. In Handled with discretion, edited by John Kleinig. Lanham, MD: Rowman and Littlefield, pp. 145-58.

Arrow, Kenneth J. 1972. Models of job discrimination. In Racial discrimination in economic life, edited by A. H. Pascal. Lexington, MA: D. C. Heath, pp. 83-102.

Arrow, Kenneth J. 1998. What has economics to say about racial discrimination? Journal of Economic Perspectives 12:92-100.

Ayers, Ian, and Peter Siegelman. 1995. Race and gender discrimination in bargaining for a new car. American Economic Review 85:304-21.

Babcock, Linda, Henry S. Farber, Cynthia Fobian, and Eldar Shafir. 1995. Forming beliefs about adjudicated outcomes: Perceptions of risk and reservation values. International Review of Law and Economics 15:289-303.

Babcock, Linda, and Sara Laschever. 2003. Women don't ask: Negotiations and the gender divide. Princeton, NJ: Princeton University Press.

Cornell, Bradford, and Ivo Welch. 1996. Culture, information, and screening discrimination. Journal of Political Economy 104:542-71.

Curley, Shawn P., and J. Frank Yates. 1985. The center and range of the probability interval as factors affecting ambiguity of preferences. Organizational Behavior and Human Decision Processes 36:273-87.

Davis, Douglas D. 1987. Maximal quality selection and discrimination in employment. Journal of Economic Behavior and Organization 8:97-112.

Fehr, Ernst, Erich Kirchler, Andreas Weichbold, and Simon Gachter. 1998. When social norms overpower competition: Gift exchange in experimental labor markets. Journal of Labor Economics 16:324-51.

Fershtman, Chaim, and Uri Gneezy. 2001. Discrimination in a segmented society: An experimental approach. Quarterly Journal of Economics 116:351-77.

Gneezy, Uri, and John A. List. 2006. Are the physically disabled discriminated against in product markets? Unpublished paper, University of Chicago.

Goldberg, Pinelopi Koujianou. 1996. Dealer price discrimination in new car purchases: Evidence from the Consumer Expenditure Survey. Journal of Political Economy 104:622-34.

Griffin, Dale, and Amos Tversky. 1992. The weighing of evidence and the determinants of confidence. Cognitive Psychology 24:411-35.

Harless, David W., and George E. Hoffer. 2002. Do women pay more for new vehicles? Evidence from transaction price data. American Economic Review 92:270-9.

Kahneman, Daniel, and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. Econometrica 47:263-91.

Ladd, Helen F. 1998. Evidence on discrimination in mortgage lending. Journal of Economic Perspectives 12:41-62.

Lang, Kevin. 1986. A language theory of discrimination. Quarterly Journal of Economics 101:363-81.

List, John A. 2004. The nature and extent of discrimination in the marketplace: Evidence from the field. Quarterly Journal of Economics 119:49-89.

Loury, Glenn C. 1998. Discrimination in the post-civil rights era: Beyond market interactions. Journal of Economic Perspectives 12:117-26.

Lundberg, Shelly J., and Richard Startz. 1983. Private discrimination and social intervention in competitive labor markets. American Economic Review 73:340-7.

Niederle, Muriel, and Lise Vesterlund. 2007. Do women shy away from competition? Do men compete too much? Quarterly Journal of Economics 122:1067-1101.

Neumark, David. 1999. Wage differentials by race and sex: The roles of taste discrimination and labor market information. Industrial Relations 38:414-45.

Oaxaca, Ronald L., and David L. Dickinson. 2005. The equivalence of panel data estimators under orthogonal experimental design. Unpublished paper, University of Arizona.

Phelps, Edmund S. 1972. The statistical theory of racism and sexism. American Economic Review 62:659-61. Smith, Vernon L. 1965. Experimental auction markets and the Walrasian hypothesis. Journal of Political Economy 73:387-93.

Smith, Vernon L. 1982. Microeconomic systems as an experimental science. American Economic Review 72: 923-55.

Tversky, Amos, and Daniel Kahneman. 1973. Availability: A heuristic for judging frequency and probability. Cognitive Psychology 5:207-32.

Wilson, William Julius. 1996. When work disappears: The worm of the new urban poor. New York: Alfred A. Knopf.