SIEPR Discussion Paper No. 08-24

The Dynamics of Collective Reputation

By
Jonathan Levin
Stanford University

February 2009

The Dynamics of Collective Reputation*

Jonathan Levin†

February 2009

Abstract

I present a stochastic version of Tirole's (1996) collective reputation model. In equilibrium, group behavior is persistent due to a complementarity between the group's reputation, which depends on the past behavior of the group's members, and current incentives. A group can maintain a strong reputation even as conditions become unfavorable, while an improvement in the environment may not help a group with a poor reputation. I also connect the model to the theory of statistical discrimination and show that the same mechanism can explain why discrimination might persist over time.

*This is a revised version of a paper titled "Career Concerns and Collective Reputation" that I drafted a number of years ago. I remain thankful for the helpful suggestions of Richard Levin, Stephen Morris, Ilya Segal and Steve Tadelis and for financial support provided by the Cowles Foundation and the NSF.

†Department of Economics, Stanford University, Stanford CA 94305–6072; jdlevin@stanford.edu.

1 Introduction

Why do certain groups of individuals develop reputations, and why do these reputations persist over time? One explanation comes from theories of statistical discrimination that point to the self-fulfilling nature of expectations. In the models of Arrow (1972, 1973) and Coate and Loury (1993), low expectations about the skills of a certain workers can reduce their incentive to acquire human capital. The resulting low investment justifies the low expectations. Tirole (1996) suggested that a related mechanism might explain patterns of behavior in professional groups, such as the prevalence of hard work or corruption. He also argued that collective reputations might persist because history shapes expectations about the future. Tirole's specific model captures this idea at least partially. Under certain conditions, there are multiple steady-state equilibria but the efficient steady-state is unreachable if initial conditions are poor.

The persistence of collective reputation can be captured sharply in an environment where conditions evolve over time. I develop this point in a model that has roughly the same ingredients as Tirole's. Workers belonging to a given group (e.g. a professional or racial or ethnic group) interact with employers. Each worker has an incentive to work hard and develop an individual reputation. When information is imperfect, however, the strength of this incentive depends, through beliefs, on the behavior of a worker's peers. To understand why such a mechanism might arise, consider the reputational cost incurred by a physician who is sued for malpractice. If patients believe that only the least attentive doctors are sued, the penalty may be severe and create a strong incentive to avoid being sued. If patients believe malpractice suits are common and not very informative, the penalty may be reduced.

When the environment is stable, this mechanism can lead to alternative equilibria in which workers behave quite differently. This suggests that a change in expectations might trigger a sudden shift in behavior. Persistence is easier to understand when the environment is changing. I introduce this by allowing the cost of effort to evolve randomly, which pins down expectations in the manner of Burdzy and Frankel (2005). In equilibrium, workers behave well when conditions are objectively favorable and when the group enjoys a favorable reputation. Over time, behavior is persistent: a group with a good reputation can maintain it in poor conditions, while the costs of effort may have to fall dramatically before the members of a poorly-regarded group are motivated.

With a few twists, the collective reputation model can be interpreted as a dynamic version of standard statistical discrimination models (e.g. Phelps, 1972; Arrow, 1973; Coate and Loury, 1993; Moro and Norman, 2004). In these static models, otherwise homogenous groups may end up with different equilibrium wages or levels of human capital investment. The dynamic version illustrates how history can be decisive in structuring expectations and influencing behavior at any point in time. The model suggests that to escape a vicious cycle of low expectations and low investment, policy changes that lower the cost of human capital acquisition may need to be relatively large.

Blume (2006) presents a interesting dynamic analysis of statistical discrimination that uses ideas from evolutionary game theory. He points out that although there may be multiple static equilibria, in a finite population model where workers are subject to random small shocks, the economy will spend almost all of the time at one of the equilibria. The approach here is a bit different. I formulate both the statistical discrimination model and the collective reputation model as games with strategic complementaries, and characterize equilibrium dynamics using the approach of Frankel and Pauzner (2000), Burdzy, Frankel and Pauzner (2001), and Burdzy and Frankel (2005). Indeed, one can view this paper as identifying some new settings where their insights about coordination dynamics apply.

2 The Model

The model consists of a group of workers of unit mass and a corresponding set of employers. Parties share a common discount rate r and time is continuous. In each interval, $[t, t+dt)$, a fraction $\eta \cdot dt$ of workers are matched with employers and perform a service in exchange for compensation. For simplicity, imagine that each transaction is completed instantly, although it also would be possible to think of an employment spell lasting until the agent is re-matched, with the salary and performance level chosen at the outset.

Workers can perform with either low or high effort. A fraction $1-\gamma$ of workers are *unproductive* and always exert low effort. The remaining fraction γ are *productive*. Productive workers incur a cost $c(\theta_t)$ from exerting high effort. The parameter θ_t reflects (possibly time-varying) factors that make high effort more or less onerous, for instance the availability of resources or the opportunity costs of time. A higher value of θ corresponds to a lower cost

$c(\theta)$.

Employers naturally prefer high effort, but cannot effectively motivate an employee using financial incentives. Instead employers pay a salary based on expected performance. I model the salary process in a very simple way, assuming that if an employer expects high effort with probability $\rho \in [0,1]$, she will pay a salary $R(\rho)$, where R is some increasing function.

Workers may have a reputational incentive to work hard because future employers will observe an imperfect signal of their behavior. Specifically, I assume that each worker has a track record, or *individual reputation*, that is updated after every employment episode. An individual reputation can be either "Good" or "Bad". If a worker enters employment with a record $z \in \{B, G\}$ and exerts effort $e \in \{L, H\}$, his updated record will be "Good" with probability $\phi_{ze} \in [0,1]$. Of course, high effort is good for one's reputation, $\phi_{zH} \geq \phi_{zL}$, with strict inequality for $z = G$.

A simple example of this kind of reputation mechanism is that each employer observes a binary signal of a new hire's effort in his previous job. I allow for a bit more generality in that individual records can be sticky. To the extent that they are, however, let's assume that a good record improves prospects going forward, so $\phi_{Ge} \geq \phi_{Be}$, and does not reduce the returns to effort, i.e. $\phi_{GH} - \phi_{GL} \geq \phi_{BH} - \phi_{BL}$.

Each worker knows his own record, and everyone knows the fraction of workers that currently have good records. Define $\pi_e = \phi_{Be}/(1 - \phi_{Ge} + \phi_{Be})$ to be the steady-state probability of having a good record for a worker who always chooses effort $e \in \{L, H\}$. The fraction of unproductive workers with good records will always be π_L. Let ξ_t be the fraction of productive workers that have good records. Then the overall fraction of workers with good records is $x_t = \gamma\xi_t + (1-\gamma)\pi_L$. Because of the one-to-one relationship between x_t and ξ_t, we can think of them more or less interchangeably as measuring the group's *collective reputation.* It is convenient in the analysis below to focus on ξ_t.

When an employer meets a worker, she determines his salary by making an inference about his productivity. Using Bayes' rule, the probability than an employer hiring at time t will assign to a worker being productive is

$$\mu_{Gt} = \frac{\gamma\xi_t}{\gamma\xi_t + (1-\gamma)\pi_L} \quad \text{or} \quad \mu_{Bt} = \frac{\gamma(1-\xi_t)}{\gamma(1-\xi_t) + (1-\gamma)(1-\pi_L)}, \tag{1}$$

depending on the worker's record. Note that there is a complementarity between individual and group reputation. When ξ_t is higher, a good record is a stronger signal of productivity: μ_{Gt} is higher, while μ_{Bt} is reduced.

The worker's salary will depend on the employer's belief about current effort choices. If a worker hires an agent with record z at time t, she will expect high effort with probability

$$\rho_{Gt} = \mu_{Gt} h_{Gt} \quad \text{or} \quad \rho_{Bt} = \mu_{Bt} h_{Bt}, \tag{2}$$

where h_{zt} is the probability the employer assigns to getting high effort from a productive worker hired at time t with record z. Employers always expect unproductive workers to exert low effort. From these calculations, it follows that workers will a good record at time t can expect a compensation premium of $R(\rho_{Gt}) - R(\rho_{Bt})$. This premium is the key driver of worker incentives.

3 Collective Reputation Steady-States

I now describe potential outcomes in the case where the environment is unchanging. As in Tirole's model, there can be multiple steady-state equilibria due to the feedback between expectations and incentives.

Proposition 1 *There are between one and three pure-strategy steady-state equilibria. In a low incentive steady-state, all workers exert low effort and compensation does not depend on a worker's record. In a high incentive steady-state, productive workers always exert high effort. In an intermediate steady-state, productive workers exert high effort only if they have a good record to maintain.*[1]

Now consider the possibilities.

Steady-State with Low Effort. If productive workers exert low effort regardless of their records, then $h_{Gt} = h_{Bt} = 0$, and so $\rho_{Gt} = \rho_{Bt} = 0$ for all t. As a consequence, $R(\rho_{Gt}) = R(\rho_{Bt})$, and workers receive the same compensation regardless of their records. This means there is no return to building a good record and the only reason to choose high effort would

[1] If there are multiple pure-strategy equilibria, there will also be mixed equilibria in which at any point in time a fraction of workers exert high effort. In principle, one could also interpret these as pure strategy equilibria with non-anonymous strategies.

be a short-term benefit. So a necessary and sufficient condition for there to be a low incentive steady-state equilibrium is that $c(\theta) \geq 0$. In this steady-state, $\xi_t = \pi_L$.

Steady-State with High Effort. If productive workers always exert high effort, $h_{Gt} = h_{Bt} = 1$, and in the steady-state, $\xi_t = \pi_H > \pi_L$. Using the formulas above,

$$\rho_{Gt} = \frac{\gamma \pi_H}{\gamma \pi_H + (1-\gamma)\pi_L} \qquad \text{and} \qquad \rho_{Bt} = \frac{\gamma(1-\pi_H)}{\gamma(1-\pi_H) + (1-\gamma_t)(1-\pi_L)}.$$

The flow benefit to having a good record is then $\eta(R(\rho_G) - R(\rho_B))$, where we can drop the t subscript, and the following incentive condition ensures high effort by productive agents:

$$(\phi_{zH} - \phi_{zL}) \frac{\eta}{r + \eta(\phi_{BH}/\pi_H)} \left[R(\rho_G) - R(\rho_B)\right] \geq c(\theta).$$

If this condition holds for $z \in \{G, B\}$, there is a high-effort steady state equilibrium. At equilibrium, the high effort of productive workers makes individual records highly informative, giving workers a means and an incentive to signal their type.

Steady-State with Intermediate Effort. Finally, it may be possible to have possible to have a steady-state equilibrium in which productive workers exert effort to maintain a good record but not to improve a bad one. In such an equilibrium, $h_{G,t} = 1$ and $h_{B,t} = 0$, and in the steady-state $\xi_t = \xi_I = \phi_{BL}/(1 - \phi_{GH} + \phi_{BL}) \in (\pi_L, \pi_H)$. Employers assign probability $\rho_B = 0$ or $\rho_G = \gamma\xi_I/(\gamma\xi_I + (1-\gamma)\pi_L)$ to receiving high effort from workers with good and bad records. In equilibrium, the incremental present value of having a good rather than a bad record is

$$\Delta_I = \frac{\eta}{r + \eta(\phi_{BL}/\xi_I)} \left[R(\rho_G) - R(0) - c(\theta)\right].$$

The following incentive condition states that productive workers want to exert effort if and only if they have good records:[2]

$$(\phi_{GH} - \phi_{GL})\Delta_I \geq c(\theta) \geq (\phi_{BH} - \phi_{BL})\Delta_I.$$

This equilibrium can only exist if $c(\theta) > 0$, so there is also a low effort steady state.

[2] In Tirole's model, a condition of this sort is satisfied because agents can never undo past misdeeds. This is what allows for his "high-trust" equilibrium.

The model captures a particular type of external effect between workers in a given group. When some workers exert effort, it raises the expectations of employers about all members of the group. There is spillover effect, therefore, in the form of increased compensation. There is also the potential for increased incentives, leading to a complementarity between the effort choices of group members. As we will see in the next section, this complementarity causes equilibrium behavior to be persistent in a dynamic environment.

4 The Dynamics of Collective Reputation

This section introduces a stochastic element into the model by allowing the cost of effort to evolve over time. The resulting dynamic equilibrium captures persistence in an intuitive fashion, so that as conditions improve or deteriorate exogenously, the group's reputation and behavior may follow slowly if at all.

To proceed, we need a few assumptions on workers' cost functions and the way that costs evolve. First, let's assume the cost parameter θ_t follows a driftless Brownian motion, from some initial condition θ_0, with positive (but possibly small) instantaneous variance. Second, suppose there are values $\underline{\theta} < \overline{\theta}$, such that if conditions deteriorate below $\underline{\theta}$, no worker will want to provide effort regardless of how they expect other workers to behave, while if conditions improve above $\overline{\theta}$, it will be similarly dominant to exert high effort (so $c\left(\overline{\theta}\right) \leq 0$).[3] Finally, I assume that $R(\rho)$ and $c(\theta)$ are lipshitz continuous and also $c'(\theta) < k_c < 0$ on $[\underline{\theta}, \overline{\theta}]$.

In the dynamic setting, the state of play is summarized by the current cost conditions θ_t and the group's current reputation ξ_t.[4] The choice variables are the probabilities $h_G(\theta_t, \xi_t)$ and $h_B(\theta_t, \xi_t)$ that productive workers exert high effort (where convenient, I'll suppress the arguments and write $h_{zt} = h_z(\theta_t, \xi_t)$).

The workers' effort decisions determine how the group's reputation evolves. To derive this relationship, let $\psi_{z,t}$ denote the probability that a productive worker who is hired at t with a record z will have a good record after the transaction:

[3] Sufficient conditions for the dominance relationships assumed here are that $c(\underline{\theta}) \geq (\eta/r)\,(R(\pi_H) - R(\pi_L))$, and the same holds for $-c(\overline{\theta})$. Note also that the assumption of stationary, driftless Brownian motion can be relaxed considerably (see Burdzy and Frankel, 2005).

[4] In general, the agents' strategies could be conditioned on the entire history $(\theta_\tau, x_\tau)_{\tau\in[0,t]}$ as well as on time, but allowing this does not expand the set of equilibria.

$$\psi_{zt} = \phi_{zL} + h_{zt} \left(\phi_{zH} - \phi_{zL} \right). \tag{3}$$

Then the group reputation ξ_t will evolve according to

$$\dot{\xi}_t = \eta \left[-(1 - \psi_{Gt})\xi_t + \psi_{Bt} \left(1 - \xi_t\right) \right], \tag{4}$$

starting from its initial condition ξ_0.

When a worker meets an employer at time t, the employer pays a salary based on expected effort, as described above. We can let $\rho_G(\theta_t, \xi_t)$ and $\rho_B(\theta_t, \xi_t)$ denote the probability employers assign to receiving high effort at time t from workers with good and bad records. In equilibrium, these beliefs must be consistent with the workers' strategies and with Bayes' rule, as described above in equations (1) and (2).

For worker incentives, the key variable is the compensation premium $R(\rho_{Gt}) - R(\rho_{Bt})$. In particular, at time t, the flow value to having a good record rather than a bad record is:

$$\eta \left[\left(R(\rho_{Gt}) - R(\rho_{Gt}) \right) - \left(h_{Gt} - h_{Bt} \right) c(\theta_t) \right]. \tag{5}$$

Therefore the present value of having a good record rather than a bad one, starting at time t, is:

$$\Delta \left(\theta_t, \xi_t \right) = \mathbb{E} \int_t^\infty \left\{ \begin{array}{c} \exp\left(-\int_t^\tau \left(r + \eta(1 - \psi_{Gs} + \psi_{Bs}) \right) ds \right) \times \\ \eta \left[\left(R(\rho_{G\tau}) - R(\rho_{B\tau}) \right) - \left(h_{G\tau} - h_{B\tau} \right) c(\theta_\tau) \right] \end{array} \right\} d\tau. \tag{6}$$

This expression is the difference between two optimized present values, so $h_{G\tau}$ and $h_{B\tau}$, and by extension $\psi_{G\tau}$ and $\psi_{B\tau}$, reflect optimal time τ behavior, taking as given the behavior of other workers and the beliefs of employers. These latter quantities determine the evolution of the group's collective reputation ξ_t and the resulting wages. The expectation in (6) is over the possible paths $\theta_{\tau \geq t}$, starting from the present position θ_t.

For a worker at time t, high effort is optimal if and only if

$$\left(\phi_{zH} - \phi_{zL} \right) \Delta \left(\xi_t, \theta_t \right) \geq c(\theta_t), \tag{7}$$

where z is the worker's record. Note that if it is optimal to exert high effort with a bad record, it must also be optimal to exert high effort with a good record.

A *dynamic equilibrium* consists of strategies $h_G(\theta_t, \xi_t)$ and $h_B(\theta_t, \xi_t)$ for the productive agents, and beliefs $\rho_G(\theta_t, \xi_t)$, $\rho_B(\theta_t, \xi_t)$ for employers with the following properties: *(i)* agent behavior is optimal according to (7); *(ii)* the aggregate reputation of the group evolves according to (4); and *(iii)* employer beliefs are correct and follow (1) and (2).

The dynamic equilibrium will be unique and relatively easy to describe if the workers' effort choices are strategic complements. One sufficient condition for this, which I assume for the remainder of the section, is that workers use strategies h_G, h_B that are increasing in (θ, ξ), and that $R(\rho) = \underline{R}$ for all $\rho \leq \gamma$. Restricting attention to increasing strategies seems fairly mild: if θ is constant, it still permits all the steady-state equilibria in the previous section. The salary condition means that employers pay a minimal wage when the probability of effort is sufficiently low, and in particular will pay minimally for a worker with a bad record.[5]

To understand why these conditions create strategic complementarity, suppose workers use strategies h_G, h_B that are increasing in θ, ξ. Then the present value of having a good record, $\Delta(\theta, \xi)$ will be increasing in both θ and ξ. An increase in ξ, for instance, will raise the group's collective reputation along any path of $\theta_{\tau \geq t}$, and also the salary premium, increasing the value to having a good record. Consequently any given worker will do best to use an increasing strategy himself. Moreover, if other workers raise their effort, i.e. adopt a strategy $h'_z(\theta, \xi) \geq h_z(\theta, \xi)$, this also will improve the group's collective reputation and the salary premium going forward. So it will be optimal for a given worker to raise his own effort in response.

Proposition 2 *Under the above conditions, there is a unique dynamic equilibrium. In equilibrium, agent behavior follows a threshold rule: $h_z(\theta_t, \xi_t) = 1$ if and only if $\theta_t \geq Q_z(\xi_t)$, where $Q_G(\cdot)$, $Q_B(\cdot)$ are strictly decreasing and $Q_G(\xi) \leq Q_B(\xi)$ for all ξ.*

The proposition follows by adapting the arguments of Frankel and Pauzner (2000), Burdzy, Frankel and Pauzner (2001), and Burdzy and Frankel (2005) to the present setting. The details are a bit involved, but the basic intuition for why equilibrium is unique is reasonably straightforward. It resembles "global game" logic. In situations where θ becomes

[5] There are other sufficient conditions for strategic complementarity. For example, in the case where employers see only a noisy signal of the previous period's effort and compensation is linear $R(\rho) = \kappa\rho$, we have strategic complements so long as workers use increasing strategies and do not condition their effort on their record, which is irrelevant in terms of their incremental returns to high effort.

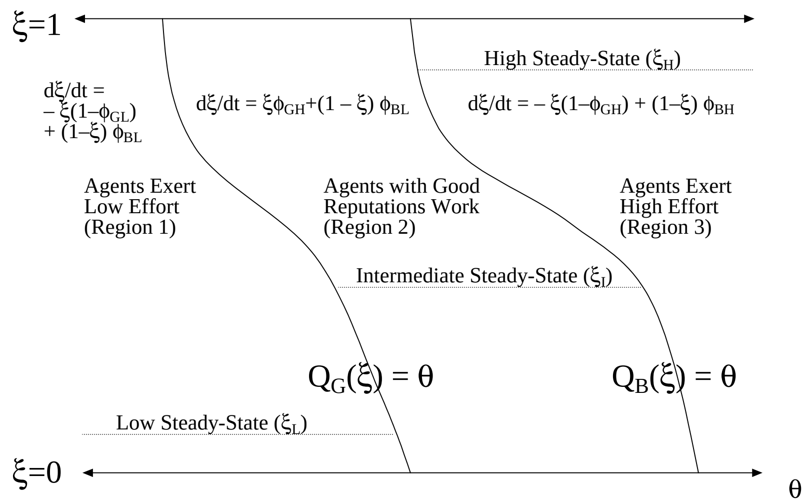

Figure 1: Dynamic Equilibrium

sufficiently high or low, behavior is uniquely determined. This places a limit on how optimistic or pessimistic one can be about worker behavior, and hence on the returns to building a good record. Iterating the best-response mapping places tighter and tighter bounds on rationalizable behavior, converging to a unique dynamic equilibrium.

More interesting than uniqueness per se is the description of behavior and the resulting dynamics. Figure 1 represents equilibrium behavior in (θ, ξ) space, in the case where $\phi_{GH} - \phi_{GL} > \phi_{BH} - \phi_{BL}$, so that the marginal effect of high effort on a worker's reputation is higher when the worker already has a high record. Workers exert effort when θ is high (current cost conditions are favorable), and when ξ is high (the group has a stronger collective reputation).

In Region 1, where $\theta < Q_G(\xi)$, productive workers choose low effort and employers pay $\underline{R}$ regardless of an agent's record. In Region 2, where $Q_G(\xi) < \theta < Q_B(\xi)$, employers pay workers with bad records $\underline{R}$, but offer a premium to workers with good records, anticipating that productive workers will exert effort to maintain their good reputations. Finally in Region 3, where $\theta > Q_B(\xi)$, workers with good records again command a premium and conditions are sufficiently favorable that all productive workers exert effort.

Dynamics of Collective Reputation. The evolution of the group's collective reputation depends on the current behavior of its members. When conditions are bad (Region 1), ξ_t trends down toward the low-reputation steady-state (i.e. toward $\xi_L = \pi_L$). In intermediate

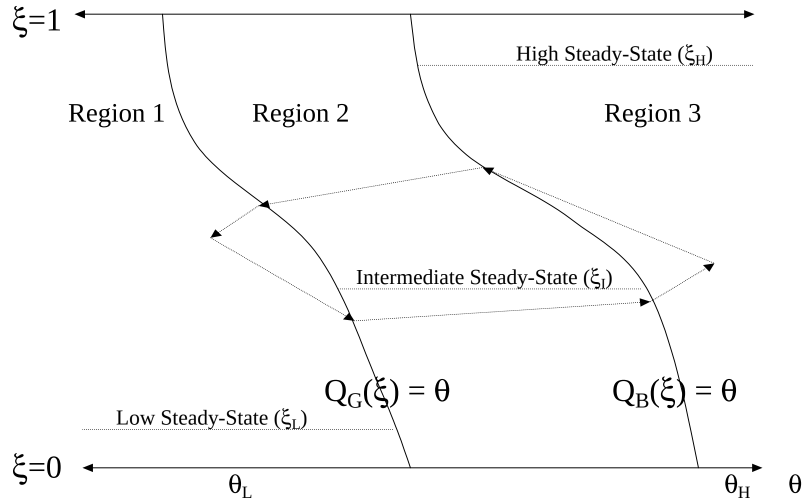

Figure 2: Equilibrium Dynamics, an Example

conditions (Region 2), the group's reputation evolves toward the intermediate steady state ξ_I. Finally, when conditions are good (Region 3) the collective reputation drifts toward the high-reputation steady-state ($\xi_H = \pi_H$). The compensation premium for having a good reputation tracks the group's reputation in Regions 2 and 3, while compensation is minimal in Region 1.

Persistence of Reputation. A highly-regarded group of workers can withstand worse conditions without their behavior deteriorating than can a group that initially has a poor reputation. To see this, observe that $Q_G(\xi)$ and $Q_B(\xi)$ are both strictly decreasing. Suppose that initial conditions are in Region 3. If θ declines so that high effort becomes more costly, then whether or not there is a fall-off in effort depends critically on the initial level of ξ. For instance, imagine θ oscillates back-and-forth across the interval $[\theta_L, \theta_H]$. Then ξ will follow a path that resembles the one depicted in Figure 2.

Selection of a Long-Run Steady-State. The persistent randomness of θ_t means that the model has no true steady-state, but nevertheless one can ask what happens as the instantaneous variance of θ_t approaches zero. In this case, an equilibrium selection result is obtained where the initial values of θ_0, ξ_0 determine the steady-state behavior toward which the group will evolve. There is still history dependence in this limiting case in the sense that for many values of θ_0 the selected steady-state will depend on ξ_0.

5 Persistence of Statistical Discrimination

There is a close connection between the collective reputation model and the classic statistical discrimination models of Arrow, Coate and Loury, and others. Recall that in the latter models, workers choose whether to invest in human capital. Employers observe a noisy signal of this investment and each worker are compensated based on the noisy signal, rather than his actual investment. If employers place more weight on the signal when they are optimistic that workers are investing, expectations can be self-fulfilling, leading to multiple pareto-ranked equilibria.

To re-cast the current model in this light, suppose that when a worker completes employment, he exits the economy and is replaced by a new worker. So new workers are arriving at a constant flow rate η. Each new worker chooses whether to invest in human capital. This investment is parallel to the effort decision above. The investment cost is infinite for a fraction $1-\gamma$ of workers, and equal to $c(\theta_t)$ for the remaining fraction γ. After making their investment decisions, new workers enter the pool of workers searching for employment and are matched to jobs at a flow rate η.

When a worker meets an employer, the employer observes a noisy signal of the agent's human capital. The signal is either "Good" or "Bad," and the probability it is good is ϕ_H if the worker has invested and ϕ_L if not. Therefore the probability an employer assigns to a worker with a signal z being skilled will be

$$\rho_{Gt} = \frac{\phi_H \xi_t}{\phi_H \xi_t + \phi_L (1-\xi_t)} \quad \text{or } \rho_{Bt} = \frac{(1-\phi_H)\,\xi_t}{(1-\phi_H)\,\xi_t + (1-\phi_L)(1-\xi_t)},$$

where ξ_t is now the fraction of workers in the pool searching for jobs who have invested in human capital.

As before, we assume that salary is based on expected productivity, which in this case depends only on the agent's human capital — there is no "on-the-job" effort. The salary will be $R(\rho_{Gt})$ if the employer observes a good signal and $R(\rho_{Bt})$ if the employer observes a bad signal. So the compensation premium for having a good signal is $R(\rho_{Gt}) - R(\rho_{Bt})$.

To understand the dynamics, let h_t denote the probability that a worker born at t invests in human capital. The fraction of skilled workers will evolve according to:

$$\dot{\xi}_t = \eta\,(h_t - \xi_t)\,. \tag{8}$$

For a worker born at t it will be optimal to invest if and only if:

$$(\phi_H - \phi_L)\,\Delta(\theta_t, \xi_t) \geq c(\theta_t), \tag{9}$$

where $\Delta(\theta_t, \xi_t)$ is the net present value of being able to present a good signal to an employer, rather than a bad signal. This present value can be expressed as:

$$\Delta(\theta_t, \xi_t) = \mathbb{E}\int_t^{\infty} \{\exp(-(r+\eta)(\tau - t))\,\eta\,[R(\rho_{G\tau}) - R(\rho_{B\tau})]\}\,d\tau. \tag{10}$$

Therefore any optimal strategy for workers satisfies:

$$h(\theta_t, \xi_t) = \begin{cases} 1 & \text{if } (\phi_H - \phi_L)\,\Delta(\theta_t, \xi_t) > c(\theta_t) \\ 0 & \text{if } (\phi_H - \phi_L)\,\Delta(\theta_t, \xi_t) < c(\theta_t) \end{cases}. \tag{11}$$

What types of equilibria are possible? If the cost of human capital investment is stable, so that $\theta_t = \theta$, the steady-state equilibria match those in standard statistical discrimination models. Provided that $c(\theta) > 0$, there is a low-investment equilibrium in which workers do not acquire skills, employers correctly anticipate this, and there is no compensation premium for a high signal. Provided that $\frac{\eta}{r+\eta}(R(\pi_H) - R(\pi_L)) \geq c(\theta)$, there is a high-investment equilibrium in which workers invest if their costs are not prohibitive, employers realize this, and this sustains a compensation premium for workers with high signals. These equilibria are analogous to the high and low steady-states of the collective reputation model.[6]

When the environment is evolving, the same forces that lead collective reputations to persist can also lead to persistent discrimination, perhaps even as a discriminated group receives greater access to resources that lower the cost of acquiring human capital. The next Proposition establishes this for the case where θ_t evolves as described above, following a driftless Brownian motion, and we impose the same assumptions as above on c and R.[7]

Proposition 3 *There is a unique dynamic equilibrium of the stochastic statistical discrimination model. In equilibrium, the cohort of workers entering at time t invest in human capital if and only if $\theta_t \geq P(\xi_t)$, where P is strictly decreasing in ξ.*

[6] If both the optimistic and pessimistic steady-states exist, there will also be a mixed equilibrium in which just a fraction of the workers in each cohort become high skill.

[7] That is, c and R are lipshitz, $-c'(\theta) > k_c > 0$, and $R(\rho) = \underline{R}$ for $\rho < \gamma$. Other assumptions will suffice to ensure that $R_G - R_B$ is increasing in ξ, for instance that $\phi_H = 1$. We also can drop the immediate focus on increasing strategies as no nonincreasing strategy will be rationalizable.

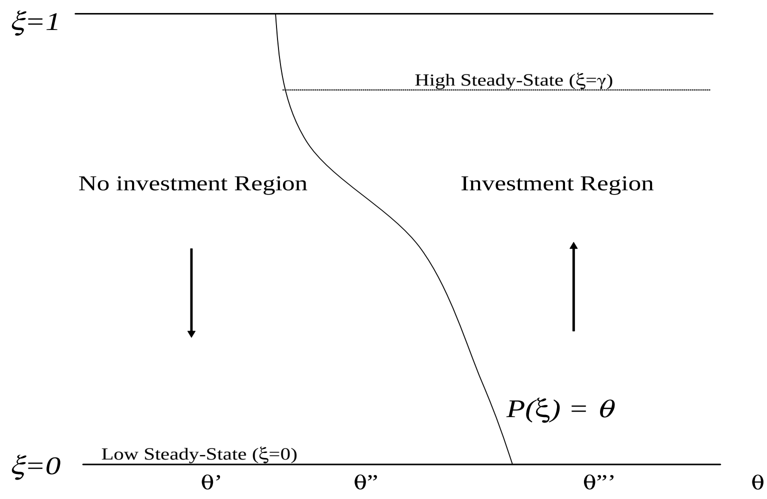

Figure 3: Dynamics of Statistical Discrimination

Figure 3 displays a potential equilibrium, and an example of persistence. Suppose increased access to resources raise θ from θ' to θ'', lowering the cost of investment. Even though for a group with costs $c(\theta'')$, it is possible to stay in a high-investment steady state, if the group has had higher costs for a significant amount of time, the current pool of workers searching for jobs will be low-skill, reducing the returns to entering workers. As a result, investment will not increase, and the pessimistic expectations will persist. A much larger increase in resources, for example to θ''', is needed to induce the entering cohort to invest and for employer expectations to begin to trend up.

As in the collective reputation model, there are really two features of the statistical discrimination model that combine to generate persistence. The first is the complementarity between employer expectations and worker incentives. The second is that employers cannot perfectly distinguish the dates at which potential employees invested. It is the latter source of imperfect information that makes the behavior of preceding cohorts relevant for decisions made today.

6 Conclusion

This paper has followed the lead of Tirole (1996) and explored dynamic incentives when agents care about their own individual reputation and about the behavior and reputation of their peers. When the environment evolves stochastically, group reputations can persist because the past, present and future actions of group members are strategic complements. The same mechanism explains why statistical discrimination of the sort analyzed by Arrow (1973) and Coate and Loury (1993) may persist even if policies are enacted that improve access to resources for a racial or ethnic minority. In contrast to the standard models of statistical discrimination that suggest a sudden change in beliefs could radically shift behavior, the present model shows that a group's history may have lingering effects.

Appendix: Omitted Proofs

The proofs of Propositions 2 and 3 closely mirror arguments in Frankel and Pauzner (2000), Burdzy, Frankel and Pauzner (2001) and Burdzy and Frankel (2005). In what follows, I provide a proof of Proposition 2, omitting a number of technical details that are available as a Supplementary Appendix. I then explain the small modifications needed to establish Proposition 3.

Proof of Proposition 2. There are three steps. The first step establishes properties of the best response correspondence for workers, in particular that the game has strategic complementarities. The second step uses the argument of Milgrom and Roberts (1990) to identify a highest and a lowest equilibrium. The third step uses the argument of Burdzy, Frankel and Pauzner to show that equilibrium is unique, with the monotonicity properties claimed in the proposition

Step 1: *Properties of Best Responses.*

Fix an increasing strategy $h_z(\theta, \xi)$ for the productive workers. Let $\rho_z(\theta, \xi; h)$ be the induced employer beliefs, i.e. those satisfying (1)-(2), and $R_z(\theta, \xi; h) = R(\rho_z(\theta, \xi; h))$ the resulting compensation. The salary $R_B(\theta, \xi; h) = \underline{R}$ is a constant irrespective of h, while $R_G(\theta, \xi; h)$ is increasing in all its arguments. Therefore the premium $R_G(\theta, \xi; h) - R_B(\theta, \xi; h)$ is weakly increasing in θ, ξ and h.

Recall that from an initial state (θ_0, ξ_0), ξ_t will satisfy the law of motion (4). Because $\dot{\xi}_t$ is increasing in h_{Gt}, h_{Bt}, and worker strategies are increasing in (θ, ξ), higher values of (θ_0, ξ_0) or higher strategies h_G, h_B will all lead to a higher realized path of ξ_t corresponding to a given path of θ_t.[8] Therefore an increase in θ_0, ξ_0, h_G or h_B will also increase the future

[8]Note that in comparing the evolution of ξ_t for higher and lower values of θ_0, I am comparing identical

salary premium at each date (for a given path of θ_t).

Now, define $\Delta(\theta,\xi;h)$ as in (6) to be the present value of having a good record rather than a bad one given current state (θ,ξ), and given that workers use the strategy h and employers expect this. The relative value Δ is continuous in θ,ξ and h. Application of the Milgrom and Segal (2002) envelope theorem shows that it is also weakly increasing in θ,ξ and h. This follows because the future salary premium is increasing in θ,ξ, and h. (Note that an increase in θ also decreases future effort cost, and here we can use the fact that there is always an optimal policy where effort is weakly higher with a good record to see that this effect will also increase Δ).

Finally, define the best response correspondence

$$BR_z(\theta,\xi;h) = \arg\max_{e\in[0,1]} e\left[(\phi_{zH}-\phi_{zL})\Delta(\theta,\xi;h) - c(\theta)\right].$$

Denote the highest and lowest best-responses by $\overline{BR}$, $\underline{BR}$. By Topkis' Theorem, both are increasing in θ,ξ, and h, and also satisfy $\overline{BR}_G \geq \overline{BR}_B$ and $\underline{BR}_G \geq \underline{BR}_B$.

Step 2: *Identifying Highest and Lowest Equilibria.*

Let h^0 to be the lowest strategy consistent with the dominance assumptions. That is, $h^0_z(\theta,\xi) = 1$ if $\theta > \overline{\theta}$ and 0 otherwise. This strategy is increasing. Iteratively define $h^{n+1}_z(\theta,\xi) = \underline{BR}_z(\theta,\xi;h^n)$. This gives a sequence of increasing strategies h^n, $n = 0,1,2,....$ Because h^0 is the lowest undominated strategy, $h^1_z(\theta,\xi) \geq h^0_z(\theta,\xi)$ and so by induction $h^{n+1}_z(\theta,\xi) \geq h^n_z(\theta,\xi)$. Moreover, no increasing strategy less than h^n is rationalizable. Let $\Delta^n(\theta,\xi) = \Delta(\theta,\xi;h^n)$ denote the sequence of relative returns associated with h^n.

Both sequences h^n and Δ^n are increasing. Denote their limits by h and Δ^∞. Of course h is an increasing strategy. By continuity, h and Δ^∞ must also satisfy the optimality condition (7). Finally, it follows from Burdzy and Frankel's (2005) Lemma 8 that $\Delta^\infty(\theta,\xi) = \Delta(\theta,\xi;h)$. So h is a best-response to itself, and hence the lowest equilibrium. An exactly analogous procedure starting from the highest undominated strategy identifies the highest equilibrium strategy $\hat{h}$.

Step 3: *Uniqueness and Monotonicity Properties.*

The last step is to show that the highest and lowest equilibria coincide, and that the unique equilibrium has the monotonicity property claimed in Proposition 2. Observe that because h is increasing, we can describe it as in Proposition 2: high effort if and only if $\theta \geq Q_z(\xi)$. The function $Q_z(\xi)$ is defined implicitly so that when $\theta = Q_z(\xi)$, we have

$$(\phi_{zH}-\phi_{zL})\,\Delta(\theta,\xi) = c(\theta).$$

incremental progressions of θ from the different starting points. Also, these comparative static conclusions implicitly assume that the path of ξ is uniquely determined by the path of θ, given an increasing strategy h. Burdzy and Frankel prove this will be so provided h satisfies an additional lipshitz property. Verifying that this property holds at each stage of the best response iteration described below requires some calculations that are reported in the Supplementary Appendix.

Define $\hat{Q}$ similarly to describe $\hat{h}$, i.e. using the appropriate relative value $\hat{\Delta}$ in place of Δ.

Both Q_z and $\hat{Q}_z$ are *strictly* decreasing functions for $z = G, B$. This follows because Δ and $\hat{\Delta}$ are weakly increasing and c is *strictly* decreasing, at least in the relevant region $[\underline{\theta}, \overline{\theta}]$. The functions Q_z and $\hat{Q}_z$ are also continuous because $\Delta, \hat{\Delta}$ and c are. Moreover, $Q_z(\xi) \geq \hat{Q}_z(\xi)$.

Define $dQ = \max_{\xi,z} Q_z(\xi) - \hat{Q}_z(\xi)$ to be the maximum distance between Q and $\hat{Q}$. Let ξ_d be the point at which this maximum is attained. Now define $\tilde{Q}_z(\xi) = Q_z(\xi) - dQ$. The idea of the proof will be to compare three alternative strategies — the lowest equilibrium strategy h, the highest equilibrium strategy $\hat{h}$, and the still higher *translation* of h defined by $\tilde{Q}$, and denoted $\tilde{h}$ — and show that they must be equal.

For each strategy, $h, \hat{h}$, and $\tilde{h}$, denote the correct employer beliefs by $\rho, \hat{\rho}$, and $\tilde{\rho}$, and the resulting compensation, as a function of (θ, ξ) by $R_z(\theta, \xi)$, $\hat{R}_z$ and $\tilde{R}_z$. Of course, $R_B = \hat{R}_B = \tilde{R}_B = \underline{R}$.

Now, fix starting points on each of the three isoquants $Q, \hat{Q}$, $\tilde{Q}$ as follows. Set $\xi_0 = \hat{\xi}_0 = \tilde{\xi}_0 = \xi_d$. And set $\theta_0 = Q(\xi_0)$, and $\hat{\theta}_0 = \tilde{\theta}_0 = \hat{Q}(\xi_d) = \tilde{Q}(\xi_d)$. Now fix an incremental brownian progression of $\theta_t, \hat{\theta}_t$ and $\tilde{\theta}$. So for all t, $\tilde{\theta}_t = \hat{\theta}_t$ and $\theta_t = \tilde{\theta}_t + dQ$. Define the corresponding paths ξ_t, $\hat{\xi}_t$, and $\tilde{\xi}_t$. Because $\tilde{Q}(\xi)$ is a translation of $Q(\xi)$ by dQ, and $\tilde{\theta}_t$ is an identical translation of θ_t, it follows that for all t, $\tilde{\xi}_t = \xi_t$. The path $\tilde{\xi}_t$, however, will be above the path $\hat{\xi}_t$, i.e. $\tilde{\xi}_t \geq \hat{\xi}_t$, because despite having the same starting point $\hat{\xi}_0 = \tilde{\xi}_0$, and comparing the same realization $\hat{\theta}_{t\geq 0} = \tilde{\theta}_{t\geq 0}$, we have $\tilde{h} \geq \hat{h}$. An immediate consequence is that if we consider the corresponding time t compensation, $R_{Gt} = \tilde{R}_{Gt} \geq \hat{R}_{Gt}$.

Finally, we compare $\Delta(\theta_0, \xi_0)$, $\hat{\Delta}(\hat{\theta}_0, \hat{\xi}_0)$, and $\tilde{\Delta}(\tilde{\theta}_0, \tilde{\xi}_t)$. Because of the ranking of flow compensation and the cost parameters, we have $\Delta(\theta_0, \xi_0) \geq \tilde{\Delta}(\tilde{\theta}_0, \tilde{\xi}_0) \geq \hat{\Delta}(\hat{\theta}_0, \hat{\xi}_0)$. So we have established that

$$
\begin{aligned}
(\phi_{zH} - \phi_{zL})\Delta(\theta_0, \xi_0) - c(\theta_0) &\geq (\phi_{zH} - \phi_{zL})\tilde{\Delta}(\tilde{\theta}_0, \tilde{\xi}_0) - c(\tilde{\theta}_0) \\
&\geq (\phi_{zH} - \phi_{zL})\hat{\Delta}(\hat{\theta}_0, \hat{\xi}_0) - c(\hat{\theta}_0).
\end{aligned}
$$

The first inequality holds with equality only if $dQ = 0$, and the second only if $\hat{Q} = \tilde{Q}$. Because Q and $\hat{Q}$ correspond to equilibria, however, the first and third expression are by definition both equal to zero. Therefore $Q = \tilde{Q} = \hat{Q}$, or equivalently the highest and lowest equilibria coincide.

Proof of Proposition 3. The proof follows the same steps as above. Step 1 is actually a bit simpler. Newly arriving workers do not have an individual history to condition on, so they choose a strategy $h(\theta, \xi)$. As in the collective reputation model, the compensation premium $R_G - R_B$ is increasing in (θ, ξ, h), in this case because R_G is strictly increasing in ξ and independent of θ or h, and $R_B = \underline{R}$. With this in mind, similar arguments to above establish that $\Delta(\theta, \xi; h)$ is continuous and increasing in its arguments, and hence $\overline{BR}(\theta, \xi; h)$ and $\underline{BR}(\theta, \xi; h)$ are both increasing in θ, ξ and h. From there, steps 2 and 3 of the proof are identical to above.

References

Arrow, Kenneth, "Some Mathematical Models of Race Discrimination in the Labor Market," in A. H. Pascal ed., *Racial Discrimination in Economic Life*, Lexington, MA: D. C. Heath, 1972.

Arrow, Kenneth, "The Theory of Discrimination," in Orley Ashenfelter and Albert Rees, eds. *Discrimination in Labor Markets*, Princeton: Princeton University Press, 1973.

Blume, Lawrence, "The Dynamics of Statistical Discrimination," *Economic Journal*, November 2006, Vol. 116, pp. F480 - F498.

Burdzy, K., David M. Frankel, and Ady Pauzner, "Fast Equilibrium Selection by Rational Players Living in a Changing World," *Econometrica*, 2001, Vol. 68, pp. 163-190.

Coate, Stephen and Glenn Loury, "Will Affirmative Action Policies Eliminate Negative Stereotypes," *American Economic Review*, 1993, Vol. 83(5), 1220–1240.

Frankel, David M. and Burdzy, Krzysztof, "Shocks and Business Cycles," *Advances in Theoretical Economics*, 2005, Vol. 5(1), Article 2.

Frankel, David M. and Ady Pauzner, "Resolving Indeterminacy in Dynamic Settings: The Role of Shocks," *Quarterly Journal of Economics,* 2000, 115, pp. 283-304.

Milgrom, Paul and John Roberts, "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities," *Econometrica*, November 1990, Vol. 58, pp. 1255-77.

Milgrom, Paul and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, March 2002, Vol. 70(2) pp. 583-601

Moro, Andrea and Peter Norman, "A General Equilibrium Model of Statistical Discrimination," *Journal of Economic Theory*, 2004, Vol. 114, pp. 1-30.

Phelps, Edmund, "The Statistical Theory of Racism and Sexism," *American Economic Review*, 1972, Vol. 62(5), 659–661.

Tirole, Jean, "A Theory of Collective Reputations (with applications to the persistence of corruption and to firm quality)," *Review of Economic Studies*, 1996, Vol. 63(1), pp. 1–22.