

CLASSROOM GAMES

Experience-Based Discrimination

Roland G. Fryer, Jr., Jacob K. Goeree, and Charles A. Holt*

February 2002

Abstract: This paper presents a simple classroom game in which students are randomly designated as employers, purple workers, or green workers. This environment may generate “statistical” discrimination if workers of one color tend not to invest because they anticipate poor job assignments, and these beliefs are self-confirming as employers learn that it is more profitable to assign workers of that color to the non-managerial jobs. Such discriminatory equilibria may arise even when workers are ex-ante identical and the employer has no prior information regarding potential workers. The exercise typically generates a lively discussion about discrimination and how it may be addressed by alternative public policies.

Keywords: discrimination, experimental economics, classroom experiments

INTRODUCTION

There is little disagreement that some economic inequities may have arisen from historical and cultural factors. Such inequities would be less worrisome if they tended to diminish as attitudes and laws changed, but economic theory offers the disturbing possibility that experience-based (“statistical”) discrimination may be self-reinforcing, even in the absence of continuing biases and asymmetric opportunities. The theory (see Arrow, 1985, pgs. 143-164) is based on an informational asymmetry between a principal (i.e. employer, credit lender, etc.) and an agent (i.e. worker, borrower, etc.). For example, suppose that a job applicant must decide whether or not to make a costly investment in “skills” (training, education, etc.). These skills are not perfectly observable by the employer, who relies on a test that is more likely to yield a positive signal when the worker has invested. After observing the signal and other observationally identifiable physical markers (i.e. race, gender, etc.) the employer decides whether or not to hire the applicant. We say an equilibrium exhibits statistical discrimination if workers of one type tend not to invest because they anticipate a lower chance of being hired, and

* University of Chicago and the American Bar Foundation (Fryer), University of Amsterdam (Goeree), and

these beliefs are self-confirming in the sense that employers are more reluctant to hire workers from categories that have lower investment rates.

This paper describes a simple web-based classroom exercise that may produce such a pattern of experience-based discrimination. In each period, workers and employers are randomly paired. Each worker is given a cost of investment in “training,” which varies randomly across workers. After observing their own investment costs, workers decide whether or not to invest. The worker’s investment decision is not observed by the employer, but it does increase the chances that the worker will score higher on a “pre-employment test” that is observed. The employer uses this signal and the worker’s color (green or purple) to make a job assignment: “regular” or “managerial,” with payoffs that give employers an incentive to try to put only workers who have invested in the managerial jobs. Some repetition of this process is needed to generate patterns that indicate discrimination in the sense that job assignments are affected by both color and the test result, and worker’s investment decisions differ by color, after controlling for investment cost. However, the instructor should keep in mind that there are symmetric equilibria, and asymmetries between green and purple workers need not occur.

This exercise is appropriate for microeconomics courses at any level, and for courses in applied areas such as public economics, sociology, anthropology, education, public policy, and labor economics. The simple intuition can be explained in a principles class, whereas a presentation to an advanced undergraduate or graduate class can be more closely tied to the Nash equilibrium calculations.¹ In our experience, the discussion is unusually lively, as people can discuss their experiences in the experiment, which allows them to approach the more emotionally charged issues of discrimination with more objectivity.

SETTING UP THE EXPERIMENT

Although this experiment can be done “by hand,” the web-based program is much faster and easier, which makes it possible to involve more people and to go through more periods quickly, an important factor in spotting discriminatory patterns in the data. The instructor begins by running the setup program from any browser connected to the Internet, by going to:

University of Virginia (Holt). This work was funded in part by the National Science Foundation (SES-0094800).

¹ The equilibrium calculations are quite simple. We include a technical appendix to aid instructors in advanced undergraduate or graduate classes.

<http://veconlab.econ.virginia.edu/admin.php>². The setup for the Statistical Discrimination program (SD) is found on the Asymmetric Information menu. If you click on “SD – Statistical Discrimination,” you will see a short description of the program, and the next page in the “admin” sequence is the setup page that allows the instructor to specify the number of participants, which must be a multiple of four (one green worker, one purple worker, and two employers). One easy way to deal with the “extra” students is to have groups of two or three sit at each PC. The instructor also specifies the number of periods; which depends on the time available. We have found that it is possible to read the instructions and go through 20 periods in about 30 minutes, which leaves some time for discussion. In addition, there are many choices that can be made with respect to payoffs etc., but defaults are provided. After the parameters are selected, and submitted, you get the results pages for the Employers and the Workers, where the instructor can watch results of the experiment, as they are collected. These results pages can also be projected later during the class discussion.

After setting up the experiment, ask the students to login from PCs connected to any browser via the site: <http://veconlab.econ.virginia.edu/login.htm>. The login menu has the various categories of experiments, and the Asymmetric Information link takes you to a menu with “SD – Job Assignment Game,” which is linked to the login page for this particular game. Then the student enters their name (or names if several are sharing a computer). In addition, the students must enter the instructor’s session name. After logging on, a series of instructions pages introduce the participants to the structure of market. Although these instructions are self-paced, it is better to read them aloud so that students go through them at the same pace. This reading also serves to remind the instructor of the details that may be important in the class discussion of the experiment’s results. The instructor can read the instructions conveniently from the results page where the data will later be displayed as decisions are made.

The first period begins with each worker seeing their randomly determined investment cost (with each penny amount from \$0.00 to \$1.00 being equally likely). Workers decide whether or not to invest, while the employers see a “wait” window. After a worker decides to

² The “Guide to Experimenters” link at the bottom of this menu has a description of how to send an email message to obtain a “session name,” which will set aside a data base table for the exclusive use of that instructor. The “Guide to Experimenters” menu also has a listing for “Hints in Using the Software,” which will supply some details that are not covered here. After obtaining a session name, the instructor can go to the Veconlab Admin page listed above, and the menu shows the major types of experiments that are available: Auctions, Bargaining, Decisions, Asymmetric Information, Markets, and Public Choice.

invest or not, a test result (described below) is generated and passed to the employer matched with that worker. Although investment is costly, it increases the chances that the worker will do well on the pre-employment test. The employer sees the test result and the worker's color, but not the id number or the actual investment decision, and decides whether to assign the worker to a "managerial" job or a "regular" job. The worker sees a "wait" window while this decision is being made, after which the period-specific and cumulative earnings results are displayed privately for both worker and employer.

Recall, employers will not be able to see whether or not a worker has invested. In order to avoid using probabilistic terminology, the pre-employment test is explained in the context of draws of colored marbles. The program uses random numbers to draw marbles from a (virtual) cup, and the cup used by those who invest has a higher fraction of Blue marbles that represent good outcomes. In particular, the no-investment cup contains 5 Red marbles and 1 Blue marble and the investment cup contains 3 Red marbles and 3 Blue marbles. Notice that the chances of drawing a Blue marble are three times as high with the investment cup, although the employer cannot see which cup is being used, and hence does not know for sure whether or not the worker invested in the current period. The computer program makes the draws "with replacement" so that one marble is drawn and returned to the cup before a second draw is made³. Finally, note that the employer only knows the test result and the worker's color, not the worker's ID number or anything about the individual worker's past record of investment or not. Some "historical" information is provided in subsequent periods. Although workers and employers do not see each other's ID numbers or individual histories, the program posts the average aggregate investment and managerial job assignment percentages for workers of each color, green or purple. Workers and employers may use this information in making their subsequent decisions. To summarize, all periods have the same structure: workers see their own randomly generated investment costs and make investment decisions, and employers see the workers' colors and the test results that

³ When doing this by hand, you would have to find colored marbles or poker chips without distinguishing marks, and all draws must be made from the same physical container, so that students cannot guess the investment decision by looking at marks on the container. Alternatively, you may use draws from a stack of playing cards, with one suit (e.g. Diamonds) being designated as a bad outcome. This process will be time consuming, so you would probably have to work with only two workers and two employers, although each of these may be represented by a group of students. The workers have to be visually isolated from the employers, to avoid extraneous signals of investment decisions.




Figure 1: Experiment 1 Five Round Investment Averages

are correlated with the unobserved investment decisions, prior to making the job assignment decisions that determine wages and employer earnings.

For the particular experiments that we will discuss, the parameters were selected so that workers' earnings were higher in the managerial job, which pays \$3.00, to the regular job, which pays \$1.50, regardless or whether they have incurred the investment cost that is subtracted from the wage. Employers' earnings were selected so that it is better to assign the worker to the managerial job if the employer thinks investment is more likely than not; this is because the employer's earnings are \$2.00 for any worker assigned to a regular job, \$4.00 if the worker assigned to a managerial job invested, and \$0.00 if the worker assigned to the managerial job did not invest. It is not necessary to pay students in a classroom experiment, although it helps stimulate interest to announce that one person will be selected at random, ex post, and will be paid some fraction (e.g., 1/4) of their earnings.

CLASSROOM DISCUSSION

As with most experiments, it is impossible to predict in advance exactly what will happen. In some of our experiments, we have seen multiple equilibria arise as a function of the workers' group identity. In other experiments there has been very little separation. We have found that students tend to notice discrimination, when it exists, even if the patterns are fairly subtle in the data.

A fairly strong pattern was observed in the first experiment we conducted in an upper-




Figure 2: Experiment 1 Five Round Managerial Placement Averages

level economics class at the University of Virginia. Data from this experiment are represented in figures 1 and 2. In this class, discrimination against Purple workers emerged rather quickly. Although investment costs were about the same for both groups in the first period, costs were about 25 cents higher on average for purple workers in periods 2 and 3.⁴ This may have been a factor that kept investment rates much higher for Green workers in most periods, as shown in Figure 1.

The difference in investment behavior seems to have had a large impact on employers' responses to test results. The employers assigned every worker, regardless of color, to the managerial job when they received a "good" test score (BB). However, employers were more liberal with Green workers who received a mixed test score (BR or RB); they received managerial jobs 100 percent of the time, whereas Purples with mixed scores only received managerial jobs 78 percent of the time. The effect was even sharper following a bad outcome (RR): employers still placed Green workers in managerial jobs 64% of the time as opposed to 15% for Purple workers⁵. The period-by-period rates of assignment to the managerial job are shown in Figure 2. Overall, Green workers invested 85 percent of the time and Purple workers only 44 percent.

In another experiment at the University of Virginia, we changed the setup in order to

⁴ Over all, average investment costs for the Greens were actually higher in 8 of the 15 periods.

⁵ In the first two periods, all employers assigned Green workers with "bad" test scores to the regular job. However, after three periods (when the social statistics regarding Green workers were relatively optimistic) employers hired Green workers with "bad" test scores 75 percent of the time.




Figure 3: Experiment 2 Five Round Investment Averages

investigate the effects of historical discrimination. To do this, in the first five rounds of the experiment, the cost of investment for purple workers was higher than that of green workers, and the cost were drawn from identical distributions for the final fifteen rounds. We thought this would cause purple workers to choose not to invest early, thereby causing them to obtain the managerial job less often. In the first five rounds, the purple workers had very high investment costs, but still chose to invest. Because they chose to invest, they were assigned to managerial jobs at a rate slightly below the cost advantaged green workers. Once the barrier (higher investment cost) was removed at the end of round 5, the purple workers invested most of the time and were assigned to managerial jobs at a very high rate. The opposite was true for the green workers. To see this, consider figures 3 and 4. In the final five periods, 73% of purple workers invested as opposed to only 33% of green workers. Similarly, 93% of the purple workers were assigned to the managerial job and 54% of green workers were assigned. A similar reversal of anticipated separation was observed in a non-classroom (research) experiment conducted at the same location.

This exercise can generate an unusual amount of discussion, even in comparison with other classroom experiments. The best way to stimulate discussion is *not* to begin by explaining the concept of statistical discrimination, but rather to let the students draw their own lessons from the data, assisted by a series of questions (Socratic method). By engaging the students, one can get a feel for how they perceived the game, what strategies they used, etcetera. For the first experiment, when the green workers were favored, we began by asking for Purple workers to explain what they had been trying to do. A young woman with this role




Figure 4: Experiment 2 Five Round Managerial Placement Averages

exclaimed “this game is fixed,” assuming that we had fixed the parameters in such a way as to induce Purple workers not to invest. We ensured her that there were no systematic differences between the two groups. We then addressed the same question to employers, and a young man said “Purple workers just can’t be trusted...they won’t invest.” A young woman then retorted, “I stopped investing because no one would place me in the managerial job.” To this, the young man stated, “I did not put you in the managerial job because you did not invest.” These comments capture the essential intuition behind statistical discrimination. Because of their personal involvement in the exercise, it seems the students were able to firmly grasp this intuition. Several of the students with employer roles realized that Purple workers were not investing, so they stopped “giving them the benefit of the doubt.” Suspecting this, some workers reacted in different and interesting ways. One student chose to “free-ride.” He indicated that, given he was a Green worker; he knew he would be hired, so why invest all the time. A close look at the data gives a better depiction of this student’s strategy. He chose to invest enough to ensure that employers continued to give green workers the benefit of the doubt, but used several occasions to “shirk” and avoid the investment cost. On the other hand, another student announced “I invested every time—even when costs were high—because I felt confident that I would get the managerial job—because *I am green.*”

It is natural to link this experiment to racial profiling, labor market discrimination, and credit market discrimination. Within these areas, the instructor may want to engage the students in an open discussion as to what policy recommendations may help “break” the discriminatory equilibria. In our exercise, the students were eager to engage in policy relevant discussions.

One student recommended that the employers be told that the underlying distribution of talent was the same for both types of workers. However, another student quickly replied, saying, “I would still discriminate against Purple workers because they are not investing...why should I care about their talent, I am trying to make money.” Some of the more insightful comments suggested policies such as (1) subsidies to investment cost for Purple workers to help alleviate the gap in training, (2) better information technology, and (3) probationary hiring periods. Somewhat surprisingly, nearly all students agreed that results-oriented policies that required equal representation among workers would not be an adequate solution. Employers indicated that they did not want to be forced to hire Purple workers, and workers indicated that their incentives to invest would decrease.

In addition to providing insights about experience-based discrimination, the exercise can be used to demonstrate the calculation of a Bayesian Nash equilibrium in more advanced courses⁶. This equilibrium involves a critical level of investment cost that is higher for one color than for another, which in turn generates differential responses to test results. In equilibrium, these differential test scores coupled with observationally distinct physical markers can serve to validate “color-based” thinking on the part of employers

FURTHER READING

The environment described above is motivated by the well-known theory of statistical discrimination. Arrow [1973] develops a model that shows employers can (rationally) discriminate against a group even when they are ex-ante identical, and the employers themselves are psychically unbiased. The model proposed by Phelps [1972] is similar, but he assumes that minorities emit noisier signals, and therefore employers (rationally) discriminate against them in equilibrium. In contrast, Arrow does not need any assumed structural asymmetries. He notes that when some employee characteristics are endogenous, an employer's *a priori* beliefs can be self-fulfilling. The classroom experiment reported here is based more closely on the Coate and Loury (1993) model with worker-specific investment costs. See Fryer (2001b) for a recent survey of economic models of discrimination.

Although there are no direct experimental tests of these influential theories of statistical discrimination, there have been some related experimental studies. Davis (1987) justifies

⁶For an explicit calculation of discriminatory equilibria for the game, see Fryer, Goeree, and Holt (2002)

asymmetric *a priori* beliefs about two groups with identical talent distributions with the assumption that employers may observe a higher maximum quality from candidates from a majority group, since there are more signals obtained from that group. Of course, a larger group will have a lower minimum, but some discrimination may arise if there is a behavioral tendency to focus on the maximal quality, as seems to be the case in the experiments reported. Anderson and Hauptert (1999) describe a classroom exercise in which employers have the chance to buy perfect information about a worker's productivity. In this setup, it may be rational for employers to only buy information about a particular group, and thus, discriminatory equilibria follow.

REFERENCES

Anderson, D., and Hauptert, M., (1999) "Employment and Statistical Discrimination: A Hands on Experiment" *Journal of Economics* V. 25 no. 1. p. 85-102

Arrow, KJ. (1985) *Applied Economics* (Collected Papers of Kenneth J. Arrow), V. 6 Harvard University Press.

Coate, S., and Loury, G., (1993a) "Will Affirmative Action Eliminate Negative Stereotypes?" *The American Economic Review*, V. 83 no. 5 p. 1220-40

Davis, D., (1987) "Maximal Quality Selection and Discrimination in Employment" *Journal of Economic Behavior and Organization* Vol. 8 p. 97-112

Fryer, R., (2001a) "A Dynamic Theory of Statistical Discrimination". mimeo. The University of Chicago.

----- (2001b) "Economists' Models of Discrimination: An Analytical Survey" Unpublished Monograph. The University of Chicago.

-----, Goeree, J., and Holt, C., (2002) "An Experimental Test of Statistical Discrimination" Discussion Paper. The University of Virginia

Phelps, E., "The Statistical Theory of Racism and Sexism," *American Economic Review* (September 1972): 659-61.