

Lecture Note: The Economics of Discrimination – Theory

David H. Autor
MIT 14.661 Fall 2003

November 24, 2003

1 ECONOMIC MODELS OF DISCRIMINATION

An enormous literature, starting with Becker's 1957 book *The Economics of Discrimination*, explores the economics of discrimination. Economic models of discrimination can be divided into two classes: competitive and collective models. Competitive models study individual maximizing behavior that may include discrimination. In collective models, groups act collectively against each other. Almost all economic analysis has focused on competitive models, and we'll do similarly here. Competitive models can further be divided into taste-based and statistical models of discrimination. Becker's model studied the former case, and we'll start with that.

First, we need to define discrimination. For our purposes, discrimination is when members of a minority are treated differently (less favorably) than members of a majority group with identical productive characteristics. Let the wage Y be equal to

$$Y = X\beta + \alpha Z + e, \tag{1}$$

where X is a vector of exogenous productivity characteristics and Z is an indicator variable for membership in a minority group. Assuming that $X\beta$ fully captures the set of productive characteristics and their returns and/or Z is uncorrelated with e , then discrimination is a case where $\alpha < 0$.

We already face difficulties just using this simple definition.

1. 'Productivity' may directly depend on Z – for example, in the entertainment industry or a market in which customers value Z in workers (e.g., discriminating customers). If customers will pay more to see a white actress or a black athlete, is this a legitimate component of productivity?
2. There is also a question of whether β – the production technology – is truly exogenous. For example, fire fighting equipment requires considerable physical strength and size to operate, an argument against the entry of women should not enter this profession. But these requirements are engineered attributes and probably could be altered. If humans were 20% less physically strong, presumably we could still fight fires. It's likely that in Japan, fire-fighting equipment demands a smaller physical stature.

3. The X 's could also be endogenous. Pre-market discrimination – or expectations of future discrimination – could reduce X 's for members of the minority group. (Examples: poor schools, or a rational belief among minorities that education will not be rewarded by the market.)

Point (1) is one we may be able to examine directly. Point (2) and (3) are much harder to test. But whether or not these are relevant, it can still be the case that $\alpha < 0$ conditional on both X and β , which would constitute discrimination.

2 TASTE-BASED DISCRIMINATION

Becker's 1957 book introduced the first economic model of discrimination. In this model, employers hold a 'taste for discrimination,' meaning that there is a disamenity value to employing minority workers. Hence, minority workers may have to 'compensate' employers by being more productive at a given wage or, equivalently, by accepting a lower wage for identical productivity. The basic insights of this model require almost no formalization, but we will formalize slightly.

- Let A denote majority group membership and B denote minority group membership.
- Employers will maximize a utility function that is the sum of profits plus the monetary value of utility from employing members of particular groups.
- Let d be the taste parameter of the firm, which Becker called the "coefficient of discrimination."
- Firms will maximize

$$U = pF(N_b + N_a) - w_a N_a - w_b N_b - d N_b,$$

where p is the price level, F is the production function, N_x is the number of workers of group $x = \{a, b\}$, and w_x is the wage paid to members of each group.

- Employers who are prejudiced ($d > 0$) will act as if the wage of b group members is $w_b + d$. Hence, they will only hire b group members if

$$w_a - w_b \geq d.$$

- Let $G(d)$ denote the CDF of the prejudice parameter d in the population of employers.
- The optimal number of workers hired at each firm is determined by the solutions to

$$\begin{aligned} pF'(N_a) &= w_a, \\ pF'(N_b) &= w_b + d. \end{aligned}$$

- Treating p as fixed and aggregating across firms in the economy leads to the market demand functions $N_a^d(w_a, w_b, G(d))$, $N_b^d(w_a, w_b, G(d))$ for each worker type. Wages are determined by

$$\begin{aligned} N_a^d(w_a, w_b, G(d)) &= N_a^s(w_a), \\ N_b^d(w_a, w_b, G(d)) &= N_a^s(w_d), \end{aligned}$$

where $N^s(\cdot)$ are the supply functions for the worker types.

- Notice the main point that comes out of this setup is this: A wage differential $w_b < w_a$ will arise if and only if the fraction of discriminating employers (or discriminating jobs) is sufficiently large that the demand for B workers when $w_b = w_a$ is less than the supply.
- In other words, *discrimination on average does not mean discrimination at the margin*. If there are enough non-discriminating employers, then discrimination is competed away. This also implies that minority workers don't work for discriminating employers.
- If the share of prejudice employers is sufficiently large, then some b group members will work at $d > 0$ employers, and this implies that $w_b < w_a$. In this case, the strength of prejudice at the margin (that is d for the marginal employer of b workers) is what affects determines the size of the wage gap.
- With free entry or constant returns to scale (CRS), these employers may be competed out of business. In a competitive market, each worker must earn his marginal product. Under CRS, non-discriminating firms would simply expand to arbitrage the wage differential born by minority workers. In equilibrium, discriminating employers must fund the cost

of their distaste out of their own pockets; they cannot pass the cost onto the minority worker.

So, to summarize:

- In partial equilibrium, minority workers must ‘compensate’ employers by being more productive at a given wage or, equivalently, accepting a lower wage for equivalent productivity.
- These tastes create incentives for segregation. It is potentially Pareto improving for minority workers to work in their own businesses and similarly for majority workers – then no one bears the cost of the distaste.
- In general equilibrium, these tastes can only be indulged at a positive cost to the employer.

Key testable implications of this model are:

- Wage differentials: Minority workers earn less than majority workers of identical productivity.
- Preferential hiring: Employers are less likely to hire minority workers of identical productivity.
- But these implications may not apply in equilibrium at the margin – so it’s not clear that we should observe them.

Is it a necessary implication of this model that firms employing minority workers have higher productivity? Not necessarily.

- Must distinguish: Discrimination at the margin from mean differences.
- Let’s say minorities are on average less productive than majority workers, and employers engage in taste-based discrimination. Then the marginal minority hire is underpaid relative to a majority hire, but it could still be the case that productivity is higher at non-minority workplaces.

- However, profitability – that is output minus wages paid to factors – should be higher at firms with greater minority employment.
- This again assumes that competition is not so ‘strenuous’ that all taste-based discrimination is eliminated.

2.1 WORKER AND CUSTOMER DISCRIMINATION

- Assume that some members of the a group are prejudiced against b workers and demand a premium to work alongside them. This is similar to the case above, and leads to segregation.
- Assume instead that customers discriminate against b workers and so get lower utility from purchasing services from a firm if they have to interact with a b worker. This will lower the labor market return to b workers to working in jobs with customer contact.
- In this case, it is not clear that consumer discrimination will be competed away by the market. This is because there is not an obvious way for one consumer to arbitrage the prejudice of another. Yes – consumers who are prejudiced may have to pay higher prices for goods. But this can be true in equilibrium; customers may be willing to bear these costs.
- Hence, this model suggests that customer prejudice may actually present a more enduring source of labor market discrimination than employer prejudice. The intriguing paper by Holzer and Ihlanfeldt on your syllabus presents evidence on this point. Ashely Lester will discuss the evidence on customer discrimination in Recitation on Friday.

2.2 STATISTICAL DISCRIMINATION

Most economic analyses of discrimination since Phelps (1972) and Arrow (1973) have focused on the statistical theory of discrimination, rather than taste-based discrimination. The premise of the statistical discrimination literature is that firms have limited information about the skills of job applicants. This gives them an incentive to use easily observable characteristics such as race or gender to infer the expected productivity of applicants (if these characteristics are

correlated with productivity). The Aigner and Cain (1977) article on your syllabus is the classic reference on this topic. We'll develop the statistical discrimination model, and then talk about its implications.

Statistical discrimination is the solution to a signal extraction problem. If an employer observes a noisy signal of applicant productivity and also has prior information about correlates of productivity (let's say a group-specific mean), then the expectation of applicant productivity should place weight on both the signal and the mean. Two cases are commonly exposted, and we'll also look at a third. We'll use normal distributions for simplicity, but this is not substantively important.

2.2.1 CASE 1: DIFFERENCE IN MEANS

- Assume that when workers apply for jobs, the employer sees race of the applicant $x = \{a, b\}$ and some error-ridden signal $\tilde{\eta}$ of productivity.
- Assume that employers have learned from experience that

$$\begin{aligned} \eta_x &\sim N(\bar{\eta}_x, \sigma_\eta^2) \text{ with} \\ \bar{\eta}_a &> \bar{\eta}_b, \text{ and } \sigma_\eta^2 \text{ identical for } a \text{ and } b. \end{aligned}$$

Hence, b group members are less productive on average, but the dispersion of productivity is the same for both groups. Notice that we can write $\eta_i = \eta_x + \varepsilon_i$.

- The productivity signal is error ridden in the following sense:

$$\begin{aligned} \tilde{\eta}_i &= \eta_i + \iota_i \text{ where} \\ \iota &\sim N(0, \sigma_\iota^2), \text{ with } \sigma_\iota^2 > 0. \end{aligned}$$

Hence:

$$\tilde{\eta}_i = \bar{\eta}_x + \varepsilon_i + \iota_i,$$

and $E(\tilde{\eta}_i | \eta_i) = \eta_i$, meaning that the signal is unbiased.

- What is the expectation of η given $\tilde{\eta}$ and x ? This is simply the regression equation,

$$\begin{aligned} E(\eta | \tilde{\eta}, x) &= \bar{\eta}_x (1 - \gamma) + \tilde{\eta} \gamma, \\ &= \bar{\eta}_x + (\tilde{\eta} - \bar{\eta}_x) \gamma. \end{aligned} \tag{2}$$

where $\gamma = \sigma_\epsilon^2 / (\sigma_\eta^2 + \sigma_\epsilon^2)$, which is the coefficient from a bivariate regression of η on $\tilde{\eta}$ and a constant (estimated separately by group). Note that $\gamma_a = \gamma_b$ in this example; all that differs is that $\bar{\eta}_a > \bar{\eta}_b$.

- This equation immediately implies that for a given $\tilde{\eta}$, the expected productivity of b applicants is below a applicants – even though $\tilde{\eta}$ is an unbiased signal for both workers. In particular

$$E(\eta|\tilde{\eta} = k, x = a) - E(\eta|\tilde{\eta} = k, x = b) = (\bar{\eta}_a - \bar{\eta}_b) \times (1 - \gamma).$$

This expression will always be positive provided that $\sigma_\epsilon^2 > 0$. (Draw yourself a picture.)

Two key points

- Notice the main insight/paradox of statistical discrimination:

$$\begin{aligned} E(\tilde{\eta}_i|\eta_i, x) &= \eta_i \\ \text{but } E(\eta_i|\tilde{\eta}_i, x) &\neq \eta_i \text{ (unless } \tilde{\eta}_i = \bar{\eta}_x). \end{aligned}$$

- Is there equal pay for equal productivity in this model? No, not in general. There is equal pay for equal *expected* productivity.
- Consider an a and b group worker who both have productivity $\eta = k$. So, $E(\tilde{\eta}_i|\eta_i, x) = k$ for both workers. But, using equation (2), it is clear that $E(\eta|\tilde{\eta}_i = k, x = b) < E(\eta|\tilde{\eta}_i = k, x = a)$.
- Hence, for some workers, statistical discrimination is ‘discrimination’ in the sense of equation (1).
- But this will not be true on average within each group; expected productivity equals true productivity *on average*. Be certain that you are clear on this point.

2.2.2 CASE 2: DIFFERENT VARIANCES

- Now take a case where $\bar{\eta}_b = \bar{\eta}_a$ and $\sigma_{\eta a}^2 = \sigma_{\eta b}^2 = \sigma_{\eta}^2$. So, the groups are identical in expectation.
- But consider that the signal $\tilde{\eta}$ may be more informative for one group or another. This would arise if for example a group managers were inaccurate judges of b group ability (or you could assume the reverse). This means that $\sigma_{ia}^2 \neq \sigma_{ib}^2$.
- It follows that $\gamma_a \neq \gamma_b$ since

$$\gamma_x = \frac{\sigma_{\eta}^2}{\sigma_{ix}^2 + \sigma_{\eta}^2}.$$

- For whichever group has lower σ_{ix}^2 , $\tilde{\eta}$ will be more informative; employers will put less weight on the mean for this group and more weight on the signal. See Figures 1a and 1b of Cain and Aigner.
- Depending on whether you are above or below the mean of your group, you are differentially helped or harmed by a steeper γ_x . If you are above the mean, you want the signal to be as informative as possible. If you are below the mean, you prefer an uninformative signal.
- Contrasting groups a and b as is done in the Cain and Aigner figures, you will see that the expectation of η given $\tilde{\eta}$ will cross at $\bar{\eta}$ for the two groups, and the relative steepness of the a versus b slope will depend positively on $\sigma_{ib}^2/\sigma_{ia}^2$.

2.2.3 CASE 3: RISK AVERSE EMPLOYERS

- Cain and Aigner also discuss a third plausible case that is rarely examined in the literature: Employer risk aversion.
- Here, there are diminishing returns to worker ability. Perhaps bad workers tend to destroy machines by accident, but really good workers are only slightly more productive operatives than average workers.

- In this case, lower testing precision is harmful to all workers, not just those below the mean.
- The estimation error of equation (2) is

$$e = E(\eta|\tilde{\eta}_i, x_i) - \eta = \frac{\iota\sigma_\iota^2 - \varepsilon\sigma_\eta^2}{\sigma_\iota^2 + \sigma_\eta^2},$$

which has variance:

$$V(e) \equiv \sigma_e^2 = \frac{\sigma_\eta^2\sigma_\iota^2}{\sigma_\eta^2 + \sigma_\iota^2}.$$

This variance is increasing in σ_ι^2 since $\partial\sigma_e^2/\partial\sigma_\iota^2 = \sigma_\eta^4/(\sigma_\iota^2 + \sigma_\eta^2)^2 > 0$. Lower signal precision (a higher value of σ_ι^2) is harmful if employers are risk averse.

- This is another reason to be concerned if *a* group employers are less able to accurately assess the productivity of *b* than *a* applicants – even if they are right on average.

2.2.4 TESTING STATISTICAL DISCRIMINATION

It's difficult to test statistical discrimination because it may be impossible to observe how employers form expectations. Almost any observed racial/gender difference in pay or hiring can be attributed to statistical discrimination (which is a problem for the theory, not a virtue). Almost all tests of statistical discrimination are therefore indirect. We'll talk about these in the 'Evidence' lecture.

2.2.5 STATISTICAL DISCRIMINATION: EFFICIENCY, LEGALITY, FAIRNESS,

Efficiency It's interesting to speculate on why economists have focused so much more attention on statistical than taste-based discrimination. My guesses:

1. The Becker model employs a modeling trick that many economists consider the last refuge of scoundrels – adding arguments to the utility function. This is a pretty undisciplined technique. By changing the utility maximand, you can pretty much get whatever you want. Arguably, however, it is also the most natural approach here. Casual empiricism says that much prejudice takes the form of 'distaste.'

2. Unlike taste-based discrimination, statistical discrimination is not competed away in equilibrium. So, we can be reasonably confident that we should be able to find it in a general set of cases.
3. Closely related to (2), statistical discrimination is ‘efficient.’ That is, statistical discrimination is the optimal solution to an information extraction problem. Economists would generally say that employers ‘should’ statistically discriminate. It is profit-maximizing, it is not motivated by animus, and it is arguably ‘fair’ since it treats people with the same expected productivity identically (though not necessarily with the same actual productivity). Hence, many economists might endorse statistical discrimination as good public policy.

Legality That said, statistical discrimination is generally unlawful. It is illegal in the U.S. to make hiring, pay or promotion decisions based on predicted performance where predictions are based on race, sex, age or disability. Because minorities, women, those over age 40, and the disabled are ‘protected groups,’ employers are not permitted to hire and fire them ‘at will.’ (An employer presumably can statistically discriminate among non-disabled, white males under age 40.) Statistical discrimination is probably difficult to detect, however, and so it is plausible that despite the law, it occurs frequently.

Fairness Leaving aside legality, it is worth asking whether statistical discrimination accords with most commons notions of fairness. Here it’s useful to take a loaded example: racial profiling. Say you are the New Jersey State Police, and there are some number of drug runners who travel on your highways. You have a limited amount of resources to expend on stopping cars, so you want to maximize your productive resources.

- Let’s go back to statistical discrimination Case 1: Difference in means. We are going to recast η as a latent index of criminality.
- Assume that when police officers observe cars on the highway, they see the race of the driver $x = \{a, b\}$ and some error-ridden signal $\tilde{\eta}$, corresponding to a latent index that that the driver is running drugs. For simplicity, define this assessment on the real line

$\tilde{\eta} \in [-\infty, \infty]$. A lower value of $\tilde{\eta}$ corresponds to a lower likelihood, and a higher value to a higher likelihood. (If you like, you can map these latent index values, $\tilde{\eta}$, into the CDF of the normal to get probabilities.)

- Assume that experience has taught the police that

$$\begin{aligned}\eta_x &\sim N(\bar{\eta}_x, \sigma_\eta^2) \text{ with} \\ \bar{\eta}_a &< \bar{\eta}_b, \text{ and } \sigma_\eta^2 \text{ identical for } a \text{ and } b.\end{aligned}$$

b group members are more likely to be running drugs, though the variance of the latent index is the same for both groups. As above, $\eta_i|x_i$ can be written as $\eta_i = \bar{\eta}_x + \varepsilon_i$.

- The signal for any given car/driver is error ridden.

$$\begin{aligned}\tilde{\eta}_i &= \eta_i + \iota_i \text{ where} \\ \iota &\sim N(0, \sigma_\iota^2), \text{ with } \sigma_\iota^2 > 0,\end{aligned}$$

which is noisy but unbiased.

- How should the police allocate enforcement resources? The optimal decision rule will involve a threshold value of η^* . Police will stop cars with $E(\eta) \geq \eta^*$. One can formalize this rule with a tiny search model, but I will not. The key point is that the police would ideally like to stop only the highest probability cars. But waiting has an opportunity cost, so it would be foolish to await only for cars with $\tilde{\eta}_i \rightarrow \infty$. Hence, the optimal decision rule will choose some η^* cutoff lower than ∞ : Stop any car that meets a critical value η^* . The threshold depends on the distribution of η , the arrival rate of cars, and the opportunity cost of waiting. We'll assume (with justification) that η^* will exist, and, with somewhat less justification, that the New Jersey State Police can solve for it.
- Assume that $\eta^* > \bar{\eta}_a, \bar{\eta}_b$. Hence, only a minority of cars from either group should be stopped.
- Since the police do not observe η for any car, they must form an expectation for this value. Using the equations above, the expected value of η given $\tilde{\eta}$ and x is:

$$E(\eta|\tilde{\eta}, x) = \bar{\eta}_x + (\tilde{\eta} - \bar{\eta}_x) \left(\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\iota^2} \right),$$

with estimation error,

$$e = E(\eta|\tilde{\eta}_i, x_i) - \eta = \frac{\iota\sigma_\iota^2 - \varepsilon\sigma_\eta^2}{\sigma_\iota^2 + \sigma_\eta^2},$$

which has variance

$$V(e) \equiv \sigma_e^2 = \frac{\sigma_\eta^2\sigma_\iota^2}{\sigma_\eta^2 + \sigma_\iota^2}. \quad (3)$$

- The variance of the expectation of η given x is the variance of true productivity minus estimation error. Define $\nu = \varepsilon - e$. We have,

$$V(E(\eta|\tilde{\eta}, x)) = V(\nu) = \sigma_\eta^2 + \sigma_e^2 - 2\text{Cov}(\varepsilon, e) = \sigma_\eta^2 - \sigma_e^2.$$

This expression underscores a crucial point. So long as $\sigma_\iota^2 > 0$ (the signal is error-ridden), the expectation of η given $\tilde{\eta}, x$ has lower variance than η . Therefore, this estimate ‘shrinks’ the true range of the underlying variable. **(Draw a picture.)**

- This is the essence of ‘racial profiling.’ Since the police know from experience that group b members are more likely to be running drugs than group a members, it is efficient to use this information in determining which cars to stop. You can demonstrate that this is efficient by confirming that the marginal η stopped is the same for both a and b . In fact, this *is* the decision rule.
- Note that for an individual with a $\iota_i = 0, x_i = x$, the police will necessarily under- or overestimate her true criminality unless η_i is equal to the group specific mean $\bar{\eta}_x$.
- Two important points follow from this rule.

1. The share of b cars stopped exceeds the share of a cars stopped. This can be seen as follows:

$$\begin{aligned} \Pr(\text{Stop}|x) &= \Pr(E(\eta|x) > \eta^*) \\ &= \Pr(\nu > \eta^* - \bar{\eta}_x) \\ &= \Pr\left(\frac{\nu}{\sigma_\nu} > \frac{\eta^* - \bar{\eta}_x}{\sigma_\nu}\right) \\ &= 1 - \Phi\left(\frac{\eta^* - \bar{\eta}_x}{\sigma_\nu}\right) \end{aligned} \quad (4)$$

where $\Phi(\cdot)$ is the cumulative density function of the standard normal distribution, and the quantity in parentheses $(\eta^* - \bar{\eta}_x)/\sigma_\nu$ is the ‘effective screening threshold’ for group x . This quantity is the standardized difference between the group’s mean and the threshold, scaled by screening precision. The lower is screening precision, the smaller is σ_ν , and the larger is the effective screening threshold.

– Differentiating (4) with respect to the group mean gives

$$\partial \Pr(\text{Stop}|x)/\partial \bar{\eta}_x = (1/\sigma_\nu) \phi\left(\frac{\eta^* - \bar{\eta}_x}{\sigma_\nu}\right),$$

where $\phi(\cdot)$ is the pdf of the standard normal distribution. Since $\phi(\cdot) > 0$, a higher value of $\bar{\eta}_x$ implies greater odds of being stopped.

– So, drivers from the b group are stopped more often.

2. The average η of b cars stopped exceeds that of a cars stopped. In other words, the level of criminality (or the fraction of criminals) is higher among stopped b cars. You can see this as follows:

$$\begin{aligned} E(\eta|\text{Stop}, x) &= \bar{\eta}_x + E(\varepsilon_\eta | \nu > \eta^* - \bar{\eta}_x) & (5) \\ &= \bar{\eta}_x + \sigma_\eta E\left(\frac{\varepsilon_\eta}{\sigma_\eta} \middle| \frac{\nu}{\sigma_\nu} > \frac{\eta^* - \bar{\eta}_x}{\sigma_\nu}\right) \\ &= \bar{\eta}_x + \rho_{\eta\nu} \sigma_\eta \lambda\left(\frac{\nu}{\sigma_\nu} \middle| \frac{\nu}{\sigma_\nu} > \frac{\eta^* - \bar{\eta}_x}{\sigma}\right) \\ &= \bar{\eta}_x + \frac{E[\varepsilon_\eta(\varepsilon_\eta - e)]}{\sigma_\eta \sigma_\nu} \sigma_\eta \lambda(Q_x) \\ &= \bar{\eta}_x + \frac{\sigma_\nu^2}{\sigma_\eta \sigma_\nu} \sigma_\eta \lambda(Q_x) \\ &= \bar{\eta}_x + \sigma_\nu \lambda\left(\frac{\eta^* - \bar{\eta}_x}{\sigma_\nu}\right), \end{aligned}$$

where $\lambda(Q) = \phi(Q)/(1 - \Phi(Q))$ is the Inverse Mills Ratio (IMR) and $\phi(\cdot)$ is the density function of the standard normal distribution.

– Differentiating this expression with respect to $\bar{\eta}_x$ gives

$$\partial E(\eta|\text{Stop}, x)/\partial \bar{\eta}_x = 1 - \lambda'(\cdot).$$

- Since $\lambda'(z) < 1$ for finite z , the expected criminality of those stopped is increasing in, $\bar{\eta}_x$, the group mean.
 - The reason is that the optimal stopping rule equates the *marginal* return to stopping an a versus b driver. Because there is more mass to the right of η^* for the b group, this means that the *average* criminality of stopped b drivers will be higher than a drivers .
- Most economists would view this stopping rule as ‘fair.’
 - The marginal car stopped has the same expected criminality for both groups.
 - There is no animus motivating these choices (i.e., taste-based discrimination).
 - Moreover, the average criminality is actually higher among the b group than the a group, despite the higher frequency of stops of b members.
 - Resources are efficiently deployed.
 - So, why do civil libertarians complain? And why do b group members get upset about being stopped for “DWB” (‘Driving While Black’)?
 - One possible answer is that well-intended liberals don’t understand basic statistical principles.
 - Another answer is that this system will seem demonstrably unfair to group b members who are *not* criminals.
 - Consider two citizens, one from group a , the other from group b , who have the same $\eta = k$. Assume that $k < \eta^*$, so neither ‘should’ be stopped. What is the likelihood that each is stopped?

$$\begin{aligned} \Pr(\text{Stop}|\eta = k, x) &= \Pr(E(\eta|x, \eta = k) > \eta^*) \\ &= \Pr\left(\tilde{\eta} > \frac{\eta^* - \bar{\eta}_x(1 - \gamma)}{\gamma} | x, \eta = k\right), \end{aligned}$$

where $\gamma = \sigma_\eta^2 / (\sigma_\epsilon^2 + \sigma_\eta^2)$, which is the regression coefficient from above.

- Notice that the expectation of the left-hand side of this expression is identical for both a and b since $\eta = k$. But the right-hand side is not. Because the effective screening threshold is declining in the group's mean $\bar{\eta}_x$, the b group driver with $\eta = k$ is more likely to be stopped than the a group driver $\eta = k$.
- Another way to see this: For any true level of criminality k , $E(\tilde{\eta}|\eta = k, x = a) = E(\tilde{\eta}|\eta = k, x = b)$ is identical for a and b . But $E(\eta|\tilde{\eta} = k, x = b) > E(\eta|\tilde{\eta} = k, x = a)$: the police will not treat these individuals identically. The b driver is more likely to be viewed as a criminal.
- Substantive point: Although racial profiling is an efficient way to apprehend criminals, it does impose a cost on all group b members, including the innocent. There are more Type I errors for group b , i.e., more innocent motorists stopped. If you are a b group member, that may seem patently unfair.
- Though this is rarely discussed, statistical discrimination does pose an equity-efficiency trade-off.
- This may explain why economic and lay intuitions on the virtues of statistical discrimination are typically at odds: I suspect that most economists might favor statistical discrimination as policy, but it is illegal.
- One possible contributor to this discrepancy: Economists may not generally recognize that statistical discrimination is inequitable on average, even if it is 'fair' at the margin.
- Another point of this example: Many important economic decisions are up/down, yes/no decisions (stop or not, hire or not, promote or not). In these cases, the crucial choice variable is whether some expectation exceeds a critical value, and this can have large distributional consequences. Imagine for example in our racial profiling model that $\bar{\eta}_a < \eta^* < \bar{\eta}_b$ and $\sigma_i^2 \rightarrow \infty$. That is, the signal $\tilde{\eta}$ is uninformative and so the decision rule puts full weight on the group mean. In this extreme case, only b group members are stopped, no a group criminals are apprehended, and all innocent b group members are inconvenienced.