

Discrimination: Evidence from Psychology and Economics

Experiments

Lisa R. Anderson, Roland G. Fryer, Jr. and Charles A. Holt*

Prepared for the *Handbook on Economics of Discrimination*, William Rogers, editor

March 2005

I. Introduction

Measuring the intensity and impact of overt racism and discrimination has been the subject of much debate over the 20th century. Few disagree that historical discrimination and other forms of social ostracism partly explain current disparities on a myriad of economic, social, and health related outcomes. Disagreement arises when ferreting out the underlying reasons behind the discriminatory treatment.

The difficulty in uncovering mechanisms behind discriminatory actions lies in the difficulty of measuring attitudes about race, gender, and other characteristics that often serve as a basis for differential treatment. Therefore, laboratory experiments have been particularly useful in the study of discrimination under conditions where experience, perceived status, and group identity can be partially measured and controlled. For example, cleverly designed experiments allow one to distinguish the effects of underlying biases in preferences for one's in-group from the effects of information-based forms of discrimination (statistical profiling, social categorization, and so on). This paper surveys laboratory studies of discrimination in psychology and economics.

* College of William and Mary, Harvard University Society of Fellows and NBER, and University of Virginia, respectively. This research was supported in part by a grant from the National Science Foundation SBR 0094800 and the University of Virginia Bankard Fund.

There is a long tradition of experimental studies in psychology that examine the effects of observed characteristics, like status or group identity, on the way subjects treat others. The division into groupings can be based on survey responses or on observed traits like eye color. If deception is permitted, then the groupings can be random, even though people are told that the groupings are determined by some task or questionnaire. The goal of such manipulations is to determine the extent to which people with high status or of one's own group are treated differently in exercises like a money division task.

Laboratory experiments in economics also reveal that status affects behavior in some market contexts. A number of economics experiments are motivated by formal equilibrium models in which discrimination arises from self-fulfilling expectations. For example, if workers in one group anticipate being discriminated against, they will be less likely to invest in acquiring skills and, as a result, employers will observe systematic differences in investment decisions. Feed-back effects can cause discrimination to become entrenched, as noted by Tajfel (1970): "For example, economic or social competition can lead to discriminatory behavior; that can then in a number of ways create attitudes of prejudice; those attitudes can in turn lead to new forms of discriminatory behavior that create new economic or social disparities, and so the vicious circle is continued." These discriminatory equilibria may persist even when the two populations are *ex ante* identical (e.g. Arrow, 1973; Coate and Loury, 1993), and this theoretical possibility can be investigated in the laboratory.

Experiments can also be effective in classroom settings, since the relatively neutral context and commonly shared classroom experience may allow for a more objective discussion of otherwise sensitive issues. In fact, some of the earliest discrimination experiments were done in classroom settings. One example is the case, discussed below, where a grade school teacher

based preferential treatment on eye color.

II. Early Discrimination Experiments

In response to the 1968 assassination of Martin Luther King, a third grade teacher named Jane Elliott devised a simple classroom exercise to facilitate discussion of discrimination.¹ Students were divided into two groups based on eye color, and it was announced that brown-eyed people would be superior to blue-eyed people that day, and that the roles would be reversed the following day. Ms. Elliott further explained: “what I mean is that brown-eyed people are better than blue-eyed people. They are cleaner than blue-eyed people. They are more civilized than blue-eyed people. And they are smarter than blue-eyed people.” The brown eyed children got to sit in the front of the room, to go to lunch first, and to have more time at recess. Blue-eyed students slumped in their chairs, as though they accepted their inferior positions. These behavioral differences were reversed when the roles reversed the next day.

Psychologists have conducted similar experiments in more controlled settings. The focus is often on the effect of group affiliation towards members of one’s own group and members of other groups. The group divisions are often selected in the laboratory on the basis of a seemingly “objective” criterion like whether the length of a line was under-estimated. In reality, the division is random. Once people are grouped in this manner, they are asked to split money between two other people, only one of whom is in their group. For example, Vaughn, Tajfel and Williams (1981) divided 7 and 11 year old children into “red” and “blue” groups based on their preferences for a set of paintings. Neither group was described as being superior to the other. The children were then asked to divide “some pennies” between other people in their class. All

they were told about the others was their group identity (red or blue). They were also told that they would get money from other people making similar decisions. (In reality, the money promised to the children was later used instead to have a class party.) The kids in both age categories consistently gave more money to members of their own group. This “in-group bias” persisted even when the participants were told that the group membership was randomly determined (Billig and Tajfel, 1973).

The effect of group affiliation is less clear when additional focal points like self-interest or attitude similarity are incorporated into this design. In particular, Turner (1978) told a group of 14-16 year old boys that their preferences for paintings would be used to determine groupings. The boys were asked to divide money between themselves and an unidentified other person in the class. The other person’s group affiliation had no effect in this context, as self-interested behavior dominated. In a similar study, Diehl (1988) had 13-15 year old high school students perform two classification tasks. First, the students had to estimate the lengths of lines drawn on paper and were told that they were separated into groups of over-estimators or under-estimators. Second, they completed an attitude questionnaire with the understanding that it could be used to compare their attitudes on a variety of issues with those of others in the group. Subjects were told that they were grouped based on their line length estimation and their responses to the attitude questionnaire, but groupings were actually randomly determined. Finally, subjects were asked to divide money between two other people, who were identified by group (over- or under-estimator) and by attitude classification relative to the subject making the allocation decision. Group affiliation only had a strong effect when it was consistent with the attitude similarity classification. In-group members who had been designated as having similar attitudes were

¹ This exercise is described in *A Class Divided* (William Peters, 1971).

awarded more than other in-group members, and were awarded more than members from the other group with similar attitude designations. However, the effect of attitude similarity dominated group affiliation when they conflicted; people with similar attitudes in a different group were given more money than people with dissimilar attitudes in a subject's own group.

Psychologists have also studied in-group bias in situations where there is status associated with group affiliation. Turner and Brown (1978) classified undergraduate students as either Arts or Sciences based on their major course of study. They had them meet in groups of three for a 20 minute discussion of the following statement: "No individual is justified in committing suicide." At the end of the discussion, one of the three made a tape-recorded summary of their views. They were told that the purpose of the discussion was to evaluate their "reasoning skills." Once the tape was recorded, they were asked to comment on how well they did relative to another group. The same "other group" tape was played for all of the subjects. Arts students were told that the comparison group was from Sciences, and vice versa. In the status treatment of this study, the experimenter singled out one group (Arts or Sciences) as having better reasoning skills than the other group. Therefore, in the status treatment subjects always compared themselves to people with different status. The authors concluded that all subjects were biased in favor of their own group, and that groups identified as superior were more biased in favor of their own group.

Klein and Azzi (2001) replicated this finding in a different environment. They had college students take a trivia quiz, which the students were told would be used to divide people into groups. Actually, all participants were put in the same group, which was announced to be the superior group in one treatment and to be the inferior group in the other. Then participants had to rate the creativity of sentences written by fictitious "other students." Before making the

ratings, they were told that the other students would receive a reward for a high score and that there was no demonstrated relationship between creativity and scores on the trivia quiz. In this treatment, both inferior and superior groups gave higher scores to people in their own group.

In summary, early discrimination experiments generally revealed that subjects were biased in favor of groups designated as being superior and/or similar in some dimension. These results were based on experiments with hypothetical incentives. This raises this issue of whether such behavioral patterns would persist in market experiments where participants are also concerned with the financial consequences of their actions. But before we review economics experiments with financial incentives, we will summarize a relevant psychology literature on categorization and recognition.

III. Psychological Experiments

Social psychologists have amassed an impressive literature on experimental approaches to measure discrimination, perceptions, and implicit bias in associations. Fundamental to their approach is the important notion of social categorization. As the distinguished social psychologist Gordon Allport (1954) noted, “the human mind must think with the aid of categories. We cannot possibly avoid this process. Orderly living depends upon it.” To this, most (if not all) psychologists would agree. More importantly, there is a long tradition in social psychology that treats discrimination, stereotyping and prejudice as inevitable consequences of social categorization (see Allport (1954), Hamilton (1981), Tajfel (1969), or Fiske (1998) for a recent review).²

In demonstrating the effect and automaticity of categorical cognition, Devine (1989)

draws a sharp distinction between stereotypes and prejudice. One hundred and twenty nine students enrolled in an introduction to psychology course at Ohio State University participated in the experiment for course credit. Participants took the Modern Racism Scale to determine their prejudice level. Participants were then shown subliminal images (appearing on a computer screen for less than 30 milliseconds) of words associated with the social category “black” (i.e. black, poor, ghetto, negroes, and so on), using the stimuli priming method developed by Bargh and Pietromonaco (1982). Subjects were told that the experimenter was interested in how people form impressions. They were asked to read the famous “Donald paragraph” – a twelve sentence paragraph that has Donald engaging in ambiguously hostile behaviors (i.e. withholding rent until an apartment is painted, demanding money back at a retail store, etc.). The results were startling: participants who were given the subliminal images rated Donald as significantly more hostile, and this was true for all prejudice levels. This experiment demonstrates the power of implicit associations and one of the clever ways that such associations have been measured.

Categorization and stereotyping also manifest themselves in other ways. There is growing literature in psychology on racial and ethnic differences in facial recognition. The terms “cross-race recognition deficit,” “cross-race effect,” and “own-race bias” all describe the frequently observed performance deficit of one ethnic group in recognizing faces of another ethnic group compared with faces of one’s own group (see Sporer 2001 for a detailed review). In other words, “they all look alike to me” is a reasonable caricature of how members of one group categorize another.

Own-race bias in the recognition of facial stimuli is observed when two factors, ethnic group of participant and ethnic group of stimulus face, interact. This effect is significant in the

² There are two models of categorization in economics, see Mullainathan (2001) and Fryer and Jackson (2003)

expected direction: both ethnic groups recognize members of their own better than members of other ethnic groups.³ Models of categorization predict that individuals with more inter-group contact will be better at distinguishing subtle features about other groups than individuals with less inter-group contact. There is substantial evidence in this regard (see Sporer 2001, Table 2). As Meissner and Brigham (2001) report: “Several studies demonstrate that adolescents and children living in integrated neighborhoods are better at recognizing novel other-race faces those living in segregated neighborhoods.”

An interesting experiment testing the relationship between contact with other groups and facial recognition is Li, Dunning, and Malpass (1998). They demonstrate that white “basketball fans” were superior to white “basketball novices” in recognizing black faces. The idea is that basketball fans watch the National Basketball Association games on a regular basis, which provides frequent exposure to black faces, given that a sizeable majority of the players are black. Participants were black and white men and women. They were presented with black and white faces on a video monitor. The subjects were informed that they would be tested on their ability to recognize the faces viewed. Performance of basketball fans that were white was indistinguishable from blacks in their ability to recognize black faces, whereas the white subjects who were not basketball fans performed at a significantly worse level. In recognizing white faces, there was no difference between basketball fans and novices.

In summary, there is a long tradition in social psychology of quantifying the prevalence and impact of social categorization on prejudice, stereotyping, and discrimination. The

³ Meissner and Brigham (2001) provide a detailed meta-study of the last thirty years of literature investigating the own-race bias in facial recognition. They review 39 articles involving the responses of over 5,000 subjects. There are a few studies that fail to find a cross-race effect. The overwhelming consensus among social psychologists, however, is that these effects not only exist, but are quite large (Meissner and Brigham 2001).

experimental literature has demonstrated the automaticity of categorical processing of information and the resulting biases in decision making. One depressing interpretation of this literature is that “people will be prejudices so long as they continue to think” (Billig(1985).

IV. Economics Experiments

Some economics experiments follow directly from the psychology literature in the sense that status or group identification is induced by laboratory manipulations such as award ceremonies (Ball et. al., 2001). Other experimental economics studies on discrimination are motivated by the large literature on employment discrimination and information economics. The main theories that have been tested experimentally include statistical discrimination (Anderson and Hauptert, 1999; Davis, 1987; Fryer, Goeree, and Holt, 2001 and 2005), asymmetric pair-wise tournaments (Schotter and Weigelt, 1992), and price-preference auctions (Corns and Schotter, 1999). In what follows, we provide a description of these theories and a review of the related experimental papers.

Status and Group Identification

Ball et al. (2001) investigated the impact of status on behavior in a market setting using undergraduates. Demand was induced for buyers by giving them a “redemption value” for a good that could be acquired through trade. Buyers earned the difference between the redemption value and what they paid for the good. Similarly, sellers were told the cost of supplying the good and that they would earn the difference between the cost and price for which the good was sold. The number of buyers’ units was equal to the number of sellers’ units. Thus there was a vertical overlap of the supply and demand curves, with a range of market-clearing prices between the

sellers' costs and the buyers' values. This overlap makes price indeterminate, which leaves more of a chance of observing the effects of non-economic factors like fairness. Trading took place via a double auction, where sellers could call out "ask" prices and buyers could call out "bid" prices. The bid prices tended to increase as buyers out-bid each other, and ask prices tended to decline as sellers under-cut each other. In this sense, there were two auctions at the same time, with prices rising in one and falling in the other. A trade occurred when these two processes met, i.e. when a buyer accepted a seller's ask or a seller accepted a buyer's bid.

Two treatments were used to introduce status in this market context. In the "earned" status treatment, subjects took a trivia quiz and were told that getting a high score on the quiz earned them a gold star. In the random status treatment, subjects observed as people were randomly picked to receive stars. In both cases, stars were actually awarded randomly and were distributed in a ceremony where non-star people were told to applaud the star people, who were moved to a special seating area. People with stars were buyers in some sessions and sellers in others. Those with stars earned significantly more of the available surplus, regardless of whether they were buyers or sellers or whether the status was earned or not. In addition, males earned more than females.

Ball and Eckel (1996 and 1998) used the same method to introduce status (both "earned" and random) into a bargaining situation known as the ultimatum game. After participants were paired, one person in each pair was chosen to propose a split of money or candy, and the other person was given the chance to accept or reject the proposed split. An acceptance finalized the split, and a rejection resulted in zero earnings for both. When college students were asked to split Hershey's kisses with an anonymous partner, both star and non-star proposers offered more kisses to respondents with stars. When students had to divide a \$10 prize, there was no

significant difference in offers made to responders with or without stars. Thus we see that the role of incentives is potentially important in mitigating the effects of discrimination.

Fershtman and Gneezy (2001) also used bargaining games to evaluate the effects of non-economic factors. Instead of inducing group membership in the laboratory, they deliberately recruited people from different ethnic backgrounds. All subjects were Jewish Israeli undergraduates with typical ethnic last names, which were revealed in the experiment so subjects could identify other participants as being either of Ashkenazic or of “Eastern” origin. They used a “trust game” in which one person has the option to pass all or part of a money endowment to the other. The money passed is increased by a pre-announced proportion. The responder then decides how much (if any) of this augmented sum to keep and how much to return to the original sender. Trust is measured by the proportion of the endowment that is originally sent to the responder. The authors report a systematic mistrust of men of Eastern origin. There are two possible explanations for this behavior: Either people have a preference for lower earnings for the members of this group or the senders fear that the men of Eastern origin will not reciprocate by returning some of the money. These two hypotheses were evaluated using a “dictator” game where one person simply decides how much of a fixed endowment of money to keep and how much to give to another person (without augmentation or opportunities for reciprocity). Dictated divisions were not systematically affected by the recipient’s ethnic background. The authors concluded that behavior in the trust game was driven by the fear that generosity would not be reciprocated by men of Eastern origin.

In conclusion, the effect of status on behavior in economics experiments depends on the origin of the status and the nature of the incentives. In the Ball et al. (2001) study with a range of equilibrium prices and real financial incentives, price moved in the direction favored by the high status group independent of whether status was perceived to be earned or randomly awarded. In

contrast, Ball and Eckel (1996 and 1998) report that assigning status to one group has no effect on bargaining outcomes in ultimatum games with real financial incentives, although some status effects were observed in experiments involving the division of Hershey's kisses. Finally, Fershtman and Gneezy (2001) report no discrimination based on ethnicity in a dictator game, but they find that in a trust game, lower amounts of money were passed to individuals in a particular ethnic group. They conclude that this discrimination is based on mistrust rather than a desire to lower the earnings for those individuals. In both the market experiments and the ultimatum games, the presence of discrimination affects the distribution of earnings but has no effect on overall welfare. However, the lower amount of money passed to particular individuals in the trust game generates a welfare loss, since passed money is tripled.

Statistical Discrimination

The theory of statistical discrimination has become a valuable tool in the study of many labor market phenomena. Kenneth Arrow (1973) and Edmund Phelps (1972) developed the theory independently. The basic framework relies on the fact that employers do not perfectly observe investments in human capital. For simplicity, it is assumed that workers who invest are "qualified" and those who do not are "unqualified." Conditional on the worker's investment decision, employers observe a noisy signal of the worker's qualification level (i.e. an interview or a pre-employment test). Finally, employers decide whether or not to hire the worker on the basis of the signal and other characteristics like race or gender.

Within this framework, Phelps (1972) assumes that the signal emitted by minorities is "noisier" than that of non-minorities. It follows directly from this assumption that minorities who emit low signals are paid a wage above their majority counterparts, and minorities with

relatively high signals are paid below their majority counterparts. In Phelps's model, however, there need not be any discrimination "on average." The assumption that it is harder to evaluate the qualifications of some ethnic groups has been questioned. Recognizing this, Arrow (1973) provides an alternative model in which some worker characteristics are endogenous, and an employer's *a priori* beliefs can be self-confirming.

To see this, consider two groups, A's and B's. Now suppose that an employer has a prior belief that B's are less likely on average to invest in pre-market human capital relative to A's. The signaling technology is imperfect, but qualified workers are more likely to emit a higher signal (e.g. pass a test or make a good impression in an interview). These models typically have an equilibrium in which the employer hires workers with signals that exceed a threshold level that can depend on the worker's group. Since the employer is relatively pessimistic about B's, the threshold for these workers is higher than that for A's. This affects the worker's investment decision: B workers (who are held to a more exacting standard) have less incentive to invest relative to A workers (who are held to a more forgiving standard). This behavior by workers confirms the employer's initial asymmetric beliefs that B workers are less likely to invest. The beauty of Arrow's theory is that the employer's biased initial beliefs are confirmed in equilibrium, even though the populations were *ex ante* identical.

One issue that arises with statistical discrimination experiments is how these biased perceptions are generated in the laboratory. Obviously, different perceptions may arise from learning and past experiences. For example, past discrimination might limit investment opportunities for one group (Davis, 1987). From an experimental perspective, however, it is also interesting to use two populations with identical ability distributions, since the emergence of discrimination under such conditions would be especially noteworthy.

The earliest economics experiment on this topic is Davis (1987), who studied the effects of the relative sizes of “majority” and “minority” populations. The intuition behind the experiment is that if more sample observations are drawn from the majority population, then this population is more likely to generate a higher *maximum* observation. If employers tend to focus on the maximal draw from each population, then this could result in an employer bias in favor of the larger group. It is not implausible that the best candidates would be more likely to be remembered, since many job searches involve narrowing consideration from a large number of applicants to a final short list. In the baseline treatment, subjects saw random realizations drawn from identical normal distributions of monetary prize values, with about 80 percent of the draws coming from the “majority population.” In the final period, subjects were free to decide what proportion of draws would come from each population, so that a bias away from equal numbers of draws from each population could be taken as evidence that one population is perceived as being better on average. Even though the two distributions or draws were identical in this baseline treatment, subjects selected about 60 percent of the draws from the population that was previously sampled more extensively. This effect was characterized as being “weak,” and was only significant at about a 10 percent level. A somewhat heavy handed treatment actually provided subjects with a “tab” sheet listing the maximum draw from each population for each prior period, and this information seemed to have an effect, raising the percentage of final-period majority population draws to about 70 percent, a significant increase. This study is interesting in that it suggests a mechanism whereby a bias might arise, even when the two populations are identical. As Davis notes, such a bias would be even stronger in the presence of some underlying inequality.

As was the case in the psychology literature on status and group effects, a number of economics experiments were developed to stimulate class discussion. This approach is especially

effective, since participants often come to realize that they are discriminating on the basis of prior experience or statistical knowledge rather on the basis of a personal bias against one type of person. For example, Anderson and Hauptert (1999) used colors (“green” and “yellow”) to identify two types of workers in a classroom exercise. Workers were represented by green or yellow index cards, with productivity numbers written on the back of each card. Subjects played the role of employers who were required to hire a specified number of workers, with some incentive to pick workers with the highest productivities. The participants knew the distributions of productivities for each color, but had to pay an “interview cost” in order to observe the productivity on a specific card. A stack of 20 cards, 10 of each color, was shuffled and presented in sequence, with the requirement that 8 workers be hired. For each card, the employer decided whether to interview (pay to observe the productivity), but the decision of whether to hire that worker could be delayed until all interview decisions had been made. In markets where the average productivity was lower for one color, employers tended to hire less of that color. The explanation is that, in the absence of an interview, the employer tends to rely on the population average, which is a type of “statistical discrimination.” Discrimination against the less productive group of workers was somewhat diminished when the interview cost was reduced, since this allowed employers to search for the most productive workers, regardless of color.

The experiments reported in both Davis (1987) and Anderson and Hauptert (1999) have the common feature that differences between the two types of workers are exogenous. As noted above, much of the theoretical literature on statistical discrimination pertains to models in which inter-group differences are endogenously determined by workers’ investment decisions. Such models are of interest because of the possibility that systematic productivity differences may arise even when the two groups are *ex ante* identical. These situations may persist in equilibrium if employers come to expect that members of one group are less likely to invest in skills, and hence tend to offer less

attractive job assignments to members of that group. The flip side of this story is that workers from the “disadvantaged” group anticipate reduced job assignment opportunities and hence tend not to invest, which in turn tends to confirm employer expectations.

The experiments reported in Fryer, Goeree, and Holt (2001) were conducted in a setting with endogenously determined worker productivities. Half of the workers were randomly designated as being Purple and the other half were Green. Each worker began a round by observing a randomly determined investment cost. Then workers decided whether or not to invest. The employer observed a test outcome (red or blue), with blue being more likely when the worker invested, as in the Coate and Loury (1993) model. Finally, the employer decided whether or not to hire a worker knowing the worker’s color and the test score, but not the investment decision.

One of the treatments for this experiment involved having the investment cost draws for the two types of workers come from different distributions for the first ten rounds and then removing this asymmetry for the final 50 periods. This initial asymmetry was not announced, since subjects were only told that the costs would be “randomly determined” amounts between \$0.00 and \$1.00. In fact, the Green workers were drawing from a uniform distribution on [\$0.00, \$0.50], and the Purple workers were drawing from a uniform distribution on [\$0.50, \$1.00]. After round 10, all draws were from a uniform distribution on [\$0.00, \$1.00]. Greens invested more and were hired more often than Purples in the first 10 periods in all sessions, but what happened next was sometimes quite interesting and surprising. The next several investment cost draws would tend to look relatively attractive to the Purples and relatively unattractive to the Greens, so the investment rates surged for the initially disadvantaged Purples and fell for the Greens. In one session, this caused a crossover effect where the Purples invested more often than Greens and were hired more often for the remaining periods. This reversal of the originally induced inequity also occurred in a classroom

experiment (Fryer, Goeree and Holt, 2005).

The initial asymmetry did produce a lasting effect for the session shown in Figure 1. The period is shown on the horizontal axis, and the switch to symmetric costs is indicated by the vertical dashed line at period 10. The surge in investments by Purples is immediate, as indicated by the dashed line in the left panel of the figure. This increase in the tendency for Purples to invest is not reflected in hiring rates until about 10 periods later, since employers are not able to observe the investment decisions of workers who are not hired. Even after employers begin to hire Purple workers more often (around round 20), the investment and hiring rates for Greens remain higher. An analysis of the individual decision sequences indicates that four of the six employers used a color-based strategy. Two of the employers in this session tended to hire Greens with bad signals, but not to hire Purples. Two other employers hired Greens but not Purples when the signals were a mix of good and bad elements. The remaining two employers did not appear to use color in making their decisions. While the data patterns do not conform closely to any theoretical prediction, they are qualitatively similar to the predictions of the Coate and Loury model, with biases in employer hiring that have a feedback effect on investment decisions, causing Purples to invest less often than Greens. The differences in investment tendencies have a larger effect on the hiring decisions.

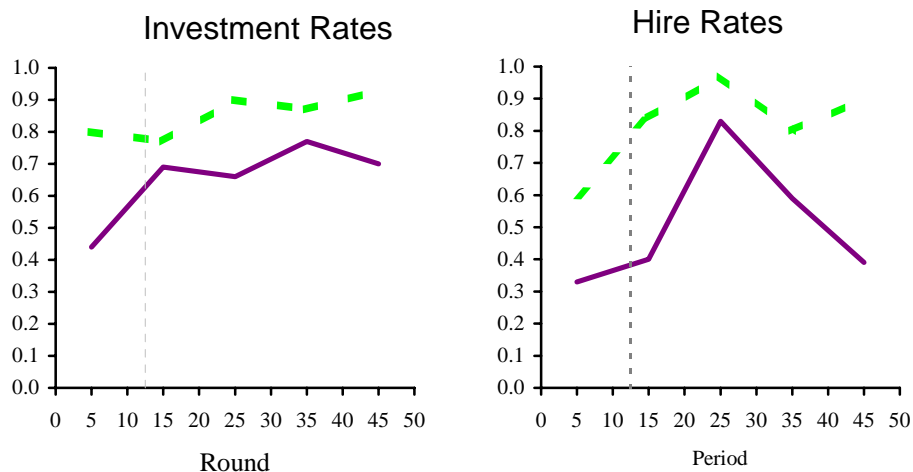


Figure 1. Ten-Period Average Investment and Hire Rates for Green (Dashed Line) or Purple (Solid Line) with the End of the Initial Cost Asymmetry Shown by the Dashed Vertical Lines. (Fryer, Goeree, and Holt, 2005)

In summary, statistical discrimination can occur both in situations where workers types are exogenously determined and in situations where workers make their own productivity investment decisions. In Davis (1987) the exogenous difference was the group size, which afforded employers more “draws” from the majority population. As a result, the maximum draw tended to be higher for the majority population, which caused subjects to make more draws from the majority population when they were able to decide which workers to hire. In the Anderson and Hauptert (1999) classroom experiment, the exogenous difference was the productivity distribution for each population. As expected, students hired more workers from the high productivity group, although this discrimination was mitigated by a reduction in the “interview cost” of finding out a worker’s productivity *ex ante*. When asymmetries arise endogenously, discrimination can also persist as a result of a self-confirming cycle of low employer expectations and low employee aspirations, which results in low rates of investment in human capital for workers in a particular group. The low

investment rates for some workers generates a reduction in overall welfare.

Tournament Theory and Price-Preference Auctions

Often licenses and government contracts are allocated on the basis of an auction or contest. In such cases, the public officials involved may wish to promote the participation of a particular group (e.g. small business or minority owned businesses) that would be underrepresented, perhaps because of size of past discrimination. This raises the issue of how affirmative action policies affect outcomes in auctions and tournaments.

In a simple pair-wise tournament, two people compete for a prize of known value. Each person knows their own private effort cost, which is randomly determined. The agent who exerts the most effort wins the prize. Alternatively, asymmetry can be introduced in the value of the prize, with constant effort costs across participants. Myerson (1981) shows that a seller may maximize revenue by subsidizing some high cost agents, thereby increasing competition. This is in conflict with most economic intuition that any interference with the competitive process must be costly. Similarly, Fryer and Loury (2003) show that providing subsidies for disadvantaged groups can raise the expected effort levels of the winner and the winning rate of disadvantaged groups in a general tournament model. These results illustrate an interesting possibility, i.e. that some degree of affirmative action (generally mild) has the dual effect of increasing minority employment and increasing efficiency as the tournament becomes more competitive.

This is the motivation for experiments conducted by Schotter and Weigelt (1992), who used a standard tournament-theoretic framework to show that affirmative action need not yield a cost/efficiency trade off. In their experiment, 20 students were each given an envelope when they entered the room. Each envelope contained a card with a random number generated from a

uniform distribution. Students were randomly assigned seats and an anonymous partner for the experiment. Once the experiment began, participants were asked to choose a number between 0 and 100, which was referred to as their “decision number.” Further, they were told that decision numbers had associated costs in the sense that their cost increased with the decision number chosen. After participants recorded their decision number, they opened the envelopes containing their random numbers. This random number was added to their decision number to generate their total amount of effort. The individual with the highest effort in each pair won the prize. Schotter and Weigelt (1992) noted two interesting patterns in the data: there is a slight over supply of effort, and affirmative action clearly benefits disadvantaged agents. Further, they find that affirmative action works best when there is a considerable difference between the advantaged and the disadvantaged group.

Corns and Schotter (1999) conducted a similar experiment in an auction setting. Four students were selected to be type B (low cost) bidders, and two students were selected to be type A (high cost) bidders. At the beginning of each of 20 rounds, the experimenter walked around the room with two bags marked A and B. Each bag contained chips corresponding to the costs for the two groups. After observing a cost chip, each participant wrote down a bid. The experimenter collected the bids and publicly announced who won, the price paid, and the group identity (A or B) of the winner. In some rounds, high cost bidders were offered price subsidies. The authors found that modest (five percent) price preferences for disadvantaged groups led to increases in minority representation and cost efficiency. They also found that price preferences that were too high (ten to fifteen percent) were not cost effective, even though they increased minority employment’s share.

The experiments reviewed here confirm the theoretical prediction that affirmative action

need not entail a cost efficiency/minority representation tradeoff. In particular, a subsidy to members of a high cost group that makes them more competitive can increase total effort from all participants. In this sense, affirmative action has the dual impact of benefiting a disadvantaged group and increasing overall efficiency.

Audit Studies and Field Experiments

In contrast to the relatively small literature on laboratory experiments with discrimination, there is a large literature on field experiments that are designed to detect discrimination in its primal form. Field experiments differ from laboratory experiments in a number of important ways, which are described nicely in Harrison and List (2004). Typically, researchers do not control the group identification or the underlying biases and attitudes of participants in a field study. In some (perhaps most) cases, participants in field studies do not even know that they are involved in an experiment. In addition, field experiments are generally conducted with relevant samples from non-student populations such as personnel directors or real estate agents. However, it is important to emphasize that field experiments are not flawless. Heckman and Siegelman (1993) identify five major threats to the validity of results from audit studies: (1) problems in effective matching; (2) the use of “overqualified” testers; (3) limited sampling frame for the selection of firms and jobs to be audited; (4) experimenter effects; and (5) the ethics of audit research. The field experiments reviewed here fall into three broad categories: labor markets, housing markets and product markets.

A small literature using audit studies involving resumés (Jowell and Prescott-Clarke 1970, Hubbick and Carter 1980, Brown and Gay 1985, Bertrand and Mullainathan 2004) provides some evidence of differential treatment in the initial hiring process. These studies send resumes of fictitious applicants to potential employers. The main difference between the two resumés is that on

one resumé the applicant has a distinctively black name and on the other the applicant has a traditionally white name. Such studies repeatedly have found that resúmes with white names are substantially more likely to lead to job interviews than the identical resúmes with distinctively black names. The most recent of these experiments, Bertrand and Mullainathan (2004), estimate that the response rate is fifty percent higher for resumes with “White” names controlling for quality, and there is a greater return to quality of the resumes for White applicants than for Black applicants. Interestingly, however, names such as Ebony and Latonya – which are clearly Black, but not necessarily associated with lower socio-economic status – received call-backs similar to the white mean. This raises the plausibility that names on resumes confound race and social class as a discriminating factor.

Housing market field experiments typically involve pairs of “equivalent” buyers who differ only in attributes such as sex, race or ethnic origin. The paired individuals each contact a designated landlord or real estate agent, and the response of that agent is recorded. Specifically, these paired studies measure the extent to which opportunities are denied or diminished for minority applicants (e.g., by showing them fewer apartments or by offering them less favorable rental terms). Many of these studies are audits conducted by government organizations and the results rarely are made publicly available. Galster (1990) gained access to a large number of audits conducted in the U.S. and Europe and reported significant discrimination based on race, ethnicity and gender.

Product market field experiments focus on differences in negotiation strategies when dealing with members of specific demographic groups. List (2003) studies dealer behavior in a sports card market. Subjects were recruited to act as either a buyer or a seller of a commonly available sports card with a predetermined redemption value (for buyers) or cost (for sellers). Each participant was instructed to approach a specific dealer and to negotiate a trade. Buyers earned the difference

between the purchase price they negotiated and the redemption value given to them previously by the experimenter. Similarly, sellers earned the difference between the negotiated sale price and the pre-determined cost of the card. The dealers were not aware that their initial and final price offers were being recorded and the non-dealer buyers and sellers were not aware that the purpose of the experiment was to evaluate possible discrimination. List (2003) reports that initial and final offers were less favorable for minority (women, non-whites and older) traders. Furthermore, he concludes that this discrimination was statistical in the sense that dealers were responding optimally to differences in the bargaining strategies of minority and majority traders.

These field experiments provide evidence consistent with racial discrimination in a wide range of economic markets, though other explanations such as class-based discrimination are also plausible. In at least some cases, the differential treatment is consistent with a profit maximizing response to differences in behavior between majority and minority groups. Moreover, List (2003) finds that the most experienced sports card dealers fully exploit these differences in bargaining strength. Riach and Rich (2002) surveys the studies discussed here and a large number of additional field experiments.

V. Conclusion

There are many theories that explain how discrimination might arise and persist in a variety of different situations. Since these theories typically rely on specific assumptions about beliefs and behavior, they are difficult to test with naturally occurring data. A large number of field experiments clearly document differential treatment of some groups based on certain demographic characteristics. In addition, laboratory experiments provide credible support for a number of theoretical insights. Group identification that results in discrimination can be induced in laboratory

experiments. The economic consequences of this discrimination include the obvious negative outcome for some members of the group discriminated against and the less obvious, but equally important, potential reductions in social welfare that arise. Examples abound: the efficiency loss caused by lower amounts of money passed to those from a particular ethnic group in a trust game; or the loss in worker productivity caused by experience-based discrimination, which may produce a cycle of low expectations, low aspirations, and inferior outcomes. Finally, experiments reveal that affirmative action can increase efficiency if subsidies to disadvantaged groups increase the competitiveness of effort-based allocations.

Appendix: A Guide to Conducting Discrimination Experiments in Economics

The most straightforward way to evaluate discrimination in the laboratory is to group people based on actual characteristics (e.g., major, home state or astrological sign) and to compare behavior and earnings with baseline experiments where group identification is not available. For example, subjects in a bargaining experiment could be told whether their partner is from the same group or not. Alternatively, market roles (e.g., buyer and seller) could be assigned based on some earned status condition (e.g., performance on a trivia quiz). Here the issue is whether to make actual group assignment on a random basis, which controls for individual differences, but involves elements of deception that are generally avoided in economics experiments. The alternative is to make the group assignments based on preferences that are not likely to be related to decision making skills in an economic context. For example, psychologists have used preferences for painting to make group assignments. Once the assignment method has been selected, the instructions for any standard economics experiment can be adapted and used. For example, Holt (2005) contains instructions for 20 experiments covering a wide range of topics. Instructions for other experiments are also widely available (e.g., voting, Anderson and Holt, 1999; information cascades, Anderson and Holt, 1997; public goods, Holt and Laury, 1997; and a market for pollution permits, Anderson and Stafford, 2000). These experiments can be conducted with a minimal set of props like cards and dice and a handout with instructions and record sheets. In addition, all of these experiments can be run over the Internet using the software associated with this book at <http://veconlab.econ.virginia.edu/admin.htm> This set of programs includes the statistical discrimination game used by Fryer, Goeree and Holt (2005). For a general introduction to the mechanics of conducting a research experiment, see Chapter 1 of Davis and Holt (1993).

References

- Allport, Gordon W. (1954) *The Nature of Prejudice*, Reading, MA: Addison-Wesley.
- Anderson, Donna M. and Michael J. Haupt (1999) "Employment and Statistical Discrimination: A Hands-On Experiment," *The Journal of Economics*, 25(1), pp. 85-102.
- Anderson, Lisa and Charles Holt (1999) "Agendas and Strategic Voting," *The Southern Economic Journal*, pp.622-629.
- Anderson, Lisa and Charles Holt (1997) "Information Cascades in the Laboratory," *The American Economic Review*, pp.847-862.
- Anderson, Lisa and Sarah Stafford (2000) "Choosing Winners and Losers in a Permit Trading Game," *The Southern Economic Journal*, pp.-212-219.
- Arrow, K.J. (1973) "The Theory of Discrimination," in Orley Ashenfelter and Albert Rees, eds., *Discrimination in Labor Markets*, Princeton University Press, Princeton NJ, pp. 3-33.
- Ball, Sheryl B. and Catherine C. Eckel (1998) "Stars Upon Thars: Status and Discrimination in Ultimatum Games," working paper, Virginia Tech.
- Ball, Sheryl B. and Catherine C. Eckel (1996) "Buying Status: Experimental Evidence on Status in Negotiation," *Psychology and Marketing*, 13 (4), pp. 381-405.
- Ball, Sheryl, Catherine Eckel, Philip J. Grossman and William Zame (2001) "Status in Markets," *Quarterly Journal of Economics*, pp. 101-188.
- Bargh, J.A. & Pietromonaco, P. (1982) "Automatic Information Processing and Social Perception: The Influence of Trait Information Presented Outside of Conscious Awareness on Impression Formation." *Journal of Personality and Social Psychology*, 43, pp. 437-449.
- Bertrand, Marianne and Sendhil Mullainathan (2004) "Are Emily and Greg More Employable Than Lakisha and Jamal: A Field Experiment on Labor Market Discrimination," *American*

- Economic Review*, 94 (4), pp. 991-1013.
- Billig, M. (1985) "Prejudice, Categorization, and Particularization: From a Perceptual to a Rhetorical Approach," *European Journal of Social Psychology*, 15, pp. 79-103.
- Billig, M. and H. Tajfel (1973) "Social Categorization and Similarity in Intergroup Behaviour," *European Journal of Social Psychology*, 3, pp. 27-52.
- Coate, Steven and Glenn Loury (1993) "Will Affirmative Action Eliminate Negative Stereotypes?" *American Economic Review*, 83(5), pp. 1220-1240.
- Corns, Allan, and Schotter, Andrew (1999) "Can Affirmative Action Be Cost Effective? An Experimental Examination of Price Preference Auctions," *American Economic Review*, 89, pp. 291-305.
- Davis, Douglas D. (1987) "Maximal Quality Selection and Discrimination in Employment," *Journal of Economic Behavior and Organization*, 8, pp. 97-112.
- Davis, Douglas D. and Charles A. Holt (1993) *Experimental Economics*, Princeton: Princeton University Press.
- Devin, P.G. (1989) "Stereotypes and Prejudice: Their Automatic and Controlled Responses" *Journal of Personality and Social Psychology*, 56, pp.5-18.
- Diehl, Michael (1988) "Social Identity and Minimal Groups: The Effects of Interpersonal and Intergroup Attitudinal Similarity on Intergroup Discrimination," *British Journal of Social Psychology*, 27 (4), December, pp. 289-300.
- Fershtman, Chiam, and Uri Gneezy (2001) "Discrimination in a Segmented Society: An Experimental Approach," *Quarterly Journal of Economics*, 116(1), February, pp. 351-377.
- Fiske, S.T. (1998) "Stereotyping, Prejudice, and Discrimination," Chapter 25 in *Handbook of Social Psychology*, 2, Gilbert, D.T., Fiske, S.T. and Lindzey, G., eds., Oxford: Oxford University

Press.

- Fryer, Roland G., Jacob K. Goeree, and Charles A. Holt (2001) "An Experimental Test of Statistical Discrimination," Discussion Paper, University of Virginia.
- Fryer, Roland G., Jacob K. Goeree, and Charles A. Holt (2005) "Experienced-Based Discrimination: Classroom Games," *Journal of Economic Education*, forthcoming.
- Fryer, Roland G., and Matthew O. Jackson (2003). "Categorical Cognition: A Psychological Model of Categories and Identification in Decision Making," *NBER Working Paper No. 9579*.
- Fryer, Roland G., and Glenn C. Loury (2002) "Categorical Redistribution in Winner-Take-All Markets," *NBER Working Paper No. 10104*.
- Galster, G. (1990) "Racial Discrimination in Housing Markets During the 1980s: A Review of the Audit Evidence," *Journal of Planning Education and Research*, vol. 9, pp. 165-175.
- Hamilton, D.L., ed., (1981). *Cognitive Processes in Stereotyping and Inter-group Behavior*. Hillsdale, NJ: Erlbaum.
- Harrison, Glenn and John List (2004) "Field Experiments," *Journal of Economic Literature*, 52, pp. 1009-1055.
- Holt, Charles and Susan Laury (1997) "Classroom Games: Voluntary Provision of a Public Good," *Journal of Economic Perspectives*, 11(4), pp. 209-215.
- Klein, Olivier and Assaad Azzi (2001) "Do High Status Groups Discriminate More? Differentiation Between Social Identity and Equity Concerns," *Social Behavior & Personality*, 29 (3), pp. 209-221.
- Li, J., Dunning, C., and Malpass, R. (1998) "Cross-racial Identification Among European-Americans: Basketball Fandom and the Contact Hypothesis," working paper.
- List, John (2003) "The Nature and Extent of Discrimination in the Marketplace: Evidence from

- the Field,” working paper, University of Maryland.
- Lundberg, Shelly and Richard Startz (1983) “Private Discrimination and Social Intervention in Competitive Markets,” *American Economic Review*, 73(3), June, pp. 340-347.
- Meissner, C. and Brigham, J. (2001) “Thirty Years of Investigating the Own-Race Bias in Memory for Faces: A Meta-Analytic Review,” *Psychology, Public Policy, and Law*, 7 (1), pp. 3-35.
- Myerson, Roger (1981) “Utilitarianism, Egalitarianism, and the Timing Effect in Social Choice Problems,” *Econometrica*, 49 (4), June, pp. 883-97.
- Mullainathan, Sendhil (2001) “Thinking Through Categories,” Mimeo, Harvard University.
- Peters, Williams (1971) *A Class Divided*, Doubleday and Company.
- Phelps, Edmund (1972) “The Statistical Theory of Racism and Sexism,” *American Economic Review*, 62, pp. 659-661.
- Schotter, Andrew, and Weigelt, Keith, (1992) “Asymmetric Tournaments, Equal Opportunity Laws, and Affirmative Action: Some Experimental Results,” *Quarterly Journal of Economics*, 107, pp. 511-39.
- Sporer, S. (2001) “Recognizing Faces of Other Ethnic Groups: An Integration of Theories,” *Psychology, Public Policy, and Law*, 7 (1), pp. 36-97.
- Tajfel, Henri (1969) “Cognitive Aspects of Prejudice,” *Journal of Social Issues*, 25(4), pp. 79-97.
- Tajfel, Henri (1970) “Experiments in Inter-Group Discrimination,” *Scientific American*, November, pp. 96-102.
- Turner, J. (1978) “Social Categorization and Social Discrimination in the Minimal Group Paradigm,” in Henri Tajfel, ed., *Differentiation Between Social Groups*, London:: Academic Press,.
- Turner, J. and R. Brown (1978) “Social Status, Cognitive Alternatives and Intergroup

Relations,” in Henri Tajfel, ed., *Differentiation Between Social Groups*, London: Academic Press.

Vaughan, Graham M., Henri Tajfel, and Jennifer Williams (1981) “Bias in Reward Allocation in an Intergroup and an Interpersonal Context,” *Social Psychology Quarterly*, 44 (1), pp. 37-42.