The Dynamics of Statistical Discrimination[†]

Lawrence E. Blume
Department of Economics, Cornell University
and
The Santa Fe Institute

June 2004

This Draft is Preliminary
and Incomplete

[†] This research was supported by the John D. and Catherine T. MacArthur Foundation, The National Science Foundation, The Pew Charitable Trusts and the Santa Fe Institute. The author is grateful to Chris Barrett, Buz Brock, Stephen Coate, Steven Durlauf, David Easley and Shelly Lundberg for useful discussions on this topic (and many others).

Abstract

To Come.

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

1 Introduction

Contemporary economics looks at discrimination in labor markets through the lens of modern analytical constructs. Thus Becker (1957) analyzes the role of a taste for discrimination — aversion to the outgroup — in determining labor market outcomes. So-called rational discrimination models, such as those of Phelps (1972) and Spence (1973), attribute the persistence of discrimination to a market failure caused by information asymmetries. Sociologists and social psychologists understand race, ethnic and gender-based wage and employment inequality to be a consequence of social processes that work throughout society and not just in markets. These explanations range from holistic theories of conflict between large social entities such as races or ethnic groups for scarce economic and power resources that explain inequality as an immediate consequence of inter-group conflict, to micro-social theories of stereotyping and stigmatization that see inequality as a durable consequence of micro-social interaction. While the macro-social theories are antithetical to the to the stance of methodological individualism which is central to all modern economics, the micro-social theories share much in common with contemporary equilibrium theory.

The micro-social theories highlight the role of negative racial attitudes in creating social outcomes that in turn help perpetuate the negative attitudes. Thus discrimination and segregation in labor markets, for example, is part of an equilibrium of beliefs and behaviors. Sociologists point to Merton (1948) as a source for this view. A similar view of discrimination was put forward by Myrdal (1944), who referred to an equilibrium of beliefs and behaviors as 'the principle of cumulation' and 'the vicious circle'. Myrdal observes that different equilibria may be more or less robust to perturbations. Thus remedial policies might work by destabilizing one equilibrium in favor of another.

The contemporary economic version of this reasoning is best represented by the statistical discrimination models of Arrow (1972) and (1973, section 4), which explicitly model labor markets as a rational expectations equilibrium of beliefs and behaviors. Employers' beliefs about workers' skill levels determine their willingness to hire, which determines the rate of return on human capital investment, which determines workers' actual skill levels. The model is closed by assuming that expectations are rational. Arrow observes that this system can have multiple equilibria, some exhibiting high employment and wage levels and others exhibiting discrimination. One difference between Arrow on the one hand, and Merton and Myrdal on the other, is that whereas the older literature emphasizes the decision making of the dominant view, conditioned by its stereotypes of minority attributes, the modern literature on statistical discrimination is more concerned with interaction of dominant and outgroup decision making. Thus Myrdal (1944, p. 75) writes, 'White prejudice and discrimination keep the Negro low in standards of living, health, education, manners and morals. This, in its turn, gives support to white prejudice.' Arrow observes that one consequence of employer discrimination is that the rate of return on human capital investment is lower for the victims of discrimination. They rationally choose to invest less in human capital, with the effect of reifying employees' beliefs about the low productivity of outgroup workers.

Multiple equilibrium models explain differential outcomes simply by asserting that one analytical subject is in one equilibrium while another subject is in a different equilibrium. Thus the US and Rwanda are in different equilibria of an endogenous growth model, and Asians and Blacks are in different labor-market equilibria of a statistical discrimination model. Such explanations beg the question because they say nothing about how this comes to be. Instead, the explanation for this sorting is at most a vague appeal to 'history' . Underlying this assertion is the implicit claim that the equilibrating process lies entirely outside the domain of theoretical analysis.

Modern evolutionary game theory has taught us that this claim is wrong. A central lesson of the stochastic population games literature is that dynamic processes select among equilibria.[1] A second lesson of this literature is that fluctuations matter. This is to say that time paths of labor market outcomes are not fully determined by the model; that from the point of view of the model the indeterminate component of the path can be viewed as random fluctuation; and that the fluctuations are instrumental to understanding equilibrium selection. This paper adds a learning dynamic to a simple and otherwise static account of equilibrium statistical discrimination. The assignment of workers to firms and the opportunity for firms to experiment generates a random data process from which firms learn. I will describe long-run stable patterns of discrimination that appear in the data process, and relate them to the equilibria of the static model. There are two points to the analysis here: First, that long-run patterns of discrimination can be identified with particular equilibria. Second, that although diffferent patterns corresponding to the different equilibria are possible, generically only one will be salient for any given specification of parameters. Finally, I will investigate how different assumptions about the learning process effect the duration of different long-run patterns.

2 The Model

The model to be presented here is a simplified version of the model described in Coate and Loury (1993). There are M workers and N firms. It is important that these populations are *finite*. Although ultimately a large-numbers analysis will be undertaken, a common mistake in the analysis of multiple-equilibrium models is to move to the large numbers limit too quickly. This will be discussed below. All proofs will be found in the appendix.

All workers are members of the outgroup. This is to say, the model presented here has only one group, and makes no attempt to model intergroup competition for jobs. The workers participate in a market for skilled jobs. They can only carry out the work successfully if they have acquired the necessary skills.

The worker population consists of three types of workers, classified according to the cost a worker pays to acquire skills. The most common type of worker, type c, can acquire skills necessary for work at a cost $c > 0$. A second type of worker, type 0, is naturally endowed with the skills or can acquire them for free. The final group of workers, type ∞, is unteachable. That is, the cost for them of acquiring the necessary skills for the skilled labor market is infinite.

The assignment of type to a worker is random, and the assignment of each worker to a type is independent of the assignment of others. A given 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 by type c; that is, ρ_0 and ρ_∞ are small. To summarize, the total number of workers is fixed at M, but the numbers of each type are random. The skill level of workers of the no-cost and infinite-cost types is fixed *ex ante*. The skill level of those who can acquire skills at cost is endogenous to the model. In making the skill-acquisition decision, workers will hold some beliefs about employment possibilities. The probability of getting a skilled job is ν, which is held in common by all workers.

Each of the N firms wants one skilled worker. Firms are of two types: Type θ values a skilled worker at $\theta > 0$. Type σ values a skilled worker at $\sigma \gg \theta$. Both type assign 0 value to unskilled workers. The assignment of types to firms is iid, and the probability of type σ is ϵ.

A worker's skill set is not observable at the time of hire. Workers must be put through an apprenticeship or training program or observed on the job before their skill level is observed. The cost to the firm of hiring an unskilled worker is $\eta > 0$. Although the skill of a given worker is unobservable *ex ante* employment, firms have beliefs about the likelihood that a typical worker is skilled. Expectations will be held in common across firms. Let π denote the probability any firm assigns to the event that a given worker is skilled. Beliefs π, which should in some way relate to the distribution of types, are determined in equilibrium in the static model and through learning in the dynamic models. Workers have no opportunity to signal their skill. The only marker they exhibit is their outgroup membership. In a multiple-group model we would assume that π is group specific: one value for whites, another for blacks, and so forth.

The market model is very simple, in order to facilitate the computations to come. Workers are assigned to firms in such a way that each firm sees no more than one worker. Let $q = \min\{N/M, 1\}$ denote the probability that a worker is matched with any firm. The firm with a worker at the door must decide whether to hire the worker or not. The wage rate for skilled workers is fixed at w, and $c < w < \theta$.[2] For workers, the labor market offers two possibilities: A skilled worker is matched with a firm, and earns the wage w. A worker who is not offered a job or who is fired from a skilled job goes immediately to the unskilled market, where the return is normalized to 0.[3]

3 Static Equilibrium

Firms' hiring decisions and the skill acquisition decision of those workers who can acquire skills at cost c 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.

A worker shows up at the door, and the firm must decide whether or not to hire her. The profit from not hiring is 0. The type-θ profit-maximizing firm will hire the worker if and only if

$$\pi\theta - \big(\pi w + (1-\pi)\eta\big) \geq 0$$

The firm's *reservation belief* π^* is that belief about 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 o hiring are exactly 0. The reservation belief is uniquely determined by the model parameters w, θ and η.

Type σ firms undertake a similar calculation. Assume, however, that σ is so large that type σ firms are always in the market, Type θ firms may or may not be, depending on market conditions.

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 believe with probability ν that they will be offered a job, and so the expected return to skill acquisitiion is $\nu w - c$. The return to having no skill is 0. The type c workers' *reservation belief* ν^* is that belief about employment at which those worker's who can acquire skills at cost c are indifferent about whether or not to do so. That is, ν^* solves

$$\nu w - c = 0$$

When $w > c$, reservation beliefs ν^* lie strictly between 0 and 1.

In an equilibrium of the static model, firms maximize profits subject to the possibility of experimentation, workers make a return-maximizing skill-acquisition decision, and all beliefs are correct. An equilibrium can be 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.

Definition 1. *An equilibrium is a pair* (ρ_f, ρ_w) *of action probabilities such that*

1. ρ_f *maximizes* $\rho_f\big(\pi\theta - \pi w - (1-\pi)\eta\big)$,
2. ρ_w *maximizes* $\rho_w(\nu w - c)$,
3. $\pi = \rho_0 + (1 - \rho_\infty - \rho_0)\rho_w$,
4. $\nu = (1-\epsilon)q\rho_f + \epsilon q$,

This definition allows for randomization in the event that firms or workers are indifferent over their choices, but generically it will be the case that equilibria are pure, that ρ_f and ρ_w take on the values 0 or 1. There are two possible types of 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* $\nu = q$. *An* underemployment 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* $\nu = q\epsilon$.

The analysis of this model is straightforward. Rather than go through all possible parameter combinations, the following theorem (and the remainder of this paper) treat only the most important cases.

Theorem 1. *Assume that* $\theta > w > c$, *that* $\nu^* < q$ *and that* $\rho_0 < \pi^* < 1 - \rho_\infty$. *If* $q\epsilon < \nu^*$, *then both a full-employment and an underemployment equilibrium exist, and these are the only pure equilibria. If* $\nu^* < q\epsilon$, *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. With parameter values in this range, and $q\epsilon < \nu^*$, the model exhibits multiple equilibria. Both full- and underemployment equilibria exist. 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.

4 Learning Dynamics

In this section learning dynamics are introduced. In this section two learning models will be introduced: An empirical learning model, which takes current beliefs to be the empirical distribution of some finite past history, and a Bayesian model which perforce weights equally the entire past history of labor market outcomes. Neither model is meant to be realistic, but both models demonstrate how learning dynamics can force a kind of equilibrium selection, and how different learning models can affect the character of long-run market behavior.

4.1 Data

Both models work off the same data. The data generation is described in the following figure. Workers

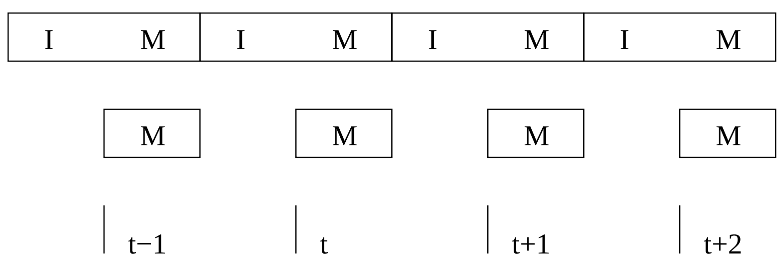

are dated not by the year of their birth but the year in which they come to the job market. A worker's life has two parts: A skill-acquisition period, denoted I (for *investment* in Figure 1) and a subsequent market period, denoted M in the diagram. At the beginning of period t, 'old' workers arrive in the market with any skills they have acquired, and are matched with firms. Based on their beliefs π_t, firms will make job offers. Of the M workers in the market at date t, K_t will receive jobs, and J_t of those with jobs will in fact have skills. It will be more convenient to normalize by the number of workers so as to work with population fractions rather than body counts. Let $k_t = K_t/M$ and $j_t = J_t/M$. The pair of numbers (j_t, k_t) is the market data generated at date t. The goal of this paper is to describe the stochastic process $\{(j_t, k_t)\}_{t=0}^{\infty}$. From this new data, firms and workers will update their beliefs, respectively, to π_{t+1} and ν_{t+1}. At this point, in the second half of date t, 'young' date $t+1$ workers make their skill acquisition decisions based on their beliefs ν_{t+1}. All data is public, and known to both workers and firms.

Both learning models considered here have firms forecasting the date $t+1$ skill frequency outcome based on previous data. Since this data is publicly available, workers can accurately predict the firms' forecast, π_{t+1}, and then use the firms' decision rule to construct their own forecast ν_{t+1} of the likelihood of employment.

$$\nu_{t+1} = \begin{cases} q & \text{if } \pi_{t+1} \geq \pi^*, \\ q\epsilon & \text{otherwise.} \end{cases}$$

The manner in which data is turned into firms' beliefs depends upon the particular learning rule firms use, but for any learning model the data generation process is the same, and is described in the following table:

	$\pi_t < \pi^*$	$\pi_t \geq \pi^*$
$\nu^* > q\epsilon$	$k_t \sim b(qM, \epsilon)$ $j_t \sim b(k_t, \rho_0)$	$k_t \sim b(qM, 1)$ $j_t \sim b(qM, 1-\rho_\infty)$
$\nu^* < q\epsilon$	$k_t \sim b(qM, \epsilon)$ $j_t \sim b(k_t, 1-\rho_\infty)$	

where $b(n, p)$ is a random variable distributed binomially from a population of size n with success probability p. If $\pi_t < \pi^*$ then the only workers who will have jobs are those who are offered a job by the type-σ firms. Among these workers, the type assignment is independent and the probability that any one worker is skilled is ρ_0, the probability that she is naturally gifted. When $\pi_t \geq \pi^*$, all matched workers will get jobs, and the only workers who will not be skilled are those with infinite acquisition costs. When $\pi_t < \pi^*$, type c workers can respond in one of two ways, depending on the relationship between ν^* and $q\epsilon$. If $\nu^* < q\epsilon$, employment only by type σ firms is not enough to induce these workers to acquire skills, while if $\nu^* > q\epsilon$, these workers will acquire skills. Notice that when $\nu^* > q\epsilon$ and $\pi_t < \pi^*$, the mean of the distributions for j_t and k_t correspond to the underemployment equilibrium, and when $\pi_t \geq \pi^*$ the means correspond to the full-employment equilibrum.

4.2 Empirical Distribution Learning

First suppose that workers and firms both predict the probability of a given worker's being skilled according to the empirical distribution of skilled data over some fixed past number of market periods. For clarity we will take the number of periods to be one, but any fixed finite horizon would do. From the date t data both workers and firms estimate the likelihood that a date $t+1$ worker will be skilled according to the rule

$$\pi_{t+1} = j_t / k_t.$$

This empirical rule determines firms' expectations.

Empirical learning makes the stochastic process $\{(j_t, k_t)\}_{t=0}^{\infty}$ Markov. The distribution of the date $t+1$ variables is determined entirely by the ratio j_t/k_t. The Markov process has two transition regimes. If $j_t/k_t < \pi^*$, then the joint distribution of (j_{t+1}, k_{t+1}) is described by one of the two leftmost columns of the table, depending on the relationship of the parameters ν^* and $q\epsilon$, each of which gives a distribution for k_{t+1}, and a conditional distribution for j_{t+1} given k_{t+1}. If $j_t/k_t \geq \pi^*$, the right column of the table describes the transition probabilities in a similar fashion. 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 when $j_t/k_t < \pi^*$ and H otherwise is also a Markov process. This process records the regime of the j,k process: s_t is low when date t expectations are low, and High when expectations are high.

Suppose that the j,k process is in the low regime at date t. The mean employment rate will be $q\epsilon$ and the mean frequency of skill acquisition is ρ_0. As M and N go to infinity in such a way that N/M converges to q, the weak law of large numbers implies that $k_{t+1}/qM \to \epsilon$ in probability. Similarly, $j_{t+1}/k_{t+1} \to \rho_0$. Thus, for large M, N, the process in the low regime hovers near the underemployment equilibrium. Similarly, the weak law implies that the (j,k) process hovers near the full-employment equilibrium in the high regime.

Theorem 2. *For all* $0 < q \leq 1$ *and* $\delta > 0$, *if* $\pi < \pi^*$ *and* $\nu^* > q\epsilon$, *then as* $M, N \to \infty$, $\max\{N/M, 1\} \to q$,

$$\mathrm{Prob}\left\{\left|\frac{k_{t+1}}{M} - q\epsilon\right| > \delta \,\Big|\, \frac{j_t}{k_t} = \pi\right\} \to 0 \quad \textit{and} \quad \mathrm{Prob}\left\{\left|\frac{j_{t+1}}{k_{t+1}} - \rho_0\right| > \delta \,\Big|\, \frac{j_t}{k_t} = \pi\right\} \to 0\,.$$

If $\pi < \pi^*$ *and* $\nu^* < q\epsilon$, *then*

$$\mathrm{Prob}\left\{\left|\frac{k_{t+1}}{M} - q\epsilon\right| > \delta \,\Big|\, \frac{j_t}{k_t} = \pi\right\} \to 0 \quad \textit{and} \quad \mathrm{Prob}\left\{\left|\frac{j_{t+1}}{k_{t+1}} - (1-\rho_\infty)\right| > \delta \,\Big|\, \frac{j_t}{k_t} = \pi\right\} \to 0\,.$$

If $\pi > \pi^*$, *then*

$$\mathrm{Prob}\left\{\left|\frac{k_{t+1}}{M} - q\right| > \delta \,\Big|\, \frac{j_t}{k_t} = \pi\right\} \to 0 \quad \textit{and} \quad \mathrm{Prob}\left\{\left|\frac{j_{t+1}}{k_{t+1}} - (1-\rho_\infty)\right| > \delta \,\Big|\, \frac{j_t}{k_t} = \pi\right\} \to 0\,.$$

This theorem describes the short-run behavior of the stochastic labor market process. If the process is in the full-employment regime, it is likely to remain in that regime, and within that regime, near the equilibrium, and increasingly so as the number of firms and workers becomes large. What happens in the underemployment regime depends upon the relationship between ν^* and $q\epsilon$. If ν^* is large, the process, once in the underemployment regime, will tend to remain there, and nearer the equilibrium as the market grows large. If ν^* is small relative to $q\epsilon$, the market will tend to jump to the full-employment regime.

We begin the asymptotic analysis by considering the multiple equilibrium case of $\nu^* > q\epsilon$. Although transitions between regions are rare, they do happen. With probability 1 there will always be another transition coming at some future date. The labor market process does not lock in to one regime forever.[4] Formally, transitions occur infinitely often. The long run behavior of the process is describing the long-run frequency with which each region is visited. This frequency is given by the invariant distribution for the s_t-process. Since this is a two-state process, it is straightforward to get at the transition probabilities.

$$\mathrm{Prob}\{s_{t+1} = H \mid s_t = L\} = \mathrm{Prob}\left\{\frac{j_t}{k_t} \geq \pi^* \mid \frac{j_{t-1}}{k_{t-1}} < \pi^*\right\}$$
$$\mathrm{Prob}\{s_{t+1} = L \mid s_t = H\} = \mathrm{Prob}\left\{\frac{j_t}{k_t} < \pi^* \mid \frac{j_{t-1}}{k_{t-1}} \geq \pi^*\right\}$$

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

$$\frac{\mu(H)}{\mu(L)} = \frac{\mathrm{Prob}\{s_{t+1} = H \mid s_t = L\}}{\mathrm{Prob}\{s_{t+1} = L \mid s_t = H\}} \tag{1}$$

The exact expressions for the right hand side appear to be intractable, but they can be simply estimated with the tools of *large deviation theory*.

Estimating the denominator is straightforward. When $s_t = H$, all firms hire workers, so the number of employed workers is qN,so $k_t = q$. Assign to each worker a random variable which takes on the value 1 if he is unskilled and 0 if he is skilled. These random variables take on the value if the worker has an infinite skill-acquisition cost, and 0 otherwise. The type assignment is independent, and so the random variables are independent, and the probability that a given random variable equals 1 is ρ_∞. The variable s_{t+1} will take on the value L if and only if among the qN workers with jobs, at least fraction $1 - \pi^*$ are unskilled.

For p and q between 0 and 1, recall the *relative entropy* of p with respect to q:

$$I(p, q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q} .$$

The basic result of large deviation theory is Cramer's Theorem, which estimates the probability of observing a sample mean much bigger (smaller) than the population mean. An immediate application of Cramer's Theorem gives the following lemma:

Lemma 1. *If* $s_t = H$,

$$\lim_{\substack{M,N\to\infty \\ N/M\to q}} \frac{1}{M} \log \mathrm{Prob}\left\{\frac{j_t}{k_t} < \pi^*\right\} = -qI(1-\pi^*, \rho_\infty)$$

The numerator is more complicated to estimate, because in state L, k_t is random. Nonetheless there is a rate function. The following lemma is proved in the appendix.

Lemma 2. *If $s_t = L$,*

$$\lim_{\substack{M,N\to\infty \\ N/M\to q}} \frac{1}{M} \log \mathrm{Prob}\left\{\frac{j_t}{k_t} > \pi^*\right\} = q \log(1-\epsilon+\epsilon \exp -I(\pi^*, \rho_0))$$

Combining these two estimates gives the central result about the long run behavior of the labor market with empirical learning.

Theorem 3. *Suppose that $q\epsilon < \nu^*$ and that $\theta > w > c$. Then*

$$\lim_{\substack{M,N\to\infty \\ N/M\to q}} \log \frac{1}{M}\frac{\mu(H)}{\mu(L)} = q \log(1-\epsilon+\epsilon \exp -I(\pi^*, \rho_0)) + qI(1-\pi^*, \rho_\infty) \tag{2}$$

Theorem 3 answers the following question: Starting from a given state at time 0 in a large market, what is the long-run distribution of states. That is, for a given market size, the distribution in equation (2) is an approximation to the true long-run distribution, and the approximation improves as the market size grows. Formally, equation (2) is the limit $\lim_{M,N\to\infty} \lim_{t\to\infty} \mathrm{Prob}\{s_t = X \mid s_0 = Y\}$ for states X and Y. This question, of computing the long run behavior of a large stochastic system, arises in many literatures, including some macroeconomic business cycle models, the labor market search literature, and evolutionary game theory. A very common approach is to look at the mean dynamics of the system; that is, to compute the limit $\lim_{t\to\infty} \lim_{M,N\to\infty} \mathrm{Prob}\{s_t = X \mid s_0 = Y\}$. The result of the interior limit is to remove the uncertainty by passing to the large numbers limit, which can be interpreted as the mean behavior of the system, and then to look at the asymptotic behavior of the deterministic dynamical system. This approach is justified by the law of large numbers, which is interpreted to say that the large system is approximately deterministic. It is often asserted, often implicitly, that the two limiting procedures give the same answers. This is wrong. The application of the law of large numbers in this way is incorrect, and the counterexample is provided by comparing the distribution in Theorem 3 to the limits of Theorem 2. The application of Theorem 2 to the long run would say that, in large markets, the system would stay for sure in the region in which it begins. It would suggest two long-run limit distributions, one at each of the two equilibria. The law of large numbers is useful for exploiting the short-run behavior of the system, as Theorem 2 suggests, but the limit behavior of the system is not nicely continuous in the transition probabilities.[5]

The theorem states that, so long as the right hand side of the equation is not 0, the invariant distribution will converge to placing probability 1 on one and only one region as the population sizes grow large in the prescribed manner. For large populations, the labor market will spend most of its time in the vicinity of one and only one equilibrium, that determined by the theorem. If the right hand side of (2), the *normalized log invariant odds ratio*, is positive, the invariant distribution converges to mass 1 on region H. If it is negative, it converges to mass 1 on region L.

The relative entropy $I(p, q)$ is a measure of distance between the probabilities p and q. The function is non-negative, convex, and takes the value 0 only when $p = q$. The function $-\log\big(1 - \epsilon + \epsilon \exp -I(p, q)\big)$, the rate function for the underemployment regime, is not convex, but it is non-negative, and 0 only when $p = q$. Figure 1 plots the normalized log invariant odds ratio against ϵ, the probability that a firm is type-σ. In this figure $\rho_0 = \rho_\infty = 0.1$. Figure 1 demonstrates

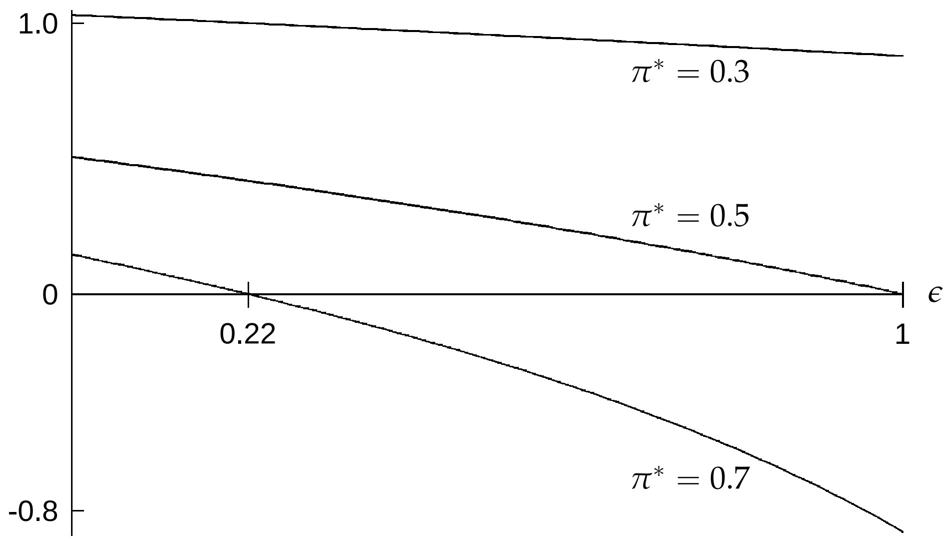

several properties of the model. First, the asymptotic behavior of the model exhibits a discontinuity as parameters change such that the continuous normalized log invariant odds ratio changes sign. If it is quite negative, policies which can change parameter values will have little effect on the long-run behavior of the market, but if it is only slightly negative, a small policy change can have a huge impact. Policies have critical points. Only if they cross these critical points can they have a marked impact on long run outcomes.

It is easily seen both in Figure 1 and by computation that the normalized log invariant odds ratio is declining in ϵ. As ϵ increases, the typical sample size of employed workers becomes larger, and so a large deviation becomes less likely. But ultimately $q\epsilon$ will exceed ν^*, and just because of experimentation workers will have an incentive to acquire skills. We might understand affirmative action hiring policies as policies which increase ϵ. As in Coate and Loury (1993), the effect of affirmative action hiring policies can increase rather than decrease the stability of the underemployment equilibrium, but for a different reason. Coate and Loury show that affirmative action policies can change investment incentives for workers for the worse. Here no such effect is possible, but affirmative action, by putting more workers into jobs, makes it harder to shift firms' expectations. Only when the policy effect is so large, making $q\epsilon > \nu^*$, will affirmative action hiring policies make the full employment policy significant. This significance can be seen in the long run analysis of the case $\nu^* < q\epsilon$. The transtion from the underemployment regime to full employment is no longer a rare event, but a transition from full employment to underemployment is. Calculations identical to those for the previous Theorem show:

Theorem 4. *Suppose that $q\epsilon > \nu^*$ and $\theta > w > c$. Then*

$$\lim_{\substack{M,N\to\infty \\ N/M\to q}} \log \frac{1}{M}\frac{\mu(H)}{\mu(L)} = qI(1-\pi^*, \rho_\infty) \tag{3}$$

This number is always positive for $\rho_0 < \pi^* < 1-\rho_\infty$. The invariant distribution always piles up in the full-employment regime as the market becomes large, because leaving the underemployment regime is typical while leaving the full-employment regime is rare.

4.3 Bayesian Learning

In the Bayesian learning model, firms assume that there is a fixed probability ψ that any given worker is employed. This is a boundedly rational learning model because the likelihood function is incorrectly specified, since ψ is dependent on the state of the market: $\psi = \rho_0$ if $\pi_t < \pi^*$, and $1-\rho_\infty$ otherwise. Suppose that firms believe that ψ is permanently fixed, and takes on one of these two values. They have common prior beliefs which assign probability p_t to the event that $\psi = \rho_0$, that workers who can acquire skills choose not to do so, and they update these beliefs in a Bayesian manner:

$$\log\frac{p_{t+1}}{1-p_{t+1}} = \log\frac{p_t}{1-p_t} + \log\Big(\frac{\rho_0}{1-\rho_\infty}\Big)^{j_t}\Big(\frac{(1-\rho_0)}{\rho_\infty}\Big)^{k_t-j_t}$$

The actual forecast is the predicted distribution:

$$\pi_{t+1} = p_{t+1}\rho_0 + (1-p_{t+1})(1-\rho_\infty)$$

The following Theorem describes the asymptotic behavior of the labor market process with this learning rule. Unlike Theorem 3, this is not a law of large numbers. The claims of the Theorem hold for all M and N. It should be clear that if $\nu^* < q\epsilon$, firms can almost surely learn that the mean fraction of skilled workers is $1 - \rho_\infty$, since although the size of the sample depends upon firms' beliefs, the probability of a worker being skilled is independent of the regime. Thus in this section attention is limited to the case of $\nu^* > q\epsilon$.

Theorem 5. *Suppose that $q\epsilon < \nu^*$ and that $\theta > w > c$. Then:*

1. *There exists a random variable p on the space of all sample paths of the (j, k) process which takes on only the values 0 and 1, and such that $p_t \to p$ almost surely.*

2. *Almost surely there is a time T such that for all $t > T$, $s_t = s_{t-1}$.*

3. *On the event $\{p_t \to 0\}$, $\pi_t \to 1 - \rho_\infty$ and $\nu_t \to 1$, the full-employment equilibrium. On the event $\{p_t \to 1\}$, $\pi_t \to \rho_0$ and $\nu_t \to q\epsilon$, the underemployment equilibrium.*

4. *For all $p_0 \neq 0, 1$, $\mathrm{Prob}\{p_t \to 0\}$ and $\mathrm{Prob}\{p_t \to 1\}$ are both positive.*

The final Theorem states that, as the market becomes large, it becomes impossible for beliefs to leave their initial region.

Theorem 6. *Suppose that $q\epsilon < \nu^*$ and that $\theta > w > c$. For $p_0 < p^*$, $\lim_{\substack{M,N\to\infty \\ N/M\to q}} \mathrm{Prob}\{p_t \to 0 | p_0\} = 1$, and for $p_0 > p^*$, $\lim_{\substack{M,N\to\infty \\ N/M\to q}} \mathrm{Prob}\{p_t \to 1 | p_0\} = 1$.*

Evidently the behavior of the labor market process is quite different with Bayesian learning than with empirical distribution learning. The (j, k) process is not Markov; it is the projection onto the first two coordinates of the Markov process $\{j_t, k_t, p_t\}_{t=0}^{\infty}$. The most striking difference between the two processes is that with empirical distribution learning the market process moves back and forth between the full-employment and underemployment regimes. With Bayesian learning the market process ultimately locks into one regime and remains there forever. However, it is possible to end up in either regime starting from any prior beliefs which reflect uncertainty about the true state of the world. Finally, the probability of locking into the region in which the market process starts converges to 1 as the market size becomes large. Both equilibria can arise in the large numbers limit, depending on initial conditions.

The market process can be thought of as a regime switching model. Let p^* solve $p\rho_0 + (1 - p)(1 - \rho_\infty) = \pi^*$. Thus p^* is that belief about models such that the forecast distribution is exactly π^*, which leaves firms indifferent between hiring workers and not. For $p_t < p^*$, the market process is governed by the probabilistic rules of the underemployment regime; while for $p_t > \pi^*$ the process evolves according to the full-employment rules. To understand why Theorem 5 is true, neglect the regime switching for a moment, and suppose that the process evolved according to the rules of the underemployment regime. Then two things are true. First, the log posterior odds process is a random walk since the log of the likelihood ratios of the date t data are iid. Second, the log posterior odds is diverging to $+\infty$ almost surely. These two facts imply that for any $p_0 > p^*$, the probability that for all t, $p_t > p^*$ is positive. This is to say, for any p_0 in the underemployment regime, there is a positive probability that p_t will remain inside the underemployment regime, and if it does it will converge to 1. This probability can be bounded away from 0 uniformly across the regime. A similar argument holds for the full employment regime. Now consider the possibility of regime switching. For any p_0 in any regime, there is a probability greater than some $\delta > 0$ that p_t never leaves the regime. From this it will follow that almost surely the process will be captured by one regime or the other.

4.4 Policy Analysis

Although the dynamics of the two learning models differs in fundamental ways, they have the same basic structure. The dynamics work on beliefs. Beliefs are partitioned into two regimes whose boundary is determined by π^*, firms' threshold belief. There is always a tendency to stay in the full-employment regime once having arrived. When workers' threshold belief is large enough, $\nu^* > q\epsilon$, there is a similar tendency to remain in the underemployment regime as well. Long run dynamics are determined by how easy it is to slip out of one region or another, and this is determined by the location of the boundary.

Any policy which lowers π^* increases the long-run salience of the full-employment regime. In the empirical learning model this means an increase in the long-run frequency with which that state occurs. In the Bayesian model this means an increase in the probability of locking into the full-employment regime. In both models the effect of a change in policy regime has a discontinuous effect in the large-economy limit. In the large-numbers limit, there is a threshold $\bar{\pi}$, which just sets to 0 the right hand side of equation (2), such that if $\pi^* < \bar{\pi}$ the long run probability of full employment is 1 while if $\pi^* < \bar{\pi}$ the long run probability of full employment is 0. Lowering π^* is an effective policy if and only if it crosses this line. With Bayesian learning in the large numbers limit, prior beliefs below π^* converge to underemployment and those above converge to full employment. Lowering π^* is an effective policy if it falls below current beliefs. A lower π^* will result from policies that lower effective firm wages, provide hiring subsidies or lower screening costs.

Policies such as education subsidies or training programs reduce the cost of skill acquisition. Wage subsidies, guaranteed jobs and similar programs increase the return to skill acquisition. All of these policies work through lowering ν^*. In both models lowering ν^* has a threshold effect in markets of any size. For $\nu^* > q\epsilon$ the underemployment equilibrium is a short-run attractor. In the Bayesian model, lock-in to the underemployment regime occurs with positive probability. For $\nu^* < q\epsilon$ the underemployment equilibrium disappears, and the only long-run stable state is the full-employment regime.

Other policy experiments are possible and some, such as the effects of ϵ are quite interesting, but they require an extensive further analysis of the market process which will not be undertaken here.

5 Conclusion

This paper investigates the effects of learning in a dynamic statistical discrimination model. Its purpose is to see how learning dynamics reinforces or significantly modifies conclusions drawn from the analysis of static statistical discrimination models. Accordingly, the simplest possible statistical discrimination model is constructed in order to clearly illustrate the different effects learning can have.

The paper develops a framework which in its static version admits the possibility of multiple equilibria with different employment levels. When learning dynamics are introduced, it is seen that the effects of noisy data can influence equilibrium selection in significant ways, which depend upon the precise nature of how firms revise beliefs. The two learning models considered here differ in how they treat the distant past. Empirical distribution learning considers only a finite horizon of market history, while Bayesian learning makes all historical data equally influential.

With empirical distribution learning, the stochastic process describing labor market outcomes is Markov.[6] Although the process bounces back and forth across different equilibrium regimes, generically one regime occurs with the highest frequency, 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*. The market process tends to remain near an equilibrium regime, but will from time to time wander to another equilibrium regime. The time it remains near any given equilibrium regime before drifting away grows to infinity with the market size, but the time spent near the stochastically stable grows large at an exponential rate compared to the time spent near the other equilibria. Were there other non-equilibrium regimes, as there is when $\nu^* < q\epsilon$, the time spent in these regimes does not grow large with market size.

Market dynamics with Bayesian learning are quite different. Ultimately the market process will lock in to one regime or another. The selection effect is observed differently, by investigating the probability that, from a given initial condition, the market process can lock in to a particular equilibrium. The large numbers analysis with Bayesian learning is quite distinct from that of empirical learning. In the latter case history does not determine the long-run outcome. The long-run behavior of the market process is independent of initial conditions. It is described by an invariant distribution which puts weight on both equilibria. Furthermore, it is typically the case that as the market becomes large, the invariant distribution concentrates all its mass on one equilibrium. In this sense a large market can be said to select a particular equlibrium. The Bayesian model behaves quite differently. History matters. The market ultimately locks into one equilibrium. The probability of locking into a particular equilibrium depends upon the initial conditions. Finally, in large markets history can never be undone. The market remains in the regime in which it started. The policy intervention necessary to move the market from the underemployment regime to the full-employment regime will in general be much larger with Bayesian learning than with empirical distribution learning.

In general the phrase 'Bayesian learning' connotes a degree of rationality. Here this is not true since the likelihood functions are correctly specified only asymptotically. The likelihood functions do not allow firms to understand that the regime can switch. Of course with lock-in, the switching stops, and this is the sense in which the likelihood functions are correct in the limit. Both learning models reflect a degree of bounded rationality in firms' understanding of the market process. The difference between them is the persistence of data. One lesson of the comparison between the two learning models is the unsurprising fact that if history persists in the minds of the market participants, it will persist in market outcomes. The comparison of the two models suggests two interesting questions which will not be addressed here. First, what is the boundary between persistence and transience of data in learning models? The empirical learning model weighs a finite number of previous periods equally, and neglects the rest of the past. The Bayesian model treats the entire past symmetrically. One can imagine learning models which consider the entire past, but put more weight on recent on observations than those further back in time. The boundary between persistence and transience should appear in the rate at which thee rates decay. A second question has to do with mixed regimes. How does the long run look in a market consisting of both Bayesian learners and empirical learners? When the entire firm population is Bayesian, history matters to the long run. When the entire population is empirical, history fades away. Where in the distribution of the population between these two types is the boundary?

Statistical inference in this model is very one-sided. Firms learn directly from data. Either firms make their future hiring plans public to workers before the young make a skill-acquisition investment decision or workers see the same data and infer firms' hiring plans. Workers never make an inference directly from the previous cohort's employment experiences. One could imagine each group learning from its own experiences or one could imagine learning models where workers mix their inference of firm behavior with direct observation of previous labor market outcomes. Such models could be analyzed with similar methods to those used here.

So long as only skill frequencies are estimated, the roles of π^* and ν^*, the reservation beliefs of firms and workers, respectively, are quite different. The firms reservation belief determine a threshold between two market regimes: Full-employment and underemployment. The full-employment regime contains an equilibrium for all parameter values. The underemployment regime may or may not contain an equilibrium, and this is determined by ν^*. When $\nu^* > q\epsilon$ is sufficiently large there is an underemployment equilibrium, but not for $\nu^* < q\epsilon$, the probability of getting a job when only type σ firms are hiring.

It would be straightforward to treat a finite number of types of workers who can acquire skills at different costs. The model would then have multiple underemployment regimes with different employment rates. Some of these regimes could fail to have equilibria. The analysis of empirical learning would show that those regimes with equilibria are metastable, that there is typically a unique stochastically stable regime, and that those regimes without equilibria are not metastable. The analysis of the Bayesian regime would look very similar. In particular, the market would ultimately lock into one regime, and as the market grows large, the tendency to lock into the initial regime converges to 1. The analysis of continuous type distributions would look quite different from the simple calculations presented here, but there is reason to believe that, at least in models with a finite number of equilibria, the conclusions would not look too different.

Readers of the stochastic evolutionary game theory literature will recognize the analysis of empirical learning as a 'basin-hopping' exercise like those of Blume (1993), Kandori, Mailath, and Rob (1993) and Young (1993). These papers examine stochastically perturbed best- or better-reply adjustment dynamics in finite games, and relate the invariant distribution (the long run outcomes) to the Nash equilibria of the static game. Two criticisms of this literature have emerged: How sensitive are the results to the specification of the stochastic perturbation, and how long is the long run? Here the answer to the first question is, 'very much!' This in fact is the point, that how people learn makes a difference to the long run saliency of particular equilibria. The answer to the second question is, 'it's hard to know.' In the empirical learning model one can compute from the transition probabilities the waiting-time to leave a region, and they get larger with the size of the market. On the other hand, Bayesian learning speeds up with the larger sample sizes that result from larger markets. But these answers rely too much on the detailed specification of the model. In general this is an empirical question, whose answer depends upon details of the model that have been abstracted away from or simply ignored here. For instance, it is well understood in the game theory literature that if information is local in nature, transitions speed up. See, for instance, Young (forthcoming). We would expect with learning from local information to find stochastic stability results for empirical learning and lock-in with Bayesian learning, but the details could be quite different. But in any case a model with sufficient attention to realistic detail to seriously address this fundamental issue is beyond the scope of this paper.

The simple model explored here can be extended in a number of interesting ways. Analysis is still feasible when different groups compete for jobs. This extension may not be useful for modelling the distinct labor market experiences of in- and outgroups, because the important issue of how tags such as color or ethnicity acquire meaning, and why those meanings might differ for ingroup designators and outgroup designators, is not addressed here. But there is something to be learned about competition among outgroups from this model, such as the effect of group size. It is also possible to extend the model to allow for endogenous wage determination. Another issue to examine is the effect of environmental randomness. Suppose that wage rates fluctuated with macroeconomic conditions. Does the stochastic effect of environmental shifts help or hinder the stability of the full-employment regime? All these questions, even unanswered, serve to make the point that the analysis of multiple equilibria models should not end with a static analysis and a vague appeal to history.

Appendix: Proofs

Proof of Theorem 1: The proof of Theorem 1 is just a simple check. Suppose that $\nu^* > q\epsilon$. To see that full employment is an equilibrium, suppose that $\rho_f = 1$, and that firms as workers both have correct expectations. Then $\nu = q$. Since $\nu^* < q$, all type c workers will acquire skills, so $\pi = 1 - \rho_\infty$. Since $\pi^* < 1 - \rho_\infty$, $\rho_f = 1$ maximizes profits.

For the underemployment equilibrium, suppose that $rho_f = 0$. Then correct worker expectations has $\nu = q\epsilon < \nu^*$. So type c workers will choose not to acquire skills; that is, $\rho_w = 0$. Correct firm expectations are $\pi = \rho_0 < \pi^*$, so it is rational for firms not to employ workers, $\rho_f = 0$.

If $\nu^* < q\epsilon$, then type c workers will choose to acquire skills no matter what firms do, $\rho_w = 1$. Firm beliefs must be $\pi = 1 - \rho_0 > \pi^*$, and so firms will hire workers, $\rho_w = 1$. □

Proof of Theorem 2: This theorem is just the immediate application of the weak law of large numbers to the relevant distributions. □

Before undertaking the proof of Lemma 2, it is worthwhile to see why one should think it would be true. For arbitrary $k \geq 0$, Chernoff's bound gives

$$\mathrm{Prob}\left\{\frac{j_t}{k_t} > \pi^* \,|\, k_t = k\right\} \leq \exp -kI(\rho_0, \pi^*)$$

The random variable k_t is the sum of iid Bernoulli random variables. Expecting over k gives

$$\mathrm{Prob}\left\{\frac{j_t}{k_t} > \pi^*\right\} \leq \left(1-\epsilon+\epsilon\exp -I(\rho_0,\pi^*)\right)^{qM}. \tag{A-1}$$

Chernoff's bound actually gives the large deviation rate function for sample means of iid random variables with non-random size, so this calculation suggests that the right hand side of equation (A-1) should give the correct rate function. This requires showing that the right hand side provides a large numbers lower bound. Unfortunately this intuition seems to have nothing to do with producing a proof, since I have not been able to find a direct proof of the lower bound property. The proof proceeds indirectly.

Proof of Lemma 2: Let N denote a population size. Let $\{K_N\}_{N=1}^{\infty}$ denote a sequence of sample sizes; $K^N \sim b(N,\epsilon)$. For any T, let $S_T = \sum_{t=1}^{T} X_t$ denote a sum of T iid random variables X_t. The lemma claims that

$$\mathcal{I}(r,p,\epsilon) = \log\Big(1-\epsilon+\epsilon\exp -I(r,p)\Big) \tag{A-2}$$

is the rate function for the sequence of independent random variables $\{S_{K^N}/K_N\}_{N=1}^{\infty}$, the sample means of randomly sized samples from an ever-increasing population. Define the random variables $\hat{S}_K^N = (1/N)S_{K^N}$ and $\hat{K}^N = K^N/N$. The method of proof is to get a rate function for the pair of random variables $\left(\hat{S}_K^N, \hat{K}^N\right)$. The random variable of interest is $\hat{S}_K^N/\hat{K}_N$, so the rate function of interest can then be computed using the contraction principle. The relevant material on large deviations is surveyed in Dembo and Zeitouni (1992).

The rate function for the sequence $\left(\hat{S}_K^N, \hat{K}^N\right)_{N=1}^{\infty}$ is shown to exist and is computed using the Gärtner–Ellis Theorem. The scaled cumulant generating function (*scgf*) is

$$\lambda(\theta,\phi) = \lim_{N\to\infty}\frac{1}{N}\log E\{\exp N(\theta\hat{S}_K^N + \phi\hat{K}_N)\}.$$

The expectation can be computed by first conditioning on the sample size and then expecting it out. That is, the expectation is computed as

$$E_{\hat{K}_N}\Big\{E\{\exp N(\theta\hat{S}_{Nk} + \phi k) \mid \hat{K}_N = k\}\Big\}..$$

The conditional expectation is

$$\begin{aligned}
E\{\exp N(\theta\hat{S}_{Nk}+\phi k)\} &= \exp\{Nk\phi\}E\{\exp N\theta\frac{1}{N}\sum_{t=1}^{Nk}X_t\} \\
&= \exp\{Nk\phi\}E\{\exp\theta\sum_{t=1}^{Nk}X_t\} \\
&= \exp\{Nk\phi\}(1-p+pe^{\theta})^{Nk} \\
&= \exp Nk\{\phi+\log(1-p+pe^{\theta})\}\,.
\end{aligned}$$

Next, expect out the sample size. On the event $\hat{K}_N = k$, $K_N = Nk$. The sample size K_N is a sum of iid Bernoulli random variables Y_y with success probability ϵ, and so

$$\begin{aligned}
E\{\exp Nk\{\phi+\log(1-p+pe^{\theta})\}\} &= E\{\exp\big(\phi+\log(1-p+pe^{\theta})\big)\sum_{t=1}^{N}Y_t\} \\
&= \left(1-\epsilon+\epsilon e^{\phi+\log(1-p+pe^{\theta})}\right)^N \\
&= \left(1-\epsilon+\epsilon e^{\phi}(1-p+pe^{\theta})\right)^N.
\end{aligned}$$

Finally the scgf pops out:

$$\lambda(\theta.\phi) = \log\big(1-\epsilon+\epsilon(1-p)e^{\phi}+\epsilon p e^{\theta+\phi}\big)\,.$$

Because it is a scgf, $\lambda(\theta,\phi)$ is convex. It is defined on all of $\mathbf{R}^2$, finite everywhere and differentiable. This is more than enough for the Gärtner–Ellis Theorem. The conclusion of this Theorem is that the sequence $\left(\hat{S}_K^N,\hat{K}^N\right)_{N=1}^{\infty}$ satisfies the large deviation property with a rate function given by the Legendre transformation of $\lambda(\theta,\phi)$:

$$\Lambda(x,y) = \sup_{\theta,\phi}\{x\theta+y\phi-\lambda(\theta,\phi)\}\,.$$

A little calculus and much algebra later,

$$\begin{aligned}
\Lambda(x,y) &= x\,\log\frac{(1-p)x}{p(y-x)} - \log\frac{1-\epsilon}{1-y} + y\log\frac{(1-\epsilon)(y-x)}{\epsilon(1-p)(1-y)} \\
&= x\log\frac{1-p}{p} + y\frac{1-\epsilon}{\epsilon(1-p)} + x\log x - (1-y)\log(1-y) + (y-x)\log(y-x)
\end{aligned}$$

For $0 \le x \le y \le 1$. Recall that $\lim_{z\to 0} z\log z = 0$.

Fortunately this rate function is only an intermediate product. It is continuous, and one can see that for all $\alpha \geq 0$ the sets $\{x, y : \Lambda(x,y) \leq \alpha\}$ are compact. This is enough to apply the *contraction principle*. Let $f(\hat{S}_K^N, \hat{K}^N) = \hat{S}_K^N / \hat{K}^N$. The right hand side is the sample mean of interest. The contraction principle states that for a large deviation r,

$$\mathcal{I}(r, p, \epsilon) = \inf\{\Lambda(x,y) : f(x,y) = r\}\,.$$

That is, $\mathcal{I}(r, p, \epsilon) = \inf_{y \in [0,1]} \Lambda(ry, y)$. The rate function Λ is a Legendre transformation of a convex function and thus convex. The minimand is convex in y, and so the minimum is characterized by the first order conditions. Much more algebra gives the answer, which turns out to be $\mathcal{I}(r, p, \epsilon)$ defined in (A-2). This proves the lemma. □

Proof of Theorem 3: The quantity $(1/qM) \log \mu(H)/\mu(L)$ can be estimated from equation (1) using the estimates of Lemmas 1 and 2. □

Proof of Theorem 4: As in the proof of Theorem 3, use the large deviation estimates to estimate the normalized log of the transition probability ratios. □

Proof of Theorem 5: To prove 1, suppose that the process begins in the underemployment region. Consider the stochastic processes S_t with

$$S_0 = 0 \quad \text{and} \quad S_t = S_{t-1} + u_t \log \frac{\rho_0}{1 - \rho_\infty} + (1 - u_t) \log \frac{1 - \rho_0}{\rho_\infty}$$

where the u_t are iid $0, 1$-valued random variables, and the probability that a $u_t = 1$ is ρ_0. Let $n_0 = 0$ and $n_k = n_{k-1} + \tau_t$ where the τ_t are iid and distributed $b(qM, \epsilon)$. So long as the market process stays in the low region, beliefs at date t are described by the log-odds process

$$\log \frac{p_t}{1 - p_t} = \log \frac{p_0}{1 - p_0} + S_{\tau_t}.$$

If the market process leaves the underemployment regime at time $t - 1$ and returns at time t, then the log-odds of beliefs at time $t + 1$ are the sum

$$\log \frac{p_{t+1}}{1 - p_{t+1}} = \log \frac{p_t}{1 - p_t} + S_{\tau_1},$$

and so forth, where the new S- and τ-processes are independent copies of the original.

Let A_k denote the event that the belief processes leaves the underemployment region for the kth time. I claim first that $\text{Prob}\{A_k\, i.o.\} = 0$. The process $\{S_t\}_{t=0}^{\infty}$ is a random walk with a positive drift: $E\{S_t - S_{t-1}\} > 0$. Thus the probability that it never returns to 0 after time 0 is $\eta < 1$. In order for the belief process to leave for the first time, event A_1, S_t must return to 0, so $\text{Prob}\{A_1\} \leq \eta$. In order to leave for the second time, it must leave for the first time, then return, and then the second S_t process must return to 0, so $\text{Prob}\{A_2\} \leq \eta \cdot 1 \cdot \eta = \eta^2$, and so forth. Thus $\sum_k \text{Prob}\{A_k\} \leq \eta/(1-\eta) < \infty$, and the claim follows from the Borel-Cantelli lemma.

There are two ways that A_{k+1} can fail to happen: Either the process remains in the full-employment region after A_k, or upon return it never leaves the under-employment region again. Suppose the second case. Since S_n is a random walk with positive drift, $\lim_n S_n = +\infty$. Since $\lim_{k\to\infty} n_k = \infty$ almost surely, $\lim_k S_{n_k} \to \infty$, and $p_{n_k} \to 1$. In the other case, an identical analysis of the process on the full employment regime shows that $p_t \to 0$ almost surely. Finally, the same analysis works if the process begins in the full employment regime.

Claims 2 and 3 of the Theorem are immediate consequences of 1. If $p_t \to 1$, then $\pi_t \to \rho_0$, and so $\nu_t = q\epsilon$ for t large enough. If $p_t \to 0$, $\pi_t \to 1 - \rho_\infty$, and so $\nu_t = 1$ for t large enough.

To prove claim 4 of the Theorem, suppose that the process starts in the underemployment regime. Then $\text{Prob}\{A_1\} < 1$, and on the event A_1^c, $p_t \to 1$. On the other hand, for any initial beliefs, there is a T such that if $u_t = 1$ for all t and if $n_1 = t$, then $p_1 > p^*$. This is a positive probability event, so $\text{Prob}\{A_1^c\} > 0$. But once in the full employment regime, the probability of returning to the underemployment regime is less than 1. On the event that the process is forever in the full-employment regime, $p_t \to 0$. □

Proof of Theorem 6: Consider again the $\{S_t\}_{t=0}^{\infty}$ process. Since the mean increment is positive, $\lim_t S_t = +\infty$. Let τ denote the *last time* $S_t = 0$. Then τ is almost surely finite. Let F_τ denote its cdf. Then $\text{Prob}\{p_t \to 1\} \geq \text{Prob}\{\tau < n_1\}$. As $M, N \to \infty$, $n_1 \to \infty$ almost surely, so $F(n_1) \to 1$. □

Notes

[1]See Blume (1993), Kandori, Mailath, and Rob (1993) and Young (1993).

[2]Were this assumption not met, either workers would never acquire skills or firms would never hire workers. Only when these three model parameters are ordered in this way can anything interesting happen.

[3]I have implicitly assumed that the value to an unskilled worker of being hired and then fired is the same as that of having never been offered a job. One might imagine that part of the cost η is a payment to the worker. This gives the unskilled worker an incentive to participate in the job market. This would change the calculation of the reservation belief ν^* defined below, but otherwise has no effect on the analysis.

[4]If $\rho_\infty = 0$ and $\rho_0 > 0$, then full employment would be the unique absorbing state. If $\rho_\infty = 0$ and $\rho_0 > 0$ then underemployment would be the unique absorbing state.

[5]The stationary distribution varies upper hemi-continuously with the transition probabilities, and the failure is the blowing up of the correspondence in the limit, a failure of lower hemi-continuity.

[6]T-step Markov if the empirical distribution has a horizon of length T.

References

ARROW, K. J. (1972): “Some Mathematical Models of Race in the Labor Market,” in *Racial Discrimination in Economic Life*, ed. by A. Pascal, pp. 187–204. Lexington Books, Lexington MA.

——— (1973): “The Theory of Discrimination,” in *Discrimination in Labor Markets*, ed. by O. Ashenfelter, and A. Rees, pp. 3–33. Princeton University Press, Princeton NJ.

BECKER, G. S. (1957): *The Economics of Discrimination*. University of Chicago Press, Chicago.

BLUME, L. E. (1993): "The Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, 4, 387–424.

COATE, S., AND G. C. LOURY (1993): "Will Affirmative-Action Policies Eliminate Negative Stereotypes?," *American Economic Review*, 83(5), 1220–1240.

DEMBO, A., AND O. ZEITOUNI (1992): *Large Deviations Techniques and Applications*. Jones and Bartlett, Boston.

KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation and Long Run Equilibrium in Games," *Econometrica*, 61, 29–56.

MERTON, R. K. (1948): "The Self-Fulfilling Prophecy," *The Antioch Review*, 8, 193–210.

MYRDAL, G. (1944): *An American Dilemma: The Negro Problem and Modern Democracy*. Harper & Row, New York.

PHELPS, E. S. (1972): "The Statistical Theory of Racism and Sexism," *American Economic Review*, 62, 659–61.

SPENCE, A. M. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 87(3), 355–74.

YOUNG, H. P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.

——— (forthcoming): "The Diffusion of Innovations in Social Networks," in *The Economy as a Complex Evolving System III*, ed. by L. E. Blume, and S. Durlauf. Oxford University Press, Oxford UK.