

Session: Profiling: Normative, Theoretical, and Empirical Issues

Paper Title: Learning and Statistical Discrimination

Correspondent:

Professor Lawrence Blume

Department of Economics

Uris Hall

Cornell University

Ithaca NY 14850

Email: LB19@CORNELL.EDU

Fax: 607-255-2818

JEL Classification: D63, D82, J71

Ms. Source: This manuscript was prepared with LaTeX, specifically pdfTeX version 3.141592-1.11b-2.1, with the *AMSTeX* macro packages and fonts.

Word Count: 2463 including the cover page.

Learning and Statistical Discrimination

Lawrence E. Blume[†]

Economists and sociologists have highlighted the role self-perpetuating negative racial attitudes play in creating discriminatory social outcomes in labor markets. Robert Merton (1948) recognized this, as did Gunnar Myrdal (1944), who referred to an equilibrium of beliefs and behaviors as ‘the vicious circle’. Kenneth Arrow’s (1972) statistical discrimination model is a contemporary expression of this idea. He describes rational expectations equilibria in labor markets when workers’ skill-investment decisions are endogenous. The models exhibit multiple equilibria, some with high levels of outgroup employment and others with high levels of discrimination.

Multiple equilibrium models ‘explain’ differential outcomes by claiming that different instances of the phenomenon being studied are in different equilibria. These claims are not explanations because they say nothing about how this comes to be. Instead, this sorting is explained by a vague appeal to ‘history’. In this appeal is the implicit claim that the equilibrating process lies entirely outside the domain of theoretical analysis. Evolutionary game theory shows that the aforementioned claim is false in game-theoretic models. Dynamic processes select among equilibria, and selection is accomplished by the random fluctuations inherent in any equilibrium adjustment process. (See Blume (1993), Kandori, Mailath and Rob (1993) and Young (1993).)

Here I explore the implications of this kind of analysis for labor market equilibrium (as opposed to game theoretic) models. I present a dynamic version of the Steven Coate and Glenn Loury (1993) statistical discrimination model in which firms learn about worker investment decisions. Sampling variance in the distribution of worker and firm characteristics is the source of random perturbations

which favor selection for a particular equilibrium. This selection appears as a WLLN for the stationary distribution of the Markov learning process. Myrdal observed that equilibria may be more or less robust to perturbations; remedial policies might destabilize one equilibrium in favor of another. I demonstrate this effect: Policies can change which of multiple equilibria is selected for without making any of the equilibria disappear. I present the *simplest possible* version of a statistical discrimination learning model just to illustrate the aforementioned points.

I The Model

A market for skilled labor contains M workers, all members of the outgroup, and N firms. The model presented here makes no attempt to model intergroup competition for jobs. The matching model is quite simple: At most one worker appears at a firm at each date, and the firm must decide whether or not to hire her. Following Coate and Loury (1993), the market wage is exogenously fixed at w .

A worker's skill set is not observable at the time of hire, but only on the job. The cost to the firm of hiring an unskilled worker is $\eta > 0$. The benefit of a skilled worker depends upon the firm's type. The value of a skilled worker is θ . With probability $(1 - \epsilon)$ a firm is a *calculator*; it hires only if the expected benefit is positive. With probability $\epsilon > 0$ a firm *experiments* and hires regardless of the expected benefit. Calculators will hire an outgroup worker only if the (commonly held) probabilistic belief, π , that a worker is skilled, is sufficiently high. The probability π is endogenous.

Workers are classified into three types, according to the cost they pay to acquire skills. Type c workers, the most common, can acquire skills necessary for work at a cost $c > 0$. Type 0 workers are

naturally endowed with the skills or can acquire them for free. Type ∞ workers are unteachable. The cost to them of acquiring the necessary skills for the skilled labor market is infinite. This model differs from Coate and Loury (1993) in that the population is finite and the distribution of types is discrete.

The assignment of each worker to a type is random and independent of others' types. A worker is type 0 with probability ρ_0 , type ∞ with probability ρ_∞ , and type c with probability $1 - \rho_0 - \rho_\infty$. We expect most workers to be type c ; that is, ρ_0 and ρ_∞ are small. Type 0 workers are always skilled, type ∞ workers are never skilled, and type c acquire skills or do not, depending on their view of the return on the investment. Type c workers will choose to acquire skills only if the (commonly held) probabilistic belief, ν , that a worker will be employed is sufficiently high. The probability ν will be determined endogenously. The value of a job is the wage w . The value of being unemployed is pegged at 0. Assume that $c < w < \theta$.

II Static Equilibrium

Firms' hiring decisions and the skill acquisition decision of type c workers are determined in equilibrium. In any equilibrium, workers must maximize their expected returns in making the skill-acquisition decision, firms must maximize their expected profits in making hiring decisions, and expectations must be correct.

When a firm and a worker are matched the firm must make a hiring decision. The profit from not hiring is 0. The calculating firm hires the worker if and only if $\pi\theta - (\pi w + (1 - \pi)\eta) \geq 0$. The firm's *reservation belief* π^* is the probability of a worker being skilled which makes the

firm just indifferent between hiring and not; that is, the belief at which the expected profits of hiring are exactly 0. The reservation belief is uniquely determined by the model parameters w , θ and η . Experimenters always hire.

A worker must make a decision to invest in skill acquisition. For types 0 and ∞ this decision is trivial: Invest and not invest, respectively. Type c workers compare the expected benefit of the investment in skills, $v w - c$, to the return from having no skill, 0. The type c workers' *reservation belief* $v^* = c/w$ is the probability of employment at which those worker's who can acquire skills at cost c are indifferent about whether to do so or not.

In a static equilibrium, calculating firms maximize profits, type c workers make a return-maximizing skill-acquisition decision, and all beliefs are correct. An equilibrium is described by two variables: ρ_f , the probability that a firm offers a worker a job, and ρ_w the probability that a type c worker acquires skills. Let $q = \min\{N/M, 1\}$ denote the probability that a worker is matched.

Definition 1. *An equilibrium is a pair (ρ_f, ρ_w) of action probabilities such that ρ_f maximizes $\rho_f(\pi\theta - \pi w - (1 - \pi)\eta)$ and ρ_w maximizes $\rho_w(v w - c)$; and $\pi = \rho_0 + (1 - \rho_\infty - \rho_0)\rho_w$ and $v = (1 - \epsilon)q\rho_f + \epsilon q$,*

Although this definition allows for randomization, it will typically be the case that in equilibrium ρ_f and ρ_w take on the values 0 or 1. There are two possible types of these *pure equilibria*.

Definition 2. *A full-employment equilibrium is an equilibrium in which those workers who can acquire skills, and all who are matched are offered jobs; $\rho_f = 1$, $\rho_w = 1$, $\pi = 1 - \rho_\infty$ and $v = q$. An under-employment equilibrium is an equilibrium in which workers who can acquire skills at cost c choose not to, and firms do not offer jobs unless they experiment; $\rho_f = 0$, $\rho_w = 0$, $\pi = \rho_0$ and $v = q\epsilon$.*

For a range of parameter values, both full- and under-employment equilibria exist.

Theorem 1. *Assume that $\theta > w > c$, that $v^* < q$ and that $\rho_0 < \pi^* < 1 - \rho_\infty$. If $q\epsilon < v^*$, then both a full-employment and an underemployment equilibrium exist, and these are the only pure equilibria. If $v^* < q\epsilon$, then the only pure equilibrium is a full-employment equilibrium.*

Both ρ_0 and ρ_∞ should be small, so $\rho_0 < \pi^* < 1 - \rho_\infty$ is the ‘typical case’. This assumption will be maintained for the remainder of the paper. Most statistical discrimination models stop here, having noted the possibility of different social configurations, and explaining the observed configuration by reference to forces which lie entirely outside the model.

III Learning Dynamics

At the beginning of each date, workers make their skill-investment decision. Then they go to market and firms make their hiring decisions. Firms will construct beliefs from the previous date’s empirical distribution of outcomes. Workers observe the same data, compute firms’ beliefs, and thus whether or not calculators will offer them a job. (This is known as ‘anticipatory learning’. See Reinhardt Selten (1990)). There is a continuum of learning models to choose from. The key to the results presented here is that data is not persistent. Learning has a finite memory. Bayesian learning behaves quite differently. See Lawrence Blume (2004).

Of the M workers in the market at date t , K_t will be offered jobs, and J_t of those with jobs will have skills. It will be convenient to work with fractions off the worker population: Fraction $k_t = K_t/M$ of the workers will be offered jobs, and fraction $j_t = J_t/M$ of the worker population will have both jobs

and skills. The market data generated at each date is the pair (k_t, j_t) , and it is publically observable. Firm's beliefs at date $t + 1$ will be the fraction of employed workers who were skilled at date t , $\pi_{t+1} = j_t/k_t$. Worker's beliefs v_{t+1} will be q if $\pi_{t+1} \geq \pi^*$ and $q\epsilon$ otherwise.

Although the manner in which data is turned into firms' beliefs depends upon the particular learning rule firms use, for any learning model the data generation process is the same: If $\pi_t < \pi^*$, the only workers who will have jobs are those who are offered a job by experimenters and the probability that any one worker is skilled is ρ_0 , the probability that she is naturally gifted: $K_t \sim b(qM, \epsilon)$ and $J_t \sim b(K_t, \rho_0)$. Otherwise, all matched workers will get jobs, and the only workers who will not be skilled are those with infinite acquisition costs: $K_t \sim b(qM, 1)$ and $J_t \sim b(K_t, 1 - \rho_\infty)$, where $b(n, p)$ is the binomial distribution with population size n and success probability p . Since beliefs at date t are determined entirely by (k_{t-1}, j_{t-1}) , the market data process is Markov.

When $\pi_t < \pi^*$, $E\{j_t\} / E\{k_t\}$ is the under-employment equilibrium probability that a worker is skilled, and when $\pi_t \geq \pi^*$ this ratio is the full-employment probability. The Markov process has two transition regimes: The 'high' expectations (H) regime in which data is being generated by beliefs $\pi_t \geq \pi^*$, and the 'low' regime (L), with data generated by beliefs $\pi_t < \pi^*$. Note that starting from one regime, the probability of transiting to the other regime is independent from where in the regime one starts. Thus the stochastic process $\{s_t\}_{t=1}^\infty$, defined such that s_{t+1} takes on the value L in the low regime and H in the high regime, is also an ergodic Markov process.

The invariant distribution of the process gives the long-run frequency with which each region is visited. This frequency is given by the invariant distribution for the $\{s_t\}$ -process. The transition probabilities are the probabilities of observing a sample mean j_t/k_t on the other side of π^* from the current region. For any numbers of firms and workers, the probability of observing data which moves

the process from one region to the next is positive, and so the s_t -process is ergodic. Thus it has a unique invariant distribution μ . A calculation shows that the stationary odds ratio $\mu(H)/\mu(L)$ is the ratio of the probability of transiting from L to H to the probability of transiting back. These transition probabilities can in turn be estimated using results from large deviation theory. The end result is a rate function for the odds ratio. We compute the *normalized stationary log odds ratio* $(1/M)\mu(H)/\mu(L)$. Let $I(p, q)$ denote the *relative entropy* of p with respect to q . The central result about the long run behavior of the labor market with empirical learning is a large deviation rate function:

Theorem 2. *Suppose that $\theta > w > c$. Then*

$$\lim_{\substack{M, N \rightarrow \infty \\ N/M \rightarrow q}} \frac{1}{M} \log \frac{\mu(H)}{\mu(L)} = \begin{cases} q \log(1 - \epsilon + \epsilon \exp -I(\pi^*, \rho_0)) + qI(1 - \pi^*, \rho_\infty) & \text{if } q\epsilon < v^*; \\ qI(1 - \pi^*, \rho_\infty) & \text{if } q\epsilon \geq v^*. \end{cases}$$

When multiple equilibria exist, then as the market becomes large, the invariant probability of state H converges to 0 or 1, depending upon the sign of the value of the rate function. When only a single full-employment equilibrium exists, the invariant distribution converges to point mass on state H , since leaving the underemployment regime is typical while leaving the full-employment regime is rare.

Policy effects can be seen, among other places, in the behavior of the invariant distribution with respect to ϵ , the fraction of experimenters, and π^* , firms' reservation belief. Affirmative action hiring incentives will increase ϵ , while π^* can be decreased, for instance, by tax incentives for hiring outgroup members. For a fixed ϵ , there can be a critical value $0 < \hat{\pi} < 1$ such that the rate function is negative for large $\pi^* > \hat{\pi}$ and positive on the other side. If $\pi^* > \hat{\pi}$ small changes in π^* will have little effect on the invariant distribution. The probability of the H regime is near 0. But as π^* crosses

$\hat{\pi}$ the rate function changes sign and the probability of H is near 1. This change is discontinuous in the limit. Policy effects are *critical phenomena*, and $\hat{\pi}$ is a *critical point* for π^* .

The situation is more complicated with ϵ . For a given π^* there can be a critical $\hat{\epsilon}$ with the rate function positive for ϵ to the left and negative for ϵ to the right. So with multiple equilibria the high regime is favored by small ϵ . Changes in ϵ are a critical phenomenon, with critical point $\hat{\epsilon}$. But policies which increase the fraction of experimenters work against full-employment. They lead to larger worker sample sizes, which decrease the probability of large deviations, and in the low regime the mean fraction of unskilled workers is low. This adverse effect is different from the adverse affirmative action effect reported in Coate and Loury (1993). It is also possible that ϵ can be increased enough that $q\epsilon > v^*$. In this case the low equilibrium disappears. The probability of getting hired is sufficiently high that type c workers have reason to become skilled. So the policy response is insignificant for small changes, bad for larger changes and good for very large changes. This effect depends critically on the anticipatory learning assumption.

IV Conclusion

The stochastic process describing labor market outcomes is Markov. Although the process returns to every equilibrium regime, one regime occurs most frequently, and this frequency converges to 1 as the size of the market grows. Such states are said to be *stochastically stable*. Although the static model has multiple equilibria, the remaining equilibria appear as *metastable states* — short run, but not long-run persistent. This raises the question, how long is the long run? The question is empirical, and the answer depends upon details of the model that have been ignored here. For instance, it is

well understood in the game theory literature that if information is local in nature rather than global, as assumed here, transitions speed up.

The precise nature of the policy effects described above will vary depending on the specification of the learning models. But so long as the process is Markov and driven by a data-generation process like that described here, the general features persist: Policy moves will be effective only insofar as they cross a critical point or change the structure of equilibrium, and the effects of effective policy changes are discontinuous.

References

- Arrow, Kenneth J.** "Some Mathematical Models of Race in the Labor Market." in Anthony Pascal, ed., Racial discrimination in economic life, Lexington MA: Lexington Books, 1972, pp. 187–204.
- Blume, Lawrence E.** "The Statistical Mechanics of Strategic Interaction." Games and Economic Behavior, July 1993, 4(3), pp. 387–424.
- _____. "Learning Dynamics and Statistical Discrimination." Working paper, Cornell University, 2004. Cornell University.
- Coate, Stephen and Loury, Glenn C.**, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" American Economic Review, December 1993, 83(5), pp. 1220–1240.
- Kandori, Michihiro; George Mailath, and Rafael Rob.** "Learning, Mutation and Long Run Equilibrium in Games." Econometrica, January 1993, 61(1), pp. 29–56.
- Merton, Robert K.** "The Self-Fulfilling Prophecy." The Antioch Review, 1948 (8), pp. 193–210.
- Myrdal, Gunnar.** An American dilemma: The negro problem and modern democracy, New York: Harper & Row, 1944.
- Selten, Reinhardt.** "Anticipatory Learning in Games," in Reinhardt Selten, ed., Game equilibrium models, vol. 1, evolution and game dynamics, Berlin/New York: Springer-Verlag, 1990.
- Young, H. Peyton.** "The Evolution of Conventions." Econometrica, January 1993, 61(1), pp. 57–84.

† Cornell University Department of Economics and The Santa Fe Institute. This research was supported by the John D. and Catherine T. MacArthur Foundation, The National Science Foundation and The Pew Charitable Trusts. The author is grateful to Buz Brock, Stephen Coate, Steven Durlauf, David Easley, Shelly Lundberg and Larry Samuelson for useful discussions.