

Statistical Discrimination in Labor Markets: An Experimental Analysis.

David L. Dickinson

Dept. of Economics
Appalachian State University

Ronald L. Oaxaca

McClelland Professor of Economics
Dept. of Economics
University of Arizona

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 would occur 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 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 discrimination estimates based only on first-moment 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.

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

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, voting the party ticket in elections, differential premiums 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 first-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 risk of the distribution of group productivity (or default rates, accident rates, 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 discrimination is empirically documented, then existing measures of statistical discrimination are biased 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. We choose a labor market context for the laboratory environment because many existing empirical studies of statistical discrimination examine labor markets. However, the insights we gain extend to other contexts. 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 possibility of less-than-expected (i.e., average) profits, in particular, lowers the wage contract made by the average employer. 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 (2004) examines statistical discrimination in sports cards markets and finds that statistical discrimination explains observed differences in negotiations with minorities better than 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 the “risk” 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, 2004), 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 risk 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 “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* 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 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 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. For the statistical analysis discussed next, we also create independent variables that isolate the effects of changes in each distinct measure of distributional risk.

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, and we also include round dummy variables for rounds 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 term (see Dickinson and Oaxaca, 2005, 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 samples 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. Across rounds, the estimated coefficients indicate that wage contracts converge towards equilibrium in later rounds of each treatment.

Table 3 presents treatment effects comparisons (i.e., coefficient comparisons) from within the uncertain productivity treatments. Treatment 3 versus treatment 2 picks up the combined effects of greater variance and greater probability of less-than-expected profits. These combined effects are negative in all samples 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 samples 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 samples but statistically significant only for the male employer sample. If employers consider expected profits to be a reference point, then less-than-expected profits may be considered a loss by subjects. 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 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, 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 sub-samples.¹ 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. The magnitude of *Loss Prob* at -.25 indicates, for example, that wage contracts were 12.5 cents lower in treatment 4 than in treatment 1 (19% lower given estimated average wage contracts of 65 cents in treatment 1). As is the case in all estimated Table 2 models, Table 4 results shows evidence that outcomes converge towards equilibrium in later rounds of each treatment.

Results for male versus female employer wage contracts again show intriguing differences in individual's response to the incentives of the different productivity distributions. Male employers significantly decreased wage contracts when *Loss Prob* and *Support* are higher, while they increased wage contracts for high *Variance* treatments. Together, the magnitude of the effects is strongest for those risk factors that cause male subjects to decrease wage contracts ($p=.08$ on the Wald test of equal coefficients on *Support* and *Variance*). On the other hand, female employers did not significantly alter wage contracts in response to changes in *Loss Prob* or *Variance*. The only significant risk variable in the Female employer sample model, *Support*, implies *higher* wage contracts to workers with a larger difference between highest and lowest possible worker productivity. This is somewhat puzzling, and may point towards risk preferring behavior that is at odds with earlier evidence consistent with risk neutrality among female employers. The result is also consistent with female optimism as to the likely productivity draw

¹ 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

from a distribution with larger support. Recall that this is the opposite of the male subject response to changes in *Support*. Therefore, if beliefs as well as risk preferences are important determinants of wage contracts, there may be systematic differences in both of these across genders (e.g., males being either less optimistic, or having risk aversion that dominates any optimism towards the likely productivity draw).²

Another possible explanation for the gender difference in our results is that female employers negotiate worse outcomes (i.e., higher wage contracts) in general. Existing research on gender differences has shown that females are generally less driven by competition and more averse to negotiations than males. 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.

Suppose that female subjects in our sample are risk-averse, 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.³

Aversion to negotiations may also manifest itself in gender-matching patterns, with female

Dickinson and Oaxaca, 2005).

² 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.

³ 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 ($p=.84$). Also, we also find statistically insignificant effects of gender-composition dummy variables. These results are available from the authors on request.

employers more likely to contract with a female worker. In our sample, single-gender contracts—male-male or female-female agreements—are statistically significantly more likely than mixed-gender contracts (306 to 252 individual wage contracts— $p=.01$ for the one-sided binomial test). This pattern is also consistent with female aversion to competitive negotiations.

So, while there is some evidence that females may be more averse to negotiating with male subjects, 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 distinct measures of the productivity risk than females. Though we cannot fully explain the nature of this result, the overall significance is that we find evidence for statistical discrimination that is not based on average group differences. Considering this labor market context, 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.⁴ 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, the subjects' statistically discriminating behavior is consistent with loss aversion.

Concluding Remarks

This paper has examined a very simple framework for studying second-moment statistical discrimination. Despite the strong competitive equilibrium convergence properties of

⁴ This result is due to the single-period framework we utilize. In a multi-period framework where market participants can have repeated interactions, this result may not hold.

the double-auction institution, we were able to uncover indications of statistical discrimination, mainly among male subjects. 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 on wage contracts (these results are available on request).

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 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 under-estimated when one ignores more hidden forms of this type of discrimination.⁵ 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

⁵ 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.

may require re-assessment if the reason behind the discrimination has a different motive than typically thought.

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.
- Babcock, Linda, and Sara Laschever. *Women Don't Ask: Negotiations and the Gender Divide*. New Jersey: Princeton University Press, 2003.
- 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*, 2004, vol. 119(1): 49-89.
- 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.
- Oaxaca, Ronald L., and David L. Dickinson "The equivalence of panel data estimators under orthogonal experimental design." Working paper, University of Arizona.
- 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	<i>Description</i> Productivity (probability)	Productivity Mean	Productivity Variance	Productivity distribution support	Likelihood of Productivity < mean productivity
1	3 (1.00)	3	0	0	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. (st. error)	Male Employer Sample (N=240) Coef. (st. error)	Female Employer Sample (n=320) Coef. (st. error)
Constant	.756 (.036)***	.832 (.059)***	.698 (.045)***
T2	-.029 (.029)	-.127 (.041)***	.045 (.041)
T3	-.078 (.029)***	-.137 (.041)***	-.035 (.041)
T4	-.115 (.029)***	-.213 (.041)***	-.041 (.041)
Round 2	-.119 (.029)***	-.116 (.041)***	-.121 (.041)***
Round 3	-.143 (.029)***	-.153 (.041)***	-.135 (.041)***
Round 4	-.149 (.029)***	-.167 (.041)***	-.135 (.041)***
R ²	.067	.115	.055

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

TABLE 3

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

Comparison	Full Employer Sample (N=560) Difference (st. error)	Male Employer Sample (N=240) Difference (st. error)	Female Employer Sample (N=320) Difference (st. error)
T3-T2	-.049 (.029)*	-.010 (.041)	-.080 (.041)*
T4-T2	-.086 (.029)***	-.086 (.041)**	-.086 (.041)**
T4-T3	-.037 (.029)	-.076 (.041)*	-.006 (.041)

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

TABLE 4
Random Effects Results
Dependent Variable=*Wage Contract*

Variable	Full Employer Sample	Male Employer Sample	Female Employer Sample
	(N=560)	(N=240)	(N=320)
	coef. (st. error)	coef. (st. error)	coef. (st. error)
Constant	.756 (.036)***	.832 (.059)***	.698 (.045)***
Variance	.001 (.038)	.087 (.052)*	-.064 (.053)
Support	.053 (.127)	-.294 (.175)*	.313 (.177)*
Loss Prob	-.252 (.084)***	-.482 (.117)***	-.079 (.118)
Round 1	-.119 (.029)***	-.116 (.041)***	-.121 (.041)***
Round 2	-.143 (.029)***	-.153 (.041)***	-.135 (.041)***
Round 3	-.149 (.029)***	-.167 (.041)***	-.135 (.041)***
R ²	.067	.115	.055

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

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 \$4, 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 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 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 EARNING ARE RELATED TO YOUR EXPERIMENTAL EARNINGS BY THE FOLLOWING EXCHANGE RATE:

\$1 EXPERIMENTAL=\$ 1 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 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.

Decision Sheet for _____

Employer ID#_____

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			

TOTAL PROFITS FOR THIS DECISION SHEET _____

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 earning 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 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 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 EARNING 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
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 then 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!

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.

Decision Sheet for _____

WORKER ID#_____

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			

TOTAL PROFITS FOR THIS DECISION SHEET_____