

Experience-Based Discrimination: Classroom Games

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

Abstract: The authors presented 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 lower opportunities in the labor market, and these beliefs are self-confirming as employers learn that it is, on average, less profitable to hire workers of that color. 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.

Key words: classroom experiments, discrimination, experimental economics

JEL codes: A2, C9, J7

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 1973) is based on an informational asymmetry between a principal (employer, credit lender, etc.) and an agent (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 to hire the applicant. An equilibrium exhibits statistical discrimination if employers are reluctant to hire a certain type of worker based on expectations of low investment rates, and these employer beliefs are self-confirming in the sense that workers of this type tend not to invest because they anticipate a lower chance of being hired.

Roland G. Fryer, Jr., is a junior fellow, Harvard Society of Fellows and NBER, Harvard University. (e-mail: rfryer@fas.harvard.edu). *Jacob K. Goeree* is an associate professor of economics, CREED University of Amsterdam, The Netherlands. *Charles A. Holt* is the A. Willis Robertson Professor of Political Economy, University of Virginia. This work was funded in part by the National Science Foundation (SES-0109196 [Fryer] and SES-0094800 [Holt and Goeree]).

We describe 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 preemployment test that is given by the employer. The employer uses the test result and the worker's color (green or purple) to decide whether to hire the worker, with payoffs that give employers an incentive to hire only workers who have invested. Some repetition of this process is needed to generate patterns that indicate discrimination in the sense that hiring decisions are affected by both color and the test result and that workers' investment decisions differ by color, after controlling for investment cost. When running the experiment, the instructor should keep in mind, however, that symmetric outcomes without color effects are possible both in theory and practice.

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. In our experience, the discussion is unusually lively, as people can relate to their experiences in the experiment, which allows them to approach the emotionally charged issues of discrimination with more objectivity.

SETTING UP THE EXPERIMENT

Although this experiment can be done "by hand," the Web-based program, using the 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: <http://veconlab.econ.virginia.edu/admin.htm>; the Guide to Experimenters Menu at the bottom of the page has instructions for how to obtain a "session name" (database storage area) by e-mail. After obtaining a session name, click on the Asymmetric Information Menu and from there, select the Statistical Discrimination program (SD). If instructors click on SD—Statistical Discrimination, they will see a short description of the program, and the next two pages in the "admin" sequence allow the instructor to specify setup parameters. In particular, the number of participants must be a multiple of four (one green worker, one purple worker, and two employers). One easy way to deal with the remaining 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 and so for the, but defaults are provided. After the parameters are selected and submitted, instructors 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 the instructor sets up the experiment, the students log in 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 students to a menu with SD—Job Assignment Game, which is linked to the login page for this particular game. Then the students enter their name (or names if several are sharing a computer). In addition, the students must enter the instructor's session name so that they all are connected to the same experiment. After logging on, a series of instructions introduce the participants to the structure of the labor market. (For a complete transcript of the instructions, see the Appendix.) 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 a random pairing of workers and employers. Half of the workers are green and half are purple, with the worker's color being known by both the worker and the employer. Workers are only told that their randomly determined investment costs would be between \$0.00 and \$1.00 and that costs are independent for each worker "so that one worker's cost is not likely to be equal to that of another, and investment costs will differ randomly from period to period." The instructions do not say whether each penny amount in this interval is equally likely, nor do they say whether both colors are drawing from the same cost distribution. After seeing their own investment costs, workers decide whether to invest, and the employers see a "wait" window. After a worker decides to invest or not, a test result (described later) 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 preemployment 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 hire the worker. 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 a worker has invested, but the employer does administer a preemployment test. To avoid using probabilistic terminology, the test is explained in the context of two 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 from the investment cup, although the employer cannot see which cup is being used and thus does not know for sure whether or not the worker invested in the current period. The computer program makes the two draws "with replacement" so that

one marble is drawn and returned to the cup before the next marble is drawn.¹ Because the investment cup has more blue marbles, a draw of two blue marbles (BB) makes it likely that the worker invested. Conversely, the draw of two red marbles (RR) is a strong signal that the worker did not invest. The employer's beliefs on the bases of prior experience may have a greater weight when the test outcome is mixed with one red and one blue marble (RB or BR), which is the reason that we use two draws for the test technology.

Finally, 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. 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 hiring percentages for workers of each color, green or purple. Workers and employers may use this information in making their subsequent decisions. This aggregate information may correspond to the types of performance statistics that employers collect and share in some industries (e.g., redlining), but this type of information may not be available in some cases. We speculate that the effect of aggregate performance information may be to reinforce statistical discrimination when it arises.

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

For the particular experiments that we discuss, we selected the parameters so that workers' earnings were higher if they were hired (\$3.00) than was the case if they were not hired (\$1.50), regardless of whether they had incurred the investment cost that was subtracted from the base earnings amount. Employers' earnings were selected so that it was better to hire the worker if the employer thought investment was more likely than not; this was because the employer's earnings were \$2.00 if the worker was not hired, \$4.00 if a worker who invested was hired, and \$0.00 if a worker who did not invest was hired. 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 that investment and hiring decisions differed sharply on the basis of the worker's color. 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.

We observed a fairly strong pattern in the first experiment we conducted in an upper-level economics class at the University of Virginia. Data from this experiment are represented in Figures 1 and 2, which show data points averaged over

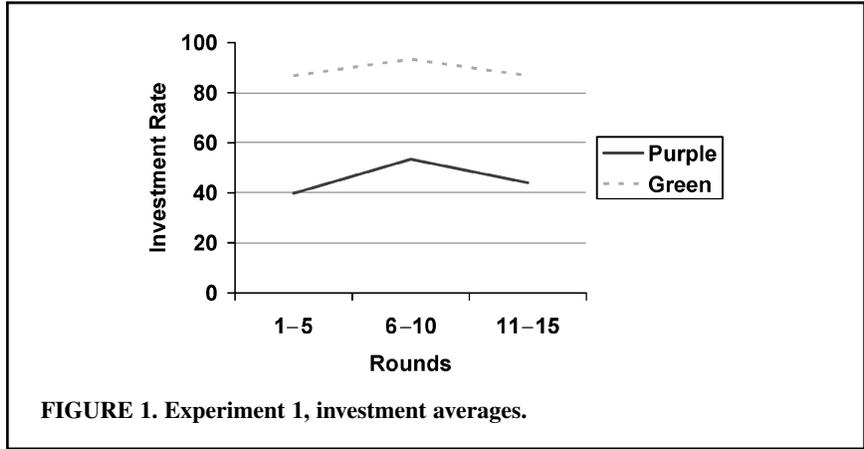


FIGURE 1. Experiment 1, investment averages.

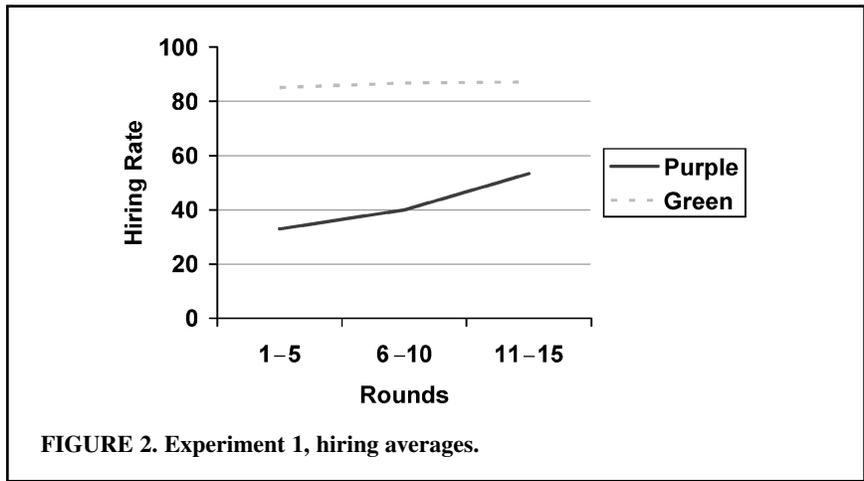


FIGURE 2. Experiment 1, hiring averages.

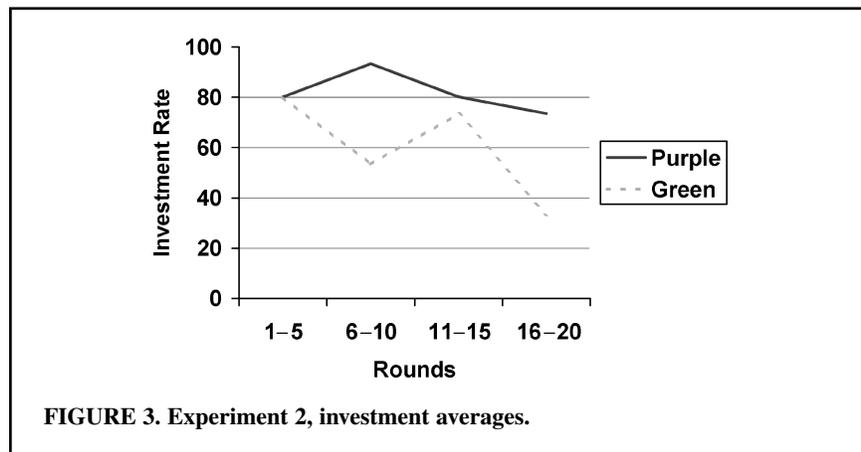
five periods. In this class, discrimination against purple workers emerged rather quickly. In this exercise, the investment cost draws were drawn from the same uniform distribution [\$0.00, \$1.00] for each color. Although investment costs were about the same for both groups in the first period, costs were about 25 cents higher on average for purple workers in periods 2 and 3.² This may have been a factor that kept investment rates much higher for green workers in most periods (Figure 1). Cost asymmetries, however, may have surprising effects. In the second experiment, the color with higher initial costs was not the one that ended up investing less.

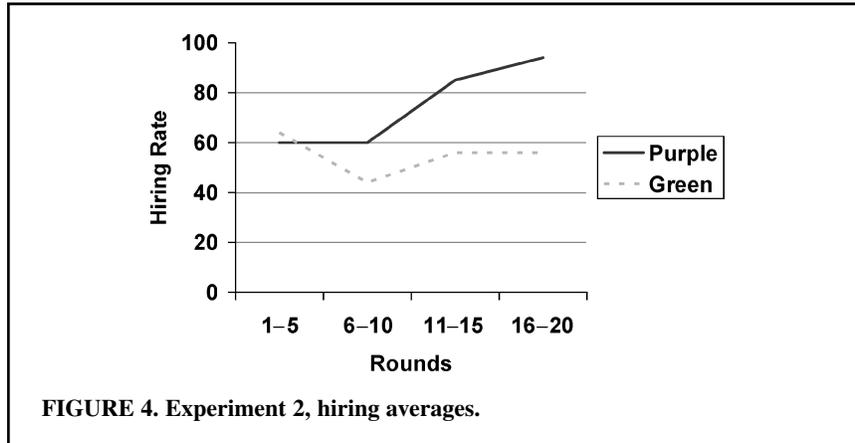
The difference in investment behavior seems to have had a large impact on employers' responses to test results. The employers hired every worker, regardless of color, when they received a good test score (BB). However, employers were more liberal with green workers who received a mixed test score (BR or RB); they hired, these workers 100 percent of the time, whereas they hired

purples with mixed scores only 78 percent of the time. The effect was even sharper following a bad outcome (RR): employers still hired green workers 64 percent of the time as opposed to 15 percent for purple workers³ (hiring rates are shown in Figure 2). Overall, green workers invested 85 percent of the time and purple workers only 44 percent.

In a second experiment at the University of Virginia, we changed the setup to investigate the effects of historical discrimination. To do this, in the first 5 rounds of the experiment, all investment cost draws for purple workers were from the top half of the [\$0.00, \$1.00] range, whereas all draws for the green workers were from the bottom half. Then in the final 15 rounds, all draws for both colors were from identical distributions (uniform on the whole range).⁴ We thought this would cause purple workers to choose not to invest early, thereby causing them to be hired less often. In the first 5 rounds, the purple workers had very high investment costs but still chose to invest. In fact, investment levels for purple workers were such that they were hired at a rate only 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 hired at a very high rate. The opposite was true for the green workers (see Figures 3 and 4). In the final 5 periods, 73 percent of purple workers invested as opposed to only 33 percent of green workers. Similarly, 93 percent of the purple workers were hired, and 54 percent of green workers were hired. For a classroom experiment, we prefer to use the initial cost asymmetry of the second experiment to increase the chances of observing interesting patterns of separation.

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, we could get a feel for how they perceived the game, what strategies they used, and so forth. 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 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 assured 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 hire me.” To this, the young man stated, “I did not hire you because you did not invest.” These comments captured the essential intuition behind statistical discrimination. Because of their personal involvement in the exercise, it seemed that 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 be hired—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 about 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 (as was the case for the first experiment). 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

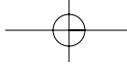
as: (1) subsidies to investment cost for purple workers to help alleviate the gap in training, (2) better information technology, and (3) probationary hiring periods. Somewhat surprisingly, nearly all students who openly expressed an opinion agreed that results-oriented policies that required equal representation among workers would not be an adequate solution.⁵ Employers indicated that they did not want to be forced to hire purple workers, and workers indicated that their incentives to invest would decrease. Finally, the instructor may want to point out the possible role of certification of an educational attainment as a way of breaking the cycle of statistical discrimination. An effective certification procedure has the effect of making the investment decision visible to the employer. The students may want to discuss whether grades in school provide an effective certification procedure or whether noise in grading and course selection processes make student transcripts more analogous to the noisy test technology used in the experiment.

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 courses. 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 observable distinct physical markers can serve to validate color-based thinking on the part of employers.

FURTHER READING

The environment described in this article is motivated by the well-known theory of statistical discrimination. Arrow (1973) developed a model that showed employers could (rationally) discriminate against a group even when they were ex-ante identical, and the employers themselves were psychically unbiased. The model proposed by Phelps (1972) was similar, but he assumed that minorities emitted noisier signals, and therefore employers (rationally) discriminated against them in equilibrium. In contrast, Arrow did not need any assumed structural asymmetries. He noted that when some employee characteristics were endogenous, an employer's *a priori* beliefs could be self-fulfilling. The classroom experiment reported in this article was based more closely on the Coate and Loury (1993) model with worker-specific investment costs. For a recent survey of economic models of discrimination, see Fryer (2002).

Although there are no direct experimental tests of these influential theories of statistical discrimination, there have been some related experimental studies. Davis (1987) justified 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 because 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) described 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 buy only information about a



particular group, and thus, discriminatory equilibria follow. In addition, there is a well-established social psychology literature on the effects of discrimination in simulated interactions (see the survey by Anderson, Fryer, and Holt 2002).

APPENDIX: INSTRUCTIONS

- **Rounds and Matchings:** The experiment sets up markets that are open for a number of rounds. Note: In each round, you will be matched with another person selected at random from the other participants. There will be a new random rematching each round.
- **Interdependence:** The decisions that you and the other person make will determine your earnings.
- **Decisions:** One person in each matched pair will have the role of a “worker” who makes an investment decision, and the other will be an “employer” who decides whether to hire the worker after seeing some information about the worker.
- **Earnings:** The worker earns the difference between the wage (if any) and the cost (if any) of investment. The employer earnings depend on the hiring decision and whether the worker invested or not, as explained below.
- **Role:** Each of you has an ID number; yours is _____. These IDs have been used to make some random role assignments. Half of you have been randomly assigned to be workers, and half have been randomly assigned to be employers. Your role will be that of a (worker or employer).
- **Worker Colors:** Half of the workers have been randomly assigned to be **Green**, and half have been randomly assigned to be **Purple**. Each worker’s color type will stay fixed throughout all rounds of the experiment. If you are a worker, your assigned color is (Green or Purple).
- **Preview:** Next we will explain how workers may decide whether or not to incur a randomly determined **investment** cost in order to acquire some skills. The employer cannot observe the worker’s investment decision, so a **pre-employment test** will be given. The test provides some information about the investment decision, but it is not perfect. Seeing the test and the worker’s color (but not the actual investment decision), the employer decides whether or not to hire the worker. The worker earns more if hired, but investment is costly and the test is imperfect. The employer earns more if a worker who is hired has made an investment.
- **Investment Cost:** At the start of each period, each worker must decide whether or not to spend money and purchase some training. The cost of this “investment” decision is randomly determined to be between **\$0.00** and **\$1.00**. The random investment cost is determined independently for each worker in each period, so one worker’s cost is not likely to be equal to that of another, and investment costs will differ randomly from period to period.
- **Investment Decision:** After seeing the investment cost, the worker must indicate a decision (**invest** or **not invest**). If a worker invests, then that worker’s investment cost is deducted from the wage received. If the worker does not invest, then no cost is incurred. The worker’s wage will be **\$1.50** if not hired and **\$3.00** if hired. Thus the worker will earn **\$1.50**, minus investment cost (if any) if not hired, or **\$3.00**, minus any investment cost if hired.

Investment Cup

Used if the worker invests
 Contents: B B B R R R
 3 blue ball(s) and 3 red ball(s)

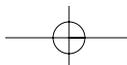
No-Investment Cup

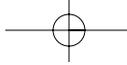
Used if the worker does not invest
 Contents: B R R R R R
 1 blue ball(s) and 5 red ball(s)

- **Pre-employment Test:** The test given by the employer will be made by a virtual draw of a colored ball from a cup. Each draw will either be a good outcome, Blue, or a bad outcome, Red. The cup (used by the computer program) for a worker who invested has a higher fraction of Blue balls, as can be seen by comparing the investment and no-investment cups above. Notice that the chances of drawing a Blue ball are 3 times as high with the investment cup. The employer cannot see which cup is being used, and hence cannot know for sure whether or not the worker invested in the current period.

Investment Cup: B B B R R R

No-investment Cup: B R R R R R





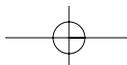
- **The Random Draws:** The test will consist of two independent draws from the same cup “with replacement” before the second draw, so the contents of the cup used do not change. Thus the possible test outcomes are: RR, RB, BR, BB. Finally, note that the employer will know the test result and the worker’s color, but not the worker’s ID or anything about the worker’s current or past record of investment or not.
- **Hiring Decision:** With this information, the employer decides whether or not to hire the worker. An investment by the worker only benefits the employer if the worker is hired, because the employer earns:
 - \$4.00, if a worker who invested is hired.
 - \$0.00, if a worker who did not invest is hired,
 - \$2.00, if the worker is NOT hired, whether or not the worker invested.
- **Earnings:** After each period, each worker will find out whether they were hired and the resulting earnings, which equal the relevant wage minus the investment cost (if any). Similarly, each employer will find out whether or not the worker invested and the resulting earnings for the period. The program will keep track of your cumulative earnings for all periods.
- **Market Information:** In addition to their own experiences, all participants will be able to see some market averages. The program will calculate the percentages of Green and Purple workers who invested, averaged over all periods up to the present one. You will also see the overall percentages of Green and Purple workers who were hired. This information will be provided at the start of each new period.
- **Matchings:** After each period, there will be a new random matching of workers and employers, so the person who you are matched with in one round may not be the same person you will be matched with in the subsequent round. Matchings are random, and you are no more likely to be matched with one person than with another.

INSTRUCTIONS SUMMARY

- **Roles:** Half of you are employers and half are workers.
- **Color:** Half of the workers are **Green** in all rounds, and half are **Purple** in all rounds.
- **Matchings:** Employers and workers are randomly matched in each round.
- **Investment Cost:** In each round, workers see their own random investment costs, which are between **\$0.00** and **\$1.00**.
- **Worker Decision:** Each worker decides whether to invest in skills or not.
- **Pre-employment Test:** Each employer sees the worker’s color and two draws (with replacement) from the relevant cup, Investment Cup: BBBRRR, No-investment Cup: BRRRRR.
- **Employer Decision:** The employer decides whether to hire the worker.
- **Worker Earnings:** The worker earns **\$1.50** if not hired or **\$3.00** if hired, minus the investment cost (if any).
- **Employer Earnings:** The employer earns **\$2.00** (if worker not hired), **\$0.00** (worker hired, without investment), **\$4.00** (worker hired, with investment).

NOTES

1. When doing this by hand, the instructor would have to modify the instructions in the Appendix in a straightforward manner. We have used two throws of a 10-sided die to generate a random investment cost. In addition, instructors may use colored marbles, poker chips, or playing cards without distinguishing marks for the test. 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. This process is time consuming, so the instructor would have to work with small numbers of workers and 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.
2. Over all, average investment costs for the greens were actually higher in 8 of the 15 periods.
3. In the first two periods, all employers did not hire green workers with bad test scores. 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.



4. The Veconlab software has an option to set up two treatments, giving the instructor the flexibility to change the cost distributions during an experiment, holding all else constant. The participants were not told anything about the cost distributions other than that all costs would be between \$0.00 and \$1.00.
5. This observation is based on our impressions about classroom discussion; we have no systematic evidence about student attitudes concerning affirmative action.

REFERENCES

- Anderson, D., and M. Haupt. 1999. Employment and statistical discrimination: A hands on experiment. *Journal of Economics* 25 (Winter): 85–102.
- Anderson, L., R. Fryer, and C. A. Holt. 2002. Experiments in discrimination. Working Paper. Williamsburg, VA : William and Mary University.
- Arrow, K. J. 1973. The theory of discrimination. In O. Ashenfelter and A. Rees, eds. *Discrimination in labor markets*. Princeton, NJ: Princeton University Press.
- Coate, S., and G. Loury. 1993. Will affirmative action eliminate negative stereotypes? *American Economic Review* 83 (5): 1220–40.
- Davis, D. 1987. Maximal quality selection and discrimination in employment. *Journal of Economic Behavior and Organization* 8 (1): 97–112.
- Fryer, R. 2002. *Economists' models of discrimination*. Unpublished monograph. The University of Chicago.
- Phelps, E. 1972. The statistical theory of racism and sexism. *American Economic Review* 62 (4): 659–61.