Utah State University

DigitalCommons@USU

Economic Research Institute Study Papers

Economics and Finance

2004

Statistical Discrimination in Labor Markets: An Experimental Analysis

David L. Dickenson
Utah State University

Ronald L. Oaxaca

Follow this and additional works at: <https://digitalcommons.usu.edu/eri>

Recommended Citation

Dickenson, David L. and Oaxaca, Ronald L., "Statistical Discrimination in Labor Markets: An Experimental Analysis" (2004). *Economic Research Institute Study Papers*. Paper 278.

<https://digitalcommons.usu.edu/eri/278>

This Article is brought to you for free and open access by the Economics and Finance at DigitalCommons@USU. It has been accepted for inclusion in Economic Research Institute Study Papers by an authorized administrator of DigitalCommons@USU. For more information, please contact digitalcommons@usu.edu.

Economic Research Institute Study Paper

ERI #2004-04

**STATISTICAL DISCRIMINATION IN LABOR MARKETS:
AN EXPERIMENTAL ANALYSIS**

by

DAVID L. DICKINSON

**Department of Economics and
Management and Human Resources Department
Utah State University
3530 Old Main Hill
Logan, UT 84322-3530**

RONALD L. OAXACA

**Department of Economics
University of Arizona
Tucson, AZ 85721**

March 2004

**STATISTICAL DISCRIMINATION IN LABOR MARKETS:
AN EXPERIMENTAL ANALYSIS**

David L. Dickinson, Assistant Professor

**Department of Economics
Utah State University
3530 Old Main Hill
Logan, UT 84322-3530**

Ronald L. Oaxaca, McClelland Professor of Economics

**Department of Economics
University of Arizona
Tucson, AZ 85721**

The analyses and views reported in this paper are those of the author(s). They are not necessarily endorsed by the Department of Economics or by Utah State University.

Utah State University is committed to the policy that all persons shall have equal access to its programs and employment without regard to race, color, creed, religion, national origin, sex, age, marital status, disability, public assistance status, veteran status, or sexual orientation.

Information on other titles in this series may be obtained from: Department of Economics, Utah State University, 3530 Old Main Hill, Logan, UT 84322-3530.

Copyright © 2004 by David L. Dickinson and Ronald L. Oaxaca. All rights reserved. Readers may make verbatim copies of this document for noncommercial purposes by any means, provided that this copyright notice appears on all such copies.

STATISTICAL DISCRIMINATION IN LABOR MARKETS:

AN EXPERIMENTAL ANALYSIS

David L. Dickinson and Ronald L. Oaxaca

ABSTRACT

Statistical discrimination occurs when distinctions between demographic groups are made on the basis of real or imagined statistical distinctions between the groups. While such discrimination is legal in some cases (e.g., insurance markets), it is illegal and/or controversial in others (e.g., racial profiling and gender-based labor market discrimination). “First moment” statistical discrimination occurs when, for example, female workers are offered lower wages because females are perceived to be less productive, on average, than male workers. “Second moment” discrimination occurs when risk averse employers offer female workers lower wages based not on lower average productivity but on a higher variance in their productivity. Empirical work on statistical discrimination is hampered by the difficulty of obtaining suitable data from naturally-occurring labor markets. This paper reports results from controlled laboratory experiments designed to study second moment statistical discrimination in a labor market setting. Since decision-makers may not view risk in the same way as economists or statisticians (i.e., risk = variance of distribution), we also examine two possible 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 individuals do respond to these alternative measures of risk, and employers made statistically discriminatory wage offers consistent with loss-aversion in our full sample (though differences between male and female employers can be noted). If one can transfer these results outside of the laboratory, they indicate that labor market

discrimination based only on first moment discrimination is biased downward. 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.

STATISTICAL DISCRIMINATION IN LABOR MARKETS: AN EXPERIMENTAL ANALYSIS¹

Introduction

When membership in a particular group conveys valuable information about an individual's skills, productivity, or other characteristics, an agent with no personal prejudice may still find it rational to statistically discriminate. Examples of statistical discrimination appear in a variety of settings such as wage or hiring decisions in labor markets, racial profiling in law enforcement, determinants of loan approval rates, differential premia for insurance, or even choosing friends or new church members. In some settings, statistical discrimination is legal and acceptable (e.g., insurance rates), whereas in other settings it is controversial and/or illegal (e.g., racial profiling and employment discrimination). Existing research on statistical discrimination has focused on 1st-moment statistical discrimination. That is, discriminatory wage offers to females or lower loan approval rates for individuals from minority racial groups are based on *average* productivity and default rates, respectively. Agents attribute the average characteristics of the group to each individual from that group when it is costly to gather information.

In this paper, 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 the riskiness of the distribution of group productivity (or default rate, accident rate, etc.). Using labor markets as an example, employers may make lower wage offers to females based on a higher productivity variance, even though average productivity may be identical to male employee productivity. If such variance-based statistical

¹The authors are grateful for research funding made possible by the McClelland Professorship. Valuable comments were provided by Bob Slonim and participants at the Economic Science Association meetings in Tucson.

discrimination is empirically documented, then existing measures of statistical discrimination are under-estimated and measures of prejudiced-based discrimination may be over-stated. In other words, 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. Four labor productivity distribution treatments are examined. In a given treatment, all workers belong to the labor pool and labor productivity is determined by an ex post random draw with probabilities based on the common knowledge productivity distribution. The productivity distribution of the labor pool differs across treatments, but average productivity is constant across all treatments. We find that subject-employers make significantly different wage offers as a result of various measures of risk that do not alter the average productivity of workers. The implication of our results is that statistical discrimination may be more pervasive than previously thought.

Statistical Discrimination

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. Lang (1986) argues that statistical discrimination can be caused also by a differential cost of communication with different groups—the minority group would bear the cost of the communication. In a somewhat similar vein, Cornell and Welch (1996) argue that statistical discrimination can result from a filtering situation in which employers, for example, find it less costly to assess workers with similar backgrounds to the employer's, thus creating “screening” discrimination. 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; the latter considers that the range of a probability distribution affects individual preferences). For risk-averse individuals, it seems clear that a less-risky outcome distribution would be preferred to a more risky distribution, although the riskiness of an outcome distribution may be defined in several different ways.

Empirical evidence alluding to statistical discrimination can be found in a variety of settings, though 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, observable marginal productivity is required to correctly identify statistical discrimination. 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 show that discrimination *not* based on productivity characteristics is observed, and it is attributed to poorer information about the discriminated-against group. In contrast, Altonji and Pierret (2001) utilize an econometric technique designed to identify statistical discrimination, and find little evidence for statistical discrimination based on race.

In credit markets it is illegal for lenders to discriminate against borrowers of a protected class, even if class turned out to be a good proxy for unobservable risk factors. Ladd (1998) reports evidence consistent with at least some amount of statistical discrimination in mortgage lending. Ayres and Siegelman (1995) and Goldberg (1996) use an audit study approach to examine discrimination in price negotiations at new car dealerships. The data reveal statistical discrimination, the argument being that dealers may infer different reservation values on

individuals buyers based on their race or gender. Similarly, List (2003) examines statistical discrimination in sports cards markets and finds that statistical discrimination explains observed differences in negotiations with minorities better than does prejudiced-based discrimination. Race also appears to affect law enforcement decisions (Applebaum, 1996), as is noted in the discussion in Loury (1998), who also emphasizes the difficulty in attributing causation to such race-based decisions.

Given some of the identification and causation issues inherent in field data approaches to examining discrimination, some have used controlled experiments to examine statistically-based discrimination. Anderson and Hauptert (1999) examine statistical discrimination where employers must decide whether or not to purchase additional information on workers (i.e., statistical discrimination based on imperfect information). Davis (1987) shows how maximal quality selection may imply that groups from which the employer draws fewer observations may lead to an inference of lower average productivity. Thus, statistical discrimination is shown to result from an incorrect inference about the productivity distribution of certain groups of workers. Finally, Fershtman and Gneezy (2001) examine behavior in simple economic experiments and find evidence that (incorrect) ethnic stereotypes—a type of statistical discrimination—are responsible for some of the observed patterns in the data.

Our paper adopts a laboratory approach to examine more hidden forms of statistical discrimination that are often difficult to examine from field data. Rather than study first-moment statistical discrimination, we focus on statistical discrimination that is more difficult to examine in the field. Average worker productivity in our experiment is identical, but what differs across treatments is “riskiness” of the worker-group’s productivity. Our focus is motivated by existing research that shows the potential importance on cognitive assessment of risk of not only the

distributional variance, but also the support of the distribution (Tversky and Kahneman, 1973; Curley and Yates, 1985; Griffin and Tversky, 1992; Babcock et al, 1995) and the potential for loss (Kahneman and Tversky, 1979). While others have found field evidence of statistical discrimination based on higher-order moments of a distribution (e.g., Ayers and Siegelman, 1995; Goldberg, 1996; List, 2003), our contribution is that we examine multiple measures of distributional risk, not just distributional variance. Additionally, our approach provides a more controlled environment in which to *precisely* manipulate the productivity distribution of the workers. Though this approach is less externally valid than field experiments or audit studies, the trade-off is necessary in order to precisely manipulate the “risk” variable in our design.

We employ a full information environment to examine the existence of higher-order statistical discrimination. Average worker productivity is identical, causation can only go one direction in our design, 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 distinct measures of riskiness of the worker-pool productivity distribution.

Experimental Design

We implement a two-sided auction market design to simulate a labor market. 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).

The baseline design we use is simple in that it generates clear equilibrium predictions. Specifically, the demand side of the experimental market consists of 5 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 3 units of output (each unit of output sells for \$1 experimental), and so the demand for labor is perfectly elastic at \$3.00 up to 5 units of labor. The supply side of the market consists of 10 workers, each with reservation wage of \$.50, and each able to sell at most one unit of labor services in each experimental market round. As such, the supply curve is perfectly elastic at \$.50 up until 10 units of labor. The predicted market wage is \$.50, and the predicted market quantity of labor traded is 5 units. We used the labels such as “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. 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 1, 2, 3, 4, or 5 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 (i.e., you cannot contract in the next round with John Doe to ensure productivity of 5 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* specific 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 1, 2, 3, 4, or 5 units of output with probability 20% for each possible outcome. In treatment 4, productivity of the labor pool is either 2 or 4 units of output with probability 50% for each.

The expected competitive employer profit is \$2.50 experimental dollars since the expected revenue is \$3.00 and the competitive wage is \$0.50. 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 sixteen times. Hence, we have a panel with 560 observations.

Table 1 describes the experimental design in terms of how each of the treatments varies with respect to various 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 less-than-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-than-expected 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.

Results

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. The results are random effects estimates, which account for differences in wage contracts across employers and possible correlation in the error terms across rounds for an individual employer's wage contracts. This random effects specification seems a reasonable approach to our panel data. OLS estimation is

rejected in favor of fixed effects and random effects. However, given our particular design, the coefficient estimates from the random effects specification are identical to those from a fixed effects or an ordinary least squares estimation with a single constant terms (see the Appendix for details). The random effects and fixed effects estimator for the treatment effects are identical and differ from OLS only in the estimated standard errors.

The Table 2 results show that, for the full sample, treatments 3 and 4 significantly lower wage contracts offered to workers, but the results from the gender-specific sample show that this is due entirely to the behavior of the male employers. Male employers offered significantly lower wage contracts in each of the 3 uncertain productivity treatments relative to certain productivity of workers in treatment 1. The largest decrease in wage contract occurred in treatment 4 for the male sample, in which wage contracts were 21 cents lower than in the certain productivity treatment. This amount is about 32% lower given the average wage contract level of about 65 cents). Female employers, on the other hand, did not offer significantly different wages across treatments. This is consistent with female employers being risk neutral.

Table 3 presents treatment effects comparisons from within the uncertain productivity treatments. Treatment 3 versus treatment 2 picks up the combined effects of greater variance and probability of less-than-expected profits. These combined effects are negative in all cases but statistically significant only for the full sample and the female sample. Treatment 4 versus treatment 2 reflects the combined effect of the smaller support but higher probability of less-than-expected profits in treatment 4. In all cases the combined effect is negative and statistically significant. This reflects the dominance of the loss aversion motive. Treatment 4 versus treatment 3 picks up the joint effect of a lower variance, a smaller support, and a higher probability of less-than-expected profits in treatment 4. The joint effect was negative in all cases

but statistically significant only for the male employer sample. Apparently for males the loss aversion motive dominates both of the other measures of lower risk when comparing treatment 4 with treatment 3.

Though these results presented thus far offer some initial evidence of statistical discrimination based on a measure of distributional risk, it is also the case that the treatment effects specification does not strictly 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. 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 variable, where loss probability is measured relative to expected (competitive) profits. As in Table 2, the Table 4 results are from a random effects specification, and results are presented for the entire employer sample as well as the gender-based employer subsamples.² From Table 4 we can see that for the overall sample, the only significant predictor of wage contract differences is the probability of profits less than average. The magnitude of *Loss Prob* at -.25 indicates, for example, that wage contracts were 12.5 cents lower in treatment 4 and in treatment 1 (19% lower given estimated average wage contracts of 65 cents in treatment 1).

Results for male versus female employer wage contracts again show intriguing differences in individual's response to the incentives of the different productivity distributions.

²As 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 will differ (see Appendix).

Male employers significantly decreased wage contracts in the *Loss Prob* treatment (by about 33%), while female employers did not significantly alter wage contracts in response to a change in the probability of less-than-average profits. Female employers did, however, offer *higher* wage contracts to workers in the treatments with a larger difference between highest and lowest possible worker productivity. This is somewhat puzzling, and it is consistent with risk preferring behavior that is at odds with earlier evidence consistent with risk neutrality among female employers. Male employers offered significantly lower wage contracts the larger the distribution support, though this result is marginally insignificant. Neither gender of employer significantly altered wage contracts as the variance in the productivity distributions changed (again the result is only marginally insignificant for male employers).

Our experimental data indicate that males may react more significantly to distinct measure of the riskiness of the productivity distribution than females. Overall, what is significant is that we find evidence for statistical discrimination in our sample that is not based on average group differences. Considering this labor market example, our full data sample show evidence that one variable in particular—a higher potential for less-than-average payoffs—can significantly decrease the wage that an employer would pay to individuals from the more risky labor pool.

Concluding Remarks

This paper has examined a very simple framework for second-moment statistical discrimination. Despite the strong competitive equilibrium convergence properties of the double-auction institution, we were able to uncover indications of statistical discrimination, mainly among male employers. At this point we have no explanation as to why there should be a gender difference. Although we do not report the results here, 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.

The next step in this line of research is to have two groups of workers with different productivity risks compete simultaneously in the market. This corresponds more naturally to field labor market institutions. We would also consider the implementation of upward sloping labor supply curves to add external validity to our design. Nonetheless, even at this initial stage there is an important message emerging from the 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 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 under-estimated when one ignores more hidden forms of this type of discrimination.³ Furthermore, measures of prejudice-based discrimination may be *over*-estimated 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 re-assessment if the reason behind the discrimination has a different motive than typically thought.

³This assumes that groups with lower average productivity are the same groups that have riskier distributions. Otherwise, these two forms of statistical discrimination would counteract each other in the data.

References

- Aigner, Dennis J., and Glen G. Cain. "Statistical Theories of Discrimination in Labor Markets." *Industrial and Labor Relations Review*, January 1977, vol. 30: 175-87.
- Altonji, Joseph G., and Charles R. Pierret. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics*, 2001, Vol. 116, No. 1(February), 313-350.
- Anderson, Donna M., and Michael J. Hauptert. "Employment and Statistical Discrimination: A Hands-On Experiment." *The Journal of Economics*, 1999, vol. 25(1): 85-102.
- Applebaum, Arthur Isak. "Racial Generalization, Police Discretion and Bayesian Contractualism." In John Kleinig, ed. *Handled with Discretion*. Lanham, MD: Rowman and Littlefield, 1996.
- Arrow, Kenneth J. "Models of Job Discrimination." In A.H. Pascal, ed. *Racial Discrimination in Economic Life*. Lexington, MA: D.C. Heath, 1972: 83-102.
- Arrow, Kenneth J. "What has Economics to Say About Racial Discrimination?" *Journal of Economic Perspectives*, 1998, Vol. 12(2): 92-100.
- Ayers, Ian., and Peter Siegelman. "Race and Gender Discrimination in Bargaining for a New Car." *The American Economic Review*, vol. 85(3): 304-21
- Babcock, Linda, Henry S. Farber, Cynthia Fobian, and Eldar Shafir. "Forming Beliefs about Adjudicated Outcomes: Perceptions of risk and reservation Values." *International Review of Law and Economics*. 1995, vol. 15: 289-303.
- Cornell, Bradford, and Ivo Welch. "Culture, Information, and Screening Discrimination." *Journal of Political Economy*, 1996, Vol 104(3): 542-71.
- Curley, Shawn P., and J. Frank Yates. "The Center and Range of the Probability Interval as Factors Affecting Ambiguity of Preferences." *Organizational behavior and human decision processes*, 1985, vol. 36: 273-87.
- Davis, Douglas D. "Maximal Quality Selection and Discrimination in Employment." *Journal of Economic Behavior and Organization*, 1987, vol. 8: 97-112.
- Fershtman, Chaim., and Uri Gneezy. "Discrimination in a Segmented Society: An Experimental Approach." *The Quarterly Journal of Economics*, 2001 (February): 351-77.
- Goldberg, Pinelopi Koujianou. "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, Vol 104(3): 622-34.

- Griffin, D. and Amos Tversky. "The Weighing Of Evidence and the Determinants of Confidence." *Cognitive Psychology*, 1992, vol. 24: 411-35.
- Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk" *Econometrica*, 1979, Vol. 47, No. 2(March), 263-91.
- Ladd, Helen F. "Evidence on Discrimination in Mortgage Lending." *Journal of Economic Perspectives*, 1998, vol. 12(2): 41-62.
- Lang, Kevin. "A Language Theory of Discrimination." *Quarterly Journal of Economics*, 1986, volume 101(2): 363-81.
- List, John A. "The Nature and Extent of Discrimination in the Marketplace: Evidence From the Field." *The Quarterly Journal of Economics*, forthcoming.
- Loury, Glenn C. "Discrimination in the Post-Civil Rights Era: Beyond Market Interactions." *Journal of Economic Perspectives*, 1998, vol. 12(2): 117-126.
- Lundberg, Shelly J., and Richard Startz. "Private Discrimination and Social Intervention in Competitive Labor Markets." *American Economic Review*, June 1983, vol. 73(3): 340-47.
- Neumark, David. "Wage Differentials by Race and Sex: The Roles of Taste Discrimination and Labor Market Information." *Industrial Relations*, July 1999, vol. 38(3): 414-45.
- Phelps, Edmund S. "The Statistical Theory of Racism and Sexism." *American Economic Review*, Sept. 1972, vol. 62: 659-61.
- Smith, Vernon L. "Experimental Auction Markets and the Walrasian Hypothesis." *The Journal of Political Economy*, 1965 (August), volume LXXIII(number 4): 387-93.
- Smith, Vernon L. "Microeconomic Systems as an Experimental Science." *American Economic Review*, 1982, 72(5): 923-955.
- Tversky, Amos, and Daniel Kahneman. "Availability: A Heuristic for Judging Frequency and Probability." *Cognitive Psychology*, 1973, vol. 5: 207-232.
- Wilson, William Julius. *When Work Disappears: The World of the New Urban Poor*. New York: Alfred A. Knopf, 1996.

FIGURE 1: Experimental Design

TABLE 1 Experimental Treatment Design					
Treatment	Description Productivity (probability)	Productivity Mean	Productivity Variance	Productivity Distribution Support	Likelihood of Productivity < mean Productivity
1	3 (1.00)	3	0	--	--
2	1,2,3,4,5 (.1,.1,.6,.1,.1)	3	1	1-5	.20
3	1,2,3,4,5 (.2,.2,.2,.2,.2)	3	2	1-5	.40
4	2,4 (.5,.5)	3	1	2-4	.50

TABLE 2
Random effects estimation
Dependent Variable=Wage Contract

Variable	Full Employer Sample (N=560) Coef. (p-value)	Male Employer Sample (N=240) Coef. (p-value)	Female Employer Sample (n=320) Coef. (p-value)
Constant	.653 (.00)***	.723 (.00)***	.601 (.00)***
T2	-.029 (.34)	-.127 (.00)***	.045 (.28)
T3	-.078 (.01)***	-.137 (.00)***	-.035 (.41)
T4	-.115 (.00)***	-.213 (.00)***	-.041 (.32)
R ²	.023	.066	.015

*, **, *** denote significance at the .10, .05, and .01 level, respectively, for the one-tailed test

TABLE 3
Binary Comparisons Among the Uncertain Productivity Treatments
(coefficient comparisons from Table 2 results)

Comparison	Full Employer Sample (N=560) Difference (P-value)	Male Employer Sample (N=240) Difference (P-value)	Female Employer Sample (N=320) Difference (P-value)
T3-T2	-.049 (.10)*	-.010 (.81)	-.080 (.06)*
T4-T2	-.086 (.00)***	-.086 (.04)**	-.086 (.04)**
T4-T3	-.037 (.22)	-.076 (.07)*	-.006 (.89)

*, **, *** denote significance at the .10, .05, and .01 level, respectively, for the one-tailed test .

TABLE 4
Random Effects Results
Dependent Variable=Wage Contract

Variable	Full Employer Sample (N=560) coef. (p-value)	Male Employer Sample (N=240) coef. (p-value)	Female Employer Sample (N=320) coef. (p-value)
Constant	.653 (.00)***	.723 (.00)***	.601 (.00)***
Variance	.001 (.99)	.087 (.11)	-.064 (.23)
Support	.053 (.68)	-.293 (.11)	.313 (.08)*
Loss Prob	-.252 (.00)***	-.482 (.00)***	-.079 (.51)
R ²	.023	.066	.015

*, **, *** indicate significance at the .10, .05, and .01 levels, respectively, for the two-tailed test.

APPENDIX: Statistical Design

The four treatments form an orthogonal design. If we let T_1, T_2, T_3 , and T_4 denote dummy variables vectors identifying the treatments, the observation matrix for the treatment variables is simply (T_1, T_2, T_3, T_4) . Our unit of observation is the contract, of which there are 560 wage contracts (5 employers per experimental session x 7 experimental sessions x 4 treatments per experimental session x 4 periods per treatment). Each treatment will be observed 140 times (5 employers per experimental session x 7 experimental sessions x 4 periods per treatment). The treatment cross product matrix appears as

$$\begin{bmatrix} T_1' T_1 & 0 & 0 & 0 \\ 0 & T_2' T_2 & 0 & 0 \\ 0 & 0 & T_3' T_3 & 0 \\ 0 & 0 & 0 & T_4' T_4 \end{bmatrix} = \begin{bmatrix} 140 & 0 & 0 & 0 \\ 0 & 140 & 0 & 0 \\ 0 & 0 & 140 & 0 \\ 0 & 0 & 0 & 140 \end{bmatrix}$$

If, say, the first treatment variable is replaced by a vector of 1's denoted by ι , the cross product matrix is given by

$$\begin{bmatrix} \iota' \iota & \iota' T_1 & \iota' T_2 & \iota' T_3 \\ \iota' T_1 & T_2' T_2 & 0 & 0 \\ \iota' T_2 & 0 & T_3' T_3 & 0 \\ \iota' T_3 & 0 & 0 & T_4' T_4 \end{bmatrix} = \begin{bmatrix} 560 & 140 & 140 & 140 \\ 140 & 140 & 0 & 0 \\ 140 & 0 & 140 & 0 \\ 140 & 0 & 0 & 140 \end{bmatrix}$$

The implication of this design is that any characteristic of the particular employer such as age or gender will have zero correlation with the treatment dummy variables. In particular, the set of dummy variables identifying each of the 35 employers in the experiment will be uncorrelated with the treatment dummy variables. This means that the estimated treatment effects will be invariant with respect to control variables such as age, gender, and employer. The latter case is the fixed effects model, which is demonstrated below.

Let Y denote the vector of wage contracts, X denote the observation matrix for treatments T_2, T_3 , and T_4 , and let D represent the observation matrix on the 35 employer dummy variables. The simple OLS model of treatment effects with a single constant term is given by $Y = \beta_1 \iota + X\delta + \varepsilon$, where ι is a vector of 1's. Without the constant term the model can be written as $Y = T_1\beta_1 + X\beta + \varepsilon$, where $\beta = \delta + \beta_1$ and $T_1 = 1 - T_2 - T_3 - T_4$. Given the orthogonality among the treatment dummy variables, the OLS estimator of β is $\beta_{ols} = (X'X)^{-1}X'Y = (x'x)^{-1}x'y$, where the lower case letters denote the variables in deviation form. Let N_1 =the number of employers per experimental session; N_2 =the number of experimental sessions; N_3 =the number of treatments per experimental session; and N_4 =the number of periods per treatment. Given that each employer appears exactly $N_3 \cdot N_4$ times out of a total of $N_1 \cdot N_2 \cdot N_3 \cdot N_4$ observations, the mean of each employer dummy variable is simply $\frac{1}{(N_1 \cdot N_2)}$. Each treatment appears exactly $N_1 \cdot N_2 \cdot N_4$ times. The number of times a given treatment interacts with each employer is N_4 , while the number of times a given treatment does not interact with a given employer is $(N_1 - 1)N_4 + N_1(N_2 - 1)N_4$. The term $(N_1 - 1)N_4$ is the number of times the given treatment interacts with the other employers in the same experimental session, while the term $N_1(N_2 - 1)N_4$ is the number of times the given treatment appears in the remaining experimental sessions. Let D_i denote the dummy variable for the i th employer and let d_i denote D_i in deviation form. Whenever $D_i = 1$, then $d_i = 1 - \frac{1}{(N_1 \cdot N_2)}$ and whenever $D_i = 0$, then $d_i = -\frac{1}{(N_1 \cdot N_2)}$.

We are interested in the sample covariance between any treatment T_k and employer D_i :

$$S_{T_k D_i} = \frac{T_k' d_i}{N_1 N_2 N_3 N_4}, \text{ which is proportional to } T_k' d_i. \text{ The row vector } T_k' \text{ consists of 1's for}$$

observations involving the k th treatment and 0's otherwise. These are multiplied by the

corresponding observation on d_i which equals 1 if the observation is on employer i and $1 - 1/(N_1 \cdot N_2)$ otherwise. The sum of these products can be decomposed into 3 terms. The first term is the value of $T'_k d_i$ for the sum of the observations in the experiment involving both employer i and treatment k : $T'_k d_i = N_4 \left(1 - \frac{1}{N_1 N_2} \right)$. The second term is the value of $T'_k d_i$ for the sum of the remaining observations in the experiment which correspond to either other employers or other treatments: $T'_k d_i = -\frac{N_4(N_1 - 1)}{N_1 N_2}$. Finally, the third term is the value of $T'_k d_i$ for the sum of the observations in the remaining experiments which do not involve employer i :

$$T'_k d_i = -\frac{N_4 N_1 (N_2 - 1)}{N_1 N_2}. \text{ Therefore, summing over all of the observations in the database shows}$$

that the sample covariance between any treatment T_k and employer D_i is proportional to

$$T'_k d_i = N_4 \left(1 - \frac{1}{N_1 N_2} \right) - \frac{N_4(N_1 - 1)}{N_1 N_2} - \frac{N_4 N_1 (N_2 - 1)}{N_1 N_2} = 0. \text{ This result establishes the zero}$$

correlation between the treatment dummy variables and the employer fixed effects. Consider the LSDV formulation of the employer fixed effects model: $Y = X\delta + D\alpha + \varepsilon$, where D is the observation matrix for the employer dummy variables. Given that the treatment effects in X are uncorrelated with the employer dummy variables in D , the LSDV fixed effects estimator $\hat{\beta}_{fe} = \hat{\beta}_{ols}$. The standard errors will, in general, differ because of the presence of fixed (employer) effects. It is clear that the treatment effect $\hat{\beta}_1 = \bar{D}' \hat{\alpha}$, where $\bar{D}'$ is the vector of mean values of the employer dummy variables.

The same sort of reasoning shows the equivalence of OLS and random effects as well. In fact, the identity of the random effects and the fixed effects estimators yields a value of zero for the Hausman statistic.