

CLASSROOM GAMES

Experience-Based Discrimination

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

April 2001

Abstract. This paper presents a simple classroom game in which students are randomly designated as either 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 change, 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, pp. 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

* Pennsylvania State University and the American Bar Foundation (Fryer) and the University of Virginia (Goeree and Holt). This work was funded in part by the National Science Foundation (SBR-9818683).

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. An equilibrium exhibits statistical discrimination if workers of one type tend not to invest because they anticipate a lower chance of being hired, 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.¹ 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 workers’ 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.

This exercise is appropriate for microeconomics courses at any level, and for courses in applied areas such as public economics, sociology, anthropology, education and public policy, and labor economics. The simple intuition can be explained in a principals class, whereas a presentation to a graduate class can be more closely tied to the Nash equilibrium calculations. In our experience, the discussion is unusually lively, as everyone can discuss their experiences in the experiment, which

¹ The exercise can also be conducted without internet access, using instructions and handouts available on <http://www.people.virginia.edu/~cah2k/teaching.html>.

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 generating discriminatory patterns in the data. The instructor begins by running the setup program from any Javascript-enabled browser connected to the internet, by going to: http://128.143.144.247/sd/setup_experiment.php. This page 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 “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 15 periods in about 40 minutes, which leaves time for discussion. The instructor then must input the name of the database table that will be created. It is best to have the table name contain unique information (like the instructor’s initials), since a table cannot be created if one by that name already exists. After the instructor presses the submit button at the bottom of the setup page, there should be a message that the database has been created, and the instructor can press the submit button again to go to the page where experimental results will be displayed. The final step in the setup is to get the students to login from PCs connected to any Javascript-enabled browser via the site: http://128.143.144.247/sd/login_page.php. After logging in, the students will see their “id” numbers and role assignments (employer, green worker or purple worker) at the top of the first instructions page. Although these instructions are self paced, the instructions go more quickly if you read them aloud so that students go through them at

the same time. This practice also serves to remind the instructor of the details that may be important in the class discussion of experiment results. Since the program will only permit the specified number of players to log in, you may simply read the instructions while looking over the shoulder of one of the participants.

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 invest or not, a test result (described below) for that worker 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 wait windows vanish and the period-specific and cumulative earnings results are displayed privately for both the worker and the employer.

Recall, employers will not be able to see whether or not a worker has invested, and so a pre-employment test (e.g. an interview) will be given. In order to avoid using probabilistic terminology, the 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 made.² 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 investments.

Workers' earnings are such that they prefer the managerial job, which pays \$3.00, to the regular job, which pays \$1.50, regardless of whether they have incurred the investment cost (which is subtracted from the wage). Employers' earnings are such 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, but the employer earns \$4.00 if a worker assigned to a managerial job invested, and \$0.00 if the worker did not invest. It is not necessary to actually 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.

Some “historical” information is provided in subsequent periods. Although worker's 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. This information may be used by workers and employers 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 worker's color and the test result that is correlated with the unobserved investment decision, prior to making the job

² 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.

assignment decisions that determine wages and employer earnings.

Classroom Discussion

As with most experiments, it is impossible to predict in advance exactly what will happen. We have found that students tend to notice discrimination even when the patterns are fairly subtle in the data. A fairly strong pattern was observed in an experiment we conducted in an upper level class at the University of Virginia. In this class, a pattern of 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 3.³ 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. Employers assigned every worker, regardless of color, to the managerial job when they received a "good" test score. 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 workers.⁴ 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.

³ Over all , average investment costs for the Greens were actually higher in 8 of the 15 periods.

⁴ 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.

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. We approach the discussion in a Socratic method. By engaging the students, one can get a feel for how they perceived the game, what strategies they used, etc. We began by asking for a Purple worker to say 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 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.” The 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.*”

In our experience, students are quick to spot perceived patterns of discrimination, even when

they are much less obvious. In another experiment, the investment rate for Green workers was about 15 percent *below* that for Purple workers, and a Green worker remarked that “Greens were stupid; they should have invested a lot in early periods to raise the confidence of employers. When I realized that other Greens weren’t doing this, I stopped investing too.”

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) “probational” 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 courses.⁵ 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 “race-based” thinking on the part of employers.

Further Reading

The environment described above encompasses 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

⁵To illustrate a discriminatory equilibria, let π_p (resp. π_g) denote the employer's prior belief about purple (resp. green) workers. After observing the worker's pre-employment test, the employer calculates the posterior probability ($\xi(\pi_i, \theta)$) that the worker is qualified, where

$$\xi(\pi_i, \theta) = \frac{\pi_i \text{prob}(\theta | \text{qualified})}{\pi_i \text{prob}(\theta | \text{qualified}) + (1 - \pi_i) \text{prob}(\theta | \text{unqualified})},$$

$i = b, w$ and $\theta \in \{RR, RB, BR, BB\}$

Suppose $\pi_p < \pi_g$ and the employment standard (t) is (RR) for purples and (BB) for greens.

Expecting these standards, green workers only invest if $c < \frac{2}{3}$ and purples invest if $c < 0$.

Thus, in equilibrium, $\pi_g = G(\frac{2}{3}) > \pi_p = G(0)$, which confirms the employers initial beliefs.

worker-specific investment costs. See Fryer (2001) 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 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.

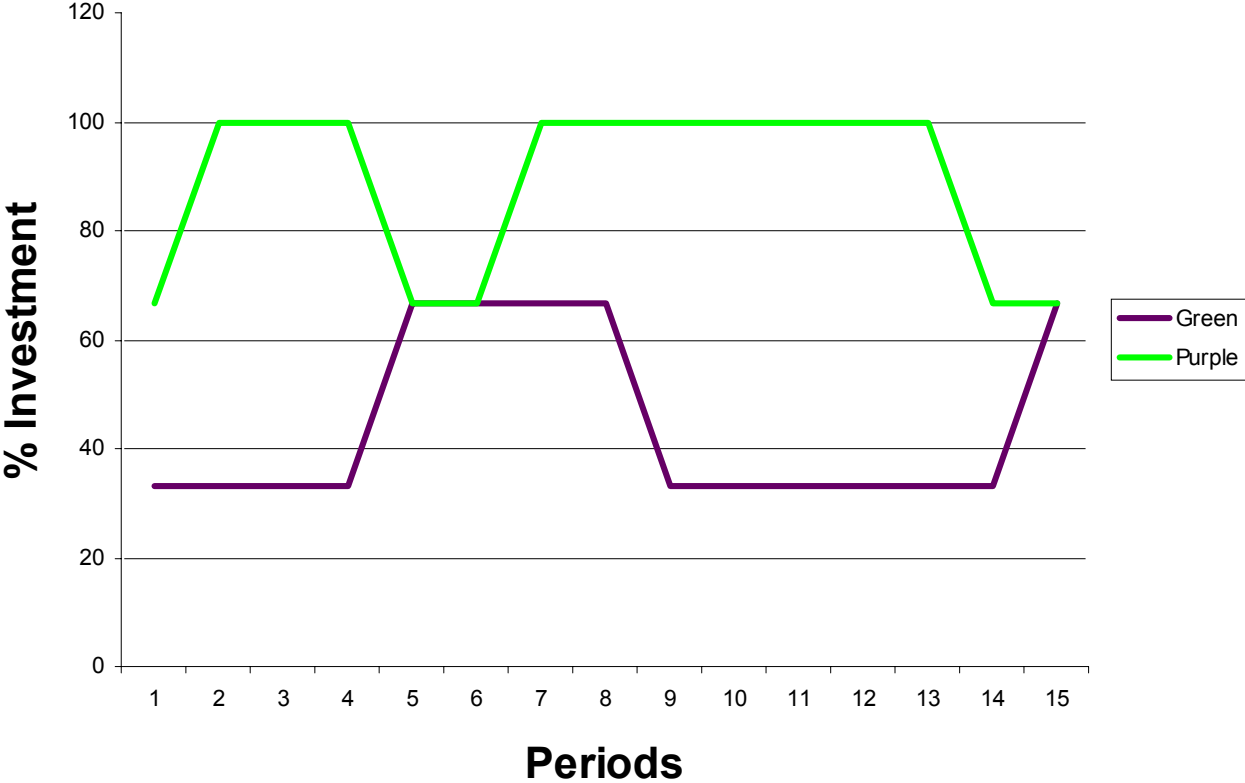
REFERENCES

- 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 Manuscript. The University of Chicago.


Lundberg, S., and Startz, R., (1983) "Private Discrimination and Social Intervention in Competitive Markets" *The American Economic Review* V. 73 no. 3 p. 340-47

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

Average Investment Rates



Average Managerial Assignment Rate



Scraps (instructions and tables to go on the web)

INSTRUCTIONS

ROLE ASSIGNMENTS

In this exercise, you are either a worker or an employer; your role is indicated at the top of the record sheet that is attached to these instructions. Workers decide whether to make a costly investment in training, and employers decide workers' job assignments, after observing test results and other information. We have already randomly assigned a color identification to each worker, which is either Purple or Green, and an ID number, as indicated on your record sheet. The class has been divided into equal numbers of workers and employers, and the number of Purple workers is equal to the number of Green workers. The experiment will consist of a number of "periods," and you will have the same role, Purple worker, Green worker, or employer, in all periods. All of you have ID numbers written on the top of your record sheets, and we will use these numbers to match each worker with an employer in each period.

WORKER'S INVESTMENT DECISIONS

At the start of a period, all workers must decide whether or not to spend money and purchase some training. The cost of this "investment" decision is randomly determined for each worker by two throws of a ten-sided die. Here I have such a die, with sides marked 0, 1, ... 9. One of us will come to each worker's desk and throw the die twice, with the first throw being the "tens" digit and the second throw being the "ones" digit. In this manner, two throws of this ten-sided die will produce a cost number between 0 and 99 pennies, with any integer in this interval being equally likely. After seeing this investment cost, each worker should record it in the second column of their record sheet and decide whether or not to invest, which is recorded in the third column as an I for "Invest" and N for "Not Invest". If you do not invest, you do not incur any cost, and if you do invest, this cost is deducted from the wage you receive from the employer. Investment is costly but it may help you do better on a pre-employment test administered by the employer with whom you are matched, as explained below.

THE PRE-EMPLOYMENT TEST

Employers will not be able to see whether or not a worker has invested, and so a test will be given. The test will have either a good outcome "Blue" or a bad outcome "Red," but it is not perfect. Nevertheless, the bad outcome will be more likely for a worker who decides not to invest. The test will be conducted by drawing colored marbles from a cup, and the cup used by those who invest has a higher fraction of Blue marbles. In particular, there are two cups:

No-investment cup: 5 Red marbles and 1 Blue marble

Investment cup: 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. You might think of it this way. 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. (In fact, we will enter the contents of the relevant cup into a common "Draw Cup" that is used for all draws.) Then we will draw one marble, show it to the employer, put the marble back into the cup and shake it before drawing again. Thus there will be two independent draws from the same cup "with replacement," so the contents of the cup do not change. 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.

HIRING DECISIONS AND EARNINGS

Each employer will be randomly matched with a worker in each period. Here I have two buckets. There are ___ balls in the workers' bucket, each one marked with the ID number of one of the Purple or Green workers, and similarly there are ___ balls in the employers' bucket. The first balls drawn from each bucket will determine the first worker-employer match, and so forth.

Each worker will find out their investment cost for the period and will decide whether or not to invest. After investment decisions are made, the relevant cup (investment or no-investment) will be used to draw two marbles, with replacement, and these test results will be shown to the employer with whom the worker is matched in that period. The employer then decides whether to assign the worker to the managerial job, which pays \$3.00, or to the regular job which pays \$1.50 per period. The employer's profit (after the wage has been paid) depends on the job assignment and whether or not the worker invested:

EMPLOYER'S EARNINGS

\$4.00 if the employer assigns a managerial job to a worker who invested,

\$0.00 if the employer assigns a managerial job to a worker who did not invest,

\$2.00 if the employer assigns a regular job, whether or not the worker invested.

Thus a worker's investment in training only benefits the employer if the worker gets assigned to a managerial job. The worker's earnings depend on the job assignment and the investment cost, which is between 0 and 99 pennies

WORKER'S EARNINGS

\$3.00 - investment cost if the worker invests and gets a managerial job,

\$1.50 - investment cost if the worker invests and gets a regular job,

\$1.50 if the worker who does not invest and gets a regular job

INFORMATION FEEDBACK

Employers never see the ID number of the worker with whom they are matched nor do they know the past history of investment decisions of any individual worker. Likewise, workers do not know the ID number of the employer with whom they are matched or the past job assignment decisions

of any individual employer. However, after each period is finished we will record the average investment and job assignment decisions on the blackboard. In other words, after each period we will list the percentage of Green workers that invested in a training, the percentage of Purple workers that invested in a training, the percentage of Green workers that was assigned a managerial job, and the percentage of Purple workers that was assigned to a managerial job.

EARNINGS RECORD

The 10-sided die will be thrown twice for each worker at the start of every period. After seeing the investment cost for the period, the worker should record it in the second column, just to the right of the period number. The worker can use the information written on the blackboard about previous job assignments for Green and Purple workers. Then the worker decides whether or not to invest, which is indicated by writing I (invest) or N (not invest) in the third column of their record sheet. We will collect the record sheets and match workers with employers, by drawing numbered ping pong balls. Then we go to the employers' desk with the relevant cup (investment or non-investment) to draw two marbles that determine the test result for that worker. The employer sees the worker's color and the test result, which we will mark in the second and third columns of their record sheet. In addition the employer can use the information written on the blackboard about the average investment decisions made by Green and Purple workers in previous periods. With this information, but without knowing whether the worker has invested or not, the employer makes the job assignment decision: regular (R) or managerial (M), recorded in the fourth column. Finally, we calculate earnings and record results for all of you, and then we will return your Record Sheets. This process will be repeated for a sequence of periods. Finally, note that the employer never sees the worker's ID number or past history of investments, so the job assignment is made only on the basis of the test result and color, Purple or Green. There is an equal number of workers of each color. In each period, the random matching of workers and employers will be done again with draws of numbered ping pong balls, so it is unlikely that you will be matched with the same person in consecutive periods.

All periods will have the same structure, except that some of the worker investment decisions are "forced" in some of the periods. Each worker can look in the third column and see the pre-decided decisions (if any) for specific periods. If your decision is forced, we will come to your desk to throw the dice that determine your investment cost, although it will not be relevant for your earnings if your decision is forced to be N for that period.

Are there any questions?

TABLES

Worker Record Sheet ID: _____ Color: _____

Period	Investment Cost	Investment Decision (I for Invest, N for Not Invest)	Test Result	Job Assignment (Regular or Managerial)	Earnings
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					

Employer's Record Sheet ID: _____

Period	Worker's Color (P for Purple, G for Green)	Test Result (B for Blue, R for Red)	Job Assignment (R for Regular, M for Managerial)	Worker's Investment Decision (not known at time of job assignment)	Earnings
1					
2					
3					
4					
5					
6					
7					
8					
9					
10					